Ontology Management and Selection In Re-Use Scenarios

Kim Finney (B.Sc, M.Sc)

Submitted in fulfilment of the requirements for the degree of Doctor of Philosophy

University of Tasmania

School of Computing and Information Systems

October 2012

This thesis contains no material which has been accepted for a degree or diploma by the University or any other institution, except by way of background information and duly acknowledged in the thesis, and to the best of my knowledge and belief no material previously published or written by another person except where due acknowledgement is made in the text of the thesis, nor does the thesis contain any material that infringes copyright.

The research associated with this thesis abides by the international and Australian codes on human and animal experimentation, the guidelines by the Australian Government's Office of the Gene Technology Regulator and the rulings of the Safety, Ethics and Institutional Biosafety Committees of the University.

Acknowledgements

Several people require special mention who have supported me to deliver this thesis.

First, is my partner, Leanne Wilkes who has lived through this episode vicariously and who probably deserves the award more than I do. Without her critical eye, robust comments and stamina beyond belief in reading and re-reading this dissertation I would never have finished.

Second are my parents. My father, Bob Finney a man of prodigious talent whose enquiring mind inspired me to undertake this research in the first place, offered much useful criticism and advice. My mother, Brenda Finney as always practical and supportive, helped me push on when the end wasn't clearly in sight.

Last, but not least are my supervisors, Simon Milton, Chris Keen and Peter Marshall whose reviews, useful insights, experience and comments helped me to make a thesis.

I would also like to dedicate this dissertation to Elizabeth Bell (1916 – 2009), who didn't understand the Web, had never heard of an ontology and who disliked computers, but encouraged me none-the-less, just as she had done for everything else I've ever attempted in my life.

Abstract

One of the main impediments to realising the Semantic Web vision is that most scientific data, even those data deployed on the web, are not generally expressed or encoded in an unambiguously defined, machine-interpretable manner. This is particularly the case for Antarctic-themed data. Ontologies that are linked to datasets via semantic annotation are required to achieve semantic-enablement of scientific data infrastructure. In scientific communities that adhere to the Open Geospatial Consortium Service-Oriented-Architecture (Web services) paradigm, Feature Catalogues are the repositories intended to manage and publish descriptions of dataset concepts. This thesis explores how Feature Catalogues can be ontologically-grounded to facilitate semantic annotation and in doing so addresses the lack of guidance in current standards about how to configure an ontologically grounded Feature Catalogue and how best to access the resources it contains for the semantic annotation of Web services. Also investigated is how ontology selection and evaluation is currently taking place in practise because ontology evaluation methodologies mentioned in the literature are resource intensive to apply, often requiring a high level of ontological expertise. Both contributions seek to lower barriers for ontology uptake and reuse within scientific communities.

To address these issues, two scientific communities of practise (i.e., AODN and SCAR) were used as case studies within a Design Science research method to ground-truth the design and to prototype an ontologically grounded, service-enabled Feature Catalogue. To address research questions pertaining to ontology selection and evaluation practise, fourteen experts (from outside of the AODN and SCAR communities) with experience in building semantically-enabled scientific infrastructure, were surveyed and interviewed to ascertain what ontology evaluation methods and criteria are being used in practise. A hierarchical evaluation model was established from analysed expert data using Template Analysis (Crabtree and Miller, 1992; King, 2004). The Analytical Hierarchical Processing (AHP) technique (Saaty, 1980), was then harnessed to establish the relative importance given by experts to each of the model elements.

The contributions arising consisted of an enhanced ISO 19110 Feature Catalogue model which accommodated additional concepts necessary to describe the observation-centric dataset paradigms of the two case study communities. The extended conceptual model was semantically grounded using the DOLCE (upper ontology) and expressed in both OWL and SKOS. Demonstration REST-based service interfaces (and REST query patterns) were created for serving Catalogue content to requesting Web clients. To the author’s knowledge, no other Feature Catalogue implementation, founded on the ISO 19110 conceptual model, has attempted to model the Catalogue as an ontology, or permits access to Catalogue content via REST-based service interfaces. This thesis also delivers a

“practical” framework for evaluating and then selecting reusable ontological content which encompasses weighted model elements (indicating relative levels of importance), coupled with expert-derived evaluation metrics. Although the evaluation criteria listed in the framework are not novel in themselves, identifying which criteria are of most utility to experts who are operating in real-world scenarios, is an important contribution to practise.

Table of Contents

List of Figures

List of Tables

Chapter 1.

Introduction

The scientific data landscape is expanding rapidly in both scale and diversity. For example, high-throughput gene sequencing platforms are capable of generating terabytes of data in a single experiment (Goble and De Roure, 2009), the Large Hadron Collider is producing 13 petabytes of stored physics data each year (Brumfiel, 2011) and global climate models generate tens of terabytes of data with each model run (Williams, 2011). Such high data volumes necessitate collaborative data mining and data analyses and because of improved accessibility, increasingly data are being re-used by scientists not originally connected with data generation. Such situations require attention to dataset documentation. For scientists dealing with heterogeneous data drawn from diverse and distributed sources, understanding what the data means is essential for its use and for enabling machine-assisted data integration and processing (Fox and Hendler, 2009).

This thesis is concerned with developing technological approaches, methods and guidance that collectively have the capacity to improve the ease with which the meaning of a dataset and its component parts can be documented. Unambiguously defined datasets are the foundation of semantically-enabled data exchange scenarios in which a dataset's descriptive elements permit automated communication and cooperation between machines to operate on these data. Ontologies are the technology that provides the type of description necessary for machine to machine communication because they specify how a given scientific community interprets and encodes their vocabularies. Semantic-enablement establishes a reference between the vocabularies used in the exchanged data and the ontologies developed to describe those vocabularies. The process of establishing these links is commonly referred to as semantic annotation (Maue, 2009).

This introductory chapter sets the scene for the research which follows by explaining the problems currently encountered which are impeding progress with formulating machine interpretable dataset descriptions (i.e., semantic annotations) and highlights the challenges which are restricting semantic-enablement of scientific data infrastructure. The motivation for asking the research questions posed in this thesis are also explained; and the chapter provides an overview of the methods which are used to conduct the research; at the same time giving a broad outline of the contributions made to theory and practise. This chapter concludes by presenting the reader with a general guide to the thesis structure.

1.1 Research Motivation and Questions

The explosion in scientific data generation is leading to unprecedented amounts of data available on the Web. Scientific communities are increasingly reliant on Web-based applications for the communication of raw research data and derived products. Sophisticated data access portals are now evident across the spectrum of scientific disciplines (e.g., see GEOSS (GEO, 2012); GCMD (NASA, 2012a); INSPIRE Geoportal (European Commission, 2012)). But most scientific resources are not yet deployed in a Semantic Web context (e.g., see Parsons *et al*. (2011) for a précis of the State of Polar Data). There are, however, some notable semantically-enabled exceptions now occurring, particularly in the biology and earth physics realms (e.g., see BioPortal (NCBO, 2012) and VSTO (VSTO, 2012)).

The Semantic Web, as envisioned by Berners-Lee *et al*. (2001) is one where there is data interoperability across applications and organizations, using a set of interoperable standards for knowledge exchange, and where there is an architecture that supports interconnected communities and vocabularies. This vision is an attractive one for science, not the least because of the opportunities which could be realised through integrating and mining existing data sources currently residing in thousands, if not millions of information silos.

But one of the main impediments to realising the Semantic Web vision is that most scientific data, even those data deployed on the Web, are not generally expressed or encoded in an unambiguously defined, machine-interpretable manner (Gil *et al*., 2006; Manning *et al.,* 2009). There are many reasons for why this is the case:

- Much scientific data have in the past been captured in non-digital form and can be expensive to digitise, re-interpret and publish online, particularly if supporting context documentation (i.e., metadata) is absent (Blue Ribbon Task Force on Sustainable Digital Preservation and Access and Smith Rumsey, 2010).
- Languages, tools and development environments necessary to create dataset descriptions have been, until relatively recently, the province of the Artificial Intelligence community (inclusive of logicians, linguists and computer programmers). It is only in the past half decade that practitioners in the broader science community have had the capability to more readily access technologies and experts that can facilitate the tasks of semantic data encoding. However, the skills required to adequately manage data in a semantic context are still undervalued (ICSU, 2011).

- Whilst semantic technologies are much more pervasive now than they have been, many pieces of the technology stack (inclusive of standards) are still not user-friendly, fit-to-task, robust, or are absent (Lefort, 2009; Manning *et al.,* 2009; Garcia-Castro and Gomez-Perez, 2011).
- Science has traditionally been highly discipline-focussed. Many early scientific web-based data delivery systems were created to service a very specific user-community who implicitly shared a common lingua franca (and usually tools), geared around informal and implicit standards. These groups weren't able to justify the investment required to make dataset vocabulary definitions explicit and saw little need to do so. The culture of science has changed in recent times towards much more collaborative, inter-disciplinary studies propelled by a need to solve complex global environmental and social problems (e.g., climate change and disease control) and is being facilitated by the availability of the internet and an ever-increasing array of new data sampling technologies. As a result, the benefits of enabling the sharing of data across disciplines are now easier to establish, but this was not always the case (Sidi, 2010; European Commission, 2010).
- There are different vocabulary uses and definitions, i.e., naming and cognitive heterogeneity across scientific disciplines (and often within a single discipline), that may also happen to change over time, making the formal semantic encoding task and the management of these resources particularly difficult (Fox *et al.,* 2009; Manning, 2009).

This thesis is therefore concerned with investigating practical ways in which some selective aspects of these afore-mentioned barriers can be overcome in order to encourage and assist scientific communities to deploy semantic datasets. This has required the formulation of a range of research questions (RQ) in response to the challenges that are evident. To identify the specific research questions being addressed, each research question discussed in this chapter is presented in *italics* and is labelled with a number in red (e.g., RQ1). These labels are then referred to throughout the thesis in order to help the reader understand the association between the questions posed, the methods used to address them, various aspects of data analysis and the specific contributions to theory and practise that are ultimately made.

1.1.1 Challenges In Creating Semantic Repositories For Feature-Centric Services

The word "semantic" has already been used several times. It is loosely defined here-in as the study of meaning (Farlex, 2012a). Formalising semantics is therefore concerned with understanding linguistic meaning by constructing precise mathematical models of the principles that people use to define relations between expressions in a natural language and the world that supports meaningful discourse (Wikipedia, 2012a). Creating semantic datasets involves scientific communities agreeing upon the terms inherent in exchanged datasets, their usage and their encoding. A formal way of organizing this type of knowledge is by the use of ontologies. Guarino (1998) describes an ontology as “an engineering artefact, constituted by a specific vocabulary used to describe a certain reality, plus a set of explicit assumptions regarding the intended meaning of the vocabulary words”. Ontologies relate concepts, or categories of things that exist or may exist, in some domain, to one another through logically defined relationships (Gruber, 2008).

An ontology can be considered as a graph of information, with terms (concepts) as nodes of the graph and relationships as the links that connect the terms. Relationships may be directed, which means that they are only true in one direction (e.g., the “retina” is *part of* an “eye”, but an “eye” is not *part of* a “retina”). Generally, ontologies are hierarchical in structure (e.g., an “eye” is a top node and below that is a “retina” node, this “retina” node may then have as its parts, “rods” and “cones” [both third level nodes] and so on). The most common type of hierarchical relation is the “*is a*” relation between two concepts (e.g., “A” *is a* “B”) indicating that concept “A” is a sub-type of “B”. However, the relationships used in an ontology are not predetermined, so any real-world relationship can be logically defined and used to connect terms and reflect reality. This makes ontologies a flexible framework for modelling many different kinds of data (Washington and Lewis, 2008).

Formal ontologies are encoded using axioms and definitions stated in logic, or in some computer-oriented language that can be automatically translated to logic. The utility of this encoding is that the defined relationships, between ontological terms, makes it possible for a computer to use logic to discover asserted or inferred relationships between concepts (and hence deduce meaning). Asserted relationships are those where there is a direct link deliberately constructed between the concepts in an ontology (e.g., “A” *is a* “B”), whilst inferred relationships are found by logically drawing a connection between two nodes, usually through intervening nodes and relationships (Washington and Lewis, 2008).

Semantic Web ontology languages are RDF (Minola and Miller, 2004), RDFS (Brickley and Guha, 2004) and OWL (McGuinness and van Harmelen, 2004). The underlying structure of any expression in RDF is a collection of triples, each consisting of a subject, a predicate and an object. The Simple Knowledge Organization System (SKOS – Miles and Brickley, 2005), although not a method for formalising ontologies, is a common data model for sharing and linking knowledge organization systems via the Semantic Web which uses both RDF and OWL to represent thesauri, taxonomies, and other controlled vocabularies. All are based on XML (Bray *et al*., 2008), which makes XML the most predominantly used mark-up language to encode and serialise ontologies deployed in Web environments.

Creating semantic descriptions of scientific datasets therefore relies on the assembly of ontologies in an appropriate encoding language, by communities-of-practice who agree explicitly to the use of these ontological formalisms. These ontological descriptions must be recorded, managed and then made accessible so that they can be referenced and used by both humans and machines. The generally accepted technology for achieving these latter tasks is via an ontology repository.

Ontology repositories (and standards for their implementation and communication) are, however, still in their infancy (Baclawski and Schneider, 2009). Hartmann *et al.* (2009) reported that the process of identifying and accessing ontological resources, which can be summarized as ontology retrieval, is confounded by an absence of mechanisms and procedures for storing and representing ontologies. Some headway, however, has been made in these areas of late, mainly through efforts driven by the Open Ontology Repository initiative (OOR, 2012; Baclawski and Schneider, 2009) and in the development of tools such as "Cupboard" (d'Aquin and Lewen, 2009). Cupboard, however, is more of a system to host ontology repositories, than it is an ontology repository.

Much more commonly deployed in scientific communities, to facilitate the description and use of datasets, are metadata repositories (also known as metadata catalogues). Metadata repositories often draw upon controlled vocabularies, taxonomies, thesauri and data dictionaries to describe dataset resources and have not traditionally been ontologically-grounded (Gil *et al.,* 2006). These repositories more recently also encompass content schemas for describing the services that encapsulate the data, i.e., Web services metadata (e.g., see Whiteside, 2005).

Descriptive metadata (whether ontologically grounded or not) are ideally meant to be closely coupled with published data and various components of a metadata description are necessary at many points in a Web services-based infrastructure (Senkler *et al.,* 2004; Nebert *et al.,* 2007). However, the majority of existing scientific metadata contained in metadata repositories are shallow, that is, contain descriptions at a level of granularity suitable for describing an entire dataset, rather than descriptions that can be anchored to individual components within a dataset (Liao and Hong, 2005). They are often also de-coupled, i.e., physically separated from the data that they describe (e.g., in the GCMD polar metadata portal only 54% of metadata records are linked to the datasets they describe (SCADM, 2011)). Shallow metadata provides information such as the overall quality of a dataset, dataset lineage, keywords for tagging dataset content, custodial contact information, and so on. In contrast, deep, granular metadata includes descriptive elements about the various individual items a dataset might contain (Craglia *et al.,* 2007), for example data capture information repeated for each type of listed parameter and/or units of measure per parameter. This deep metadata is modular so that it can be packaged with, or ascribed to, individual items in a dataset.

To achieve semantic interoperability scientific communities must therefore deploy their metadata using ontologies. They must also create this ontologically-grounded metadata at an appropriate level of granularity, depending on where it is deployed within the Web services architecture, mindful of its intended purpose. Some of the recently published, internationally accepted metadata standards such as ISO 11179 (ISO/IEC, 2010) and its spatially-focused offspring - ISO 19115 (ISO, 2003), provide hooks (through certain types of metadata elements) which could include, or reference ontological content. Most scientific communities have yet to capitalise on these constructs primarily because there is a lack of guidance, little standardisation and few exemplary implementations from which to draw upon (Schuurman and Leszczynski, 2006; Kubik and Iwaniak, 2010).

This thesis addresses a current gap in practise and theory relating to the use of ontologies for creating scientific dataset descriptions by examining the needs of the Antarctic scientific community. Antarctic science is a multi-disciplinary program of studies, comprising researchers who generally investigate physical phenomena occurring within the Antarctic region (inclusive of marine, terrestrial, cryospheric and atmospheric research). Many groups within this very broad, international community conduct observational studies aimed at understanding Antarctic ecosystems, physics and phenomena and the role that these all play in the Earth System. Temporal and spatial aspects of the data that are collected are often of high significance. For this reason many research groups within the Antarctic community have subscribed to the ISO TC 211 suite of digital geographical data standards to which the ISO 19115 metadata standard also belongs (ISO, 2012). This thesis therefore focuses on ontology use for dataset description within data infrastructures that have already chosen to adopt ISO TC 211 dependent standards.

The ISO TC211 family of digital geographic data standards are unified by a conceptual reference model to enable compliant application systems to inter-operate and share conforming geographic data. This reference model is called the ISO 19101 General Feature Model (ISO, 2002). In this model a feature instance is an identifiable object in the world, or the digital representation of it. Features can therefore be considered equivalent to “concepts” in ontological terms. In ISO 19101, features are classified into feature types on the basis of common sets of characteristics or properties and Feature Catalogues are envisaged to contain shared community definitions of feature types (see Figure 1.1). It is contended in this thesis that Feature Catalogues (ISO, 2005b) should operate analogous in many respects to an ontology repository and house much of the information necessary to support semantic dataset descriptions.

Although the TC211 standard (ISO 19110) for development of Feature Catalogues has been published since 2005, implemented examples of operational on-line Catalogues are still relatively uncommon (see the Scientific Committee on Antarctic Research Feature Catalogue (AADC, 2012a), UK Digital National Framework Catalogue (DNF, 2012) and the Intergovernmental Committee on Surveying and Mapping Feature Catalogue (ICSM, 2008) for some exemplars). Additionally, those catalogues that do exist generally do not use Web service interfaces (i.e., self-contained, self-describing, modular application interfaces that can be published, located, and invoked across the Web).

The Open Geospatial Consortium (OGC – OGC, 2012a), the main international standards body for geospatial content and services, has adopted the ISO TC211 feature-centric model as the corner-stone of its abstract specifications that are concerned with serving and accessing geospatial data (ISO, 2009b).

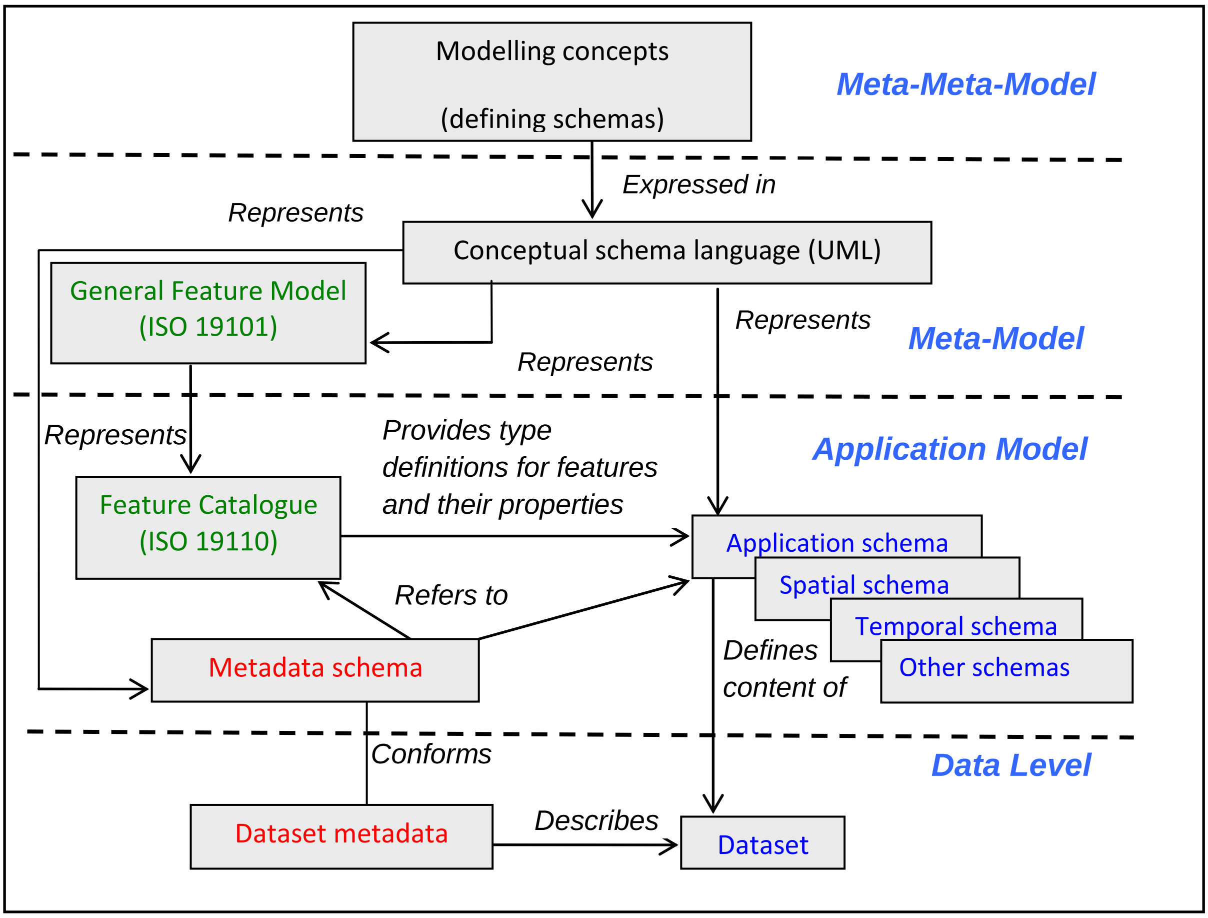

Figure 1.1 A Feature Catalogue's Role in the ISO General Feature Model (Kerherve, 2006)

The OGC and ISO/TC 211 share the objective of providing a framework for the development of domain applications using geospatial resources (see Figure 1.2 for the relationship between key OGC framework elements). Both of these bodies have a working arrangement that often results in virtually identical OGC and ISO standards (Kresse and Fadaie, 2004).

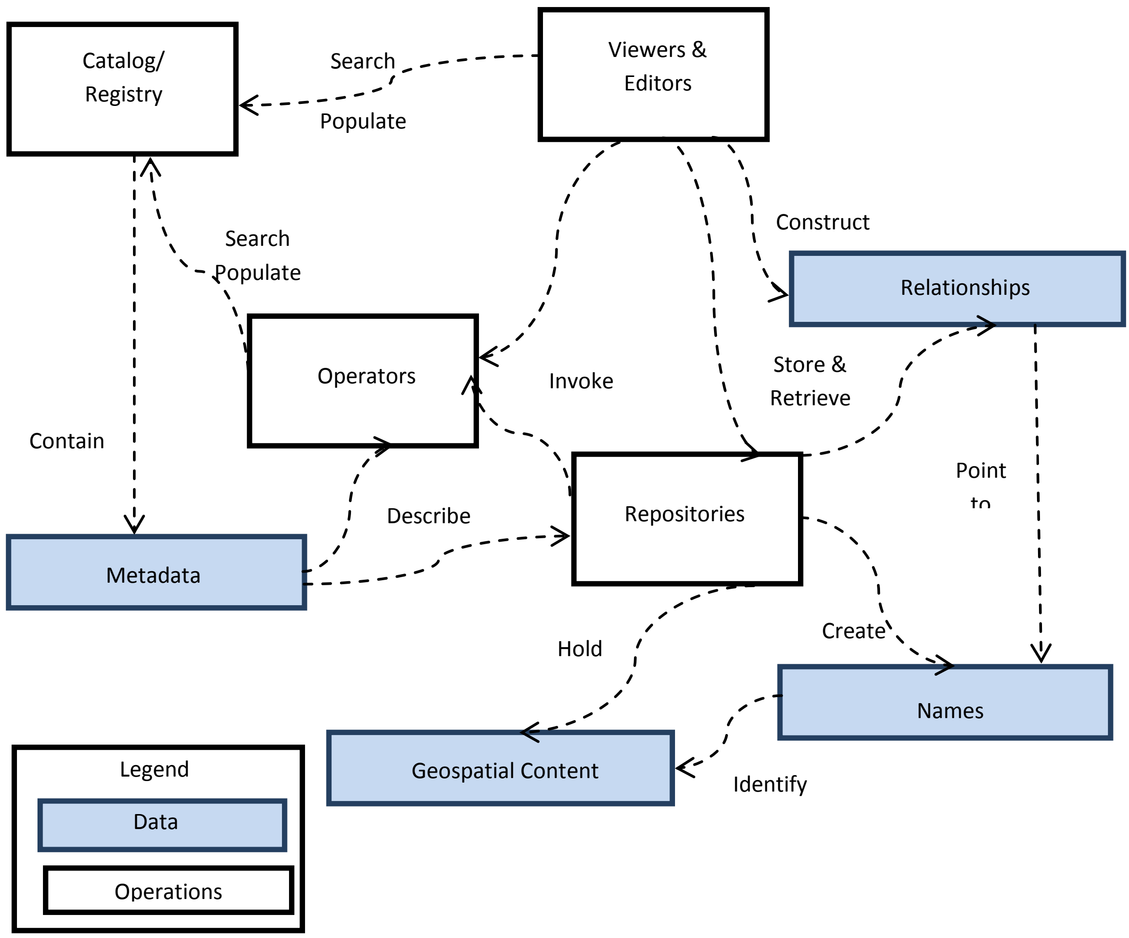

Figure 1.2 Contextual Schema Showing Relationship Between OGC Web Services (from Figure 1, Reed, 2010.) Note: in this figure "Operations" are a specification of a transformation, or query that a service may be called to execute. Operations are therefore synonymous with the various classes of service as listed in Table 1.1.

Many Antarctic communities now use OGC Web service interfaces to publish and exchange data (see Table 1.1 for the OGC standards commonly used). Given the number of scientific communities using OGC Web services, the very low number of Feature Catalogue exemplars is all the more surprising since the OGC also directs that Feature Catalogues should be the source of Feature Type definitions for features deployed in OGC Web services (OGC, 2011).

The tendency for large existing repositories of on-line metadata to describe data at a "dataset", rather than at a "feature instance" level, means that existent metadata rarely includes references to community-agreed Feature Types. Semantic matching across metadata records and across referenced datasets is usually performed via lexically comparing keywords (or other elements), drawn from community-specific controlled vocabularies. For those communities who want to pursue semantically assisted data search and data integration, there is almost no guidance about how to configure an ontologically-grounded Feature Catalogue. An as yet unresolved question is therefore what would *characterise an ontologically-grounded Feature Catalogue that can support Antarctic science data publication through Web services* (RQ1.1)? To fully answer this question it is first necessary *to identify the types of use-cases and data models that the Feature Catalogue must support* (RQ1.1.1). These gathered requirements then *need to be assessed in light of existing ISO and OGC conceptual models to determine if these models actually meet current needs* (RQ1.1.2), particularly given the rapidly evolving nature of Web technologies. Regardless of the outcome of such an assessment it would still remain *to semantically ground these conceptual models,* which of itself is an open research question (RQ1.1.3).

There is also a paucity of information and a lack of examples demonstrating how Feature Catalogue content can be delivered to describe datasets within the service-oriented-architectures underpinning many current scientific data exchange activities. This raises the further question of *what methods are best suited to extract re-usable content from an ontologically-grounded Feature Catalogue* (RQ1.1.4).

Table 1.1 OGC Web Service Interfaces Subscribed to By Antarctic Communities

Web Service Interfaces	Abbr.	Since	Current Version	Purpose
Web Map Service	WMS	2000	1.3.0 (de la Beaujardiere, 2006)	Portrayal
Web Feature Service	WFS	2002	2.0.0 (Vretanos, 2005)	Download features
Web Coverage Service	WCS	2003	2.0.0 (Baumann, 2010)	Download coverages
Catalogue Service for the Web	CSW	2004	2.0.2 (Nebert *et al.,* 2007)	Resource and service discovery
Web Processing Service	WPS	2005	1.0.0 (Schut, 2007)	Remote service invocation
Sensor Observation Service	SOS	2006	1.0.0 (Na *et al.*, 2007)	Download observations

1.1.2 Challenges In Selecting and Evaluating Ontologies

Apart from the investment required to build tools such as Feature Catalogues (as ontology servers), creating the actual (ontological) Feature Catalogue content is a highly resource-intensive task. Many knowledge and Web engineers therefore extol the virtues of re-using existing ontologies, or their ontological components wherever practical (e.g., Uschold *et al.,* 1998; Annamalai & Sterling, 2003; Blomqvist *et al.,* 2006). Assuming re-use is a more cost-effective method of ontology development, and noting that the re-use of ontologies should also result in a higher degree of semantic interoperability (Simperl, 2009), scientific communities should be adopting this approach wherever possible.

But cost-effective re-use of ontologies necessarily implies that there also exists efficient and well-developed methods for selecting and evaluating ontologies that are candidates for re-use. Ontology selection is the process that enables the identification of one or more ontologies, or ontology modules that satisfy certain criteria. The actual process of checking whether an ontology satisfies nominated criteria, is an ontology evaluation task (Sabou *et al.,* 2006). According to Orbst *et al.* (2007), evaluation criteria, by which the different qualities of an ontology can be judged, are manifold and can include:

- an ontology's coverage of a particular domain and the richness, complexity and granularity of that coverage;
- specific use cases, scenarios, requirements, applications, and data sources an ontology was developed to address;
- formal properties such as the consistency and completeness of an ontology and the representation language in which it is modelled;
- whether they are mappable to some specific upper (more general) ontology or are assessed on the basis of their underlying philosophical theory about reality, and
- the types of reasoning (inference) methods that can be invoked on an ontology.

It should be noted that evaluation criteria differ from evaluation measures. Ontology evaluation measures are a quantitative means of assessing various aspects of an ontology (Yu *et al.,* 2007). Gomez-Perez (2001), Brewster *et al.* (2004) Tartir *et al.* (2005), Gangemi *et al.* (2005) and Vrandecic, (2010) have all proposed various types of measures useful in measuring a range of criteria. Yu (2008) asserts, however, that measures deemed important for one application may not necessarily be important for others and he promotes a requirements-based ontology evaluation incorporating a relevant set of measures that he states must are tailored, concrete, relevant and meaningful.

Unfortunately there isn't a comprehensive, global and definitive approach to the selection and evaluation problem, despite a range of literature on the topic. Many of the multi-faceted evaluation methodologies mentioned in the literature (e.g., Ontometric: Lozano-Tello and Gomez-Perez, 2004; OntoClean: Guarino and Welty, 2004; EvalExon: Spyns (2005)) are resource intensive to apply, often requiring a high level of ontological expertise (Hartmann *et al.,* 2005b; Kalfoglou & Hu, 2006; Blomqvist *et al.,* 2006). Minimising the amount of effort and skill necessary to re-use ontologies, or ontological components should lower barriers for ontology re-use, particularly in scenarios where most data publishers are domain or IT experts, rather than ontological engineers.

In a previous study by Paslaru Bontas-Simperl and Tempich (2006), which examined ontology-related projects, it was reported that only a small percentage of the sampled projects demonstrated a commitment to using any systematic ontology development method. Since ontology evaluation is a component of an ontology development life-cycle, it is highly possible that formal methodologies and/or criteria for ontology selection and evaluation, particularly those outlined in the literature, are not those being used by domain communities developing semantic infrastructure capabilities. It is conjectured here that a more detailed look at what is actually taking place in practise may reveal selection and evaluation criteria and/or methods which are easier and faster to implement, but which still address community requirements. Of specific interest in this study is what are the most important evaluation criteria, in a relative sense, when a community is choosing an ontology for re-use. By only using those criteria that are most highly rated, the process of evaluation can be streamlined.

The research outlined in this thesis is focussed on how ontology use and re-use can be fostered within scientific communities, particularly those that subscribe to the ISO TC 211 and OGC spatial services standards suites, where it is expected that ontological expertise is limited. This study therefore sought to construct a selection and evaluation framework grounded in current practise, which could be applied by scientific experts (who may have few ontological skills) and who are responsible for describing, managing and deploying scientific datasets. Because the delivered framework will be developed directly from the knowledge and experience of people who have been actively engaged in scientific semantic infrastructure enablement projects, it was anticipated that any techniques emerging would be both pragmatic and proven in practise. By being aware of the elements of an expert-grounded framework, even communities without ontology development capacity can make an early assessment (when embarking upon semantic-enablement activities) about the types of skills and practises which may need to be imported, or cultivated within the community to help instantiate and use semantic dataset descriptions.

The key research question framed to investigate current ontology selection and evaluation practise is therefore: *what typifies an expert-grounded ontology selection and evaluation framework that can support multi-disciplinary Antarctic science communities using Web services* (RQ1.2) ? To answer this question three sub-questions needed to be posed: *what ontology selection and evaluation criteria are currently used across multi-disciplinary scientific communities [and are selection and evaluation methods consistent with those reported in the literature]* (RQ1.2.1) ?; *Is it feasible to derive a* *weighted evaluation criteria model in which criteria are rated according to importance* (RQ1.2.2) ?; and *what evaluation measures can be used to assess evaluation criteria* (RQ1.2.3) ?

To derive a weighted evaluation criteria model (RQ1.2.2), suitable for deployment in a multi-disciplinary scientific environment across varying use cases, research methods had to be developed to establish *whether evaluation criteria were considered by experts to have differing or equal weight* (RQ1.2.2.1) and to ascertain *whether these assigned weights differed depending on an expert's scientific domain, or any other discernible factor* (RQ1.2.2.2).

The over-arching research question which encapsulates all of the research questions already mentioned in sections 1.1.1 and 1.1.2, covering matters of both tooling and practise is: *how can Antarctic science communities practically manage and select domain ontologies for use in semantically-enabled data exchange scenarios, given feature-centric Web service design patterns* (RQ1) ? Figure 1.3 summarises the various questions that drive investigation of this problem and shows their relationship.

1.2 Methods and Contributions Overview

The research methods employed in this thesis were designed to develop a deeper understanding of the role that tools like Feature Catalogues can play in promoting ontology use in dataset descriptions and to examine whether current Feature Catalogue standards are capable of adequately supporting scientific requirements.

They are also geared towards investigating how ontology evaluation methods currently influence ontology selection decisions made within scientific communities with a view to providing practical guidance for those communities not yet engaged in semantic dataset description activities, or for those communities who may have commenced such initiatives, but as yet have no template to follow.

1.2.1 Design Science Research

All of the research questions concerning the design and use of a Feature Catalogue, as an ontological repository (i.e., all RQs stemming from the research question RQ1.1 – *What characterises an ontologically-grounded Feature Catalogue that can support Antarctic science data publication through Web services*), have been addressed through the application of Design Theory (Hevner *et al.,* 2004).

Studies employing Design Research (also called Design Science Research) "involve the analysis of the use and performance of designed artefacts to understand, explain and very frequently to improve on the behaviour of aspects of Information Systems" (Vaishnavi and Kuechler, 2004). Gregor and Jones (2007) in their articulation of Design Theory state that the phenomena of interest in Design Research include: instantiations or material artefacts (i.e., artefacts that have a physical presence or the physical actions/processes/interventions that lead to an artefact with a physical presence); theories or abstract artefacts (e.g., constructs, models and methods) and human understanding of artefacts (e.g., conceptualisations described in abstract terms). The methods used in this thesis were concerned with developing artefacts that were mainly models (e.g., conceptual data models, class and ontological models) for an ontologically-grounded Feature Catalogue, based on stated Antarctic community needs. The conceptual models were used to benchmark the relevance of existing ISO/OGC Feature Catalogue related standards.

Limited-in-function concrete instantiations of an ontologically-grounded Feature Catalogue (using both traditional RDBMS and semantic technologies) were developed as proof-of-concept designs and to enable the material testing of developed Catalogue content interface methods and access query patterns. A beneficial outcome from applying the Design Research paradigm was also "process" knowledge gained from reflections during the research which are now transferable as process understanding and implementation guidance for those communities wishing to embark upon a Feature Catalogue build task.

In Design Science, the research commences with an "*awareness of a problem*" (Vaishnavi and Kuechler, 2004), which for this study is an absence of understanding about what characterises an ontologically-grounded Feature Catalogue. *Suggestions* for a problem solution are then abductively (Pierce, 1931) drawn from the existing knowledge/theory base for the problem area. This is generally followed by the implementation of an artefact. In this study implementations of an ontologically-enabled Feature Catalogue and a set of interface methods were developed as prototypes (post modelling). This stage is not surprisingly known as "*development"*. Development is followed by *evaluation* (against explicit and implicit aspects of specifications identified in the *suggestion* phase). The *development* and *evaluation* phases sometimes lead to further *suggestions* and this triad is recursively visited until, after a suitable amount of *circumspection*, the process is deemed to be complete. It should be noted that a potential limitation of this study was that *development* did not lead to an operational artefact and stopped at the prototyping phase. Although this is an acceptable practise in Design Science, the level of validation that can be achieved by instantiating a fully integrated and functional system, is not attainable to the same degree in the prototyping situation.

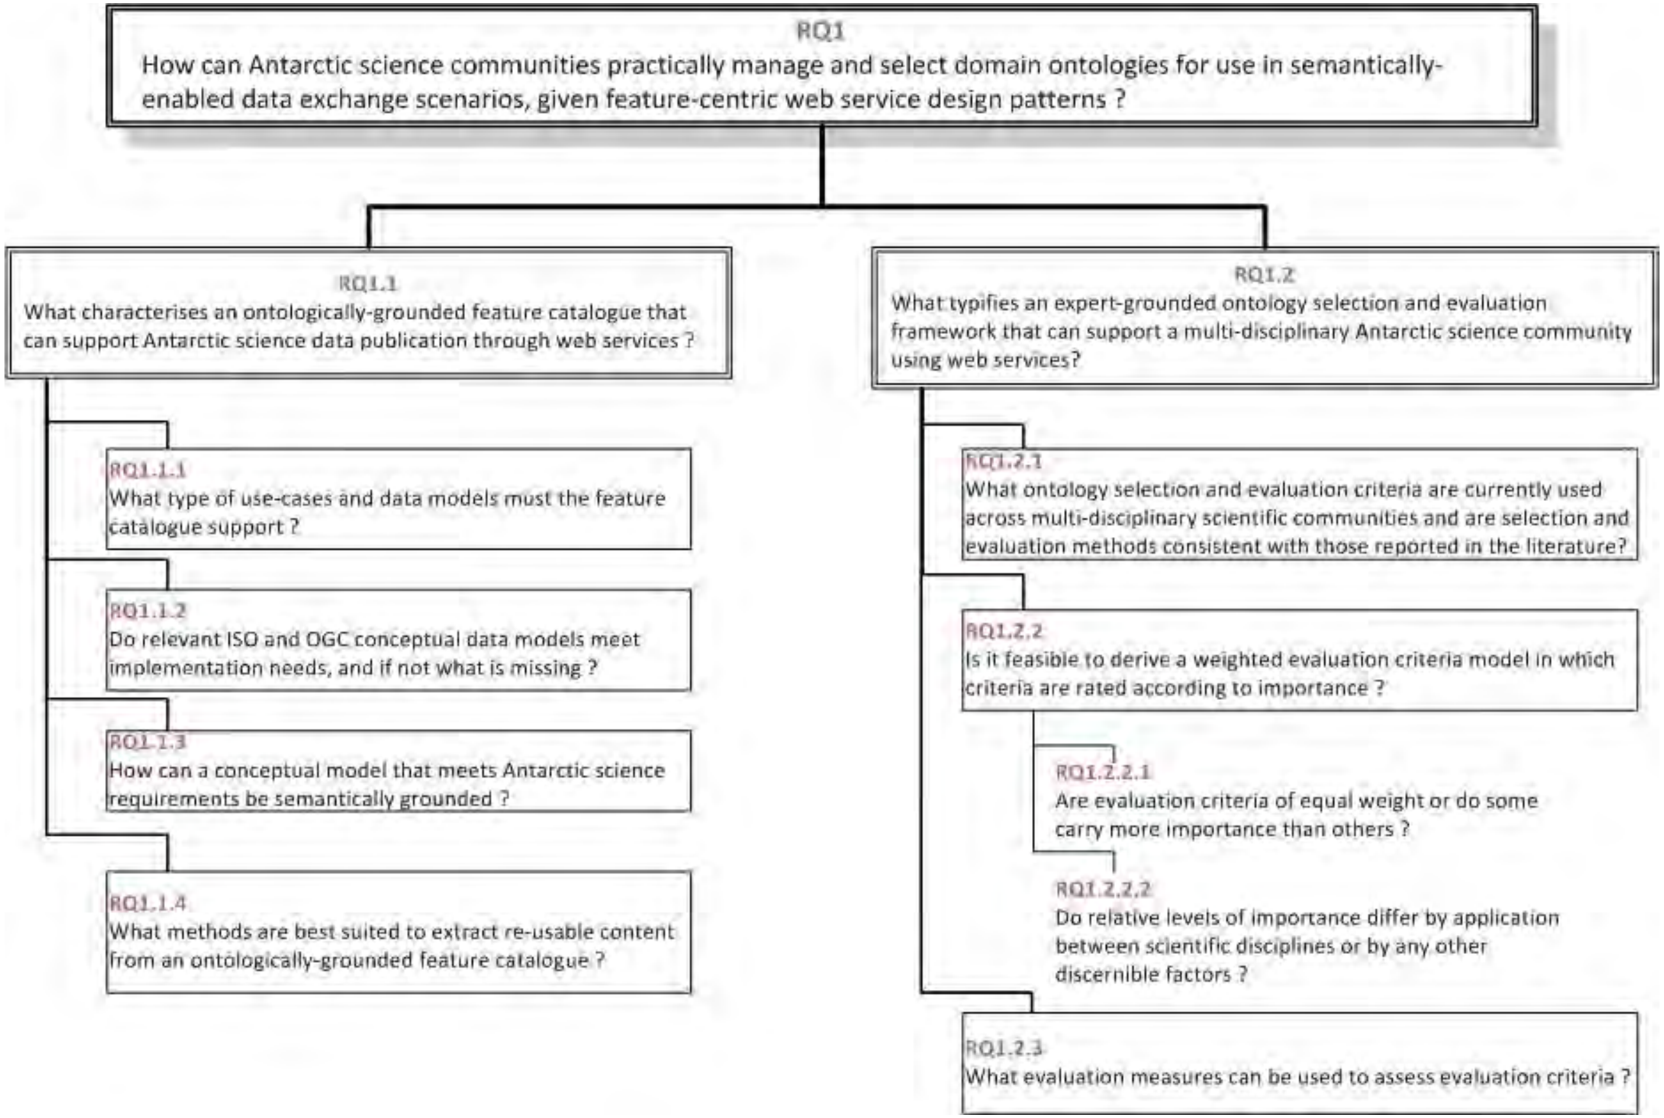

Figure 1.3 Summary of Thesis Research Questions

Similar to a full development cycle, however, prototyping activity encompassed those tasks normally found in traditional systems development initiatives, such as requirements analysis, design, programming, debugging and testing and was informed by the author's prior design experience, creativity, existing knowledge (literature), external feedback and informal discussions with community practitioners and other experts. Prototyping was particularly useful in helping to build new knowledge as specific design approaches were trialled.

The key outcome of the Design Science Research component conducted in support of this thesis was a suggested extension to the ISO 19110 (standard) Feature Catalogue Model and the casting of this extended model in a semantic context (i.e., its expression in OWL and SKOS based on the upper ontology DOLCE; Masolo *et al.,* 2003) in order to suit the semantic requirements of the Antarctic scientific community. Experimental data modelling with Antarctic-themed datasets indicated that data captured by Antarctic communities are primarily based on an observation-centric paradigm (i.e., one in which an observation act results in the estimation of the value of a feature property, and involves application of a specified procedure, such as a sensor, instrument, algorithm or process chain (Cox, 2006)). An observation entity in this thesis is considered to be a complex (or multi-featured) Feature Type. This research is considered novel in that, to the author's knowledge, no other Feature Catalogue implementation, founded on the Generalised Feature Model (ISO, 2002), has attempted to model a Feature Catalogue as an ontology.

Methods for integrating Catalogue content into existing examples of community deployed data delivery infrastructure were also demonstrated by establishing REST-based (Fielding, 2000) interfaces capable of serving Catalogue content to requesting clients, at differing levels of information granularity using URI templates (Gregorio *et al.,* 2010). The enhancements made to the ISO 19110 standard (in this thesis) to accommodate the requirements that emerged from investigating community datasets, use-cases and systems are thought to be those typically required by most multi-disciplinary scientific communities who are building infrastructure to exchange observed or measured data with spatio-temporal components. In this regard this study's enhanced Feature Catalogue model is thought to be highly generalisable and its development has identified gaps in the ISO (19110) standard when attempting to apply the Feature Catalogue Model to observation-centric data.

The research also showed how semantic annotation methods and standards can be leveraged to dynamically link descriptions of Feature Catalogue concepts to vocabulary elements in community-based metadata, service descriptions and services (encoded in XML). As the need for more automated forms of on-the-fly data integration increases, it is suggested that aspects of an observation's sampling context (e.g., sampling instrumentation, units of measure, measurement scales, sampling and measurement related datums) will need to be intrinsic components of a domain concept's semantic description.

1.2.2 Qualitative and Quantitative Research Methods

All of the research questions concerning the derivation of an ontology selection and evaluation framework, (i.e., all RQs stemming from the research question RQ1.2 – *What typifies and expert-grounded ontology selection and evaluation framework that can support a multi-disciplinary Antarctic science community using Web services ?*), have been addressed through a combination of qualitative and quantitative research methods.

Qualitative survey techniques (Denzin, 1978; Dey, 1993; King, 2004; Miles and Huberman, 1994; Jansen, 2010) were applied in this thesis to a sample of expert ontology practitioners, in order to assess the diversity of approach to the practise of ontology selection and evaluation. All experts were, or had been, actively involved in building scientific data exchange infrastructure and had practised ontology re-use. As with all qualitative studies, sampling design, data collection and data analysis techniques had to be matched to the problem domain and the goals of the research (Neuman, 1999; Jansen, 2010).

A Screening Survey, essentially a type of pilot study (Maxwell, 1992) was created as an online survey tool with a mixture of open-ended and closed questions. This Survey was developed to both identify suitable study participants and to generate an understanding of how experts viewed the meaning of some commonly used concepts that would be recur throughout the study. Answers to these types of survey questions gave early valuable insight into some perspectives that informed an expert's views on ontology selection and evaluation, expressed subsequently during in-depth interviews.

Having identified and recruited suitable study participants via the Screening Survey, each expert was interviewed for approximately one hour. Consideration was given to the extent to which the interview should be pre-structured. Structured approaches help ensure the comparability of data across sources and are useful in answering questions that deal with differences between things and the explanation for these differences (Maxwell, 2008). They can also target issues of interest. Cognisant of this it was decided that the same standardised, but open-ended questions would be asked in each interview. The purpose of these interviews were: to establish information about the expert's background and experience; become familiar with the communities with which they were involved and the ontology projects they had, or were currently supporting; understand what methods they had used to select and evaluate ontologies; to elicit opinions they had about these techniques and specifically to identify what evaluation criteria and measures they considered important/less important. Each interview was recorded; transcribed and then analysed through a process of coding and thematic (template) analysis (Crabtree and Miller 1992; Dey, 1993; Miles and Huberman 1994; King, 2004).

As a result of coding and thematic (template) analysis, a three-tiered hierarchical expert-grounded evaluation model was subsequently constructed comprising of five dimensions at the top tier (i.e., 'Structure', 'Functional Relevance', 'Usability', 'Maintenance' and 'Governance'), which decompose into thirteen sub-categories in the second tier and forty-two individual evaluation criteria at the lowest level. This model then became the basis of a quantitative method, that used an AHP-based (Saaty, 1980) pair-wise comparison exercise. In this exercise, a second questionnaire elicited weights from individual experts, for each model element (pair-wise compared). In AHP, each criterion, sub-category and dimension is assigned a rating by an expert, during pair-wise comparison, from a scale of absolute numbers (i.e., 1 to 9). This particular type of numbering scale has been proven in practice and has been validated by physical and decision problem experiments (Saaty 1980, 1994). The resultant individual preferences are then converted into ratio scale weights, framed as matrices of preferences and an Eigenvalue equation operating on these comparison matrices is used to compute estimates of the relative importance of the various model decision criteria (Genest and Zhang, 1996).

The AHP process is based on the well-defined mathematical structure of consistent matrices and their associated right eigenvector's ability to generate true or approximate weights (Merkin, 1979) and three axioms (Saaty, 1980). The first axiom, the reciprocal axiom, requires that, if $P_C(E_A,E_B)$ is a paired comparison of elements A and B with respect to their parent, element C, representing how many times more the element A possesses a property than does element B, then $P_C(E_B,E_A) = 1/P_C(E_A,E_B)$. For example, if A is 5 times larger than B, then B is one fifth as large as A. The second, or homogeneity axiom, states that the elements being compared should not differ by too much, else there will tend to be larger errors in judgment. The third axiom states that judgments about, or the priorities of, the elements in a hierarchy do not depend on lower level elements. This axiom is required for the principle of hierarchic composition to apply (Forman and Gass, 2001).

AHP is a multi-attribute decision analysis support tool and has been applied to a very wide range of decision-based problems. It has also been used previously to evaluate ontologies (Lozano-Tello and Gomez-Perez, 2004). The Ontometric technique, developed by Lozano-Tello (2002), is a formal application of AHP that uses 160 evaluation criteria for ontology assessment. This method inspired the use of AHP in this study, not for its capacity to act as an ontology evaluation method, but because of its potential use as a tool for analysing expert preferences concerning individual ontology evaluation criteria.

Information provided during in-depth interviews was also used to identify practical metrics for performing evaluation criteria assessments and these were linked to criteria in the hierarchical evaluation model. The weighted evaluation model, associated evaluation measures and some suggested methods of application, taken together constitute the expert-grounded selection and evaluation framework delivered in this thesis. Data emerging from expert interviews also provided descriptions of practise relating to methodological issues and matters associated with ontology and community governance. These are discussed and presented in detail in Chapter 7.

A limitation of the qualitative method, as applied in this research, was the relatively small expert population sample size (14 experts initially dropping down to 8 for later facets of the study) that was eventually used. Ideally, a qualitative sample should represent the diversity of the phenomenon (in this case ontology selection and evaluation approaches) under study within the target population (Jansen 2010). Whilst diversity of approach was considered adequately captured by this sample (as evidenced by later triangulation with the literature), due to participant drop-out in later stages of the research, there was little replication of expertise within the scientific disciplines covered during the pair-wise comparison exercise. This lowered confidence in some conclusions reached about patterns found in the preference data (with respect to how experts rated the importance of ontology evaluation criteria). This limitation is also discussed in detail in Chapter 7.

1.3 Reader's Guide

Figure 1.4 graphically depicts the relationship between the chapters in this thesis and outlines in skeletal form the main issues being addressed. Each chapter commences with an overview of the main issues to be discussed and ends with a summary.

In Chapter 2 the reader is introduced to various topics and literature that set the scene for the research which follows. A definition of service-oriented-architecture (SOA) environments is provided with a special focus on how the OGC has developed its Web services and reference models within the SOA paradigm. This expose is important because it helps the reader understand the role that "features" and "Feature Types" play in the Web service patterns of interest and explains how various pieces of the OGC-related technology stack are structured and inter-relate.

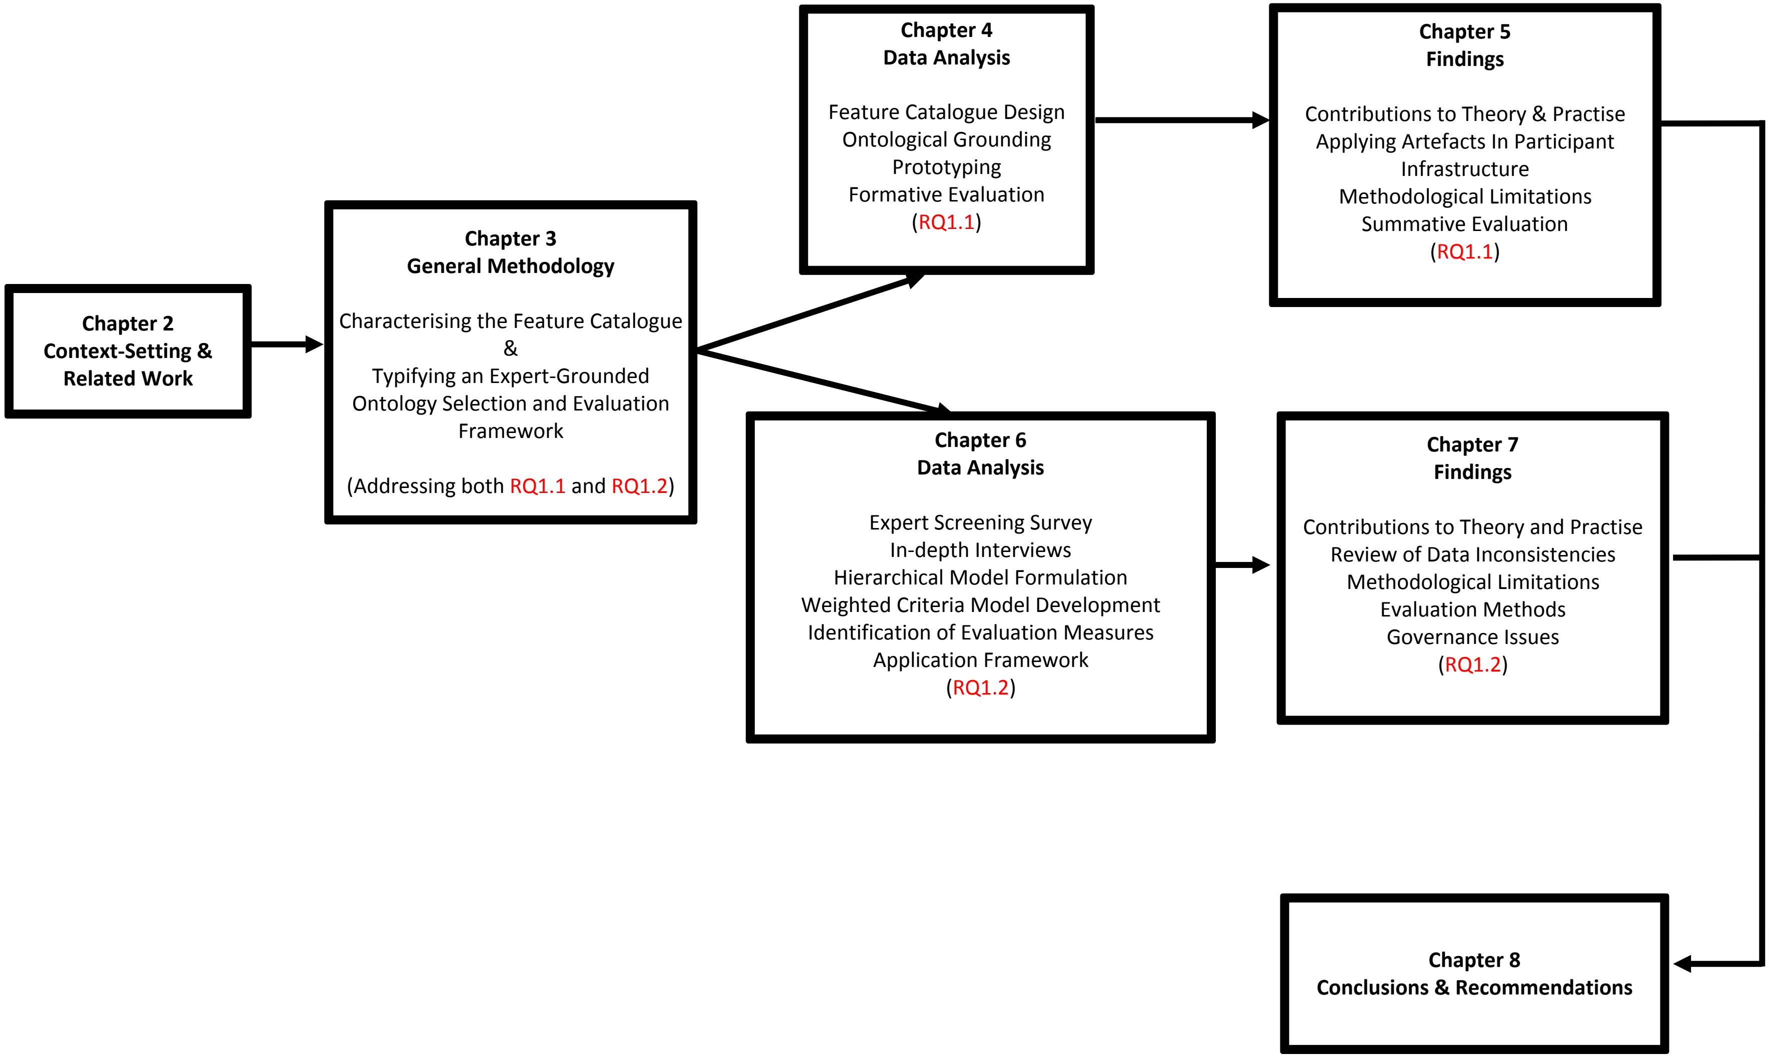

Figure 1.4 – Thesis Chapter Summary

This syntactically interoperable infrastructure is the context in which an ontologically-grounded Feature Catalogue (repository) must function. There is also a review of current practise pertaining to ontology repositories.

Also included in this chapter is an expanded description of the characteristics of ontologies and the various types of ontologies that are being created by scientific communities to support data search and data integration tasks. This description helps provide a general overview of ontologies, ontological design and ontological constructs so that ontology modelling activities presented in a later chapter, and the choices made during that modelling activity, can be followed and appreciated.

The existing state of play with respect to ontology selection and evaluation criteria, methods and metrics is canvassed and critically evaluated.

Chapter 3 covers the thesis research methods adopted in this thesis in some detail. It articulates and justifies the broad philosophical stances adopted in undertaking the study. It frames the research conducted as a mixed model of Design Science Research and other forms of qualitative and quantitative methods. This chapter then provides a general précis of the individual research methods that have been employed to address each of the specific research questions.

Chapter 4 presents the research conducted to develop a semantic Feature Catalogue that is based on the Generalised Feature Model. It introduces the reader to the two scientific collectives (the Scientific Committee on Antarctic Science and the Australian Ocean Data Network) whose data exchange requirements ground this research. Both of these collectives conduct Antarctic research and are therefore considered to be part of the Antarctic Scientific community.

Community requirements drive the Catalogue design and the derived Feature Catalogue model is ultimately instantiated as a semantic repository. This chapter explains how two types of repository are developed for comparative purposes, using the same Feature Catalogue model. One is an Oracle 11g™ semantic store and the other is built using a traditional relational database management system (RDBMS). SPARQL (Prud'hommeaux and Seaborne, 2008), an RDF query language and REST-based interfaces were trialled as access methods to Catalogue content.

The research presented is sequenced through phases of design, development and evaluation. The summative evaluation phase, however, is primarily undertaken in Chapter 5.

Chapter 5 reflects on the research results presented in Chapter 4 and discusses the role that the extended and semantically enabled Feature Catalogue model can play within observation-centric scientific data exchange infrastructures currently under development. Importantly, this chapter explains how various existing metadata and semantic service standards can be harnessed to incorporate Feature Catalogue content through the use of semantic annotation utilising the REST-based interfaces developed as part of this research.

Chapter 6 presents the results of the qualitative research methods that were used to investigate how ontology development experts select ontologies for re-use in real-world development activities. The survey instruments and the AHP analytical techniques applied to the survey data are explained in detail and a weighted ontology evaluation model is derived. Formative assessment of the model is undertaken through the development process. The weighted evaluation model is complemented with evaluation metrics and simple deployment methods, that together form the expert-grounded, practical framework for assessing ontologies for re-use.

Chapter 7 reflects on the results outlined in the previous chapter. Most experts did not harness evaluation methodologies derived from academic origins but did use many evaluation criteria that are evident in the literature. There was a high level of inconsistency detected in the pair-wise comparison preference data, both ‘within’ individual expert supplied information and that supplied ‘between’ experts. This chapter discusses possible reasons for these apparent discrepancies. The various evaluation methods used by experts are discussed and matters of governance, which emerged from interviews, is presented. Shortcomings with some aspects of the methodology used are also noted.

Chapter 8 summarises the research findings and contributions made and mentions any perceived limitations, draws together conclusions, makes a number of recommendations and suggests areas that would benefit from additional research.

Chapter 2.

Context Setting & Related Work

Sharing scientific data online requires standardisation of heterogeneous information systems such that there is interoperability between the physical (IT components) and soft (services and content) elements of the underlying infrastructure. With an adequate level of interoperability, despite a level of ongoing system heterogeneity, scientific datasets of comparable type and characterisation can be integrated. Interoperability is defined here-in “as being able to accomplish end-user applications using different types of computer and operating systems, and application software, interconnected by different types of local and wide area networks” (O'Brien and Marakas, 2007). This inter-connectedness is generally achieved by specifying and then publishing inherent communication interfaces. Wikipedia (Wikipedia, 2012b) provides a further elaboration on the concept of interoperability and qualifies the term by framing it from two different perspectives: “syntactic” and “semantic”. In systems that support spatio-temporal data, however, “framework” interoperability could also be added (Brodeur *et al.,* 2003).

In syntactically interoperable infrastructure two or more systems are capable of communicating and exchanging data. This requires specification of data formats and communication protocols. Syntactic interoperability is a pre-condition for semantic interoperability. Semantic interoperability is the ability to automatically interpret the information exchanged meaningfully and accurately in order to produce useful results as defined by the end users of both systems. As has already been foreshadowed to achieve semantic interoperability, both sides must refer to a common information exchange reference model (or ontology). The content of the information exchange requests are then unambiguously defined: what is sent is the same as what is understood (Wikipedia, 2012b).

When dealing with spatio-temporal-based data entities, concepts can be represented using different geometrical and temporal schemas (Parent *et al.*, 2006). For example, a particular instance of a glacier might be represented as a point or a polygon and it may, or may not have a temporal component (e.g., associated with its transgression or regression over time) which is portrayed using various temporal techniques. A mismatch between the spatio-temporal representation of concepts can lead to a lack of framework (or structural) interoperability.

Framework interoperability problems occur when versions of the same entities on the Earth’s surface are represented differently in terms of their data model, scale, level of abstraction, or reference framework. Scale conflicts occur when attributes of a particular entity have different units or are represented using different measuring scales (e.g., a glacier measured in km vs m, or two glaciers measured in m, but the baseline from which the measure is taken differs for both glaciers). Different precisions/resolutions can also impede interoperability when the same entity appears to have the same structural type and measurement unit but the scale at which the entity was captured varies between two datasets (Lemmens *et al.,* 2007).

Research in the field of interoperability, associated with systems that exchange data which have both spatial and temporal facets, therefore appear to cover three general (and sometimes overlapping) fields of interest (encompassing the interoperability issues mentioned above). They are those involving:

(a) solving problems that are manifest because of the abstract and conceptual frameworks that are required to represent spatio-temporal data (e.g., Walter and Fritsch 1999; Chen *et al.,* 2003);
(b) issues of semantic-enablement (e.g., Bishr, 1998; Brodeur *et al.,* 2003; Kalfoglou and Schorlemmer, 2003; Kuhn, 2003; Kavouras *et al.,* 2005; Kuhn, 2005; Lutz and Klien 2006; Klein *et al.,* 2006; Euzenat and Shvaiko, 2007; Cruz and Sunna, 2008; Schwering, 2008; Staub *et al.,* 2008; Vaccari *et al.*, 2009; Maue and Schade, 2009; Stock *et al.,* 2009a; Hossein *et al.,* 2010; Zhang *et al.*, 2010; Stock *et al.,*2010), and
(c) syntactic service composition and service orchestration (e.g., Percivall, 2002; Whiteside, 2005; Lemmens *et al.,* 2006; Wehrmann *et al.,* 2011).

Figure 2.1 depicts these three key fields of interest, further delineated into sub-categories which are then useful for partitioning the research space. In dimensionalising the interoperability problem, topics of interest in this thesis (denoted by the shaded boxes) predominantly fall under the categories of: “Semantic Data Annotation”; “Managing & Reusing Semantic Resources” and “Service Discovery”. There is also interest in “4D Framework Reference Issues” but only in so far that framework issues (such as scale, measurement, geometric and temporal representation) need to be accounted for in semantic annotation.

To improve the capacity for semantic data annotation and the deployment of semantic dataset descriptions it is necessary to review existing standards which purport to provide guidance for scientific communities wishing to semantically-enable their spatio-temporal data-centric Web services. Feature Types (or real-world concepts) and associated data Application Schema provide the main semantic currency of OGC Web service standards. Since a Feature Catalogue is the conceptual and logical entity (or repository) that the ISO and OGC standards mandate should hold the definitions for, and relationships of, Feature Types, a better understanding of a Feature Catalogue's intended role in semantic data (and potentially service) annotation is necessary, specifically within the context of the overall OGC standards framework (also known as the services stack). There is also a requirement to review current practises with respect to semantic repositories 'in general', i.e., those repositories that are not necessarily managing spatio-temporal data, but managing concepts that serve a wide range of applications and communities. Lessons learned about these non-spatial repositories may be translatable to managing concepts with spatio-temporal attributes, and of particular interest are the interfaces that can support transactions over these non-spatial types of repositories.

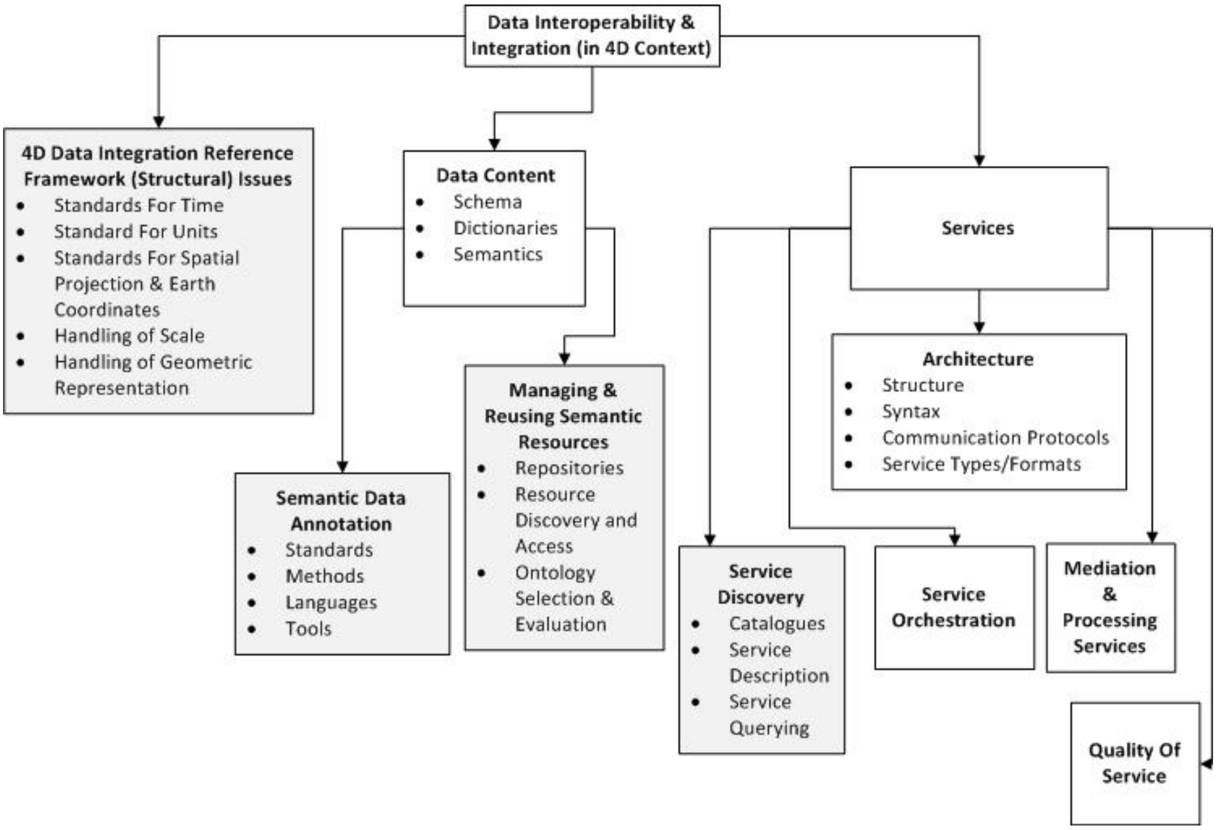

Figure 2.1 Inter-connected research space surrounding systems interoperability.

Semantic-enablement to support systems interoperability and data integration by definition implies the use of formal encodings for concept descriptions (i.e., ontologies). Deriving such formalisms requires significant community effort and skill (Blomqvist *et al.*, 2006). The re-use of semantic components is therefore advisable (Uschold *et al.*, 1998) to reduce the development overhead. The selection and evaluation techniques that are available to community practitioners for choosing these semantic components also therefore requires investigation. Given the relative paucity of semantically described data emanating from the science domain (which was an issue raised in Chapter 1), it is of interest to understand whether the literature reported selection and evaluation techniques are being used by scientific data infrastructure developers. If the reported methods are not those being used, what methods are being used in their stead ?

In this chapter the focus is on presenting existing research in areas of interest, noting gaps in current understanding that have led to the formulation of the main research question in this thesis "(RQ1) *how can Antarctic science communities practically manage and select domain ontologies for use in semantically-enabled data exchange scenarios, given feature-centric Web service design patterns"*. Given the complexity and breadth of the interoperability issues which this question spans, this chapter also provides some introductory material for the reader, so that the topics which are a focus can be appreciated relative to the overall interoperability research space (as outlined in Figure 2.1). The chapter is therefore divided into four sections. The first section gives a preliminary overview of the OGC/ISO and IT standards stack that supports spatio-temporal data transactions. Most of the OGC (as opposed to IT) standards covered in this context-setting section provide the syntactic under-pinning for interoperability. It is only recently (i.e., 2009 onwards) that the OGC has formally begun to implement standards which can cover semantic aspects of their services stack. This overview of OGC standards also highlights the fundamental importance played by 'Feature Types' as carriers of conceptual meaning within the OGC framework and explains the role that is articulated for Feature Type Catalogues. Section 2.2 introduces the reader to ontologies as a basis for formalised semantic description, describes the various types of ontologies encountered and their structural characteristics. Section 2.3 presents current progress and research into the establishment of Feature Catalogues within the science domain, noting that most semantic enablement pilot projects are not currently using Feature Type Catalogues as containers for ontological concept definitions. This section also reviews open issues in ontology repository research that are of relevance to establishing Feature Catalogues as ontology repositories. Section 2.4 presents methods and criteria for ontology selection and evaluation, highlighting the short-comings of existing approaches that may be affecting ontology take-up for use in developing semantic dataset descriptions.

2.1 OGC/ISO and IT Standards (Services) Stack

The main technological infrastructure to support Web service publication, discovery, selection and composition is based on a Service Oriented Architecture (SOA). This architecture is rapidly becoming the standard in the domain of distributed systems (Sholler, 2008). The SOA framework has been adopted by the OGC and adapted to accommodate services that transact data with spatio-temporal attributes (OGC, 2011). Recall that the OGC's remit is to deliver spatial interface and encoding specifications for the delivery of data with a spatial component.

SOAs are particular architectural patterns that involve building applications using reusable services. The services are self-contained and do not generally depend on the context or state of another service. The services communicate with each other in a distributed system architecture where the services are deployed at different locations in a network within or outside of an enterprise, and they communicate through well-established protocols. The World Wide Web Consortium (W3C) describes a Web service as: a software system identified by a URI, whose public interfaces and bindings are defined and described using XML (Champion *et al.,* 2002). In an SOA the service's definition can be discovered by other software systems. These systems may then interact with the Web service in a manner prescribed by the service's definition using XML based messages conveyed by internet protocols (Champion *et al.,* 2002). See Figure 2.2.

Messages passed between actors in the SOA paradigm can include zero or more headers in addition to data. The header part of a message can include information such as security, transaction context, orchestration, or message routing information. The data part of a message contains the message content, or data (Marks and Bell, 2006).

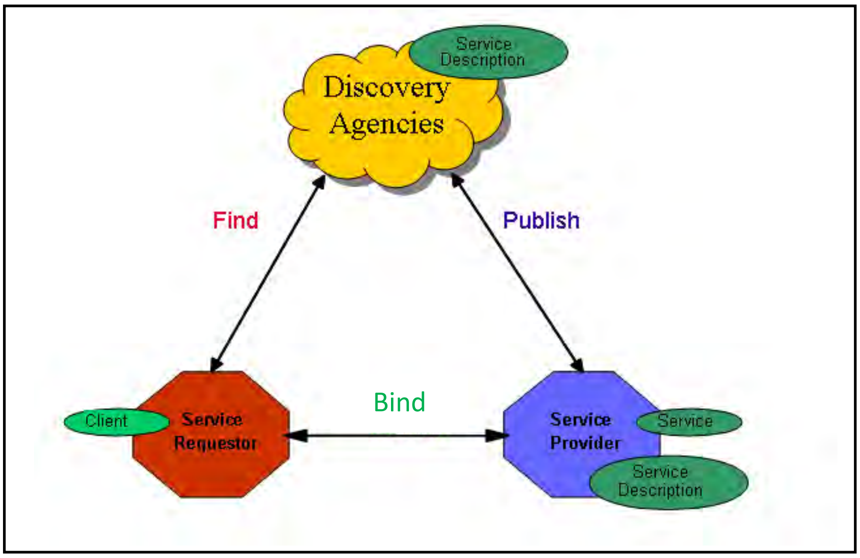

Figure 2.2 SOA Architecture (from Champion *et al.*, 2002).

A Web service is described using a standard, formal XML notation called its "service description" that provides all of the details necessary to interact with the service, including message formats (that detail the operations), transport protocols, and its location. The nature of the interface hides the implementation details of the service so that it can be used independently of the hardware or software platform on which it is implemented and independently of the programming language in which it is written (Champion *et al.*, 2002).

The SOA pattern is also often referred to as the Publish-Find-Bind paradigm. In this architecture a service provider publishes to a services registry (labelled Discovery Agencies in Figure 2.2) and also performs the role of delivering services directly to requesting clients. A service requestor attempts to find available services by interrogating a services registry which holds service descriptions. The requestor then uses the service descriptions in the registry to bind directly to the service provider's service offerings.

The OGC has developed a range of standards (OGC, 2012b) that conform to the SOA paradigm. In earth system sciences, i.e., those fields of science that contribute directly or indirectly to understanding the earth as a system, much of the data captured and exchanged has a spatial and temporal context. Many scientific communities therefore leverage the OGC standards to exchange data and it is these standards and those promulgated by the ISO 211 Technical Committee that set the context for the semantic enablement of scientific data services in this research.

2.1.1 Feature-Centricity Of OGC Services

In adopting the SOA paradigm, the OGC has taken a feature-centric information viewpoint in establishing its architectural model. A feature, as has already been discussed, is an abstraction of any real-world phenomenon. Most features transacted according to OGC standards are geographic features, i.e., they are associated with a location relative to earth (ISO, 2005a). From an OGC perspective a digital representation of the world is viewed through the lens of sets of features (Kottman and Reed, 2009).

A feature instance is a discrete representation of some specific phenomena (e.g., Botany Bay), whilst Feature Types aggregate feature instances into types or classes (e.g., Bay). "Botany Bay" is therefore an instance of a "Bay" Feature Type. It is anticipated that different communities will define sets of Feature Types (representing real-world phenomena) and within a particular community there will be only one canonical form of the Feature (type), which will have a unique identifier (Kottmann and Reed, 2009). This Feature Type may, however, be represented through different views. For example, a "Bay" Feature Type might be represented as a feature with polygonal geometry and a depth attribute, which by convention is measured at the centre of the polygon. In the same community another view of a "Bay" is that it is a 3-dimensional solid, with multiple depth attributes distributed throughout the solid. Both views are of the same Feature Type (i.e., a Bay), but as previously discussed they have a different frame of reference. The characterisation of a Feature Type provides the semantic under-pinning for OGC services (OGC, 2011).

A Feature Type Catalogue is the OGC-preferred method for registering and managing Feature Type characteristics (OGC, 2011). The conceptual model (and textual description) for such a catalogue is the ISO 19110 Geographic Information – Feature Cataloguing Methodology (ISO, 2005b). There is meant to be a close relationship between the Feature Types defined and managed within a community Feature Type Catalogue and the feature instances that are encoded in service instance documents and which are served to clients according to OGC-based Application Schema (OGC, 2011). Application Schema provide the formal description of data structure and content for OGC data services in much the same way that XML Schema provides the formal definition and description of XML instance document content (ISO 19109 (ISO, 2005c)). However, many science domain services transact feature-centric data without reference to a Feature Type Catalogue. This often leaves the intended meaning and description of the feature unclear. Additionally, most existing Feature Type Catalogue repositories have been built by communities who have traditionally been involved in surveying and mapping, utilities and transport management (as a result of using GIS technologies), rather than by communities involved in scientific data exchange. This issue and its implications will be discussed in more detail later, in section 2.3.

2.1.2 SOA-Related Metadata

Different components and activities in an OGC specified SOA require different levels of description in order for the architecture to function as envisaged. Descriptions can occur at a feature-level (i.e., at the most atomic level) within a data service, or can describe the entire dataset that is being made accessible through the service. Both of these types of descriptions are generally considered to be associated with “dataset level metadata”. Two other types of metadata are also used by the OGC standards, i.e., “high level service metadata” and “registry metadata” (Lesage, 2007). From a semantic perspective, all of these metadata variants (i.e., “feature”, “dataset”, “service” and “registry” metadata) are important because each has the capacity to reference ontological content (defined later in section 2.2). Some of this ontological content should be drawn from a Feature Type Catalogue as will be discussed and demonstrated later in Chapter 4 of this thesis.

Dataset Level Metadata

Although a feature can be annotated by metadata, more usually in current OGC-based service implementations, metadata is used to describe an entire dataset, or groups of datasets, rather than the individual features that might make up a dataset (Batcheller *et al.,* 2007). In most cases the metadata record sits separate from the data, which it then references (Batcheller *et al.,*2007). If the metadata is applied to individual features that are encoded according to Application Schema the metadata is embedded in the schema instance document (that represents the encoded data). Regardless of how metadata is applied, the OGC implements the ISO 19115 metadata abstract specification for a metadata (MD_Metadata) entity (OGC, 2011).

In reality ISO metadata is a collection of related entities (see Figure 2.3 for the UML diagram showing the various metadata components). These high level components include descriptors for information such as dataset contact details (in MD_Identification); data distribution and licensing requirements (in MD_Distribution); data quality and lineage (in DQ_DataQuality) and data maintenance (in MD_MaintenanceInformation). ISO 19115 also provisions in a few places (e.g., within MD_ContentInformation) for the identification of specific Feature Types that are inherent in a dataset and can also encompass references to Feature Type Catalogues that are the source of Feature Type definitions (ISO, 2003). These latter facilities are not well exercised currently in the science domain and there is conflicting guidance about how to use the available Feature Type semantic provisions in the ISO 19115 metadata (Maue, 2009).

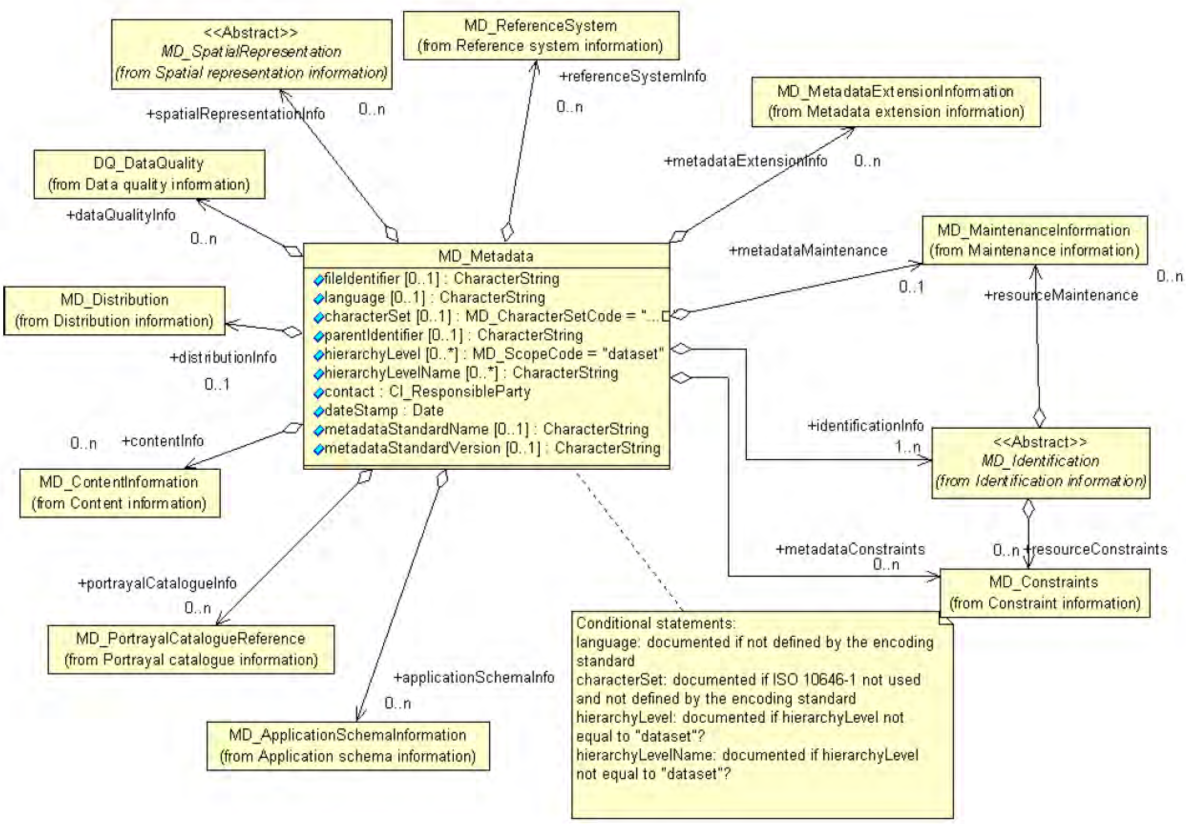

Figure 2.3 ISO 19115 Metadata Conceptual Schema (from Figure 14: Percivall, (2003))

Service Level Metadata

As explained earlier, all services deployed in a SOA must carry a service description (i.e., services metadata). Services that conform to OGC standards describe themselves using high level 'Capability Documents' which are accessed by clients issuing a "GetCapabilities" service request (Whiteside and Greenwood, 2010). Service 'Capability Documents' describe the service interface in sufficient detail so that an automated process can read the description and invoke an operation that the service advertises. It also describes the data content of the service (or the data it operates on) in a way that enables service requestors to dynamically compose requests for service. This content description component is optional, depending on whether the service contains or operates on data. Additional description units provide information specific to particular types of services as well as specific instances of services (Percivall, 2003). As is the case for dataset level metadata, there is a lack of guidance in the OGC standards about where Feature Type information should be stored in service descriptions.

It should be noted here that outside of the OGC standards environment, the Web Services Description Language (WSDL: Christensen *et al.,* 2001) is commonly used by the software industry in conjunction with the Simple Object Access Protocol (SOAP: Gudgin *et al.*, 2007) and XML Schema to describe and orchestrate Web services, instead of Capability Documents. Whilst WSDL is not a mandatory component of OGC service descriptions, the OGC services suite now also accommodates the inclusion of WSDL and SOAP (Whiteside and Greenwood, 2010).

Registry Level Metadata

Registry metadata on the other hand is used to describe "registry objects" and how they are organised in the (registry) repository. Registry objects are any registered resources. The OGC information model for defining this metadata is a profile of the OASIS (2002) ebXML registry information model (ebRIM). This registry profile, coupled with the OGC Catalogue Services Specification (Nebert *et al.,* 2007), specifies the information content and the framework for developing interfaces and the service bindings required to publish and access digital catalogues (registries) of metadata. Implementations of the OGC Registry (Catalogue) services are often referred to as OGC Catalogue Services for the Web (CSW: Nebert *et al.,* 2007). A registry is typically maintained by an authorized registration authority who assumes responsibility for complying with a set of policies and procedures for accessing and managing registry content (ISO 11179-6 (ISO, 2005d)).

The application interface of the OGC CSW(ebRIM) profile supports multiple query patterns including browse and drill-down (by category), or filtered queries against specified registry objects. Service offerings are the main resources managed in the registry and these are associated with other objects that help to provide a full and flexible description of the service (e.g., service taxonomies that classify the service in order to distinguish different service types, interface definitions, dataset descriptions and application schemas). Arbitrary relationships among catalogued items can be expressed by creating links between any two resource descriptions. MD_Metadata type entities are permissible and routinely included registry objects (Martell, 2007). There is a specific CSW profile, called the "ISO Metadata Application Profile" (Voges and Senkler, 2007) which uses ISO 19115 metadata. Not surprisingly, and as for the other levels of metadata, the encoding of Feature Type information in registry metadata is also lacking a standardised approach.

2.1.3 OGC Service Type Standards

Having introduced the main types of descriptive metadata and the points within an OGC-influenced SOA where metadata are generally used, an overview is provided of the OGC service types that are routinely deployed by scientific communities to exchange data. How these services function, further serves to highlight the feature-centricity of OGC services. There are four main service types: Web Mapping Services (WMS: de la Beaujardiere, 2006); Web Feature Services (WFS: Vretanos, 2005); Web Coverage Services (WCS:Baumann, 2010) and a Sensor Observation Services (SOS: Na *et al.,* 2007).

The simplest of the four is the WMS which standardises the way a client can publish and request maps (i.e., server-rendered pictures). Maps are requested through a WMS by naming a "map layer" that is made available, as a response via a particular WMS instance document. It is also possible, however, to ascribe information to individually identified 'features' (delineated via coordinates) within the map and a client can also issue a query to access any assigned feature-specific information. The main output of the service, however, is an image. The remainder of the services to be described are concerned with the delivery of data (as opposed to images).

A WFS is transactional in that it supports the "insert", "update", "delete", "query" and "discovery" of features. The payload of a WFS are data and metadata encoded in Geography Mark-Up Language (GML: Portele, 2007). GML is an XML grammar written in XML Schema for the description of Application Schemas and is used for the transport and storage of geographic information (in instance documents). The main operations of non-transactional WFS' (i.e., those not allowing update and delete) are invoked using “GetCapabilities”, “DescribeFeatureType” and “GetFeature” interfaces (emphasising the fundamental position that the concept of a Feature Type occupies in OGC services).

A WCS supports the transfer of a feature’s properties at specific geo-locations. Data transferred using this service resembles map layer data except that individual values for each data point at each geographic location is transferred, rather than a simple ‘picture-based’ rendering.
Datasets transferred via WCS are usually referred to as coverages. This type of service is routinely used to transfer data often used in scientific modelling scenarios (e.g., regular and irregular gridded data). Its services are invoked using “GetCapabilities”, “GetCoverage” and “DescribeCoverageType” interfaces. In OGC parlance, a “Coverage” is a special type of feature (ISO 19123 (ISO, 2005e)).

SOS provides an interface for managing deployed sensors and for retrieving sensor data and specifically “observation” data. Whether from in-situ sensors (e.g., water monitoring) or dynamic sensors (e.g., satellite imaging), measurements made from sensor systems contribute most of the geospatial data by volume used in geospatial systems today (Na and Priest, 2007). Similar to a WFS, a SOS can be transactional. The mandatory core operations are “GetCapabilities”, “DescribeSensor” and “GetObservation”. An observation is an event involving a feature of interest (Cox, 2006).

An SOS organises collections of related sensor system observations into Observation Offerings. An Observation Offering is analogous to a “layer” in a Web Map Service because each offering is typically a non-overlapping group of related observations. Each Observation Offering is constrained by a number of parameters including the following:

- specific sensor systems that report the observations;
- time period(s) for which observations may be requested (supports historical data);
- phenomena that are being sensed;
- geographical region that contains the sensors, and
- geographical region that contains the features that are the subject of the sensor observations (may differ from the sensor region for remote sensors).

The SOS is one of a family of bundled standards that make up the OGC activity called “Sensor Web Enablement” (or SWE). Currently, the other ratified specifications that pertain to SWE are Sensor Model Language (SensorML: Botts, 2007), Observations and Measurements (O&M: Cox, 2006), Sensor Planning Service (SPS: Simonis and Echterhoff, 2011), and the Transducer Markup Language (TML: Havens, 2006).

2.1.4 Semantic Service Description and Orchestration

Using any of the service standards described above (in section 2.1.3) it is not possible to determine the exact semantics (meaning) of the Web service's function, or the data it uses, and it is not possible to use the information that is available to automate the process of Web service discovery, execution and orchestration (Stock *et al.,* 2011). Service orchestration relates to the execution and mediation of specific business processes. For example, orchestration might encompass integration of the output of an OGC WFS service and an OGC SOS, or involve chaining these services together end-to-end so that output from one becomes input to the other.

In order to formally define the semantics of Web services, ontologies and ontology languages must be used. An introduction to ontologies follows shortly in section 2.2. However, a detailed understanding of ontologies is not required for the reader to appreciate where ontologies need to play a role within the SOAs that are used by groups such as the OGC, nor is it required to comprehend that there is often more than one way to introduce ontologies into the architecture. Figure 2.4 summarises the major functions that must be performed in an idealised semantic Web services framework and shows the relationship of ontologies to these functions.

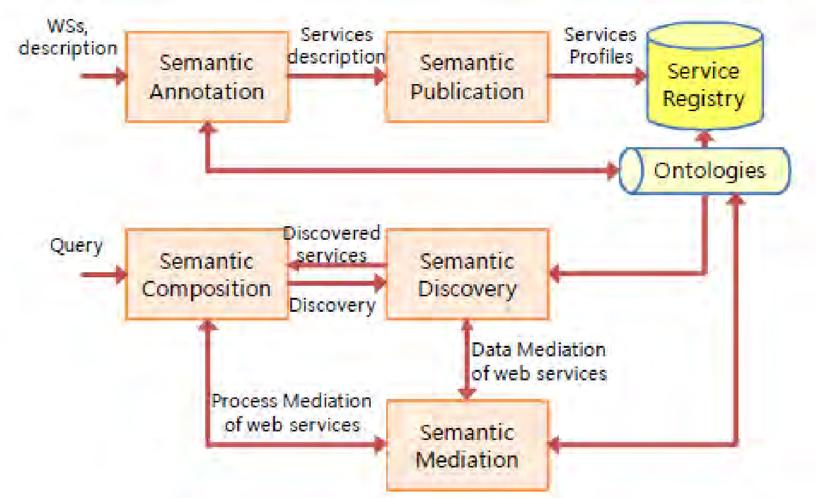

Figure 2.4 Elements Of A Semantic Web Services Framework (from Fig.1. Yoo *et al.*, 2010)

Recognising the desirability of formalising the semantics of registry metadata there are now draft OWL-based profiles of the OGC CSW available (Stock, 2009b; Dogac *et al.*, 2010). This new ability to include OWL ontologies in OGC geospatial registries has the benefit of enabling representation of richer semantic information to describe resources that will ultimately assist with discovery tasks and provide opportunities for Web services orchestration (Yue *et al.,* 2010). This development is particularly important given broader IT industry efforts to develop OWL-based ontologies and mark-up languages for the task of semantic Web service description. In OGC services a Feature Type Catalogue should be able to function as a source of ontologies for both registry metadata and direct dataset annotation in Web services (as will be shown later in Chapter 5 of this thesis).

The most often mentioned IT industry semantic Web service description specifications are OWL-S (Martin *et al.*, 2004), Web Service Modelling Ontology (WSMO: Roman *et al.*, 2004) and Semantic Annotation For WSDL (SAWSDL: Farrell and Lausen, 2007). Note that these types of service description and orchestration languages go beyond lexical and syntactic service description (which were canvassed earlier in the form of Capability Documents and WSDL), because of the inclusion of ontologies.

OWL-S supplies Web service providers with a core set of mark-up language constructs for describing the properties and capabilities of their Web services in unambiguous, computer-interpretable form (Stock *et al.,* 2011). It includes the semantics of Web service behaviour and the semantics of the static information objects with which the service interacts. OWL-S mark-up of Web services is intended to facilitate the automation of Web service tasks including automated Web service discovery, execution, interoperation, composition and execution monitoring (Martin *et al.*, 2004). It is an ontology of service concepts based on OWL that breaks the service description into three areas (each described by a different sub-ontology): a 'process model', a 'profile' and 'grounding'. The service 'process model' describes how a service performs its tasks. It includes information about service inputs, outputs, preconditions, and results. The service 'profile' is related to service annotation and provides a general description of a Web service that is intended to be published and shared to facilitate service discovery. This (profile) component also includes information such as a service name, contact information, service category, service classification, service product, and a textual description. The service 'grounding' ontology provides the details of how to access the services, by mapping to message parts of WSDL (Yoo *et al.*, 2010).

Stock *et al.* (2011) have experimented with using OWL-S Web service ontologies to represent OGC-compliant geospatial Web services by automatically populating a Web service ontology from OGC 'GetCapabilities' documents. Their experience revealed serious short-comings in how OWL-S can be applied to OGC service descriptions. OWL-S was found to be a cumbersome way to describe Web services. In particular, there was significant repetition required to fully express the elements of the 'GetCapabilities' document in OWL-S. This occurred particularly in the definition of parameters. Input and output parameters had to be defined in three places: for the 'processes' that the Web service operations implement; for the 'grounding' of that Web service and finally in the 'profile' for the Web services operation. It was also difficult to describe some of the cardinality constraints associated with input and output parameters using OWL-S (Stock *et al.,* 2011). Interestingly, Stock *et al.,* (2011) did, however, find that the OWL-S ontology was useful for describing the detail of the data made available by 'WFS' and 'WMS' (i.e., the Feature Types and map layers respectively).

WSMO is a major competitor to OWL-S for the description of Web service semantics. Lara *et al.* (2004) claim that in comparison, WSMO is defined more precisely and unambiguously, and it provides a more complete framework for aspects and challenges arising within Semantic Web Services. It is built using the Web Services Modelling Framework (WSMF: Fensel and Bussler, 2002). WSMF consists of four different elements for describing semantic Web services: ontologies that provide the terminology used by other elements; goals that state the intentions that should be solved by Web services; Web services descriptions that define various aspects of a Web service, and mediators which resolve interoperability problems. In this regard, Lara *et al.* (2004) claim that WSMO provides the ontological specifications for the core elements of Semantic Web services.

SAWSDL is yet another service description offering and is a set of extensions for WSDL. SAWSDL defines an annotation mechanism for specifying the data mapping of XML Schema types to, and from an ontology. To accomplish semantic annotation, SAWSDL defines extension attributes that can be applied both to WSDL elements and to XML Schema elements. The extensions take two formats: 'modelReferences' that point outwards to semantic concepts and 'schemaMappings' that specify data transformations between a message's XML data structure and the associated ontology model. SAWSDL is often used in conjunction with WSMO. A 'Lifting Schema' mapping is called the up-cast mapping and transforms XML data from a Web service into a concept of an ontology. A 'Lowering Schema' mapping is called a down-cast mapping and it transforms a concept of an ontology into XML data (Farrell and Lausen, 2007).

The use of OWL-S, WSMO and SAWSDL is in its infancy with respect to the roles that they can play in deployed OGC services and all are in the process of being trialled through various pilots and test-bed implementations (Lieberman, 2006; Stock *et al.*, 2009a). Later, in Chapter 5, it will be demonstrated that SAWSDL can be most effectively used to map Feature Type Catalogue content into data services, since it is relatively easy to embed 'modelReferences' (which point to external ontologies) in encoded dataset documents and the more complex issues of service orchestration (handled by OWL-S and WSMO) are not of direct interest in this thesis.

Before leaving the topic of semantic service description and orchestration, the OGC's own service orchestration specification should be mentioned for completeness purposes only, i.e., the Web Processing Service (WPS: Schut and Whiteside, 2005). This is essentially a non-semantic interface specification. A WPS provides client access across a network to pre-programmed calculations and/or computation models that operate on spatially referenced data. The calculation can be extremely simple or highly complex, with any number of data inputs and outputs. The specification does not identify the individual processes that could be implemented by a WPS. Instead, it specifies a generic mechanism that can be used to describe and web-enable any sort of geospatial process (Schut, 2007).

2.2 Ontologies

It has been consistently claimed throughout the preceding section that the ‘metadata’ created to describe certain facets of data and services should include the use of ontologies. An explanation is therefore required as to what an ontology is and how an ontology helps with machine-mediation of data exchange activities.

The word “Ontology” was taken from the discipline of Philosophy and dates back to the ancient Greeks (in the 4^{th} and 5^{th} centuries BC). Ontology is the philosophy of ‘being”. It is concerned with identifying things and asking questions about the “essence” that remains inside things even when they change (e.g., when they change colour or size). It also asks whether things really exist outside of our mind and how we can classify them. The art and practise of ontological engineering emerged much later in the late 20^{th} and early 21^{st} centuries when it became a research area in computer science (particularly in the Artificial Intelligence (AI) field). In AI and ontological engineering, research focuses on activities associated with developing ontologies and the methodologies, tools and languages that could be used for these activities (Gomez-Perez *et al.,* 2004).

2.2.1 Characteristics Of An Ontology

There are many definitions for the term “ontology” and one version cited is the description already given in the introductory chapter by Guarino (1998). But the most often quoted definition is that of Gruber (1993), who coined an ontology as a “formal specification of a conceptualisation”. In plain language, an ontology specification is a formally described, machine-readable collection of terms and their relationships expressed with a language in a document file. A “conceptualisation” refers to an abstract model of a domain that identifies concepts and their relationships (Guarino and Giaretta, 1995). A domain is a specific subject area (or areas) of knowledge that is typically the focus of a particular community-of-interest.

Formal ontologies are sometimes divided into two parts called Terminological and Assertion components. A Terminological component or “TBox” is in the form of a terminology and is built through declarations that describe general properties of concepts. An Assertion component or "ABox" contains the extensional asserted facts, or ontological instances/individuals (Gomez-Perez *et al*. 2004). Taken together TBox and ABox allow one to represent a domain of interest in terms of *concepts* and *roles,* where concepts model classes of individuals, and roles model relationships between classes. Abox components are also referred to as knowledge-bases. The TBox contains the definitions of concepts and roles, while the ABox contains the definition of individuals (Baader *et al.,* 2003).

In many respects an ontology (TBox and ABox) functions similarly to a populated database schema. The Tbox is similar to the database schema, where there are entities represented through tables, and relationships are created between tables and their respective attributes, via foreign keys. Filling the database schema with instance data creates something approaching an Abox knowledgebase. The differences between a formal ontology and a database, however lies in how the respective conceptual models are created, in how instance data are handled in both technologies, in how the data can be queried and most importantly, ontologies can be reasoned over (this is not possible with a traditional Relational Database Management System - RDBMS). "Reasoned over", means deriving facts that are not explicitly expressed in an ontology through the process of logical inference.

With respect to logical inference, database schema are said to conform to a closed-world assumption, whilst ontologies operate according to an open-world assumption. In closed-worlds everything we don't know is deemed to be false, while in an open-world assumption the lack of knowledge does not imply falsity (it may be true until we find some reason to prove that it isn't true) (Motik *et al.,* 2006). In a database if facts are not asserted (i.e., are absent) the inference is that they don't exist.

The TBox/ABox conceptualisation is the foundation for a type of logical formalism called Description Logics (DL). In this system ontologies are represented by three things: concepts, roles and individuals. Concepts represent classes of objects with similar characteristics; roles describe binary relations between concepts, which allows for the description of properties of concepts; and individuals represent instances of concepts. Concepts and roles are described with terminological descriptions, which are built from pre-existing terms and with a set of constructors (Gomez-Perez *et al.,* 2004). Constructors describe the types of allowable relationships between concepts, roles and individuals (see Figure 2.5). Roles are also sometimes referred to as properties. The expressivity of a particular DL ontology can be gauged by the number, type and complexity of constructors it implements (Kepler *et al.*, 2006).

The smallest unit of knowledge within an ontology is referred to as an axiom (Vrandecic , 2010). Axioms are used to associate (concept) class and property (role) IDs with either partial or complete specifications of their characteristics, and to give other logical information about classes and properties. Staab and Maedche (2000) present a classification of the main types of axioms commonly found in ontologies. The characterisation by Vrandecic (2010) however is used here and it groups axioms into those that are classed as a 'Fact', 'Terminological', or as an 'Annotation'. Facts can be instantiations where an individual is asserted to be of a particular type of concept. Facts may also be property (role) assertions, where an individual property **P** relates two individuals *a* and *b*.

Terminological axioms cover semantics associated with properties and classes. These axioms can express equivalence, disjointedness or subsumption. For example, it is useful to be able to express that one class is equivalent to another (i.e., they are exactly the same thing). Contrastingly, an ontology may assert that a class is the opposite, or disjoint to another class. In such cases one individual cannot simultaneously be an individual of two classes that have been declared disjoint (Kepler *et al.,* 2006).

Symbol		Name
$\sqcup$	$\mathcal{U}$	Union
$\sqcap$	$\mathcal{AL}$	Intersection
$\neg$	$\mathcal{C}$	Concept negation (complement)
	$\mathcal{A}$	Atomic negation
$\forall$	$\mathcal{AL}$	Value restriction
$\exists$	$\mathcal{E}$	Existential quantification
$\geq, \leq, =$	$\mathcal{Q}$	Qualified number restriction
	$\mathcal{F}$	Functional number restriction
	$\mathcal{N}$	Unqualified number restriction
I	$\mathcal{O}$	Nominals (instance enumeration)
	$\mathcal{I}$	Inverse roles
	$\mathcal{H}$	Role hierarchies
	$\mathcal{D}$	Data type restriction

Figure 2.5 Common Types Of Constructors (From Table 1. Kepler *et al.,* 2006).

Taxonomic hierarchies of concepts (which are also Terminological class axioms) allow for the description of specialisation and generalisation relationships. These relationships are usually expressed through an 'is_a' construct implying inheritance, subsumption, generalisation and specialisation type relations (Horridge, 2011). Declaring a concept to be a 'subclass' through an 'is_a' relation is called subsumption and means that all instances of a sub-class can be inferred to be instances of its super-class. Concepts can be defined through simple subsumption (e.g., *ShallowBay* is a subclass of *Bay*) or through more complex expressions. A simplified example of a more complex expression is one where a concept such as '*Estuary'* is defined as the 'equivalent class' of two concepts, ‘*ShallowBay’* and ‘*RiverOutlet’.* This would be achieved by adding a property restriction on ‘*Estuary*’ such that every ‘*Estuary*’ instance must have at least one member from ‘ShallowBay’ and ‘*RiverOutlet*’. By inference all instances of ‘Estuary’ are also instances of ‘*Marine GeomorphologicalFeature’*. Specialisation occurs as concepts are subclassed. The most generalised concepts are those at the top of an “is_a” concept hierarchy.

Terminological property axioms on the other hand usually define the relationship between two individuals (or the relation between an individual and a data value). They can also define a property type, or define a property’s domain and/or range (Horridge, 2011). A property’s domain limits the individuals to which the property can be applied. It is the class of individuals that are the subject of the relation (expressed through the property) and the range of a property limits the individuals that the property may have as its value. For example, in the ontological (tuple) statement “*Bay* **hasShorelineCircumference** *Circumference”,* the class *‘Bay’* is the domain of the property named “**hasShorelineCircumference”** and the range of the property is the class named ‘*Circumference’.* Properties can be of several types including: functional, inverse functional, reflexive, irreflexive, symmetric, asymmetric and transitive. See Table 2.1 for brief definitions of each property type.

Table 2.1. Property Types (collated from Lacy, 2005)

Property Type	Definition
Functional	Can have at most one (unique) value for a particular subject individual. If a property is a FunctionalProperty, then it has no more than one value for each individual (it may have no values for an individual). FunctionalProperty is shorthand for stating that the property's minimum cardinality is zero and its maximum cardinality is 1. For example, hasPrimaryEmployer may be stated to be a FunctionalProperty. From this a reasoner may deduce that no individual may have more than one primary employer. This does not imply that every Person must have at least one primary employer however.
Inverse functional	This is the opposite of a Functional property. It is like a foreign key field in a RDBMS. It is used to identify properties whose values uniquely identify the subject instance of the property. If you know the value of the inverse functional property, you will know which subject it belongs to. If a reasoner identifies two URI resources and both have the same value for an inverse functional property, the reasoner can infer that the two URIs are in fact referencing the same individual.
Reflexive	Is a property that defines a relationship where all elements in the tuple are related to themselves and where the property holds “true”. To understand this, assume the subject and object are the same real number *X*. Since every real number is equal to itself, “*X* **is equal to** *X”* holds to be true.

Irreflexive	Is the opposite of a reflexive relation. It is a binary relation on a set where no element is related to itself. Take for example the real numbers *X* , using the property "unequal to", it can be seen that the tuple will be false. This is because no real number is unequal to itself (e.g., "*X* **is unequal to** *X*" is false but "*X* **is unequal to** *Y*" is true).
Symmetric	A symmetric property implies a relationship exists in two directions. These properties are used for inference short-cuts. This means that "x aProperty y" implies "y aProperty x". Instead of having to specify that a property holds in both directions between a subject and an object (by declaring for example "*individual A* **P** *individual B*" and then "*individual B* **P** *individual A*"). A statement can be made about the subject along the property to the object once (e.g., *individual A* **P** *individual B)* and a reasoner can infer that the relationship also holds in the opposite direction.
Asymmetric	A type of property that defines an asymmetric relationship. In this sense an asymmetric relation is a binary relation which is not a symmetric relation. A property **P** with subject *A* and object *B* is asymmetric if *A* is said to be related to *B* through **P** but *B* is not related to *A* (i.e., the relation does not hold in both directions).
Transitive	The subjects and values of transitive properties can be chained together. They are commonly used in part/whole relationships. If a property **P** is transitive and you have a tuple "*individual A* **P** *individual B*" and you also have a tuple *"individual B* **P** *individual C"* it is possible to infer that "*individual A* **P** *individual C*" (i.e., *individual A* also has the *individual C* in relation to property **P** even though this was not explicitly declared).

Annotation axioms consist of an annotation property whose domain can be a concept, another property, an individual or an ontology. The range of the annotation property is a data literal or a URI, which is called the annotation value (Lacy, 2005). This is a special type of property since it has no impact on how a reasoner parses an ontology. It is a highly useful property never-the-less because it allows for the attachment of metadata (such as when a concept or property was created, what version an ontology is up to, who created an ontology concept, property or individual etc).

Figure 2.6 diagrammatically demonstrates some of the ontology characterisations discussed in the preceding paragraphs.

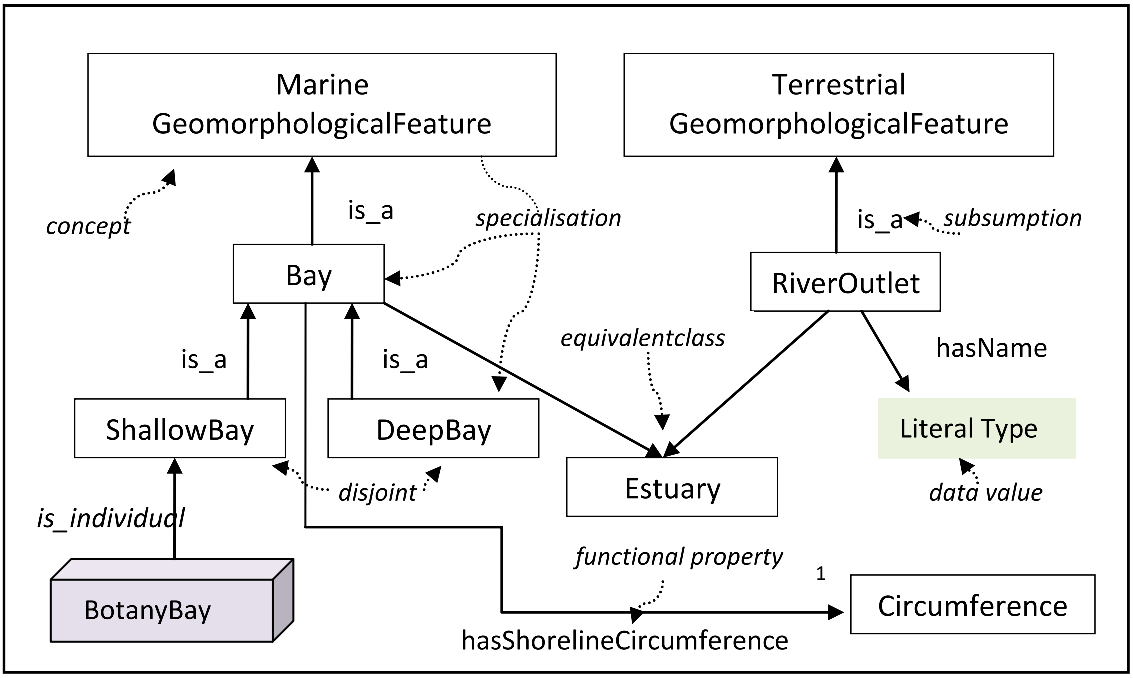

Figure 2.6 Elements Of A Hypothetical Ontology

In DL ontology systems, reasoning is mainly based on concept subsumption (but depends on the constructs used in the supporting language). Reasoning is used to determine whether an ontological description is satisfiable (non-contradictory) and minimally redundant. In other words reasoning derives whether an ontology has a consistent model and whether it entails that a particular individual is an instance of a given concept description (Baader *et al.,* 2004).

2.2.2 OWL Ontologies

Ontologies have been implemented in a variety of languages (e.g., CycL (Lenat and Guha, 1990), Knowledge Interchange Format (Genesereth and Fikes, 1992), RDFS (Brickley and Guha, 2004), OIL (Horrocks *et al.*, 2000) and DAML+OIL (Horrocks , 2002)). The expressivity of these languages determines the formalism of the ontology. Today the most commonly used (formally expressive) language for Web-deployed ontologies is the Web Ontology Language or OWL (Dean and Schreiber, 2003). Since this thesis is about the Semantic Web, descriptions concerning ontologies from this point on, that necessitate discussion of any language-specific constructs, will be confined to ontologies constructed in OWL. OWL has more facilities for expressing meaning and semantics than other Web-centric languages such as XML, RDF (Manola and Miller, 2004), and RDF-S and therefore OWL goes beyond these languages in its ability to represent machine interpretable content on the Web.

There are three sub-languages of OWL 1: OWL Lite, OWL DL and OWL Full. OWL Lite supports those users primarily needing a classification (is_a) hierarchy and simple constraints. For example, while it supports property cardinality constraints, it only permits cardinality values of '0' or '1' and has no provision for explicit negation or union constructors. OWL DL supports those users who want the maximum expressiveness while retaining computational completeness (i.e., all conclusions are guaranteed to be computable) and decidability (all computations will finish in finite time). OWL DL is the most popular language variant because of its reasoning capability. It includes all of the OWL language constructs, but they can be used only under certain restrictions. OWL Full provides users with no constraints on how vocabularies are used except that they must be legal RDF. OWL Full is not used for reasoning (McGuinness and Harmelen, 2004).

OWL 2 was released in 2009 and it is backwards compatable with OWL version 1, with a number of additional features. It has increased expressive power for properties (e.g., both qualified and unqualified cardinality restrictions (see Table 2.2 for main OWL restriction types) and support for reflexive, irreflexive and asymmetric property types). Extended support for datatypes, simple metamodeling capabilities, extended annotation capabilities, and keys has also been provided (Golbreich and Wallace, 2009). Keys enable the unique identification of individuals of a given class by values of (a set of) key properties. In OWL 1 a strict separation was required between the names of classes and individuals but OWL 2 DL relaxes this separation to allow different uses of the same term for both a class and an individual. This is called punning. OWL 2 also defines several profiles which are OWL 2 language subsets that may better meet certain performance requirements of users, or which may be deemed easier to implement (Golbreich and Wallace, 2009).

Table 2.2. OWL Restriction Types (derived from Golbreich and Wallace, 2009).

Notation	Term	Meaning
$\exists$	Existential, someValuesFrom	"Some" "At least one"
$\forall$	Universal, allValuesFrom	"Only"
$\ni$	hasValue	"Equals x"
$=$	Cardinality	"Exactly n"
$\geq$	MaxCardinality	"At most n"
$\leq$	MinCardinality	"At least n"
	QualifiedCardinality	"Restrains the *class* or *data range* of the instances to be counted e.g., for specifying the class of persons that have at least three children who are

		girls”
	QualifiedMaxCardinality	“Same as for QualifiedCardinality and allows for the assertion of maximum qualified cardinality restrictions.”
	QualifiedMinCardinality	“Same as for QualifiedCardinality and allows for the assertion of minimum qualified cardinality restrictions.”

In OWL 2 classes (concepts), datatypes, object properties, data properties, annotation properties, and named individuals are entities, and they are all uniquely identified by a Unique Resource Identifier (URI). Classes represent sets of individuals; datatypes are sets of literals such as strings or integers; object and data properties can be used to represent relationships in the domain; annotation properties can be used to associate non logical (i.e., properties not used by reasoning engines) information with ontologies, axioms, and entities; and named individuals can be used to represent actual objects from the domain. The language also provides for literals, which consist of a string called a “lexical form” and a datatype specifying how to interpret this string (Motik *et al.,* 2009) See Figure 2.7.

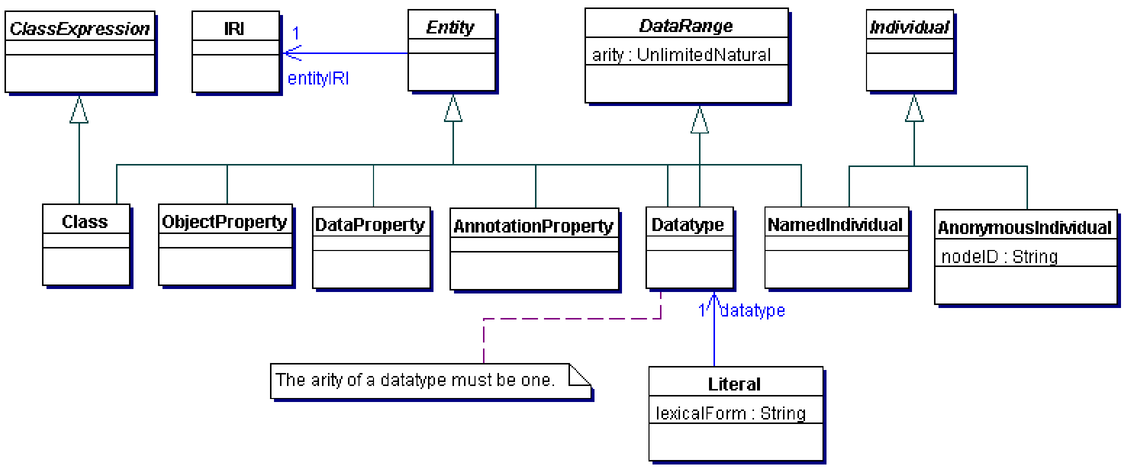

Figure 2.7 Components of OWL2 (from Figure 2; Motik *et al.*, 2009)

Some examples of OWL2 functional syntax (Motik *et al.*, 2009) for describing axioms are given below:

(a) Describing the class *Bay* as a subclass of *TerrestrialGeomorphologicalFeature* from the domain called “examples” (abbreviated to ex) can be expressed as:

SubClassOf(ex: *Bay* ex:*TerrestrialGeomorphologicalFeature*)

(b) An ObjectProperty can be used to associate two classes expressed as:

ObjectPropertyAssertion(ex: **hasShorelineCircumference** ex: *Bay* ex:*Circumference*) which states *Bay* **hasShorelineCircumference** of *Circumference*.

(c) A DataProperty ex:**hasName** has a literal as its range, expressed as:

DataPropertyRange(ex: **hasName** rdfs:Literal)

(d) *BotanyBay* can be asserted as an individual of a *Bay*, expressed as:

ClassAssertion(ex: *Bay* ex:*BotanyBay*)

OWL functional syntax is relatively concise and readable. The syntax and grammar does, however, become more difficult to interpret as expressions become more complex. It should be noted that Web-based ontologies can be serialised (expressed) in a range of standards, e.g., RDF/XML (Beckett, 2004), OWL Abstract Syntax (Patel-Schneider *et al.*, 2004) and the Manchester Syntax (Horridge *et al.*, 2006). For a review of the various syntaxes and their strengths and weaknesses see Horridge (2010). The ontology developed in this thesis is described using the Manchester Syntax.

2.2.3 Ontology Types

Ontologies can be classified according to different perspectives, for example by the richness of their internal structure (Lassila and McGuiness, 2001), or by the subject of their conceptualisation (Van Heijst *et al.*, 1997), or according to their level of dependence on a particular task (Guarino, 1998). No one viewpoint is adequate. As Borges (1964) has observed "all classifications of reality are by nature conjectural and fictional" and ontologies often easily fit into more than one category.

With respect to structural richness, ontologies can range from simple controlled vocabularies and glossaries of terms to Description Logic-based ontologies such as those just discussed. Agreement as to whether thesauri and simple taxonomies are a type of ontology is contentious and classification is highly context-dependent (Van Rees, 2003; Vanopstal *et al.,* 2009). Many practitioners differentiate by simply referring to a formal view of ontology (Uschold and Gruninger, 2004; Gruber, 1995) in which an ontology is formal if it is able to be reasoned over.

The most popular categorical labels given to ontologies generally fall into one of the following: **Upper (Top-Level or Foundation); Core; Domain; Representation; Task; Application or Reference ontology**. Roussey (2005) notes that ontology classification often occurs according to one of three dimensions: according to purpose; their level expressiveness and/or degree of specificity. In this author's opinion the differences between ontology categories, regardless of the dimension of classification, are difficult to comprehend from the available literature. A sample of definitions and descriptions of often used ontology types is provided below to illustrate this point.

Upper ontologies describe very general concepts and sit above other types of ontologies (which are subsequently meant to derive from an upper ontology's root concepts). Gangemi *et al.* (2002), state that 'Upper' ontologies are "axiomatic theories about domain-independent top level categories such as object, attribute, event, part-hood, dependence and spatiotemporal connection. They amount to repositories of highly general information modelling concepts that can be reused in the design of application ontologies for all kinds of domains". Unfortunately, the literature reveals that there are a number of competing upper ontologies (e.g., SOWA (Sowa, 2010), DOLCE (LOA, 2012), BFO (IFOMIS, 2011), Cyc (Cycorp, 2012), UMBEL (Bergman and Giasson, 2012) and COSMO (COSMO, 2012)) all with slightly different underpinning philosophical origins. To counter this situation Niles and Pease (2001) began a project to bring together elements of existing upper ontologies in order to create the definitive upper ontology. This ontology (still under construction) is called the Suggested Upper Merged Ontology (SUMO). Niles and Pease (2001) claim that SUMO will provide definitions for general-purpose terms, and it will act as a foundation for more specific domain ontologies. It is estimated that it will eventually contain between 1000 and 2500 terms with approximately ten definitional statements for each term.

Core ontologies appear to fit somewhere between an upper ontology and a domain ontology. The boundary between the two is often unclear (Roussey, 2005). Doerr *et al.* (2002) argue that the goal of a core ontology is to provide a global and extensible model into which data originating from diverse sources can be mapped and integrated. This canonical form can then provide a single knowledge base for cross domain tools and services such as resource discovery, browsing, and data mining. A complete and extensible ontology expressing basic concepts that are common across a variety of domains can provide a shared foundation for more specialized vocabularies and domain-specific concepts.

Valente and Breuker (1996) further explain that core ontologies should consist of a clear, theoretical framework for the selection of elements of the domain and principles for their definition. In their paper they propose that core ontologies be constructed using four main principles (they should contain enough concepts, but only those concepts which are strictly necessary; they should not be a simple hierarchy of terms, but a theoretical framework that describes what the domain is about; core ontologies should not aim at the specification of the most common terms, but of basic categories of domain knowledge; and they should be coherent).

It is difficult to assess from either of the two former descriptions of core ontologies how one would differentiate what concepts and axioms should go into an upper ontology vs a core ontology, or even a domain ontology (mentioned next).

Domain ontologies are envisaged as highly reusable resources within a specific universe of discourse. They provide vocabularies of concepts and the relationships between these concepts covering activities taking place in the domain (Gomez-Perez *et al.,* 2004). Although Gomez-Perez state there is a clear distinction between a domain ontology and an upper ontology, their rationale refers to the fact that domain terms are specialisations of concepts that are found in an upper ontology. This is a rather circular and self-referencing argument and in practise many domain ontologies exist that simply do not reference an upper ontology. Guarino (1997) has defined a domain ontology as a set of knowledge modelled for obtaining “task independent” aspects of a conceptualisation.

Jean *et al.* (2006) define a domain ontology as a formal and consensual dictionary of categories and properties of entities of a domain and the relationships that hold among them. By ‘entity’ they mean anything that can be said to be in the domain. The term dictionary emphasizes that any entity of the ontology and any kind of domain relationship described in the domain ontology may be referenced directly, for any purpose and from any context, independently of other entities or relationships, by a symbol. This identification symbol may be either a language independent identifier, or a language-specific set of words. But, whatever the symbol is, and unlike in a linguistic dictionary, this symbol denotes directly a domain entity or relationship, the description of which is formally stated providing for automatic reasoning and consistency checking.

Jean *et al.* (2006) further discriminate between domain ontologies that are “linguistic” i.e., those whose scope is the representation of the meaning of the words used in a particular universe of discourse, in a particular language and those that are “conceptual”. In conceptual ontologies the goal is the representation of the categories of objects and of the properties of objects available in some part of the world. This leads them to define three forms of domain ontology: a canonical conceptual ontology (CCO); a non-canonical conceptual ontology (NCCO) and a linguistic ontology (LO).

In canonical forms each domain concept is described in a single way, using a single description that may include ‘necessary’ conditions. As a consequence, CCOs include ‘primitive’ concepts only. NCCO’s contrastingly focus on defined concepts. In OWL, ‘primitive’ concepts have only a label, an annotation and the properties that can apply to its individuals (also known as ‘necessary’ conditions). ‘Necessary’ conditions represent the conditions that must be fulfilled by individuals that are known to be members of the class in question. Necessary conditions alone are not enough to determine that any random individual that fulfils the conditions is a member of the class in question. ‘Primitive’ concepts contrast with defined concepts (which have both ‘necessary’ and ‘sufficient’ conditions). These ‘necessary’ and ‘sufficient’ conditions are not only ‘necessary’ for class membership, but also ‘sufficient’ to be able to determine that any random individual that meets these conditions can be inferred to be a member of the class in question (Fitting, 2011).

LO’s define words, or the contextual usage of words through word relationships (synonym, hynonym, homonym etc). These ontolgies are intended to be used as sophisticated thesauri. CCO’s are most often used by data processing communities. Malone and Parkinson (2010) refer to canonical (or reference) ontologies that are orthogonal and do not overlap. NCCO usage focuses on inference and concept equivalence and LO’s (like WordNet (Princeton University, 2012)) were designed for computational linguistics. Interestingly, although the focus of the paper by Jean *et al.*(2006) was on domain ontologies, they have cited WordNet in the context of describing domain ontologies, yet WordNet is routinely classified by other practitioners as an upper (lexical) ontology.

(Knowledge) Representation ontologies define the primitives used to formalise knowledge under a given knowledge representation paradigm. These include primitives for concepts such as ‘class’, ‘relation’ and ‘attribute’ (Gomez-Perez *et al.,* 2004). Guarino (1997) argues these are meta-level ontologies. Examples of representational ontologies include the Frame ontology (Gruber, 1995) and the OWL ontology (Dean and Schreiber, 2003).

Task ontologies, according to Guarino (1998) describe the vocabularies related to a generic task or activity (like diagnosing, scheduling, selling etc) by specialising the terms of upper ontologies. Tasks may cross domain boundaries. Chandrasekaran *et al.* (1998) state that task ontologies link the reasoning process to domain factual knowledge. “They are models of partitions of reality preserving the context that determines the semantics of the concepts within the partition”. Neither definition makes clear the boundary between concepts meant to exist in a task ontology vs those in a domain ontology.

An **Application** ontology is not surprisingly considered application-dependent and may specialise the vocabulary of domain and/or task ontologies. Often practitioners will include task ontologies within the definition of an application ontology (e.g., Yu, 2008). An application ontology is an ontology engineered for a specific use or application focus and whose scope is specified through testable use cases (Malone and Parkinson, 2010). An example of an application ontology is one which defines the workflow execution of a specific modelling application.

It should be relatively clear from the preceding descriptions that ontology classification is not a field of consensus. The idea of what constitutes an ontology is highly variable along the continuum from simple thesauri through to those ontologies enabled with Description (and other forms of) Logic. There is perhaps even less consensus and clarity surrounding the terminology used to define certain types of ontology, i.e., one man's domain ontology is another woman's core or task ontology.

Suffice to say that ontologies are applied in a wide range of disciplines, to meet various goals and use-cases. In this thesis the focus is on those ontology types that can be considered to fall into the broad classifications of upper, domain and application ontologies, although it is acknowledged that there is considerable fluidity in the definitions that have been provided.

2.3 Feature Catalogues

A key concern in this research is the role being played by a Feature Type Catalogue as a container for ontologies that can provide some of the semantics necessary to support data transactions in scientific OGC-standards-based data infrastructures and in particular those used within the Antarctic community. Using the broad definitions articulated above, a Feature Type Catalogue would primarily be considered to be a repository for domain ontologies. Although a Feature Catalogue is seen as the source of semantics for Feature Type definitions (OGC, 2011), most existing Catalogues are relational database management, or XML Schema-based implementations, that have no ontological underpinnings (e.g., see the International Hydrographic Organisation (IHO) Feature Catalogue – Powell, 2011).

In the last 10 years many investigators have identified the importance of semantically annotating services that deliver scientific data via heterogeneous Web environments (Bechhofer *et al.*, 2002; Reitsma and Albrecht, 2005; Uren *et al.*, 2006; Lutz, 2007; Klien *et al.*, 2007; Maue and Schade, 2009; Noy *et al.* 2009). Discovering and integrating data from distributed services that lack semantic annotation is highly problematic since there is usually poor agreement between data service providers regarding the use of terminology (both within and external to a domain discipline). Whilst adhering to standards, such as those promulgated by the OGC provides for a level of syntactic conformance (Percivall, 2002; Whiteside, 2005; Whiteside and Greenwood, 2010), semantic interoperability requires that service content be defined in an explicitly declared and shared *lingua franca* to minimise naming and cognitive heterogeneity (Klien *et al.*, 2006).

Some of the earlier attempts at semantically-enabling scientific spatial data infrastructure (SDI) focussed on using ontologies to describe aspects of dataset or registry level metadata. These ontologies were not directly coupled to individual metadata registries (catalogues) and didn't target

Feature Type descriptions. In such pilot projects ontologies were generally mapped to metadata attributes evident in one or more registries through a set of 'virtual' terms (e.g., Gil *et al*., 2006). See Figure 2.8.

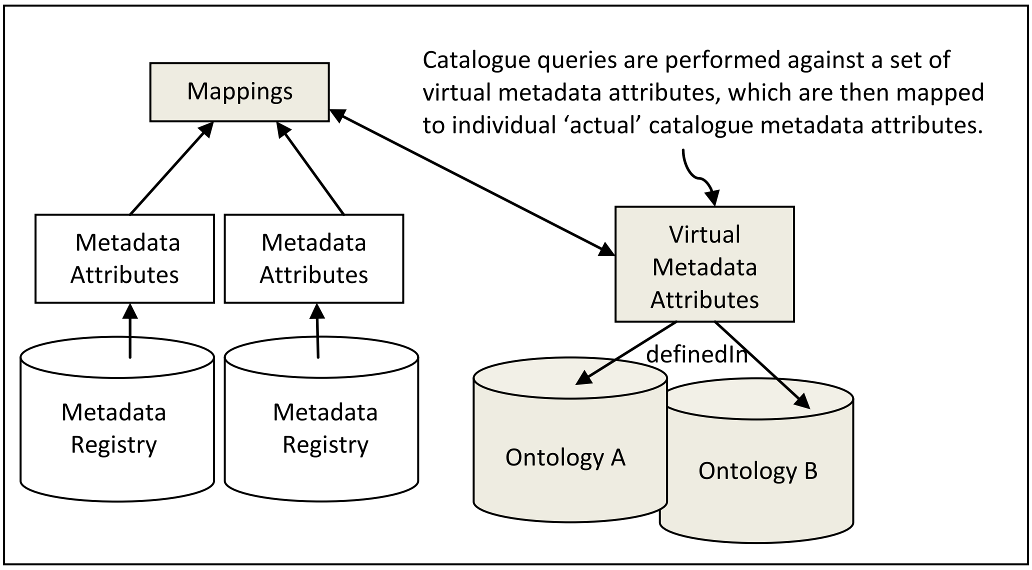

Figure 2.8 – Architecture For Early Types Of Semantic Enablement

In cases where Feature Type descriptions were instantiated as part of semantic enablement, Bakillah and Mostafavi (2010) believe that many of the older approaches were tailored towards relatively simple representations of concepts, where spatio-temporal feature attributes are either not treated, or treated as for any other generic (non-spatio-temporal) features. They also argue that important contextual information was not defined and the ontological structuring of the concept was often poor. However, this author would argue that in the more recent and advanced semantic data infrastructure test-beds issues of context and spatio-temporal characteristics have been a focus for treatment, particularly in pilot studies involving sensor networks (e.g., see Broring *et al.,* 2011).

The most advanced semantic enablement activities encompass semantic annotation involving one, or all of: creating an ontological description of dataset Feature Types (Green *et al.,* 2008, Batcheller and Reitsma, 2010; Zhang *et al.,*2010); embedding these descriptions in the registries that publish dataset service offerings (Bernard *et al.,* 2004; Schade *et al.*, 2004; Lutz and Klien, 2006) and linking to these descriptions in the service or dataset metadata (Pschorr *et al.,* 2010).

However, in almost all of these most current approaches no consideration has been given to using a standards-based Feature Type Catalogue as the container for the ontological content, rather, a customised and specific ontology (or ontologies) is usually developed to sit as a mediator, or mapping aid, somewhere in the infrastructure. This situation is probably manifest because there is a lack of standard(s) and guidance regarding how features should be ontologically described (beyond the ISO 19110 Feature Cataloguing standard) and there are no standards for how these feature-level descriptions should be managed, governed and accessed within distributed, service-based environments. This absence of guidance has led directly to a key research question in this thesis RQ1.1 *"What characterises an ontologically-grounded Feature Catalogue that can support Antarctic science data publication through Web services*".

The most recent work in this space, however, has begun to explore possible models for how Feature Type Catalogues, Service Registries and ontologies could be more tightly linked. This work is now reviewed.

2.3.1 OGC-CSW (ebRIM) Embedded Feature Catalogue

Stock *et al.* (2010) discussed embedding a Feature Type Catalogue in a Web Catalogue Service (CSW)-OGC standards compliant service registry (Nebert and Whiteside, 2005) using the main Catalogue elements (i.e., 'Feature Types', 'Attributes', 'Associations' and 'Operations') drawn from the ISO 19110 Feature Cataloguing Standard. They achieved this by extending the OASIS /ebXML Registry Information Model (ebRIM), through specialising ebRIM's ExtrinsicObject class (Stock, 2008).

In their approach Stock *et al.* (2010) and Stock (2008) effectively encapsulated all of a Feature's static semantics (i.e., 'Feature Type definition', 'Attribute definitions' and 'Feature Type associations') as well as its 'operational behaviour' (in this case, the known services that use it) in a service registry. Operational behaviours were included by linking any associated Web service implementation information to an operation's description in the embedded Feature Type Catalogue. Although the Web service implementation details were stored elsewhere in the CSW (ebRIM) Registry model, they were intrinsically linked to the Feature Type Catalogue content component.

Although the Stock *et al.* (2010) approach is a significant step forward in terms of recognising the role that can, and should, be played by a Feature Type Catalogue in a scientific SDI, Stock *et al.* (2010) argued that their Feature Type Catalogue implementation did not need to harness an ontology to be effective. They were particularly critical of current ontological approaches that had the propensity to divide ontologies up according to purpose (e.g., domain vs application vs Web service ontologies) and claimed that this trend of separating out ontological function did not sit well with their concept of a full semantic model which encompasses both static feature-level descriptions and feature behaviour (as exhibited by services), combined in one place.

In the author's opinion, the ebRIM Feature Catalogue solution (Stock, 2008), whilst a significant embellishment to the CSW Registry, suffers from a lack of constructs to express axioms and therefore cannot, on its own support reasoning. Nor does the model actually semantically define the behaviour of operations, it merely names the operations and if applicable links to an implementable service. Janowicz *et al.* (2008) also indicated that the inability to directly link Feature Types and ontological concepts using the ebRIM Feature Catalogue extension was a limiting factor in using the Registry for semantic similarity matching tasks.

There are also issues with the ebRIM Feature Catalogue extension associated with the implementation of 'inheritance' for linked Web services. Stock *et al.* (2010), whilst not subscribing to an ontological paradigm, do follow an object-oriented approach which includes the concept of inheritance (Worboys, 1994). In theory, a Feature Type that inherits operations from its parent should also inherit the linked Web service (if using either an object-oriented, or ontological paradigm). But this inheritance relationship is not necessarily straight-forward when dealing with linked Web services. Because a service may not have been designed with an inheritance relationship in mind, a service designed to operate over one type of feature may not also operate over those features considered to be its children. In this sense, the coupling of operations and static aspects of Feature Types, into one bound information model has some issues.

Stock *et al.* (2010) contrasted their encapsulated, registry-integrated approach to managing operation descriptions with competing alternatives, many of which are service-focussed and use OWL-S or WSMO to define service inputs, outputs, preconditions and results. In these latter types of semantic annotation methods, which are primarily designed to describe a service, links are created between the service and the domain objects (for example Feature Types) which are managed separately. The main drawback identified by Stock *et al.* (2010) with these types of service-based annotation approaches is that the semantics of the Feature Type can only include a Feature Type's behaviour if that behaviour is being implemented as a Web service.

It would therefore seem (to the author) that including "operations" as a component of an object's semantics poses a range of issues that are problematic whether they are described within a Catalogue, or as part of a service description. This particular issue will be examined later in this thesis, when reviewing relevant standards and the ISO 19110 Feature Cataloguing methodology (which currently includes 'operations' as part of a Feature Type Catalogue's semantic elements).

2.3.2 Web Ontology Services

If, as asserted earlier, reasoning support is required, Feature Type semantics must be delivered in a language that supports inference (preferably OWL DL). Various ontology-based enablement layers have been suggested as a potential solution to some of the missing elements required to support semantically enabled spatio-temporal data services (Lieberman, 2006; Stock *et al.,* 2009b; Janowicz *et al.,* 2009; Jirka *et al.,* 2009, Fox *et al.,* 2009). It is variously argued that this layer should complement and extend the current OGC services suite, rather than be a completely new offering (Janowicz *et al.,* 2009; Stock *et al.,* 2009b).

To this end Stock *et al.* (2009b) have developed an OWL application profile for the basic OGC CSW registry and Stock *et al.* (2009a) have demonstrated its use within a project called COMPASS (i.e., The COAstal and Marine Perception Application for Scientific Scholarship). The CSW ontology-based registry (developed in Stock *et al.,* 2009b), as opposed to the ebRIM Feature Catalogue extended profile (Stock *et al.,* 2010), stores all content in the form of ontologies, of which there are multiple (including OWL-S service descriptions; a location ontology; domain and application-specific ontologies; scientific methods and models ontologies; and an information source ontology). Reasoning support is also embedded within this registry.

Although the COMPASS prototype was considered a success (technologically), evaluation of the project highlighted that the work involved in building the necessary domain, application and service ontologies was considerable. This was despite only a very small part of the marine domain having been modelled (i.e., a description of some marine instruments) and most of that work had to be completed by ontology engineers. This observation indicates the obvious need for more tools that are focussed around the development and management of domain and application ontologies in order to encourage the ongoing development of Feature Type descriptions by communities-of-practise so that once-off and possibly unsustained, big bang approaches to ontology creation aren't the norm. It may also be suggestive of a need to lessen the complexity associated with semantic annotation and to decouple components of the infrastructure so that communities (at least at this early level of semantic evolution) can work with technologies that are more specifically focussed. The ontology registry described by Stock *et al.* (2009a) is relatively sophisticated, highly integrated and as a result complex from a domain practitioner perspective. Scientific communities may be better served by having an application that simply focuses on managing and serving Feature Type semantics. An ontology-based registry should then be able to harvest the necessary domain descriptions from independently maintained ontologically-grounded Feature Type Catalogues.

Another problem cited by Stock *et al.* (2009a) was that for the COMPASS project to be successful, domain ontologies with large (disciplinary) scope were required to allow users to select both detailed concepts from a particular specialised domain and also more general concepts applicable across multiple domains. This particular issue will be relevant to most scientific infrastructure implementations and as such anchoring to an upper ontology will be explored in this thesis to improve the scope for having both specialised and generalised concepts (the latter of which may readily transcend domains).

Janowicz *et al.* (2009) have taken a slightly different approach to that piloted in the COMPASS project and they describe two new services needed to support semantic enablement: a Web Ontology Service (WOS) for managing and accessing ontologies and a Web Reasoning Service (WRS) for providing reasoning functionality. They suggest creating the WOS as a profile of the OGC CSW and the WRS as a profile of the OGC Web Processing Service and indicate that both may be loosely or tightly coupled. In their model the WOS encapsulates existing ontological repositories and simple ontology-based files through an access, look-up and retrieval service. It is not clear from the description provided by Janowicz *et al.* (2009), however, whether the encapsulation of ontological information is by replication or by reference, or a combination of both.

Lassoued *et al.* (2008) have implemented a limited-in-scope type of WOS to create a virtual solution for integrating Coastal Web Atlases. In this pilot they created a global ontology as part of a "super atlas" that maps to ontological descriptions contained in local CSW-based ontologies. This is performed through the use of multiple mapping or mediating ontologies for each atlas service offering. Using this method, local atlases are not integrated or copied, instead they remain at their hosted locations and are remotely accessed, harmonised and integrated on the fly depending on a user's request. This model and that of Vidal *et al.* (2009), who demonstrate a similar mediated method over aeronautical data, rely on data (service) providers supplying or generating domain and/or application ontologies. Hobona and Brackin (2011) also used a WOS-like approach to semantically mediate between topographic datasets from remote sources, where Feature Type domain ontologies were constructed in OWL and linked to high-level Feature Type metadata stored in a basic CSW Registry. None of these approaches, however, have conceived of a Feature Type Catalogue as the source of domain ontologies.

These latter types of mediated and mapped semantic infrastructure models, however, can sit comfortably with the idea of establishing independent ontologically-enabled Feature Type Catalogue repositories which could then be accessed by one or many WOSs. Such separately established Catalogues could then be used as the point of independent reference for ontological descriptions contained in a wide range of annotated services and use-cases.

2.3.3 General Ontology Repository Characteristics

It is not surprising that there is a lack of standards specifically relating to the role of Feature Catalogues as ontology repositories, because the general processes for identifying and accessing ontological resources are themselves poorly standardised (Hartmann *et al.,* 2009). Hartmann *et al.* (2009) have suggested the following definition for an ontology repository and its associated management system and have proposed a Generic Ontology Repository Framework (GORF) as a starting point to address current short-comings:

> "An Ontology Repository (OR) is a structured collection of ontologies (schema and instances), modules and additional meta knowledge by using an Ontology Metadata Vocabulary. References and relations between ontologies and their modules build the semantic model of an ontology repository. Access to resources is realized through semantically-enabled interfaces applicable for humans and machines. Therefore, a repository provides a formal query language. Software to manage an ontology repository is known as Ontology Repository Management System (ORMS). An ORMS is a system to store, organize, modify and extract knowledge from an Ontology Repository."

Hartmann *et al.* (2009) argue that the main driving motivation for creating ontology repositories is to support knowledge access and reuse for humans and machines. Hence ontology repositories on the one hand act as a storage facility and on the other provide access to knowledge through defined interfaces and policies. To achieve these goals, comprehensive facets must be considered by an ontology repository when handling ontologies. The GORF extends conceptually the SEAL (SEmantic portAL) framework (Hartmann and Sure, 2004) and preliminary work on ontology repositories, as described in Hartmann *et al.* (2005a). The idealised OR Framework comprises of five conceptual knowledge layers (as distinct from an ORMS):

- **Access**: the repository must provide adaptable views on the stored knowledge, as shown in Hartmann and Sure (2004), involving different ways to view and query the knowledge.
- **Processes and services**: the repository should include facilities for ontology: evaluation, rating, mapping, security and engender trust.
- **Organisation**: the repository needs to cover factors such as ontology metadata/annotation, lifecycle management, modularisation, validation, registries and indexes.
- **Storage**: the repository needs to be scalable in terms of its: query processing ability, consistency of response and propensity for replication.
- **Sources**: the repository should include the ability to harvest from heterogeneous sources.

Nyulas *et al.* (2009) and d'Aquin and Lewin (2009) are amongst the few groups who have exercised many of the conceptual Framework facets described in GORF, the former through 'BioPortal' (a collection of biomedical ontologies) and the latter, through 'Cupboard' (a system to upload, expose and explore ontologies, as well as to find, select, reuse and assess other users' ontologies). 'Oyster' (Palma and Haase, 2005) and 'Onthology' (Hartmann and Sure, 2004) are current examples of ORMS.

Both Baclawski and Schneider (2009) and Hartmann *et al.* (2009) point out that most existing ontology repositories tend to have many of the same features, such as registration, submission and upload; browsing and search; description and documentation; metrics and statistics, but that there are many missing and desirable features. The most noticeable is the lack of structure among the hosted ontologies (Allocca *et al.,* 2009). The ontologies in a repository are treated as independent entities. Another missing feature is the lack of sufficient metadata annotating the hosted ontologies. While ontologies are claimed to be a mechanism for interoperability and communication between data sources, ontologies themselves are nearly always built in isolation (Baclawski and Schneider, 2009). There is no common representation of metadata annotations of ontologies and no common ways to identify versions of ontologies. The various extant ontology repositories use a variety of techniques and do not enforce any standard conventions for content description or for communication interfaces. As such discovering and querying repository content is problematic. A strong candidate ontology metadata description, however, is the Ontology Metadata Vocabulary (OMV: Hartmann *et al.,* 2005a).

To address the problem of common browsing across repositories, Viljanen *et al.,* (2010) have proposed creating a network of Linked Open Ontology Services (LOOS) consisting of ontology repositories that publish their content using a shared API. This LOOS API is implemented in a demonstrator as a lightweight, stateless, and cacheable HTTP GET based API that returns data using the JSON format. In many respects this API is similar to the REST API used in BioPortal (Noy *et al.,* 2009). However, the goal of LOOS was to build a network of ontology services whereas BioPortal's focus is to publish an API for accessing its own repository's full functionality.

In terms of lessons that could be learned from the stand-point of policies and best practise, Kendall (2009) has shown the following factors to be well correlated with ontology repository reuse:

- small development teams with larger user communities;
- commitment to users and to continuous improvement;
- publication of maintenance policies; URI naming conventions and protocols; and useful documentation.

Baclawski and Schneider (2009) claim it is therefore important to have well specified policies for vocabulary management, metadata, and provenance specification, particularly to enable trust. It is also critical to have a commitment to forming, accommodating, serving, and working with a community of users. This emphasizes the importance of outreach and education, including the identification and promotion of best practices.

All of the previously mentioned open issues and activities in general ontology repository research are of relevance and interest in helping to characterise a Feature Catalogue, which is essentially a repository to manage domain concepts in scientific OGC-standards based infrastructure.

Of particular significance in the work that has just been described is the assertion by Hartmann *et al.,* (2009) that ontology 'rating' and 'evaluation' are important facets of an idealised GORF. In a somewhat poorly justified argument they separate 'rating' from 'evaluation' on the basis that 'rating' is a subjective assessment of an ontology, whilst ontology 'evaluation' can be seen as an assessment of the quality and the adequacy of an ontology, or parts of it regarding a specific aim, goal or context. In the author's opinion 'rating' is actually a type of 'evaluation' and doesn't necessarily deserve to be singled out as a separate element in GORF. Also, as will be shown in the next section, ontology evaluation, as an activity and depending on what is being assessed, can be just as subjective as ontology rating.

Hartmann *et al.* (2009) envisaged that selected evaluation strategies could readily be implemented in "an evaluation component and applied in a large repository". Whilst the author agrees that ontology evaluation is an important part of ontology management and re-use, the inference in Hartmann *et al.* (2009) is that this is a relatively straightforward exercise of picking strategies from an available list of techniques and applying them. As will be shown next, the field of ontology selection and evaluation comprises many diverse evaluation strategies, many of which when put to the test do not have necessarily measurable criteria, or where there are measures these are not experimentally verified, or worse, what is being measured is not an indicator of the criteria it purports to measure.

2.4 Ontology Selection and Evaluation

A brief overview of the topic of ontology selection and evaluation is provided next in order to give some context for the decisions made regarding the choice of research questions outlined in Chapter 1 and the research methodologies that are detailed later in Chapter 3. Recall that this thesis is concerned with the issue of fostering ontology use within scientific communities, ultimately to increase the capacity for computer assisted data search and data integration. Whilst an ontologically-grounded Feature Type Catalogue can go part of the way to assist communities with this issue, it only addresses half of the problem. A desired research outcome of this thesis is a contribution towards reducing barriers in scientific communities for ontology uptake. Ontology re-use is fundamental to achieving this goal, so processes for selecting existing ontologies that can be used, in whole or in part, to seed a community-based repository are of central interest (thus RQ1.2 *What typifies an expert-grounded ontology selection framework that can support multi-disciplinary Antarctic science communities using Web services"*). Importantly, as Hartmann *et al.* (2009) note, evaluation is central to a repository management framework.

2.4.1 Evaluation Criteria

Ontology "evaluation" is a task within the ontology selection process where assessment against any given selection criteria takes place (Sabou *et al*., 2006). Ontology "selection" is the process that is used to determine if one or more ontologies, or ontological components satisfy certain, pre-determined criteria. According to Davidson (2005) the purpose of evaluation is ''to find areas for improvement and/or to generate an assessment''. This is achieved by ''the systematic determination of the quality, or value of something'' (Scriven, 1991).

While there appears to be general agreement on the broad categories of methods (discussed later), there are differing opinions about how to dimensionalise the criteria actually used for ontology evaluation. Sabou *et al.* (2006) categorise criteria that are used to select ontologies as those based on an ontology's popularity, the richness of semantic data provided and topic coverage. Gangemi *et al.* (2005) have grouped criteria according to structural measures, as well as functional measures that are related to intended use and usability related criteria. Lozano-Tello & Gomez-Perez (2004) posit five dimensions. These include ontology content and organisation of that content; language of implementation; development methodology; software tools used to build and edit the ontology and ontology development and usage costs.

Brank *et al.* (2006) focus less on criteria and alternatively distinguish different "levels" of evaluation, where "level" is not a particularly well-defined concept. A "level" appears to be a categorisation based on an arbitrary mixture of measurement types, evaluation methods and intended ontology use. It is instructive, however, to note the different types of levels determined by Brank *et al.* (2006), which include a:

- *Lexical, vocabulary, or data layer*: where the focus is on evaluating the presence or similarity of concepts, instances and vocabularies between any given ontology and domain data requirements.
- *Heirarchy or taxonomy*: mainly an evaluation of how closely related concepts are, through various techniques for examining "is-a" relationships.

- *Other semantic relations:* as above, but including an examination of other types of relationships e.g., “part-of”.
- *Context or application level:* which assesses how the ontology affects the performance of an application in which it is to be deployed, or how well it fits within a particular context.
- *Syntactic level:* evaluation of how well an ontology matches the syntactic requirements of a particular language in which it is encoded.
- *Structure, architecture, design level:* which often involves a manual assessment of conformance with pre-defined criteria, or requirements.

There are various other, often overlapping categorisations, posited by ontologists in the literature. The more common of these cited include:

a) **Consistency**: which covers conditions related to the logical properties of the ontology and its constituent parts (including formal and informal descriptions). The main purpose of assessing consistency is to check for inferred contradictions (Gomez-Perez *et al.,* 2004). A consistent model is satisfiable (Stecher and Nideree, 2005). This class could also subsume issues of representational coherence such as assessments of inconsistent naming and modelling styles, which for example hinder (human) ontology interpretation (Stecher and Niederee, 2005). Obrst *et al*. (2007) cover this latter aspect and labels it ‘intelligibility’.

b) **Clarity**: covers criteria associated with how an ontology communicates the intended meaning of terms. Where possible, a complete definition (a predicate defined by ‘necessary and sufficient’ conditions) is preferred over a partial definition (defined by only ‘necessary’ or ‘sufficient’ conditions). Entities should also be described with natural language annotation (Gruber, 1995).

c) **Completeness**: a class of criteria concerned with whether all knowledge that is expected within an ontology is either explicitly stated, or alternatively can be inferred from included descriptions (Gomez-Perez *et al.,* 2004).

d) **Conciseness**: criteria fall into this class that are associated with checking for unnecessary and redundant axioms (Gomez-Perez *et al.,* 2004).

e) **Correctness:** refers to whether the concepts, instances, relationships and properties that are modelled, correlate with those in the world being modelled. Correctness is context-sensitive in that it depends on the frame of reference that the ontology is based on (Corcho *et al.,* 2006).

f) **Expandability**: this alludes to the ability of an ontology to be able to incorporate new axioms without changing the existing underlying semantics of the ontology (Gomez-Perez, 2004). Gruber (1995) characterised this as 'extensibility' and explained that new terms should be able to be added without the need to revise existing axioms.

g) **Sensitiveness**: this class is related to the former definition but is centred around an ontology's capacity to absorb small changes in included axioms before changes in semantics become evident (Gomez-Perez, 2004).

h) **Flexibility**: encompasses criteria associated with the degree to which an ontology can be adapted to suit multiple views (Gangemi *et al.*, 2005). This class of criteria is also related to issues of ontology modularity. **Modularity** implies that merging ontologies should be "safe" in the sense that they do not produce unexpected results such as new inconsistencies or subsumptions between imported components. It should be possible to compose complex ontologies from simpler (modular) ontologies in a consistent and well defined way, in particular without unintended interactions between the component ontologies. Grau *et al.* (2007) have defined semantic definitions for modular ontologies which can be used to assess ontology modularity. These criteria (encompassing modularity) would then also cover issues of 'mappability' raised by Obrst *et al.* (2007) in which they suggest assessing how well an ontology maps to an upper level ontology, or other ontologies.

i) **Coverage**: which addresses issues of the richness, complexity and granularity of the vocabularies used in a particular domain of discourse (Obrst *et al.,* 2007). It is concerned with congruence of fit, rather than issues of completeness. Hovy (2002) separates coverage into two areas, i.e., coverage of terms and coverage or completeness of instances.

j) **Minimal encoding bias**: assessable criteria in this class focus on whether an ontology has been constructed to unnecessarily leverage aspects of the encoding language (or implementation framework) and thus hampers its re-usability in different representational styles or implementation environments (Gruber, 1995).

k) **Minimal ontological commitment:** An ontology should make as few claims as possible about the world being modelled, allowing the parties committed to the ontology, freedom to specialize and instantiate the ontology as needed. Since ontological commitment is based on consistent use of vocabulary, ontological commitment can be minimized by specifying the weakest theory (allowing the most models) and defining only those terms that are essential to the communication of knowledge consistent with that theory (Gruber, 1995).

l) **Organisational fitness**: these criteria cover the ease with which an ontology can be deployed within an organisational context and also relates to the '**coverage**' class already described (Gangemi *et al.,* 2005).

m) **Popularity**: these criteria cover different mechanisms for establishing the 'popularity' of concepts in ontologies, either for the purpose of determining if a given ontology adequately covers a domain of discourse or to establish its 'ranking' or popularity across the Web. For example, there may be measures that determine how many terms, present in a local ontology, are being instantiated (through reference) by other entities across the Web and this therefore indicates the *popularity* of the terms in the ontology (if not the ontology itself) (Patel *et al.,* 2003; Zhang *et al.,*2007; Peroni *et al.,*2008; Zhang *et al.*,2009).

In his assessment of the literature on evaluation, Vrandecic (2010) has condensed all of these classes of criteria into a set comprising:

- *Accuracy*: which covers criteria that relate to whether axioms of the ontology adequately comply with the knowledge of stakeholders about a domain of discourse. Axioms should constrain the possible interpretations of an ontology so that resulting models are compatible with the conceptualisations of the users.
- *Adaptability*: which measures how well the ontology anticipates its uses in that an ontology should offer the conceptual foundation for a range of anticipated tasks and be extendable and specialisable, without the need to remove axioms.
- *Clarity*: which measures how effectively the ontology communicates the intended meaning of the defined terms. Definitions should be objective and independent of the context. Names should be understandable and unambiguous. Classes should be defined not simply described.
- *Completeness*: which measures if the domain of interest is appropriately covered.
- *Computational efficiency*: which measures the ability of used tools to work with the ontology, in particular the speed that reasoners need to fulfil required tasks, be it query answering, classification or consistency checking.
- *Conciseness*: is about measuring to what extent an ontology includes irrelevant elements with regards to the domain covered, or includes redundant elements.
- *Consistency*: is concerned with detecting whether the ontology includes any contradictions (both formal – i.e., logical and informal).
- *Organisational fitness*: aggregates several criteria that are designed to assess how easily an ontology can be deployed within an organisation (and covers topics such as tools, libraries, data sources, desired alignments).

Yu *et al*. (2009), identify that some criteria, such as those surrounding 'clarity' and ontology 'expandability', can be difficult to evaluate as there are no means in place to determine them. Moreover, while the 'completeness' of an ontology can be demonstrated, it cannot be proven. They also note that some criteria can be particularly challenging to evaluate as they may not be easily quantifiable. These more challenging criteria require manual inspection of the ontology. For example, 'correctness' requires a domain expert or ontology engineer to manually verify that the definitions are correct with reference to the real world. Yu *et al*. (2009), argue that this may not always be feasible for a large ontology, or a repository covering many ontologies.

Vrandecic (2010) also points out that some evaluation criteria can be contradictory, for example criteria for assessing 'conciseness' may be contradictory to those assessing 'completeness'. Both Yu (2008) and Vrandecic (2010) stress the importance of choosing criteria that are relevant for any given evaluation. Yu (2008) goes on to propose a complete methodology (ROMEO, mentioned in the next section) for undertaking such targeted evaluations.

In addition to his dimensionalising of the evaluation criteria space, Vrandecic (2010) also defines what he calls ontology 'aspects. These 'aspects' are ones which it is claimed are amenable to the automatic, domain and task-independent evaluation of an ontology. They include assessments of:

- *Vocabulary*: Which deals with the different choices made with regards to the use of URIs or literals for all sets of names within an ontology and the types of criteria that can be used to evaluate various associated issues such as: the well-formed-ness of URIs; the correctness of resolvability; use of naming conventions; problems associated with punning and superfluous blank nodes.
- *Syntax*: Addresses the types of syntax used and the criteria that can be applied to assess aspects involving schema validation and correct application of syntax rules.
- *Structure*: Structural matters assessed through the evaluation of ontologies as graphs. Graphs are sets of RDF triples (i.e., statements containing a subject (resource) concept with a property that links to a property value or object concept). In graphs, nodes are the subject and object concepts and the property is referred to as an arc.
- *Semantics*: Evaluations surrounding the formal meanings inherent within an ontology.
- *Representation*: Deals with the relationship between ontology structure and its semantics.
- *Context*: This aspect covers features of the ontology as compared to other artefacts in its environment (e.g., a data source that an ontology describes, or competency questions).

Vrandecic (2010) asserts that few reported evaluation methods currently place emphasis on 'vocabulary', 'syntax' and 'structure' and instead focus almost exclusively on 'semantics' and 'context'. This is a particularly pertinent observation, since many of the ontology characteristics assessed in these categories (e.g., syntax errors, un-escaped characters, encoding problems, broken links, ambiguous identifiers, undefined vocabulary terms, mismatched semantics, unintended inferences) are some of the more often encountered issues and are relatively easy to detect (and repair). The degree to which these types of issues (or criteria) are being evaluated by science domain practitioners is an open question. This thesis examines which evaluation criteria are being used by experts in the field (i.e., RQ1.2.1) and in doing so, will provide a vehicle to assess the findings of Vrandecic (2010).

2.4.2 Evaluation Methods and Measures

For ontology evaluation to be meaningful, ontology evaluation criteria must be measurable. In software engineering, measurement has been defined as "the empirical, objective assignment of numbers, according to a rule derived from a model or theory, to attributes of objects or events with the intent of describing them" (Kaner and Bond, 2004). Whilst this definition leans towards a qualitative emphasis, others, e.g., IEEE Standards Association, appear to open the door to the possibility of quantitative measures. For example, IEEE (1998a) defines measurement as "the act or process of assigning a number or category to an entity to describe an attribute of that entity." In either case an important consideration when assigning numbers or categories, is 'construct validity' (Fenton and Melton, 1996). Construct validity is concerned with ensuring that the metric used actually measures the attribute which was the intentional target of the measurement. This concept is important, because as will be demonstrated, some of the evaluation methods which will be discussed have either not specified measurable attributes for the criteria under assessment, or the measures used may not follow the rule of construct validity.

Additionally, of the relatively wide range of suggested measures, often associated with the more academically-oriented ontology evaluation and development exercises, many measures are difficult to extract, particularly for domain practitioners with little ontological experience (Hartmann *et al.*, 2005b). Given these short-comings of some of the methods on offer it would be of considerable interest to better understand how the relatively small number of practitioners in the science domain, who are evaluating ontologies for reuse, are framing their evaluation criteria and measuring them.

Usually the metrics found in the literature are associated with specific types of evaluation methods (i.e., systematic processes designed for conducting evaluation exercises). Brank *et al.* (2006) have summarised the various evaluation methods currently available as falling into one of four broad categories:

- those based on comparison with a Gold Standard;
- those based on using the ontology in an application and evaluating the results;
- those involving comparisons with a source of data about the domain to be covered, and
- those where evaluation is performed manually by people who assess how well an ontology meets pre-defined criteria, standards and requirements.

These types of methods can be automated, semi-automated or manual in terms of the assessment process (Obrst *et al.* 2007; Yu, 2008; Vrandecic, 2010). Additionally, these methods might be applied at different stages in the ontology life-cycle (e.g., at design time, post development or during whole-of-life). The metrics used within the different ontology evaluation methods differ significantly dependent upon the purpose of the evaluation, on the type of evaluation methods used, on whether one ontology or a group of ontologies is being assessed and on what tools are harnessed in the process. Table 2.3 summarises reported evaluation methods and lists what is broadly being measured (in terms of the main classes of criteria discussed earlier) and the metrics that are being used.

Table 2.3 – Summary of Reported Evaluation Methods

Method/Tool Name	Criteria Focus	Measures	Reference
EvaLexon	Content	Recall[a], precision[b], accuracy[c] and coverage[d] are metrics used in this method by comparing mined text with a list of domain relevant words.	*Spyns (2005)
Ontometric	Structure, syntax, function, content, usability	160 characteristics that are assessed across the 5 dimensions of content; language; development methodology; software environment and cost using AHP.	*Lozano-Tello & Gomez-Perez (2004)
OntoClean	Structural	Uses measures associated with applying the ontological concepts of rigidity[e], unity[f], identity[g] and dependence[h]. An appropriately tagged ontology can then be compared with a predefined ideal taxonomical structure to detect inconsistencies.	*Guarino & Welty (2004)
OntoKhoj	Popularity	Counts links and referals between ontologies (PageRank method)	Patel *et al.* (2003)
OntoSelect	Popularity, structure and content	Corpus coverage measured by assessing the number of labels for classes and properties that can be	Buitelaar *et al.* (2004)

		matched to a supplied corpus document. Structure measured by number of properties. Connectedness measured by the number of imported ontological components	
OntoEval	Structural, functional and usability	Develops a metaontology (O^2) to characterise ontologies, complements this with an ontology of evaluation and validation (*oquα*l) and uses it to assert tests for certain types of evaluation tasks. What is provided is a design pattern for *quality-oriented ontology descriptions* (*qoods*).	Gangemi *et al.* 2005
ActiveRank	Structural	Uses 4 main measures which are later aggregated and summed taking into account weights. Measures include: *Class Match Measure*: coverage of search terms; *Centrality Measure*: measures how close a concept is to mid level of a concept hierarchy. *Density Measure*: how many attributes and siblings a concept has – used as a surrogate for detailed concept representation. *Semantic Similarity Measure*: uses formula based on the shortest-path measured between concepts.	Alani & Brewster (2005)
ODEval	Syntax and content	Measures to detect circularity, redundancy and partition errors using graphs.	*Corcho *et al.* (2004)
OntoManager	Popularity	Usage statistics	*Stojanovic *et al.* (2003)
Swoogle	Popularity	Counts links and referrals between ontologies (PageRank method)	UMBC http://ebiquity.umbc.edu/
OntoSearch	Structure	Search front-end for selecting ontologies. Uses ActiveRank for evaluation.	Thomas *et al.* (2005)
OntoQA	Structural and functional	Schema relationship richness (% relations not in an *is-a* relation), attribute richness (average no. of attributes per class) and inheritance richness (average no. of sub-classes per class). Also includes measures for ontology instances and readability measures (inclusion of comments, labels or captions)	Tartir *et al.* (2005)
OntoCAT (Protégé plug-in)	Structural and functional	Measures of semantic similarity and relatedness (similar to OntoQA)	Cross & Pal (2006)
Gold Standard Based Approach	Structural	A given ontology is assessed against an *ideal* ontology using measures that capture the arrangement of ontological concepts and their hierarchy (lexical similarity measures). Alternatively, similarity measures can	*Maedche & Staab (2002) and Brank *et al.* (2006)

		be captured for actual instance data and may use distance measures based on clustering algorithms.	
CORE	Structural, popularity and functional	Uses lexical and taxonomic measures proposed by Maedche & Staab (2002) but also includes manually assessed functional measures of correctness, readability, flexibility, level of formality and type of model. Post structural analysis users of CORE can set minimum thresholds that an ontology must meet for each manual criterion and CORE will compute which ontologies then best match criteria.	Fernandez *et al.* (2006)
ROMEO	Requirements	The ROMEO methodology helps ontology engineers to determine relevant ontology evaluation measures for a given set of ontology requirements. ROMEO links these requirements to existing ontology evaluation measures through a set of questions.	Yu (2008)

[a]Recall: Is a value derived for comparing a reference retrieval with a computed retrieval returned by a system. It is the fraction of the concepts (or triples) that are relevant to the query that are successfully retrieved.

[b]Precision: Is a value derived for comparing a reference retrieval with a computed retrieval returned by a system. It is the fraction of the retrieved concepts (or triples) that are relevant to a search query.

[c]Accuracy: A value obtained by averaging the coverage percentage for the relevant frequency class. A frequency class is relevant if it contains more than 60% of typical vocabulary words considered as characteristic of the text on the basis of statistical calculations.

[d]Coverage: A value obtained by counting, for each frequency class of mined triples the number of words, constituting the triples, that are identical with words from a reference text for that frequency class.

[e]Rigidity: A rigid property is one essential to an object and all of its nuances. An ant-rigid property is not essential to any of an object's instances.

[f]Identity: Metrics related to the problem of distinguishing a specific instance of a certain class from other instances of the same class by means of a characteristic property which is unique to it.

[g]Unity: Metrics related to the problem of distinguishing the parts of an instance from the rest of the world by a unifying relation that binds the parts, and only those parts together.

[h]Dependence: A property is dependent when it is necessary that another class must exist for it to exist (e.g., a mother only exists if she has a child).

Hartmann *et al.* (2005b) reviewed a sub-set of the evaluation methods listed in Table 2.3 (see asterisked methods) and identified that all came from an academic realm and could not be considered methods regularly used by real-world domain practitioners. Although they considered some methods more immediately useful in practical applications than others, all in their opinion necessitated some form of adaptation to be applicable for industrial (practical) use. Although tools such as OntoCAT, OntoManager and OntoSelect are aimed at practitioners, they only address some of the evaluation criteria that will be of interest to potential ontology consumers.

In real-world ontology development exercises, ones in which an ontology is being developed to meet a genuine use-case expressed by a community-of-interest for operational purposes, it is assumed that the ontology developers have skills skewed towards domain expertise rather than ontological engineering skills. This is presumed to contrast with a developer's ontological engineering skill level in the more "academically-slanted" development exercises, where an ontology is primarily developed to meet hypothetical use-cases and/or to support some aspect of basic ontological research. In these latter cases it is anticipated that ontological engineering skills are much higher. This assumption about the nature of practical developments is an important one because it should influence the development and evaluation methods selected by community ontology builders who may be constrained in their methodological choices by their skill-base.

Apart from a broad ontology engineering survey conducted by Paslaru Bontas-Simperl and Tempich (2006), this topic has not been extensively investigated. Paslaru Bontas Simperl and Tempich (2006) indeed found that only a "small fraction" of the sampled ontology builders engaged in real-world development tasks had received any ontological engineering training. Importantly the survey also highlighted that only a "small percentage" of practical ontology-related projects follow any systematic approach to ontology-building and even less commit to a specific methodology. The research in this thesis will investigate to what extent this is still the case, given that Paslaru Bontas Simperl and Tempich conducted their study in 2006. In addressing RQ1.2.1- "*what ontology selection and evaluation criteria are currently used across multi-disciplinary scientific communities (and are selection and evaluation methods consistent with those reported in the literature)"* this thesis will examine whether there is any connection between an interviewed expert's stated skill level and the methods they have chosen for ontology evaluation. Likewise, in answering RQ1.2.2.2 – where the focus will be to *"to ascertain whether an expert's assigned* (relative importance) *weights* (for different evaluation criteria) *differed depending on an expert's scientific domain, or any other discernible factor*", consideration will be given to assessing whether skill level is a differentiating factor.

Independent reviews of ontology evaluation methods that are presented in the academic literature also consistently express an opinion that many of these academically-slanted evaluation techniques are complex (Hartmann *et al.*, 2005b), too narrow in what they are assessing (Lozano_Tello and Gomez-Perez, 2004), can be resource-intensive to apply (Kalfoglou & Hu, 2006) and require reasonable ontological engineering skills to execute (Blomqvist *et al.*, 2006). Yu *et al.* (2005) for example, carried out an application of 'OntoClean' and found that it was able to clarify modelling assumptions made in an ontology. However, the most crucial part of this ontology methodology is that each class in the ontology is required to be tagged with formal meta-properties and this must be done manually. Yu *et al.,* (2005) lamented that the annotation activity was very time-consuming and tedious. Additionally, the classifications given to the annotations could be highly subjective and depended on interpretations by the knowledge engineer about the concept being labelled.

If the reflections above are an accurate depiction of the methods currently being promoted, these factors could be hampering re-use in real-world development scenarios. By developing guidance in this thesis in the form of an evaluation framework, grounded in practise, assistance could be provided to those scientific domain experts just entering the semantic-enablement field.

It is also claimed that many of the methods canvassed in Table 2.3 have not been properly designed, defined, implemented and experimentally verified (Vrandecic, 2010). A specific example involves the method of Gangemi *et al.* (2005) where a number of criteria are listed and some metrics are provided that purport to be either positive or negative ontology design indicators. However, no experiments are reported that investigate whether these metrics correlate with their design conclusions. Vrandecic (2010) cites a case where Yu *et al.* (2007) describes a higher degree of 'tangledness' as being beneficial for increasing the efficiency of an ontology designed for browsing tasks, whereas this factor is a negative design facet according to Gangemi *et al.* (2005). Tangledness is where ontologies have subsumption hierarchies with more than one super-class per concept.

Yu (2008) has pointed to a potential lack of construct validity in ontology evaluation research and indicates that only relatively recently has the practice of rigorous definition and validation of ontology measures been adopted (Orme *et al.,* 2006; 2007).

Other criteria such as 'consistency' and 'satisfiability' are difficult for non-logicians to assess. Existing ontology development environments provide some limited support, in conjunction with a reasoner, for reporting errors in OWL ontologies. Typically these are just restricted to the detection of unsatisfiable concepts. However, the diagnosis and resolution of such a bug is often poorly supported. For example, no explanation is given as to why the error occurs (e.g., by pinpointing the root clash, or axioms in the ontology responsible for the clash) or how dependencies between classes cause the error to propagate (Kalyanpur *et al.,* 2005). Criteria such as 'clarity' and 'expandability' are very difficult to evaluate as there are currently no measures for them (Yu, 2008).

Another problem, according to Yu (2008) is that although there are aspects of existing research into ontology evaluation that relate to each other, often methods for ontology evaluation are discussed and implemented quite independently of each other. For example, Tartir *et al.* (2005) propose measures for ontology evaluation, of which some measure identical aspects as measures found in Gangemi *et al.* (2005). Thus, there is overlapping work in this area that can be extremely confusing for domain practitioners without a high level of ontological skill. Yu (2008) also laments that in current ontology evaluation research, few examples are found outlining detailed examples of evaluation for real world data (or scenarios), further supporting the view of the author that a review of practise is required.

2.5 Summary

In this chapter an overview of the OGC interoperability standards stack was presented, highlighting the central role played by Feature Types as carriers of semantic information. It was explained that the OGC stack primarily focuses on issues of syntactic interoperability and supplementation is required from general IT standards to semantically-enable scientific data infrastructure. In particular, the current 'state of play' was discussed with respect to research in the area of semantic-enablement for feature-centric Web services. It was also shown that although Feature Type Catalogues are designated, in OGC standards, as the repositories for domain semantic information, most current scientific infrastructure do not use Feature Catalogues as repositories for ontologically-grounded content. Addressing this gap is a primary driver for research in this thesis. Innovative advances in the realm of ontology repositories and ongoing research areas were canvassed, given that Feature Catalogues can be considered a specialised type of ontology repository. Lessons that could be learned from the general ontology repository space can be used to help characterise a standards-based Feature Catalogue – which is one of the two main research questions being addressed in this thesis.

A general introduction to ontologies was provided, mainly as a readers guide, so that work in subsequent chapters can be appreciated and placed in context.

This chapter concluded with an appraisal of reported ontology selection and evaluation criteria and their potentially associated evaluation methods and measures. This review highlighted what knowledge exists, but also focussed on current limitations in the field. It was obvious from this review that further investigations are warranted into the current ontology evaluation practises of people who are involved in building today's scientific semantic data infrastructure in order to better understand how they are applying the current knowledgebase on ontology selection and evaluation. This review is integral to addressing the second major research question in this thesis aimed at comparing current ontology evaluation practise with that reported in the literature.

The results of the expert-grounded investigations on ontology evaluation, that will be conducted in this study, in combination with knowledge from the literature, should be able to provide best-practise guidance that can facilitate selection of reusable semantic content suitable to seed an ontologically-grounded Feature Type Catalogue.

Chapter 3.

General Methods

The studies outlined in this chapter are designed to investigate how scientific communities of practise can improve the semantic description and therefore interoperability of datasets shared via the Internet. Despite significant improvements in the evolution of scientific data exchange networks, a range of fundamental software, process and governance elements are still lacking in most national and global data infrastructures (Finney, 2007; Moeller, 2010; Yang *et al.,* 2010). Three of these often missing elements are the tooling, practical standards and practises necessary to enable communities to readily agree upon, deploy and manage the terms (and their relationships) that semantically define exchanged data, particularly in the context of a Web services environment. It has already been argued, in Chapter 1, that designing an appropriately configured and implemented Feature Type Catalogue could fill part of the current void. The first half of this study is therefore devoted to examining what an "appropriately configured and implemented" Catalogue would look like and how such a Catalogue could readily interface with other key elements that generally make up existing scientific data infrastructures.

However, simply creating a Feature Type Catalogue artefact does not solve the problem (for scientific communities) of what to put in it. So, the remainder of this research examines how communities of practise currently evaluate and then select ontological components, to describe their datasets, so that there is guidance regarding Catalogue population. As has already been described, of particular interest in this research are scientific communities who are developing ontologies in order to facilitate the exchange and/or integration of domain data that includes spatial and temporal dimensions and who intend to implement their ontologies within a Web services framework, according to either Open GIS Consortium (OGC) and/or the ISO TC 211 standards suites. For these communities, manipulation and integration of compatible four-dimensional (3D-space & time) data is a common requirement.

The various research questions designed to address these issues were presented in Chapter 1 (Figure 1.2). As will be demonstrated in this chapter, the problem space addressed by the main research question: "*how can Antarctic science communities practically manage and select domain ontologies for use in semantically-enabled data exchange scenarios, given feature-centric Web service design patterns* (RQ1) ?", is explored by a mixed model of research spanning Design Science, qualitative and quantitative methods.

In Chapter 2 the reader was provided with an in-depth examination of Web service contexts of specific interest to this study and was introduced to ontology fundamentals primarily as background material. A literature review covering Feature Type Catalogues; ontology repository implementations and the topic of ontology selection and evaluation also formed a substantive part of Chapter 2 so that the reader can subsequently make informed judgements about the value of the various contributions developed progressively throughout this thesis.

In this Chapter (3) the focus is on expanding upon the overall philosophy behind the research paradigms adopted. The individual research methods used to assess the questions previously stated are outlined in detail and an explanation is provided of how the research outcomes are validated.

3.1 Research Philosophy and Validation

Paradigms

A paradigm refers to a typical pattern to be followed, or model, or exemplar (Scott and Marshall, 2005) and in science it implies a conceptual and philosophical framework incorporating the methods and tools by which researchers conduct investigations into areas of interest, the problems of interest, and the means by which research is evaluated. Guba and Lincoln (1994) argue that a paradigm is not a conflation of a body of “knowledge” and “method”, but refers to the epistemological, ontological, and methodological set of assumptions through which knowledge can legitimately be derived through rigorous, scientific research.

The two major research paradigms, or models which are commonly subscribed to by research professionals are the positivist/empiricist/scientific and the constructivist/phenomenological/ postmodernist approaches (Tashakkori & Teddlie, 1998; Cherryholmes, 1992; Guba & Lincoln, 1994, Howe, 1988). The former paradigm underpins what are called quantitative methods, and the latter underlies qualitative methods. Qualitative and quantitative research paradigms both have their own philosophical foundations, or assumptions.

Lincoln and Guba (1985) have identified a number of salient areas related to research in which the naturalistic (qualitative) and positivistic (quantitative) paradigms differ. In terms of the nature of reality, according to the positivistic model, there are human characteristics and processes that constitute a form of reality in that they occur under a wide variety of conditions and thus can be generalised to some degree. Furthermore, different variables related to a complex process may be studied independently. In contrast, the naturalistic model holds that there are no human characteristics or processes from which generalisations can emerge. Thus, each subject or phenomenon is different and can only be studied holistically. In terms of the relationship of the researcher to the research project, the positivists maintain that the researcher can function independently of the subject to a major degree, while the naturalistic model holds that the researcher and research subjects should interact to influence one another. Quantitative researchers seek causal determination, prediction, and generalization of findings, while qualitative researchers seek instead illumination, understanding, and extrapolation to similar situations (Hoepfl, 1997).

Positivist researchers prefer precise quantitative data and often use experiments, surveys and statistics. They seek rigorous, exact measures and objective research and they test hypotheses by analysing numbers from these measures (Neuman, 1999). Naturalistic (or interpretivist) researchers often use participant observation, field research, analyse transcripts of conversations, and/or study videos of behaviour (Neuman, 1999). Ostensibly these two worldviews are opposed.

Brewer and Hunter (1989) and Patton (1990), however, argue that it is feasible to subscribe to research that encompasses 'mixed methods' (or a mixed paradigm model) which contain elements of both qualitative and quantitative approaches. Indeed this thesis uses a mixture of methods that have their foundations in both the qualitative and quantitative paradigms, and additionally also incorporates a possible third paradigm – that of Design Science.

Design Science has its roots in engineering and the sciences of the artificial (Simon 1996). According to Osterle *et al.* (2011) design-oriented Information System (IS) research aims to develop and provide instructions for action (i.e., normative, practically applicable means-ends conclusions) that allow the design and operation of IS and innovative concepts within IS (instances). Thus, for each specific IS that is developed, design-oriented IS research builds upon a 'to-be' conception and then searches for the means to construct the system according to this model while taking into account given restrictions and limitations. It is a problem-solving 'paradigm' that involves the development, evaluation and communication of innovative artefacts, namely constructs (e.g., concepts, terminologies, and languages), models, methods, and instantiations (i.e., concrete solutions implemented as prototypes or production systems) in a rigorous manner. Concrete manifestations of such artefacts could be axioms, guidelines, frameworks, norms, patents, software (with open source code), business models and enterprise start-ups (Osterle *et al.,* 2011; Gregor and Jones, 2007; Hevner *et al.*, 2004; March and Smith, 1995; March and Storey, 2008; Pries-Heje and Baskerville, 2008; Vaishnavi and Kuechler, 2004). Knowledge and understanding of the problem domain is achieved through artefact construction (Hevner *et al.,* 2004), which must have novelty and utility in the application environment (Hevner and Chatterjee, 2010; March and Storey, 2008; Simon, 1996).

Hevner *et al.* (2004) have argued that Design Science is a paradigm in its own right (see Figure 3.1 for the elements and interactions of the framework proposed). This view is also shared by Vaishnavi and Kuechler (2004) and Iivari (2007). Vaishnavi and Kuechler (2004) justify their argument for Design Science as a paradigm by concluding that Design Science ontology, epistemology and axiology cannot be derived from any other existing paradigm and that it is sufficiently different to warrant separate classification. They consider that Design Science research by definition changes the state-of-the-world through the introduction of novel artefacts. Thus, Design Science researchers are comfortable with alternative world-states. This contrasts with a positivist ontology view where a single, given composite socio-technical system is the typical unit of analysis. The multiple world-states of Design Science ontology is also at odds with the multiple realities of the interpretive researcher because many, if not most design science researchers, believe in a single, stable underlying physical reality that constrains the multiplicity of world-states. Vaishnavi and Kuechler (2004) state that epistemologically, "the Design Science researcher knows that a piece of information is factual and knows what that information means through the process of construction/circumscription. An artefact is constructed. Its behaviour is the result of interactions between components. Descriptions of the interactions are information and to the degree the artefact behaves predictably, the information is true. Its meaning is precisely the functionality it enables in the composite system (i.e., artefact and user). What it means is what it does".

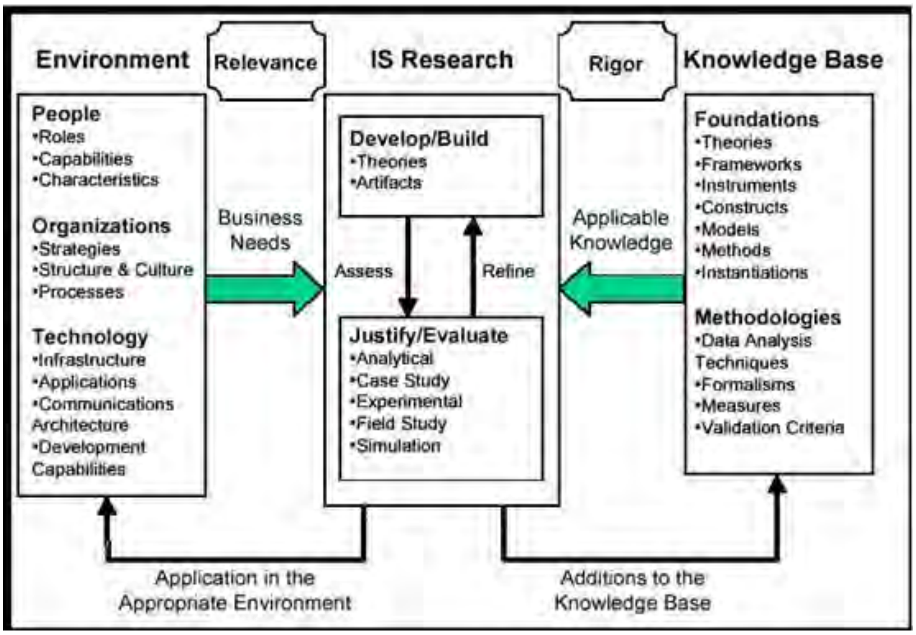

Figure 3.1. Information Systems Research Framework (From Figure 2. Hevner *et al.* 2004).

McKay and Marshall (2005; 2008) in their examination of Design Science on the other hand, argue that Design Science is not a paradigm, but a body of knowledge, most usefully built through the application of a variety of methods. Gregory (2010) agrees that Design Science is not a paradigm and instead suggests that Design Science is a research approach, “something in between a hands-on research method and a more general philosophy of science, or research paradigm”, adding that many different research methods can be used within a Design Science research project.

Osterle *et al.* (2011) agree with this latter observation and point out that an important principle applied in design-oriented IS research is deductive reasoning. The ideal case is to formally deduct (i.e., mathematically) or to use semiformal (i.e., conceptual) instruments; however in most cases design-oriented IS research takes advantage of natural-language (i.e., argumentative) deduction, taking into account existing theories and models. Design oriented IS research thus contributes substantially to the structuring and integration of the body of knowledge. But design-oriented IS research also uses inductive reasoning (e.g. when inferring from single case studies).

As can be seen there is considerable debate about the paradigmatic status of Design Science Research. Regardless of whether Design Science is classed as a “paradigm”, a “body of knowledge” or a “type of research method”, it was used as a guiding framework and employed in this thesis to develop a prototype, semantically-grounded Feature Catalogue (see RQ 1.1 – Figure 1.2). Simultaneously, and in parallel, a stream of additional activity, motivated by the need to populate the Catalogue (see RQ 1.2 – Figure 1.2), was carried out which involved both qualitative and quantitative methods (that were sequentially executed). Figure 3.2. diagrammatically provides an outline of the overall research design.

Most research supporting mixed methods is designed such that weaknesses of one method can be offset by utilising the strengths inherent in the other (Greene *et al.,* 1989; Greene 2007). Those who study mixed method research tend to try to classify research according to different dimensions, for example by:

- the degree to which the different methods used are conceptualized, designed, and implemented independently or interactively;
- determining when the mixing happens – at the end of the study or during the study;
- the priority given to one methodology or another; and/or according to whether the different methods were executed sequentially or concurrently (Green *et al.,* 1989; Teddlie & Tashakkori, 2009).

This classification aids in identifying the applicability of the research techniques chosen to address research problems (Onwuegbuzie and Collins, 2007).

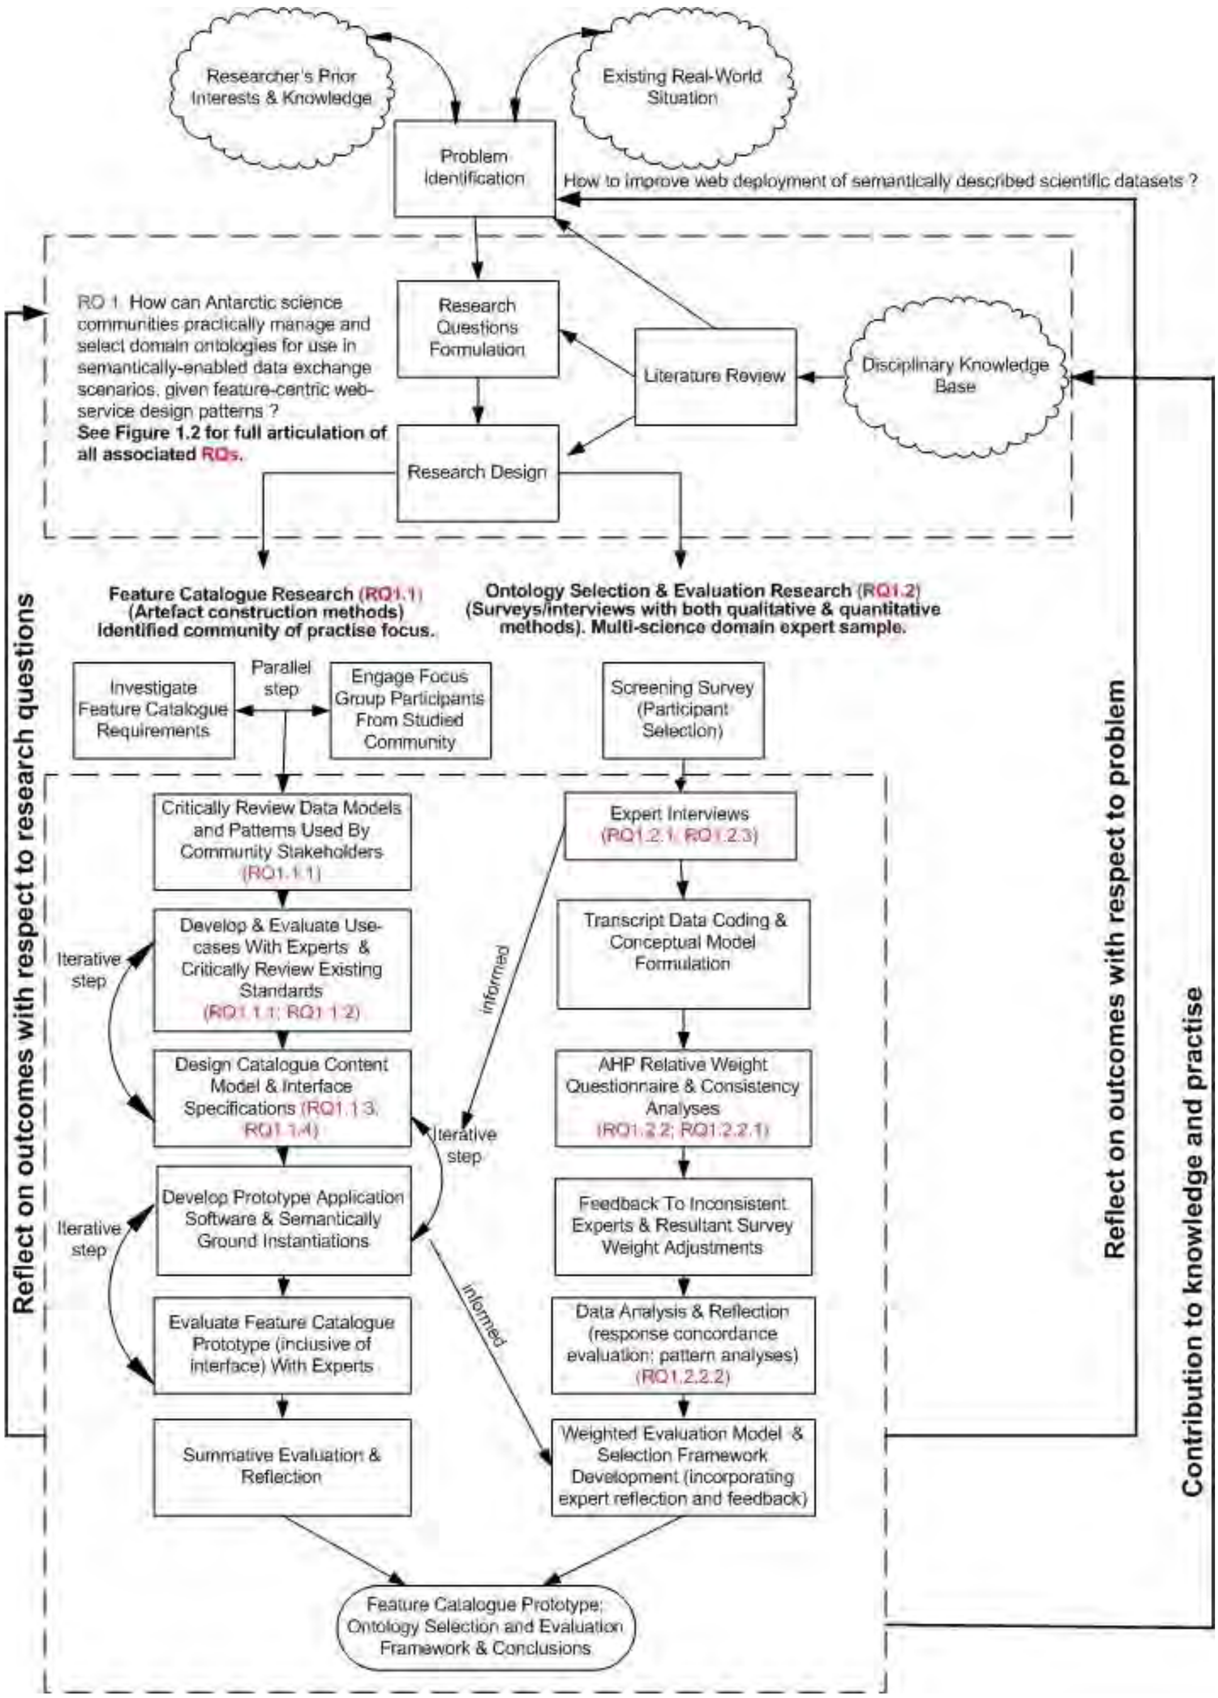

Figure 3.2 – Research Design

In this study different methods were used primarily for expansionist purposes (Green *et al.,* 1989) in that by using multiple analytical strands for different study components, the breadth and range of what was possible to investigate was increased. Although there were two clear parallel streams of research (i.e., a Design Science influenced strand and a mixed qualitative/quantitative strand), the data emanating from both strands was essentially separate, with each providing an understanding of the problem under investigation. These "understandings", taken together, formed the integrated knowledge taken from this thesis that addressed RQ1.

As can be seen from Figure 3.2, analysis of qualitative interviews, via transcripts, provided concepts arranged into a conceptual hierarchical framework. These results were then incorporated into a questionnaire which was subsequently analysed using quantitative techniques associated with AHP (Saaty, 1980). In this type of design the results of the first (qualitative) method was used to inform the second (quantitative) method. This sequential application of the qualitative and quantitative paradigms is used frequently in many other types of studies (Green *et al.,* 1989).

Whilst the data generated from both strands were subjected to different treatments, it should be recognised that since the author performed both activities in parallel, the knowledge acquired through 'doing' inevitably, and at times quite consciously, helped shape directions taken and assessments made (in each parallel strand). For example, information uncovered during the process of interviewing experts about ontology selection and evaluation methods played a role in informing how the author could select and evaluate ontologies for re-use in semantically enabling the Feature Catalogue. Likewise, the act of prototyping a semantic Feature Catalogue, gave the author a better appreciation for the ontology evaluation criteria used and discussed by experts, which helped the author's interpretative skills in building the Selection and Evaluation Framework.

From an axiological perspective, because the proposed research design involved survey techniques and expert group consultations, ethics approval was sought. The University of Tasmania has a structured process for guiding and approving studies involving human subjects (UTAS, 2012). After an initial risk assessment by the University, approval was granted to engage study participants. All of those participants that were engaged to provide expert opinion through interview, or survey, were supplied with detailed information on the research.

All experts subsequently signed Consent Forms indicating their willingness to collaborate in the study. Progress throughout was assessed annually by the University through formal written evaluation to ensure ethical practises were being exercised.

Validation

Having discussed paradigms, the issue of research validity requires illumination particularly given the mixed model nature of the design. The term "validity" has broad meaning and is most often used in connection with assessment of whether the knowledge claims resulting from research are warranted. From a mainly positivist perspective, “internal validity” refers to proof of a causal link between a treatment and effect and “external validity”, refers to the generalisability of research findings, i.e., assuming that there is a causal relationship in a particular study between the constructs of the cause and the effect, can we generalise this effect to other persons, places or times (Trochim, 2006). In discussing validity, as interpreted across both quantitative and qualitative paradigms, Golafshani (2003) quotes Healy and Perry (2000) who assert that the quality of a study in each paradigm should be judged by its own paradigm's terms. For example, while the terms “reliability” and “validity” are essential criterion for quality in quantitative paradigms, in qualitative paradigms the terms “credibility, neutrality or confirmability, consistency or dependability and applicability or transferability” are to be the essential criteria for quality (Lincoln & Guba, 1985).

If validity is to be established in studies using a mixed model, such as in this research, how is validity to be judged ? Maxwell (1992) observes that the degree to which an account is believed to be generalisable is a factor that clearly distinguishes quantitative and qualitative research approaches. In quantitative studies generalisability is often demonstrated through specific tests applied to numerical data, whereas in qualitative methods, triangulation is often used to infer generality. Triangulation, according to Yin (1994), is used to assess construct validity by using multiple sources of evidence to provide multiple measures of the same phenomenon. It is not suggested, however, that triangulation is a method only used in the qualitative domain. Quantitative methods frequently use triangulation of several data sources, where any exception could lead to a refutation of the hypothesis. But in contrast, when triangulation is used in qualitative research these exceptions can be used to modify the theories and are considered fruitful departures (Barbour, 1998).

In the mixed model used here-in, most conclusions drawn from data analyses in this thesis were not validated to a statistically measured level of certainty and often the corroboration of conclusions rely on triangulation. In this study, for example, structured surveys provided both quantitative and qualitative information which could be checked against data provided through open-ended interviews and through documentation provided by study participants. These lines of enquiry were also assessed against published literature to establish corroboration or negation of any of the conclusions reached. Through the process of triangulation, a key finding from the results of investigations surrounding RQ1.2 (involving both qualitative and quantitative methods) concluded that the selection and evaluation framework developed in this thesis was of utility and significance to a broader range of domain practitioners and experts, than just those involved in the study. This assertion was founded on the basis of argumentation and extrapolation to other groups of practitioners operating in similar circumstances, under similar constraints, with overlapping needs.

It should be noted, however, triangulation was not used as a tool for corroboration between findings from the different parallel strands of the mixed method research model, because as previously stated these streams of research endeavour were essentially independent. One strand was not meant to be a corroboration of findings in the other. The two strands were 'informative' not corroborative.

Some techniques used, were however designed to quantitatively explore the data and did utilise mathematical validation techniques. AHP (Saaty, 1980), was the key quantitative technique harnessed to investigate the preferences of experts for specific ontology evaluation criteria (arranged in a hierarchical model), which emerged as a result of qualitative investigations (involving expert interviews). The robustness of responses provided through AHP analysis was assessed through AHP-specific consistency criteria (discussed later in this chapter) which lead to adjustments of the data before a final (weighted) ontology evaluation model was proposed. Other quantitative methods (i.e., Kendall's co-efficient of concordance (Kendall and Gibbons, 1990)) were also used to assess the degree of agreement amongst expert responses derived from the AHP questionnaire. Multi-dimensional scaling was also employed to investigate any similarities in patterns amongst experts in how they rated ontology evaluation criteria within the hierarchical model. All of these techniques anchor their epistemology in a positivist perspective and lend rigour to the research, but in themselves do not constitute validation of the main research conclusions reached.

Validation of research strands which were primarily framed in the Design Science "methodology", which has positivist leanings, but which is, as previously argued a distinct epistemology to that of the former paradigms, was perceived differently. Validation and rigour in Design Science is still an evolving topic (Piirainen *et al.,* 2010). Contributions made through Design Science research require identifying a relevant organizational information and communication technology (ICT) problem, demonstrating that no solution exists, developing an ICT artefact that addresses this problem, rigorously evaluating the artefact, articulating the contribution to the ICT knowledge-base and to practice, and explaining the implications for ICT management and practice (March & Storey, 2008).

In Design Science, according to Hevner *et al.* (2004) rigour is achieved by appropriately applying existing foundations and methodologies from the knowledge base. The utility, quality, and efficacy of a design artefact must be rigorously demonstrated via well executed evaluation methods. Evaluation is therefore a fundamental component of the Design Science research process. The organisational environment (in this case the Antarctic community) establishes the requirements upon which the evaluation of the artefact is based. This environment includes the technical infrastructure which itself is incrementally built by the implementation of new IT artefacts. Piirainen *et al.* (2010), however, accuses Hevner *et al.* (2004) of focussing primarily on evaluation at the expense of validation.

Piirainen *et al.* (2010), consider that evaluation is concerned with the utility of the artefact, whereas validation is about the truthfulness of claims and their reliability and robustness. "Truthfulness" in their view is established by reference to the conceptual background of the artefact: the kernel and design theories. Validation should provide evidence of whether the artefact represents the theory sufficiently to give theoretical insights. They therefore assert that validation allows for critical evaluation of the theory and creates additions to the existing theoretical knowledge. However, having made these points they don't provide practical mechanisms by which rigour and validation can best be achieved. Instead, they state that the Design Science literature could benefit from more explicit discussions about good practise and conduct about what methods and tools could, or should, be used in the Design Science process (as related to validation).

In this thesis, noting current divergence within the discipline of Design Science on issues such as philosophical underpinnings and validation methodology, the processes as described by Hevner *et al.* (2004) are used to both validate and evaluate the research used. An emphasis has been placed on utility, rather than truth. Contributions made are in a form of knowledge that is validated and useful to practitioners and academic communities. Looked at this way, justified true beliefs are knowledge that will work. Hevner *et al.* (2004) have proposed a set of guidelines for the Design Science Research Framework that describes the artefact construction process (Table 3.1, columns 1 & 2). This set of guides can be used to examine the scientific grounding of an artefact. With the fulfilment of these guidelines a designed artefact could be viewed as scientifically validated knowledge. Table 3.1 (column 3) also maps approaches taken in this thesis which utilised Design Science, onto the guidelines proposed by Hevner *et al.* (2004).

Markus *et al.* (2002) explain that the design process is a sequence of expert activities that produces an innovative product (i.e., the design artefact). The evaluation of the artefact then provides feedback information and a better understanding of the problem in order to improve both the quality of the product and the design process. This build-and-evaluate loop is typically iterated a number of times before the final design artefact is generated. During this creative process, the

Design Science researcher must be cognisant of evolving both the design process and the design artefact as part of the research.

Table 3.1 Design Science Research Guidelines (Adapted from Table 3. Hevner *et al*. 2004) With Thesis Approaches Mapped.

Hevener *et al.* (2004) Guideline	**Description**	**Conformity Of Thesis To Guideline**
Design as an Artefact	Design Science research must produce a viable artefact in the form of a construct, a model, a method, or an instantiation.	The ontologically-grounded Feature Catalogue Model produced during the development phase of the research meets the criteria of an artefact, as it embodies a construct (conceptualisation of the problem), a model (description of the semantic relationships between constructed entities) and a method (in this case instantiated interface patterns for Catalogue communication using a REST protocol).
Problem Relevance	The objective of Design Science research is to develop technology-based solutions to important and relevant business problems	Both the literature and the Antarctic scientific community of practise indicate that the problem space being addressed is important and relevant.
Design Evaluation	The utility, quality, and efficacy of a design artefact must be rigorously demonstrated via well-executed evaluation methods	The artefact (in this case a Feature Catalogue design) is evaluated by establishing a prototype instantiation which is functionally assessed using community feedback and descriptively assessed in terms of its architectural and functional 'fit' within current data exchange infrastructure.
Research Contributions	Effective Design Science research must provide clear and verifiable contributions in the areas of the design artefact, design foundations, and/or design methodologies	This research identifies a clear gap in the existing IS knowledge base with respect to a Feature Catalogue's function and specification in supporting semantic data description using OGC standards. A design artefact is delivered in the form of a concrete Catalogue model and Catalogue interface patterns.
Research Rigour	Design Science research relies upon the application of rigorous methods in both the construction and evaluation of the design artefact.	This research design follows well-founded prescriptions from the IS literature concerning the use of Design science (and in other areas of the study involving qualitative and quantitative methods based on inductive and deductive methodological paradigms). The Design Science construction component was guided by community-grounded requirements, use-case development, literature survey and formative evaluation phases involving iterative development through prototyping. Summative (or goal-oriented) evaluation involved community and developer assessment of functionality and architectural 'fit'.
Design as Search	The search for an effective artefact requires utilising available means to reach desired ends while satisfying laws in the problem	The developed artefact is bounded by theories used in Description (logic)-based semantics. It is also developed within a framework of domain Web Services, modelling and IT standards adhered to by the large community of practise used as a study case

	environment	in this thesis.
Communication of Research	Design Science research must be presented effectively both to technology-oriented as well as management-oriented audiences	Aspects of the research pertaining to the Catalogue communication interfaces have already been peer reviewed and published (Finney and Watts, 2011). Material from this thesis has also been presented at practitioner (community) conferences and will also be published in the literature.

Guangzhi (2009) presents a useful diagram (Figure 3.3.) of the inputs and processes typically involved in Design Science development activities involving prototyping. In overview this approach was followed and encompassed in the overall research design used in this thesis (and is consistent with elements of the methodology as executed, already presented in Figure 3.2).

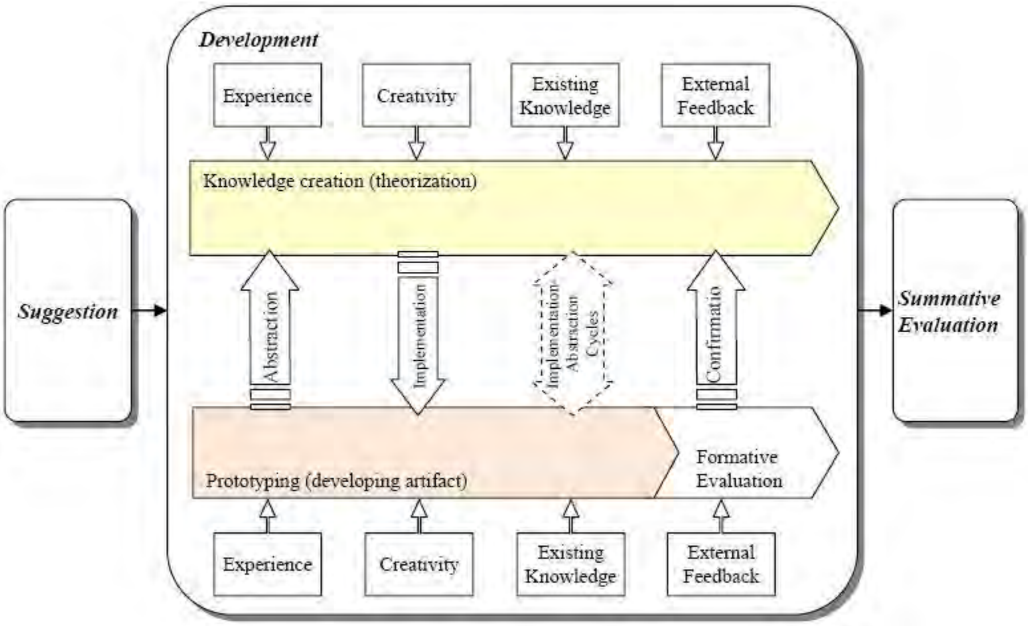

Figure 3.3 Design Science Methodology Encompassing Prototyping (From Guangzhi (2009).

Verschuren and Hartog (2005) have expanded upon the design life-cycle in more detail and have identified six stages (1. First hunch; 2. Requirements and assumptions; 3. Structural specification; 4. Prototype; 5. Implementation and 6. Evaluation). Although presented as a separate, and an apparently post development process, Verschuren and Hartog (2005) argue that evaluation actually takes place during the entire design process. So, ideally they should have really differentiated the last stage as “goal-oriented” or “summative evaluation”. Their conception of evaluation is “the comparison of separate parts of a design process with selected touchstones or criteria (in the broadest sense of the word), and to draw a conclusion in the sense of satisfactory or unsatisfactory”. Evaluation, when performed in this thesis, is viewed according to this definition. It has been useful, in framing thesis research methods, to conceptually utilise the three levels of designing as articulated by Verschuren and Hartog (2005), i.e.:

(a) the plan (on paper) of the design (covering stages 1 to 3 listed above);
(b) the realisation, or carrying out of this plan (which approximately equates to stages 4 and 5); and
(c) the effects that the “use or presence” of the artefact has (covering stage 6).

Categorising the design process in this manner provided a style for structurally presenting research results in Chapter’s 4 to 5. Chapter 4 primarily deals with aspects of the design plan and physical realisation (i.e., categories “a” and “b” above) and Chapter 5 focuses on the effects of use (i.e., category “c”). This categorisation also permitted disambiguation of the evaluation tasks into the following three classes, corresponding to the three levels of design of Verschuren and Hartog (2005):

- **Plan evaluation**: involving an assessment of the quality, adequacy and relevance of the paper design. It is a logical, ethical and empirical check of all of the design requirements and design assumptions.
- **Process evaluation**: covering consideration of how the plan is realised. Primary aims here are to improve the process, and via this, the product of designing, and
- **Product evaluation**: which includes assessment of the results of the design process in terms of the effects of the artefact. Significant aims of this type of evaluation are to determine whether to stop or continue development; to legitimise the activity for stakeholders in terms of the effort it takes and the money it costs; or to motivate the stakeholders to continue their passive or active support.

Several methods were used to undertake these various levels of evaluation at different stages in the Feature Catalogue prototype development life-cycle (see Table 3.2). Since the methods in this thesis did not go beyond a limited-in-scope prototype, no implementation stage was conducted. In such cases Verschuren and Hartog (2005) assert that testing whether the scale model or partial product satisfies the design requirements is legitimate, even though such a test is based on assumptions that relate the test results to the behaviour of the full blown product. Evaluation of the process that led to a partial product implies an assessment of the theoretical line of reasoning that leads to the conclusions that the partial product fits the structural specifications.

Table 3.2 Design Science Research Evaluation Methods Using Five Stages of Design of Verschuren and Hartog (2005).

Stage	Evaluation Type	Thesis Evaluation Methods
Hunch	**Process Evaluation**: An assessment of whether the design goal(s) really cover the desires of the stakeholders; does the designer have an appropriate overview of the social domain and the technological material and knowledge that is available from which to draw ?	1. One multi-disciplinary scientific community (i.e., Antarctic community) chosen as a study case from which three focus groups were selected and used to corroborate the design goals and the utility of Feature Catalogue development. 2. Literature review of the body of knowledge relevant to the problems space.
Requirements and assumptions	**Plan Evaluation:** Entities of interest are the design goals *(evaluated by clarity of stipulation; stakeholder consensus; feasibility; and affordability);* the requirements and assumptions *(evaluated by the acceptability of the granular articulation of requirements; statements on operational performance measures and constraints; relevance to goals and reality; a demonstrated understanding of existing theory and practise)*. **Process Evaluation**: An assessment of design guidelines imposed by existing standards and architectural fit.	1. Design goals were elicited through detailed investigation of typical data capture and data exchange scenarios, desk-top data modelling exercises and through iterative use-case construction with focus groups. 2. Relevant data standards composed in the UML pattern language were critiqued and any design guidelines (or constraints) imposed by these standards were identified, as well as any short-comings. Desired adjustments to patterns were documented in UML and through detailed data dictionaries. 3. Artificial data was used in a desk-top exercise to explore data serialisation. 4. Existing theory and practise was reviewed through literature search and aspects of the data modelling

		process were presented at a practitioner conference and peer reviewed (Finney, 2007).
Structural specification	**Plan Evaluation:** Evidence of an examination of structural alternatives; an assessment of the clarity of the specification and evidence of fitness to design goals.	1. Different approaches to Catalogue development were selected for trial (i.e., an Oracle 11g semantic data store with a SPARQL query interface and an Oracle 11g relational data store with a REST-based interface). XML; HTML and SKOS-based outputs were chosen as alternatives from the REST-based interface. 2. Specifications for data models were created as ER diagrams and concept maps for ontological models. REST-based output patterns were supplied to programmers as sample XML document and XML schema documents (using artificial data).
Prototyping (iterative)	**Plan Evaluation:** An assessment of the prototype Feature Catalogue's functionality against user requirements and contextual criteria. **Process Evaluation:** An assessment of how the designers assure that the Feature Catalogue prototype conforms to structural specifications. **Product Evaluation:** Requires an empirical test of the symbolic representation on paper and the Feature Catalogue prototype to test for mismatches.	1. One focus group participated in iteratively assessing/acceptance-testing Catalogue developments during initial prototyping and development of a test web site for the REST-based Catalogue interface. 2. Demonstration of the test interface was provided to a second focus group for evaluation, before it was finalised. 3. As the prototype was only a partial implementation, the author comparatively assessed instantiated functionality against pertinent concrete specifications and also speculated on the best

		modes for a full implementation through an evaluation of the alternative technologies and approaches used. Implementation feasibility was assessed inclusive of an analysis of existing Antarctic community capability.
Evaluation (Summative)	**Product Evaluation:** An assessment of the effects of a fully instantiated Feature Catalogue and the benefits as a consequence of use. **Process Evaluation:** Identification of what was learnt in the process of design that could be passed on as guidance during full implementation ?	1. A descriptive theoretical assessment of the "goodness of architectural fit" with existing Antarctic data infrastructure and related connected systems was undertaken by the author through argumentation, using knowledge acquired during the course of the research (through community-based participatory activities) and extensive literature review. 2. Issues of practise were documented as guidance for a community-sponsored build process beyond prototyping.

Having reviewed the fundamental paradigmatic assumptions on which the methodologies in this thesis are based, the remainder of this chapter outlines in detail, the specific techniques used. The research activity fell into three broad streams of work. The first two streams, relating to Feature Catalogue development inclusive of the Catalogue interface, are tightly bound, highly interrelated and were some-what cyclically executed. The third stream of work, conducted in parallel with the first two tasks, was executed quite independently (not-with-standing its informative value in unpacking streams one and two). This body of research leveraged a different group of experts and used a more self-contained methodology. The remainder of this chapter is therefore partitioned into two main sections:

Section 3.2 covers activities:

- Related to the overall design of the Feature Catalogue including the tasks involved in: analysing the nuances of data typically exchanged in large scientific networks; subsequently postulating generic data models that are capable of capturing the salient characteristics of exchanged data; interacting with network communities to solicit requirements; analysing current standards relevant to Feature Catalogue development; and formatively evaluating aspects of the design.

- Associated with specifying and then building two limited-in-scope Catalogue artefacts for demonstration purposes (one based in OWL and the other using a more traditional relational model), with a focus on deriving a semantically grounded content model and a widely interoperable Catalogue access interface.

Section 3.3 covers:

- The methods used to assess how domain practitioners currently select ontologies for re-use and a description of the processes harnessed by the author to develop a practical ontology evaluation framework that a community can use to select the ontological components that are necessary to populate the Feature Catalogue.

3.2 Feature Catalogue Design & Development Methods

Design Science Research differs from (ordinary) design practice. The key distinguishing characteristic is that Design Science Research attempts to solve problems that are general in nature, with generic solutions that can be applied in multiple situations. The results (artefacts) should be relevant to typical classes of stakeholders rather than to particular people or organisations. Design practice, on the other hand, solves particular, situated problems with particular stakeholders (Venable, 2009).

Venable (2009) defines a stakeholder as “a person or organisation with an interest in a problematic situation, or in actions taken to ‘improve’ the problematic situation.” Cranefield and Yoong (2007) explain that there are many things that can be ‘stakeholders’ and at many levels, including: individual people, groups, neighbourhoods, organisations, institutions, societies, and even future generations. Ulrich’s (2002) identifies four types of stakeholder that are of relevance to Design Science Research in his development of Critical System Heuristics, which is a framework for identifying what is in or out of scope for any given systems development activity. Critical System Heuristics draws on general systems theory which provides guidance that the scope of a system under consideration must be decided such that it is wide enough to prevent local decisions from causing significant problems in the larger system(s), within which the system under consideration can be thought of as a component. Ulrich’s (2002) four categories of stakeholder are client; decision-maker; professional and witness of interest. Venable (2009) examines the characteristics of each of these stakeholders and concludes that:

- Clients play the role of the person/people/organisation(s) with an interest in solving the problem, or they are the type of hypothetical people with potential future interest in solving problems of the particular type(s) to be addressed by the new solution technology. In summary they are "the set of all members of the generalised class of all people or organisations who could potentially be motivated to solve instances of the generalised class of problem(s)".
- The Decision-maker is "the set of all members of the generalised class of all who might need to decide whether to *employ* the result of the research to the solution of the generalised class of problem(s)".
- The professional is "the set of all members of the generalised class of all who could *apply* the solution technology developed in the Design Science research to the solution of an instance of the generalised class of problem(s)".
- The witness role is that of anyone who represents the interests of those who could be affected by the intervention and who are not able to represent their own interests in the intervention problem formulation, solution design and deployment process.

Given the problem domain described in Chapter 1, regarding the relative current paucity of semantically defined scientific datasets on the Internet, which is hampering systems interoperability and machine-assisted data integration, the issue for this body of research was how to bound the scope of the problem so that it was amenable to investigation. The problem statement can be reduced to a simple goal statement i.e., how can we improve Web deployment of semantically described scientific data? Since "scientific datasets" are terms potentially covering all data ever generated through the act of measurement, observation, modelling or theorising in the history of science, some means of further reducing the problem space is required. The problem was therefore bounded by: concentrating on datasets deployed via Web services conforming to specific types of standards; placing an emphasis on those datasets that are nominally generated through observation and measurement; and narrowing the scope of the type of scientific 'stakeholder' community under investigation. Narrowing the breadth of the stakeholder community was achieved by selecting the Antarctic scientific community as the unit of study. Since this group is highly multi-disciplinary, international and highly distributed, representative collectives from within this community had to be selected.

Using guidance from Ulrich (2002) and Venable (2009) representative collectives should include at least the first three types of stakeholder described above. The fourth role (that of witness) is considered here as being taken up by reviewers and editors of the research literature. The two collectives chosen and invited to participate in this study encompass individuals who variously play the role of client, decision-maker and professional, thus ensuring coverage of the main stakeholder types. At various times, as will be described later, sub-groups from these collectives were harnessed as focus groups to assist in artefact requirements elucidation and research evaluation.

The overall goal of this particular strand of research is to deliver generalisable solutions, in this case Feature Type Catalogue content and interface designs, development methodologies and potential standards enhancements. A desired outcome was for the Feature Type Catalogue prototype, and the activity surrounding its development, to be a catalyst for encouraging ontology uptake and ontology re-use within the study's participating communities for the purpose of semantic dataset description. An aim was to demonstrate the Catalogue's value as a multi-purpose vocabulary repository. This required its seamless integration with the existing data systems deployed by the scientific communities collaborating in the study.

To achieve both the goal and the outcome required working with sub-groups of domain experts drawn from two scientific communities-of-interest, the Scientific Committee on Antarctic Research (SCAR, (SCAR, 2012)) and the Australian Ocean Data Network (AODN, (IMOS, 2012a)), both of which are engaged in Antarctic science and are part of the broader Antarctic science community.

SCAR represents scientists that conduct internationally collaborative, multi-disciplinary science in the Antarctic, sub-Antarctic territories and in the Southern Ocean. Its peak data management coordinating body is the SCAR Standing Committee on Antarctic Data Management (SCADM, (SCADM, 2012)), a 25-member strong international group of national polar data centre managers responsible for building and maintaining an Antarctic data infrastructure. The author has been a member of SCADM since 2006 and was elected as its Chief Officer in 2009. In this role the author's remit is to chart and implement a strategic direction for the ongoing development of the Antarctic Spatial Data Infrastructure.

In contrast to the international membership of SCAR, the AODN is an emerging Australian-centric community focussed on building a national data infrastructure covering marine and coastal observation data. The AODN is managed under the auspices of the Integrated Marine Observing System (IMOS, (IMOS, 2012b)), a National Capability Research Infrastructure System programme, through a funded Development Office (i.e., the AODN DO).

In working with the SCAR community a substantial amount of interaction was with experts from the Australian Antarctic Data Centre (AADC, (AADC, 2012b)), Australia's national polar Data Centre. The

AADC is both a leading member of the SCAR community, part of the AODN consortium and is managed by the author.

In the AODN community, expertise was mainly leveraged from within the membership of the Australian Ocean Data Centre Joint Facility (AODC JF, (AODC, 2012)). The AODC JF is primarily a group of national marine data centre managers (from the following government institutions: CSIRO, Australian Institute Marine Science, Geoscience Australia, Bureau of Meteorology, Navy, and the Australian Antarctic Division) and observation network operators (associated with IMOS) who are seeking to improve publication of, and access to, Australian marine and coastal data. This particular group has an expert-based Technical Committee and a governing Board of which the author is a founding member. See Figure 3.4 for an overview of the relationships between these various entities.

The author's various affiliations with each of the groups mentioned above, combined with the AADC's capacity to help sponsor and promote aspects of the Feature Catalogue development work were intrinsic to fulfilling the goals of the study.

From the preceding declarations it should be obvious that the author is not a neutral observer with respect to the stakeholder communities and in this study has adopted the role of participant observer. The author has had prolonged engagement with the aforementioned stakeholders and this has afforded many opportunities to observe and participate in a variety of activities leading to an in-depth appraisal of community issues. Lincoln and Guba, (1985) assert that research findings are considered to be more trustworthy, when the researcher can show that he/she spent a considerable amount of time in the setting, as this prolonged interaction with the community enables the researcher to have more opportunities to observe and participate in a variety of activities over time.

In the participant as observer stance, the researcher is a member of the group being studied, and the group is aware of the research activity. Participant observation is the process enabling researchers to learn about the activities of the people under study in the natural setting through observing and participating in those activities. Observations made in such circumstances help the researcher to have a better understanding of the context and phenomenon of interest (Dewalt & Dewalt, 2002). This technique is generally used in anthropology and sociology and since observation is filtered through one's interpretive frames it is important, for achieving accurate reflections, that "observations are shaped by formative theoretical frameworks and scrupulous attention to detail" (Schensul *et.al.*, 1999).

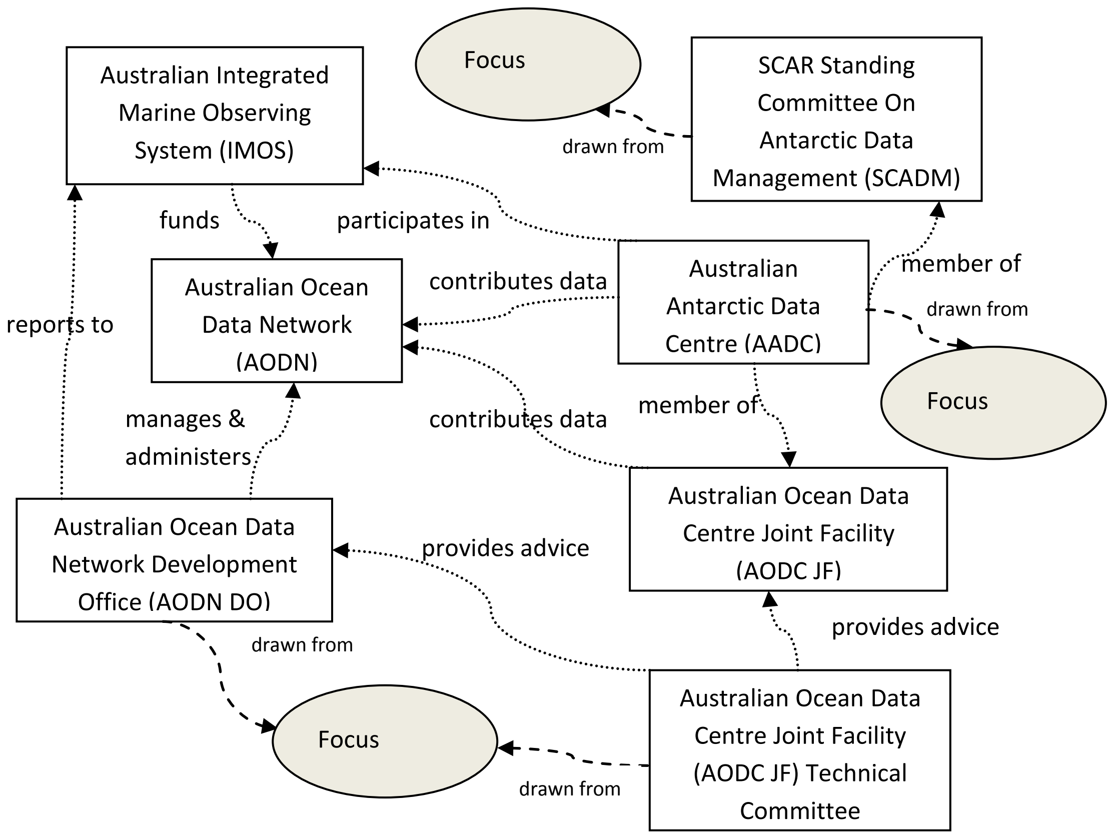

Figure 3.4 Relationships Between Stakeholders In The Antarctic Community

Merriam (1998) calls the stance of participant observer a "schizophrenic activity", because the researcher participates in the setting under study, but not to the extent that he/she becomes too absorbed to observe and analyse what is happening. The question frequently asked is, should the researcher be concerned about his/her role of participant observer affecting the situation. Merriam (1998) suggests that the question is not whether the process of observing affects the situation or the participants, but how the researcher accounts for those effects in explaining the data. In this study observations made whilst ‘participating’ in community activities (such as email lists, strategy formulation, group meetings and distributed technology development) were often corroborated by other purposive information soliciting activities using focus groups and workshops designed to both garner feedback on observations and to generate new insights. Focus groups were selected sets of individuals drawn from the AADC, AODCJF Technical Committee, the AODN Development Office and SCADM who, at various times were invited to participate in discussion sessions where their perceptions, opinions, and attitudes towards Catalogue design and development issues were obtained (see Figure 3.4). In these interactive group settings individuals were free to talk with other group members and responded to questions from the author, who facilitated discussions. Outlined next are the methods used to design, build and evaluate an ontologically-grounded Feature Catalogue which involved an interaction with the SCAR and AODN communities just discussed.

3.2.1 Feature Catalogue Design

Both the SCAR and AODN communities are currently using centralised ISO 19115 compatible metadata registries to manage information about collected scientific datasets. The SCAR metadata registry standard is the Directory Interchange Format (DIF) and is currently mapped to ISO 19115. These systems also provide access to the described data via direct hyperlink, or alternatively via (a de-coupled) reference to published data services. The SCAR system is hosted by the USA-based National Aeronautics and Space Agency (NASA), and is known as the Global Change Master Directory (GCMD, (AMD, 2012)). The AODN community is using the Metadata Entry and Search Tool (MEST, (IMOS 2012c)), which is based on open-source GeoNetwork (GeoNetwork, 2012) technology.

Both the GCMD and MEST metadata systems use controlled vocabularies to facilitate data documentation and data discovery but neither system harnesses the capabilities afforded by a service-enabled Feature Type Catalogue. The MEST tool has a rudimentary Feature Catalogue capability, by virtue of new functionality recently introduced (in 2011) into the underlying GeoNetwork software, but this functionality has not been utilised by the current AODN user community. At the time of commencing this research these flagship portal-based, data delivery systems developed by the AODN and SCAR communities did not permit data discovery based on Feature Type search criteria, despite the capacity for both systems to use OGC-based feature-centric Web services for data exchange. In addition, any remote feature-centric data services deployed by the different data providers within both communities, tended to use provider-specific Web service payload schema. There was also minimal consistency between the semantic representations of any published schema-embedded Feature Types.

Both communities appeared to collectively understand the utility of a shared set of data delivery schema (i.e., Application Schema in ISO and OGC parlance), but were relatively ignorant about how to construct them and how a Feature Type Catalogue might assist with the semantic alignment of the various data-provider services. Community experts within the AADC sub-group were an exception. Having instantiated a Web-accessible (as opposed to service-enabled) Feature Catalogue, prior to the author's arrival as Centre Manager, these community members were aware of some of the uses of such a facility. However, they had failed to gain wide polar community support for their Feature Catalogue's use. Their view of the Feature Type Catalogue, and their use of it, was slanted towards supporting in-house mapping applications and the sharing of (polar) topographic data.

Before any serious design work could begin on a generically applicable, ontologically-grounded Feature Catalogue, it was therefore necessary to seed conceptual ideas within both (the SCAR and

AODN) communities about the potential roles that a Feature Catalogue could perform in order to muster interest in its development. Ultimately, the Feature Catalogue design work relied on the specification of functional requirements, mainly through the construction of use-cases with community involvement and through independent data modelling exercises (which were later peer reviewed). A brief introduction to use-case techniques as a tool for requirements development is given next. Chapter 4 describes in detail the results of the application of these techniques.

Use Cases As Tools For Defining Functional Requirements

In systems engineering functional requirements capture the intended behaviour of a system. This behaviour may be expressed as services, tasks or functions the system is required to perform. A use case defines a goal-oriented set of interactions between external actors and the system under consideration. Actors are parties outside the system that interact with the system (Jacobson *et al*., 1992). An actor may be a class of users, roles users can play, or other systems. Cockburn (1997) distinguishes between primary and secondary actors. A primary actor is one having a goal requiring the assistance of the system. A secondary actor is one from which the system needs assistance. A use case is initiated by a user with a particular goal in mind, and completes successfully when that goal is satisfied. So, the use-case describes the sequence of interactions between actors and the system necessary to deliver the service that satisfies the goal. It may also include possible variants of this sequence, e.g., alternative sequences that may also satisfy the goal, as well as sequences that may lead to failure to complete the service because of exceptional behaviour, error handling, etc. In a use-case, the system is treated as a "black box", and the interactions with the system, including system responses, are as perceived from outside of the system. Thus, use cases capture who (actor) does what (interaction) with the system, for what purpose (goal), without dealing with system internals. Ideally, a complete set of use cases specifies all the different ways to use the system, and therefore defines all behaviour required of the system, bounding its scope.

The use cases developed in this study are written in an easy-to-understand structured narrative using the vocabulary of the domain. Cockburn (1997) argues that this textual method is engaging for users who can easily follow and validate the use cases, and the accessibility encourages users to be actively involved in defining the requirements. The benefits associated with the ease of writing and reading cancel out the disadvantages of imprecision and redundancy (Cockburn, 2000) which are hallmarks of textual descriptions (Genova *et al.,* 2005). There were situations, however, for example during a workshop in which functional requirements were being finalised and validated with an AADC focus group, that graphical Unified Modelling Language (UML) diagrams (Booch *et al.,* 1999) were used to capture aspects of system functionality, primarily because programmers from within the AADC who were to be engaged to assist with prototyping were familiar with the notation (see Figure 3.5).

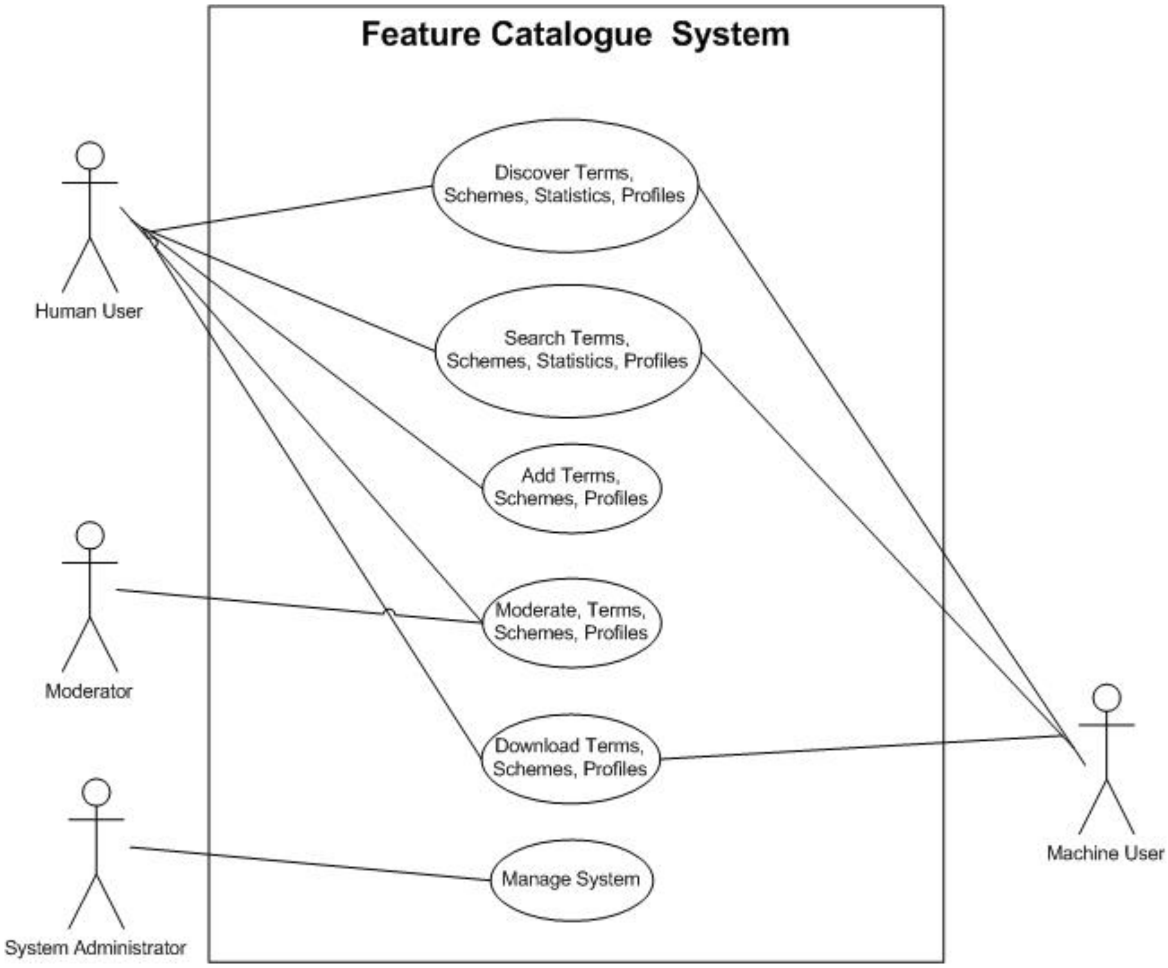

Use Case 3	**Add Terms Schemes & Profiles**
Actors:	Human User
Purpose:	Add terms, term relationships, schemes, themes and profiles
Overview:	A human user navigates an interface and is able to add: terms; relationships between terms; terms belonging to specific schemes and profiles; specific schemes and their themes.

Typical Course Of Events

Actor Action	**System Response**
1. Human logs-in to FC application	2. Presents a browse, search and administrative interface.
3. Human elects to add specific terms, schemes, profiles, term relationships.	4. Adds items to temporary cache. Marks items as in-review. Notifies moderator of new additions.
5. Human completes adding items and exits application	6. Closes application

Figure 3.5 Example of a use-case diagram used during an AADC focus group session and a high level description of use-case No.3 (Add Terms Schemes and Profiles).

Defining functional requirements through use-cases has the benefit that they can readily be used to verify and validate the system. These tasks can be performed before and after system construction. A systems development team can perform a kind of “thought experiment” by walking through the use-case steps and sequences (usually depicted in diagrammatic form) to check off various dependencies and interactions. Anda and Sjøberg (2002) and Hansen and Miller (2002) describe four approaches to verifying & validating a series of use-case models. Elements of all of these approaches were used at different stages during requirements gathering and use-case construction. They include:

- Inspections (verify & validate) – the act where an individual or a team looks at the use-cases according to pre-defined criteria to verify their adherence to standards and specifications.
- Reviews (verify & validate) – involve multiple readers examining the different use case artefacts (text, diagrams). Reviews should involve customer representatives and other stakeholders.
- Walkthroughs (validate) - a form of review where a use case or a business scenario (comprised of several use-cases interacting) is actively presented (usually by the author), and possibly role-played in-order to examine the flow of events.
- Prototyping (validate) – This is based on creating rapid prototype (often only screen mock-ups) to demonstrate to stakeholders the behaviour depicted in the use case. The advantage of this approach is the visibility of the understanding captured by the use case.

Despite the considerable benefits of the use-case approach for requirements gathering there are limitations. Genova *et al.* (2005) and Sinnig *et al.* (2009) point to the controversy, inconsistent use, and free-flowing interpretations of use case models. They assert that often experts, who are widely recognized in the systems engineering community, fail to agree on the meaning of concepts, particularly those coined in the UML use-case meta model (OMG, 2000). Consequently, they believe that use case models (graphical and textual) are often ambiguous, and there is an unnecessary divergence of practice. These limitations aside they are still one of the most widely used techniques for defining requirements and their simplicity made them amenable for use in this study, particularly given the divergent nature of the stakeholder expertise within the participant community.

Stakeholder Interactions - Use Case Development & Data Modelling

As mentioned previously, some use-cases were informed by the results of data modelling and data modelling of itself helped define structural patterns that fed into requirements. Early in the research, the author evaluated two candidate data models (i.e., Climate Science Modelling Language: Woolf, 2007 and Observation and Measurement: Cox, 2006) as a suitable basis for developing an ontological model to support the exchange of marine science data within the AODN infrastructure (Finney, 2008). At the time, the AODC JF Technical Committee was interested in better understanding the options available for developing marine community Geography Mark-Up Language (GML)-based application schema, and the author offered to investigate the topic. The decision by the AODN community to use OGC standards-based Web services for data exchange within the AODN, clearly pointed to the need for a Feature Type-centric ontology model and the author reasoned that software would be necessary to aid the community in managing these Feature Types and their associated ontologies (inclusive of any inherent vocabularies). The broad ideas about the development of a Feature Catalogue were duly exposed to AODC JF Technical Committee experts in a one day workshop in February 2008, although they had already been raised cursorily with community members and mentioned in Finney (2007).

At the same time as ideas were being canvassed within the AODN community, opportunistic discussions regarding the Catalogue had already commenced with SCAR community members, primarily during the annual meetings of the SCAR Standing Committee on Antarctic Data Management, particularly in Rome in 2007 (JCADM, 2007) and following these initial discussions, again in Amsterdam in 2009 (SCADM, 2009). The author also flagged the desirability of semantically enabling SCAR data services in the SCAR Data and Information Strategy (Finney, 2009) which in 2007 was in draft form and was the main topic of discussion at the 3-day Rome meeting (mentioned previously). A dedicated workshop was subsequently conducted at the AADC in June 2009, which finalised all of the use-cases and the desired base-line functionality of the Feature Type Catalogue. By this time a thorough critique of existing (relevant) OGC standards had also been pursued and it was then feasible to assess standards against the requirements articulated by the Antarctic community. Progressively, the author had also already crafted a draft conceptual design for the Feature Type Catalogue content model using UML notation, which was iteratively refined as requirements were bedded down.

Prior to the finalisation of the uses-cases and whilst functional requirements were still emerging, a literature review was undertaken to assess existing ontology design patterns that might be suitable for grounding the Catalogue content model so that it could be expressed in an ontological form. The possibility of using an ontological content model that dove-tailed with an upper ontology was also explored. It was reasoned that it would be useful to anchor the Catalogue's conceptual content schema in a common semantic context. At this point in the research life-cycle, the parallel strand of research activity focusing on RQ1.2 (ontology selection and evaluation), had evolved to a point where outcomes from expert interviews could be used to inform the author's choices regarding ontology patterns and issues of ontology re-use in the real-world task of Catalogue design (and development). Expert opinions and information were pertinent in helping to select DOLCE (Masolo *et al.*, 2003) as the upper ontology to anchor the Feature Catalogue semantics. The experts referred to here are those who were assembled specifically to assist in addressing RQ1.2 and had expertise in ontology development.

The processes outlined next describe in overview the techniques used to instantiate the prototype Feature Type Catalogue(s) with a view to encouraging Catalogue uptake and its further development by the AODN and SCAR communities. Although this narrative, as unpacked so far, infers a linear progression from "design" activities through to "build" tasks, there was an iterative element to the design/build phases (see Figure 3.2). Actually implementing models and functionality in some cases lead to design revisions.

3.2.2 The Feature Catalogue Build Process

The Feature Type Catalogue build-phase harnessed some existing ontology design patterns (true to the spirit of re-use) and the first Catalogue content model that was built used OWL and a specialisation of DOLCE (upper ontology) concepts. This was achieved using Oracle™ 11g's support for RDF triple storage, TopBraid Composer™, the Protege (version 4.0 beta, (Stanford University, 2012a)) ontology development tool and concept mapping software - Cmap, (version 5.04) developed by the Institute For Machine and Human Cognition and SPARQL (the RDF query language for databases). The detailed results of this exercise are discussed in Chapter 4 and a summative evaluation of these results follows in Chapter 5.

In addition to the OWL-based development, further work was undertaken to instantiate the UML Feature Type Catalogue model as a relational database (again using Oracle™ 11g), maintaining alignment with the OWL model as far as was practical. Development of this second artefact primarily provided an opportunity to assess the use of less formal semantic encodings and potentially more user-friendly services for accessing Catalogue content (than SPARQL). Results of this facet of the research are also presented in Chapter 4 and discussed and evaluated in Chapter 5.

It should be mentioned that some aspects of this study were made possible only through the support provided by staff of the AADC. The AADC is a key data and service provider in both the SCAR and AODN communities. At the commencement of this research the AADC hosted one of only a few, Web-based Feature Catalogues. The AADC Feature Catalogue is not, however, service-enabled, ontology driven, nor designed to interface with ISO 19115 metadata catalogues or associated data systems. As has already been highlighted, this existing AADC Catalogue has essentially been used to manage features found in topographic maps and is little more than a text-based data dictionary. Although it is a SCAR community endorsed product, with multi-user access, there has been very limited use of the tool by community members outside of the AADC. For example, there were only 15 new Feature Types added to the Catalogue between the years 2007-2011 and all of those that were contributed, were added at the request of AADC staff.

However, the AADC had a desire to upgrade its existing Feature Type Catalogue to take advantage of new technologies, but was uncertain about what elements required improvement and more importantly which improvements would lead to greater community uptake of both the tool and its contents. This study capitalised upon this situation and the author tasked two AADC staff members to undertake software coding for the service interface component of the prototype Feature Catalogue (which was based on a relational data store). In the main, all prototype development in this research was seen as a limited-resource investment activity to practically demonstrate to other community members many of the conceptual ideas canvassed in stakeholder workshop and community discussions.

In the case of the relational (Feature Type Catalogue) solution, various interface design specifications were reviewed by the author and ultimately a REST-based protocol was selected. The author then designed the query patterns for accessing Catalogue contents and the patterns for the delivery schema (with XML, HTML and SKOS language output options). For each type of query pattern a sample schema was prepared to help the software developers understand the coding task (and to permit verification and software evaluation). Output schema design was facilitated by the use of an XML language editor, Altova™ XMLSpy and the RDF output was validated using the W3C RDF Validator (W3C, 2007).

The REST interface component was coded in ColdFusion™ version 9. The REST service accesses a Catalogue populated with basic sample data loaded into Oracle 11g, which was facilitated by the use of Oracle's™ SQL Developer tool. The AADC infrastructure which hosts the Catalogue prototypes uses an Apache 2.2 web server running on a Windows Server 2003.

3.2.3 Feature Catalogue Evaluation

Formative evaluation techniques were used iteratively throughout the Feature Catalogue development activity to assess alignment with initial design criteria and goals. As previously explained in some cases, the act of formative evaluation, of itself, lead to some modification of the initial design criteria. The development process also anticipated some cross-fertilisation between the information gathering steps outlined next for the ontology selection and evaluation stream of work and this Feature Catalogue implementation task. It was therefore important to unpack both streams according to a timetable that permitted outcomes from one stream to become potential (information or guiding) inputs to the other.

Summative evaluation of the prototype Feature Type Catalogue involved its exposure to the AODN community, in order to gain feedback on its utility during a 2-day workshop held at CSIRO in September 2010. A formal presentation was also made to a broader group of AODN data managers at a workshop sponsored by the Australian National Data Service (ANDS) in November 2010 (ANDS, 2010). The prototype Catalogue was also presented, and discussed, at the SCAR SCADM meeting in Buenos Aires in August 2010 (SCADM, 2010). In addition, a paper was prepared and submitted for peer review on the service-based aspects of the prototype Feature Catalogue (Finney and Watts, 2011).

3.3 Practical Ontology Selection and Evaluation Methods

Since this research is concerned not only with demonstrating a Feature Catalogue artefact, but also with providing guidance to communities about how to populate it, it was necessary to analyse how leading ontology developers, within scientific communities, were (and are) approaching the problem of selection and evaluation of semantic components.

In terms of investigating community-based, practical ontology selection and evaluation processes there were four key stages in the overall methodology. The first stage involved a screening exercise to identify a range of community-based ontology developers who could later provide insight into the evaluation processes and criteria used in real-world ontology development activities, where re-use had been practised. Information gathered through the screening survey not only identified suitable study participants but also helped stratify participants to aid subsequent phases of analyses. The second stage involved the qualitative capture of ontology developer’s experiences and opinions through formal interviews, with a view to identifying a practical ontology selection and evaluation model based on interpretative analysis of the interview material. Thirdly, through another survey instrument, experts were asked to quantify the relative importance that they would place on evaluation criteria within the derived model. At this point expert consistency in responding to the survey was empirically tested and patterns in expert responses were examined. The fourth stage culminated in testing the model’s applicability and validity with study participants and matching metrics derived from expert interviews (and the literature) with each criterion. This last stage also involved the postulation of ways in which the model could be applied.

The methodologies used in each of these stages are discussed in more detail next.

3.3.1 Expert Screening Survey

The Screening Survey's primary role was to locate people who were suitable to provide insight into how practical ontology development tasks, involving re-use, in scientific data exchange scenarios are undertaken. It was anticipated that the pool of people suitable, willing and able to participate would be limited given the relatively immature nature of ontology re-use within many scientific communities and disciplines. The SCAR and AODN communities had very few individuals practised in the art of ontology development so the net was cast much more widely (and outside of these groups) to secure ontology practitioners for the study.

The survey was a mechanism to help identify those people who could provide relevant and detailed information. Because of the potentially small numbers of experts available and to transcend any jurisdictional biases, both national and international contacts were sought to maximise the study cohort size, the domain and expertise mix.

The secondary functions of the Screening Survey included:

- Introducing respondents (particularly those who would later go on to be interviewed) to the research questions, key terms and concepts that would recur throughout the research and to give these people an opportunity to start thinking about selection and evaluation issues in preparation for more in-depth questioning.
- Extracting information about an individual, their associated community(s)-of-interest and their interpretation of key concepts to help stratify (expert-derived) data which would later be collected through in-depth interview.

The survey itself consisted of a mixture of open and closed questions, conducted via a commercial Web-based survey instrument (QuestionPro, (QuestionPro, 2012)), powered by SurveyConsole™. A survey template was built using wizards available in QuestionPro and where possible automatic field validation checks were included to minimise the chance of respondents only partially completing the survey (see Figure 3.6 for a snapshot of sample form validation). Appendix 1 contains a full list of the survey questions.

The survey targeted ontology developers and key personnel in communities known to be developing ontologies designed for application in either an OGC or TC 211 standards suite environment and/or for use in scientific data exchange scenarios. Since Antarctic Science covers key disciplines such as oceanography, biology, ecology, meteorology, physics of the upper atmosphere, geology, glaciology, medicine and limnology, an effort was made to contact experts distributed across these disciplines. It wasn't important that these experts be actively engaged in Antarctic science, it would suffice for them to simply be familiar with their disciplinary practises. The screening exercise used purposive sampling (i.e., sampling where the researcher chooses the sample based on who they think might be appropriate for the study). Fourteen respondents completed the survey.

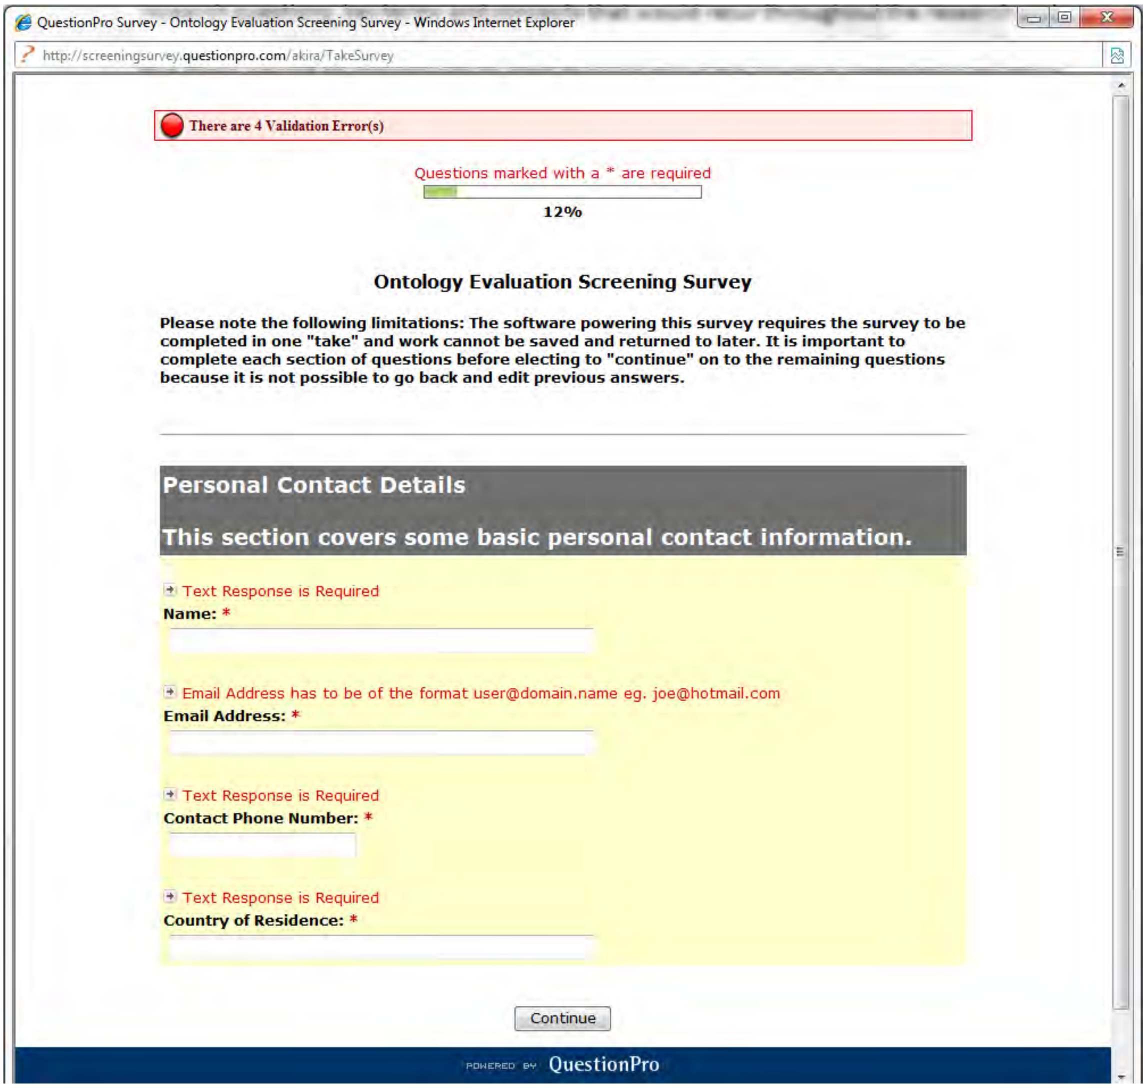

Figure 3.6 – Typical automatic form validation in the Screening Survey

All experts who responded positively to emails requesting participation, bar one, were considered suitable. The person considered unsuitable declared that they had not developed ontologies for use in scientific data exchange scenarios, nor had they practised ontology re-use.

The survey design followed standard survey design best practise (Neuman, 1999) with regards to:

- ethical considerations;

- question formulation and sequencing (e.g., avoided: slang, jargon, ambiguity, confusion, emotional language, double-barrelled questions, false premises, leading questions, overlapping categories and bias);
- survey length and ease of completion (i.e., it was kept relatively short and delivered on-line using radio buttons for option selection and question forking depending on responses in order to improve logic flow) and
- piloting with colleagues to improve structure and comprehensibility prior to use.

Initial contact with potential respondents was made by email (Appendix 2) and followed up soon after by phone to seek the contact's participation in the survey and to urge their participation.

Potential participants were given a dead-line for completing the survey and were phoned (or emailed) again as a reminder if they failed to register a response by the closing date. Because the survey provided an option for the respondent to nominate some-one else that they thought could be of assistance in this research, any nominated individuals (not already targeted) were contacted after the survey dead-line had been reached and were also asked to provide a response by a second closing date. By using this approach (i.e., one which incorporates a limited type of snowball sampling) it was possible to identify an adequately sized cohort of people to draw upon for the research. Although the sample size (of 13 experts) was considered small and some disciplines were not covered (e.g., glaciology, medicine), there was a good mix of disciplines and nationalities represented, sufficient to give confidence that the phenomena of interest would be covered by the participating group.

As explained earlier, largely for ethics purposes, each participant who was willing to take the survey was provided with an information sheet which further explained the purpose of the study and each was asked to sign a Consent Form (see Appendix 3). To ensure that the Screening Survey questions made logical sense from a respondent's perspective, the on-line survey was piloted by asking several colleagues to take the beta version of the survey and to provide feedback on both its structure and content. Using their feedback for fine-tuning purposes the survey template was modified before it was released operationally.

Primary analyses of the screening data involved establishing expert suitability for interview and their willingness to participate in a subsequent quantitative survey. Experts were also classified by community-of-interest type, community governance model, domain of practise, ontology development experience, and by self-assessed skill level. These stratification levels were thought to be useful in seeking to explain any subsequently emergent patterns in the interview and quantitative survey data (that might be of interest in addressing RQ1.2.2.2).

A review of respondent's answers to the questions posed in the survey regarding the definition of "terms" also provided another potential way in which to segment and subsequently analyse data derived from the interviews. Understanding a respondent's terminological interpretation also helped place their subsequent interview answers in an appropriate context. Paslaru Bontas Simperl and Tempich (2006) found a serious lack of terminological knowledge amongst practical ontology developers that affected interpretation of their survey results. It was hoped that by encouraging study participants to think about, and declare their understanding of key terms that were used in this study, interpretation issues that could confound results could be explicitly addressed early (e.g., during interviews).

3.3.2 In-depth Interviews With Community Ontology Developers

Once experts had agreed to participate in the study, an appointment was made to interview them. The purpose of the interview was to glean information from participants that would help to answer RQ1.2 *"what typifies an expert-grounded ontology selection and evaluation framework that can support a multi-disciplinary Antarctic science community using web services ?"* A series of open-ended questions were pre-formulated and used as prompts (see Appendix 4 for a list of these questions). These questions were only used if the expert failed to address the key issues that the author thought required discussion, otherwise the expert was encouraged to range over any material they thought relevant to the topic of ontology selection and evaluation as applied by them, or their community.

Because interviewees had all completed the Screening Survey, clarification could be sought on any ambiguous responses that they may have provided in the Screening Survey, prior to launching into the main substance of the interview. Any references that survey respondents supplied regarding community-centric information, or about ontology resources was perused prior to interview so that any questions arising regarding this material could also be addressed with the interviewee.

A research design goal was to secure at least 3 separate people for interview from each scientific domain that would be represented in this study. Ideally, this number should have been much higher for the research results to be as robust as possible, but the difficulty of (a) identifying suitable people and (b) then gaining access to these extremely busy practitioners was an anticipated challenge of the study. Although a minimum of three experts ended up being associated with each of the six major disciplines covered (i.e., hydrology, biology, oceanography, geosciences, and atmospheric physics, meteorology) this was achieved because experts had overlapping interests across more than one discipline. One participant identified themselves with a discipline they characterised as “energy” – but the majority of their work had been in two other disciplines. In summary, there was some replication of expertise in terms of disciplinary coverage but because of the propensity for experts to work across disciplines, replication of experts (as opposed to disciplinary expertise) was relatively low.

All interviews were scheduled to run for approximately one hour’s duration with the author as interviewer. All interviews were recorded with audio equipment (i.e., a SONY ICD-P620 Recorder). Hand drafted notes were also made both during and straight after the interviews to record any nuances that may not later be discernible from transcriptions of the audio recordings. Although face-to-face interviews were preferred, many participants were located inter-state or internationally with respect to the interviewer and remote communication techniques had to be employed. Initially Skype (Microsoft™, 2012) was explored as an option for video communication but poor local bandwidth precluded its use. Ultimately, one interview was able to be conducted via a dedicated video-conferencing facility, four were conducted face-to-face and the remainder were performed by phone. All interviews were subsequently transcribed into written form for later analysis.

Interviews were scheduled over a period of several months and conducted serially with time in between interviews so that information obtained from each interview had the capacity to inform discussions in each subsequent interview.

Desired outcomes from these interviews were:

- the identification of key ontology selection and evaluation criteria and their categorical relationship (i.e., dimensionalisation), plus an assessment of whether criteria and/or methods applied by experts were consistent with those described in the literature (RQ1.2.1);
- an enhanced understanding of how communities of practise manage and govern ontology development processes (involving re-use) and the measures they use to assess ontology evaluation criteria (RQ1.2 & RQ1.2.3).
- Any resulting dimensionalised model of the evaluation criteria would go on to be used to gain a deeper understanding of the importance placed on criteria by experts, when using these criteria in ontology selection processes (RQ1.2.2.1).

Qualitative coding (i.e., open, thematic, template) techniques (Neumann, 1999; Dey 1993; Crabtree and Miller, 1992; King, 2004), were used to examine the resultant interview data in order to deconstruct the information and in parallel look for any emergent concepts, concept clusters and/or concept relations. The results of coding are presented in Chapters 6 and 7. Some of these data were used to construct an ontology selection and evaluation model. Theorising resulting from data capture and analysis is presented in Chapter 7. In framing the selection and evaluation model, which ultimately needed to be used as input to a decision-making processes, several factors (Baker *et al.,* 2002) were considered:

- **Completeness:** Does the model cover all of the criteria mentioned by experts that are generally agreed as being required to make 'good' selections ?
- **Criteria clustering:** It was considered useful to be able to group together evaluation criteria into a series of sets that related to separate and distinguishable components of the overall objective for the ontology selection decision. It was felt that this would facilitate high level views of the issues, and in particular help domain practitioners better realise/visualise trade-offs that may need to be made in judging between competing options.
- **Redundancy**: Has the coding process inadvertently created duplicate or redundant criteria ?
- **Operationalisation:** Have the criteria been sufficiently well-defined in order to be able to be assessed and measured ?
- **Mutual independence of preferences:** Is it possible to assign preference scores for ontology alternatives on one criterion without knowing what the ontology alternatives' preference scores are on any other criteria? Criteria independence is important for various forms of multi-criteria decision analysis techniques (discussed in the next section).
- **Number of Criteria:** An excessive number of criteria leads to extra analytical effort in assessing input data and can make communication of the selection process more difficult. It was important to ensure that the model structure was no larger than it needed to be.

A software product called Atlas Ti™ was used to mark-up each interview transcript. The mark-up was coded and additional notes could be appended to codes and/or segments of the transcript. Microsoft™ Excel was later used to tabulate and plot data extracted from Atlas Ti™.

3.3.3 Quantitative Survey – Relative Importance of Model Evaluation Criteria

To solicit quantitative assessments from experts on the relative weight that each model criterion might exert during evaluation of ontology alternatives (RQ1.2.2 and RQ1.2.2.1), a Microsoft™ Word-based survey instrument was devised. This survey consisted of 80 questions, each of which asked the respondent to perform a pair-wise comparison of individual criterion inherent in the qualitatively derived evaluation model and to rate the dominant criterion in each pair according to a set numeric scale. The technique used to analyse these data, which were transformed into matrices, was the

Analytical Hierarchical Processing (AHP) method (Saaty, 1980). Not-withstanding valid differences of opinion within the sample group, the goal of this exercise was to arrive at a group consensus. The Delphi survey approach pioneered by the Rand Corporation (Dalkey and Helmer, 1963) and aptly demonstrated for use in ranking-type situations by Schmidt (1997) was originally favoured to iteratively control feed-back from participants in order to gain a group view. Using this approach, participants would be given the opportunity to change their weightings or rankings after viewing the weights applied by others in successive rounds of the survey. This activity could have required one, two or perhaps as many as four more passes through the survey.

By the time this phase of the research was reached, however, the response rate of participants was down to approximately 50% of the original cohort post the screening survey. The high drop-out rate was almost certainly due to the fairly onerous task of responding to 80 survey questions, particularly given each of the experts had individually identified as being time poor. The extreme risk of alienating the currently still invested cohort of experts determined that a different approach for arriving at a group view would be required. Instead of getting the group to gradually converge on a common set of values, the statistical spread of the data was evaluated and a suitable measure of central tendency was chosen as being representative (after calculating and inspecting the Geometric Mean, Arithmetic Mean, Median and Perth Measure for expert rankings across all of the dimensions and sub-categories of data).

Before any central tendency measures could be computed, however, internal consistency checks were performed for each of the eight sets of expert responses (received from those who had remained in the study). Internal consistency here refers to the degree to which each expert rated pairs of alternatives, in a logical manner. Initially, datasets with inconsistencies of greater than 0.10 (Saaty's upper recommended limit for his computed Consistency Ratio, which is discussed later) were identified. To assist experts to improve their consistency scores, an 'R' (open source statistical software package, (GNU Project, 2012)) program was developed which iteratively computed the minimum changes that could be made to data within inconsistent matrices to bring them closer to an acceptable consistency ratio. Although 0.10 is the ratio recommended by Saaty (1980), for this study 0.14 was selected as the cut-off. The justification for this figure is discussed in detail in Chapter 6. Each expert with inconsistent data was contacted by email and asked to review their responses in light of the detected inconsistency. They were each provided with a set of changes which would bring their data to an acceptable level of consistency and asked if they would accept these modifications. The degree of consistency across evaluators in arriving at priority weights was also assessed using Kendall's co-efficient of conformance (W) (Kendall and Gibbons, 1990) and the expert group's similarities/dissimilarities with respect to the ratings provided were inspected using multi-dimensional scaling (MDS: Steyvers, 2002). Kendall's (W) has been used previously in other AHP-centred research to examine between 'evaluator' consistency (Smith and Gannon, 2008; Ware, 2009) as has MDS for inspecting patterns in evaluator ratings (Chen *et al.,* 2008).

Information was gathered from (and about) participating experts throughout the course of this phase of the study and experts were progressively stratified according to a wide range of criteria (e.g., discipline; expert skill type; team experience; ontology application area). This stratification subsequently permitted a detailed evaluation of reported importance ratings aimed at detecting any explainable patterns in the way that different individuals, or clusters of experts, had responded (RQ 1.2.2.2).

Following communication with experts regarding any inconsistent data, the measures of central tendency were computed for each criteria, sub-category and dimension present in the derived model. Of these measures, the normalised geometric means were subsequently chosen as being the most representative of a "group" result. Histograms of all means (for each level in the model hierarchy) were prepared and the eight sets of expert data were over-plotted to show the distribution and variability of the individual data as compared to the means.

To provide participants with an opportunity to give feedback on the suitability of the normalised geometric mean as a group result (relative to other central tendency measures) and to enable them to review/revise their own ratings; and provide commentary on the utility of the model, these plots along with a weighted model were emailed to the eight remaining participating experts. In this email three questions were posed:

(a) Do you believe that the normalised geometric mean values (selected to represent the "group" result) are a good enough approximation (cognisant of the spread of expert values). If you don't, could you elaborate as to why (indicating how you would intuitively change either ranks, and/or actual histogram proportions in any graph, possibly in light of the weights given by others) ?

(b) Are there any aspects of the model structure (as opposed to weights) that you would change (if so, could you explain the reasoning behind any suggested modifications) ?

(c) Could you envisage using this model (regardless of whether it was anchored in an AHP implementation framework – which is a very specific technique for undertaking multi-criteria based assessments) - in exercises conducted by your community-of-interest that are designed to assess the suitability of ontologies for re-use ?

Feedback received from emails sent to experts who participated in the qualitative survey assisted with evaluation/validation of study findings.

For the purposes of additional research validation, specifically related to development of the weighted ontology evaluation model, it was decided to try to leverage those experts who did not reply to the quantitative survey to the study's advantage, by co-opting them back into the study. Five experts who did not provide comparative data from which the weights were derived, were therefore sent emails inviting them to comment on similar issues to (a) – (c) above (weight allocation, model structure and model application). Given this group of (five) people were not involved in criterion weight allocation their views were of considerable interest, particularly with respect to the generalisability of the model's application. The relatively low response rate from the invitation to comment, however, reduced the utility of this step. When it was clear that all responses had been received, the selection and evaluation model was finalised.

Analytical Hierarchy Process

Since the AHP technique played a central role in the method used to assess how experts weighted dimensionalised ontology evaluation criteria it warrants further discussion in terms of the foundations on which it is based, its potential limitations and why it was selected for use.

The problem to which AHP was applied is essentially one where a decision has to be made by a group of people about which ontologies they will choose, in whole or in part, for use in supporting the semantic description of some facet of their data. This is usually not a once-off process and is continuous through time (as more of the domain of discourse requires semantic encoding or because terminology changes over time). It is inherently a decision-making problem. To make good decisions clear objectives are required (e.g., about the requirements or benchmarks that an ontology must meet). Ideally, these objectives should be specific, measurable, agreed, realistic and sometimes time or effort-dependent. This implies that the options available for satisfying stated objectives need to be assessed in light of criteria which can reflect performance of a particular option in meeting any stated objectives. Each criterion should therefore be measurable, at least in a qualitative sense in order to aid the selection process, which ultimately requires human judgement. In ontology selection it is inevitable that many criteria will need to be assessed in order to arrive at an informed decision. This type of decision problem belongs to a class of problems which are usually addressed via multi-criteria decision analysis (Dyer *et al.,* 1992; Keefer *et al*., 2004; Bragge *et al.,* 2010).

A standard feature of multi-criteria analysis is a performance matrix, or consequence table, in which each row describes an option and each column describes the performance of the options against each criterion. Multiple criteria analysis techniques commonly apply numerical analysis to a performance matrix in two stages:

1. Scoring: the expected consequences of each option are assigned a numerical score on a strength of preference scale, for each option for each criterion. More preferred options score higher on the scale, and less preferred options score lower.
2. Weighting: numerical weights are assigned to define, for each criterion, the relative emphasis that a particular criterion should be afforded, with respect to other criterion.

Mathematical routines are then used to combine these two components to give an overall assessment of each option being appraised.

There are essentially two categories of multi-criteria decision making problems: multiple criteria discrete alternative problems and multiple criteria optimization problems. “Discrete” alternative problems are those involving sets of alternatives which are typically modestly-sized collections of choices. They are more likely to be modelled with uncertain values for criteria, than multiple criteria optimization problems. In “optimisation” problems, feasible sets of alternatives for problems usually consist of a very large, or infinitely many alternatives, defined by systems of equations and inequalities that identify the feasible region for the decision variables. In such problems because the vectors of alternatives may have many components and the number of equations and inequalities may be large, the feasible regions may be complex (Bragge *et al.,* 2010). Because of these different problem types, different families of approaches have evolved for solving them.

AHP is a technique used for solving multiple criteria “discrete” problems (e.g., choosing one ontology from a number of alternate ontologies) and it attempts to represent aspects of a decision maker’s utility, or value function mathematically and then applies these results to estimate the alternatives’ (expected) utilities (Forman and Gass, 2001; Bragge *et al.,* 2010). The main role of techniques like the AHP is to deal with the difficulties that human decision-makers have been shown to have in handling large amounts of complex information in a consistent way. As the complexity of making choices rises, people tend to simplify their decision-making processes by relying on simple heuristics (Bazerman, 1998; Iyengar and Lepper, 2000). AHP and other, similarly focussed techniques are designed to help overcome such limitations by imposing a disciplined structure which directs attention to criteria in proportion to weights that have been pre-agreed.

AHP and other linear additive methods are built upon Decision Theory (see Hansson, 2005). Decision Theory is concerned with identifying the values, uncertainties and rationality that are inherent in making optimal decisions. A main assumption embodied in Decision Theory is that decision makers wish to be coherent in taking decisions. That is, decision makers would not deliberately set out to take decisions that contradict each other. Solving general decision problems is dealt with by assessing probabilities of consequences conditional on different choices of experiments (and acts) and by assigning a utility function to the set of consequences according to some scheme of value, or preference of the decision maker (Encyclopedia Britannica Online, 2012). An optimal solution consists of an optimal decision function, which assigns to each possible experiment an optimal act that maximizes the utility, or value, and a choice of an optimal experiment. Decision Theory draws on Utility Theory (Read, 2004), which expands on this notion of coherence, or consistency of preference, and proposes some simple principles of coherent preference, such as the principle of transitivity: e.g., if A is preferred to B, and B to C, then A should be preferred to C, which is a requirement if preference is to be expressed numerically. Keeney and Raiffa (1976) extended Decision Theory so that decisions with multiple objectives could be analysed.

It is important to note that whilst AHP is built upon Decision Theory, Gass (1998) asserts that AHP does not assume transitivity and, given an intransitive situation, the decision-maker has available procedures for finding and mitigating against it, if deemed appropriate. Further, (Saaty, (1990) explains "from its axioms to its procedures, the AHP has turned out to be historically and theoretically a different and independent theory of decision making from Utility Theory". As will be demonstrated later in this discourse, this departure from Utility Theory is significant as it is at the heart of some concerns expressed about the AHP methodology.

At the core of the AHP lies a method for converting subjective assessments of relative importance to a set of overall scores or weights. It is the most widely applied multi-criteria decision analysis method (e.g., see Forman and Gass (2001) and Wallenius *et al.* (2008) for a review of application areas). The fundamental input to the AHP is the decision maker's answers to a series of questions of the general form, 'How important is criterion *A* relative to criterion *B*?'. These are termed pair-wise comparisons. Questions of this type may be used to establish, within AHP, both weights for criteria and performance scores for options on the different criteria. It is this former aspect of AHP which was considered of utility in this study, since an objective of the research was to better understand the level of importance placed by experts on the various criteria they used in ontology evaluation exercises. Pair-wise comparisons are used in AHP to derive "local" priorities (weights) of the elements in a cluster with respect to their parent. The principle of hierarchic composition allows for the multiplication of the local priorities of the elements in a cluster by the "global" priority of the parent element, thus producing global priorities throughout the hierarchy (Saaty, 1994).

Saaty's (1980) basic method to identify the value of the weights depends on relatively advanced ideas in matrix algebra and calculates the weights as the elements in the Eigenvector associated with the maximum Eigenvalue of the matrix. Nominal values (ranks) of "1" through "9" are assigned to criteria in order to derive the weights, with "1" meaning that there is no difference between two compared criteria and "9" meaning one criterion is extremely more important than the other (see Table 3.3 for the textual narrative associated with each score). Pair-wise comparisons across criteria are then used in solving an Eigenvalue equation operating on a skewed matrix of comparisons. The computed solution provides an estimate of the relative importance of decision criteria.

Table 3.3 AHP Preference Ratings

How Important Is A Relative To B	**Preference Index Assigned***
Equally important	1
Moderately more important	3
Strongly more important	5
Very strongly more important	7
Overwhelmingly more important	9

* 2, 4, 6 and 8 are intermediate values that can be used to represent shades of judgement between the five basic assessments.

In summary, Saaty's (1980) skewed (n x n) matrices (R) have the form R = (r_{jk}) with the convention that "j" and "k" are the compared alternatives and $r_{kj} = 1/r_{jk}$, for all 1 <= j and k <= n. An expert performing a pair-wise comparison is being asked to provide a ratio (w_j/w_k) which is measuring the relative dominance of alternative "j" over alternative "k" where the priority weights (w) applied are $w_1 > 0,..w_n > 0$; which by convention sum to "1" (Genest and Zhang, 1996).

The method is based on three relatively simple axioms (Saaty, 1994). First, the *reciprocal axiom*, which requires that, if P_C(A, B) is a paired comparison of elements A and B with respect to their parent element C (i.e., an element above A and B in a hierarchical decomposition of a problem), representing how many times more the element A possesses a property than does element B, then P_C(B, A) = 1/ P_C(A, B) [i.e., $r_{kj} = 1/r_{jk}$ in the notation above] . For example, if A is 5 times larger than B, then B is one fifth as large as A.

Second, the *homogeneity axiom*, which states that the elements being compared should not differ by too much in the property being compared. If this is not the case, large errors in judgment could occur. To mitigate against large departures in judgements, when constructing a hierarchy of objectives, one should attempt to arrange elements in clusters so that they do not differ by more than an order of magnitude in any cluster.

Third, the *synthesis axiom*, which states that judgments about, or the priorities of, the elements in a hierarchy do not depend on lower level elements. This axiom is required for the principle of *hierarchic composition* to apply.

The criteria to be evaluated in the AHP method are arranged in a value tree (or hierarchy) according to the principle of decomposition. Decomposition in this case is the method of classification used by the human mind in ordering experience, observations, entities and information (Whyte, 1969). The model which resulted from research activities described earlier in section 3.2.2 was therefore represented as a 3-tiered hierarchy in readiness for an analysis using AHP (similar schematically to that shown in Figure 3.7).

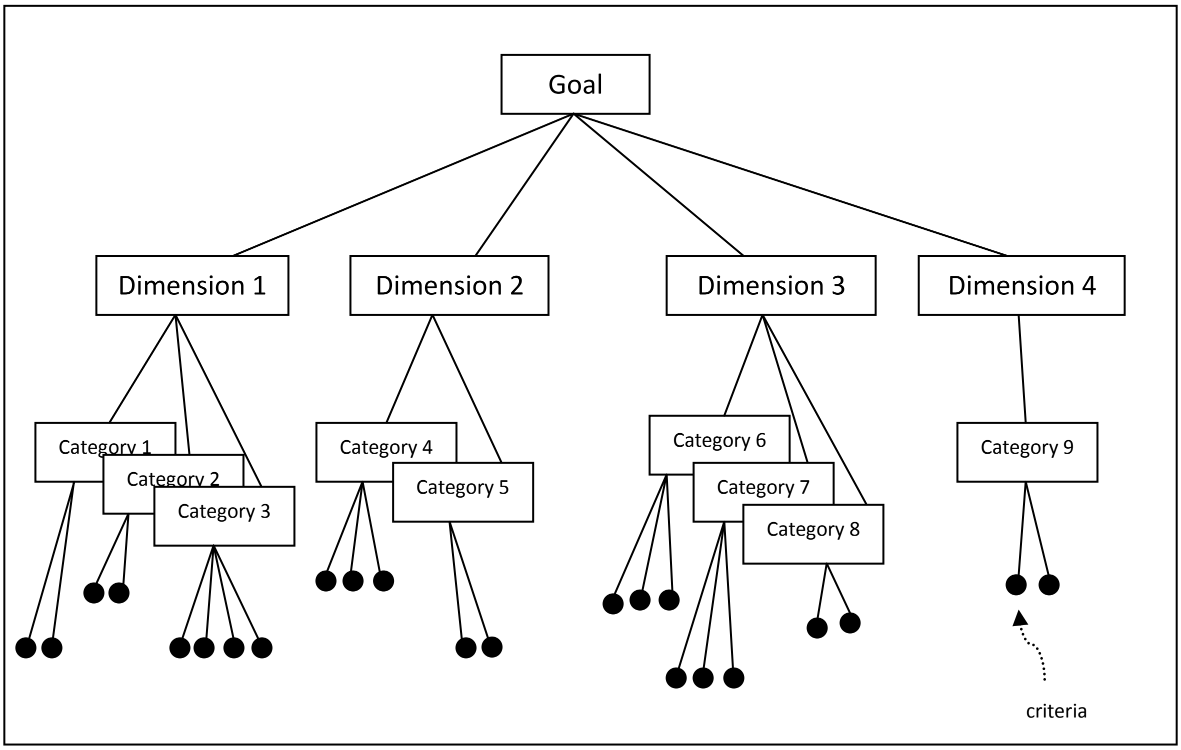

Figure 3.7 Schematic of an AHP conceptual model for multi-criteria decision-making

The AHP technique is capable of measuring the "consistency" of judgements made by an individual expert during the pair-wise comparisons, through the inclusion of internal redundancy checking and evaluation against "consistency ratios". Inconsistency, as previously discussed, usually arises when the logic applied to a set of ratings, which should be transitive, does not result in a transitive solution. For example if an expert has three options "A", "B" and "C" and they state that they prefer "B" to "A" and they also prefer "A" to "C", it follows logically that they should also prefer "B" to "C".

This is deduced as a transitive property. If they state instead that they prefer “C” over “B” there is an inconsistency in the logic applied. This is sometimes also called an ordinal inconsistency. The other type of inconsistency that routinely occurs in ranking multiple criteria is cardinal inconsistency. Consider the case where there are three distinct alternatives j, k and l, where 1 <= j, k, l <= n; if $r_{jl} = r_{jk} \times r_{kl}$ does not hold to be true, then the matrix “R” is said to be inconsistent (Li and Ma, 2006).

Saaty (1980) used the Principal Eigenvalue (which can be obtained by summation of the products between each element of an Eigenvector and the sum of columns of the reciprocal matrix of R) to compute what he termed a consistency index (CI). By proving that for a consistent reciprocal matrix, the largest Eigenvalue is equal to the size of the comparison matrix (R), or ($\lambda_{max} = n$), he computed “CI” as the deviation of the largest Eigenvalue from the size of the comparison matrix. His formula: $CI = (\lambda_{max} – n)/n – 1$ calculates the consistency index. This index is then compared with a randomly generated index “RI” (derived through boot-strapping). Using the bootstrapping technique Saaty (1980) generated a table of “RI” values for matrices of various sizes. He then proposed what he termed a Consistency Ratio (CR), which is computed as the ratio of CI/RI. Somewhat controversially, Saaty (1980) then argued a consistency ratio (CR) smaller or equal to 10% (CR<=0.1) was the acceptable limit before expert judgements need to be revised. Various authors since Saaty (1980) published his work have commented on the utility of this CR value, with some (e.g., Kauko, 2002; Kryvobokov, 2005) believing this value to be too low.

Despite the apparent utility and popularity of AHP there are a range of concerns raised in the literature about some of its characteristics which are the subject of ongoing debate. French (1988); Goodwin and Wright (1998) and Warren (2004) summarise the primary issues thus:

- The 1–9 scale has the potential to be internally inconsistent. “A” may be scored 3 in relation to “B” and “B” similarly scored 5 relative to “C”. But the 1–9 scale means that a consistent ranking of “A” relative to “C” (requiring a score of 15) is impossible. This is due to transitivity issues.
- The link between the points on the 1–9 scale and the corresponding verbal descriptions does not have a theoretical foundation.
- Weights are elicited for criteria before measurement scales for criteria have been set. Thus the decision maker is induced to make statements about the relative importance of items without knowing what, in fact, is being compared.
- Introducing new options can change the relative ranking of some of the original options. This ‘rank reversal’ phenomenon, first reported by Belton and Gear (1983) is alarming and arises

from a failure to consistently relate scales of (performance) measurement to their associated weights.

- Although it is a matter of debate among decision analysts, there is a strong view that the underlying axioms on which AHP is based are not sufficiently clear as to be empirically testable.
- In the AHP literature there is considerable ambiguity as to whether the input relative importance ratings are on an *implicit* ratio scale, or whether the derived priorities computed from the comparison matrix are on a *derived* ratio scale. Both the existence and the location of an absolute zero comprise the necessary and sufficient conditions for the presence of a ratio scale (and it is claimed this is not the case in AHP's preference scale).

Saaty (1990); Forman and Gass (2001) and Gass (2005) have responded to most of these concerns defensively. For example, Forman and Gass (2001) argue that intransitive situations are a fact of life and it is therefore appropriate for AHP to recognise this and deal with it. AHP 'deals with' this particular issue through pair-wise consistency analysis. Forman and Gass (2001) further argue that AHP does not claim to hold to the Multi-Attribute Utility theory (MAUT) and Multi-Attribute Value Theory (MAVT)(Keeney and Raiffa 1976, Edwards 1977, Edwards and Barron 1994), in which transitivity is an axiom and assert that MAUT includes the axiom of transitive relations simply because transitivity decisions are far more mathematically tractable than intransitive ones (which MAUT practitioners cling to in a type of denial of reality).

Gass (2005) and Saaty (1987, 1990) likewise argue that 'rank reversal' in decisions happens in the real world. With respect to decision procedures, rank reversal can or cannot occur, that is, it can be ruled in or out, based on the axiomatic base of the decision procedure. In general, based on a decision model's procedural decision rules and associated axioms, the decision maker can choose to rule rank reversal in or out in terms of how they establish the model. They grant that this may result in different outcomes, but they argue that the decision to follow one route or the other can be overtly made. Forman and Gass (2001) describe extensions to AHP that let it operate in either mode.

Forman (1990) and Forman and Gass (2001) also refute that there are problems with (a) the ratio scale as used in AHP or (b) the verbal scale for preferences. They explain that "whereas an interval scale is defined to be a scale that is invariant under the transformation $y = ax + b$, a ratio scale is defined to be invariant under the transformation $y = ax$. Because there is no "b" in the ratio-scale transformation, the ratio scale is said to have a "true" zero. Some have questioned whether AHP produces a ratio scale because they do not see any zero in either the fundamental verbal judgment scale used for pair-wise comparisons or the resulting priorities. This misunderstanding is partly due to the misconception that "fuzzy" verbal judgments are the only way to express relative judgments".

In AHP the supposition is that humans are more capable of making relative judgments than absolute judgments. The use of an ordinal scale of verbal measurements, is defended on the basis that, Saaty's (1980) empirical research showed that "the principle eigenvector of a pair-wise verbal judgment matrix often does produce priorities that approximate the true priorities from ratio scales such as distance, area, and brightness. This happens because, as Saaty (1980) has shown mathematically, the eigenvector calculation has an averaging effect – it corresponds to finding the dominance of each alternative along all walks of length k, as k goes to infinity. Therefore, if there is enough variety and redundancy, errors in judgments, such as those introduced by using an ordinal verbal scale, can be reduced greatly". Given that the preference scale has been used extensively and is intuitively well understood by decision-makers, Forman and Gass (2001) argue that the proof of the scale's utility is in its applied execution.

Despite these concerns, the AHP method was chosen as a tool to investigate expert preferences for the evaluation criteria used in ontology selection in this study for a number of reasons:

- It is arguably internally consistent and logically sound (despite current controversy surrounding the issues mentioned previously) and it provides a quantitative mechanism for obtaining weights as well as option preference. Obtaining weights was the primary goal of this phase of the research. This contrasted with methods based on MAUT which primarily obtain weights by qualitatively surveying stakeholders.
- It is transparent in how preferences are calculated and is based on matrix algebra.
- The method is relatively straightforward and simple to apply (albeit somewhat laborious if a large number of criteria need to be rated) for evaluators/decision-makers. This 'ease-of-use' aspect was particularly important given the expert group involved in the study was highly dispersed and could not be brought together for the data capture or data analyses phases. A technique was therefore required which could be easily understood and participated in by experts who were remote from the author's location.
- The method appeared to be easy to apply even with a relatively low level of decision-support analytical skills (on the part of the author).
- The method provided for an auditable evaluation process.
- The algorithms used for analyses were available as open source code that could be run independently or plugged into a Microsoft™ Excel software package.

3.3.4 Ontology Selection and Evaluation Framework

Having explained the rational for selecting AHP as a key tool for determining ontology evaluation criteria preferences amongst experts, it remains to explain how the grounded ontology selection and evaluation framework was derived. This was achieved by combining the results and theoretical conclusions drawn from the research described so far in sections 3.2.1 to 3.2.3.

A framework (Farlex, 2012b) is a scaffold or a set of assumptions, concepts, values and practices that constitute a way of viewing reality. In this strand of research the aim was to establish guidance for practitioners. To achieve this guidance, both qualitative and quantitative methods were used to elicit different facets of the framework. The qualitative interviews (section 3.3.2) provided not only material that was used in constructing the hierarchical ontology evaluation model (section 3.3.3), but also material that exposed community governance practises germane to ontology management and re-use. These interviews also shed light on the current usage of ontology evaluation methodologies and the metrics used in evaluating criteria. This material, triangulated with documentation provided by the expert practitioners, along with a review of current theory, lead to the postulation of a grounded ontology selection and evaluation framework.

Despite the small expert group sample size involved in the study, it is still considered that this assembled group had sufficient depth to be representative of the phenomena and issues under investigation.

3.4 Summary

In this chapter the epistemological and ontological underpinnings of this thesis have been discussed and the overall research design has been presented. It has been argued that given the mixed (i.e., multi paradigm) model used, each component of the research should be assessed for validity according to the tenets of the particular paradigm in which it is anchored. The rationale for the selection of particular methods has also been provided and justified.

The individual methods employed to address the research questions posed in Chapter 1 (Figure 1.3) have been articulated in detail and, where appropriate, the evaluation criteria applied by the author in conducting particular methods have been outlined to assist readers in establishing the research validity and rigour.

The thesis now moves on to present the results of the application of the various types of methods that were discussed in this chapter.

Chapter 4.

Feature Catalogue Design and Development Results

This chapter is concerned with presenting the research and results related to the design and development of a semantically-enabled Feature Catalogue. It directly addresses RQ1.1 – "*What characterises an ontologically-grounded Feature Catalogue that can support Antarctic science data publication through Web services ?*". The research was grounded by continuous interaction with both the SCAR and AODN scientific communities-of-interest to ensure that the research outputs (and outcomes) would be of "practical" value. This interaction spanned several years (2007-2010) during which time requirements were gathered, ideas were seeded, artefacts were developed and community feedback was received.

Four factors were considered crucial to the development of the semantic Catalogue:

1. the Catalogue had to fit architecturally as a component within the evolving infrastructures of both Antarctic-connected communities and incorporate (or interoperate with) any standards already adopted;
2. any developed access interfaces to the Catalogue had to demonstrably support dataset annotation methods capable of being executed within current Antarctic-themed scientific data infrastructure;
3. the Catalogue content model had to accommodate the types of data routinely captured and exchanged; and
4. the content model also had to provide an appropriate semantic signature for included (Feature Type) concepts.

This last factor was considered important for using the Catalogue as a domain ontology repository in data service annotation scenarios. The semantic signature of a served concept needed to include the right type of information sufficient to enable the differentiation between dissimilar Feature Types, particularly for use-cases involving automated data integration exercises.

This chapter unpacks by explaining all of the research that was associated with the Feature Type Catalogue development. It commences (in section 4.1) by characterising what is meant by some of the key terms used in the chapter. The chapter then proceeds (in section 4.2) with investigative activities aimed at acquiring a deeper understanding of the data and data models which are applicable for encoding community datasets. This section then concludes by outlining the elicited user-based requirements that a Feature Catalogue artefact should meet. The chapter then progresses (in section 4.3) with a review of existing Feature Catalogue related standards, noting any standards-based deficiencies. The discussion then moves on (in section 4.4) to define a Feature Catalogue conceptual model that better meets user requirements (than does ISO 19110) but which still conforms to the basic design of the ISO 19110 Feature Cataloguing Standard. This enhanced, ISO-standards-based Feature Catalogue model is first cast as an ontology by anchoring to the upper-level ontology DOLCE (in section 4.5) and is later concretely realised through three different serialisations (i.e., in OWL, SKOS (RDF) and RDBMS) in sections 4.6 and 4.7. The final section (4.7) of this chapter presents a prototype Feature Catalogue service interface that can be used to access sample Feature Types (and associated contextual concepts) from the RDBMS catalogue serialisations. The service architecture used (i.e., REST) ensures that Catalogue content can be accessed from any type of Web service or Web client capable of resolving a URL. By providing XML and SKOS (RDF)-based output formats for accessed resources, the Catalogue can serve as a source of linked data. Linked data, as described by Tim Berners-Lee (2006) is a method of exposing and connecting to data on the Web that is derived from different sources (using URLs as names for resources; RDF, or XML for serialising the resource information; and by providing resources with embedded hyperlinks).

To assist the reader to interpret the data models which are presented in this chapter some basic Unified Modelling Language (UML) notation is summarised in Figure 4.1

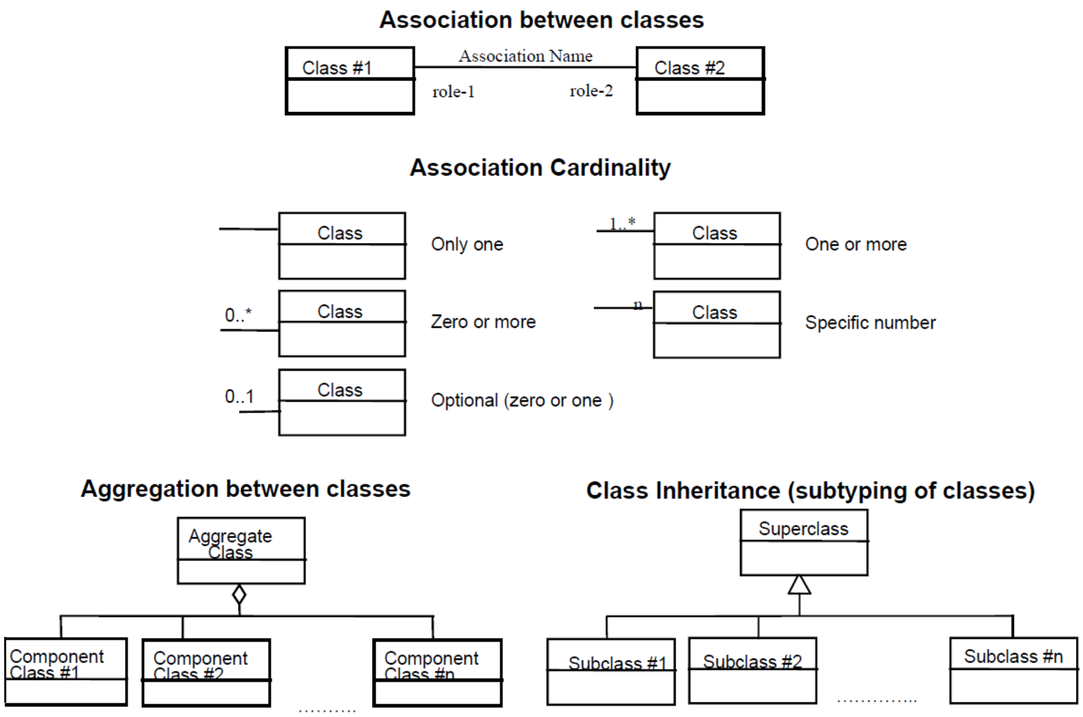

Figure 4.1 Basic UML Notation For Interpreting UML Diagrams

4.1 'Datasets', 'Features', 'Observations' and 'Features Of Interest'

It is important, for the research that follows, that we commence with some broad definitions for the central entities that will be modelled in the course of this chapter. Primarily, we are concerned with what is meant by the terms: "data or dataset", "feature", "Feature Type" , "observation" and "feature-of-interest" since understanding what these terms refer to is critical for designing an appropriate Catalogue content model and accompanying Catalogue application software.

Generally, AODN and SCAR community members have used these terms vaguely, interchangeably and sometimes erroneously when discussing data exchange issues (e.g., "*I want to publish the ship **observation** data through the portal*", "*I just uploaded the **underway dataset***", "*I plotted the **sea-surface temperature features** through the portal*" or "*I used CTD data to publish features-of-interest*"– as quoted from various AODC JF Technical Committee meetings). A lack of community coherence surrounding the use of the terms ("dataset", "feature", "observation" and "feature-of-interest) has probably helped perpetuate the current situation of divergent approaches to data publication (which results in poor data and systems interoperability).

It is instructive to examine how these terms have been defined and applied in a number of standards that are related to data exchange in order that the scope and intent of the terminology might be more clearly discerned.

The International Standards Organisation (ISO), which in unison with the OGC drives most of the standards development for spatial data exchange, defines the terms "dataset" and "feature" (ISO/TC 211, N630), rather ambiguously (and in a somewhat terminologically self-referencing manner) as follows:

- *Dataset*: "an *identifiable collection of data*".

  "Note: A dataset may be a smaller grouping of data which, though limited by some constraint such as spatial extent or feature type, is located physically within a larger dataset. Theoretically, a dataset may be as small as a single feature or feature attribute contained within a larger dataset."

- *Feature: "an abstraction of real-world phenomena"*

  "Note: A feature may occur as a type or an instance. Feature type or feature instance should be used when only one is meant."

These relatively limited definitions from ISO have possibly added to the terminological confusion that exists within the studied communities and it is unclear how the various standards, based heavily around some of these terms (particularly ISO 19109 – *Rules For Application Schema*, ISO 19110

Methodology For Feature Cataloguing and ISO 19126 Feature Concept Dictionaries and Registers (ISO (2009a)), should be practically applied as a result.

4.1.1 Datasets

The ISO 'dataset' definition takes a 'collection'-centric view in that a 'dataset' is seen as a 'collection of data'. Renear *et al.* (2010) have reviewed various definitions of 'dataset' found in technical documentation and the scientific literature and discovered that four basic concepts can be identified as common to most definitions: 'grouping'; 'content'; 'relatedness'; and 'purpose'.

Definitions which emphasise 'grouping' use terms like 'set', 'aggregation', 'container', and 'collection' to indicate that datasets are 'data' treated collectively as a unit (as in the ISO example).

Definitions which emphasise 'content' generally use qualified terms beyond simply mentioning 'data' as a constituent and imply that the constituents of a dataset are things of some particular kind. The 'data' in datasets are variously described with terms such as "observations" (Feeley *et al.,* 2004; Purchase *et al.,* 2008), "facts" (McDermott *et al.,* 2001), "values" and "records of values" (Purchase *et al.,* 2008). Typically the content of a dataset is intended to reflect the results of certain sorts of activities, such as measuring or observing. In particular what is recorded are observations (Feeley *et al.,* 2004) or "the results of" observations (Purchase *et al*., 2008). Of obvious importance is the considerable variation in the level of abstraction at which dataset 'contents' are conceived. In some places these contents appear to be abstract conceptual entities (e.g., observations, property values), and in other places particular representations of those entities (records of values, XML elements), or even lower level entities such files (Renear *et al.* 2010). Within the Antarctic communities studied in this research, this 'content' flavour of 'dataset' definition is probably the most often encountered.

In dataset definitions which are framed with respect to 'relatedness' it is evident from these definitions that datasets are thought of as grouping together constituents (data) that are related to each other in some way that goes beyond both the grouping itself, and the identification of the grouped things as all being of the same general kind of entity. In 'circumstantial relatedness' for example, a 'dataset' is sometimes thought of as consisting of 'data' related by time, place, instrument, or object of observation. Again, this is a commonly occurring frame of reference for scientists generating Antarctic-themed data. These facets draw attention to the circumstances around the creation or maintenance of a dataset, as opposed to any internal characteristics of the data, which would then be more commensurate with a type of 'syntactic relatedness'. In syntactic relatedness 'data' in a 'dataset' are typically expected to have the same syntactic structure (records of the same length, field values in the same places, etc). The last type of relatedness generally encountered is 'semantic relatedness' where 'data' in a 'dataset' may be about the same subject.

In those 'dataset' definitions that lean towards 'purpose', the characterisation implies that the datasets are clearly created in order to contribute in some way to scientific activity. This might be by providing evidence to be analysed, suggesting new hypotheses, providing refutation or confirmation of existing hypotheses, or supplying new phenomena to be explained (Renear *et al.,* 2010). An example of a purpose-based definition is one given by the digital preservation community for 'data', where 'data' is described as: "a reinterpretable representation of information in a formalized manner suitable for communication, interpretation, or processing. Examples of data include a sequence of bits, a table of numbers, the characters on a page, the recording of sounds made by a person speaking, or a moon rock specimen" (CCSDS, 2002).

In the Antarctic community all of these 'dataset' definition types have been encountered. It is clear that there is no single well-defined concept of dataset. The Catalogue content model must therefore be able to accommodate 'data' in all of the senses just described. Leveraging the work of Renear *et al*. (2010), however, the following "all encompassing" definition for 'dataset' has been coined for use in this thesis – "a dataset is a grouping of observations, measurements, values or facts that are related in some manner that is meaningful to either the dataset producer or user and which can be used for informational value".

4.1.2 Features

The ISO definition listed earlier, for a 'feature', gives the broad intent of the term but is not particularly expansive. Since the concrete data encodings of both communities will definitely be couched in terms of GML (Portele, 2007), as specified by the OGC, it is instructive to understand how GML applies the term. A Catalogue incapable of serving 'features' that accord with GML requirements would not be particularly useful to communities using OGC standards. A "concept" in GML is generally referred to as a "feature", which like a concept in traditional ontology description languages such as RDFS or OWL, can describe an abstraction of a real-world phenomena. While features are the focal elements of most GML documents, the language is actually comprised of "objects" which include features, geometries, coordinate reference systems and styles. In GML, a **feature instance** is an identifiable object in the world (or the digital **representation** of it) and feature instances are classified into **Feature Types** on the basis of common sets of characteristics, or properties (e.g., their attributes, associations and relationships, operations and/or behaviours). These definitions have been established earlier but it is useful to recap at this point.

A GML "application schema" is said to describe the "logical structure and semantic content of a dataset using a feature-based model". A detailed expose of the Generalised Feature Model perspective, as described by ISO 19101, has been reproduced in Figure 4.14 later in this chapter. This feature-based model has also been named the 'object-property model', similar to the RDF subject-property-value model – (Minola and Miller, 2004). Essentially all the model posits is that there are things called Feature Types, which are carriers of characteristics (or properties). These properties can be of type 'Attribute', 'Operation', or 'Association'. It also states that Feature Types can be 'generalised' or 'specialised' and/or 'constrained' in some undefined (i.e., by the model) manner.

Most GML application schema are currently augmented by the use of "dictionaries", i.e., external instance documents that define things like 'coordinate reference systems', 'units of measure', 'value types' and 'temporal reference systems'. In GML, however, it is left up to a community to decide which features are of importance to their domain; at what level of detail or abstraction they are described; what relationships these features have to each other; how they should be encoded, managed and accessed; what dictionaries should be built or subscribed to and how these types of entity should interoperate with Feature Type definitions as couched in terms of the ISO (19110) Feature Catalogue Model. So, whilst there is some general conceptual and syntactic guidance (within GML), much is still left undetermined and it would appear is at a community's discretion.

Kottmann and Reed (2009), in an OGC discussion paper, have recently re-iterated and elaborated upon the meaning of the term "feature" from a geospatial perspective. In this perspective the focus is on depicting things in the real-world as points, lines, or polygons plus their attributes (i.e., information about those things). They assert that that when linked together, this pairing of geometry and attributes representing one or more real-world objects, is called a "feature". Note that in this definition a feature's geometry is intrinsic to its semantic definition (i.e., the geometry has primacy and the semantics are couched as an adjunct to, or as a description of this geometrically represented entity). Since it is entirely plausible for a real-world feature to have no geometry (take for example a socially constructed feature such as a "project"), this definition appears to have narrowed the scope of features to only include those things that potentially have some type of geometry, which might be limiting.

Much of the Kottmann and Reed (2009) paper, from which the "feature" definition above is drawn, provides a framework for describing an abstract model for features. The conceptual viewpoint assumed for most of their paper is the case where a feature is modelled primarily from the vantage of its spatial extent, using a well known geometric primitive, (which should be a familiar approach to most GIS practitioners). Whilst the communities who have participated in this study often use data congruent with this viewpoint (e.g., data that includes features defining: geographic boundaries, fishing zone limits, conservation zones, coastlines, Antarctic stations inclusive of buildings and utilities, various types of "map" data etc) this is not the primary type of data that is routinely transacted. Instead, the types of datasets that this study's communities frequently exchange are those perhaps best described from an "observation" perspective, albeit cognisant of the importance of the shape and temporal dimensionality of an observation's sampling regime. In such datasets the act of 'observing' is the central unifying factor and the focus of the observation, its sampled parameters and its sampling regime are the 'things' of interest (not just the immediate feature that is the focus of the observation and its spatio-temporal geometry). Despite this apparent difference in perspective, many of the assertions made by Kottmann and Reed (2009) about features, are, however, still relevant and there are two assertions in particular which should be highlighted.

They recommend that:

> (i) "*features with spatial extent ... be conceptualized using simple primitive geometric shapes.. and that for interoperability (purposes)...the primitive shapes (should) be instances of the Well-Known-Types (WKTs) of geometry, such as polygons, line strings, polyhedrons, and other shapes...(drawn from ISO 19107)*".
>
> (ii) "*each feature type can be thought of as a 'template' to be populated at each occasion of a feature instance corresponding to that type*"; and

The conclusion drawn from (i) is that a Feature Type should explicitly include the geometry WKTs that are to be used in its semantic representation. From (ii) we can conclude that since a 'template' (Farlex, 2012c) is a "*document or file having a preset format so that the format does not have to be recreated each time it is used... and it is also used as a guide in making something accurately"*, that a Feature Type 'template' should be as fully descriptive as is possible for distinguishing the characteristics of a feature included in any encoded dataset.

To the author it seems entirely reasonable to include possible geometries in a Feature Type's definition, although whether the "geometry" should necessarily be drawn from "primitive types", or whether alternatively it could be drawn from more complex "Well-Known -Types" that are endorsed and documented 'a priori' by a community of practise (perhaps including both spatial and temporal dimensions) is arguable.

The idea of a Feature Type acting as a "template" and the inclusion of its "geometry" in its semantic description are important principles that will be carried forward in this study, regardless of whether a geospatial or observation-centric perspective is being adopted.

4.1.3 Observations and Features-Of-Interest

In contrast to the Kottmann and Reed (2009) description of a Feature Type given above, an 'observation' perspective of a feature has been conceived by Cox (2006, 2010a) which has now been adopted as an OGC standard. When this research commenced, however, the Cox (2006) model was not part of the OGC standards suite and was conceived originally as a conceptual model for describing geological datasets. Its introduction to the OGC standards stable has occurred without modification to, or revision of, existing underlying standards such as ISO 19110 (Feature Cataloguing Methodology) which, when originally created, was clearly geared towards a geospatial perspective for Feature Types. This is probably one reason why the investigations reported later in section 4.3 found deficiencies in the ISO 19110 Feature Cataloguing standard with respect to its ability to adequately capture the semantic signature of Feature Types, as expressed from an observation-centric viewpoint.

Whilst Cox (2006) did not eschew the importance that geometry needed to play in shared data models, he organised embedded dataset features from a sampling perspective. In the Cox model (Figure 4.2), an "observation" (itself a type of Feature) is an "event" whose "result" is an estimate of the value of some "property" of a "feature-of-interest" obtained using a specified "procedure". Therefore, in a simplified observation event (which results in the generation of a dataset) there are two types of feature that are important: the observation Feature Type and the observation's feature-of-interest. The latter is the representation of the real-world entity, which is the subject of observation. The observed properties (or 'Phenomenon' in the Cox model, Figure 4.2) in the dataset are directly related to the feature-of-interest through the act of observation.

In the main, the observation model takes a data-user-centric viewpoint, emphasizing the semantics of the feature-of-interest and its properties (Cox, 2010a), rather than a geometric-centric perspective (favoured by Kottmann and Reed (2009)). However, the Cox (2010a) model does provide for a range of utility sampling Feature Types that can be used in place of his model's "feature-of-interest" entity and these Sampling Features are primarily differentiated on the basis of geometric and topologic structure and not on the semantics of the observable, or the measured property. These pre-configured Cox (2010a) Sampling Features have been presented in Figures 4.3 and 4.4. Each type of Sampling Feature can be substituted individually into the place of the 'feature-of-interest' entity in Figure 4.2. All pre-configured 'Sampling Features' (in Figure 4.4.) are associated (through a 'shape' property) with Well Known GML Geometry Types (drawn from the ISO 19107 Spatial Schema Standard). When used in the place of the observation's 'feature-of-interest' these sampling Features act as a proxy for the real-world entity that is being sampled (or observed).

Much of the data exchanged by this study's communities are observations and measurements of specific phenomena related to real-world entities (i.e., feature-of-interest). So the Cox (2006, 2010a) model (which takes an observation perspective), at least at a high level, appears to be a better fit for modelling Feature Types that occur within many Antarctic datasets than a model which primarily focuses on the geometry of objects.

From the author's twenty years of experience as a professional scientific data manager, it has often been demonstrated that it is usually impossible to ignore how data have been collected (sampled), in terms of formulating sound (data exchange) data models because the types of data acquired and how these data have been collected, in part, determines how data can later be discovered, accessed and used within a data exchange network. Sampling information is also a key element in determining a dataset's fitness for use. The Cox (2006, 2010a) Observation and Measurement (O&M) model appears to provide a basic framework that can potentially capture much of the information that may be required to optimise data exploitation and re-use from a sampling, as well as a spatio-temporal perspective.

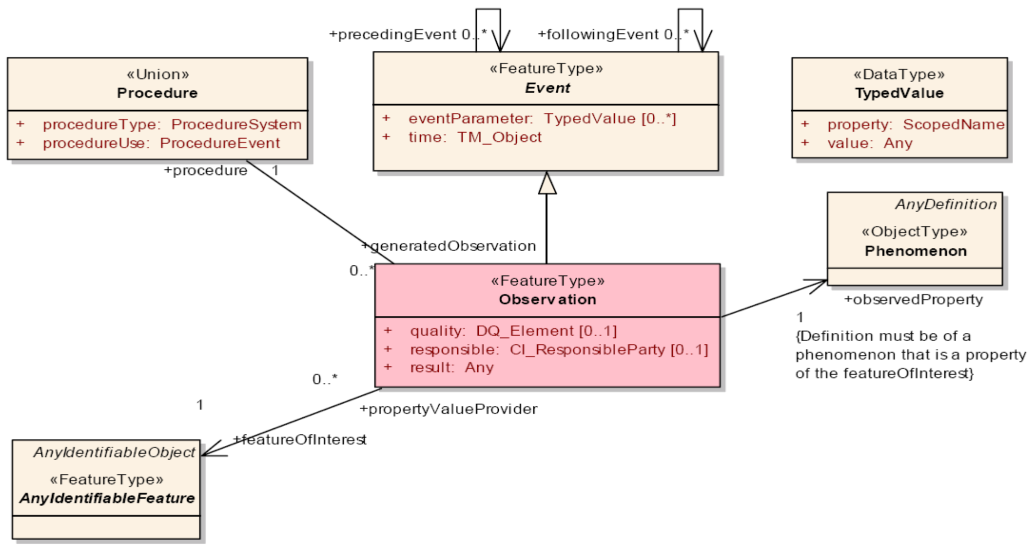

Figure 4.2 Observation and Measurement (O&M) Schema (Cox, 2006)

In summary, a 'dataset' is a very loose term that can be used to describe a very wide range of information entities that are grouped and related in a manner that makes sense from the perspective of the purposes to which the dataset is to be put. A 'feature instance' identifies a specific individual of a 'Feature Type'. A 'Feature Type' is a user-defined classification of 'feature instances' and both a feature instance and a Feature Type are 'representations' of real-world entities. A 'Feature Type' should provide a complete semantic template for the description of a 'feature instance', inclusive of its spatio-temporal attributes and an observation-centric data modelling paradigm appears intuitively to suit the majority of datasets that are generated and exchanged within the Antarctic community. A Feature Type Catalogue will therefore need to accommodate both geometric and observation-centric Feature Type semantics.

The next section moves from an abstract examination of terms like 'dataset' and 'Feature Type' and looks in more detail at the characteristics of both, from an Antarctic community perspective. Through the development of use-cases with the communities of interest, a deeper understanding of the expectations regarding Feature Catalogue functionality and purpose also emerges.

4.2 Data, Data Models and Requirements Elicitation Via Use Cases

From earlier work (Finney, 2007) the author was already aware that the flexibility, extensibility and interoperability of the AODN and SCAR (data exchange) systems would depend upon community conformance around issues of syntax, semantics and governance of Feature Types and their associated feature instances (regardless of the language of their expression). Cognisant of this, the AODC JF (who are the technical sponsors of the AODN infrastructure) sanctioned the author (in 2007/08) to perform an evaluation of potential dataset encoding patterns in order to gain a better understanding of the nuances in community data that needed expression through Feature Types and to acquire a better appreciation for how end-users might utilise a Feature Type Catalogue. The desired output as a result of this sanctioned exploration, particularly from the AODN community's immediate perspective, was an assessment (or development) of one or more encoding patterns that could be piloted as a data exchange schema within the emerging AODN infrastructure (and which could latterly be adopted by SCAR).

The work that was initially undertaken by the author in 2007/08 (Finney, 2008) and then subsequently, to evaluate existing data encoding patterns, also contributed directly to the instantiation of detailed use-cases and requirements against which a robust Catalogue, suitable for serving the communities in question, could be assessed. This work, from the author's perspective, was specifically performed to explore RQ1.1.1 – "*What type of use-cases and data models must the Feature Catalogue support ?".*

Some of the key findings from this work are presented next, which have a direct bearing on shaping the eventual characteristics of the developed Feature Catalogue.

4.2.1 Review Of The Observation and Measurement Model

As described earlier, the observation-centric model (as presented by Cox (2006)) appeared to be a logical way of expressing AODN (observing system) datasets. So, in 2007 the author evaluated whether the O&M Schema (Cox, 2006) did in fact provide a good basis for a broad dataset encoding pattern. The rationale being that the nature of the data models supported by the study's participating communities should have a strong influence on the type of content model required for a community-based Feature Catalogue.

The core model provided by Cox (2006) was found to be a highly useful framework, but was also deficient in some areas for the community's application. Some additional Feature Types which were needed were not in the model and the choice of the custom Sampling Features available didn't cover all of the community's requirements for encoding their existing observed data types. At the time of releasing the O&M model, Cox (2006) had not yet fully defined the utility "Sampling Features" mentioned earlier. These have since been fully articulated in Cox (2010a), refer again to Figures 4.3 and 4.4. He had, however, in 2006 indicated the potential use of "Coverages" as O&M "result types". Recall that an observation "result" is an estimate of the 'value' of some 'property' of a 'feature-of-interest' obtained using a specified 'procedure'.

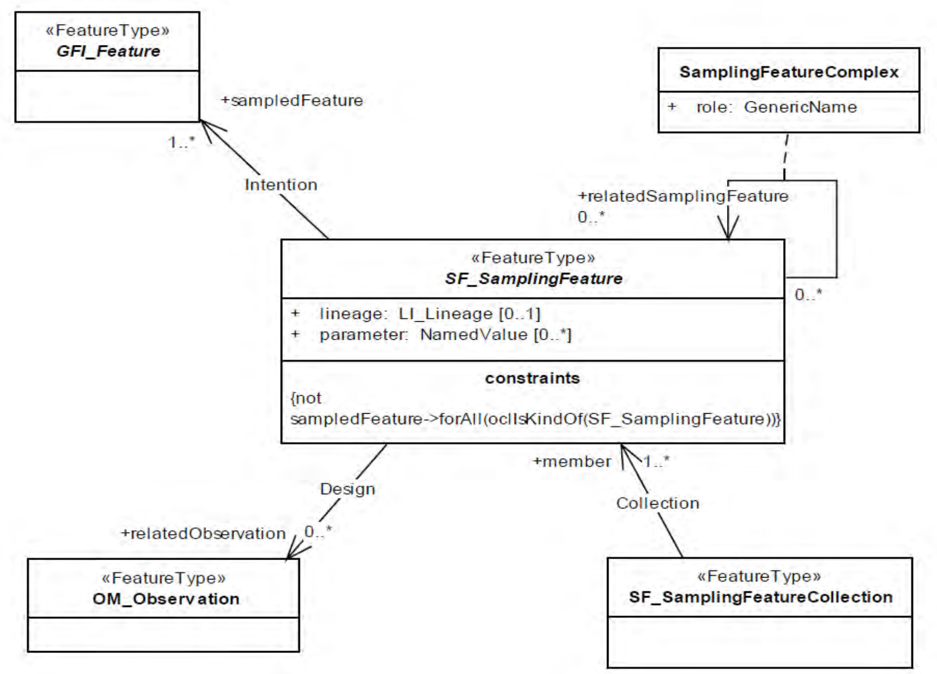

Figure 4.3 UML Model Of O&M Sampling Feature Construct (Cox, 2010a)

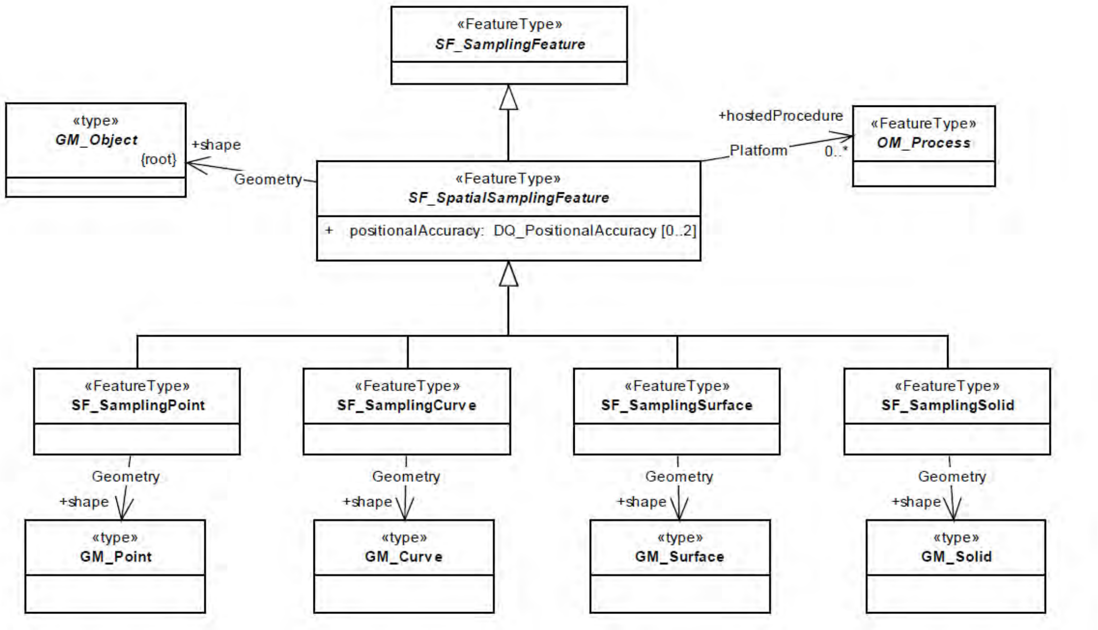

Figure 4.4 UML Model of O&M Specialised Spatial Sampling Features (Cox, 2010a)

Utility Sampling Features and Coverages are attractive modelling constructs because the former can be used to establish some well-known sampling patterns for Feature Types and the latter provide relatively efficient encoding structures for transferring observation results.

Recall (from Chapter 1) that a ‘Coverage’ is a special type of ‘feature’ that acts as a function to return values from its range for any direct position within its spatial, temporal or spatio-temporal domain. Cox (2010a) describes this more simply using tabular data (modified by the author and presented as a geologically-oriented example in Table 4.2). In Table 4.1, the highlighted row (labelled ABC-124) represents a feature-centric view of the data and the highlighted column (labelled Cu-b%) provides the Coverage view. The feature-centric perspective gives all properties for a particular sampled specimen and the Coverage view shows the variation of a single property over a range of specimens.

Table 4.1 Feature vs Coverage Perspective (adapted from Figure D.1, Cox (2010a))

Specimen	Au (ppm)	Cu-a%	Cu-b%	As (ppm)	Sb (ppm)
ABC-123	1.23	3.45	4.23	0.50	0.34
ABC-124	1.22	3.44	3.56	0.45	0.30
ABC-125	1.11	3.56	4.11	0.46	0.31

Given the O&M model's apparent deficiencies for modelling the various types of features-of-interest (in 2007) encountered by the communities in question, an embellished model was proposed by merging some aspects of another data model, called the Climate Science Modelling Language (CSML), developed by Woolf (2007) in support of the UK National Data Grid project (BODC, 2011), with the Cox model. CSML, which was firmly rooted in the Climate Sciences, had already defined 13 different types of feature primitives (equivalent to Cox's "Sampling Features", refer again to Figure 4.4), which covered data sampling contexts similar to those faced by the physical observation programs of the AODN (and SCAR) communities. Table 4.2 outlines these CSML generic Feature Type descriptions (which can be substituted in as 'features-of-interest' in the Cox O&M Schema, Figure 4.2). This Table (4.3) gives examples of the types of observations/measurements that they are designed to model. Note the inclusion of "time" as a defining attribute of these CSML Feature Types. In the CSML model, each CSML Feature Type has as its "value" (or a "result" in Cox's O&M terms), data encoded as a Coverage. It was therefore considered useful to explore the re-usability of these particular Feature Types (and their associated Coverages) for Antarctic data modelling purposes.

It should be noted that O&M and CSML have subsequently been harmonised and so the O&M specification of 2010 is able to express the spatial domains of the CSML generic Feature Types, but that wasn't the case at the time when initial modelling took place in this research. Regardless, as will be demonstrated next (in section 4.2.2), the CSML Feature Types, whilst adequate to express many types of Features found in climate science datasets, are not appropriate for modelling biological Feature Types.

Table 4.2 CSML Feature Types (Source: Woolf 2007)

Feature type	Description	Example
PointFeature	Single point measurement.	Rain-gauge measurement
PointSeriesFeature	Time-series of single datum measurements at a fixed location in space.	Tide-gauge, rainfall time-series
TrajectoryFeature	Measurement along a discrete path in time and space.	surface salinity along a ship's cruise track; atmospheric aerosols along an aircraft's flight path
PointCollectionFeature	Collection of distributed single datum measurements at a particular time	2m temperatures measured at weather stations across the UK at 0600Hz.
ProfileFeature	Single 'profile' of some parameter along a vertical line in space.	wind sounding, XBT, CTD, radiosonde
ProfileSeriesFeature	Time-series of profiles on fixed vertical levels at a fixed	vertical radar timeseries, thermistor chain timeseries

	location	
RaggedProfileSeriesFeature	Time-series of unequal-length profiles, but on fixed vertical levels, at a fixed location	repeat daily balloon soundings of atmospheric temperature from the same location
SectionFeature	Series of profiles from positions along a trajectory in time and space.	shipborne ADCP
RaggedSectionFeature	Series of profiles of unequal length along a trajectory in time and space	marine CTD measurements along a ship's cruise track
ScanningRadarFeature	Backscatter profiles along a look direction at fixed elevation but rotating in azimuth	weather radar
GridFeature	Single time-snapshot of a gridded field.	gridded analysis field
GridSeriesFeature	Time-series of gridded parameter fields	numerical weather prediction model, ocean general circulation model
SwathFeature	Two-dimensional grid of data along a satellite ground-path	AVHRR satellite imagery

The modelling results of the 2007 evaluation exercise can be seen in the UML diagram in Figure 4.5. This derived model sought to take the best elements of CSML and O&M and incorporate additional features, thought necessary for fully describing the observation program data of the Antarctic communities in question.

In Figure 4.5 the additional or modified entities that were created to augment the basic O&M model are coloured purple. In this model all observations were considered to be owned by a "Project" and conceptually can be thought of as being associated with a type of specialised O&M model "Event" entity called a "Deployment". A "Deployment" is a specialised observation in which "equipment" or "instruments" are launched, or set into action. A "Deployment" feature might be associated with (or hosted by) a "Deployment Platform" (e.g., vessel, buoy, satellite, float). Often there might be multiple deployments during the operation period of a "Deployment Platform". A "Deployment" feature might link several different types of measurements, made by different instruments.

In this augmented model all observations should be associated with an "InstrumentProcedure", a "ProcessingMethod" or various packaged combinations of both. These latter objects were specialised from the original O&M (Cox, 2006) entity - "ProcedureType". Specialisation was necessary because some Antarctic datasets can be composed of observations generated from real, or conceptual compound instruments which can produce measures for multiple phenomena. These individual phenomenon measurements are often processed, post collection, using different data processing techniques (and it was important to capture such information). Also some instruments can generate several types of measures, which can either be processed together or separately. It seemed useful therefore to be able to differentiate what processing was performed on a certain parameter that was sampled using a specific instrument.

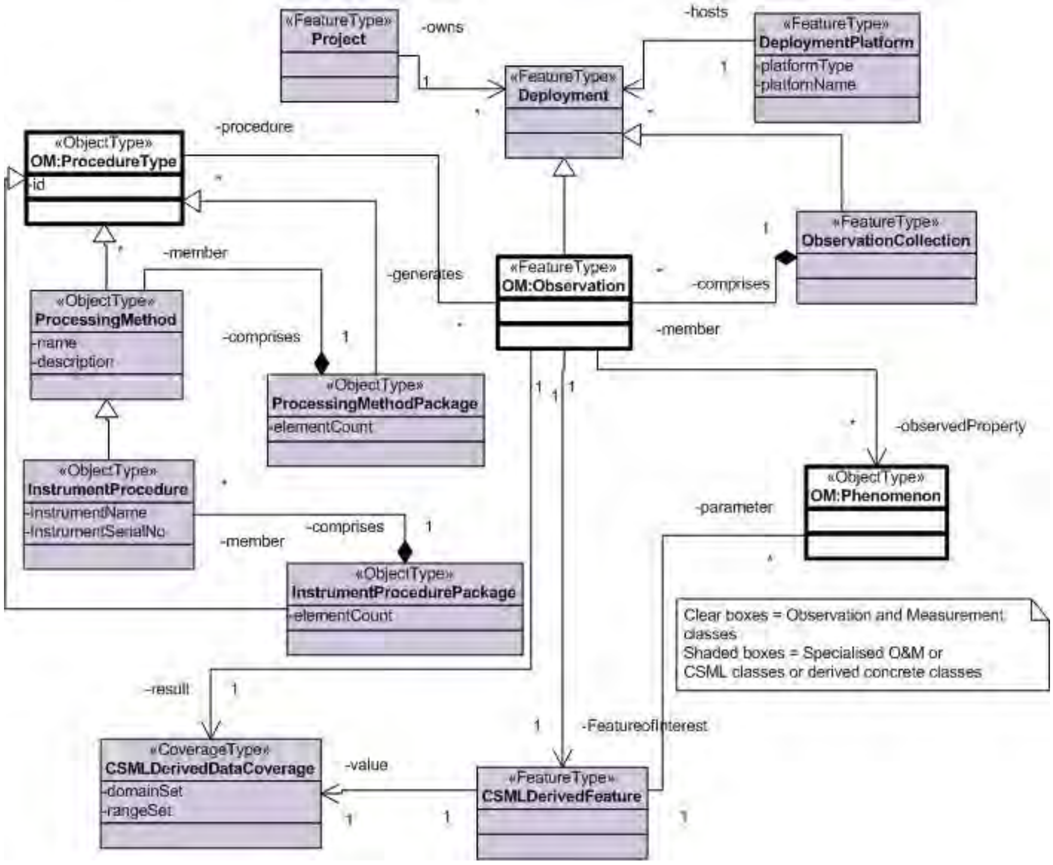

Figure 4.5 UML Diagram For The Augmented O&M Schema (From Finney 2008 – Fig 3: CSML/O&M Merged Model)

In summary, this data modelling exercise using O&M and CSML brought to light three issues that are significant for designing a Feature Catalogue:

- Some Feature Types in the model may not be directly associated with spatial footprint (e.g., Project, Deployment Platform) thus the demonstrated need to cater for both spatial and non-spatial features,
- The importance of being able to qualify the 'feature-of-interest' (and its associated parameters) by fully describing the sampling procedure,
- The CSML Feature Types did not adequately cover biological type observation raising the question of whether there were intrinsic characteristics of biological data, sufficiently

different to physical type observations that would need to be accommodated in a Feature Catalogue.

This last issue warrants justification by way of explanation and is described next in section 4.2.2.

4.2.2 Evaluation Of CSML Feature Types For Modelling Biological Data

As explained, an outcome of the 2007 modelling activity was a realisation that biological data, typically exchanged by AODN and SCAR programs and agencies, were not handled well by the attempt to use CSML Feature Type patterns embedded in an O&M framework. Biological data had to be 'shoe-horned' into relatively inappropriate data structures. This was not entirely unexpected since CSML (and to some extent the more generic O&M) had strong origins in the physical, rather than the biological sciences.

In initial investigations surrounding the merged O&M and CSML models, two types of representative observations were selected for detailed modelling. These datasets were primarily used to assess how well the encoding patterns captured the data and any nuances that related to transmission in GML.

One observation type used for evaluation was that made by towing a Continuous Plankton Recorder (CPR). The CPR Tow data (sourced from the AADC archives) was deliberately chosen because it was a relatively complex dataset and it was biologically-based.

The primary purpose of a CPR instrument is to capture plankton, which are caught by towing a CPR instrument behind a vessel. This instrument contains a silk cloth wound around a small coil which unwinds at a known rate and is exposed to seawater before being wound back around another coil, which is then immersed in preservative. Planktonic biota adheres to the cloth as it is exposed to the seawater. The cloth is returned to the laboratory, cut into equivalent length strips representing a known number of nautical miles travelled and the adhered biota are identified and usually counted. While the vessel is underway other instruments record physical phenomena such as temperature, salinity and dissolved oxygen at the sea surface. A schematic of the main elements of a CPR tow are provided in Figure 4.6.

In the modelling exercise, the CPR data was encoded using the pattern shown in Figure 4.7 (following a basic O&M pattern). An "ObservationCollection" (i.e., a collection of member Observations) was constructed to be equivalent to a CPR "tow" (i.e., one continuous recording of a CPR instrument during a voyage tracing the vessel's track) and each "Observation" equated to a "tow segment" (i.e., a numbered segment of cut cloth with entrapped biota).

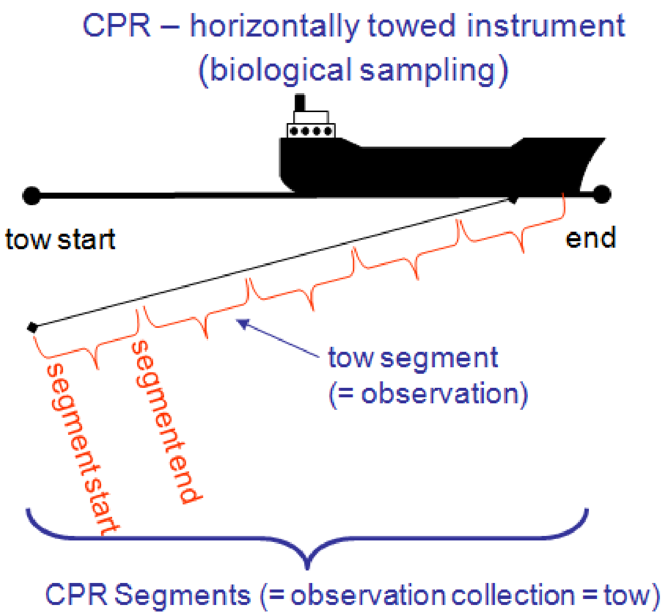

Figure 4.6 Characteristics of CPR Observations

```
<AODN:Dataset>
{include and import statements for re-using other GML-based ontology elements}
        <AODN:ObservationCollection>  {encompasses all observation members i.e., all segments and equals a CPR Tow}
        {properties describing bounding box, time, etc for all deployments}
                <AODN:Observation> {first observation member equal to a CPR Tow Segment}
                {project, platform, location, time, procedure, composite phenomena properties}
                        <AODN:FeatureofInterest> {CPR feature}
                        {CPR specific properties encoded in a specialised Feature}
                        </AODN:FeatureofInterest>
                        <AODN:Result> {CPR data values}
                                <AODN:domainSet>
                                {CPR domain properties and values encoded in a specialised Feature}
                                </AODN:domainSet>
                                <AODN:rangeSet>
                                {CPR range properties and values encoded in a specialised GML MultiCurve Coverage}
                                </AODN:rangeSet>
                        </AODN:Result>
                </AODN:Observation>
                <AODN:Observation> {second observation member – pattern repeats}
                        ..............
                        ..............
                </AODN:Observation>
        </AODN:ObservationCollection>
</AODN:Dataset>
```

Figure 4.7 Derived Model CPR Data Encoding Pattern

A problem was immediately encountered in ‘fitting’ these sample data to any of the available CSML (Version 2.0) Feature Type patterns (refer again to Table 4.2). The most suitable CSML Feature Type for a CPR deployment appeared to be a “Trajectory Feature”, represented by a distribution of irregularly timed observations embedded within a 2, 3, or 4-d compound spatio-temporal coordinate reference system. However, this CSML Feature Type did not permit a good modelling of the “tow segment” domains. As a consequence, a standard GML “MultiCurveCoverage” (OpenGIS, 2012) was used instead. This was a far better match for modelling the data as it was then possible to encode each CPR segment's domain (i.e., spatial coordinates) as a line with as many points as needed to appropriately depict the coordinates of the ship's track during collection of a particular CPR segment. The domain of the "feature-of-interest" in each "Observation" in the CPR case was a line segment (tracing part of the ship's track) as defined by spatial coordinates.

In a MultiCurveCoverage, the domain is partitioned into a collection of curves comprising of "gml:MultiCurve" data types. The Coverage function then maps each curve (member) in the collection to a value in the "gml:rangeSet" (which is an array). Figure 4.8 schematically depicts the CPR Tow Feature encoding.

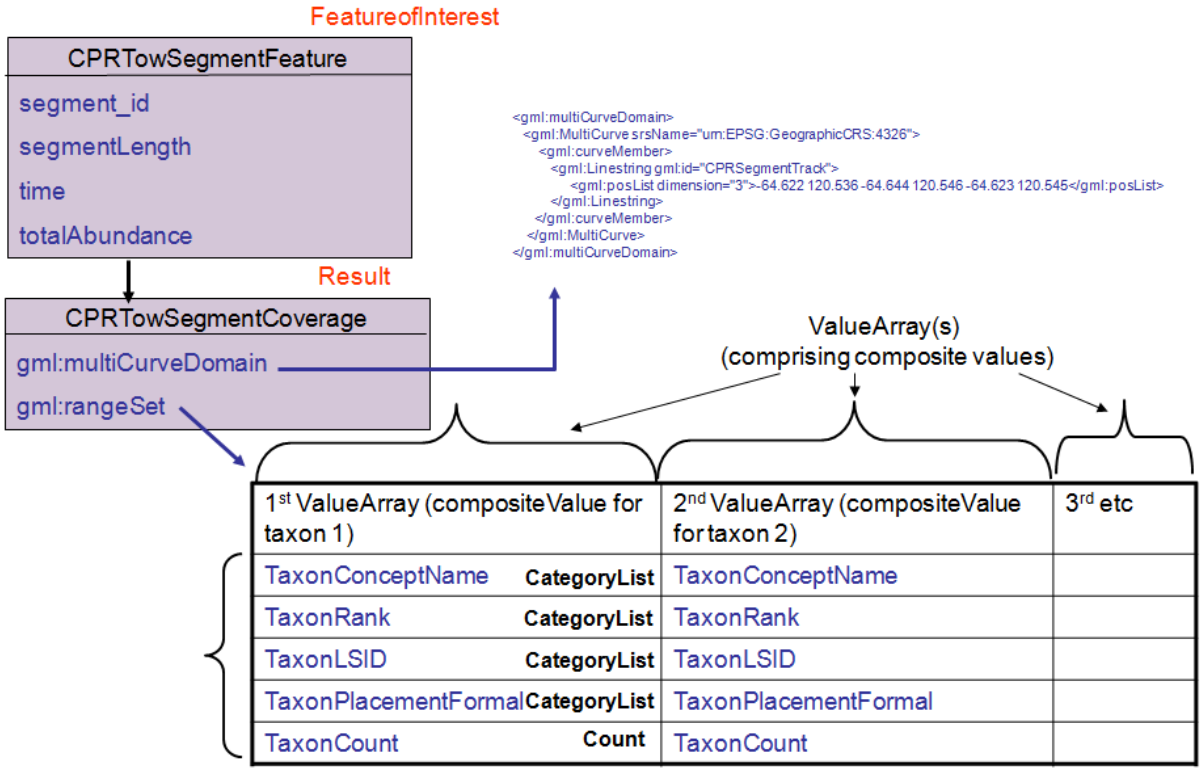

Figure 4.8 Schematic Representation Of A Portion Of The CPR Tow Segment Feature

In addition to the problem of not being able to find an appropriate CSML Feature Type to model the domain for CPR segments, it was necessary to develop specialised GML "value objects" as part of the 'gml:MultiCurveCoverage rangeSet' property to represent CPR biological data elements (refer again to Figure 4.8). 'Value' objects in GML are simply pre-defined data structures that can be used to record values or measured quantities. The biological data elements eventually used by the author were derived from concepts defined in various biological ontologies championed by the international Biological Standards Group (TDWG (TDWG, 2011a)). These elements were aggregated into a new specialised value object which was named "MarineBiotaStatistics" (see Appendix 5 - for the GML schema). The types of biological information encoded in this object included:

- **Taxon Concept Name**: A Taxon Concept is a named classification unit (or taxon) as explicitly defined in a taxonomic treatment to which individuals, or sets of other taxon concepts are assigned (TDWG, 2007b). A Taxon Concept Name can be a scientific name or a vernacular name. If it is a scientific name it is governed by a biological code of nomenclature. In the case of AODN data this could have been the Codes For Australian Aquatic Biota (CAAB (CMAR, 2010)) or the Register of Antarctic Marine Species (RAMS (Belgian Biodiversity Platform, 2012)).
- **Taxon Concept Life Science Identifier (LSID):** a globally unique identifier (GUID) provided by some type of authority that represents a stable reference to a Taxon Concept (Kennedy *et al.*, 2006). This property was factored into the AODN data model in anticipation of GUID's becoming common place in the near future.
- **TaxonomicPlacementFormal**: A comma separated list of scientific taxon names that indicate (for administrative and data exploitation purposes only) the taxonomic placement of this object. The words should represent taxa of decreasing rank (TDWG, 2007a)
- **TaxonRank**: An enumerated rank of a taxon e.g., "species", "sub-phylum", "family", "sub-variety" (TDWG, 2007c).
- **Taxon Count**: Total number of individuals of a particular taxa observed in the sample,
- **Taxon Maturation Stage**: A textual description of the dominant life stages evident in the sample.

While the "MarineBiotaStatistics" object was devised with the CPR data in mind, it was felt at the time that this construct could be made more general and re-used for other biological data types, for example in situations where biological statistics such as 'mean weight' and 'mean length' are ascribed to specific taxa.

These types of statistics are often recorded, for example, from analyses of fishing trawl and trap data. This "MarineBiotaStatistics" object was not considered appropriate, however, for observations made on individual biological specimens. Encoding the "MarineBiotaStatistics" object also required development of a new "composite phenomenon" (refer again to Figure 4.2 for phenomenon entity in O&M) which could be externally referenced from within the GML schema (as in the example provided in Appendix 6), or it could be declared locally.

The main observation drawn from this exercise, with respect to designing a Feature Catalogue, was the propensity for biological data to require complex data encodings, potentially involving observation collection Feature Types (i.e., observation Feature Types that are themselves collections of individual observation Feature Types) and the possibility that collections might also be required at the feature-of-interest level. A Feature Catalogue housing biologically-derived Feature Types therefore needs to be capable of expressing the relationships between the various constituent component Features that might make up a 'complex Feature Type'.

An analysis of the CPR data also demonstrated that the semantic description of the properties of biological Feature Types will, in many cases, require cross-referencing to other specialised external (to the Feature Catalogue) semantic repositories (e.g., species registers).

4.2.3 A Broader Review Of Biological Dataset Characteristics

Given the problems encountered with modelling just one, albeit complicated biological data type in the 2007 exercise and cognisant of changes that were taking place with respect to the O&M specification (in preparation for it becoming an OGC Standard), it was decided post this initial modelling exercise to re-investigate the nuances of biological data more generally (to better understand biological Feature Type patterns). This step was important since, as explained previously, designing an adequate Feature Catalogue content model depended upon accurately capturing the characteristics of the Feature Types that would populate it. The types of biological data investigated were therefore expanded to include datasets routinely exchanged via the SCAR network to give greater depth to the analyses, since the focus of the AODN community in 2007 was primarily on physical observation programs and SCAR encompasses significant biological science activity.

In this next set of investigations, instead of taking specific sample datasets and trying to fit them to a model (the approach taken in the previous step), a wide variety of datasets were inspected and their characteristics noted (see Appendix 7). It became apparent that a useful way of categorising the biological observations was to record the following traits:

- The **focus of the sampling**, or observation (e.g., in terms of whether the sample included single or multiple taxa),
- The **type of sampling** method used,
- The **characteristic spatio-temporal sampling dimension** (i.e., the typical geometric shape of the sampling and whether the sampling was fixed or time variant within the sampling shape),

- What the **purpose** of the **sampling** was (e.g., was the extraction of a specimen undertaken to understand more about that specific specimen or was it in fact to “sample” a wider population within a specific region),
- What the **minimal amount of sampled information** was that was **needed to codify** the observation,
- **Some examples** that could demonstrate the types of datasets that fell within a particular categorisation.

Not unexpectedly the variation in the datasets, from a data modelling pattern perspective, centred around whether the focus of the observation was a single biological taxa or multiple taxa, coupled with the spatio-temporal sampling regime. Using the O&M model as an anchor, although a wide variety of data were inspected, the variations could be expressed (at a basic level) by generating three modelling patterns. These models and their significance in formulating the Feature Catalogue design will be discussed next.

However, before examining these patterns, an important facet of biological data description came to light during the biological data review, which should be mentioned first. In contrast to many physical (Climate Science-based) observations and measurements, biological data, particularly in ecological studies is subjected to analyses that focus on the relationship of observed data to one or more additional variables (sometimes called ‘treatments’). Traditional analysis techniques used by biologists to look at these relationships include ‘regression’, ‘Analysis Of Variance (ANOVA)’, ‘logistical regression’ and ‘contingency tests’. Schneider (2009) explains that these techniques partition the data into a structural model plus an error term. The structural model is used to address issues of causality by partitioning variables. The error, or residual, helps address the question of uncertainty (i.e., given the variability in the data, does the difference between two treatments lie within the bounds of chance alone). Choosing the right type of model and selecting the most appropriate analysis technique depends on knowing the “scale” of the measured variable.

The most widely used taxonomy of scale is that of Stevens (1946, 1951) which distinguishes variables according to: a nominal scale (presence/absence); an ordinal scale (ranks); an interval scale (equal steps, such as are apparent in ^{0}C measurement) and a ratio scale (equal steps and with a known zero). This type of information is an important component of the semantics that should be encoded and communicated for biological data (and included in biological Feature Types, where appropriate). Scale is not a direct ‘property’ (or ‘attribute’) of a Feature Type, but is a quality of a Feature Type’s ‘attribute’.

4.2.4 Possible Biological Feature Type (Observation) Patterns

After a series of data modelling exercises conducted by the author, based on a general review of biological data (summarised in Appendix 7), three types of high-level biological observation pattern emerged which were capable of expressing the biological data covered, and which leveraged the core elements of the O&M model. The three resultant models are primarily differentiated on the basis of whether they are encoding data where there is a single taxon focus, where multiple taxa need to be described, or when multiple observations of multiple taxa occur (as in CPR-type datasets).

All three models use specialised “OM_Observation”, and “SF_SamplingFeatures”. Cox’s Sampling Features (refer back to Figs. 4.3 and 4.4) are used in the “feature-of-interest” role when the proximate “feature-of-interest” of an observation may not be the ultimate domain-specific feature, whose properties are of interest in the investigation, of which the observation is a part. This has been discussed earlier. Cox (2010a) explains there are two circumstances that can lead to using Sampling Features instead of domain ‘features-of-interest’: (i) when the observation does not obtain values for the whole of a domain feature; or (ii) when the observation procedure obtains values for properties that are not characteristic of the ultimate feature (of interest). In biological sampling, situation (i) is usually the norm and situation (ii) also occurs frequently enough to justify the use of the (OM) ‘SF_SamplingFeature’ pattern. For example, in biology, observations are made on one or a few organisms and those observations are then used to infer something about the entire population of that organism, or at least the sub-population comprising that organism within a particular piece of geography (at a particular time). In this situation the *Sampling* Feature is focused around observing particular biota, but the ultimate *Sampled* Feature (see Figure 4.3) is the biological assemblage at a particular location.

Pattern 1 - SingleTaxonObservation Model

The *SingleTaxonObservation* model (Figure 4.9) assumes that the object of the Sampling Feature is a single taxon, or a specimen. The Sampled Feature is generally the population that the taxon is drawn from or, in the case of tagged species which are often used in biology as mobile environmental sensors, the Sampled Feature may be the physical environment at a particular region. In this pattern the *SingleTaxonObservation* is the container for the observation identifier and perhaps details about the observer, observation sampling times, sampling location and result quality details (the latter three attributes inherited from the abstract OM_Observation). If the biological sampling has been performed using some type of platform and the act of sampling can be considered to involve a Deployment (see earlier discussions about a ‘Deployment’ Feature Type – section 4.2.1) then the

SingleTaxonObservation can be established as a specialisation of a *Deployment* (Feature Type). In this case *SingleTaxonObservation* will have a relationship with a *Project* and a *DeploymentPlatform*. It is also feasible that a *SingleTaxonObservation* may not be considered a specialisation of a *Deployment*, but data providers may still wish to associate *SingleTaxonObservation* with project information. In this case, the *Project* (Feature Type) should be mapped directly onto the *SingleTaxonObservation* entity.

In this first pattern, a *BiologicalTaxonSamplingFeature* is the "feature-of-interest" for a *SingleTaxonObservation*. The *BiologicalTaxonSamplingFeature* holds the taxonomic information for the taxa being sampled. Only the feature attributes: "id" and "taxonConceptName" are mandatory. All other properties, either measured or observed, that are made concerning the *SingleTaxonObservation* are expressed via an *ObservedProperty* (e.g., presence/absence, taxon_count,%_cover) which holds the characteristics which are said to be "carried" by the feature-of-interest. This *ObservedProperty* entity can carry composite property types, where multiple properties are observed. For convenience (and as a rule) it is assumed that in addition to having the "taxonConceptName" in the *BiologicalTaxonSamplingFeature* there must also be a "taxonConceptName" as a parameter in the *ObservedProperty* entity so that a parameter estimation (or sampling method) can be associated with the act of taxon identification (if desired).

ObservedProperties are input parameters to *ObservationProcess* Feature Types, which themselves are a component of an *ObservationSamplingProcedure*. Each *ObservedProperty*, if it has a measured value, can have a unit-of-measure ("*UoM*") entity. The *UoM* has one mandatory attribute called "description". The second attribute, "datum" is optional and serves to describe the reference frame in which a unit-of-measure is anchored. "Scale" is modelled as an attribute (or quality) of "ObservedProperty".

The *ObservationSamplingProcedure* holds information pertaining to the overall sampling processes associated with the observation event. Note that *DeploymentPlatform* (if it were to be included) is considered a Feature Type in its own right and is associated with a *Deployment* (as in the previous model outlined in Figure 4.5). A *DeploymentPlatform* is therefore not conceived of as being part of the *ObservationSamplingProcedure.* It is possible that some may consider a platform to be part of the description of a sampling procedure, which is a valid perspective. The modelling in this research, however, has taken the view that sampling procedure descriptions are focussed on the individual instruments (or components) and methods used to sense the observed properties at the level below platform.

Through its association with a Feature Type called *ObservationSystemCollection* the *ObservationSamplingProcedure* captures method and/or component (often instrument)-specific information for each parameter in *ObservedProperty*. In this context a component could be considered to be a computer program (e.g., statistical software package) if it is used to generate the property to which it pertains, or it could simply be an instrument. These procedure-based patterns have been borrowed, in part, from patterns used in Sensor Model Language (Botts, 2007). Only "processURI" and "processMethodDescription" are mandatory attributes in the *ObservationProcess* Feature Type. The *ObservationSystemsCollection* Feature Type supports one association, i.e., a "collection" association, which links the *ObservationSystemsCollection* to member *ObservationProcess* Feature Types (where the bulk of the descriptive material is expressed).

Pattern 2 - MultiTaxaObservation

The *MultiTaxaObservation* model (Figure 4.10) assumes that the objects of the Sampling Feature are multiple taxa or specimens. The Sampled Feature is generally a particular community or biological assemblages.

The difference between this pattern and pattern "1" is that the *MultiTaxaObservation* has as its feature-of-interest, a *MultiTaxaSamplingFeature* which identifies a number of taxa being described in the observation. *MultiTaxaSamplingFeature* is a type of *OM:SF_SamplingFeatureCollection*. The OM:*SF_SamplingFeatureCollection* supports one association, i.e., a "collection" that links a *MultiTaxaSamplingFeature* with a member - *BiologicalTaxonSamplingFeature*. Note that in Figure 4.10, the *Deployment* Feature Type and its associated classes have not been modelled, simply to reduce clutter in the diagram. As with the model in the first pattern, the MultiTaxaObservation could be a specialisation of a *Deployment* (if applicable).

Pattern 3 – MultiTaxaMultiObservation

The *MultipleTaxaMultiObservation* model, which is the most complex (Figure 4.11), also assumes that the objects of the Sampling Feature are multiple taxa, or specimens but in addition, the specialised observation is also a member of an observation collection. In this model you effectively have a collection of 'collections'. This is the type of model that would support serialisation of CPR-type data described earlier in section 4.2.2. As with the patterns above, the *Deployment* Feature Type has not been modelled in Figure 4.11, but can be plugged in as appropriate.

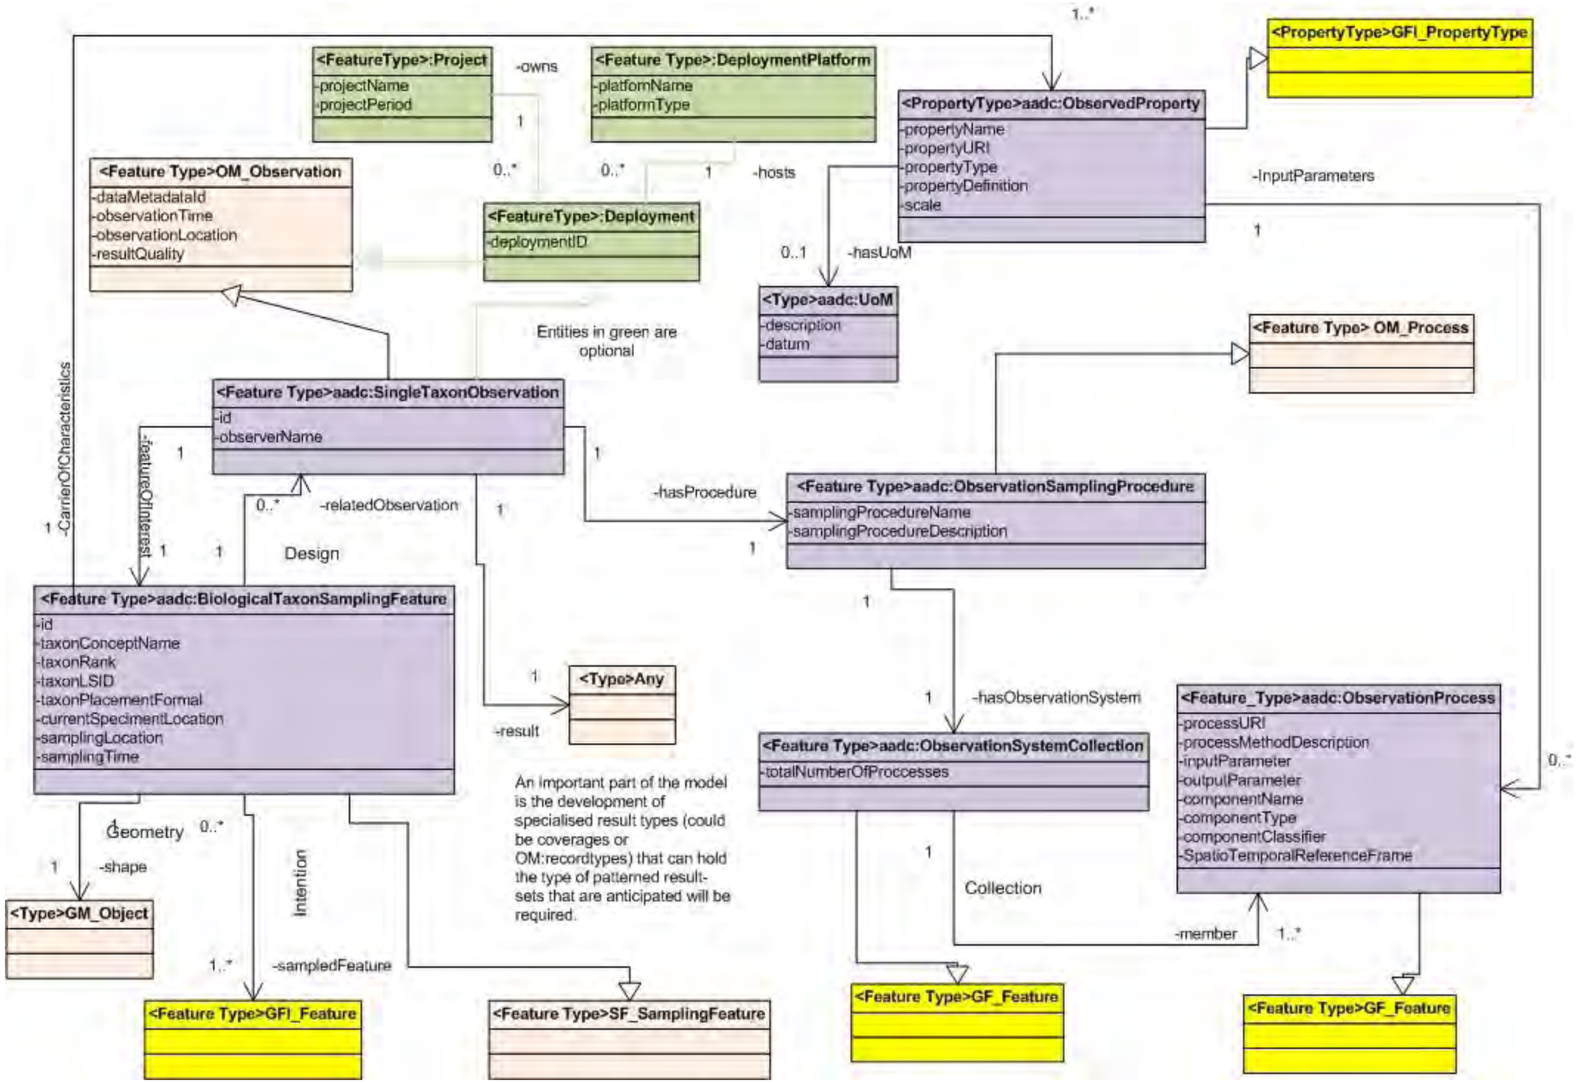

Figure 4.9 Single Taxa Observation UML Model

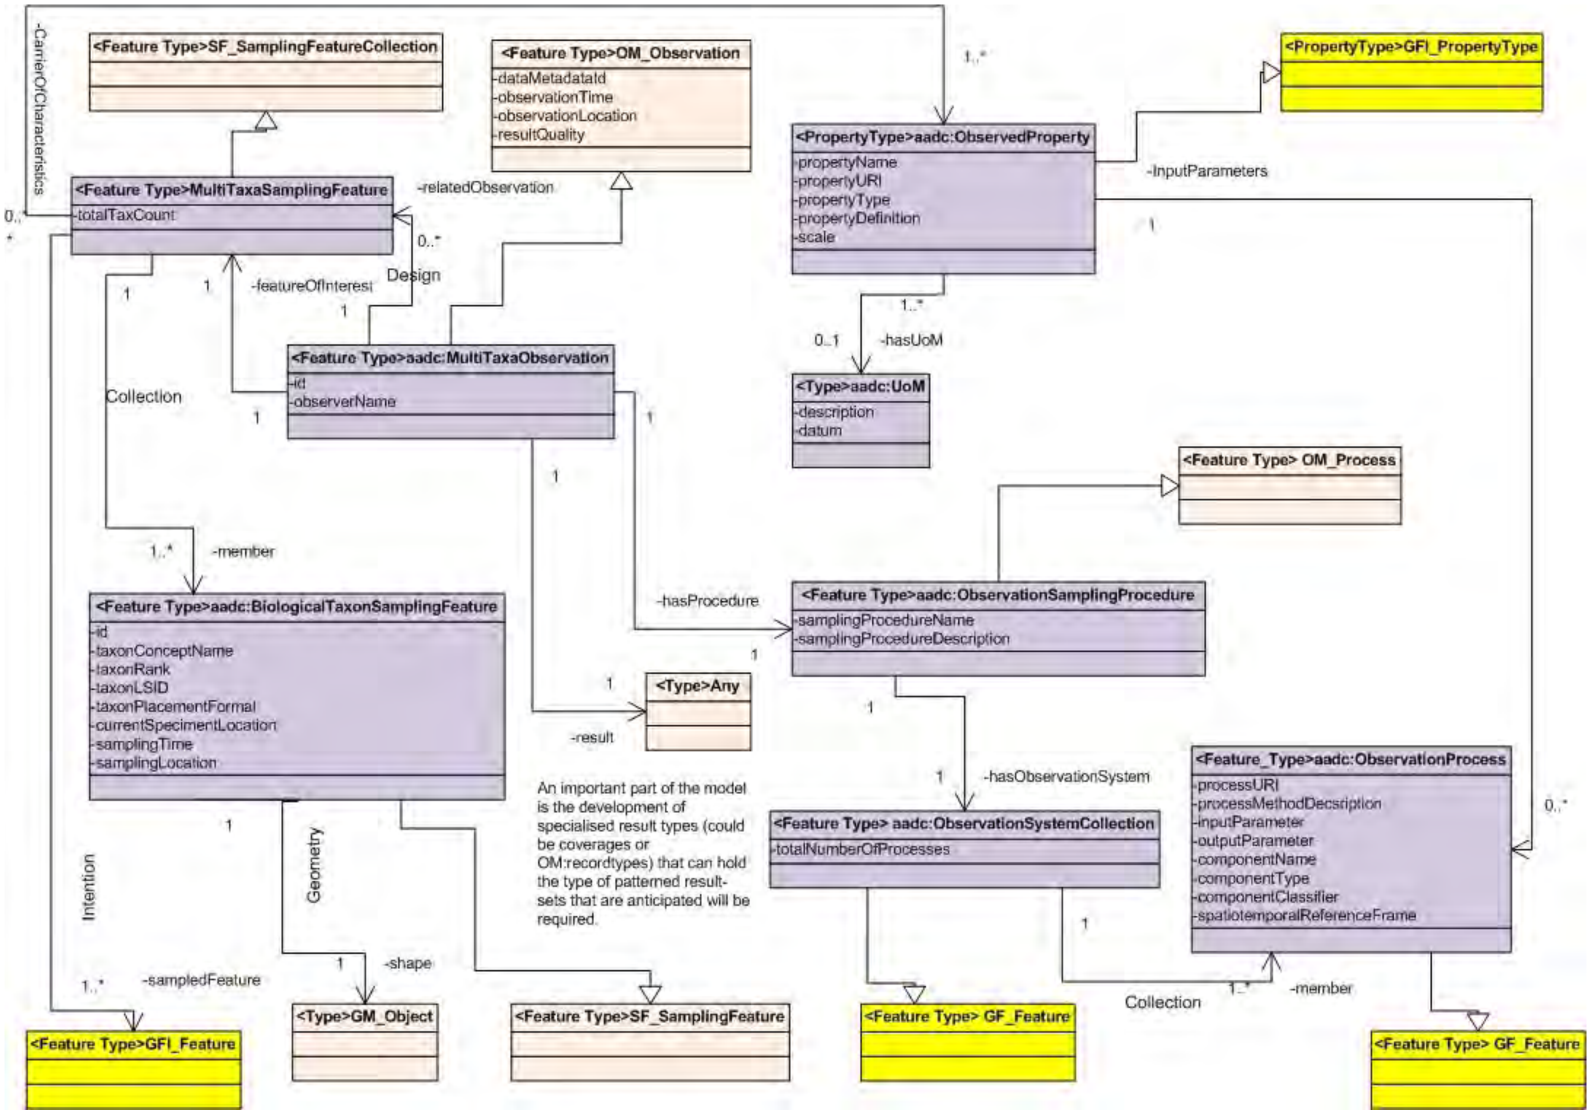

Figure 4.10 Multi Taxa Observation UML Model

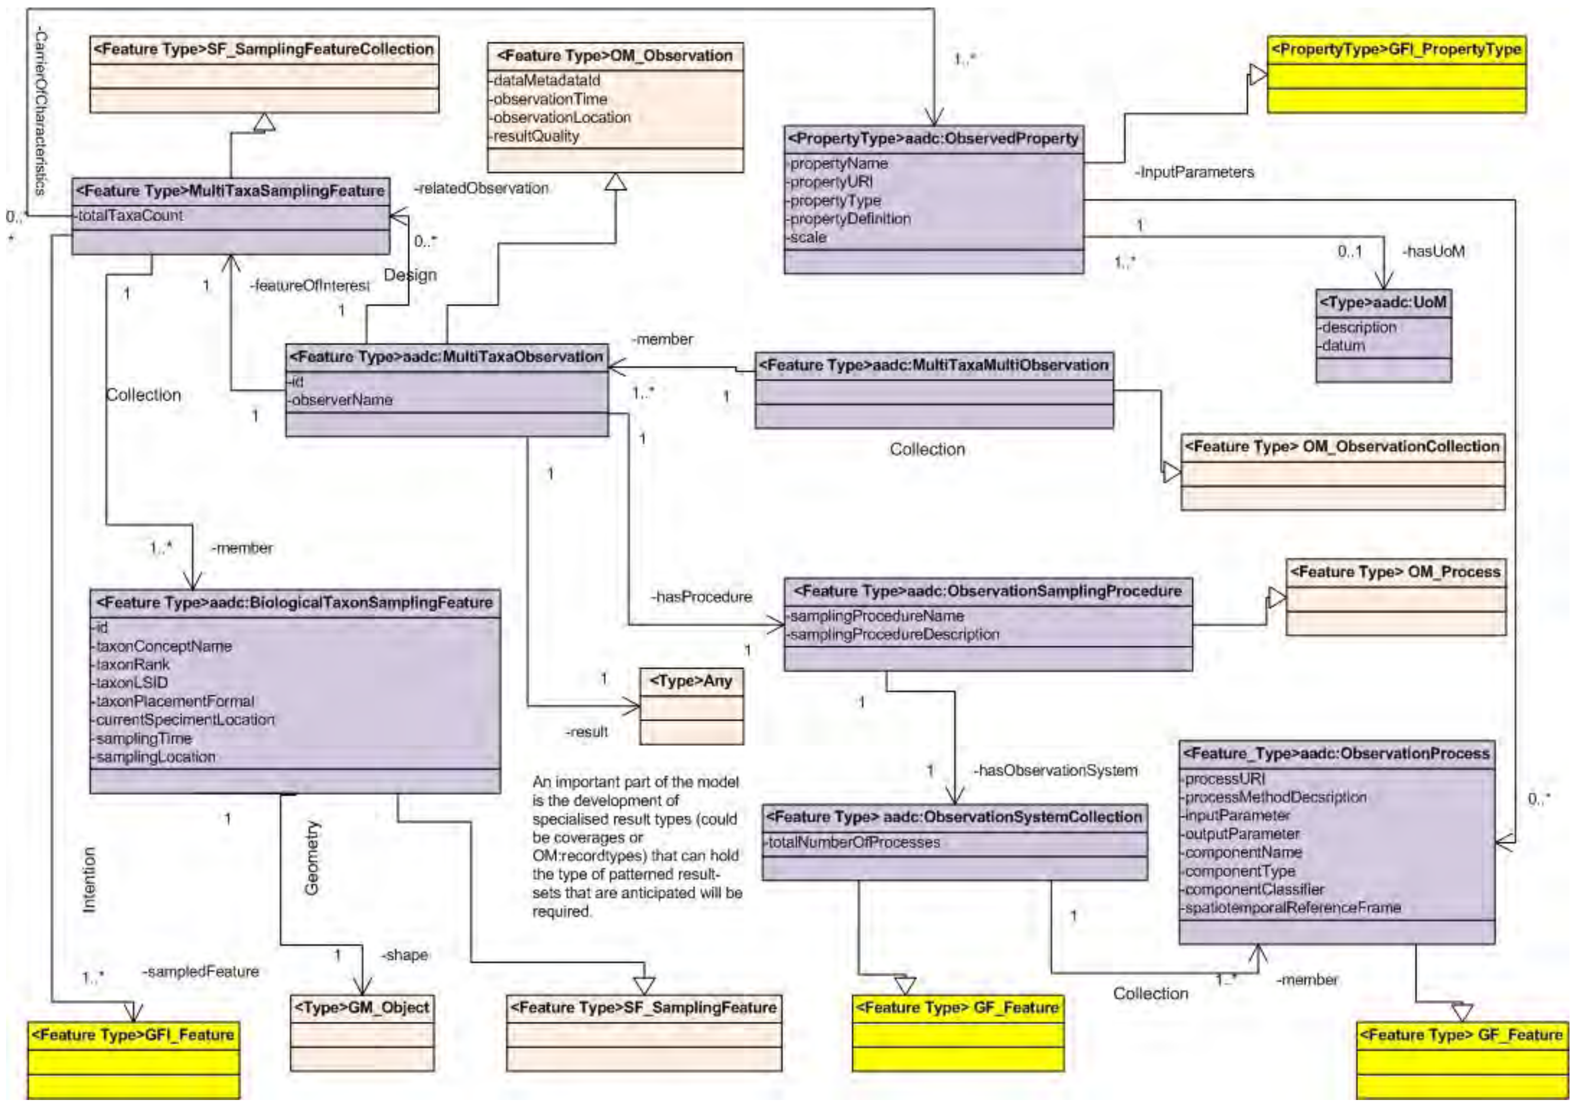

Figure 4.11 Multi Taxa Multi Observation UML Model

Specialisation Of All Three Models

The majority of the variation required within the patterns described above will relate to the methods selected for efficiently encoding the observation results. It is envisaged that there will need to be several specific result types (Feature, Coverage or Record) for the *SingleTaxonObservation* model (Pattern 1) to cover the spatio-temporal patterns identified in Appendix 7. Each result type will need to be structured to hold parameter values for the observations in scenarios where:

- location, time, depth and height are fixed, or
- location changes but time is fixed, or
- location changes as does time, or
- location changes (and depth or height) as does time.

Once developed, each result type can be plugged into the model depending on requirements.

The *MultipleTaxaObservation* model (Pattern 2) requires the same plug and play result types, however, the complication with these result encodings is that they must also cater for ascribing the listed parameters to individual taxa within the dataset. It is this pattern (and the *MultiTaxaMultiObservation*) pattern that could re-use a more generalised form of the GML-based "MarineBiotaStatistics" value object, described earlier in Figure 4.8, embedded within a suitable Coverage result type.

All three models are purposefully abstract. What is required is specialisation of the Observation entity in all patterns i.e., the *SingleTaxonObservation*, *MultiTaxaObservation* and *MultiTaxaMultiObservation*. In many cases, some specialisation of the Sampling Features will also be required, particularly to adequately factor in variations in temporal domains. See Cox (2011) for XML implementations of different O&M Observation types.

A specialisation should occur when two or more parties agree to define, maintain and govern a common Feature Type for the purposes of constraining the parameters (and their associated types) and parameter value representations in their exchanged datasets. Specialisation might be achieved according to some type of classification based on observation sampling method (e.g., *TrawlNetObservation* or *TransectObservation*), or perhaps according to certain types of parameterisation within the data (e.g., a specialised observation called *VegetationCoverObservation*). Specialisation in this context simply means that the named observation has a very specific set of mandatory and non-mandatory characteristics, but all still expressed according to the basic patterns described in Figures 4.9, 4.10 and 4.11.

It was not the purpose of this research to definitively establish all possible encoding schema, but to explore what type of patterns might suit the community's data and therefore, what types of elements would need to be managed within a community Feature Catalogue. However, a useful outcome of the research is the establishment of these generic, re-usable patterns.

More importantly, this broader modelling activity has confirmed that biological data in general, as suspected from modelling the CPR data, is often best expressed through complex Feature Types (i.e., collections of features-of-interest associated with collections of observation Feature Types). Sampling Features appear to suit many circumstances in biology where the observation feature-of-interest is not generally the ultimate domain feature being sampled (as in sampling fish in one location, but the data is being used to draw inference about a whole assemblage of fish, not necessarily even resident in the same geographic location). A Feature Catalogue must therefore be capable of modelling the relationships between 'Sampling', 'Sampled' and 'Observation' Feature Types, as well as Feature Types that play a supporting role in setting the sampling context (e.g., DeploymentPlatform and Project).

4.2.5 Defining The Feature Catalogue Use Cases

The broad requirements for the AODN information infrastructure were initially scoped by IMOS and the AODC JF data management community through a series of face-to-face meetings and a commissioned data scoping exercise, culminating in a standards workshop in March 2007. Joint AODC JF Technical Committee/AODN Development Office meetings, which subsequently guided infrastructure development, were ongoing from this period to the present day. Similarly the SCAR community outlined the development of its infrastructure, through annual meetings of its international SCADM technical reference group and through development of the SCAR Data Management Strategy (Finney, 2009). Refer to the previous chapter, Figure 3.4 for community stakeholder participants and their relationships.

To conduct research in this thesis, 'focus' groups were periodically formed drawing on participants from various SCAR and AODN community entities, to assist with either developing or confirming use-cases designed to make explicit the Feature Catalogue requirements.

Since the Catalogue has to sit as a component within the broader infrastructure, it is instructive to have a basic understanding of the function of the infrastructure itself. Three, high-level use-cases (Table 4.3) were articulated that commonly characterised the type of functionality that the AODN/SCAR infrastructures should ultimately support.

Table 4.3 AODN/SCAR High Level Use-Cases

Scenario	Description
Use Case 1: Actors: Description:	***General Data Discovery* & *Data Selection & Retrieval*** Data Users (and software client) A system user interacts with a Web-based client connected to a community-based services registry to search for and locate data of interest. The search paradigms offered encompass complex, multi-dimensional queries. Once sources for the data of interest have been found and a decision has been made to acquire the data, the client contacts the data source to request a copy of the data. The data is then sent to the client's browser.
Use Case 2: Actors: Description:	***Data Integration*** Data Users (and software client) A system user interacts with a Web-based client connected to a community-based services registry to acquire several similar datasets from a variety of sources. Having located the datasets of interest the client is able to visualise (e.g., plot) common attributes within these datasets as if these attributes were drawn from a single data source. The user can then elect to have the variably sourced datasets combined into a single dataset and a copy of this integrated dataset sent to his/her browser client.
Use Case 3: Actors: Description	***Inference*** Data Users (and a software client) A system user interacts with a Web-based client connected to a community-based services registry to acquire datasets from a variety of sources. They are presented with a list of instrument names, categorised by the types of parameters that these instruments measure and the features and properties that the instruments sample. The user wants to acquire data from any instrument that is capable of sampling sea-surface temperature, so selects all instruments listed under the client-presented navigation term "Ocean Temperature". The list in the Web-based client has been automatically built from querying an ontological data source that was capable of making assertions of the type: (a) All CTD instruments measure temperature and (b) A Seabird™ is a CTD instrument; which permits the conclusion that a Seabird™ measures temperature.

Both sets of infrastructure are currently capable of satisfying use-case 1. The AODN infrastructure is advancing towards satisfying use-case 2, but for a limited set of data where data providers have agreed to use the same data formats, data terms and Feature structures and neither community's infrastructure is currently capable of use-case 3. Neither community is currently using a Feature Type Catalogue to assist with data integration between data providers.

A focus group, drawn from the Australian Antarctic Data Centre (AADC), elaborated on the high level use-cases 1 and 2 above (in Table 4.3) by drawing on their experience in responding to requests for data (from national and international scientists). Typical information resulting from the focus group session is presented as simple statements of request. These are outlined in Table 4.4 and are separated into 'Discovery Use Cases' and 'Integration Use Cases'. All of the use-cases presented in

Table 4.4 resonated with members of the AODC JF Technical Committee and AODN Development Office, as well as with SCADM members who managed national Antarctic Data Centres.

Table 4.4 More Detailed Discovery and Integration Use Cases

Typical Discovery-Centric Use Cases
Show me all datasets that are only about taxon "X"
Show me all datasets that include taxon "X"
Show me the spatial distribution of taxon "X"
Show me the temporal distribution of taxon "X"
Show me all datasets measuring parameter "Z" for taxon "X"
Show me all datasets measuring parameter "Z"
Show me all datasets where parameter "Z" for taxon "X" is value "Y"
Show me all datasets for taxon "X" in the delimited region
Show me all datasets for all taxa in the delimited region
Show me all datasets from Principal Investigator named "I"
Show me all dataset from project named "P"
Show me all datasets that were sampled using method "M"
Show me all datasets that are in the same taxonomic class as taxon "X"
Show me all datasets captured from platform "PL"
Typical Integration-Centric Use Cases
Can I combine the values from parameter "Z" from dataset A with parameter "Z" values from dataset B (i.e., are the parameters measuring/observing the same thing, sampled using the same method and with the same units of measure).
Can I transform the values of parameter "Z" from dataset A so that they can be combined with parameter "Z" from dataset B (i.e., even though the parameters are measuring/observing the same thing, and have used the same sampling method, but their units of measure might differ).
Can I combine dataset A with dataset B (i.e., is there sufficient commonality of sampled features and parameters that the two datasets can be joined across these common points).

Armed with a basic understanding of what each community expected in terms of high level requirements from the infrastructure as a whole, the focus of investigations shifted to examining the types of scenarios in which a Feature Catalogue might be used within this infrastructure, to support the use cases articulated. Twenty four use-case vignettes were generated over a period of time (see Appendix 8), often after discussions with SCAR and SCADM community members, AADC and AODN Development Office staff and the AODC JF Technical Committee about their views regarding infrastructure requirements. These use-cases focus specifically on the anticipated functionality associated with a Feature Catalogue.

Figure 4.12 summarises the main classes of interaction that users (both human and machine) can expect to have with a community-based Feature Catalogue. The associated individual use-case vignettes (in Appendix 8) are presented and grouped according to these classes of high level interaction. Only unique requirements were expressed through the vignettes, i.e., those use-cases that essentially repeated requirements were omitted from the list of use-cases presented.

Through developing the use-cases it became clear that a single Feature Type Catalogue could be used to serve more than one community, and therefore the Catalogue must be capable of differentiating Feature Types on the basis of their community affiliation. Some Feature Types may be part of one community's lingua franca but may not be part of another community's domain of discourse. The use-cases therefore assumed the existence of some type of 'community profile', which was conceived of as a grouping of Feature Types (encompassing their associations and attributes), anchored to an explicitly declared 'community-of-interest'.

Additionally, the described scenarios supported situations in which two communities subscribing to the same Catalogue could share the same definition of a Feature Type, but differ in how they wished to express the number, range and characteristics of 'attributes' for that Feature Type. In this case we would be dealing with the same Feature Type, but each community may choose to apply the feature differently (e.g., use different units of measure for certain attributes, or include some attributes and not others). The Feature Type per se is describing the same object in both cases and is not semantically different.

The use-cases also identified that a community member should be able to retrieve the information regarding how their community has agreed to apply the Feature Type in question (consistent with the concept of a Feature Type as a 'template', refer to discussion in section 4.1). Additionally, a single community may wish to have the flexibility of defining a single Feature Type that has different characteristics depending on its spatio-temporal framework. For example, it is possible that a "mountain" Feature Type that is to be exchanged with a three dimensional (3-D) rendering might carry more attributes than a one-dimensional (1-D) mountain representation. Both descriptions are valid and both may be required.

A use-case describing access to Catalogue content indicated the requirement to browse Catalogue content via some type of classification scheme or by some aspect of a Feature Type's sampling regime, inferring that Catalogue Feature Type resources would be classified according to one or more classification schemes and intrinsically linked to contextual information associated with sampling.

Several use-cases either inferred or stated that Catalogue resources required annotation with audit and versioning information by association with 'history notes'.

Finally, from a Catalogue content model perspective one use-case anticipated the existence of complex Feature Types and indicated the need to be able to search for the various Feature Type components that made up the complex.

The vignettes present a very rich set of requirements and only the key characteristics which significantly affected the Catalogue content model design have been highlighted above. A number of vignettes associated with accessing Feature Catalogue content influenced the design and execution of the Catalogue interface discussed later in 4.3. For a detailed appraisal of all stated requirements refer to Appendix 8.

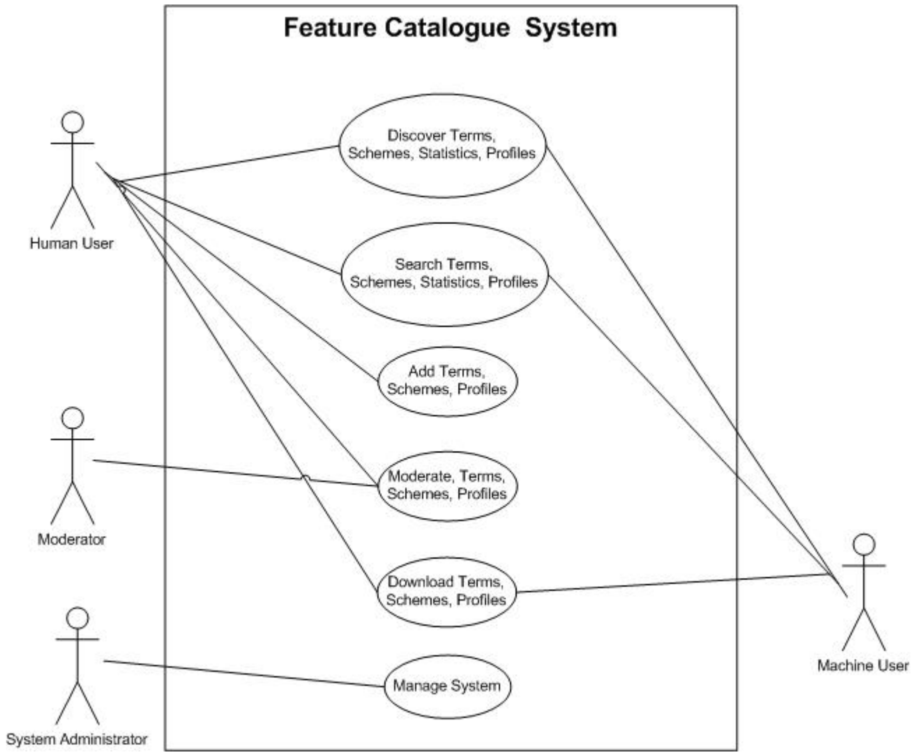

Figure 4.12 Key Types of Uses Cases

4.3 Review of Related ISO Standards For Defining Feature Types

Armed with the results of data modelling from sections 4.2.1 to 4.2.4 and the use-case scenarios referred to in 4.2.5, it was then appropriate to review the ISO 19110 *Methodology For Feature Cataloguing* standard (and the closely allied standards: ISO 19109 - *Rules For Application Schema* and ISO 19126 *Feature Concept Dictionaries and Registers*) in order to determine if this Feature Cataloguing standard met community requirements and whether it could be employed to develop a Feature Catalogue conceptual model, which could be ontologically grounded. This particular facet of the research addresses RQ1.1.2 "*Do relevant ISO and OGC conceptual data models meet implementation needs, and if not what is missing ?".*

As outlined earlier in Chapters 1 and 2, ISO 19110 provides the conceptual model for the elements that make up a Feature Catalogue. The UML Figure 4.13 shows that the key elements of a Feature Catalogue are: a Feature Catalogue entity; its constituent Feature Type(s); Feature Type Attributes which themselves may have particular 'values'; Feature Operations (i.e., operations that can be performed on Feature Types) and Feature Type Associations (which provide for the different types of relationship one Feature Type may have with another Feature Type).

In the ISO/OGC standards stack a Feature Catalogue is generally considered to be the designated repository for the semantic description of Feature Types. Feature Types are introduced in ISO's Generalised Feature Model (GFM, ISO 19101). The UML Figure 4.14 (which depicts the GFM in ISO 19101) shows how all Feature Types are essentially carriers of 'property class' characteristics which are differentiated (specialised) into 'operations', 'attributes' and 'associationrole' sub- classes. These basic classes are also directly reflected in the Feature Cataloguing standard (ISO 19110).

The ISO 19109 (*Rules For Application Schema*) standard ostensibly provides the principles by which dataset application schema should be encoded and draws on ISO 19110 for any embedded Feature Type definitions. ISO 19109 is also based on the GFM (ISO 19101). Somewhat confusingly, ISO 19126 (*Feature Concept Dictionaries and Registers)* is also a specification for a repository of semantic Feature Type definitions and has an association with both ISO 19110 (*Feature Cataloguing Methodology*) and ISO 19109 (*Rules For Application Schema*).

These standards were all reviewed with respect to their 'suitability' for meeting the Feature Cataloguing requirements of the Antarctic community. As a result, some observations are made and gaps are noted which helped to further inform the Catalogue design, but importantly, in the author's opinion these observations point to deficiencies in the Feature Cataloguing standard (which perhaps have contributed to the current lack of Feature Catalogue implementations).

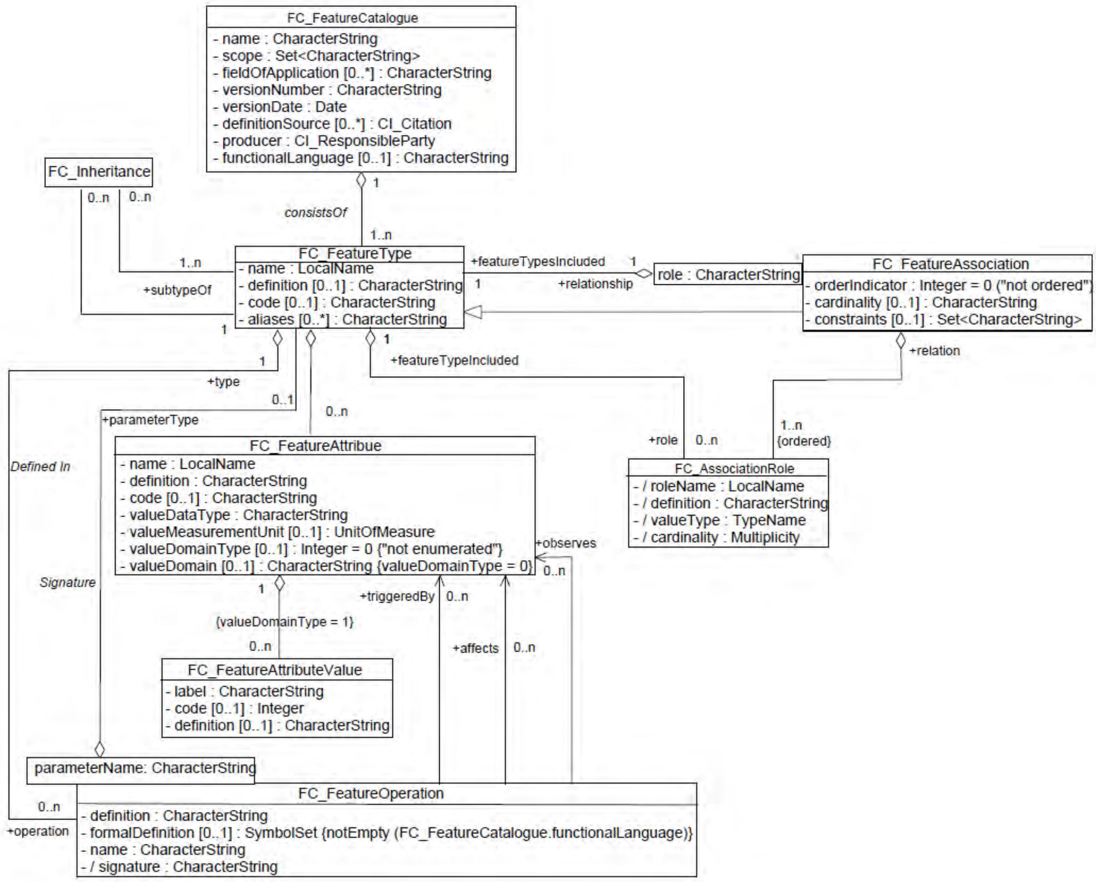

Figure 4.13 Feature Catalogue Conceptual Model in UML (source ISO 19110)

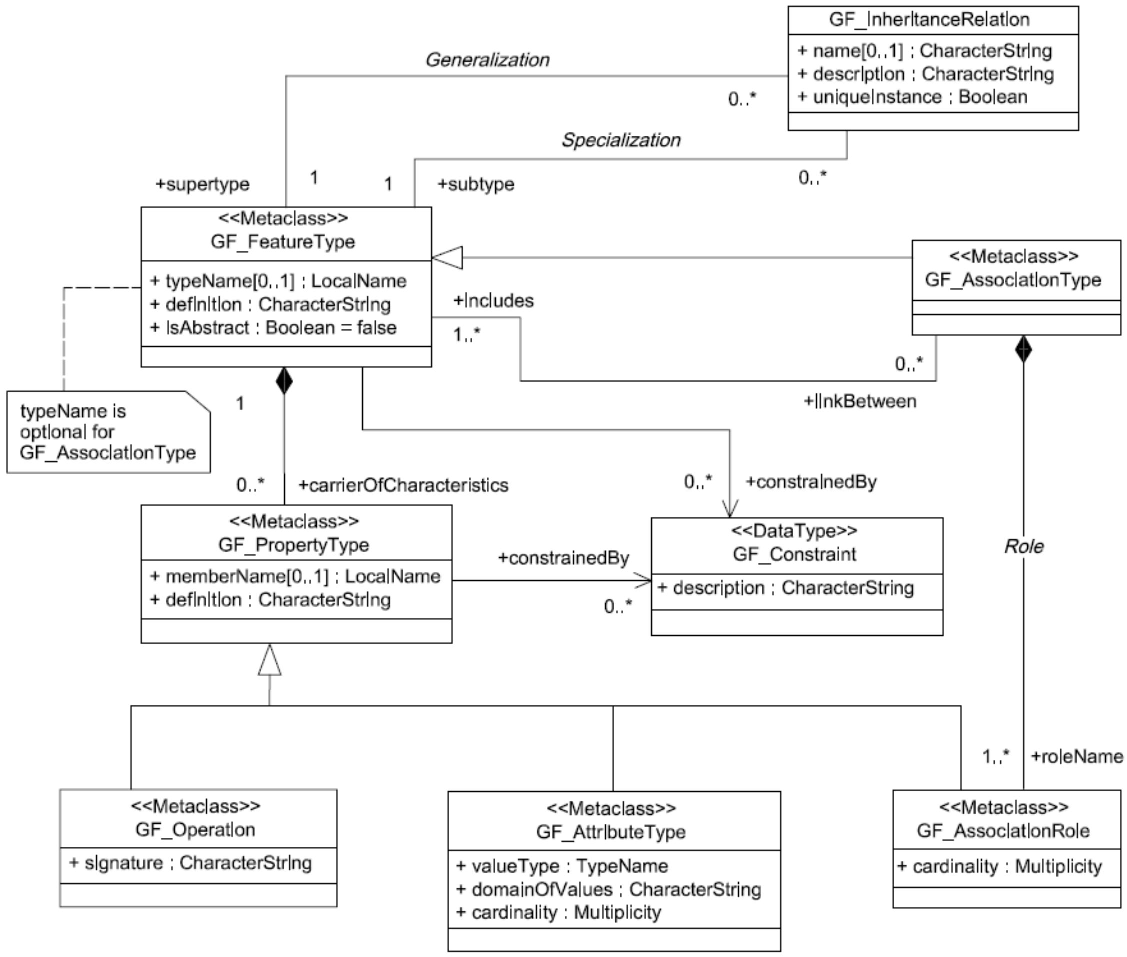

Figure 4.14 Generalised Feature Model in UML (source: ISO 19101)

Since this thesis has taken a predominantly observation-centric view of Feature Types, the extent to which the Feature Cataloguing standard accommodates important conceptual elements of the observation-centric (data model) patterns was specifically assessed during the review.

4.3.1 Apparent Overlaps and Contradictions In ISO Standards 19110, 19109 and 19126

ISO 19110 scopes the function of a Feature Catalogue in the following terms: "*Any set of geographic data is a greatly simplified and reduced abstraction of a complex and diverse world. A catalogue of feature types can never capture the richness of geographic reality. However, such a feature catalogue should present the particular abstraction represented in a given dataset clearly, precisely, and in a form readily understandable and accessible to users of the data*". The standard has a stated relationship with ISO 19109 – (*Rules For Application Schema*) whose purposes are to: "*set a framework for defining a common and correct (unambiguous) understanding of the content and* *structure of data within a particular application field, and to provide a computer readable schema for applying automated mechanisms for data management*". The distinction, at first glance, between these two standards appears relatively clear. One is setting guidance for defining a Feature Type (i.e., ISO 19110) and the other (ISO 19109) is setting guidance for encoding, for machine consumption, the entire dataset (inclusive of references to Feature Types). However, ISO 19109 goes much further than referencing Feature Types. A significant function of the standard is involved with setting up the modelling constructs necessary to adequately define Feature Types.

ISO 19126 (*Feature Concept Dictionaries and Registers*) declares its function to be: *"the provision of basic definitions and related information about a set of concepts that may be used to describe geographic features and shared across multiple application areas.* ***Elements from a feature concept dictionary may be re-used in one or more feature catalogues".*** In differentiating itself from ISO 19110, this ISO (19126) specification intriguingly also states that: *"A feature catalogue is often associated with a particular application schema, product specification, and data set* (**whilst a Concept Dictionary is not**). *It provides a* ***complete textual specification*** *of a set of feature types and their properties and relationships".* It is unclear to the author (given the definition supplied) how a 'Concept Dictionary' differs substantively from a Catalogue. It is also unclear why a Catalogue would routinely reference the semantics of its Feature Type resources (concepts) from an external source (such as a dedicated Feature Concept Dictionary) in preference to defining those concepts internally (since that is the key function of the Catalogue). Is a Feature Concept Dictionary something in competition with a Feature Catalogue ? If not, why do we have two and perhaps (as we will see later) three standards specifying how to define Feature Types? For example, Feature Concept Dictionaries and Catalogues both share the idea of 'Feature Type' (concepts) linked to properties such as 'Attributes', 'Operations' and 'Associations'. The main differences in these two standards appears to be that in a Concept Dictionary, Feature Type properties are not declared as being either mandatory or optional, but in a Feature Catalogue they must be. ISO 19126 (*Feature Concept Dictionary and Registers*), also provides for the classification of concepts into 'categories'. Recall that the ability to classify Feature Types was mentioned as a desirable quality for a Feature Catalogue in the use-cases outlined earlier, but the ISO 19110 (Feature Cataloguing) standard, unlike ISO 19126, does not provide for such classification. There are also some small differences, between the standards, in terms of the descriptors used to qualify some of the Feature Type (concept) properties.

The textual descriptions quoted earlier, from ISO 19126 (*Feature Concept Dictionary*) where the standard implies that a Feature Catalogue is often limited to only apply to one type of application schema, or dataset , appears out of step with guidance in the Feature Cataloguing standard (ISO

19110) itself. This author can find no reference in the ISO 19110 specification that implies using the Catalogue in such a limited way and it would appear highly resource-intensive (and unnecessary) to have a separate Feature Catalogue for each application schema used by a community of interest. Rather, for efficiency purposes it would be the opposite, that one Feature Catalogue would contain many Feature Types regularly used in multiple schema. This is an observation that the ISO 19110 (*Feature Cataloguing*) standard concurs with as it states: "*The availability of standard Feature Catalogues that can be used multiple times will reduce costs of data acquisition and simplify the process of product specification for geographic datasets*".

Upon deeper review, the stated functions (and scope) of all the three standards ISO 19110, 19109 and 19126), in parts, appear highly inter-dependent, and at times, over-lapping, circular, contradictory and unclear. A succinct demonstration of this point is the diagram reproduced here (as Figure 4.15) from ISO 19126 (Annex A) which is designed to show the interdependencies between ISO 19126 and other standards. The diagram is accompanied by this statement: "*Feature concept dictionaries are maintained as registers (ISO 19135); feature catalogues may be maintained as registers, but they may also be incorporated into documents such as product specifications. Both feature catalogues and feature dictionaries may be incorporated into systems of feature dictionaries and feature catalogues that reference each other*". Presumably, the problem that is trying to be addressed by having developed ISO 19126 is that the semantic meaning of a Feature Type (e.g., a mountain) can remain the same, whether it has one or twenty attributes, thus the attempt to separate out the general "meaning" of a Feature Type (concept) from all of the properties the Feature Type might possess. But ISO 19126 still includes Feature Type properties as part of its semantic definition of a Feature Type (concept). The author's argument being that this duplicates what is already in a Feature Catalogue and by enabling Feature Type concept definers to semantically qualify a Feature Type (concept) with properties, immediately limits its re-usability as an unfettered broad concept.

This ambiguity of scope and function between all three standards is further exemplified by reference to statements in ISO 19109 (*Rules For Application Schema*) which assert that: "*ISO 19110 provides a standard framework for* ***organizing and reporting the classification of features****. It also gives a* ***broader discussion on different aspects of geographic features"*** (but without clarifying what these 'broader discussions' actually are). It goes on *"This (19109) International Standard gives rules for creating application schemas, including the* ***principles for the definition of features***". It continues: "*This ....Standard* ***supports the definition of features*** *with respect to their representation in data structures defined by application schemas. ...The* ***definitions of the feature types and their***

properties**, as perceived in context of an application field, will be derived from the universe of discourse. A **feature catalogue documents the feature types". From these statements it would seem that ISO 19109 is where the Feature Types are actually 'defined' and they are then reported on, or documented in, a Feature Catalogue. It is not clear, from these statements, how detailed these 'documented' definitions need to be in ISO 19110 (*Feature Cataloguing*).

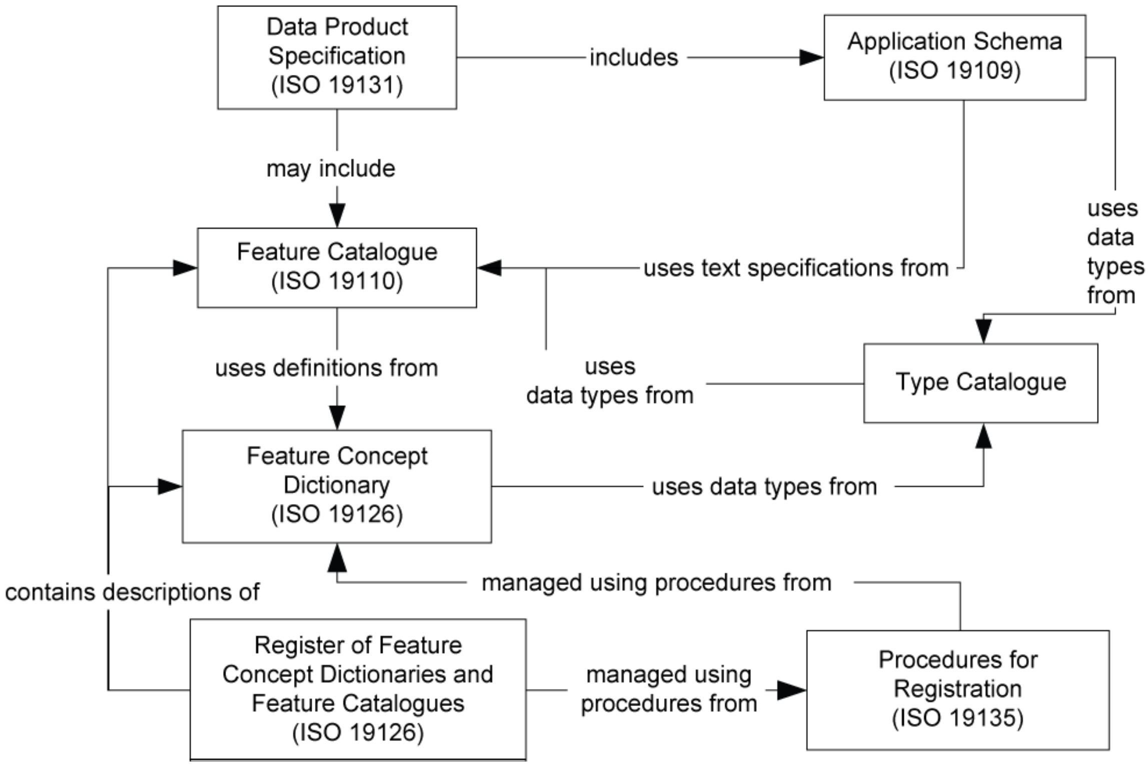

Figure 4.15 Dependencies Between ISO 19110, 19109 and 19126 (from Fig A.1, Appendix A ISO 19126)

From the texts quoted above (particularly those parts highlighted), concerning the interdependence between 19110 and 19109, the implication is that it should be possible, and indeed desirable, for an ISO 19110 conforming Feature Catalogue to be capable of documenting (**and therefore fully defining**) a Feature Type at a level of detail that is commensurate with its use within an application schema (and the Catalogue's structure should conform to the GFM). However, in ISO 19109 (*Rules For Application Schema*) the "**principles for the definition of features**" identify a richer set of constructs for describing Feature Types than are currently available in either ISO 19110 or 19126. So, a Feature Type in a Catalogue (if it was based solely on ISO 19110) would not be as semantically complete (or rich) as one appearing in an application schema. For example, the ISO 19109 standard has defined its own specialisations of attributes (see UML Figure 4.16) which are believed to be sufficient for adequately characterising features and their properties.

The specialised attributes defined are spatial, temporal, thematic and metadata-centric attributes, but these attributes are not available in ISO 19110. Importantly, the ISO 19109 standard also inserts the property type: ‘attributeOfAttribute’ which allows for further qualification of a GF:AttributeType (see Figure 4.16). An “attributeOfAttribute” links an attribute to another attribute that describes some characteristics of the first attribute. The example given in the 19109 (*Rules For Application Schema*) standard to demonstrate the use of this association is “*An attribute that carries the position of a feature may have another attribute that holds the positional accuracy (with data value of GF_QualityAttributeType) of this position*”. Why there should be a discrepancy between the expressiveness for defining Feature Types and their properties between the two standards (19110 and 19109) remains unclear to this author.

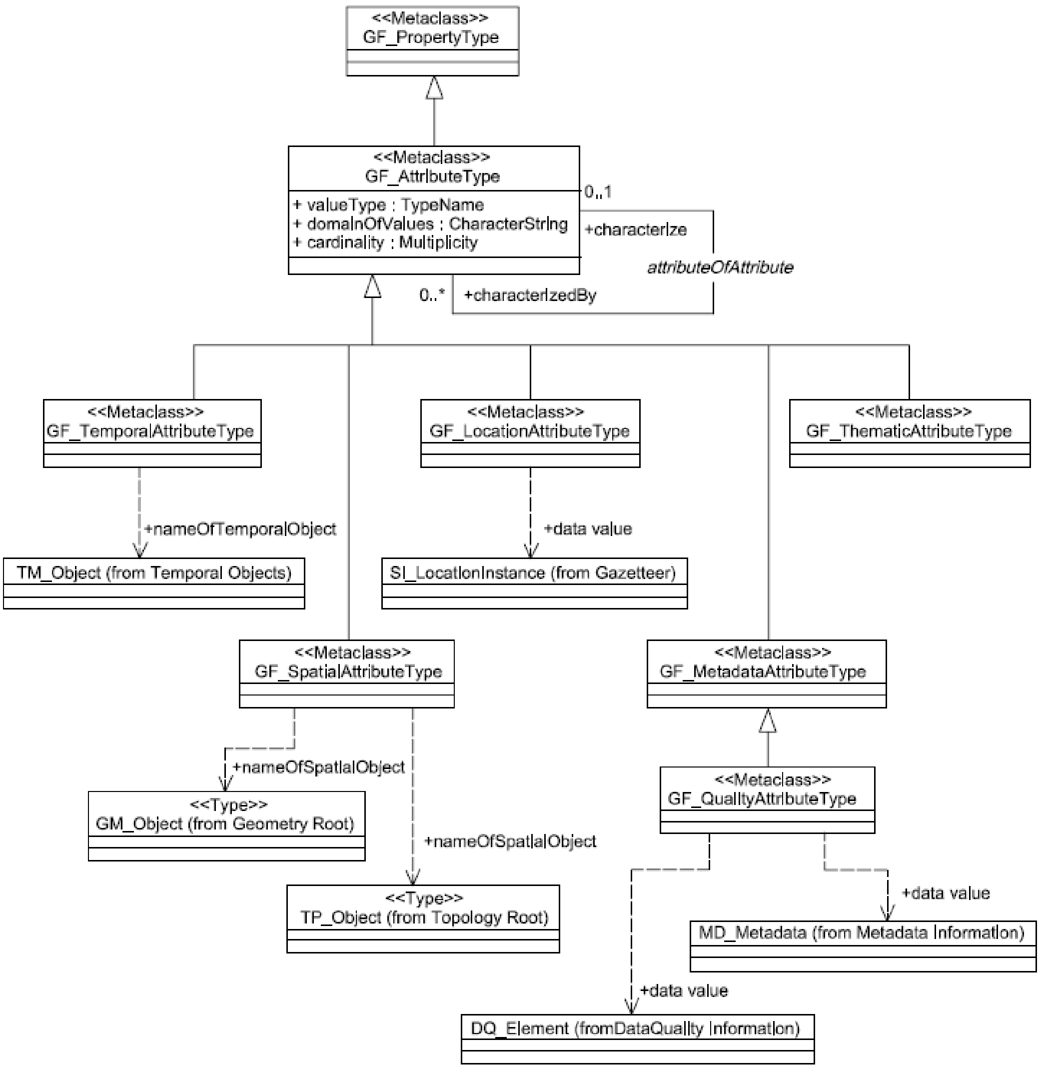

Figure 4.16 UML Diagram For Attribute Taxonomy from ISO 19109

The introduction of ‘attributeOfAttribute’ (in ISO 19109) is an extremely useful association which brings an added dimension to what can be expressed about attributes and it could be gainfully used within an enhanced Feature Catalogue Model, even though it is not part of the ISO 19110 (*Feature Cataloguing*) standard.

4.3.2 Referencing Between ISO 19110 and 19109

The interconnection between ISO 19109 (*Rules For Application Schema*) and 19110 (*Feature Cataloguing*) is made real from ISO 19109’s perspective through rules which require references from UML classes (or entities) in the application schema to an ISO 19110 compliant Catalogue. Presumably, in concrete instantiations of a schema, this is achieved via direct reference to the Feature Catalogue entry for a specific Feature Type (although the 19109 standard is silent on how this implementation should occur). In contrast, ISO 19110 makes no statement at all about how a Feature Type, documented within a Feature Catalogue might be linked to one or more example schema in which they are used. This latter type of functionality is clearly not considered in scope in ISO 19110, but linking to one or more sample schema (from within a Feature Type definition) may be beneficial for community members, because such sample schema serve as exemplars of how to construct conformant data services.

4.3.3 ISO 19110 Implemented As A Register

An important requirement mentioned in the use-case vignettes (in Appendix 8) was the ability to anchor Feature Types to particular communities-of-interest (called Feature Type Profile Users in the vignettes). The ISO 19126 (*Feature Concept Dictionary and Registers*) standard, whilst focussing on Feature Concept Dictionaries, indicates that Feature Catalogues could be established as standalone (and self-contained) entities or be established as registers. Feature Catalogues framed as registers, form an information modelling perspective, would have their contents registered to individual “register owners” and “control bodies” (ISO 19126). These particular register concepts could be also carried through, or accommodated, in standalone Feature Catalogues and will need to be considered when exploring how to conceptualise a Feature Catalogue meeting community requirements, particularly for the purposes of later ontologically grounding the Feature Catalogue model.

4.3.4 ISO 19110 – Omits ‘Collection’ Criteria

In the author’s opinion, perhaps one of the most severe limitations of ISO 19110, is that the standard specifically removes some types of descriptions from its scope:

"NOTE: The full description of the contents and structure of a geographic dataset is given by the application schema developed in compliance with ISO 19109. The feature catalogue defines the meaning of the feature types and their associated feature attributes, feature operations and feature associations contained in the application schema.

The ***collection criteria*** *used to identify individual real world phenomena and to represent them as feature instances in a dataset are not specified in this International Standard. Because they are not included in the standards, collection criteria should be included separately in the product specification for each dataset".*

Additionally, the ISO 19109 Standard also expands on what is not in scope for ISO 19110 by stating:

"NOTE: The feature catalogue provides the definition of geographic features at the type level, not the recording and representation of individual instances of each type. Therefore, it excludes ***spatial referencing, temporal referencing, and portrayal parameters*** *(see also ISO 19107, ISO 19108 and ISO 19117)".*

It is assumed that the terminology "collection criteria" used in 19110 can be assumed to be synonymous with "sampling", or "capture criteria".

It is argued here that to fully describe the "meaning" of a Feature Type (particularly one conceived from an observation-centric perspective), some elements of context are required that are currently omitted from the (19110) standard. Moreover, it is also not obvious why "feature operations" are considered important elements for inclusion in a Feature Catalogue, whilst aspects of the "collection" criteria, that may add significantly to the characterisation (or definition) of a Feature Type, are not.

The ISO 19110 standard asserts that: "*Feature operations are frequently included in the natural language definitions of the types. They are important for several reasons. First and foremost, they are the distinguishing characteristics that are embedded in the perceptions of the human beings who distinguish one type of geographic feature from another: they have psychological and behavioural significance to the people who use geographic information. Another reason is that computer systems are increasingly able to represent geographic phenomena, not just as a static set of maps, but as a dynamic representation of events occurring in geographic space in real time. Still another reason is that interoperability is an increasingly important goal in the design of geographic information systems. Functional equivalence of features is the key to interoperability of geographic information systems in the emerging open systems environment*".

If we examine the rationale given above for including “feature operations” in the Catalogue model, similar lines of argument can equally be made for the inclusion of “collection criteria”. Consider, for example, how one might describe a hypothetical TrawlNetObservation (i.e., an observation whose Sampled Feature are species assemblages, generated by the use of dragging a trawl net behind a ship for the purposes of catching biological specimens about which certain properties are recorded):

- The community’s natural language definition of this observation (Feature Type) of necessity includes reference to aspects of collection criteria (e.g., trawl net and ship), since they fundamentally characterise the observation type, and
- Inclusion of collection criteria will certainly help to distinguish this hypothetical Feature Type from other Feature Types, and
- An observation is itself an “event” which could be depicted dynamically because it inherently involves a temporal aspect, and
- Establishing “functional” equivalence, for interoperability purposes, between the TrawlNetObservation and other Feature Types, ostensibly of similar type, will rely upon articulating aspects of the collection criteria. Perhaps this would even require the inclusion of other criteria in addition to “trawl net” and “ship”, for example the specific “type” of trawl net used to sample the biota, if value is perceived as being added by doing so.

Although the 19110 standard doesn’t currently support documentation of “collection criteria” it would not be difficult to extend the standard to accommodate such elements. In fact, by borrowing from ISO 19109 (*Rules For Application Schema*), the association “attributeOfAttribute” could be included in an enhanced Feature Catalogue Model and used to attach sampling-related criteria to attributes that define observed properties.

4.3.5 ISO 19110 – Foreseeable Problems With Feature Type ‘Operations’

Whilst the “Operations” that a Feature Type can perform (or participate in) are no doubt of interest, the question arises as to whether operations are actually necessarily required to unequivocally describe a Feature Type. Additionally, one can speculate on how infinite the list of such operations might be for many real-world objects if operations are considered an intrinsic part of the semantic definition of a Type. Take for example something relatively familiar, such as a “building” Feature Type. Some hypothetical constructor functions (Operations) it might participate in (modelled on the samples provided in the Standard), might include: “building-livable” (e.g., an operation designed to estimate if it is fit for human habitation); “building-shadow” (e.g., an operation designed to estimate the shadows cast by the building through time) and “building-mass” (e.g., an operation for calculating the mass of the building from its constituent parts). It is not considered that any of these would be truly necessary for semantically defining a building and the list could be quite endless. By contrast, including certain aspects of "collection criteria" for observation features is often intrinsic to the definition of the observation (as has already been discussed).

4.3.6 ISO 19110 – Placeholder For UoM But Not For Datum

The ISO 19110 standard also seems inconsistent in its treatment of attribute qualifying concepts. In the standard there is an element called "valueMeasurementUnit", which is a characteristic of a "Feature Attribute", designed to record units of measure for attribute values. This, of itself, is not controversial but the standard does not include a placeholder for the reference frame in which this measurement is anchored (e.g., its datum). If the purpose of the Catalogue is to document the characteristics of a Feature Type so that a community can use these Feature Types as templates for constructing feature instances (served in datasets according to community agreed application schema), presumably both the 'unit of measure' and a measurement 'datum' will be required for a complete description. In the ISO 19110 standard examples, provided at the end of the standards document, a general constraint class called "FC_Constraint" is used to provide a reference frame for a "depth" feature attribute, whose "valueMeasurementUnit" is given as a "metre". This would tend to indicate that reference frames (or datums) are required and given their omission, a surrogate, in the form of the generic "FC_Constraint" class has been used to fill the void. Ideally the reference frame should be part of the standard.

4.3.7 ISO 19110 – Limitations On Temporal Referencing

It should also be noted that temporal referencing is not possible in ISO 19110 (except via the introduction of a specific "FC_Constraint" on a time-based Feature Attribute). Since all observations are events, time will almost always be a significant attribute of Feature Types in observation-centric paradigms. Given, for example, that ships as observing platforms will frequently cross time zones and date-lines during a voyage, stating something about how a community expects to record and reference time attributes within observations is critical.

4.3.8 ISO 19110 – Limitations On Referencing Feature Type Symbology

Finally, the exclusion of portrayal parameters from ISO 19110 means that it is not possible to assign symbology to Feature Types for the purpose of standardising, from a community perspective, how datasets served from distributed sources should be portrayed in Web clients. Whilst it is not expected that a Feature Catalogue would assume the function of a Symbology Catalogue, it would be beneficial if the Feature Type definition either linked out to a Symbology Catalogue, or documented sufficient symbology information to enable the communication of uniform portrayal information for specific Feature Types.

4.4 Summary Of Requirements For An Enhanced Feature Catalogue Conceptual Model

In the preceding paragraphs the adequacy of the ISO 19110 standard for documenting (and therefore defining) observation-centric Feature Types has been evaluated. Across a number of areas the ISO 19110 Feature Catalogue Model was found lacking, leading to the conclusion that enhancement and supplementation of the standard will be required. What now remains is to explore what an enhanced model should include such that it is capable of managing the types of observations (Feature Types) that have been described in earlier sections and which has the capacity to service the use-cases that have been articulated. As far as is practicable, to remain consistent with the OGC standards stack, the model should continue to be based on the Generalised Feature Model but should encompass the richer constructs established in ISO 19109 (*Rules For Application Schema*) and borrow some elements from ISO 19126 (such as the ability to classify Feature Types and then register items to specific communities-of-interest).

4.4.1 Enhanced Feature Catalogue Model

Table 4.5 distils the key Catalogue content model requirements necessary for satisfying the various use-cases that were outlined earlier (in Appendix 8) and which are needed for accommodating nuances in the data as a result of the modelling that was performed in sections 4.2.1 to 4.2.5. This table also summarises whether ISO 19110 currently supports stated requirements (as evidenced from the review that was just reported in 4.3). Figure 4.17 then depicts, in UML form, the major entities and associations that are thought necessary for an enhanced Feature Catalogue Model that is capable of supporting these requirements.

Descriptions for each entity that follow in Table 4.6 provide an indication of the types of attributes anticipated for each modelled UML class. Each attribute has been given a designation of “M” for mandatory, or “O” for optional (in the case of attributes that are desirable but not always required). For example, if a ‘ProfileUser” class is a machine user (e.g., an application such as a Portal) – a ‘URL’ attribute would be given for the application. However, this attribute has been designated as “Optional” because it is possible that a ‘ProfileUser’ is instead a human. A human user may have a Web page in which case the ‘URL’ could be completed, or they may not, in which case no ‘URL’ would apply.

Table 4.5 Summary Of Catalogue Content Requirements

Requirements	Fulfilled by ISO 19110
1. An ability to group (or link) Feature Types and their properties to more than one identifiable community.	No
2. Feature Types may be defined using one or more spatiotemporal objects. Semantically the Feature Type will be the same, except will have more than one spatiotemporal characterisation (e.g., a mountain is still a mountain whether it is represented by a point, a polygon or a solid).	No
3. A Feature Type with more than one spatiotemporal characterisation may have different attributes, or attributes with different units of measure.	No
4. Feature Types can be allocated to community-defined themes and themes will be part of community-defined classification schemes.	No
5. A Feature Type can exist in more than one theme and more than one scheme.	No
6. A Feature Type can have a named association with other Feature Types.	Yes
7. A Feature Type can reference a portrayal symbol.	No
8. A Feature Type and its properties can reference an external (ontological) definition.	Yes
9. An ability to log all updates and deprecations of Catalogue components.	No
10. It should be possible to document elements of the sampling methods used for observed or measured attributes (i.e., the observed properties of Feature Types).	No
11. It should be possible to document the units-of-measure used for observed or measured attributes and it should be possible for an attribute to have more than one unit-of-measure.	Partially
12. It should be possible to document the datum (e.g., vertical, engineering, geodetic, image or temporal) used to frame measured attributes.	Partially
13. A Feature Type can have one or multiple attributes.	Yes
14. An attribute can be listed as mandatory or optional	Yes
15. A reference to sample application schema(s) in which the Feature Type is used.	No
16. A Feature Type belongs to a Catalogue	Yes

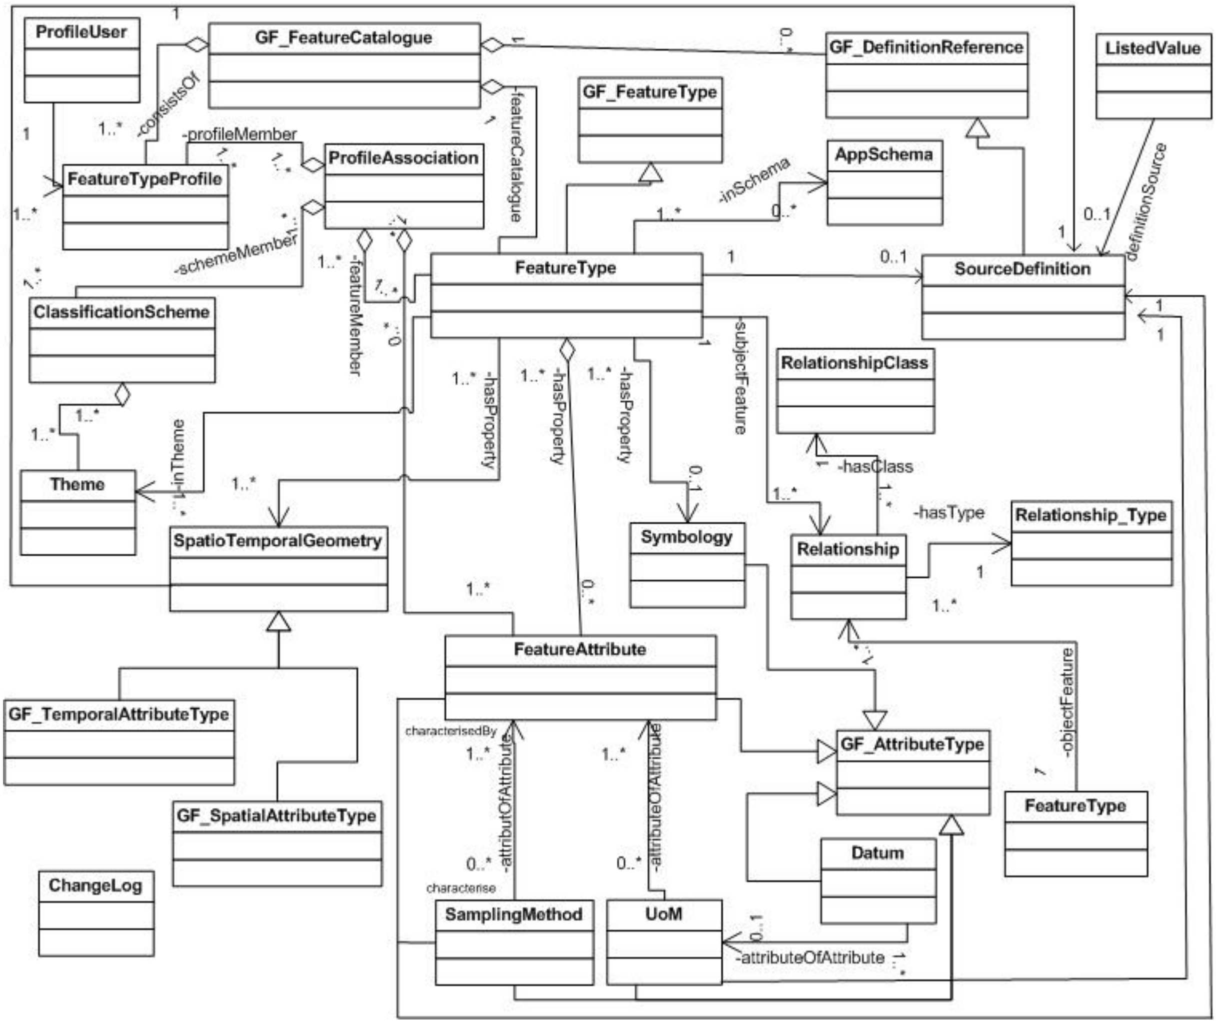

Figure 4.17 UML Class Diagram For An Enhanced Feature Catalogue Model

Table 4.6 Attributes Of The Classes Shown In Figure 4.17

FC:Class and Attributes	**Description**	**Obligation (M or O)***
Class:FeatureCatalogue	A feature catalogue contains its identification and contact information, and a description of some of the feature types and other information necessary for those definitions (as per ISO 19110).	M
Class: FeatureType	Class of real world phenomena with common properties.	
featureTypeCode	Unique ID within the Feature Catalogue	M
featureTypeName	Name given to the Feature Type	M
featureTypeDefinition	Lexical description of the Feature Type	M
isAbstract	Indicates if the feature type is abstract or not (Boolean)	M
deprecatedFlag	A binary flag indicating whether the Feature Type is no longer in use and has been	M

	deprecated.	
Class:FeatureTypeProfile	A collection of Feature Types and their properties which belong to a specified user community.	
featureTypeProfileId	Unique ID (or name) within the Catalogue.	M
featureTypeProfileName	Local name given to the Profile.	M
featureTypeProfileDecsription	Lexical description of the Profile.	M
moderatorName	Name of the person maintaining and moderating the Profile.	M
moderatorEmail	Contact email of the person maintaining and moderating the Profile.	M
profileCreationDate	Date on which the Profile was created.	M
profileUserCode	The unique ID for a given ProfileUser.	M
deprecatedFlag	A binary flag indicating whether the Feature Type is no longer in use and has been deprecated.	M
Class:ProfileUser	A class of user that owns a FeatureTypeProfile. The ProfileUser could be a named community, an application (for example the AODN Portal) or a community project.	
profileUserId	The unique ID for a given ProfileUser.	M
profileUserName	Local name given to the Profile User.	M
profileUserEmail	Point of contact.	O
profileUserURL	URL to an application.	O
profileUserDescription	Lexical description of the ProfileUser	M
deprecatedFlag	As above.	M
Class:ClassificationScheme	A scheme that classifies Feature Types on any basis. It is composed of Themes.	
classification SchemeId	The unique ID for the Classification Scheme.	M
classificationSchemeName	Local name given to the scheme.	M
schemeCreationDate	Date on which the scheme was created.	M
deprecatedFlag	As above.	M
Class:Theme	A Theme is a user-defined group to which one or more Feature Types belong. They are components of a ClassificationScheme.	
themeId	The unique ID for the Theme.	M
themeName	Local name for the Theme.	M
themeDescription	Lexical description of the Theme.	M
deprecatedFlag	As above.	M
Class:SpatioTemporalGeometry	A lexical description of the domain coordinates (temporal and spatial) of a feature-of-interest. The descriptions are drawn from a code list of predefined types. The domain coordinates set the framework for a range of properties associated with an observation result type.	
spatioTemporalGeomDescription	Lexical description of the SpatioTemporal-Geometry	M
shape	Geometry (from GF Spatial Types)	M
spatioTemporalGeomValue	The value of the SpatioTemporalGeometry code.	M
deprecatedFlag	As above.	M

Class:FeatureAttribute	A characteristic of a Feature Type.	
attributeId	The unique Id for the Attribute.	M
attributeName	The local name for the attribute.	M
attributeDescription	The lexical description for the Attribute.	M
attributeValueDataType	The type of attribute drawn from a nominated listed namespace.	M
enumeratedValueDomain	A boolean flag indicating if the attributeValue is drawn from a code list.	M
attributeListedValue	The value of the Attribute According to a code table).	M
attributeScale	Scale (e.g., interval, ratio etc) drawn from a code list.	O
attributeMandatory	A boolean flag indicating if the attribute is mandatory in the FeatureTypeProfile.	O
deprecatedFlag	As above.	M
Class:SamplingMethod	A class containing attributes that describe the method by which a property has been derived, generated or sampled.	O
methodId	The unique Id for the method.	M
methodName	Local name for the method.	M
methodDescription	A lexical description of the method.	M
instrumentId	A unique Id for an instrument.	O
instrumentName	A local name for an instrument.	O
instrumentType	A type of instrument drawn from a code list.	O
instrumentClassifier	Another attribute for further classifying instruments (drawn from a code list).	O
deprecationFlag	As above.	M
Class:UoM	A unit of measurement.	O
uomId	The unique Id for the measurement.	M
uomName	Local name for the measurement.	M
uomDescription	A lexical description of the measurement.	M
deprecationFlag	As above.	M
Class:Datum	A reference frame for a measurement.	O
datumId	The unique Id for the datum.	M
datumName	Local name for the datum.	M
datumDecsription	A lexical description of the datum.	M
deprecationFlag	As above.	M
Class:Symbology	A class containing attributes that depict how a Feature Type will be displayed in an on-line mapping system.	O
symbologyName	A local name for the symbology (e.g., grey star).	M
symbologyDescription	A lexical description of the symbology.	M
symbologyURL	A URL to a sample of the symbology or to the symbol entry in a symbology catalogue.	M
symbologyDisplayScaleRange	Some additional information about display scales for the symbol.	O
deprecationFlag	As above.	M
Class:Relationship	A general relation that links one feature type to another.	
relationshipId	The unique Id for the relation.	M
relationshipType Name	The type of relation between Feature Types	M

	(e.g., “partOf”).	M
relationshipClassName	The type of “ontological” relationship drawn from a code list (e.g., Equivalent, Symmetric, Transitive, Inverse, Functional).	M
relationshipSubjectFeatureType	The FeatureType participating as the subject of the predicate.	M
relationshipObjectFeatureType	The FeatureType participating as the object of the predicate. Could be a URL.	M
deprecationFlag	As above.	
Class:Relationship_Type	The type of relation between Feature Types (e.g., “partOf”).	
relationshipTypeNameId	The unique Id for the relationshipTypeName.	M
relationshipTypeName	The local name.	M
relationshipTypeNamedescription	The lexical description for the TypeName.	M
deprecationFlag	As above.	M
Class:RelationshipClass	The type of “ontological” relationship drawn from a code list (e.g., Equivalent, Symmetric, Transitive, Inverse, Functional).	
relationshipClassNameId	The unique Id for the relationshipClassName.	M
relationshipClassName	The local name.	M
relationshipClassNamedescripion	The lexical description for the ClassName.	M
deprecationFlag	As above.	M
Class:SourceDefinition	A description of the source of a definition which can contain links to on-line resources.	
definitionAuthority	Citation of the source.	M
onlineOntologyFlag	Boolean flag indicating if the definition includes an online ontological definition.	M
ontologyURL	Navigable URL to an ontological definition.	O
Class:ListedValue	Value for an enumerated feature attribute domain, including its codes and interpretation.	
valueLabel	Value of the attribute.	M
valueDefinition	Lexical description if required.	O
Class:AppSchema	A reference to online application documentation that contains the feature type.	O
appSchemaName	A local name for the schema.	M
appSchemaURL	A URL linking to the schema.	M
Class:ProfileAssociation	An entity that groups the elements (Feature Types, Attributes, and Schemes) of a Profile.	
featureTypeCode		M
featureTypeProfileId		M
classificationSchemeId		M
attributeId		M

*M=Mandatory, O=Optional

4.4.2 Encapsulating Observation Features Using The Enhanced Model

The enhanced Feature Catalogue content Model (just described in Figure 4.17) has been purposefully derived to cater for the adequate documentation of observation-centric Feature Types. However, it should be demonstrated how an observation (Feature Type), modelled using the types of patterns provided in Figures 4.5, 4.9, 4.10 and 4.11, could be mapped (i.e., encapsulated) into the classes depicted in the enhanced Feature Catalogue content Model (as an example). This 'round-trip' desk-top mapping exercise effectively 'checks' whether the Catalogue content model can accommodate the types of datasets being captured by the communities-of-interest (at least those that conform to the patterns previously described).

To recap, in the data models developed in this study, it is anticipated that perhaps an O&M "feature-of-interest" and/or the O&M "procedure" will be used to specialise the observation type. The observation focuses on the data collection event and the act of observation serves to assign a value to a property of a feature-of-interest. A description of an instrument, or sampling process, familiar within the application domain, is the value of the observation procedure.

Observation and measurement 'SamplingFeatures' have been used in the biological data models (e.g., refer to Figure 4.9) to sample some feature-of-interest in the application domain. Although Cox (2010a) asserts that these "*artefacts of an observational strategy have no significant function outside of their role in the observation process*", this author disagrees and considers that their description is important within a community whose focus is on monitoring and observation regimes (as in the case groups in the AODN community and many observing and monitoring programs associated with SCAR). The description of the 'Sampling Feature' provides an intrinsic element of the observation protocol, along with the observation procedure and is fundamental to the decomposition of the domain spatio-temporal geometry (particularly in the case of Coverage-valued observation results). So, 'Sampling Feature' should be important types of features in a Feature Type Catalogue.

It has already been argued that conceptually an observation (Feature Type) is always treated as a complex (i.e., a multi-part) feature, for the purpose of its description within the Catalogue. To demonstrate one option for how the various elements of an observation could be mapped into the Enhanced Feature Catalogue's Conceptual Model, a hypothetical biological observation was used in a desk-top encoding exercise. The hypothetical observation chosen was a specialisation of the simplest of the generic patterns presented earlier, i.e., the pattern identified in Figure 4.9. This hypothetical observation is a relatively straight-forward entity involving a single biological specimen. For this fictitious example we will assume that only one biological specimen was observed and collected.

This hypothetical *SpecimenObservation* uses a *BiologicalTaxonSamplingFeature* as its feature-of-interest directly (see Figure 4.9 for observation Feature Type patterns). This hypothetical observation is made by a net deployed from a ship. The main properties to be recorded are "taxonConceptName", "samplingTime" and "weight". The "weight" property will be obtained using a

"Doran 50kg digital fish hook scale". Units-of-measure for "weight" will be "kilograms". The "sampling time" will be in "hours", "minutes" and seconds according to "Coordinated Universal Time" (UTC). The Sampled Feature is the population of *Myctophidae sp*. in Prydz Bay (Antarctica).

The desk-top mapping exercise undertaken, which encodes the observation just described, into the enhanced Feature Catalogue Content Model (Figure 4.17) is presented in Appendix 9. Not all of the mandatory attributes (in the modelling pattern shown in Figure 4.9) were populated. Only those details sufficient to demonstrate the mapping pattern are provided. The mapping involves the development of five Feature Types: a Deployment Observation (SpecimenObservation), a Project (BiologicalProject), a Platform (Ship), the Sampling Feature (BiologicalTaxonSamplingFeature) and the Sampled Feature (FishPopulationFeature).

4.5 Casting The Feature Catalogue Model As An Ontology

The basic design of the enhanced ISO 19110-based Catalogue Content Model was developed in the preceding section after an extensive investigation into community requirements and was tested on paper (in Appendix 9) using a hypothetical observation. But the real value-add for communities who invest in the establishment of such a Feature Catalogue will be derived from the re-usability of Catalogue content in support of the many applications that will require, or will be enhanced by, access to qualified machine interpretable information. Simple textual definitions, as espoused through ISO 19110, are insufficient to provide for semantically-enabled machine-based processing. Ideally, Catalogue content should be accessible in a range of forms and formats, including its expression in an ontological form. The more stringent the ontology is in terms of its modelling rules, the lesser the ambiguity of the semantics. But increasing expressivity also means greater complexity in terms of Feature Type Catalogue development, its application and deployment. This section therefore addresses RQ1.1.3 –"*How can a conceptual model that meets Antarctic science requirements be semantically grounded ?".*

As discussed earlier, Lassila and McGuinness (2001) have classified ontology types according to the richness of their internal structure and present ontology classification in the form of a spectrum, with controlled vocabularies at one end of their spectrum and ontologies that express general logical constraints at the other end. The cut-over between "informal" and "formal" ontologies along this spectrum was demarcated based on how an ontology specified and applied the "is-a" relationships. Recall from earlier chapters that in "formal" ontologies if "B" is a sub-class of "A" and an object is an instance of "B", then the object is also an instance of "A". In "informal" ontologies, whilst the syntax exists to make these types of assertion, there is no “rule-base” associated with the expression that supports the truth of the statement.

Unequivocal semantic typing might be preferred over the more ambiguous informal systems, but as previously explained there is a significant overhead associated with using very formal types of expressivity. An element of this research is to explore, in informal terms, the cost-benefits associated with developing and using semantic approaches that are destined to be used in very practical scenarios. Both formal and informal mechanisms of expressivity therefore need to be examined.

The Feature Catalogue content model presented in Figure 4.17 is both a conceptual and a structural representation of the main elements that thus far make up the enhanced Catalogue design. However, in its current form the Catalogue model still lacks the formalisms that are required to unambiguously convey the meaning of, and relationships between, the modelled concepts. From a formal (ontological) perspective, what is missing is a set of formulas surrounding the concept definitions that should be always be set to “true” (using the rules that make up the formulas) and which in being “true” act to approximate the intended meaning of the concepts.

The next task, therefore, involved an exploration of what would be required to ontologically formalise the model presented in Figure 4.17. To ensure semantic interoperability across the multiple domains that may access and use the Catalogue, the casting of the model concepts into an ontological framework was performed using a top level (Upper or Foundation Ontology) as the terminological and theoretical end-point from which the concept definitions would be derived.

As explained in Chapter 2 top level ontologies provide general concepts which are theoretically used by all domains, and they can provide a common ontological foundation. Most domain-level ontologies will need to deal with “things” such as: objects, processes, properties, relations, space, time, roles, functions, categories and individuals. Top level ontologies define these parent classes of “thing”, relying heavily on collections of axioms (statements of truth) to distinguish between the various concepts. As has already been discussed there is no single standard top-level ontology. The reasons for the dichotomy of approach are often philosophical in nature. Milton (2004) provides a discussion of the philosophical foundations of realism vs naturalism vs Aristotelian common-sense realism that affect the choice of an ontological framework. In summary there is often disagreement about the basic categories of “thing” and the properties that they exhibit depending on your view of the world.

Given the proliferation of upper level ontologies a choice had to be made about which top-level ontology to use in which to anchor the enhanced Catalogue model. Mascardi *et al.* (2007) have provided an overview of seven of the main top-level ontologies, but unfortunately they did not indicate which of the reviewed ontologies best suited specific types of use-case. The Descriptive Ontology for Linguistic and Cognitive Engineering (DOLCE) upper-level ontology (Masolo *et al.*, 2003) was eventually selected.

In Chapter 3 it was explained that a group of thirteen experts, predominantly external to the Antarctic science community, was selected to assist in investigating RQ 1.2 (focussed on the issue of ontology selection and evaluation methodologies – which is un-packed in Chapters 6 and 7). Ancillary information (e.g., references to ontology development work and the literature) provided by this expert group indicated that DOLCE might be an appropriate choice on which to base a Feature Type Catalogue ontology.

Work by Brodaric and Probst (2009), for example, which was mentioned by some experts, had demonstrated the use of DOLCE as a bridging ontology in a geoscience exercise to integrate two different types of geoscience knowledge representation. Their rationale for using DOLCE was that "spatiality" is a key criterion used to distinguish DOLCE's most general categories, and they showed that DOLCE has the potential to represent aspects of scientific classification systems. Both of these issues are of concern in building a container for spatio-temporally attributed Feature Types. Additionally, Kuhn (2009) in his development of an ontology to represent observation and measurement, noted that 'observation' (as a process) holds together three (of the four) top-level branches of DOLCE in that observations are afforded by changes in the environment (stimuli) which involve DOLCE's 'endurants' and 'perdurants' and then an observation's result consists of symbols, which stand for DOLCE 'qualities' which inhere in these endurants and perdurants (see Figure 4.18 for a schematic of where these concepts sit in relation to one another). In other words, there is a very comfortable ontological fit in terms of the observation modelling paradigm adopted in this thesis and how that model can be expressed through DOLCE.

The 'spatial' conceptualisation within DOLCE is worth mentioning further because it was a key influence in the choice to use this Foundation ontology. Probst (2006, 2007) had already used DOLCE to perform a mapping of Cox's O&M schema and in doing so he explored the idea of semantic datums. Kuhn (2003) explains a semantic datum thus: "semantic datum implies that semantic spaces can be grounded in physically observable phenomena. Just like a geodetic datum grounds latitude and longitude in observable positions relative to the stars, soil types, vegetation classes, land use categories, and other qualities can be grounded in observable qualities, such as chemical composition or radiation". Kuhn (2003) also offers a simple example of a semantic datum for the quality of 'extent'. He observes that "extent serves as a basis for the semantics of all geographic extent or distance attributes, such as the length or width of a river or road segment, as well as any type of height or elevation measures. It also underlies rectangular coordinates, which are distances between coordinate axes and locations. The datum for extent simply consists of the extent quality space structured by the ratio scale, with meter or foot/yard/mile as physically grounded units".

In exploring the conceptualisation of semantic datums, Probst (2007) significantly advanced the ontology of DOLCE qualities, in particular DOLCE's spatial qualities, by relating the spatial qualities to the dimensions of their carriers (i.e., the things that the qualities inhere in). He formalised DOLCE's quality spaces and introduced reference spaces that are partitioned by symbols denoting measurement values. This particular work provided DOLCE with a foundation for how measurement scales can be semantically defined and this is of utility in defining Feature Types.

Finally, the choice to use DOLCE was also influenced by the availability of well-articulated ontological design patterns (Pisanelli *et al.,* 2003; Probst, 2006; Arora *et al.,* 2006; Presuti and Gangemi , 2008), which happened to be anchored in DOLCE, and which were used to cast the Feature Catalogue concepts into an ontological framework

The following commentary provides an overview of DOLCE and a number of ontology design patterns that were harnessed to explore a DOLCE-based Feature Catalogue content Model. A less formal approach is also suggested and later fully investigated which leverages a language called SKOS (Simple Knowledge Organisation System). SKOS lacks the formality required to support logical inference but it is now regularly used to exchange semantic information on vocabularies and requires relatively less (ontological engineering) skill to apply.

4.5.1 An Overview Of The Top-level Ontology - DOLCE

DOLCE was developed in 2001 as part of the WonderWeb project (Masolo *et al.,* 2003). It is marketed as a modular ontology of "particulars", i.e., an ontology of entities that cannot have instances. Entities that can have instances are labelled "universals". Universals are what objects have in common e.g., "a type, or a kind", "properties" and "relationships". "Particulars" are in fact instances of universals. In the philosophical underpinnings of "particulars" they can also have a 'type', 'properties' and a 'relation' but these things are said to "inhere" in a particular object (e.g., the "browness" of my dog) as opposed to the universal property of "brown" that is not related to anything specific. Another example is the "city" of "London", where "London" is considered a "particular" in that it cannot have an instance of itself, unlike the entity "city" (a "universal") which could have "London" as a particular instance.

In DOLCE, the difference between basic entities is established by ascribing to them incompatible essential properties. The main basic categories of DOLCE are shown in Figure 4.18.

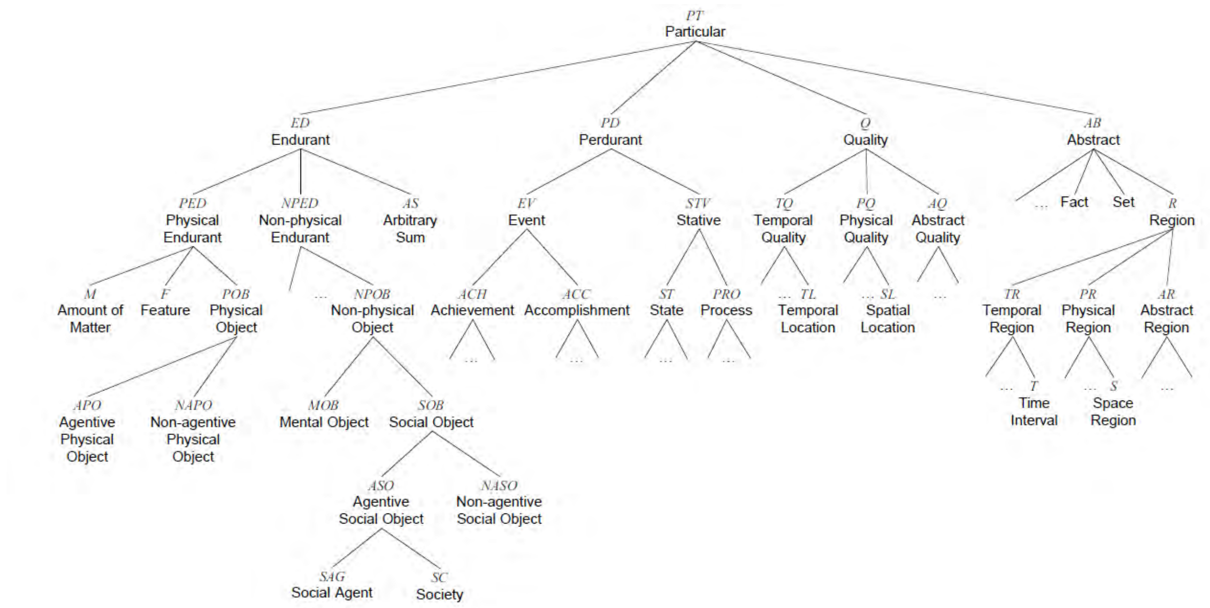

Figure 4.18 DOLCE Categories. Source (Fig 2 Masolo *et al.,* 2003)

The two main distinctions in DOLCE surround those things determined as "endurant" and those things considered "perdurant". In DOLCE terms, something is enduring if it is wholly present in time (i.e., all of their properties are present) at any time they are present. "Perdurants" by contrast happen in time (and space) and can have temporal parts. "Physical" and "non-physical" endurants are differentiated on the basis of whether they have direct spatial qualities. The next level of differentiation for perdurants is based on whether the concept represents a single temporal moment vs a cumulative collection of parts that can be temporally distinct, which are described as either "stative" or "events" respectively.

DOLCE "qualities" are the basic entities that we can perceive or measure. Qualities only exist as long as the entity in which they inhere-in exists. Qualities have types (e.g., colour, size, smell) and are characteristic for specific individuals. Note a quality is not synonymous with a "property" as DOLCE considers properties to be "universals". No two particulars can have the same quality because a quality is specific to an individual. Qualities (e.g., colour) have values or "quales" (e.g., a shade of green). Quales describe the position of an individual quality within a certain quality space. The example given by Masolo *et al.* (2003) to exemplify the usage of these concepts is that when people say that "two roses have exactly the same colour" they mean that their colour qualities, which are distinct, have the same position in the colour space, that is, they have the same colour "quale". In

DOLCE, space and time locations are considered as individual qualities like colours. Their corresponding qualia are called spatial (or temporal) regions.The final root category of DOLCE terms is “abstract” and this entity refers to any entity that does not have spatial or temporal qualities and which is not a quality. The only member of this hierarchy at present is “region”. A “region” is defined as an extent, a magnitude or a “value” of a quality.

Now familiar with the basics of DOLCE it is appropriate to look at two design patterns (which are implemented in DOLCE) that combined together enabled the expression of the UML-based enhanced Feature Type Catalogue Model in ontological terms.

4.5.2 Useful Ontological Design Patterns

The first design pattern that was explored in depth as a possible template for ontologically characterising the Feature Catalogue concepts was the mapping that Probst (2006) undertook to ground the Cox (O&M) model in DOLCE. As in this research, Probst (2006) first had to grapple with what current standards documents actually meant conceptually by the terms “feature” and “Feature Type”. He concluded that there were two main (ontological) senses in which a feature can be thought of. In the first sense the feature *is* actually any entity that exists in physical and social reality and the process of abstracting the features into categories results in the definition of Feature Types. In the second sense features are conceived of as information objects, where the information object *represents* a category, or an instance of a real-world entity. Probst (2006) used this second sense to develop his mapping and this is also the approach taken here. The ontology design will therefore provide the scaffolding to support the description of Feature Types (cast as instances), where these Types ‘represent’ a category.

In DOLCE an information object is part of an ‘Information Object’ (IO) Design Pattern (Arora *et al.,* 2006) which is itself part of the ‘extended description and situations’ (eDnS) design pattern (Pisanelli *et al.*, 2003). In the IO pattern the actual realisation of an entity is distinct from its conceptualisation. Any Web-based content can “realise” a socially constructed object, called an *Information Object* (see Figure 4.19).

In the pattern in Figure 4.19 an *entity* plays two distinct roles. One role refers to a *physical entity* (i.e., actual data or content) which realizes the *information object*. Entities playing this role are referred to as *information- realisations*. Alternatively, the role refers to something (which can be physical or virtual) which the information object is *about*. As an example, a “Bay” Feature Type in a Web-based Catalogue is an information object *about* a real-world feature called a Bay. The Feature Type “Bay” is also the *information-realisation* of the real-world entity “Bay”.

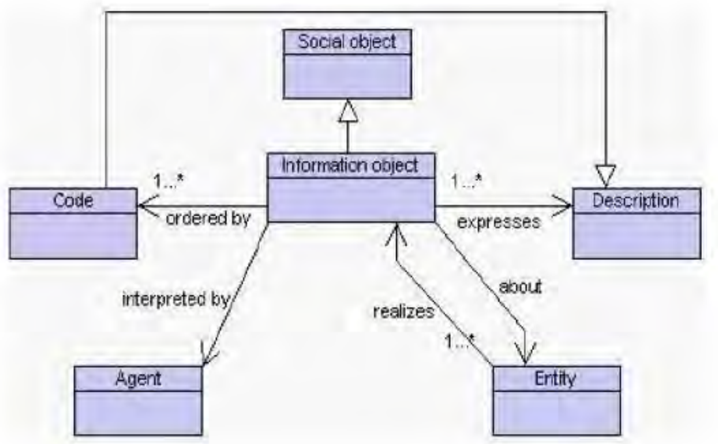

Figure 4.19 Information Object Pattern (Arora *et al.,* 2006)

The expressed *description* of an information object describes the *meaning/conceptualization* and it is feasible to have a semantic description which can be expressed in some particular *code*. This notion of an Information Object being about a real-world phenomenon sits comfortably with the definition of a Feature Type representing something in the real-world.

Presutti and Gangemi (2008) have taken the idea of resources as information objects on the Web further and suggested a complementary pattern which specifically addresses the issue of "identity" on the Web. Again, this is of interest since the idea of a Feature Catalogue, as understood in this thesis, is one of a Web-addressable semantic repository. Presutti and Gangemi (2008) call their pattern Identity and Information Resources (IRE). In this pattern they are concerned with the problem of explicitly distinguishing between the identity of a resource and its identifier (in other words they confront the semantic problems associated with "accessing" resources vs "referencing" them). Since this study's communities will be using services to access Feature Catalogue content (resources), this pattern can be of assistance.

Pressuti and Gangemi (2008) explain that there are clear distinctions which need to be made about Web resources from an identity perspective. They cite the case of the Web home page for the W3C organisation to explain the potential semantic dilemma. The W3C Web page URL (is an identifier); the (physical) place at which the page is located is another identity (but one based on place); the actual Web page (a document) has its own identity as a computational object (or file) and the subject of the Web page (W3C) has its own identity as an organisation. The Pressuti and Gangemi (IRE) pattern seeks to separate out these concerns so that semantically, the meanings of applied identity in each case, is clear. Their approach is to restrict the nature of Web resources to be '*computational objects*'. A *computational object* is a specialisation of an '*information-realisation*'. Any physical document, electronic service, file or application is considered to be a '*computational object'* (and it is '*addressable'*). They define '*Web resources'* as *computational objects* that can be placed in one or more '*abstract Web locations'*. A URI identifies one and only one *abstract Web location*. Because it is not possible to directly address a real-world entity (i.e., people can't be accessed via a URI) they envision a Web resource to be a '*proxy'* for some entity at any given time (t) about which you can make any number of assertions. The basic IRE model is depicted in Figure 4.20.

In Figure 4.20 an '*AbstractWebLocation'* is a point in the combinatorial regions identified by the URI metric. A '*Resource'* is a '*ComputationalObject'* that can be composed of other resources. A composed resource has a ***partOf*** relation with its components. It might have a location (i.e., '*AbstractWebLocation'*), the address of which is a '*URI'*. If the resource is a composed resource the identifier of its abstract location is also an approximate identifier for its parts. A '*WebResource'* is a resource that is made available on the Web, hence is accessible through a Web protocol (e.g., a document or a Web service). A '*ProxyResource'* is a Web resource which functions as a ***proxy for*** whatever entity (e.g., a personal home page, a set of metadata describing a person). Finally, a '*SemanticResource'* is a Web resource that realizes an information object through codification in a formal language for the Web (e.g., OWL) which functions as a ***proxy for*** whatever entity.

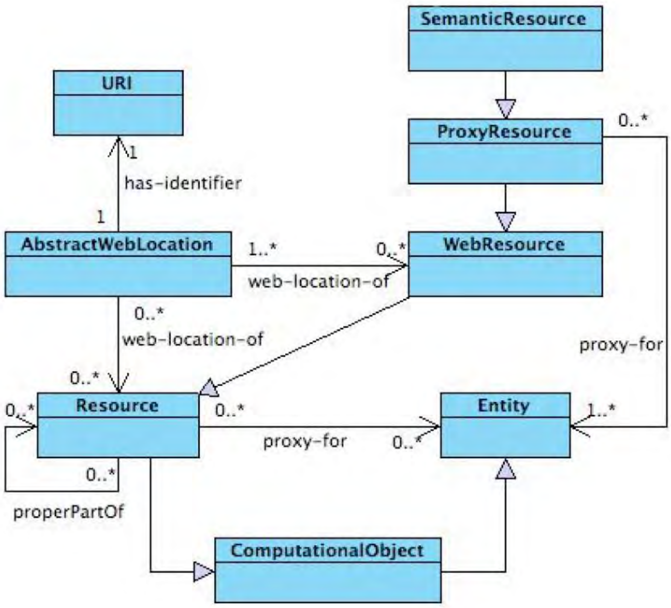

Figure 4.20 IRE Ontology Pattern (Source: Fig.1. Presutti and Gangemi, 2008)

By using a combination of IO (with classes and axioms drawn from the eDnS pattern) and IRE it was possible to design an enhanced Feature Catalogue Model, anchored in DOLCE that met community requirements.

4.5.3 Initial Design For DOLCE-based Feature Catalogue Content Model

A high level schema (in the form of a concept-map) casting the Feature Catalogue in DOLCE is provided in Figure 4.21. In this model both the *FeatureCatalogue* and *FeatureType* are considered *ProxyResources*. A *FeatureType* is a component (or "ConstituentOf") a *FeatureCatalogue*. Each *FeatureType* is accessible via a *URI*, as is the *FeatureCatalogue*.

A *FeatureType* is further defined by the relationships it has with other entities in the model. It can have two types of attribute: a *QualityAttribute* and a *NonQualityAttribute*. In essence the first attribute is a DOLCE *Quality*, but the second type of attribute is a DOLCE *Particular*. Two types of attributes have been defined to circumvent the problem that DOLCE *Qualities* are things which inhere in particular entities but attributes (as defined by the Generalised Feature Model) can also encompass characteristics such as names and other various labels for objects, which are not considered to be *inhering* characteristics (in DOLCE).

A *FeatureType* is a "proxyFor" a specific real-world (or socially constructed) *Particular* and it can have one or more *SpatioTemporalGeometries* (which themselves are the union of DOLCE temporal and spatial qualities). A *FeatureType* can have a synonym by virtue of its "synonymFor" relationship with another *FeatureType* and it can participate in a number of other types of relations that need to be predefined. A *FeatureType* also has a relationship with a *Theme* (which itself is related to a *Scheme* through an associative entity called a *FeatureSchemeContext*). A *Scheme* is a type of *ClassificationSystem* (which is a type of DOLCE *InformationEntity*).

A second associative entity included in the (ontological) conceptual model is *FeatureAttributeContext* which links *QualityAttribute*s with *Units-of-Measure*, *Datum, SamplingMethod and Instrument* (if all exist). These latter entities are used to either sense, or estimate the *value* (Quale) of, a DOLCE Quality (Attribute). The third associative entity - *ProfileRelationContext* links together the associative entity just mentioned (*FeatureAttributeContext),* a *Profile, a FeatureType, the FeatureType's SpatioTemporalGeometry* and a *Scheme* entity.

The various contextual associative entities play an important role in the DOLCE-based ontology model (they are shown in colour in Figure 4.21). Since ontologies are generally expressed through binary relations between concepts it has been necessary to use what are termed "n-ary" relationship constructs (Noy and Rector, 2006) to model the fact that some of the Feature Catalogue entities represent relations between multiple things. To resolve the situation where a property should relate to more than two individuals, these particular properties have been modelled as classes.

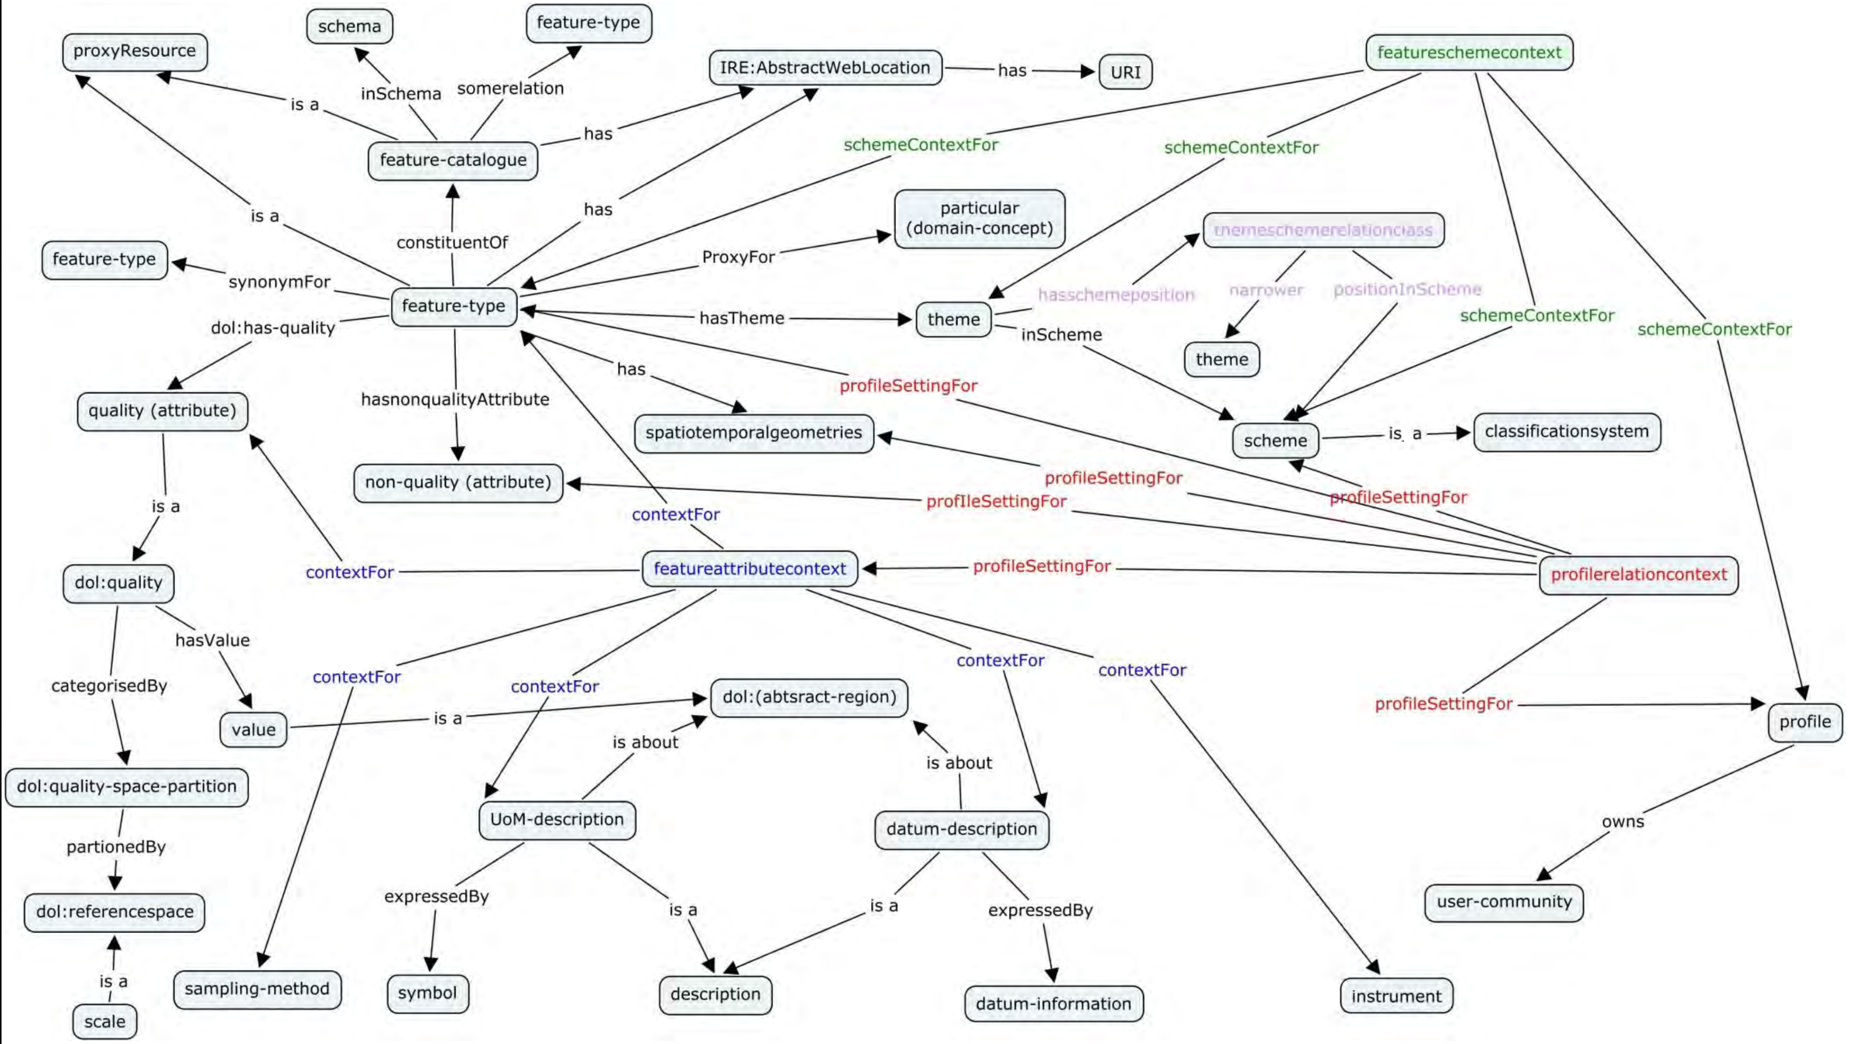

Figure 4.21 DOLCE-based Feature Catalogue Content Conceptual Mode

For example, consider the case where there is a need to model the relationship between a *FeatureType (*as defined by a specific *Profile)*, the *Theme* it is assigned to, and a *Scheme* in which that *Theme* and *FeatureType* reside. What is required is a property 'P' that can link all four entities (as shown below in Figure 4.22).

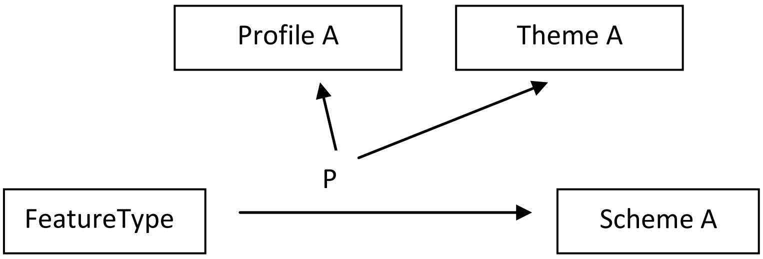

Figure 4.22 N-ARY Associations

The easiest solution is to model the relation 'P' as a class in its own right. In the case being discussed this is the class *FeatureSchemeContext.* In the types of "n-ary" relationship needed for the Catalogue content each of the individual entities generally plays a different role in the model without any single individual standing out as the subject or the "owner" of the relation. So each individual entity is simply linked by a generic relation to the new relation (modelled as a class) as shown below (in Figure 4.23). The relation is generic because it plays no role in the ontology other than to associate the various entities. The range of the generic property ('schemeContextFor') is the union of each of the individual entities involved.

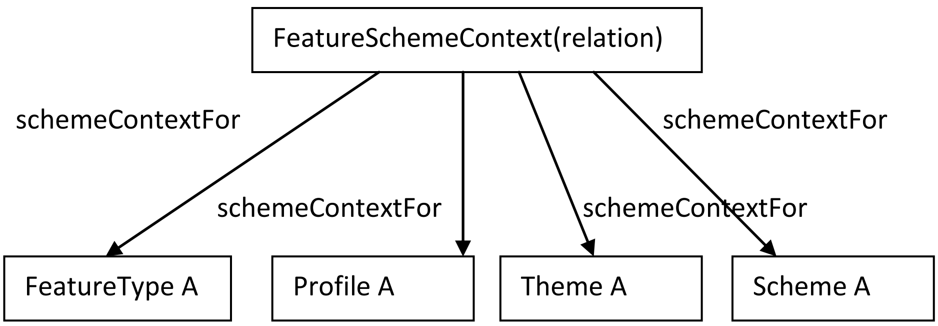

Figure 4.23 FeatureSchemeContext Association Example

Since *Themes* can belong to *Schemes* and potentially *Themes* may be nested in a hierarchy within a *Scheme* (e.g., a 'bedrock' theme may be a sub-theme of a 'geology' theme) a fourth associative entity was created (called *ThemeSchemeRelationClass*). This relation class links a *Theme* to a *Scheme* in which it appears and also to another *Theme* whose meaning it is "narrower" than. A *Theme* with no broader terms is considered to be at the first level of the *Scheme* hierarchy and a *Theme* with narrower terms can be placed appropriately, relative to other terms in the *Scheme* which it is narrower than.

This high level model just described was implemented in OWL (referred to in more detail in the next section 4.6) and its implementation is subsequently evaluated in Chapter 5. But before finalising this section on ontology design issues, the option of using the Simple Knowledge Organisation System (SKOS), either in tandem with the DOLCE-based formal approach, or on its own is briefly canvassed.

4.5.4 Less Formal Approaches – The Simple Knowledge Organisation System

SKOS is a conceptual model (see Figure 4.24) implemented using RDF that is designed to express the basic concepts and structure of taxonomies, thesauri, classification schemes and other vocabularies (Miles and Brickley, 2005). In the core of SKOS, conceptual resources are defined as concepts which can be identified with URIs, labelled with strings in one or more natural languages, documented with various types of note, semantically related to each other in informal hierarchies and association networks, and aggregated into concept schemes (Isaac and Summers, 2009). Of interest in this research is the fact that the SKOS vocabulary is easily extended to suit the requirements of any particular communities-of-interest and it is also theoretically possible to use SKOS in combination with OWL.

Whilst SKOS is not considered a language for modelling ontologies, the data model itself is in fact described as an OWL ontology. SKOS *Concept* and *Collection Scheme* are OWL Classes and particular SKOS concepts are considered instances of an OWL Class. The semantic relations and documentation properties in SKOS are "object properties" and the labelling properties are "data properties". It is therefore theoretically possible to use SKOS to annotate OWL classes. This would be highly beneficial in the case of the Feature Catalogue Ontology Model where additional "audit and maintenance" type information is required to embellish concepts, properties and individuals for the purposes of management and administration. Because this type of information (in general terms) has no real bearing on the formal semantics, or logical interpretations of the model components themselves, it would be sufficient that this "housekeeping" information be modelled as annotation.

Since a Feature Catalogue is essentially a semantic concept (or vocabulary) repository and server, the idea of using SKOS in its own right to codify the payload of a service request for Feature Type definitions was also attractive and therefore investigated as part of the Feature Type Catalogue build phase (next in section 4.6).

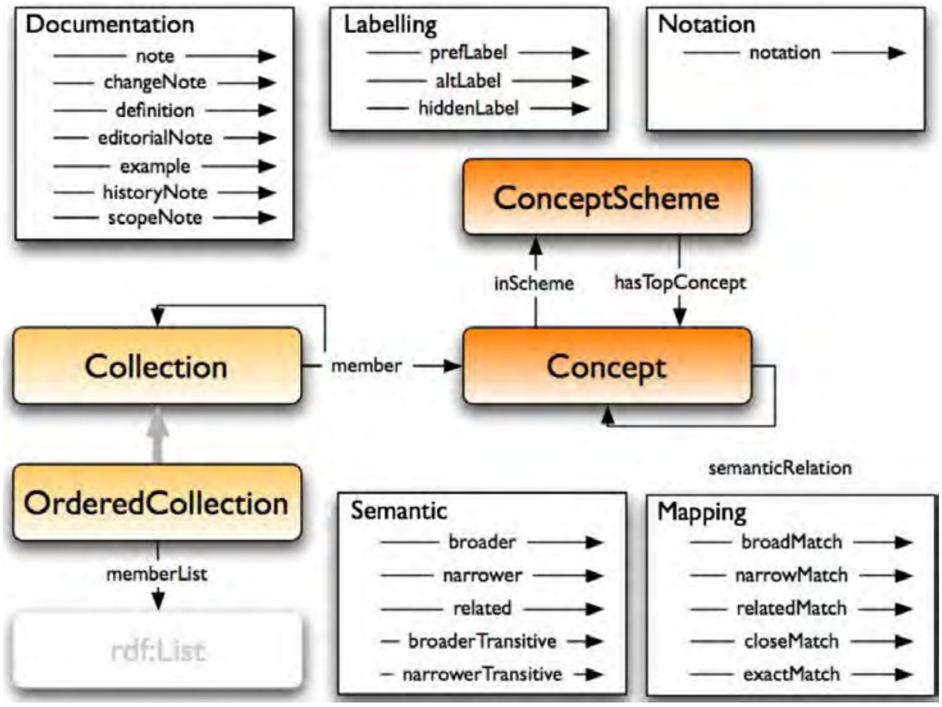

Figure 4.24 SKOS Model (source: Bechofer, 2010)

4.6 Feature Catalogue Ontological Repository Prototyping

Thus far the focus has been on defining the requirements for, and design of, a Feature Type Catalogue application. The results presented in this section, however, all relate to the tasks of building various parts of the Feature Catalogue application using proof-of-principle prototyping. The aim of prototyping various aspects of the design is to establish whether the concepts and ideas generated through the research to this point are sufficiently developed and robust to warrant going forward to a full application development (from the perspective of the community's participating in this study). Significant resources will be required to commit to the long-term hosting, development, maintenance and governance of a community-based Feature Catalogue (all of which are beyond the scope of activity described in this thesis). This research was therefore intended to provide the conceptual foundation and catalyst for a community's commitment to a "resourced" Feature Catalogue application development project. With this purpose in mind the design and prototyping activities should provide a sound framework for a subsequent operational implementation and substantive guidance on the advantages and disadvantages of certain implementation methodologies (including those related to the choice of ontology development language).

A fully implemented (semantically-enabled) Feature Catalogue repository will need, amongst other things, to leverage technologies tuned for storing and accessing content. This functionality could be secured in numerous ways using various combinations of custom-built and off-the-shelf software.

Clearly this study cannot trial all of the available combinations. However, some approaches to storage and access are considered sufficiently fundamentally different that a closer investigation of these empirically different methods is considered beneficial.

With respect to storage, the content management component could either be implemented as a semantic resource (i.e., using RDF or OWL to encapsulate the data in a triple store) or it could be implemented using the more traditional relational or object-oriented data storage paradigms. From a practical perspective, comparing what is involved in using a triple-store vs using a relational approach would be a useful exercise.

The first approach that was trialled harnessed the semantic storage and access capabilities of Oracle™ 11g coupled with TopBraidComposer™ (a proprietary ontology application development environment and a SPARQL testing tool). TopBraid was used to create a proprietary Oracle semantic Feature Catalogue data (triple) store and an OWL-based instantiation of the Feature Catalogue Content Model. Protege (an open source ontology development environment) was later used to finesse and test the OWL-based model after it was decided to switch midstream to a more modular form of the DOLCE upper ontology (i.e., DOLCE Ultralite).

In the second packaged approach Oracle 11g's native relational database structure was used to develop a relational Feature Catalogue Model that was coupled with custom application software (built in ColdFusion) which provided a REST-based service interface into the relational data store, with Catalogue content accessible in HTML, XML and SKOS. The results of this prototyping, which focussed primarily on testing Catalogue access methods, are reported in section 4.7.

Cognisant of all of the above, two different combinations of ontology language, storage and access approaches were packaged together and trialled in order to assess which one might be more practical in terms of their ease of deployment, operational use and ongoing administration. Evaluation of these packaged approaches was performed with the understanding that the lower the barrier appeared for implementation and maintenance, the greater the propensity would be for community uptake (in a full application development scenario).

4.6.1 Ontology Repository Store & Data Schema Creation Tools

Ontological models were stored as an OWL file and in Oracle™ 11g (using a semantic store and a relational schema). In addition to using its existing data dictionaries and tables, Oracle™ has added 'models', 'rule bases' and 'rule indexes (entailments, i.e., sets of triples derived via inference)' to create the triple store environment. Oracle™ 11g's semantic store treats data as triples (Subject, Predicate, Object) modelled as a graph structure. OWL and RDF rule bases are pre-loaded. SDO_RDF_TRIPLE_S is a new Oracle™ object type created to hold the RDF triples. Therefore each RDF triple: (Subject, Predicate, Object) is treated as one unique database object. Relational Oracle™ application tables contain a column of object type 'SDO_RDF_TRIPLE_S' to permit the loading and accessing of RDF triples and for storing ancillary values. RDF triple data is mapped onto an internal graph (and managed in a model) by storing subjects and objects as nodes, and properties as links. Whilst this is significantly different to how Oracle manages its relational data, a competent database practitioner can become quickly familiar with the syntax and utilities required for creating RDF schema and with populating them.

The task of creating an Oracle-based triple store for the Catalogue content was, however, greatly simplified by using TopBraidComposer™. TopBraidComposer™ connects to Oracle™ via a Jena API connection and allows a user to create schema, add content, run inference engines (Oracle's and others) and also execute SPARQL-like queries using a tool called SPARLQMotion. This environment was initially used to construct an OWL (1.1) Full ontology and to run SPARQL queries on the contained sample data (see Figure 4.25 – snapshot of TopBraidComposer™ development environment). Although TopBraidComposer™ worked well for building an Oracle-based repository, Protege (version 4) was ultimately used to refine the OWL Feature Catalogue ontology and to modify it so that it was an OWL Description Logic (DL) file.

4.6.2 Versions of DOLCE: DOLCE Lite-Plus vs DOLCE Ultra Lite

The initial ontology developed with TopBraidComposer™ imported a DOLCE 'Lite-Plus' (Gangemi, 2005) product as the foundation on which to build the Feature Catalogue ontology. DOLCE Lite (LOA, 2001) forms the back-bone of the "Lite-Plus"ontology but was extended (by the DOLCE authors) through a range of imports that are listed in Table 4.7. The "Lite-Plus" ontology, although classed as "Lite", was as a result particularly complex and voluminous.

Using DOLCE Lite-Plus meant incorporating all of the included imports which lead to a particularly bloated initial Feature Catalogue ontology. Most of the imported classes and relations were not required as anchors for the Feature Catalogue ontology, which initially had 19 locally defined classes and 24 properties. These imported entities were therefore extraneous to requirements, but had to be accommodated in the Feature Catalogue ontology never-the-less. The number of nested imported entities made the task of navigating around the ontology, even with the aid of the relatively user-friendly expansion tree interface provided by TopBraidComposer™, a long-winded process particularly if the tree structure had to be collapsed for any reason. The ontology was also relatively slow to load (several seconds on a Pentium 1.73Hz processor with 1Gb RAM). Even though this speed was considered well within acceptable bounds there was a concern that a fully populated ontology might be slow to operate on (although there were no direct tests performed to confirm this view).

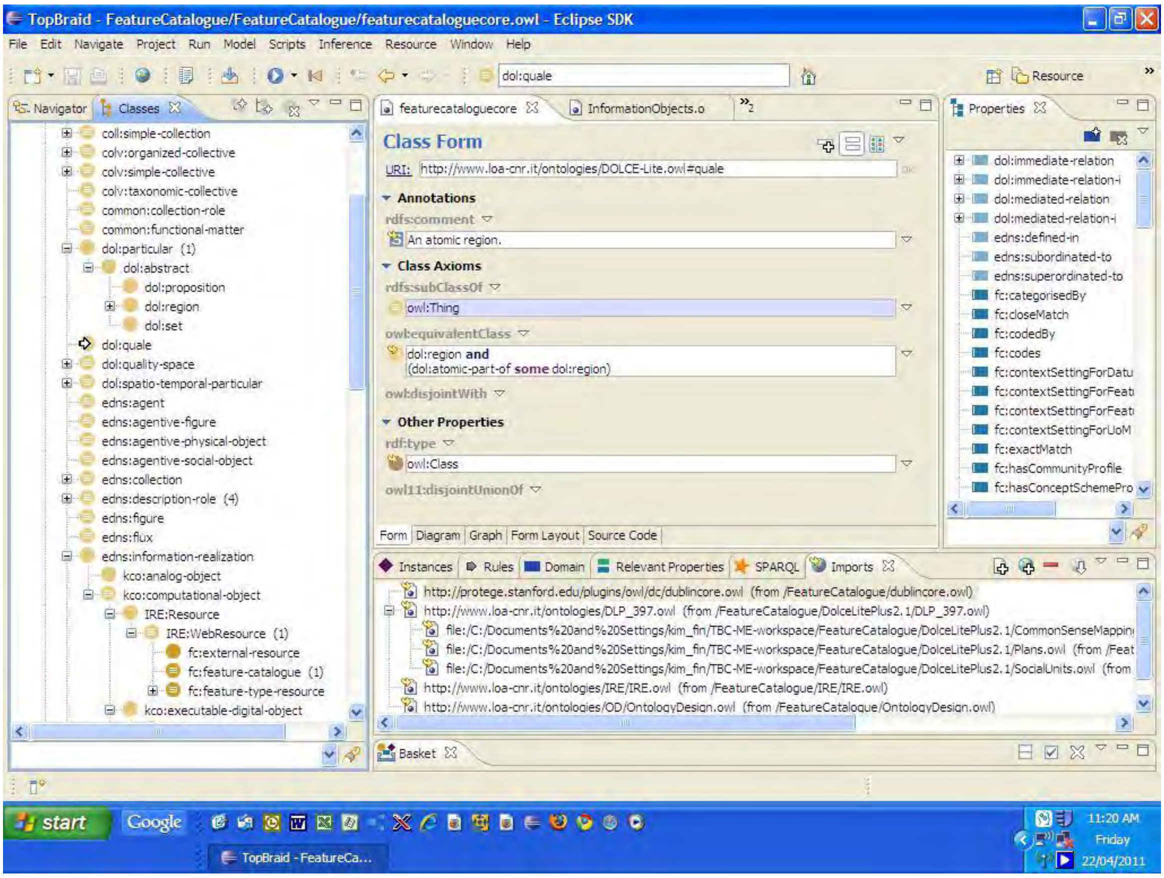

Figure 4.25 – Snapshot of TopBraidComposer Development Environment

Table 4.7 DOLCE Lite-Plus Ontology Imports

Imported Ontology URIs
http://www.loa-cnr.it/ontologies/TemporalRelations#
http://www.loa-cnr.it/ontologies/SpatialRelations#
http://www.loa-cnr.it/ontologies/DnS#
http://www.loa-cnr.it/ontologies/InformationObjects#
http://www.loa-cnr.it/ontologies/Actions#
http://www.loa-cnr.it/ontologies/SocialUnits#
http://www.loa-cnr.it/ontologies/Plans#
http://www.loa-cnr.it/ontologies/FunctionalParticipation#
http://www.loa-cnr.it/ontologies/FunctionalParticipation#
http://www.loa-cnr.it/ontologies/Collections#
http://www.loa-cnr.it/ontologies/Collectives#
http://www.loa-cnr.it/ontologies/CommonSenseMapping#
http://www.loa-cnr.it/ontologies/Systems#
http://www.loa-cnr.it/ontologies/SemioticCommunicationTheory#
http://www.loa-cnr.it/ontologies/Causality#
http://www.loa-cnr.it/ontologies/ModalDescriptions#

Fortunately, Gangemi (2007) had also released an "UltraLite" version of DOLCE which is now up to iteration 3.22. It uses friendly names and comments for classes and properties, has simple class restrictions and includes in one file the main parts of DOLCE, eDnS and other modules of Dolce Lite-Plus. DOLCE-Ultralite is an OWL-DL ontology. The architecture of the ontology is very different from DOLCE Lite-Plus in that it is pattern-based, which means that it is possible to use contained modules, called "content ontology design patterns" (NEON, 2011), which can be applied independently if only a few foundation entities are needed. Since the Feature Catalogue model needed modules from a range of patterns the entire UltraLite file was used. UltraLite's modularity and compactness of approach was a particularly attractive characteristic since the lack of modularity and the unnecessary complexity of DOLCE Lite-Plus made it difficult and at times confusing to use. It was confusing in the sense that there were so many classes and properties, many with lexically ambiguous definitions, which made it hard to discern when one type of class should be used over another.

The Feature Catalogue ontology design is based heavily, as was previously explained, on the concept of an Information Realisation (outlined earlier through the IRE Ontology design pattern). In DOLCE Lite-Plus an *InformationRealisation* class is first defined in the ontology "Information Objects" (IO), which is then imported into the "eDnS ontology" and refined by creating it as disjoint from other eDnS classes.

In DOLCE Lite-Plus this IO class is located at the same level as a DOLCE *Particular* (also known as an *Object*). In release 3.20 of DOLCE UltraLite, however, a class called *InformationEntity* was added as the superclass for an *InformationObject* (a child of a *SocialObject* class in DOLCE Lite-Plus) and an *InformationRealisation*. *InformationEntity* therefore represents the union of the *InformationObject* and *InformationRealisation* classes in DOLCE Ultralite. This particular difference is noted here because all of the classes just mentioned are used as parent classes in the Feature Catalogue Ontology and there is clearly a difference in the treatment of some of them between DOLCE Lite-Plus (that was used initially) and DOLCE UltraLite (which was used ultimately as the Foundation ontology).

4.6.3 Problems Using SKOS For Annotation

Although the original build process set out to develop a Description Logic (DL) flavoured version of an OWL ontology, progressive classification of the ontology first instantiated in TopBraidComposer™ revealed the flavour to be OWL Full. OWL DL had been preferred over OWL Full because, as previously explained, reasoners are unable to classify ontologies using OWL Full. In OWL DL the axioms permissible enable a reasoner to conclude that an individual either belongs to a class or it doesn't and there is sufficient structure and information (in OWL DL) for the reasoner to determine if a solution is "decidable".

During initial stages of development (using TopBraid™) the reason for the "Owl Full" ontology classification remained elusive. It subsequently became apparent, when SKOS was accidently deleted as an imported ontology, that including SKOS for annotation purposes was the cause of the OWL Full classification. Although SKOS is marketed as being used in combination with OWL as an annotation device (particularly for interoperability purposes, given the broad user-base of SKOS), some documented problems associated with using SKOS in this mode came to light after further investigation. In summary, the SKOS schema is considered to be OWL Full (Jupp *et al.,* 2008). SKOS "documentation" properties are OWL Object properties and SKOS "labelling" properties are OWL Datatype properties. Some SKOS documentation properties are used with data values or are used to state properties of classes within the SKOS vocabulary. Both of these situations violate the requirement that the namespaces of classes, properties and individuals are kept separated in OWL DL ontologies. SKOS (documentation) properties that were useful to this development exercise (e.g., "historyNote" and "changeNote") fell into this category and so their use resulted in an OWL Full ontology.

Although an OWL (2)-DL (W3C, 2011) compatible version of SKOS has recently (2011) been developed and is now in public use, the problem has been solved by removing the offending documentation properties from the SKOS OWL schema. Loading the SKOS OWL-DL version into Protege 4 now reveals all of the SKOS properties, except for those properties originally used in this study. So, in light of this situation, instead of using SKOS properties, an OWL version of Dublin Core (Stanford University, 2012b) was accessed and used instead for ontology annotation purposes.

4.6.4 Ontology Debugging and Tooling

Before moving on to explain in more detail how the final (OWL) ontology was developed, the issue of ontology debugging and the availability of the tools to support ontology development, will be mentioned briefly since these factors are applicable to issues of community uptake.

In the main, both Oracle™ and TopBraidComposer™ were robust tools as would be expected for relatively expensive proprietary offerings. Neither, however, produced particularly helpful error messages when ontology classification problems did arise, making problem resolution difficult. When creating ontologies most problems occur as a result of reasoning. The error messages logged and produced by an application during classification are generally just esoteric code-based messages (often expressed in Java e.g., "null pointer exceptions"). Oracle™ and TopBraidComposer™ were not alone, however, in offering poor de-bugging support. Protege 4 (or to be strictly accurate its reasoner plug-ins) was no better at offering debugging tools and bug tracking assistance. The paucity of available debugging facilities became very apparent when part way through the build process it was decided to re-factor the initial ontology using Protege 4 (instead of TopBraid Composer™). This was at the time of moving from DOLCE Lite-Plus to DOLCE UltraLite as the ontology foundation.

Protege is an excellent tool that comes for free. It has a wide range of plug-ins developed by what appears to be a solid developer base. The choice to use the most recent version of Protege (i.e., version 4), however, was perhaps not ideal since developers of the plug-ins had not yet caught up with the advances and code changes made in the latest version of Protege (designed to be compliant with OWL version 2). After commencing to use Protege it came to light that some of the existing, more useful plug-ins that were compatible with earlier versions of Protege which made it attractive to use, were not actively being ported to Protege 4. But since Protege 4 offered full support for OWL 2 (the latest version of OWL) it was the tool that was chosen.

In rebuilding the Feature Catalogue ontology in Protege, erroneously the ontology was created without iteratively classifying it along the way (unlike during the initial build). A number of significant changes had been made to the detail of the design during the rebuild and because of a failure to check regularly for logical problems, errors crept in. By the time the "hermiT reasoner" (KRR, 2012) which comes standard with Protege, was invoked to classify the ontology, the ontology was deemed inconsistent. An ontology is considered inconsistent if it contains errors in logic that prevent it from having a logical model. The only feedback from "hermiT" was that the ontology was inconsistent. An alternative reasoner "FaCT++" (Tsarkov and Horrocks, 2012) was downloaded and added into Protege 4. It also failed to classify the ontology and additionally produced errors unrelated to the ontology's inconsistency. The errors generated were related to the issue that the reasoner did not yet support all of the data types that are permitted in OWL2. This reasoner was duly abandoned and the Pellet (Clark and Parsia, 2011) reasoner was installed instead. Although Pellet was stable and effective, having trialled it on other consistent ontologies, it could not assist with locating the inconsistency errors. It did, however, have mediocre support for advising about unsatisfiable classes, a type of ontological inconsistency but one related to whether classes can have instances (as was discussed earlier in Chapter 2).

A Web search was commenced for ontology de-bugging tools. It transpired that there is highly limited support, from a tooling perspective for de-bugging ontologies (Stuckenschmidt, 2008; Wang *et al.,* 2005). The tool most often mentioned in support forums was SWOOP (MINDLab, 2012) but this tool had not been in active development since 2006 and it was unable to load the inconsistent ontology. The only solution available was to start the development process over again and incrementally rebuild the ontology, classifying the ontology at very frequent junctures using Pellet. Once a reported inconsistency occurred it was a case of working through the logic of what had been instantiated in the intervening period since the last classification and manually working out the logical inconsistency (without any reasoning support). These experiences left a strong impression that the technology base is still catching up with what users really need in order to be able to confidently and easily undertake ontology development.

4.6.5 FCATOWL (The Ontologically-grounded Enhanced Feature Catalogue Model)

The high level DOLCE-anchored Feature Catalogue design presented earlier in Figure 4.21 was augmented and modified as the build-process iteratively progressed. The initial design (in Figure 4.21) was conceptually very similar to what was eventually built in terms of modelling the *FeatureCatalogue* and *FeatureType* concepts as *WebResources* and harnessing n-ary relations to group related clusters of things.

However, a number of different DOLCE classes and relations had to be introduced to fully semantically express the final (FCATOWL) model. In some cases the 'class' and 'relation' choices made during the build process differed from that envisaged at design time because working with the DOLCE ontology gave better insight into what was actually meant by DOLCE concepts and properties than was possible to establish from reading notes on classes and axioms in published papers (as was done during the design process). Sometimes, choices made during the build differed because DOLCE UltraLite was much more transparent and navigable than DOLCE Lite-Plus and finding suitable anchors for Feature Catalogue concepts appeared that much simpler and easier. Ultimately, the fact that the ontology was actually built three times meant that there was plenty of opportunity to fine-tune the model with each iterative build. Because ontology development was also performed in tandem with interviewing ontology experts about how they selected and evaluated ontologies for re-use, it was also possible to use some of this information as practical advice and apply it in the development of the Feature Catalogue ontology.

In the next several pages the main components of the prototype Feature Catalogue ontology will be described in more detail than was supplied earlier in the concept map of Figure 4.21 to give a full appreciation for what has been delivered. All Figures were generated with Protege (Version 4, OWLViz or OntoGraf plugins). This commentary will also describe any assumptions that have been made in making certain class and property choices.

Ontology Class Hierarchy Overview

The final ontology was named "FCATOWL" and it imports two ontologies: "Dublin Core" for annotation purposes and "DOLCE UltraLite". Of the one hundred and seventeen classes in the ontology, forty-four are new classes that were defined locally to describe the Feature Catalogue repository (see class hierarchy diagram in Figure 4.26). The expressivity of the Description Logic used is described as SHIQ (D). This means an attributive language that supports: atomic negation; concept intersection; universal restrictions; limited existential qualification; role hierarchy; inverse properties; complex concept negation; qualified cardinality restrictions; use of data types; data values and data type properties. Throughout the remainder of the commentary all class names will be labelled beginning with upper-case letters and all property names will use camel-case (e.g., camelCaseWord). In the descriptions that follow, when referring to class names, the text will be *italicised*.

Description of Key Classes

The *FeatureType* class (whose parent is an IRE:*WebResource*) has relationships with many of the other classes defined in the ontology (see Figures 4.27 and 4.28). In some cases the relationship binds an individual of the *FeatureType* class with another individual of the same class. This happens for example when using the property "hasRelationshipObservationFOI" which relates a *FeatureType* individual to another *FeatureType* individual and implies that the subject *FeatureType* is the feature-of-interest (FOI) for an observation *FeatureType*.

A *FeatureType* is defined (in part) as a "proxyFor" the union of an *Object* and an *Event*. In Probst (2006) a *FeatureOfInterest* is posited as a *Role* played by either a DOLCE endurant (an *Object* in DOLCE UltraLite) or a perdurant (an UltraLite *Event*) in the course of an observation. Conceptually, a Feature Catalogue can be considered to contain features-of-interest to Catalogue users (some of which participate in observations and some of which don't). Although a *FeatureType* is a *WebResource* not a *Role*, Probst's restriction that *FeaturesOfInterest* are those things that are both an observation (a kind of *Event*), or a concrete *Object* in which qualities inhere also applies here to a *FeatureType*. There is a valid argument therefore to restrict the range of the *FeatureType* "proxyFor" relations to the union of *Objects* and *Events* and narrow its scope from that defined in IRE. The ontology therefore models the range of the "proxyFor" relation as an *Entity* (which is inclusive of DOLCE *Qualities* and *Regions* as in IRE) but specialises the relation with respect to its use with a *FeatureType* (i.e., locally restricts its range to *Objects* and *Events*). The (formal) semantic definition of a *FeatureType* entity lies in the ontological articulation of the *Object* or *Event* that is the range of the "proxyFor" relation.

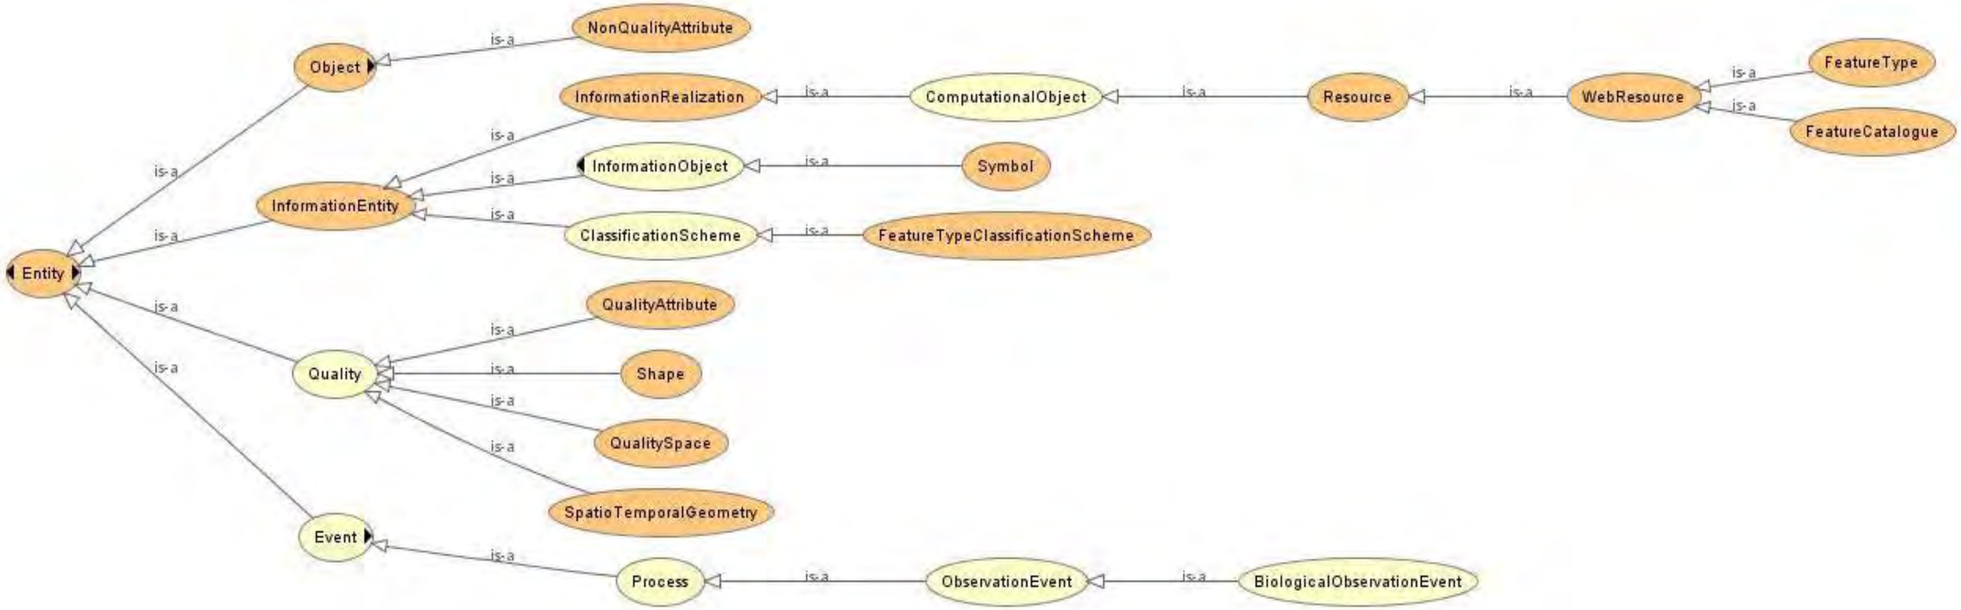

Figure 4.26 Feature Catalogue Model Class Overview
(Primitive classes are coloured yellow; defined classes are orange. Tool OWLViz plugin)

Figure 4.26 Feature Catalogue Model Class Overview (continued)
(Primitive classes are coloured yellow; defined classes are orange. Tool OWLViz plugin)

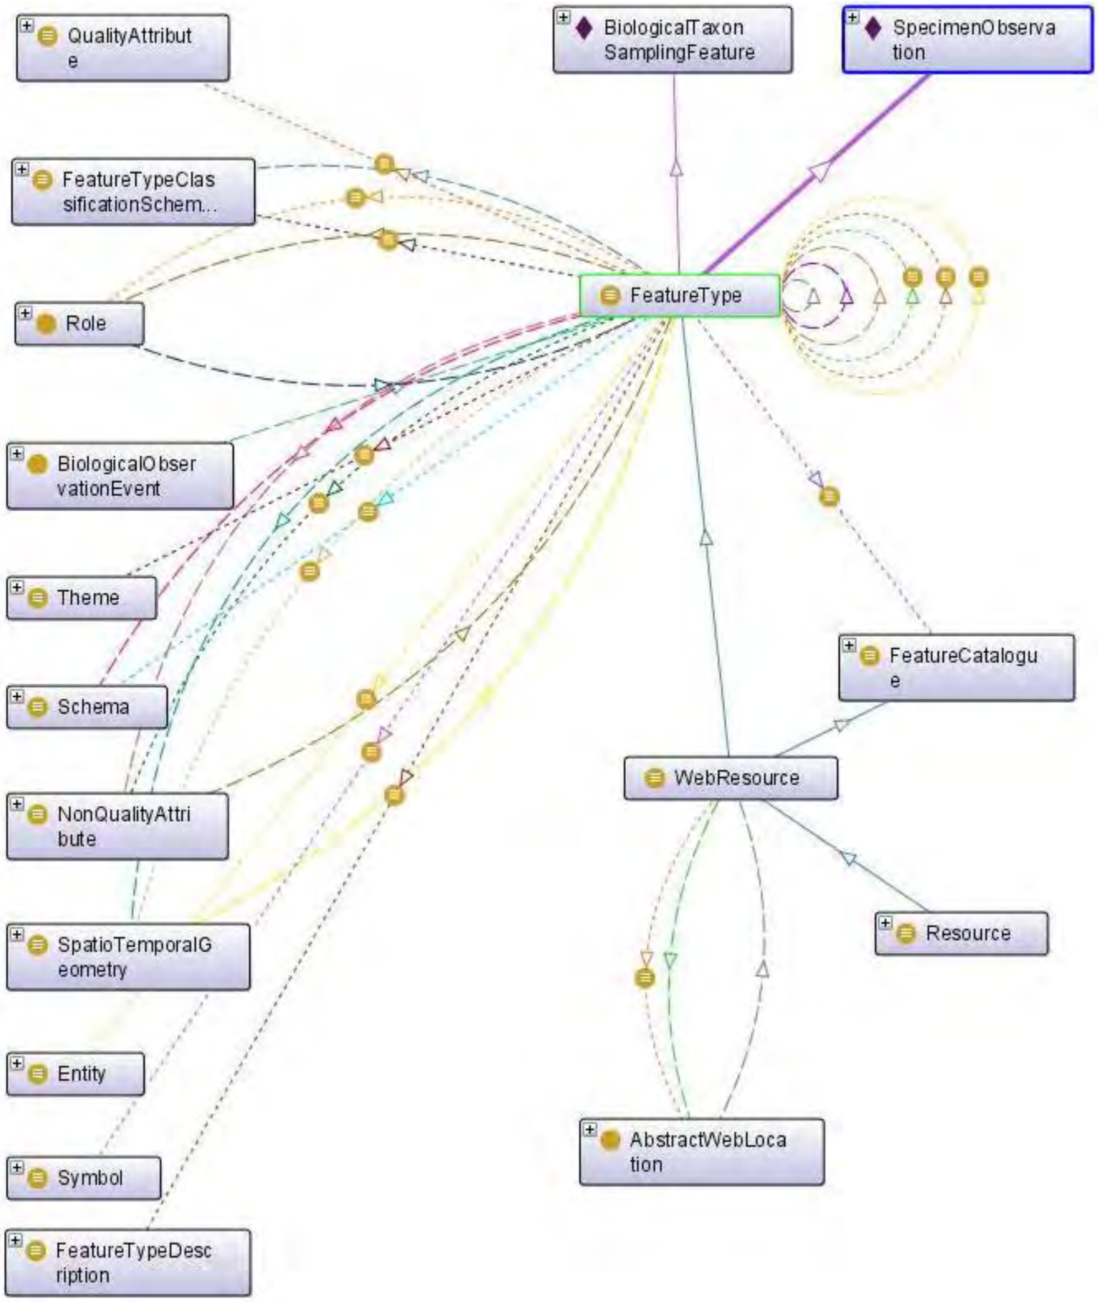

Figure 4.27 FeatureType Definition

(Generated with OntoGraf plugin. Classes prefaced by a crimson diamond are individuals; means defined class; means primitive class; dashed arcs represent different class relations; arcs with embedded arrows and defined class symbols e.g., '' can be interrogated in OntoGraf to show class-property-class association as displayed in the sample text bar below:

WebResource -- hasWebLocation(Equivalent class some) --> AbstractWebLocation **)**

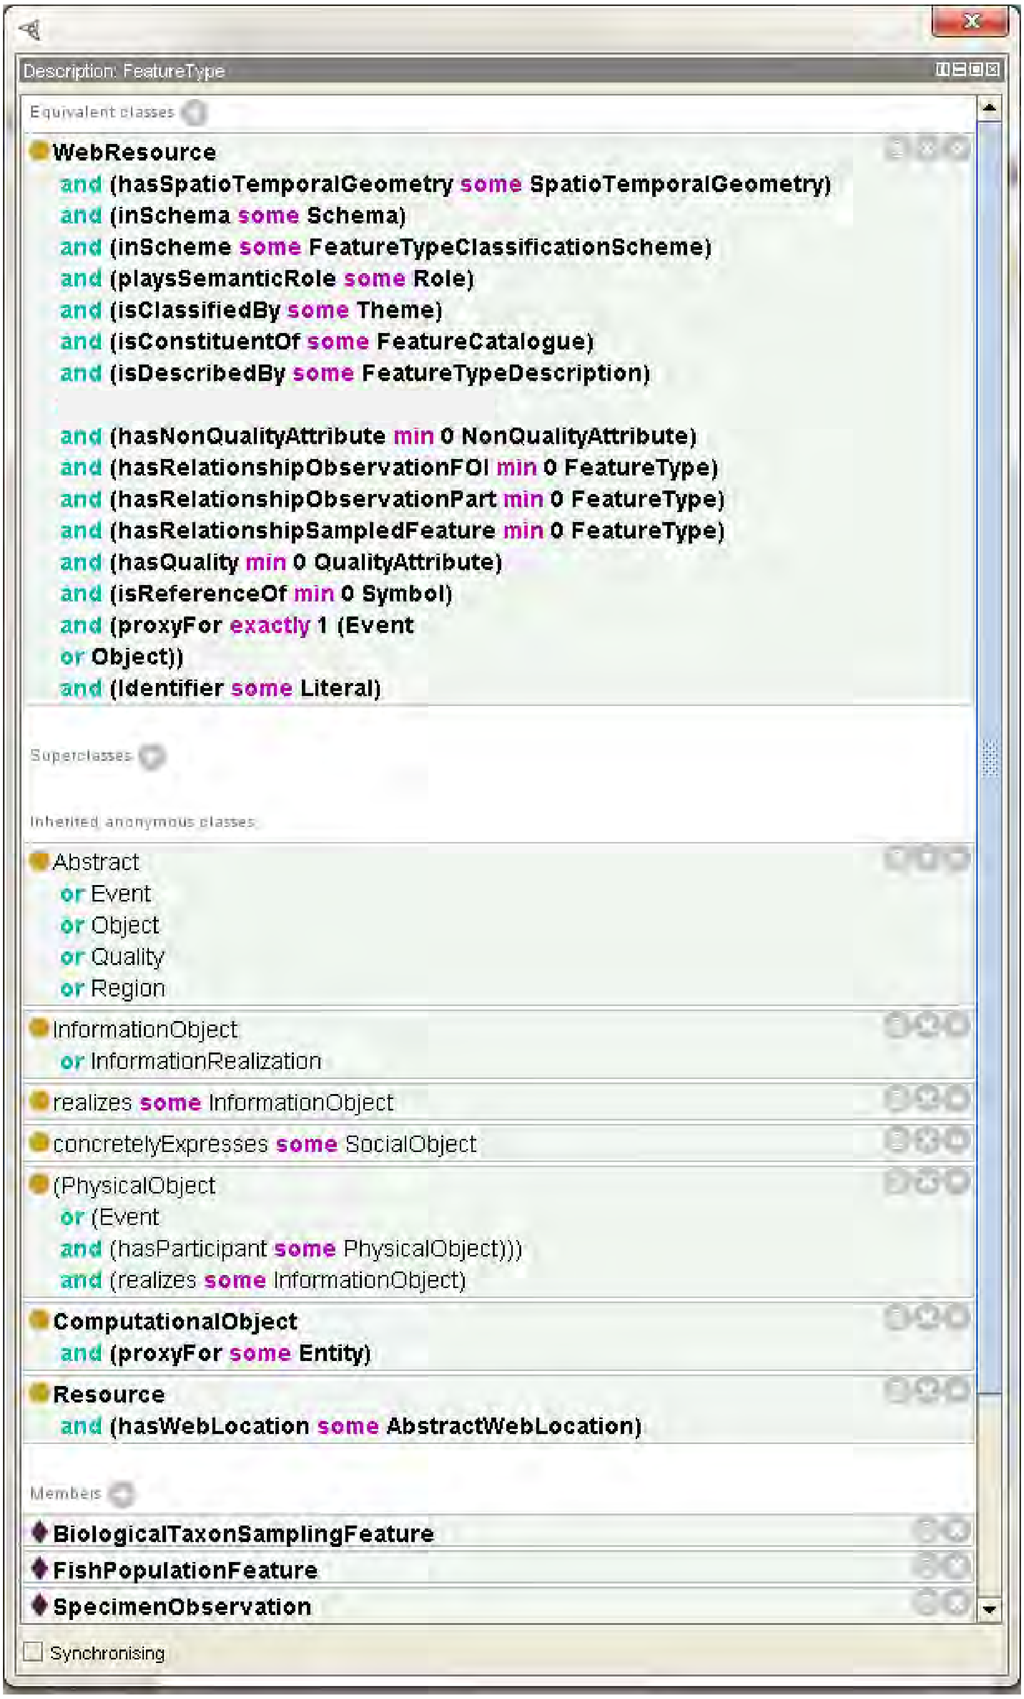

Figure 4.28 FeatureType Semantics (Note: entities prefaced by a crimson diamond are individuals. Generated with Protege)

The relationship between a *FeatureType* and its inhering *QualityAttribute* is a key facet of the Feature Catalogue model, since it is through the *QualityAttribute* that properties (in the Observation and Measurement sense) are expressed. From Figure 4.29 it can be seen that a *QualityAttribute* is a type of DOLCE *Quality* and it has a relationship with a *Value* (a kind of DOLCE *Quale or Region*). A *Value* is where the actual magnitude of an observation result would be recorded (e.g "10m" is made up of the magnitude 10 and a unit-of-measure in the form of the symbol "m" - standing for metres). Since this is a model which will be populated with "types" and not individual observations, this entity (*Value*) is placed into the model to enable the specification of other important classes that rely on the existence of the *Value* class. As such *Value* will never be populated with magnitudes but instead with placeholders such as those individuals shown in Figure 4.29 (i.e.LatitudeValue, WeightValue etc).

A *UnitOfMeasure* (specialised in this ontology as an *AttributeUoM*) is something that parameterises a *Value*. Note that a *UnitOfMeasure* does not have a direct relationship with a *QualityAttribute*, but instead the *QualityAttribute's Value*. According to Probst (2007) a *UnitOfMeasure* is part of a *Reference Region* (i.e., a DOLCE *Abstract Region).* In combination with an FCATOWL *Datum, AttributeUoM* semantically grounds the magnitude so that its meaning can be communicated unambiguously between measurement and observation users. In this ontology both the *AttributeUoM* and a *Datum* are conceived of as DOLCE *Parameters (*a *Parameter* in DOLCE being something that classifies a *Region*).

The relationship between the various semantic components of a *ReferenceRegion* and its associated *ReferenceSpace* and how these "map" onto a *Quality* and its components has been elaborately described in Probst (2007). To simplify, what is a highly complex classification and partitioning of differently named abstract regions, this ontology has taken a summarised approach.

In FCATOWL an *AttributeUoM* parameterises a *QualityAttribute*'s *Value.* A *Datum* (also a type of *Parameter*) is said to "anchor" an *AttributeUoM*. A *QualityAttribute* also has a *QualitySpace* (which is a specific location or quality space occupied by a particular *Quality*). Again, since this ontology is about individuals as "types" not about individual observations and their results, the *QualitySpace* is defined based on the taxonomy of *QualitySpace* types listed by Probst in Figure 5.1 of his PhD Thesis (2007), reproduced here as Figure 4.30.

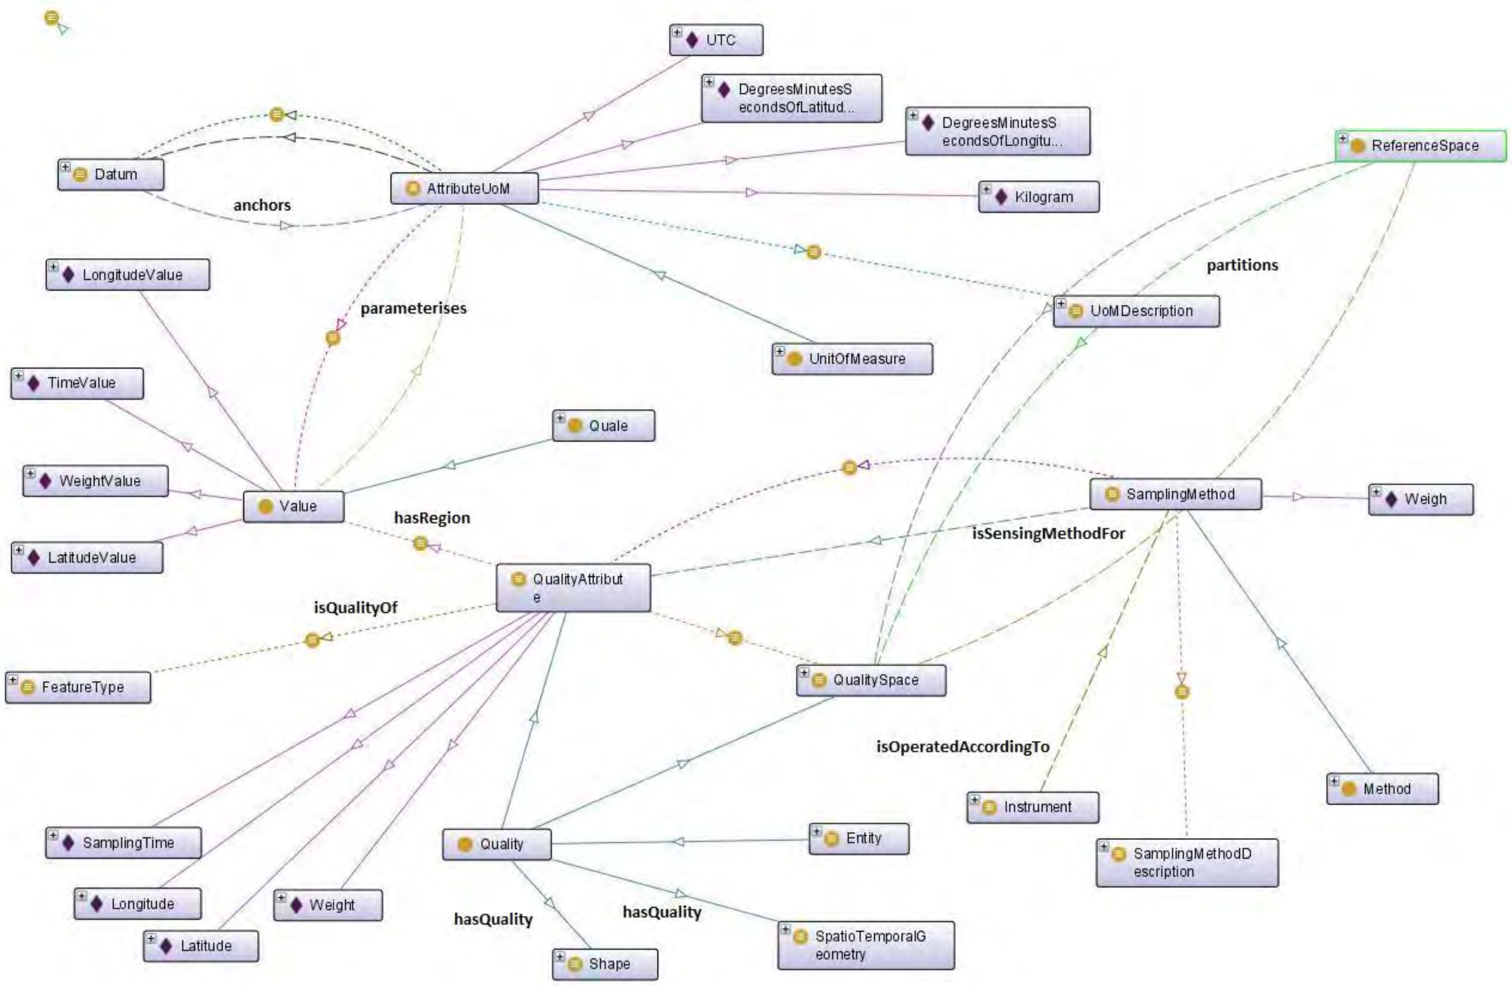

Figure 4.29 Associations involving a *FeatureType* and its *QualityAttribute*(Generated with OntoGraf plugin. Classes prefaced by a crimson diamond are individuals; means defined class; means primitive class; dashed arcs represent different class relations; arcs with embedded arrows and defined class symbols e.g., ‘ ’ can be interrogated in OntoGraf to show a class-property-class association (as described earlier for Figure 4.27. Properties (arcs) of interest have been labelled.

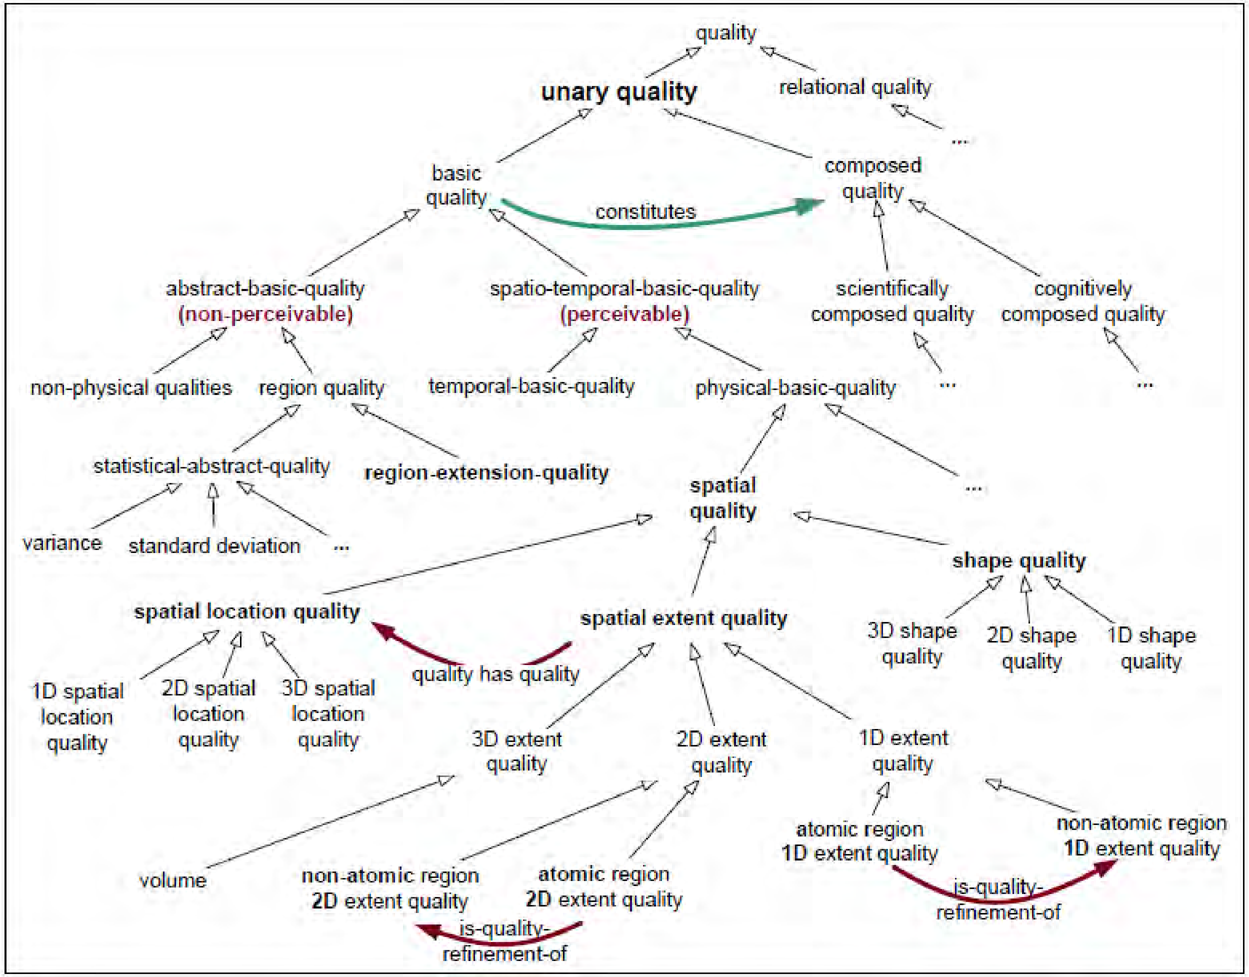

Figure 4.30 Taxonomy of QualitySpaces (from Figure 5.1, Probst, 2007) Which Could Form The Foundation For Well Known Spatio-temporal Types.

It is anticipated, however, that this taxonomy may need some embellishments in the future – particularly with respect to spatio-temporal regions if the SCAR and AODN communities decide to classify certain types of space-time combinations as "well-known" spatio-temporal shapes (as has been done in creating Feature Types in the Climate Science Modelling Language). Existing examples of elements of the Probst taxonomy include types such as "statistical abstract qualities", "basic temporal qualities" and various "spatial qualities". By understanding something about the "type" of the *Quality* (through its *QualitySpace*) it is feasible, in combination with other information, to deduce whether the *Values* of two different *Qualities* are able to be "integrated" for a specific use-case (i.e., are they qualities that are occupying a similar location in quality space).

A *ReferenceSpace* in some senses is analogous to the idea of a *QualitySpace* (except *ReferenceSpaces* are composed of *ReferenceRegions*). Examples of *ReferenceSpace* types are categories such as "ordinal", "interval", "basic" and "ratio" (Probst, 2007). Recall that these are the "scale types" that are of interest to biologists when they are deciding how to statistically treat and analyse data. A *ReferenceSpace* in FCATOWL is said to "partition" a *QualitySpace*.

The two remaining classes closely associated with a *FeatureType's QualityAttribute* are *SamplingMethod* and *Instrument*. A *SamplingMethod* is a child of a DOLCE *Method* and is characterised as "isSensingMethodFor" a *QualityAttribute*. The *SamplingMethod* has a relationship with an *Instrument* through a property called "isOperatedAccordingTo". An *Instrument* is of type DOLCE *PhysicalArtifact*. The *SamplingMethod* "isDescribedBy" a *SamplingMethodDescription* which can include any type of information pertaining to the sensing methods and *Instrument* operations used in estimating the *Quality's Value*.

All of the aforementioned classes directly associated with a *Quality* and its *Value* (except the *FeatureType*) are then grouped together through the *FeatureAttributeContext* (n-ary) class – which is a type of DOLCE *Situation*. This convenience class is necessary to ensure that a *Quality* can be linked with one or more *SamplingMethods*, or *Units-Of-Measure*, depending on a community's needs, and that this grouping can easily be recalled and re-instated. This would not be possible if the only option was to rely on the individual binary relationships between a *Quality* and its associated classes.

This *FeatureAttributeContext* class is itself part of another n-ary class called *ProfileRelationContext* (see Figure 4.31). It is through *ProfileRelationContext* that a community's preferences for individual *FeatureType* configurations are recorded. Each instance of *ProfileRelationContext* groups a *FeatureType*; a feature's *SpatioTemporalGeometry*; one or more *FeatureAttributeContexts*, one or more *NonQualityAttributes*; a community *Profile* and the *FeatureTypeClassificationScheme(s*) in which the *FeatureType* sits. Each instance of *ProfileRelationContext* is determined by a unique combination of *FeatureType*, *Profile* and *SpatioTemporalGeometry*. For example, one individual of *ProfileRelationContext* would be created for a community that wishes to add a *FeatureType* called *BiologicalTaxonSamplingFeature,* which has only one type of *SpatioTemporalGeometry* but which has several types of *NonQualityAttribute* and three *QualityAttributes* (each of which is part of a *FeatureAttributeContext* instance).

If, however, *BiologicalTaxonSamplingFeature* had two types of *SpatiotemporalGeometry* (e.g., it could be represented by a point or by a polygon) two instances of a *ProfileRelationContext* would be required to describe the various facets of the *FeatureType* and its associations.

A key class that is part of the context-setting for *ProfileRelationContext* is *Profile*. A *Profile* is "ownedBy" a *CatalogueUserCommunity*, which is a type of DOLCE *Community*. A *CatalogueUserCommunity* also plays the *Role* of *UserRole* and "isDescribedBy" a *UserDescription*.

N-ary classes do not generally have meaningful names but in populating the FCATOWL ontology with sample data some naming conventions for the various classes established as children of DOLCE *Situation* were trialled. As can be seen in Figure 4.32 two instances of *ProfileRelationContext* have been created: *PRC:P001:F001:G001* and *PRC:P001:F004:G001*. In these names "PRC" stands for *ProfileRelationContext*, "P001" identifies a *Profile* by its numeric ID, "F001" is the ID of a *FeatureType* and "G001" is also the numeric ID assigned to an individual of *SpatioTemporalGeometry*. In naming the *FeatureAttributeContexts* lexical naming was trialled. Although not used here, in both cases 'rdfs:label' could be added to the ontology for human readable labelling.

An example instance of *FeatureAttributeContext* that was created was named "*FAC:SamplingTime:UTC"*. In this naming convention "FAC" represents *FeatureAttributeContext*, "SamplingTime" is the *Quality* and "UTC" is the *AttributeUoM*. The various naming conventions were adopted to try to convey some meaningful information to a human user. The names are of course meaningless in the context of machine interpretation.

There are a range of other classes (e.g., *Theme*, *Scheme*, *FeatureSchemeContext*) which have been established so that a *FeatureType* can be appropriately classified by a community user group. The class *Schema* has also been included so that it is possible for a community to nominate an example digital schema in which the *FeatureType* is used, for reference purposes.

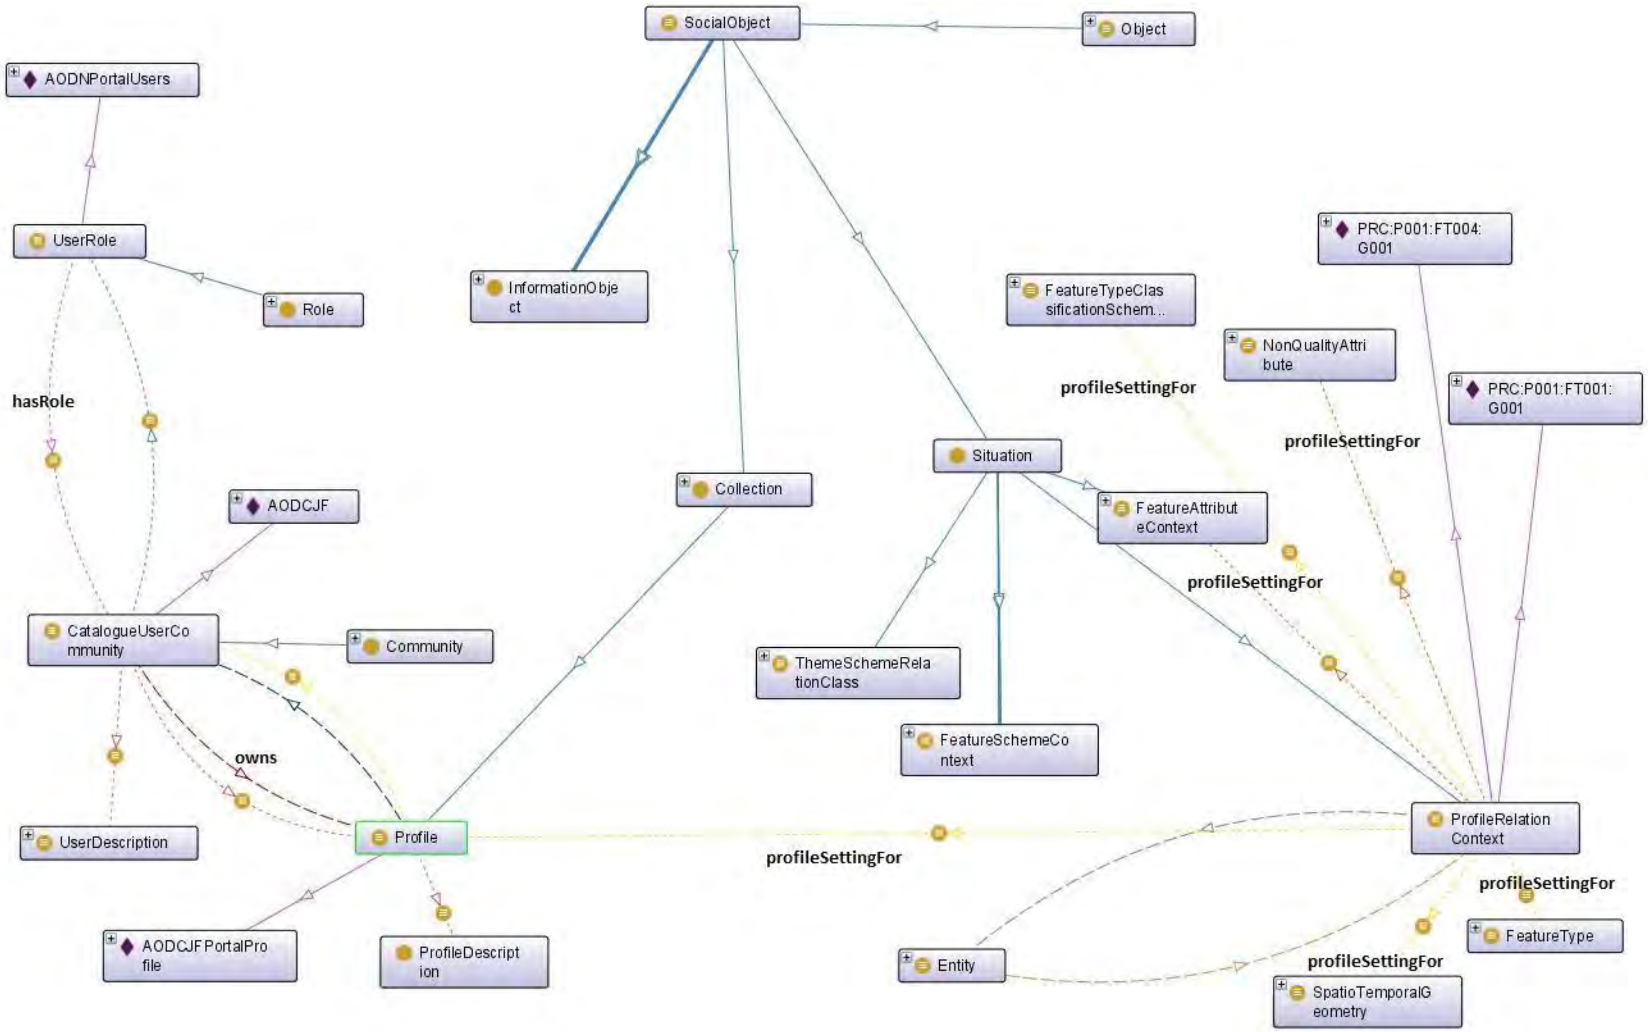

Figure 4.31 A community Profile and ProfileRelationContext and associated classes and relations.

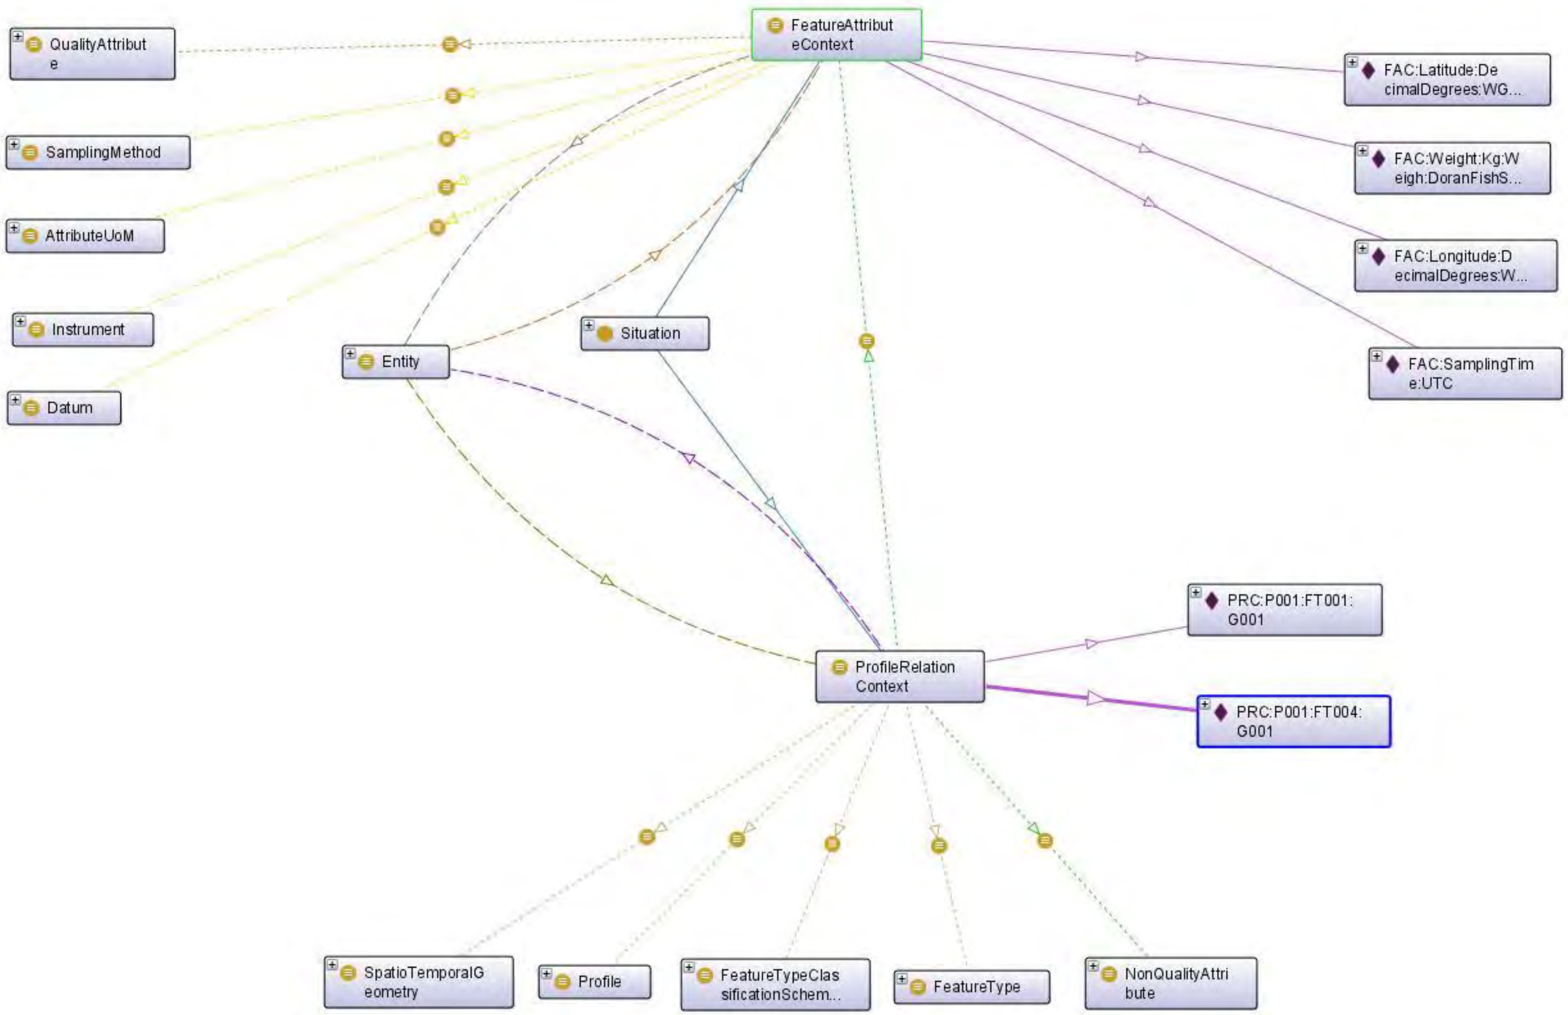

Figure 4.32 Naming Of N-Ary Contexts.

In FCATOWL, *Schema* is a type of DOLCE *Configuration*. The last class worthy of special mention is *SpatioTemporalGeometry*. This class is a type of DOLCE *Quality* and therefore also has an association with a *QualitySpace* and its *QualitySpace* is also “partitioned” by a *ReferenceRegion*. Individuals of *SpatioTemporalGeometry* will be drawn from a taxonomy of well-known spatiotemporal types (which as described earlier will need to be developed and agreed upon, and are ontologically a union of spatial and temporal quality spaces). The *SpatioTemporalGeometry* class also has a relationship with another *Quality* called *Shape*. Individuals of *Shape* will also be drawn from “well-known” taxonomies but from spatial geometry types only (i.e., spatial classifications excluding time). The full OWL ontology can be found as Appendix 10 (inclusive of a Manchester OWL Syntax document). Figure 4.33 is the final ontology schema diagram (in concept map format).

Properties

Both global and local property restrictions were used in FCATOWL. For example, in addition to (global) domain and range restrictions, Functional properties were employed for relating *FeatureTypes*, which together represent an observation and its feature-of-interest (as described below). Since a given individual can have only one value for a Functional property, it is then possible to infer that two *FeatureTypes* are the same, if they are both values for a Functional property associated with the same subject (*FeatureType*). Inverse properties were also employed (e.g., “hasProfileSetting”/”profileSettingFor”) so that it was possible to identify an individual *ProfileRelationContext* by knowing the values of its associated members and vice versa.

Local property restrictions were used to define classes and to introduce cardinality restrictions (e.g., the “proxFor” property required exactly one value from individuals in the class *Event* or *Object)*.

Ontology Population & Testing

In order to test the ontology, three *FeatureType* individuals were created: *SpecimenObservation*, *BiologicalTaxonSamplingFeature and FishPopulationFeature.* All were drawn from the previous desk-top mapping exercise in section 4.4.2 (and Appendix 9). To support semantic associations that exist between these three individual *FeatureTypes* (which together represent a complex observation), three additional properties were created: “hasRelationshipFOI”, “hasRelationshipSampledFeature” “hasRelationshipObservationPart”. The domain and range in all cases is *FeatureType* (Figure 4.34).

The first two properties are Functional and by convention the *FeatureType* playing the role of an *Observation* (in a triplet observation complex) should be placed in the domain position when using both (Functional) properties. In the third type of property (“hasRelationshipObservationPart”) the *FeatureTypes* playing the roles of *ObservationFeatureOfInteres*t and *ObservationSampledFeature*,

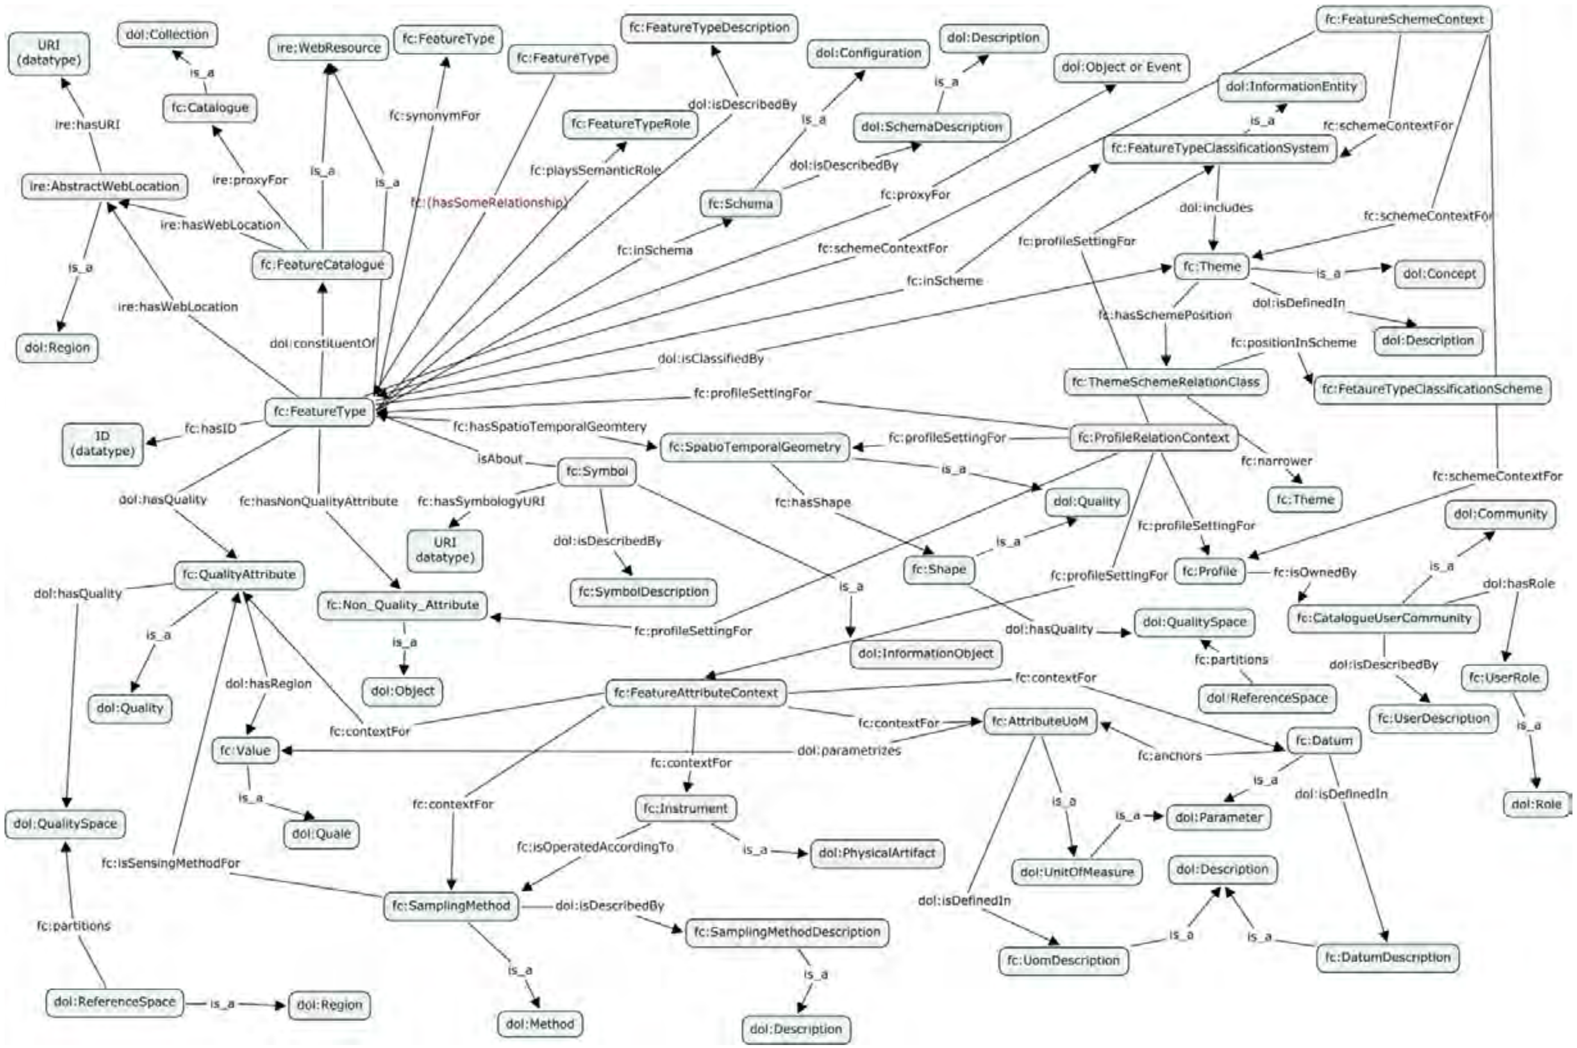

Figure 4.33 Final FCAT Ontology Schema. Note: fc:(hasSomeRelationship) is a placeholder for 'as yet to be defined' relations that might exist between two Feature Types and which may be added to FCATOWL in the future.

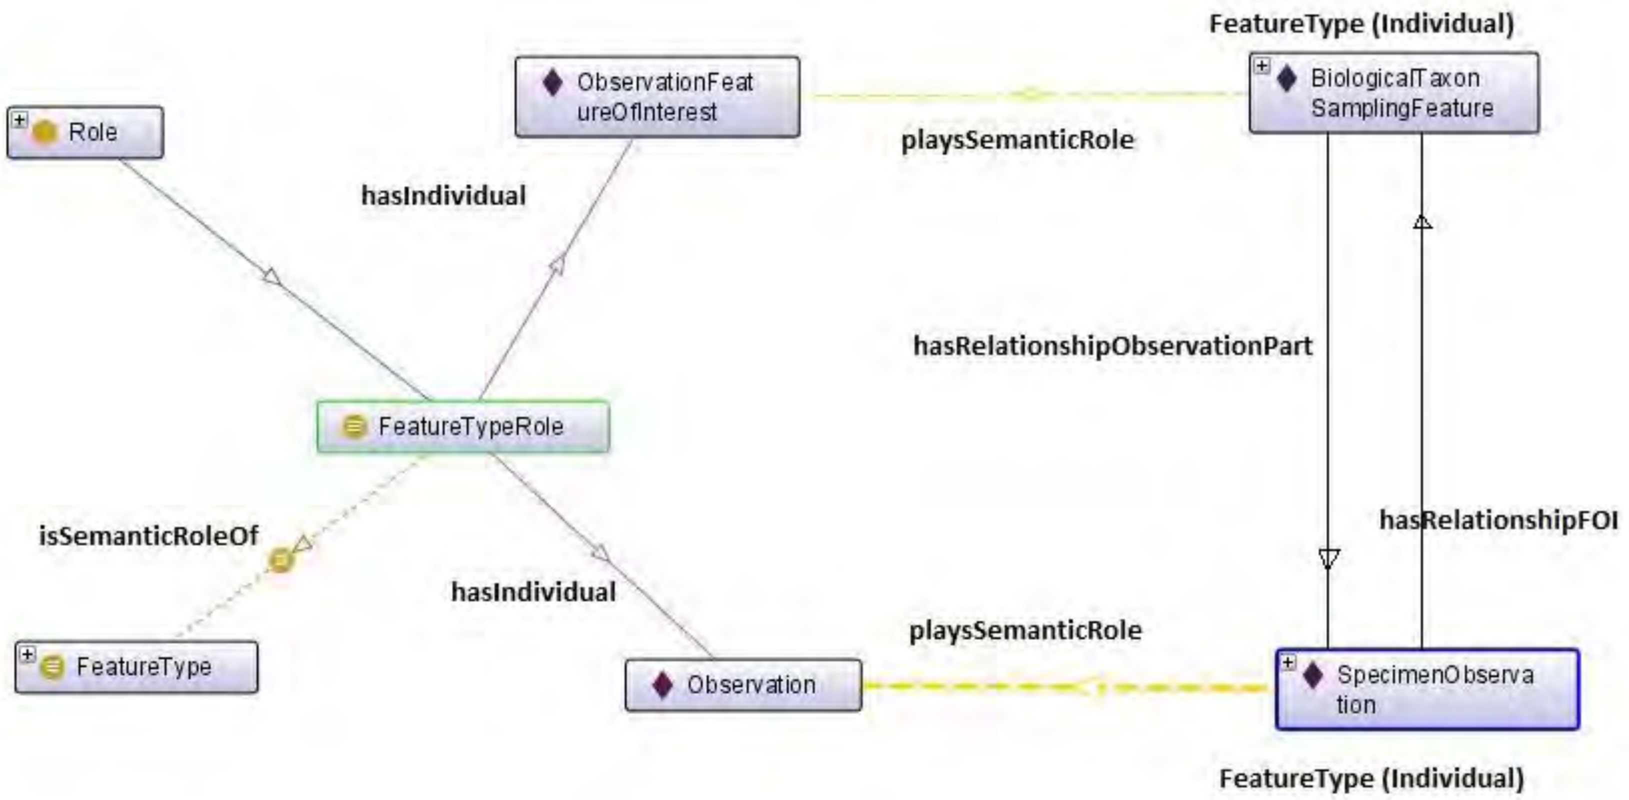

Figure 4.34 Semantic Relations Between (Example) Feature Types

respectively should be in the domain position, whilst the *FeatureType* playing the role of *Observation* is in the range position.

Semantically assigning the role that each *FeatureType* plays in the relationship is achieved by virtue of a *FeatureType's* relationship along another property called "playsSemanticRole", whose domain is a *FeatureType* and whose range is a DOLCE *Role*. Since earlier (in section 4.4.2 and Appendix 9) it was determined that a *BiologicalTaxonSamplingFeature* was a conceptual feature-of-interest for the observation called *SpecimenObservation* and that *BiologicalTaxonSamplingFeature* had a partOf type relationship with *SpecimenObservation*, these new properties (i.e., "hasRelationshipFOI" and "hasRelationshipObservationPart") have been used to create the appropriate semantic links. In a similar fashion the *FishPopulationFeature* is conceptually a SampledFeature within the *SpecimenObservation* and can be related to *SpecimenObservation* via the "hasRelationshipSampledFeature" property. Conversely a *SpecimenObservation* "hasRelationshipObservationPart" with *FishPopulationFeature (i.e., FishPopulationFeature is a* "partOf" *SpecimenObservation).*

As can be seen in Figure 4.35 the ontological model supports the required descriptions for the *SpecimenObservation* and *BiologicalTaxonSamplingFeature* entities. Mock *Themes*, *Schemes* and appropriate n-ary associations were created to which the *FeatureTypes* were associated. Figure 4.36 shows three individuals of the class *Theme* (i.e., *BiologicalObservations*, *TaxaObservations* and *ChemicalObservations*). Each *Theme* has a theme description (which in FCAT is deliberately brief). From this diagram it can also be seen that the *Theme: "TaxaObservations"* (which happens to have a numeric ID of "456") was used in defining both *SpecimenObservation* and *BiologicalTaxonSamplingFeature* and it appears in two *FeatureSchemeContexts* (*FSC:FT001:S001:P001:T456* and *FSC:FT004:S001:P001:T456*). Recall that a *FeatureSchemeContext* groups a *FeatureType* with a *FeatureTypeClassificationScheme*, a *Profile* and the *Theme* to which the *FeatureType* has been assigned by the community *Profile* owners. Figure 4.36 also depicts how the *TaxaObservations Theme* is part of a *ThemeSchemeRelationClass* (*Theme456:Scheme001*), which through its "narrower" association with *BiologicalObservations* identifies that *TaxaObservations* is below *BiologicalObservations* in a *Theme* hierarchy.

It is worth highlighting that in creating the *SpecimenObservation* it was linked to its real-world entity called *SpecimenObservationEvent* (which is sub-classed indirectly from a DOLCE *Event*), for which it is a proxy. A *SpecimenObservationEvent* is a direct child of an FCATOWL:*BiologicalObservationEvent* that has, as a "participant", a DOLCE *BiologicalObject* (see Figure 4.37). The entity that is the range of a feature type's "proxyFor" relation defines the semantic type of that Feature Type.

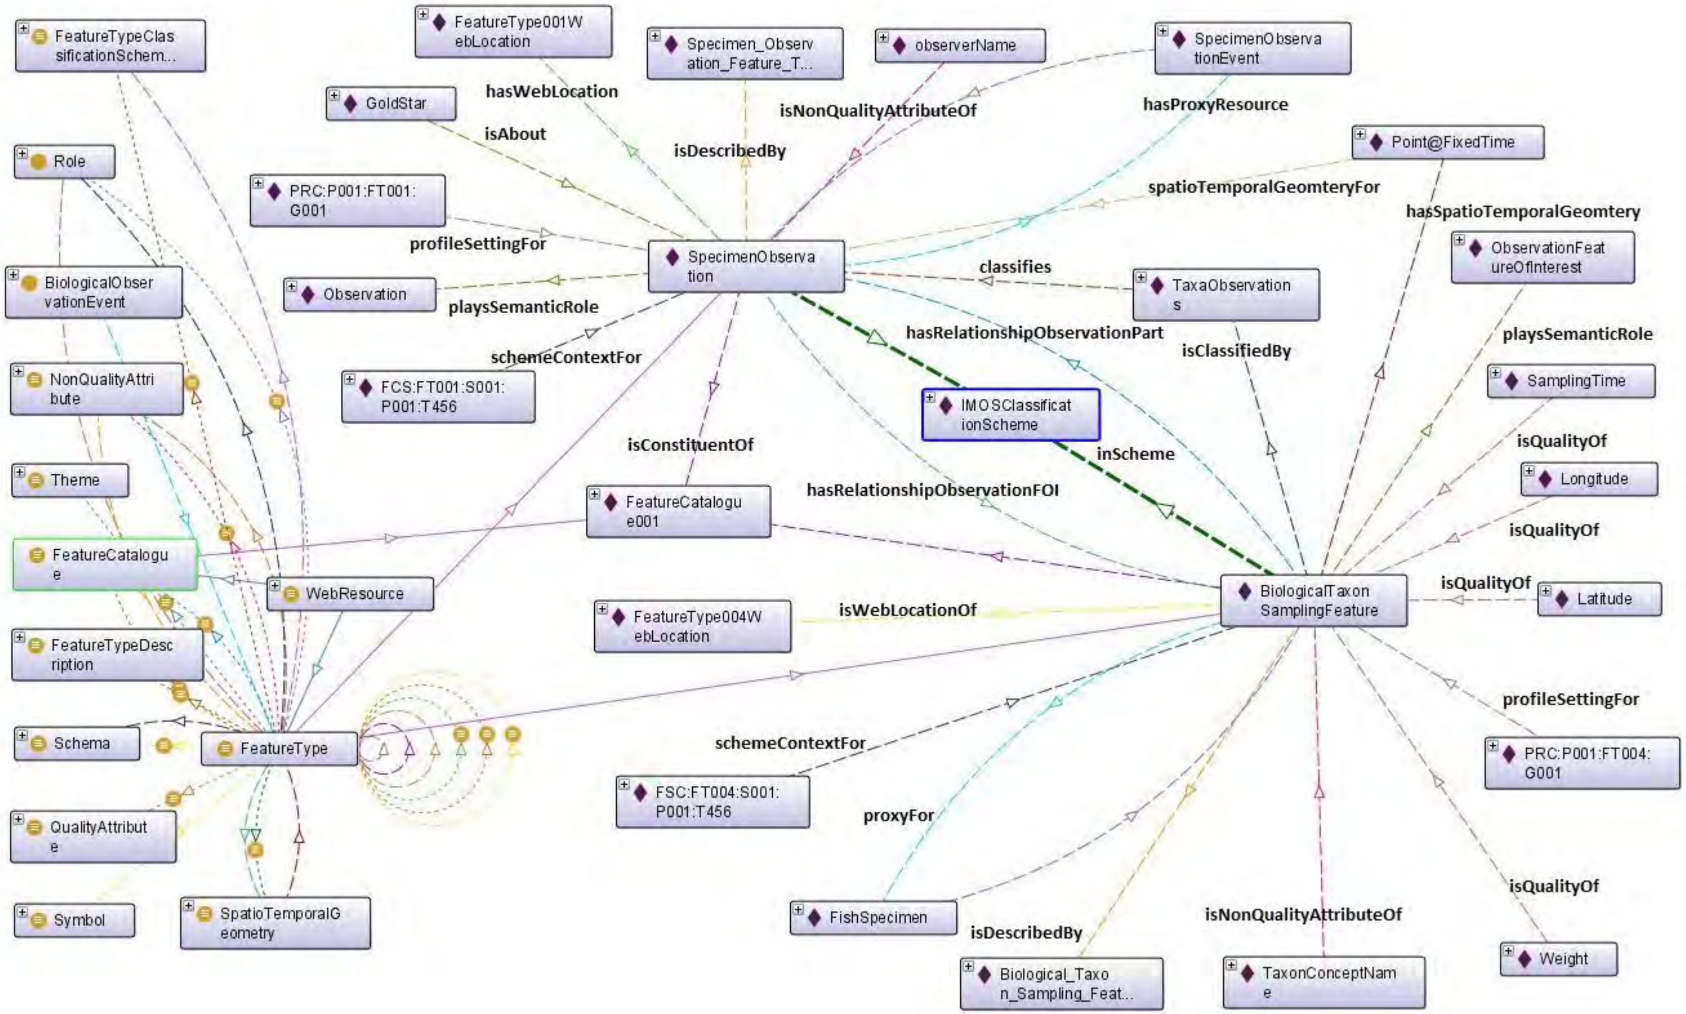

Figure 4.35 Overview of *SpecimenObservation* and *BiologicalTaxonSamplingFeatures* And Their Associations.

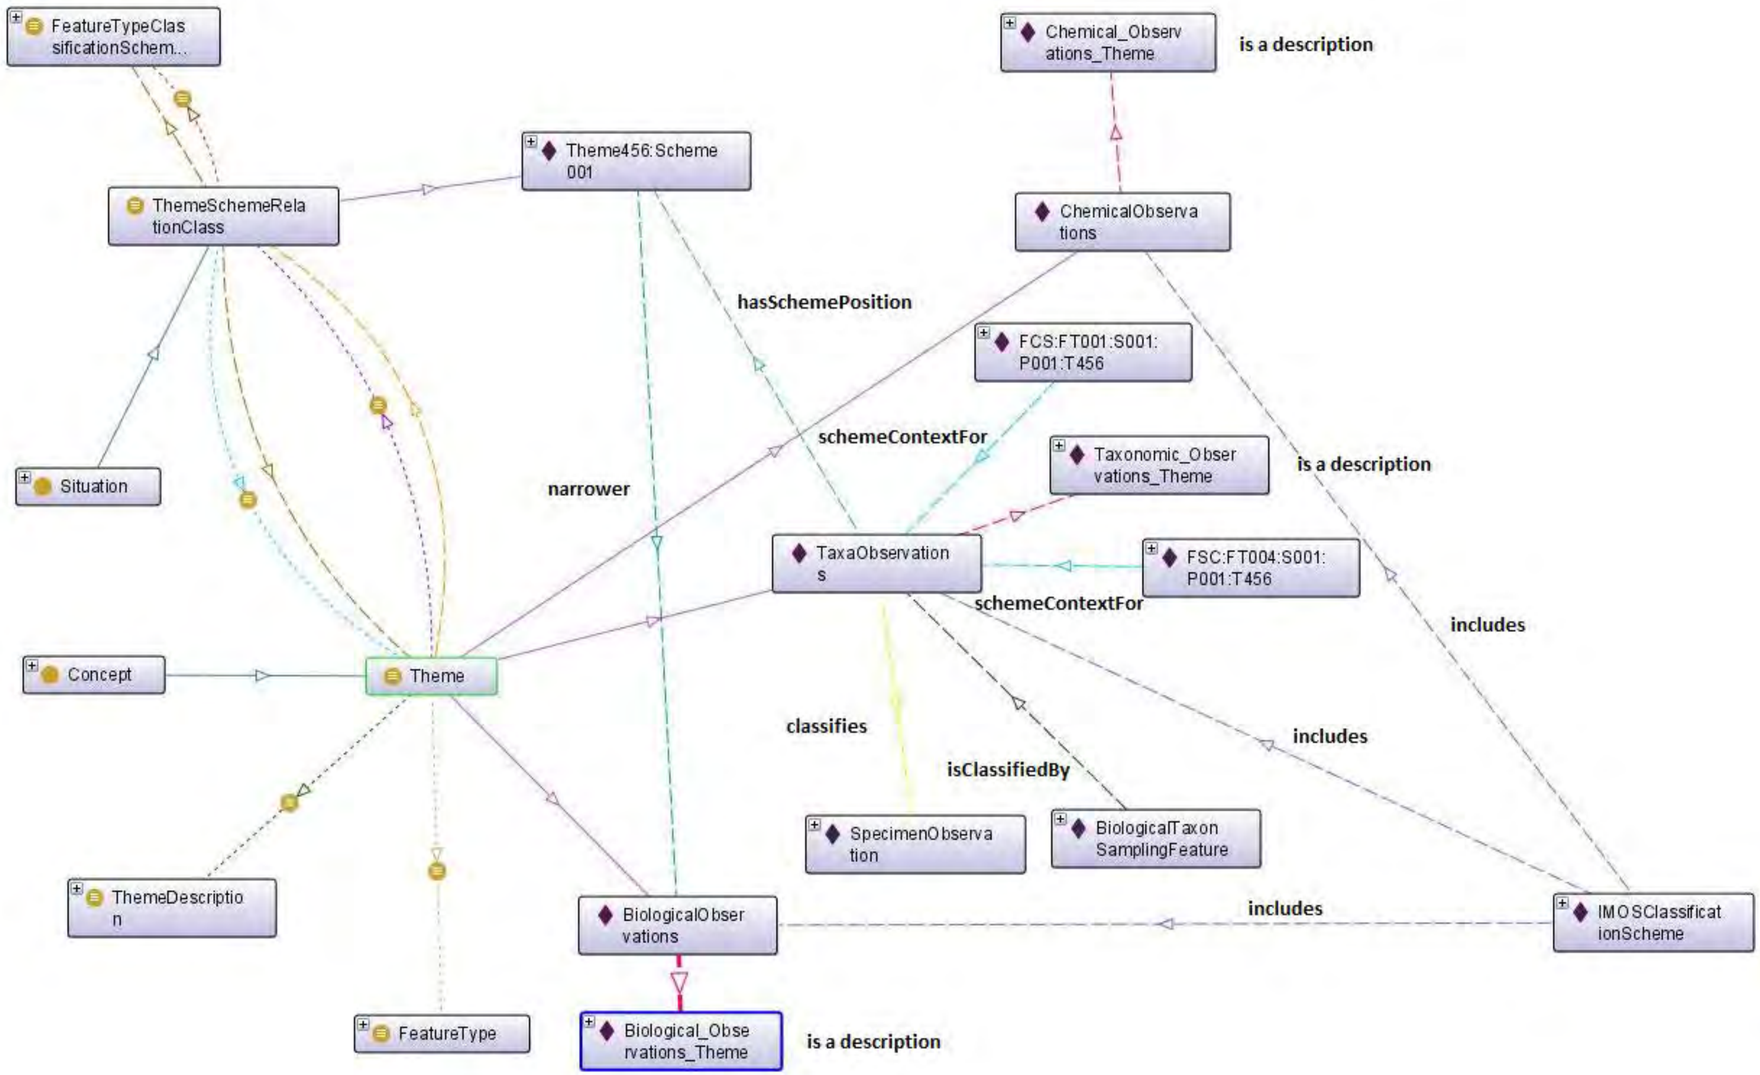

Figure 4.36 *Theme* Individuals And Their Relationships

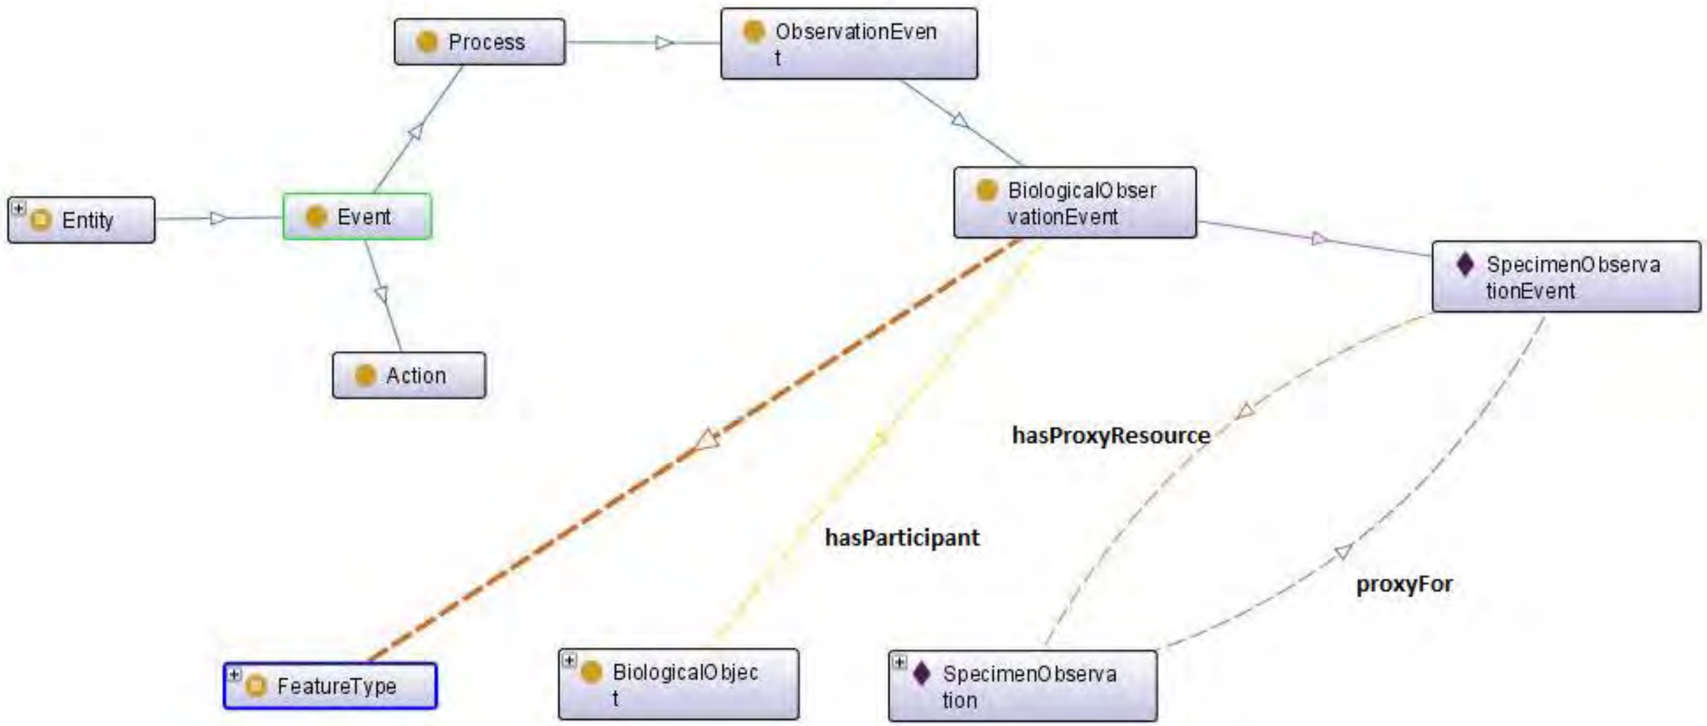

Figure 4.37 Relationships Between A *FeatureType* (*SpecimenObservation* Individual) and its real-world event (Individual)

Ontology Annotation

Annotation plays an important role in FCATOWL. Although the reasoner will ignore annotation this device was used to service a wide range of requirements articulated by the communities-of-interest. The rationale for relying on annotation as an integral part of the model was to keep the model as simple as possible and to limit an administrator's exposure to a potentially constant ontology build process.

Having imported 'DublinCore' there was a good choice of common descriptors available via the Protege 4 annotation tool (see Figure 4.38). Instead of ontologically defining things like "Definition Source"; "Creator"; "CreationDate"; "Comment"; "Deprecated"; "Identifier"; "Description" these were all routinely included as annotation.

Annotation can also be used to satisfy a requirement that the Catalogue Content Model should be capable of conveying whether a (fc:Quality or fc:NonQuality) *Attribute* is considered a mandatory Attribute. Although there is not a specific descriptor in DublinCore to express a "mandatory" condition, the Dublin Core "Comment" field could be used. Alternatively a new annotation property – called "mandatory" could be created.

Other types of annotation used less frequently, but which could prove useful included: "language" and "mappableTo".

In populating FCATOWL with sample data, "language" was used to refer to the fact that the Catalogue instance - "*Catalogue001"* was only available in English. For Catalogues such as FCATOWL which would be used internationally, an annotation property such as the "SKOS:prefLabel" could be used to provide alternate concept names, expressed in different languages.

The "mappableTo" annotation could be used as a boolean flag to indicate whether the literal value identified via the "source" property is an externally available semantic definition (as opposed to a lexical citation).

Annotation is used deliberately in FCATOWL to manage the more "administrative" components of the Feature Catalogue content model including content creation details, user contact information, update and auditing functions.

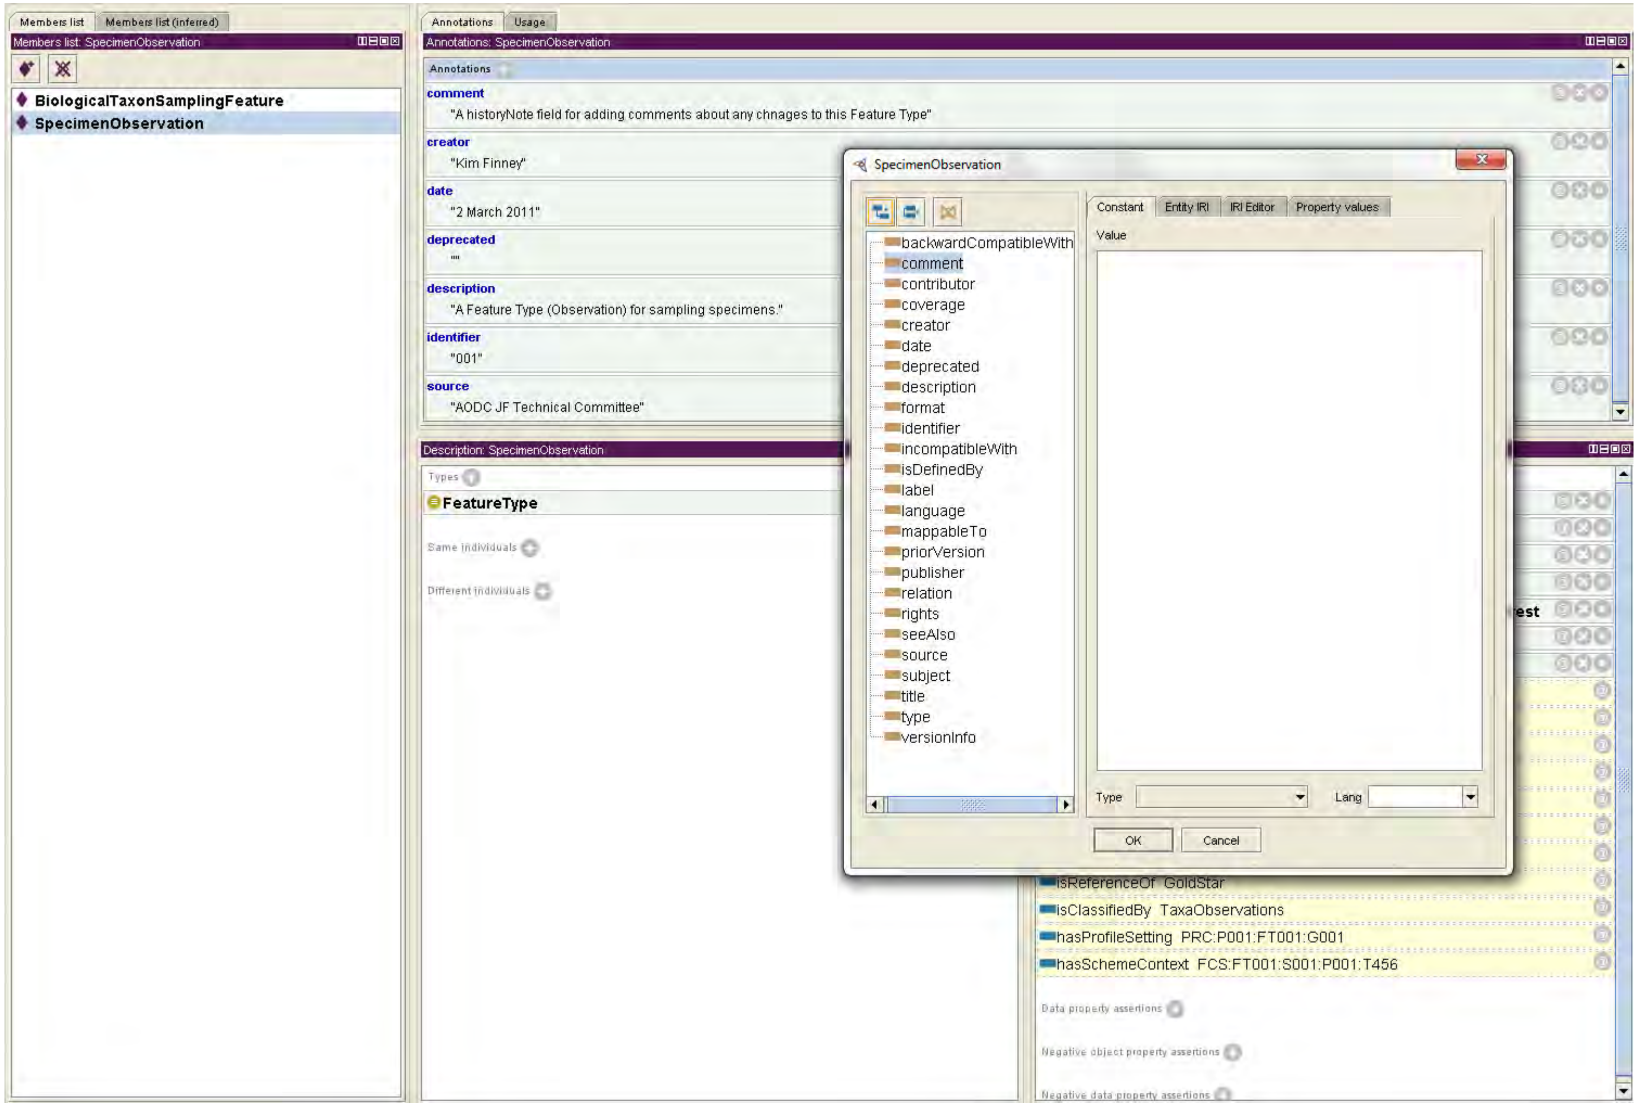

Figure 4.38 Imported Dublin Core Annotation Properties In Protege.

Using SPARQL

Before moving on to describe the prototyping process involving the establishment of a relational repository with a custom service interface based on a REST architecture, some commentary is required about SPARQL, the main access interface trialled to access content in FCATOWL. The TopBraidComposer™ client-based tools were able to extract data held remotely in Oracle™, through a client server interface (i.e., a Jena Application Program Interface and a secure Virtual Private Network connection) as well as data held locally (and natively in OWL), using SPARQL. Since TopBraid™ was designed to integrate closely with Oracle™, configuring the necessary connection was relatively simple. However, a few comments should be made about the use of SPARQL in different contexts (not involving TopBraid Composer™).

In order to connect to a populated FCATOWL ontology from anywhere on the web using a client (if it were stored within Oracle™ at a host site), Oracle™ administrators would need to establish a SPARQL endpoint. A SPARQL endpoint is a service endpoint capable of interpreting SPARQL queries and returning results over the HTTP protocol. Oracle™ uses a third-party SPARQL engine called Joseki (HP, 2009) that works in combination with a Jena Adapter.

Jena (HP, 2012) is an Open Source framework developed by Hewlett-Packard and is available under a Berkeley Software Distribution-style (open) license. Jena provides an RDF API; permits the reading and writing of RDF in RDF/XML and other serialisations such as N3 (Wikipedia, 2012c) and N-Triples (W3C, 2001); has an OWL API; has in-memory and persistent storage and a SPARQL query engine. Oracle™ uses Jena to extend its native capabilities which currently don't implement the SPARQL protocol. Although no experiments were executed to establish a SPARQL endpoint in order to connect to FCATOWL (because one was available via TopBraidComposer™), a scan of Oracle's™ documentation (Oracle, 2010a) and Oracle™ user forums (Oracle, 2010b) indicated that it is a far from trivial exercise to stand up the endpoint.

However, assuming that an endpoint is possible to establish, a Web client must issue a legal SPARQL query in order to extract content from the Feature Catalogue (FCATOWL). In most cases this will require an understanding of the ontological schema. A user must know in advance the schema class names, properties and individuals that comprise the ontology's contents. Alternatively, suitable mechanisms will be necessary to expose this information to the user, possibly in the form of a custom built query interface (Koutsomitropoulo *et al.,* 2011).

SPARQL is designed to be used with RDF graphs so in its core form doesn't support most of the OWL vocabulary. There are, however, semantic extensions (or entailments) for SPARQL (Glimm and

Krötzsch, 2010). The added complexity inherent in the syntax and grammar of queries resulting from the use of these entailments will, in most cases, be foreign to all but expert users. To combat this complexity some SPARQL-like alternatives have been proposed such as SPARQLAS (Schneider, 2010) and Terp (Sirin *et al.,* 2010) which show promise. Terp merges SPARQL grammar with Manchester OWL syntax to produce a much more user-friendly set of query constructs. A Terp parser is now available as part of the Pellet Reasoner (version 2.1) and is integrated through the Jena architecture, but it is currently in first release.

It is important to highlight that any extended protocol and query language constructs must be supported by the SPARQL engine that serves as the SPARQL endpoint. Currently many of the available engines do not support all of the SPARQL 1.1 constructs and extensions (as SPARQL 1.1 is still a W3C working draft).

Finally, a problem that is relatively high risk in establishing SPARQL endpoints is the potential for performance issues to arise associated with queries that are ill-formed by users and which as a result end up consuming too many resources at the server end (ultimately ending in a denial of service for other users).

4.7 Prototype REST-Based Feature Catalogue Access Methods

Having presented the results of the prototyping process for instantiating a Feature Catalogue based on a formal ontology (stored in a triple store and in OWL), using both proprietary and open source tools, an alternative approach is now outlined. This second approach is one based on a less formal ontological model (expressed in SKOS) that uses a more traditional technology to manage the semantic content (i.e., a relational database repository). The final research question related to Feature Catalogue development is also addressed, i.e., RQ1.1.4 -"*What methods are best suited to extract re-usable content from an ontologically-grounded Feature Catalogue*".

The methodologies possible for accessing data (content) in a real-world repository application scenario will generally depend on whether the focus of the application is intended to be through a Web-client or via services, or a combination of both. The community requirements for the Feature Catalogue repository/server indicated that both types of interface were desired. Since the choice of Web client development environment would generally distil down to two issues: compatibility of the application software development tools with the existing software infrastructure of the Catalogue host agency and the skills available in that agency to build the application; it is more sensible for this study to constrain any investigations regarding access issues to those involving services. Although services are also influenced to a degree by these factors, when using services the desire to achieve service interoperability (external to the hosting agency's context) is usually more influential in directing the shape and character of any developed services, than is the existing development environment.

In choosing service interfaces, a number of fundamental choices of method exist for service approaches. It is feasible in general terms, to embark down three different routes: (a) use proprietary services, or (b) pick a complementary suite of open standards-based Web service templates for a tailored implementation, or (c) use services based on Representational State Transfer (REST) architecture (also an open approach).

Proprietary-end points and/or custom service interfaces will often already be built into off the-shelf software (e.g., Oracle™ RDBMS software supports a SPARQL endpoint), so the choice of content management system can also influence the access paradigms chosen. However, if option (b) was the desired implementation route, some well-known and applicable open service standards are available including those found in the OGC Web services suite (that enable content transfer, query and transaction) or one could use SPARQL - the standard RDF query language. As previously explained in earlier chapters, OGC services, whilst open, are not part of the general (W3C) Web Services standard stack. This might limit their broad interoperability with systems not currently using this set of standards. Option (c) is therefore of particular interest because REST maximizes the use of pre-existing, well-defined interfaces and other built-in capabilities that are already provided by the Web (HTTP) network protocol, and additionally REST minimizes the requirement for adding new application-specific features on top of this existing ubiquitous protocol.

The choice of a REST-based interface solution provides access to Catalogue content consistent with other approaches being taken for other ontological repositories (Noy *et al.,* 2009; Viljanen *et al.,* 2010). Although the prototype instantiation to be described in this section is less formal in ontological terms (than it would be if OWL had of been used) it is still based on the conceptual design presented earlier in Figure 4.17. Since the focus in this prototyping activity was on trialling 'access methods' as opposed to ontological content models, it was decided to limit the complexity involved (and to see what benefits could be had from using a less formal ontology language). Some of the work presented next has previously been published in Finney and Watts (2011).

Ordinarily SKOS is used to manage simple classification schemes and is useful for mapping between different structured and controlled vocabularies (Isaac *et al.,* 2009). It is not able to be reasoned over, but it does force the representation of information into a particular set of hierarchical/taxonomic and associative relations that are now widely familiar (that are written in RDF). Although the inherent “semantics” are not expressed in a machine processable form, the SKOS schema provides sufficient descriptions to facilitate interoperation between applications using SKOS vocabularies. SKOS also allows for extensions to the vocabulary (e.g., through sub-properties and concept specialisation). SKOS therefore provides a vocabulary with an associated description of how that vocabulary should be interpreted, which takes us one step beyond using basic XML tagging. Although the Feature Catalogue content Model is not a simple classification scheme, it was decided to see if an extended SKOS vocabulary could be usefully deployed to describe the Feature Catalogue model.

4.7.1 The Enhanced Feature Catalogue Model In A Relational Database Form

The first step in building the relational Oracle-based Feature Catalogue repository (on which the SKOS model would be based) was to translate the Feature Catalogue content Model (given in Figure 4.17) into an appropriately normalised set of physically-realised database tables and relations. The schema developed is given in the ER diagram in Figure 4.39. The accompanying data dictionary is located in Appendix 11. A set of test data was then developed and inserted manually into the tables using ORACLE™ SQL Developer, ensuring that all integrity constraints were observed.

To check that the entities and relationships were modelled appropriately, several SQL queries were constructed to extract data from the database and the results checked to see if the data returned, matched expectations. The types of (SQL) queries indicative of the tests performed can be seen in Table 4.8.

Three associative RDBMS entities (Feature_Attribute; Attribute_UoM and Attribute_Method) that were established in the database schema (in Figure 4.39) are not essential for the model to function as desired but they were created for testing purposes so that data extractions (based primarily on the use of foreign keys in the table: Feature_Type_Attribute_Profile) could be cross-checked. The Feature_Type_Attribute_Profile entity alone is sufficient to group the requisite entities (and table attributes) and permit the necessary database extractions.

4.7.2 Overview Of REST-Based Catalogue Services and Service Descriptions

Having developed the RDBMS repository model and seeded it with test data, the likely scenarios in which the content would be used were re-visited. The three scenarios common in the use-cases provided earlier (in section 4.2.5, Appendix 8) included the repository content being accessed:

- by data providers when developing their exchange schema to check for conformity with community agreed standards on feature type definitions,

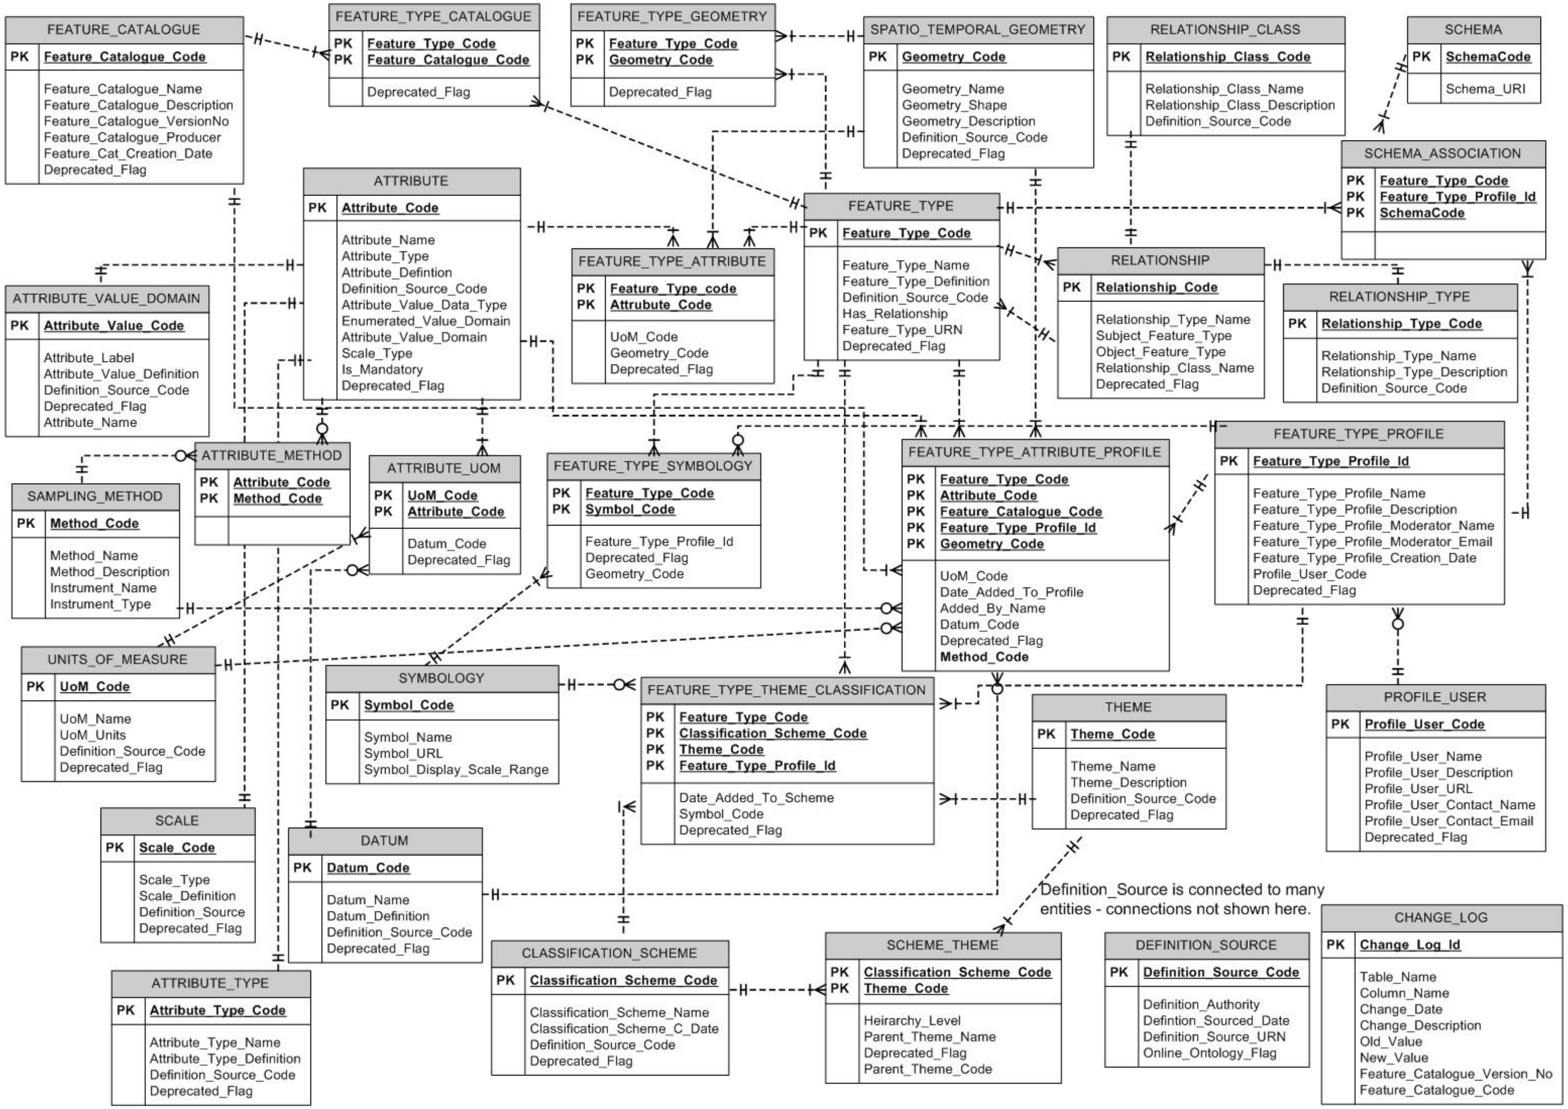

Figure 4.39 ER-Diagram For Feature Catalogue Content Model

Table 4.8 Sample SQL Queries Run Over (Feature Catalogue) Database

Scheme_Theme Query (gets all schemes with their themes)

```
select SS.CLASSIFICATION_SCHEME_NAME, SS.DEFINITION_SOURCE_CODE,
SS.DEFINITION_AUTHORITY, TS.THEME_NAME, TS.DEFINITION_SOURCE_CODE,
TS.DEFINITION_AUTHORITY
from SCHEME_SOURCES SS, THEME_SOURCES TS, scheme_theme ST
where SS.CLASSIFICATION_SCHEME_CODE = ST.CLASSIFICATION_SCHEME_CODE
and ST.THEME_CODE = TS.THEME_CODE;
```

Feature_Type_Profile Query (gets all Profiles with associated Feature Types)

```
select FTP.FEATURE_TYPE_PROFILE_NAME, FT.FEATURE_TYPE_CODE , FT.FEATURE_TYPE_NAME,
A.ATTRIBUTE_NAME,  FT.DEFINITION_SOURCE_CODE as feature_source,
G.DEFINITION_SOURCE_CODE as geom_source,
A.DEFINITION_SOURCE_CODE as att_source, G.GEOMETRY_NAME, U.UOM_NAME, D.DATUM_NAME
from attribute_sources A,   units_of_measure_sources U,
feature_type_sources FT, datum_sources D, feature_type_attribute_profile FTA,
feature_type_profile FTP, geometry_sources G
where FTP.FEATURE_TYPE_PROFILE_CODE = FTA.FEATURE_TYPE_PROFILE_ID
and FTA.FEATURE_TYPE_CODE = FT.FEATURE_TYPE_CODE
and FTA.ATTRIBUTE_CODE = A.ATTRIBUTE_CODE
and FTA.GEOMETRY_CODE = G.GEOMETRY_CODE
and FTA.UOM_CODE = U.UOM_CODE(+)
and FTA.DATUM_CODE = D.DATUM_CODE(+)
order by FT.FEATURE_TYPE_CODE;
```

Feature_Type Query (gets all Feature Types and associated Attributes)

```
select FT.FEATURE_TYPE_CODE , FT.FEATURE_TYPE_NAME, A.ATTRIBUTE_NAME,
  FT.DEFINITION_SOURCE_CODE,
A.DEFINITION_SOURCE_CODE, U.UOM_NAME, D.DATUM_NAME, G.GEOMETRY_NAME
from attribute_sources A,  attribute_uom AU, units_of_measure_sources U,
feature_type_sources FT, datum_sources D,
feature_type_attribute FA, geometry_sources G
where FT.FEATURE_TYPE_CODE = FA.FEATURE_TYPE_CODE
and FA.ATTRIBUTE_CODE = A.ATTRIBUTE_CODE
and A.ATTRIBUTE_CODE = AU.ATTRIBUTE_CODE(+)
and D.DATUM_CODE(+) = AU.DATUM_CODE
and AU.UOM_CODE = U.UOM_CODE(+)
and G.GEOMETRY_CODE = FA.GEOMETRY_CODE
order by FT.FEATURE_TYPE_CODE;
```

- dynamically by a machine client attempting to resolve and interpret an in-line reference to some repository content, where the reference is located in an exchanged dataset, and
- by a machine client requesting, or posting repository content to fulfil some component of its own internal functionality.

The additional challenges during this prototyping activity (as opposed to the previous FCATOWL ontology modelling approach) lay in creating the SKOS ontology capable of managing the desired Feature Catalogue content, and developing a service interface to the repository that permitted the implementation of the three use-cases listed, in a manner that would present a very low uptake barrier for programmers and domain specialists alike.

In true service-oriented-architectures, the service interface specification between provider and consumer must be supplied as well as the content interface (as was discussed earlier in Chapter 2). Whilst considerable effort has gone into devising and harmonising open standards supporting service-oriented-architectures, implementation remains complex, not the least because of the freedom available to devise one's own service interface (The Opengroup, 2009). Although the OGC services stack is relatively well specified, not all possible consumers of the Feature Catalogue ontology service will conform to OGC standards. Ideally the access service needs to be agnostic of the service protocols adopted by potential consumers. REST-based services present all resources to clients with one uniform interface and are stateless. Vinoski (2007) argues a significant advantage of REST is that the uniform interface constraint provides for better scalability by removing the entire interface contract term from the client–service interaction equation. Rodriguez (2008) also points out that a REST Web service application (or client) includes within the HTTP headers and body of a request all of the parameters, context, and data needed by the server-side component to generate a response. Statelessness, Rodriguez argues, in this sense improves Web service performance and simplifies the design and implementation of server-side components because the absence of state on the server removes the need to synchronize session data with an external application.

It must be acknowledged, however, that some deficiencies exist in using REST. For example, a standardized form of transferring information (as is in REST) can be less efficient than transferring information in a format created specifically for an application. Additionally important aspects of Quality of Service, security and reliable messaging are not supported by this architecture (Higashino *et al.*, 2009). On balance the advantages, which equate to increased simplicity, prompted the author to choose REST.

It has already been discussed that implemented examples of operational on-line Catalogues are not common. Those Catalogues that do exist, almost without exception, do not use Web service interfaces so it was necessary to devise a set of RESTful service patterns for accessing catalogue content de novo. The approach to developing the REST service interface involved identifying those repository "resources" that the community wished to interact with directly (as first order objects). Intuitively and from reference to the use-cases reported earlier, the first order objects drawn from the Catalogue model should be: *FeatureCatalogue*; *Profile*; *FeatureType*; (spatio-temporal) *Geometry*; Symbology; *Attribute*; *UoM*; *Datum*; *SamplingMethod*; *Instrument*; *Schema*; *Relationship; Theme* and *ClassificationScheme*.

At the outset it was decided that the task of developing a service interface would be restricted to only exposing a representation of the content, rather than permitting content modification or content deletion. This prototype build is considered a proof-of-concept activity only and these additional functions could clearly be added later during a full development effort to permit two-way transactional activity with the repository.

In addition to identifying the content that would be exposed as resources, the option of nesting references (hyperlinks) within resources to permit an iterative drill-down to finer levels of detail (i.e., to expose other resources) was considered a necessity, since this is the basic model behind a REST-based approach. Three output message encodings (XML, HTML and SKOS) were determined to be the key alternate languages for delivering the payload, with the option of implementing OWL at later date in the case of a fully resourced development effort, based on the approach demonstrated in this thesis. Extending the (various) services developed to incorporate an OWL-based output would not be difficult given that a fully compatible OWL encoding has already been developed for the purposes of trialling the former ontology-based storage and management approach (and the foundation model for both the OWL and SKOS implementations is common to both).

The task of describing how to invoke the services (i.e., the service contract) came down to a choice between three different languages, each of which were applicable for describing services in a REST-based architecture: Web Services Description Language (WSDL 2.0); Web Application Development Language (WADL (Hadley, 2009)) and hREST microformats (Kopecky *et al.***,** 2008a).

The main items to be communicated about the service interface included:

- The service's URL,
- The communication mechanisms the service understands,
- The operations the service can perform, and
- The structure of the service's messages.

Takase *et al.* (2008) have provided a good overview of the advantages and disadvantages of using either WADL or WSDL. They concluded that WADL is relatively simple (as compared to WSDL) but has limited scope. They assert that the WSDL 2.0 HTTP binding is more feature rich, but at the cost of increased complexity, and it still lacks a true resource-centric model. Whilst either of these languages could have been applied to suit the contexts confronted in this study, it is telling that there are still very few services being described using these standards (including OGC services) in the scientific data exchange context. It was speculated that their perceived complexity continues to be a barrier to uptake.

The third option examined for describing the service contract, hREST microformats, are a set of simple, open data formats built upon existing and widely adopted standards. The communities working on microformats favour simplicity of approach and try to adapt their solutions to fit with existing web usage patterns and developer behaviour. Some of the better known microformats include hCard (for publishing contact details, Celik and Suda (2012a)), hResume (for publishing names and CV (King, 2012)) and hCalendar (for publishing a semantic representation of a calendar (Celik and Suda (2012b)). A guiding principle in the design philosophy of microformat developers is to design for humans first and machines second. Believing this philosophy to sit comfortably with the goal of lowering the uptake barrier, it was decided to implement hREST microformats for describing the prototyped REST-based services.

The hREST microformat takes advantage of existing XHTML facilities such as the **class, body, div, span** and **rel** attributes to mark the fragments of interest in a Web page. Additionally it translates HTML elements into a hierarchy of objects and their properties by embedding the key classes of Service; Operation; Address; Input; Output and Label within the XHTML mark-up facilities. For example, the HTML appearing on a web page as presented in Figure 4.40 would consist of the hREST mark-up as shown in Figure 4.41. In hREST an individual service can be comprised of many (numbered) and labelled operations and is called from a particular “address”. The “method” class describes the “service” method (e.g., GET) and the “input” and “output” classes define the input and output message formats respectively. The full definition of, and syntax constraints for each of the hREST classes is available on the web (Kopecky *et al.*, 2008b).

Copies of the microformat hREST capability (XHTML) documents created to describe each of the services that were subsequently developed in this research can be found in Appendix 12. The XML output formats referenced from within the hREST capability documents, for all services developed in this prototyping activity were described using W3C Schema (and are accessible online via the hyperlinks embedded in the hREST documents).

StreamlinedFeatureCatalogueService

```
1. GET Profile n
```

To make a query to retrieve individual profiles from the Feature Catalogue using the http method GET, use the request:

```
http://data.aad.gov.au/aadc/FeatureCatalogue/?/StreamlinedCatalogueService/Catalogue/{
id}/Profiles/{id}
```

The parameter: id - (a positive integer) is the identifier of a particular catalogue and a profile within that catalogue.

The output format for the request provides profile information encompassing Feature Types, Feature Type Attributes, UoMs, Datums, Feature Type Relationships and the Classification Schemes and Themes that the profile's Feature Types belong to. The response is supplied according to an XML schema document - Profile_Package.xsd

Figure 4.40 Human Readable Web Page

```
<html>
<body>
<div class="service" id="svc">
<span class="label"><strong>StreamlinedFeatureCatalogueService</strong></span>
<div class="operation" id="1"><br/>
<code class="label"> 1. GET Profile n </code><br/>
To make a query to retrieve individual profiles from the Feature Catalogue using the http method
<span class="method">GET</span>, use the request: <code class="address">
http://data.aad.gov.au/aadc/FeatureCatalogue/?/StreamlinedCatalogueService/Catalogue/{id}/Profile
s/{id}</code> <br/> </p>
<p><span class="input"> The parameter: <code>id</code> - (a positive integer) is the identifier of a
particular catalogue and a profile within that catalogue.</span><br/></p>
<p><span class="output">The output format for the request provides profile information
encompassing Feature Types, Feature Type Attributes, UoMs, Datums, Feature Type Relationships and
the Classification Schemes and Themes that the profile's Feature Types belong to. The response is
supplied according to an XML schema document - <a href
="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/Profile_Package.xsd">Profile_Package.xsd
</a></span></p>
</div>
</div>
</body>
</html>
```

Figure 4.41 hREST machine interpretable code for Figure 4.40

4.7.3 REST-Based Feature Catalogue Service Patterns (HTML and XML Output)

The task of developing the services was split into two stages, primarily to see if anything could be gleaned from implementing the first part, which could then be applied in implementing the remaining work. It was decided to initially develop three types of RESTful service (which were called "Streamlined", "Iterative" and "Component" catalogue services respectively), each of which targeted a different use-case, and which would output content in HTML and XML only. Lessons learnt from implementing these services were used to inform development of an extended "Iterative" service that was capable of also delivering its payload using SKOS (described next in section 4.7.4). ColdFusion version 8 was used in both stages to develop the service interface software.

Two of these services ("Streamlined" and "Iterative") regardless of their output types, expose a specific *Catalogue's* resource content according to a declared community (*FeatureType*) *Profile* and the other ("Component") service provides access to a *Catalogue's* content without any pre-conditions, except for the requirement of providing the ID of the *Catalogue* instance that is to be interrogated. These three services and their patterns of implementation are discussed further below. As part of the prototyping activity a Web page (AADC, 2012c) was established to demonstrate the XML and HTML services developed in stage one. A snapshot showing the first part of the web page is provided in Figure 4.42.

It should be noted in implementing stage one of the service interfaces for the relational Feature Catalogue model, some resources (entities) were omitted from the services (namely an *Attribute's SamplingMethod* and any *Schema* encodings that a *Feature Type* might participate in). This was done purely to limit the amount of coding that would be required to develop the services. AADC programmer time was limited to work with the author on the services, so some boundaries had to be set. The service pattern outlines which follow next state what was actually delivered in these prototype services and an explanation is also given of how these service would operate if they had included the missing resources. The omissions do not affect the overall patterns devised for the services, nor the value of these demonstrations as proof-of-concept.

Streamlined Service

The "StreamlinedCatalogueService" provides resource information (see Figure 4.43) without the need to drill down iteratively to access (nested) content. This service requires that a *Catalogue* ID is provided as the first part of the service call string then a community (FeatureType) *Profile* resource must be nominated. If a specific *Profile* ID is provided, all resources requested and delivered as a result will have been assigned to this particular *Profile*. To restrict the payload to cover a particular *Feature Type*, the *Feature Type* resource must be part of the query string along with an ID.

In this service type all information associated with a particular *Profile* is delivered in one step, bar the detailed structure of any associated *ClassificationScheme*s (which requires an extra service call to obtain the full *Scheme* information). Identification details about which *ClassificationSchemes* that an individual feature has been assigned to are, however, delivered as part of the first service query. As explained previously, *Sampling Methods* associated with any Feature Type's *Attributes* and any (data encoding) *Schema* that a specific *Feature Type* participates in are not currently part of the pilot service payload.

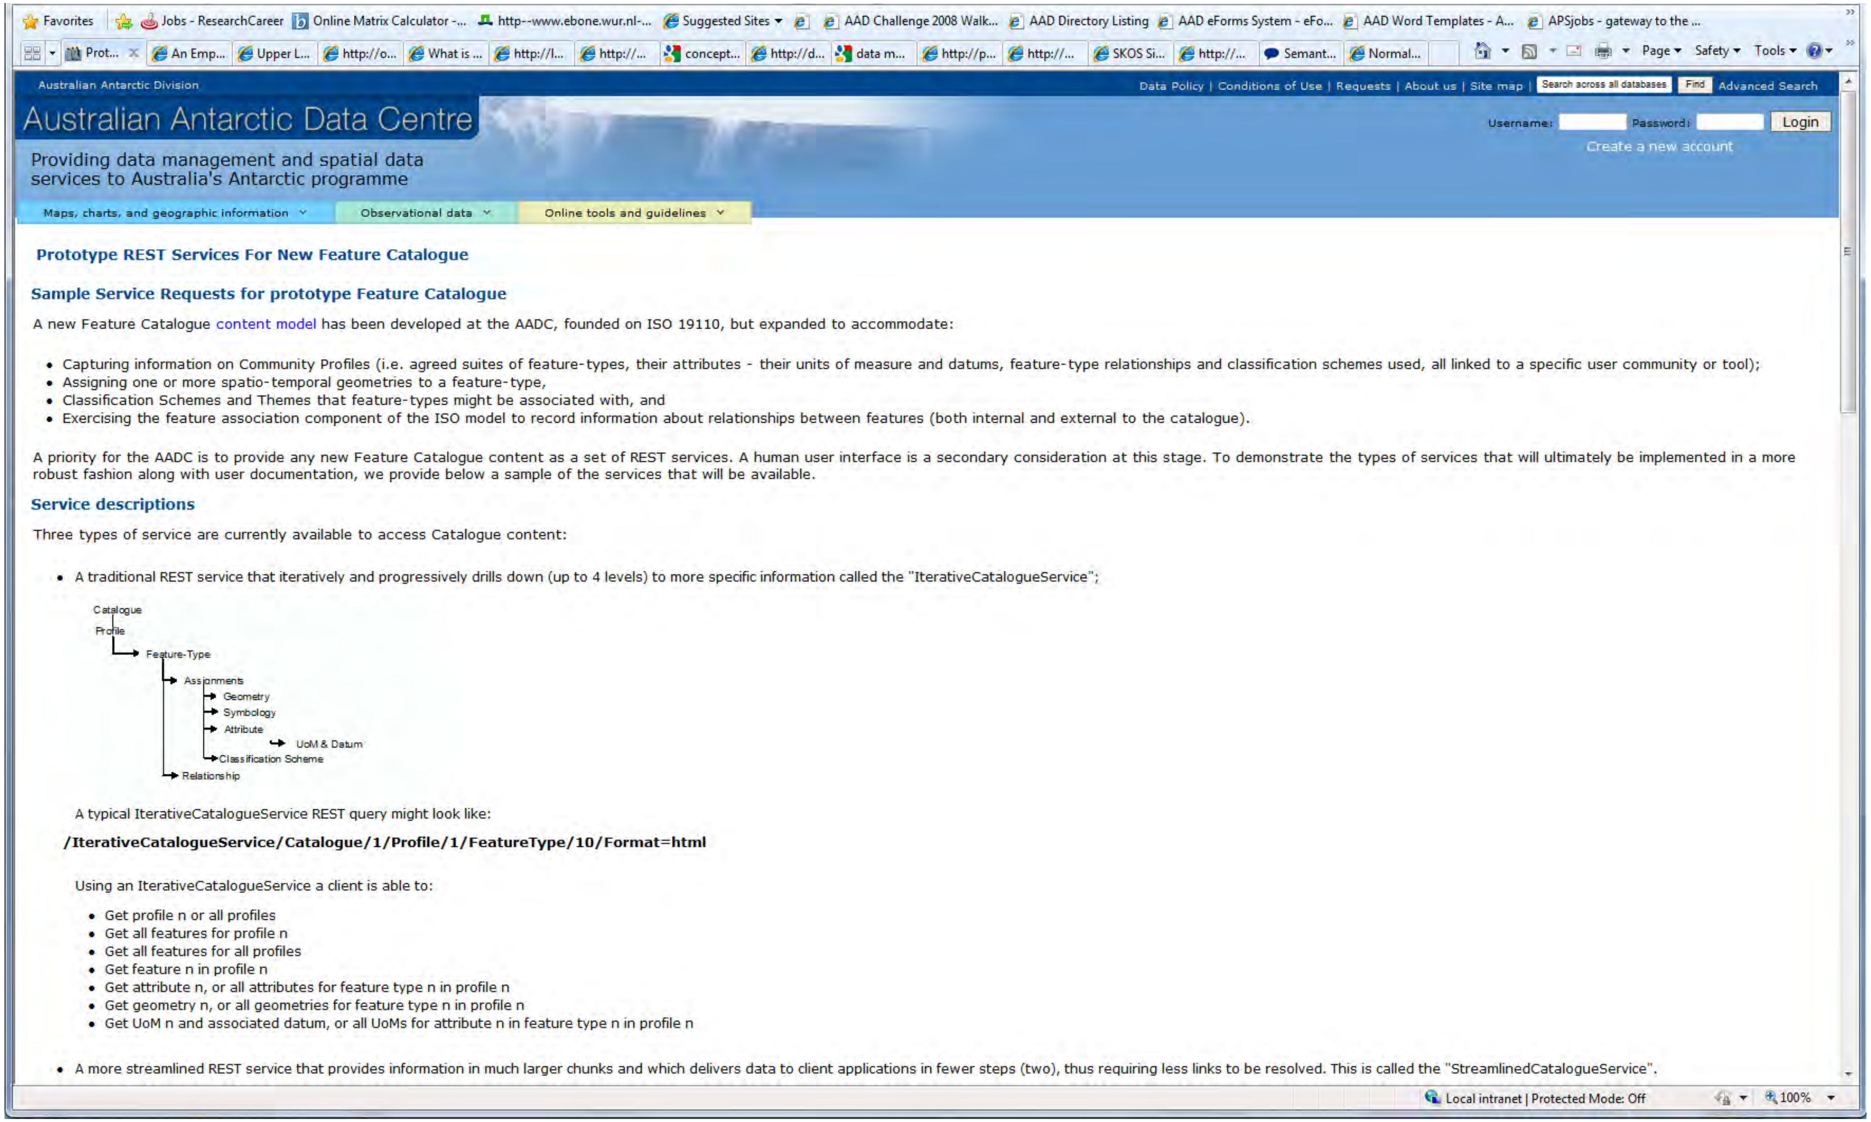

Australian Antarctic Division

Data Policy | Conditions of Use | Requests | About us | Site map | Search across all databases | Find | Advanced Search

Australian Antarctic Data Centre

Username: Password: Login

Create a new account

Providing data management and spatial data services to Australia's Antarctic programme

Maps, charts, and geographic information | Observational data | Online tools and guidelines

Prototype REST Services For New Feature Catalogue

Sample Service Requests for prototype Feature Catalogue

A new Feature Catalogue content model has been developed at the AADC, founded on ISO 19110, but expanded to accommodate:

- Capturing information on Community Profiles (i.e. agreed suites of feature-types, their attributes - their units of measure and datums, feature-type relationships and classification schemes used, all linked to a specific user community or tool);
- Assigning one or more spatio-temporal geometries to a feature-type,
- Classification Schemes and Themes that feature-types might be associated with, and
- Exercising the feature association component of the ISO model to record information about relationships between features (both internal and external to the catalogue).

A priority for the AADC is to provide any new Feature Catalogue content as a set of REST services. A human user interface is a secondary consideration at this stage. To demonstrate the types of services that will ultimately be implemented in a more robust fashion along with user documentation, we provide below a sample of the services that will be available.

Service descriptions

Three types of service are currently available to access Catalogue content:

- A traditional REST service that iteratively and progressively drills down (up to 4 levels) to more specific information called the "IterativeCatalogueService";

A typical IterativeCatalogueService REST query might look like:

/IterativeCatalogueService/Catalogue/1/Profile/1/FeatureType/10/Format=html

Using an IterativeCatalogueService a client is able to:

- Get profile n or all profiles
- Get all features for profile n
- Get all features for all profiles
- Get feature n in profile n
- Get attribute n, or all attributes for feature type n in profile n
- Get geometry n, or all geometries for feature type n in profile n
- Get UoM n and associated datum, or all UoMs for attribute n in feature type n in profile n

- A more streamlined REST service that provides information in much larger chunks and which delivers data to client applications in fewer steps (two), thus requiring less links to be resolved. This is called the "StreamlinedCatalogueService".

Local intranet | Protected Mode: Off 100%

Figure 4.42 Snapshot of The Feature Catalogue Service Prototype Page

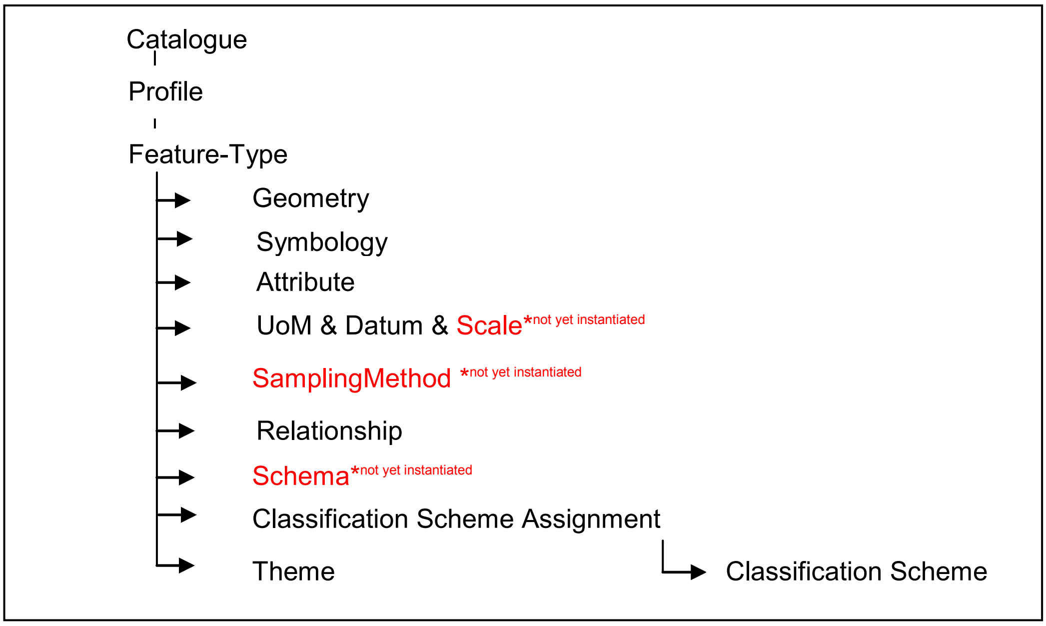

Figure 4.43 Resources available from the StreamlinedCatalogueService and access hierarchy

The “StreamlinedCatalogueService” response delivers information on the *Profile* itself (i.e., who owns it, when it was created etc); it also exposes all associated *Feature Types*; lists *Attribute* assignments for each *Feature Type*; details *Attribute* information including nominated *UnitsofMeasure* (*UoM*) and any *Datum* information; categorises *Feature Types* by nominated *ClassificationSchemes* and identifies the individual *Themes* to which *Feature Types* belong. Lastly, the service lists all of the *Relationships* in which the *Feature Types* (in a given *Profile*) participate. See Figure 4.44 for an abbreviated sample output format. Figure 4.45 expands part of the abbreviated output to show how the individual resources are described in the live service.


```
<?xml version="1.0" encoding="UTF-8" ?>
<Feature_Type_Catalogues xmlns:xlink="http://www.w3.org/1999/xlink" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/Feature_Profile.xsd">
 <Catalogue id="1">
        <Catalogue_Name>AADC Feature Catalogue</Feature_Catalogue_Name>
        <Catalogue_Description>Catalogue of Antarctic Feature Types</Feature_Catalogue_Description>
        <Catalogue_VersionNo>V1.0</Feature_Catalogue_VersionNo>
        <Catalogue_Producer>AADC</Feature_Catalogue_Producer>
        <Cat_Creation_Date>17-Jun-2009</Feature_Cat_Creation_Date>
  <Feature_Type_Profiles>
    <Feature_Type_Profile id="1">
      <Feature_Type_Profile_Name>Test_Profile_No1</Feature_Type_Profile_Name>
      <Feature_Type_Profile_Description>A test profile, number 1</Feature_Type_Profile_Description>
      <Feature_Type_Profile_Moderator> Kim Finney</Feature_Type_Profile_Moderator>
      <Feature_Type_Profile_Moderator_Email> kim.finney@aad.gov.au</Feature_Type_Profile_Moderator_Email>
      <Feature_Type_Profile_User>
        <Profile_User_Name>AADC</Profile_User_Name>
        <Profile_User_Description>Australian Antarctic Data Centre</Profile_User_Description>
        <Profile_User_Contact_Name>Ursula Harriss</Profile_User_Contact_Name>
        <Profile_User_Contact_Email>ursula.harriss@aad.gov.au</Profile_User_Contact_Email>
      </Feature_Type_Profile_User>
    <Feature_Types>
      <Feature_Type id="4">
        <Feature_Type_Name>Aircraft_Runway</Feature_Type_Name>
        <Feature_Type_Definition>a paved surface on which planes land and take off</Feature_Type_Definition>
```

```
            <Definition_Authority>Dictionary.com</Definition_Authority>
            <Definition_Source id="51" xlink:href="http://dictionary.reference.com/browse/Runway"/>
            <Assignments id="1" Geometry_CodeIDREF="5" Attribute_Code IDREF="8, 11">
              <Association_Scheme_Assignments id="1" Scheme_Code IDREF="1" Theme_Code IDREF="2" />
            </Assignments>
            <Relationship_Assignments Id="1" Relationship_Code IDREF="1" />
          </Feature_Type>
            ..........
          < Feature_Geometry Geometry_Code id="5">
          {Geometry code descriptive information for geometry with id=5, see Figure 4.45 for example expansion of descriptive content}
          </Feature_Geometry>
            .........
          < Feature_Attribute Attribute_Code id="8">
            .........
          </Feature_Attribute>
          < Feature_Attribute Attribute_Code id="11">
            .........
          </Feature_Attribute>
          < Feature_Scheme Scheme_Code id="1">
            .........
          </Feature_Scheme>
          < Feature_Theme Theme_Code id="2">
            .........
          </Feature_Theme>
          < Feature_Relationship Relationship_Code id="1">
            .........
          </Feature_Relationship>
      </Feature_Types>
            .........
    </Feature_Type_Profile>
   </Feature_Type_Profiles>
 </Catalogue>
</Feature_Type_Catalogues>
```

Figure 4.44 Abbreviated and annotated XML code snippet for the StreamlinedCatalogueService

```
<Assignments id="1" Geometry_CodeIDREF="5"......................./>
......
......
<Feature_Geometry Geometry_Code="5">
            <Geometry_Name>Trajectory</Geometry_Name>
            <Geometry_Shape>2D</Geometry_Shape>
            <Geometry_Description>A geometric figure formed by a point moving in a fixed direction along a path in time and
space.</Geometry_Description>
            <Definition_Authority>A. Woolf (2007) Climate Science Modelling Language Version 2    </Definition_Authority>
            <Definition_Source id="31" xlink:href="http://csml.featuretypes.definitions"/>
 </Feature_Geometry>
........
```

Figure 4.45 An expansion of the abbreviated XML code snippet showing the detail exposed through the Feature_Geometry resource.

To retrieve a detailed description of the *ClassificationSchemes* to which the *Feature Types* (in the Profile Resource) belong, the client rendering this resource would need to traverse and resolve an "xlink" (in the resource) pointing to the listed instance of a *ClassificationScheme* resource. So, even though a requestor gets a lot of information in one hit (from the Profile resource), they still need to resolve embedded links to get *Feature Type* classification information. The use-case supporting the "StreamlinedCatalogueService" is one where it would be undesirable from a performance overhead perspective, to continuously traverse and resolve numerous hyperlinks to access desired resource information. The payload is therefore delivered as one cohesive chunk of information.

If the *SamplingMethod* resource had been included as part of the service payload it would have been referenced in the XML “Assignments” element (see Figure 4.44 or 4.45) and expanded later in the XML document below the “Feature_Attribute” elements, via its own tagged element. Similarly, any (data encoding) *Schema* that the *Feature Type* participated in would have been allocated to a “Schema_Assignments” element and then inserted below the “Relationship_Assignments” element in the XML output.

It should be explained that in this service, because there is considerable potential for repetitive data structures within the delivered XML document, particularly when a profile’s *Feature Types* share similar *attributes* and *geometries*, XML's *ID* and *IDREF* tokenized attributes were used to reduce code replication (Morgenthal and Evdemon, 2000). *IDREF* is a special attribute that references a previously defined *ID* value (refer again to Figures 4.44 and 4.45). *IDREFS* is similar to *IDREF*, but can point to several previously defined *ID* values, each separated by a space.

A typical StreamlinedCatalogueService REST query has the following syntax:

…./StreamlinedCatalogueService/Catalogue/1/Profile/1/FeatureType/4/Format=XML

The service syntax for all three services is relatively straightforward. The service name is the first element of a REST query string (e.g., StreamlinedCatalogueService) followed by a “/”. Each query string can have only one service name. Next is the mandatory declaration of the “Catalogue” resource followed by a “/” and the Catalogue’s identifier (e.g., “1”), followed by another “/”. Other resource names follow the catalogue identification string, each separated by a “/”. If no specific resource identification number (e.g., “4”) follows the resource name, details for all resources of that type will be returned. The only exception to this rule (for XML and HTML output service versions) applies to the “Catalogue” resource which must have an identifier when used in the “StreamlinedCatalogueService” and in the “IterativeCatalogueService”.

In the “ComponentCatalogueService”, a description of which follows shortly, no identifier following the “/” after the “Catalogue” resource will provide details of all *Catalogue* instances. For the two services (StreamlinedCatalogueService and IterativeCatalogueService) that rely on requests that must always include a community *Profile* identifier, the sequence of resource names must follow an ordered pattern of hierarchical access (see previous Figure 4.43 and Figure 4.46, respectively).

The syntax used for the sample "StreamlinedCatalogueService" query listed earlier will deliver an XML document with content drawn from the *Catalogue* with ID=1, containing community *Profile* information (with ID=1) and the semantics of the *Feature Type* with ID=4. Using a "StreamlinedCatalogueService" with a nominated Catalogue ID, a client is currently able to:

- Get all *Features* for *Profile* n,
- Get all *Features* for all *Profiles*, and
- Get *Feature* n in *Profile* n.

Where "n" = an ID

Iterative Service

The "IterativeCatalogueService" is a traditional REST service that iteratively and progressively drills down (up to 4 levels) to expose resource information (see Figure 4.46). This service also requires that a *Catalogue* ID and a community *Profile* are provided as the first two parts of the service call string. The use-case supporting this type of service is one where there is a need to gradually expose details of individual resources contained in the catalogue and performance issues are not a concern. Of course it is always possible to navigate directly to a specific resource of interest if the URL is already known.

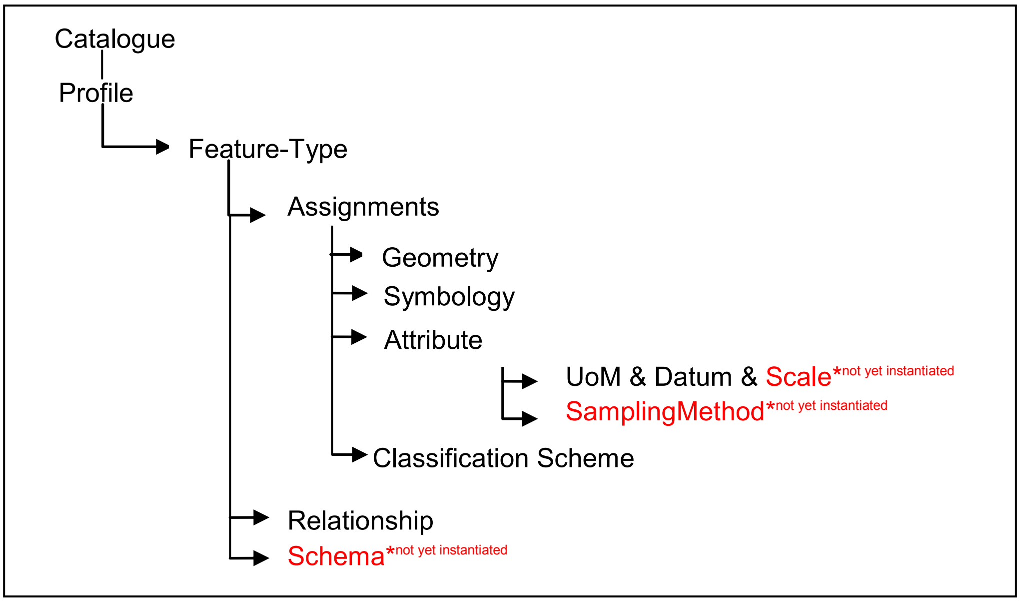

Figure 4.46 Resources available from the IterativeCatalogueService and access steps required.

A typical “IterativeCatalogueService” REST query might look like:

…/IterativeCatalogueService/Catalogue/1/Profile/1/Format=HTML

This service would deliver an HTML document containing a listing of all *Feature Types* in a *Profile* with ID=1, from a *Catalogue* with an ID=1. Each *Feature Type* listed has a hyperlink to another resource which contains more detailed information about that particular *Feature Type* (see Figure 4.47). A user is required to traverse successive embedded hyperlinks to obtain information on individual *Attributes*, *Units-of-Measure* (*UoM*) and *Datums*. Using an “IterativeCatalogueService”, having nominated the catalogue ID, a client is currently able to:

- Get *Profile* n or all *Profiles*,
- Get all *Features* for *Profile* n,
- Get all *Features* for all *Profiles*,
- Get *Feature* n in *Profile* n,
- Get *Attribute* n, or all *Attributes* for *Feature Type* n in *Profile* n,
- Get *Geometry* n, or all *Geometries* for *Feature Type* n in *Profile* n, and
- Get *UoM* n and associated *Datum*, or all UoMs for *Attribute* n in *Feature Type* n in *Profile* n

Catalogue Name: AADC Feature Catalogue
Catalogue Description : Catalogue of Antarctic Feature Types
Catalogue Version No: V1.0
Producer: AADC
Creation Date: 17-Jun-2009
Feature_Type Profile Description : A test profile, number 1
Feature_Type Profile Moderator : Kim Finney
Feature_Type Profile Moderator_Email : kim.finney@aad.gov.au
Profile User Name: AADC
Profile User Description: Australian Antarctic Data Centre
Profile User Contact Name: Ursula Harriss
Profile User Contact Email: Ursula.harriss@aad.gov.au

Feature Types

1 Aerial
2 Automatic_Weather_Station
3 Aircraft_Corridor
4 Aircraft_Runway
5 Ice_Field
6 Ice_Boundary
7 Ice_Sheet
8 Landing_Area
9 Mooring
10 Mountain
11 Mountain_Peak

Figure 4.47 HTML formatted sample code snippet for an IterativeCatalogueService

If the service had been constructed to also access *SamplingMethod* content this resource would have been hierarchically at the same level as *UoM* and *Datum* (see Figure 3.46) and would be accessible using the following syntax:

…/IterativeCatalogueService/Catalogue/1/Profiles/1/FeatureType/4/Attribute/5/SamplingMethod/1/Format=HTML

The *Schema* resource would be accessed in a similar manner, but commensurate with its place in the syntax hierarchy (as follows):

…/IterativeCatalogueService/Catalogue/1/Profiles/1/FeatureType/4/Schema/ Format=HTML

All *Schema* in which the *Feature Type* appears could be accessed (as in the sample above) or a *Schema* "ID" could be provided to list a specific *Schema*.

Component Service

The "ComponentCatalogueService" provides access to all component resources within the nominated *Catalogue*. It treats each class of resource as a separate entity and assumes nothing about the associations between *Catalogue* resources (i.e., it is agnostic of the relationship between *Profiles* and resources). The primary function of this service is to provide direct access to all resources managed within the *Catalogue*.

A typical "ComponentCatalogueService" REST query might look like:

…/ComponentCatalogueService/Catalogue/1/Relationship/5/Format=XML

This service would deliver an XML document containing semantic information for the *Relationship* resource with ID=5 from Catalogue with ID=1. Using a "ComponentCatalogueService" a client is currently able to:

- Get *Attribute* n or get all *Attributes*,
- Get *UoM* n or get all *UoM*s,
- Get *Geometry* n or get all *Geometries*,
- Get *Relationship* n or get all *Relationships*,
- Get *Datum* n or get all *Datums*,
- Get *Feature Type* n or get all *Feature Types*,
- Get *Profile* n or get all *Profiles*,
- Get *Classification Scheme* n or get all *ClassificationSchemes*, and

- Get *Catalogue* details.

4.7.4 REST-Based Feature Catalogue Service Patterns (SKOS Output)

Although SKOS is expressed in RDF, mapping the Feature Catalogue output into SKOS was a different type of task to serialising the database content into a suitable XML schema. To produce content in SKOS, a SKOS Concept Model had to be developed first that was complementary to the existing UML (Figure 4.17) and Relational (Figure 4.39) Feature Catalogue content Models. This was a similar type of task to that required for developing the OWL-based ontology model (although a simpler activity given that SKOS only has four types of pre-defined concept to leverage).

The resultant SKOS model is depicted in Figure 4.48 and is called FCATSKOS. SKOS, as is, was not semantically rich enough to fully express the Feature Catalogue content so the language concepts had to be specialised. This was undertaken by defining a SKOS extension scheme (Figure 4.49) which is then referenced along with SKOS core in any output files delivered by an "IterativeCatalogueService". The major entities in the RDBMS model have been specialised as SKOS concepts in FCATSKOS. SKOS, Dublin Core (DCMI, 2012) and VCard properties have been used to link from SKOS concepts to literals and other resources.

In the FCATSKOS model, *Instrument* information is related to the *SamplingMethod* concept via a specialised property called "operatedAccordingTo" (similar to FCATOWL). *SamplingMethod* is related to an *Attribute* through a "hasSensingMethod" property. But in contrast to FCATOWL there is a direct relationship between a *UoM* and an *Attribute* (through a "hasUoM" property). Recall in FCATOWL that *UoM* has a relationship with an *Attribute's Value* concept. Since it is not possible to reason using SKOS, the semantic definitions required (through property axioms) do not need to be so strict. Note also that *Attribute* uses the property "hasAttributeType" to link through to a concept called *AttributeType*. *AttributeType* holds the taxonomy of *QualitySpaces* as discussed in a previous section (describing aspects of FCATOWL). *Attribute* also has a relationship with *Scale*, which holds the same taxonomies as *ReferenceSpaces*, also discussed earlier.

Feature Type "relationships", other than those already defined in the SKOS extension schema, will need to be defined as they are needed and added into the extension file before they can used in service output. Ideally, identification of the need for any new relational properties should be performed as part of the normal data loading, maintenance and governance processes associated with adding content into the Catalogue. For example, if a *Feature Type* is added by a community member that has a relation "partOf" with another *Feature Type* in the Catalogue, a SKOS property "partOf" must be defined in the extension file.

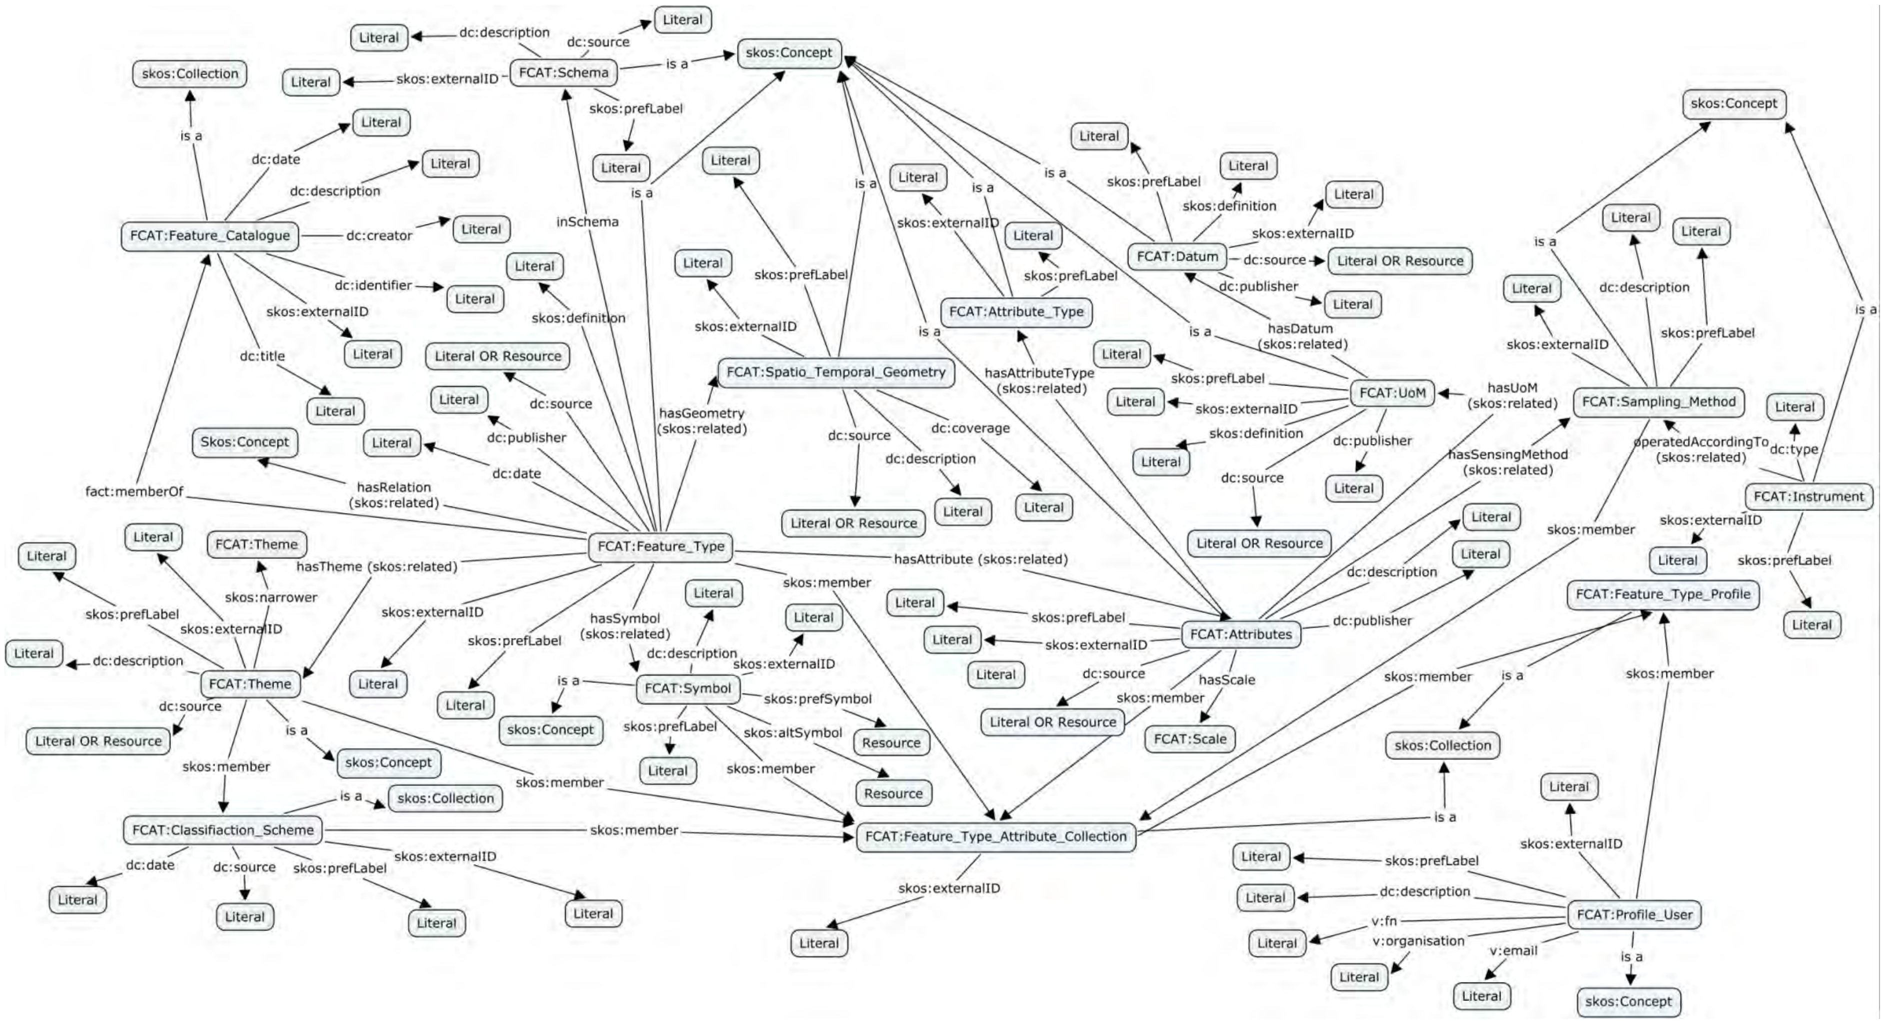

Figure 4.48 SKOS Feature Catalogue Content Model

A live demonstration of the “IterativeCatalogueService” (AADC, 2012d) capable of producing FCATSKOS output has been established. This is a different instantiation of the “IterativeCatalogueService” previously described, in part because the AADC programmer re-factored part of the code given knowledge gained from the previous exercise coding services producing the XML and HTML output. The author also wanted the previous work to remain stable, whilst different programming methods were explored to deliver the SKOS output. Therefore a cloned but edited service focussed only on SKOS output was instantiated. In an operational implementation of the “IterativeCatalogueService”, theoretically the only part of the address string which would need to change to receive Catalogue content in a SKOS format would be “/Format=SKOS” written at the end of the address string.

The data model delivered through the “IterativeCatalogueService” that provides the message response in SKOS is, however, sufficiently distinct that perhaps the SKOS and eventually OWL versions of the service should remain separate. The hREST service descriptions for these services (delivering SKOS and OWL outputs) could then leverage MicroWSMO (Kopecky and Vitvar, 2008), an extension of hRESTS that adds semantic annotations. Because the hRESTS view of services is similar to that of WSDL, it is also possible to adopt SAWSDL properties to add semantic annotations.

Recall that SAWSDL is an extension of WSDL that specifies how to annotate service descriptions with semantic information. By combining various elements of SAWSDL and MicroWSMO, a service description can be annotated with semantics. A SAWSDL “modelReference” attribute on an hREST “service” element can point to a description of the service’s functional and nonfunctional semantics; a modelReference on an hREST “operation” element points to the operation’s part of the behavioural semantic description; and a modelReference on an hREST “message” element points to the message’s counterpart(s) in the service’s information semantics ontology. SAWSDL annotations are URIs that identify semantic concepts and data transformations. These URIs can be added to the HTML documentation of RESTful services in the form of hypertext links (Kopecky *et al.,* 2009).

It is also possible to cast an hREST description in RDFa instead of microformat. An advantage of casting the hREST in RDFa is that processing microformats requires vocabulary-specific parsers (such as a custom XSLT transformation), while parsing the RDF data from RDFa is independent from any specific data vocabularies. Although the demonstration SKOS-based services have not been bundled with a service description (unlike those services developed in the first stage), Figure 4.50 shows a snippet of what the description would look like, by describing two hREST “operations” pertaining to the SKOS-based service (IterativeCatalogueService) using RDFa.

```xml
<?xml version="1.0"?>

<rdf:RDF
   xml:base="http://data.aad.gov.au/aadc/FeatureCatalogue/fcatskos-extension3#"
   xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
   xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
   xmlns:skos="http://www.w3.org/2004/02/skos/core#"
   xmlns:owl="http://www.w3.org/2002/07/owl#">
  <rdf:Property rdf:ID="memberOf">
            <rdfs:label>member Of</rdfs:label>
            <rdfs:isDefinedBy rdf:resource="http://data.aad.gov.au/aadc/FeatureCatalogue/fcatskos-extension3" />
            <skos:definition xml:lang="en">Relates a member to its collection.</skos:definition>
            <rdf:type rdf:resource="http://www.w3.org/1999/02/22-rdf-syntax-ns#Property"/>
            <rdfs:domain rdf:resource="http://www.w3.org/2004/02/skos/core#Concept"/>
            <rdfs:domain rdf:resource="http://www.w3.org/2004/02/skos/core#Collection"/>
           <!--  For OWL aware applications
                      <rdf:type rdf:resource="http://www.w3.org/2002/07/owl#ObjectProperty"/>  <owl:Class>
-                                <owl:unionOf rdf:parseType="Collection">
                                 <owl:Class rdf:about="http://www.w3.org/2004/02/skos/core#Concept"/>
                                 <owl:Class rdf:about="http://www.w3.org/2004/02/skos/core#Collection"/>
                                </owl:unionOf>
                       </owl:Class>
           </rdfs:domain>  -->
           <rdfs:range rdf:resource="http://www.w3.org/2004/02/skos/core#Collection"/>
           </rdf:Property>
  <rdf:Property rdf:ID="hasAttribute">
    <rdfs:label>hasAttribute</rdfs:label>
     <skos:definition>hasAttribute: a property that links a Feature Type concept and an Attribute concept.</skos:definition>
    <rdfs:range rdf:resource="#Attribute"/>
    <rdfs:domain rdf:resource="#Feature_Type"/>
    <rdfs:subPropertyOf rdf:resource="skos:related"/>
  </rdf:Property>
  <rdf:Property rdf:ID="hasAttributeType">
    <rdfs:label>hasAttributeType</rdfs:label>
    <skos:definition>hasAttributeType: a property that links an Attribute concept and an Attribute Type concept.</skos:definition>
    <rdfs:domain rdf:resource="#Attribute"/>
    <rdfs:range rdf:resource="#Attribute_Type"/>
    <rdfs:subPropertyOf rdf:resource="skos:related"/>
  </rdf:Property>
  <rdf:Property rdf:ID="hasDatum">
    <rdfs:label>hasDatum</rdfs:label>
    <skos:definition>hasDatum: a property that links a UoM concept and a Datum concept.</skos:definition>
    <rdfs:range rdf:resource="#Datum"/>
    <rdfs:domain rdf:resource="#UoM"/>
    <rdfs:subPropertyOf rdf:resource="skos:related"/>
  </rdf:Property>
  <rdf:Property rdf:ID="hasGeometry">
    <rdfs:label>hasGeometry</rdfs:label>
    <skos:definition>hasGeometry: a property that links a Feature Type concept and a Spatio_Temporal_Geometry
concept.</skos:definition>
    <rdfs:domain rdf:resource="#Feature_Type"/>
    <rdfs:range rdf:resource="#Spatio_Temporal_Geometry"/>
    <rdfs:subPropertyOf rdf:resource="skos:related"/>
  </rdf:Property>
  <rdf:Property rdf:ID="hasTheme">
    <rdfs:label>hasTheme</rdfs:label>
    <skos:definition>hasTheme: a property that links a Feature Type concept and an Theme concept.</skos:definition>
    <rdfs:domain rdf:resource="#Feature_Type"/>
    <rdfs:range rdf:resource="#Theme"/>
    <rdfs:subPropertyOf rdf:resource="skos:related"/>
  </rdf:Property>
  <rdf:Property rdf:ID="hasUoM">
    <rdfs:label>hasUoM</rdfs:label>
     <skos:definition>hasUoM: a property that links an Attribute concept and an UoMconcept.</skos:definition>
    <rdfs:domain rdf:resource="#Attribute"/>
    <rdfs:range rdf:resource="#UoM"/>
    <rdfs:subPropertyOf rdf:resource="skos:related"/>
  </rdf:Property>
```

```xml
 <rdf:Property rdf:ID="hasSymbol">
   <rdfs:label>hasSymbol</rdfs:label>
    <skos:definition>hasSymbol: a property that links a Feature Type concept and a Symbol concept.</skos:definition>
   <rdfs:domain rdf:resource="#Feature_Type"/>
   <rdfs:range rdf:resource="#Symbol"/>
   <rdfs:subPropertyOf rdf:resource="skos:related"/>
</rdf:Property>
 <rdf:Property rdf:ID="hasMethod">
   <rdfs:label>hasMethod</rdfs:label>
    <skos:definition>hasMethod: a property that links a Sampling Method concept and an Attribute concept.</skos:definition>
   <rdfs:domain rdf:resource="#Sampling_Method"/>
   <rdfs:range rdf:resource="#Attribute"/>
   <rdfs:subPropertyOf rdf:resource="skos:related"/>
</rdf:Property>
 <rdf:Property rdf:ID="hasSchema">
   <rdfs:label>hasSchema</rdfs:label>
    <skos:definition>hasSchema: a property that links a Schema concept and a Feature Type concept.</skos:definition>
   <rdfs:domain rdf:resource="#Schema"/>
   <rdfs:range rdf:resource="#Feature_Type"/>
   <rdfs:subPropertyOf rdf:resource="skos:related"/>
</rdf:Property>
<rdf:Property rdf:ID="hasInstrument">
   <rdfs:label>hasInstrument</rdfs:label>
    <skos:definition>hasInstrument: a property that links a Sampling Method concept and an Instrument concept.</skos:definition>
   <rdfs:domain rdf:resource="#Sampling_Method"/>
   <rdfs:range rdf:resource="#Instrument"/>
   <rdfs:subPropertyOf rdf:resource="skos:related"/>
</rdf:Property>
<rdfs:Class  rdf:ID="Attribute">
   <rdfs:label>Attribute</rdfs:label>
   <skos:definition>An Attribute concept: an inherent characteristic of a Feature Type.</skos:definition>
   <rdfs:subClassOf rdf:resource="skos:Concept"/>
</rdfs:Class >
<rdfs:Class  rdf:ID="Attribute_Type">
   <rdfs:label>Attribute_Type</rdfs:label>
   <skos:definition>An Attribute_Type concept: a description drawn from a controlled list of terms that classifies the type of an inherent characteristic (attribute) of a feature type concept.</skos:definition>
   <rdfs:subClassOf rdf:resource="skos:Concept"/>
</rdfs:Class >
<rdfs:Class  rdf:ID="Scale">
   <rdfs:label>Scale</rdfs:label>
   <skos:definition>A Scale concept: a description drawn from a controlled list of terms that classifies the scale (e.g., ordinal, ratio, interval) inherent in an attribute of a feature type concept.</skos:definition>
   <rdfs:subClassOf rdf:resource="skos:Concept"/>
</rdfs:Class >
<rdfs:Class  rdf:ID="Classification_Scheme">
   <rdfs:label>Classification_Scheme</rdfs:label>
   <skos:definition>A Classification Scheme concept: an arrangement or division of objects into groups based on characteristics which the objects have in common.</skos:definition>
   <rdfs:subClassOf rdf:resource="skos:Collection"/>
</rdfs:Class >
<rdfs:Class  rdf:ID="Datum">
   <rdfs:label>Datum</rdfs:label>
         <skos:definition>A Datum concept: a set of parameters or control points that provide a reference framework for grounding measurements.</skos:definition>
   <rdfs:subClassOf rdf:resource="skos:Concept"/>
</rdfs:Class >
<rdfs:Class  rdf:ID="Feature_Catalogue">
   <rdfs:label>Feature_Catalogue</rdfs:label>
   <skos:definition>A Feature Catalogue concept collection: a collection that contains Feature Types, Attributes, Datums, UoMs, User_Profiles, Themes, Classification Schemes, Feature Type Profiles and Feature Type Attribute Collections.</skos:definition>
   <rdfs:subClassOf rdf:resource="skos:Collection"/>
</rdfs:Class >
<rdfs:Class  rdf:ID="Feature_Type">
   <rdfs:label>Feature_Type</rdfs:label>
   <skos:definition>A feature type concept: any real-world object.</skos:definition>
   <rdfs:subClassOf rdf:resource="skos:Concept"/>
</rdfs:Class >
<rdfs:Class  rdf:ID="Feature_Type_Attribute_Collection">
   <rdfs:label>Feature_Type_Attribute_Collection</rdfs:label>
```

```
      <skos:definition>Feature_Type_Attribute_Collection: a grouping of Feature Types, Attributes, Datums, UoMs, User_Profiles, Themes, Classification Schemes.</skos:definition>
      <rdfs:subClassOf rdf:resource="skos:Collection"/>
   </rdfs:Class >
   <rdfs:Class  rdf:ID="Feature_Type_Profile">
      <rdfs:label>Feature_Type_Profile</rdfs:label>
      <skos:definition>A Feature Type Profile concept collection: a grouping of Feature Attribute Collections that are owned by a Profile_User.</skos:definition>
      <rdfs:subClassOf rdf:resource="skos:Collection"/>
   </rdfs:Class >
   <rdfs:Class  rdf:ID="Profile_User">
      <rdfs:label>Profile_User</rdfs:label>
      <skos:definition>A Profile User concept: an organisation, individual, community of practise or on-line application that owns a Feature Type Profile.</skos:definition>
      <rdfs:subClassOf rdf:resource="skos:Concept"/>
   </rdfs:Class >
   <rdfs:Class  rdf:ID="Spatio_Temporal_Geometry">
      <rdfs:label>Spatio_Temporal_Geometry</rdfs:label>
      <skos:definition>A spatio-temporal-geometry concept: a description drawn from a controlled list of terms that defines the spatial and/or temporal nature of a feature type concept.</skos:definition>
      <rdfs:subClassOf rdf:resource="skos:Concept"/>
   </rdfs:Class >
   <rdfs:Class  rdf:ID="Symbol">
      <rdfs:label>Symbol</rdfs:label>
         <skos:definition>A Symbol concept: something that represents something else by association, resemblance, or convention.</skos:definition>
      <rdfs:subClassOf rdf:resource="skos:Concept"/>
   </rdfs:Class >
   <rdfs:Class  rdf:ID="Theme">
      <rdfs:label>Theme</rdfs:label>
      <skos:definition>A Theme concept: a distinct and recurring subject that is part of a Classification Scheme.</skos:definition>
      <rdfs:subClassOf rdf:resource="skos:Concept"/>
   </rdfs:Class >
   <rdfs:Class  rdf:ID="UoM">
      <rdfs:label>UoM</rdfs:label>
      <skos:definition>A Units of Measure concept: any division of quantity accepted as a standard of measurement or exchange.</skos:definition>
      <rdfs:subClassOf rdf:resource="skos:Concept"/>
   </rdfs:Class >
   <rdfs:Class  rdf:ID="Sampling_Method">
      <rdfs:label>Sampling_Method</rdfs:label>
      <skos:definition>A Sampling Method concept: any method or process used to sample an Attribute concept.</skos:definition>
      <rdfs:subClassOf rdf:resource="skos:Concept"/>
   </rdfs:Class >
   <rdfs:Class  rdf:ID="Instrument">
      <rdfs:label>Instrument</rdfs:label>
      <skos:definition>An instrument concept: any Instrument (usually a collection of components) used to sample a property.</skos:definition>
      <rdfs:subClassOf rdf:resource="skos:Concept"/>
   </rdfs:Class >
   <rdfs:Class  rdf:ID="Schema">
      <rdfs:label>Schema</rdfs:label>
      <skos:definition>ASchema concept: a data encoding schema expressed using formal notation.</skos:definition>
      <rdfs:subClassOf rdf:resource="skos:Concept"/>
   </rdfs:Class >
</rdf:RDF>
```

Figure 4.49 The SKOS Schema Definition Extension File

```
<div typeof="wsl:Service" about="#svc"
xmlns:hr="http://www.wsmo.org/ns/hrests#"
xmlns:wsl="http://www.wsmo.org/ns/wsmo-lite#"
xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#">
<span property="rdfs:label"><strong>SKOSIterativeFeatureCatalogueService</strong></span>
<p>This service is a <a rel="model"
href="http://data.aad.gov.au/aadc/FeatureCatalogue/index2.cfm/IterativeCatalogueService">SKO
S Feature Type"</a> service.</p>
<div rel="wsl:hasOperation"><span typeof="wsl:Operation" about="#op1"><br/>
<code property="rdfs:label"> 1. GET Catalogue n</code><br/>
To make a query to retrieve individual catalogues from FCAT using the http method <span
property="hr:hasMethod">GET</span>, use the request <code property="hr:hasAddress"
datatype="hr:URITemplate">
http://data.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/{id}</code
> <br/> </p>
<p><span rel="wsl:hasInputMessage"><span typeof="wsl:Message"> The parameter:<a
rel="model" href="http://data.aad.gov.au/aadc/FeatureCatalogue/fcatskos-
extension#Catalogue"></a> <code>id</code> - (a positive integer) is the identifier of a particular
Catalogue </span></span><br/></p>
<p><span rel="wsl:hasOutputMessage"><span typeof="wsl:Message">The output format for the
request provides Catalogue information encompassing links to user profiles. The response is
supplied in RDF incorporating echema elements
from:xmlns:fcat="http://data.aad.gov.au/aadc/FeatureCatalogue/fcatskos-extension#"
xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-
ns#"xmlns:dc="http://purl.org/dc/elements/1.1/"
xmlns:skos="http://www.w3.org/2004/02/skos/core#"
xmlns:v="http://www.w3.org/2006/vcard/ns#" schema documents.</span></span></p>
</div></div>
<div rel="wsl:hasOperation"><span typeof="wsl:Operation" about="#op2"><br/>
<code property="rdfs:label"> 1. GET Profile n in Catalogue n</code><br/>
To make a query to retrieve individual profiles from the FCAT catalogue using the http method
<span property="hr:hasMethod">GET</span>, use the request <code property="hr:hasAddress"
datatype="hr:URITemplate">
http://data.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/{id}/Profile
{id}</code> <br/> </p>
<p><span rel="wsl:hasInputMessage"><span typeof="wsl:Message"> The parameter:<a
rel="model" href="http://data.aad.gov.au/aadc/FeatureCatalogue/fcatskos-
extension#Profile"></a> <code>id</code> - (a positive integer) is the identifier of a particular
Catalogue </span></span><br/></p>
<p><span rel="wsl:hasOutputMessage"><span typeof="wsl:Message">The output format for the
request provides Profile information encompassing links to Feature Types. The response is supplied
in RDF incorporating schema elements
from:xmlns:fcat="http://data.aad.gov.au/aadc/FeatureCatalogue/fcatskos-extension#"
xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-
ns#"xmlns:dc="http://purl.org/dc/elements/1.1/"
xmlns:skos="http://www.w3.org/2004/02/skos/core#"
xmlns:v="http://www.w3.org/2006/vcard/ns#" schema documents.</span></span></p>
</div></div>
```

Figure 4.50 Snippet of SKOS IterativeService Description in hREST RDFa

Differences Between IterativeCatalogueService Instantiations (XML/HTML vs SKOS)

In developing the RDF serialisation of content commensurate with the SKOS model constructs an issue arose concerning how the "IterativeCatalogueService" had been made to work in delivering XML output and how it should work in delivering SKOS output. In creating the XML/HTML output schema patterns for all services, inadequate consideration was given to ensuring that a REST query was able to deliver the most granular representation of a resource. For example, consider the case where a *Feature Type* with two different types of (spatio-temporal) *Geometry* (i.e., a line representation and a polygon) and with different sets of *Attributes*, is accessed from the "IterativeCatalogueService" developed earlier (for XML output) using the following syntax:

…/IterativeCatalogueService/Catalogue/1/Profile/1/FeatureType/4/Format=HTML

The resultset sent to the client includes all of the information for *Feature Type* with ID=4, by showing those *attributes* that are associated with the *Feature Type* when it has *geometry* types "line" and "polygon". This is achieved by using an XML "Assignments" element that has no corresponding value in the underlying database. It is just an artificial construct used in the XML/HTML output to differentiate between the two sets of *attribute* assignments that a *Feature Type* can have if it has more than one *geometry* type. See the sample output reproduced in Figure 4.51.

Catalogue Name: AADC Feature Catalogue

Description	The best Feature Catalogue in Australia managed by the AADC
Catalogue Version No	1
Producer	Australian Antarctic Data Centre
Creation_Date	17-Jun-2009

Feature_Type id="4"

Feature_Type_Name : Aircraft_Runway
Feature_Type_Definition : A surface on which aircraft can land or take-off from
Definition_Authority : kim
Definition_Source id="51" http://aadc.kims.wisdom
Assignments id="1"
    Feature_Geometry Geometry_Id="5" geometry_name="Line"
    Feature_Attribute Attribute_Id="3" attribute_name="Name"
    Feature_Attribute Attribute_Id="4" attribute_name="Latitude"
    Feature_Attribute Attribute_Id="5" attribute_name="Longitude"
    Feature_Attribute Attribute_Id="11" attribute_name="Runway_Type"
    Classification_Scheme id="1" scheme_name="TestScheme1"
        Assigned_Theme id="2" theme_name="Infrastructure"
Assignments id="2"
    Feature_Geometry Geometry_Id="4" geometry_name="Polygon"
    Feature_Attribute Attribute_Id="3" attribute_name="Name"
    Feature_Attribute Attribute_Id="4" attribute_name="Latitude"
    Feature_Attribute Attribute_Id="5" attribute_name="Longitude"
    Feature_Attribute Attribute_Id="11" attribute_name="Runway_Type"

Classification_Scheme id="1" scheme_name="TestScheme1"
Assigned_Theme id="2" theme_name="Infrastructure"
Relationship Id="1" relationship_name="part_Of" Object_Feature_type Id 3

Figure 4.51 Sample HTML Output Format For a Feature Type With Two Geometries.

On further reflection, using this approach wasn't necessarily the best way forward because it means a REST query that is translated and sent to the database will return all forms of a given *Feature Type* when really the user may only want the form of the *Feature Type* with a specific type of *Geometry*. The solution to this issue, which was adopted in the "IterativeCatalogueService" designed to deliver SKOS output, was to add in a convenience resource called the *Feature_Type_Attribute_Collection* which is essentially an associative entity that groups a *Feature Type*, its specific *Geometric* representation and associated *Attributes*. As with the XML "Assignments" element, discussed earlier, there is no 'ID' that exists for this resource in the database (because it is not an entity), therefore the service program code assigns an "ID" to a *Feature_Type_Attribute_Collection* using two numbers, i.e., a *Feature Type* resource ID in combination with a *Geometry* resource ID, separated by a "_". In the example syntax below:

.... /Catalogue1/Profile/1/Feature_Type_Attribute_Collection/7_4/ ...

"7" is the *Feature Type* ID and "4" is the *Geometry* ID joined together to form an ID for the *Feature_Type_Attribute_Collection* resource. So, a query of the form:

.... /Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/7_4/ Feature_Type/7/Format=SKOS

would return output for *Feature Type* ID=7, but only show those attributes (and other resources) that have been assigned when the *Geometry* resource, associated with *Feature Type* ID=7, has an ID=4. The query below also yields exactly the same result as the REST query above:

...Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/7_4/ Feature_Type/7/Geometry/ 4/Format=SKOS...

However, if a query is issued that only includes the *Feature_Type_Attribute_Collection* resource (with no following *Feature Type* or *Geometry* resource declarations) the information is delivered in a slightly different format (see sample output in Appendix 13) because a *Feature_Type_Attribute_Collection* is of SKOS type: *Collection* (which by the SKOS design specification has a different pattern of expression to SKOS *Concept* types).

In contrast, output received as a consequence of using the *Feature Type* and/or *Geometry* resources in the query string following the *Feature_Type_Attribute_Collection resource declaration* will use a SKOS "Concept" expression pattern (see Appendix 14 for this type of sample output).

In hindsight, the method used to provide greater granularity in this (SKOS) version of the "IterativeCatalogueService" should also have been used in the previous service instantiation delivering XML and HTML. In a fully developed solution the approach should be the same so that there is only one uniform model for "IterativeCatalogueService" operations, regardless of output type. Apart from the introduction of the *Feature_Type_Attribute_Collection* resource as a legal part of the query syntax (and a queryable resource in its own right) the "IterativeCatalogueService", does however, function as explained previously.

This SKOS-based (Iterative) Service should also exercise the *SamplingMethod, Instrument* and *Schema* resource components (associated with the *Attribute* and *FeatureType* resources, respectively) which were also omitted from the previous prototyping activity. If these resources were included the FCATSKOS (Iterative) service, the service would be capable of performing the following:

- Get profile n or all profiles,
- Get all features for profile n,
- Get all features for all profiles,
- Get feature n in profile n,
- Get attribute n, or all attributes for feature type n in profile n,
- Get geometry n, or all geometries for feature type n in profile n, and
- Get UoM n and associated Datum, or all UoMs for Attribute n in Feature Type n in Profile n
- Get SamplingMethod n and associated Instrument, or all SamplingMethods (and Instruments) for Attribute n in Feature Type n in Profile n

Although robust service prototypes for demonstrating SKOS output from the Component and StreamlinedCatalogueServices were not implemented, Appendices 15-16 lists sample SKOS output patterns for both of these service types.

4.8 Feature Catalogue Design and Development Results Summary

In this chapter the results of all research investigations into RQ1.1 – "*What characterises an ontologically-grounded Feature Catalogue that can support Antarctic science data publication through Web services ?*" and its related sub-questions have been presented.

It has been demonstrated that many of the Antarctic datasets exchanged by the communities participating in this study conform to a basic observation and measurement (O&M) pattern. In this pattern Feature Types can be any of: 'Observation Events' – which are always complex (i.e., multi-featured) Feature Types; 'domain features-of-interest'; 'Sampling Features' and/or 'Sampled Features' (all of which could also be complex Feature Types). The basic O&M pattern, as presented by Cox (2006, 2010a) requires some augmentation and specialisation in order to cover the components commonly found in Antarctic-themed datasets. Biological Feature Types have been shown to present different modelling challenges to Feature Types that model purely physical phenomena. Most importantly, an observation-centric dataset contains Feature Types whose complete semantic signature requires reference to both its spatio-temporal footprint and its sampling context (if the purpose of the semantic definition is to establish Feature Type equivalence for use in data integration exercises).

Through reference to focus groups, drawn from communities-of-interest who have participated in this study, it was possible to establish the functional requirements that a Feature Catalogue should meet. A key requirement was that the Catalogue be able to interface easily and seamlessly with the existing scientific infrastructure already developed by these groups. Other requirements mentioned, through use-case formulation, helped define the necessary conceptual content model of the Feature Catalogue. Combined with the data modelling investigations, the use-case formulation which was aided by community participants provided an answer to RQ1.1.1 – "*What type of use-cases and data models must the Feature Catalogue support ?".*

RQ1.1.2 asked "*Do relevant ISO and OGC conceptual data models meet implementation needs, and if not what is missing ?".* To address this question, a review of the current ISO 19110 *Feature Cataloguing Methodology* standard was undertaken in the context of its relationship to companion standards such as ISO 19109 (*Rules For Application Schema*) and 19126 (*Feature Concept Dictionaries and Registers*). Observations regarding all three standards were made, with a focus on determining whether ISO 19110, as is, met community requirements. ISO 19110 was found lacking in several areas resulting in the development of an enhanced ISO 19110 Feature Catalogue model, more able to meet community expectations. This enhanced model was trialled through a desk-top exercise to examine how a hypothetical dataset would be encoded using the enhanced Catalogue model.

The ISO *Feature Cataloguing Methodology* prescribes textual Catalogue descriptions for Catalogue resources, which are insufficient for either semi-formally (e.g., formalisms that can be attained by using RDFS) or formally (e.g., formalisms attained by using OWL) defining terms. The enhanced (ISO-based) Feature Catalogue conceptual model was therefore ontologically grounded (to attain the required level of formalism) by casting the conceptual model in terms of concepts and properties defined by DOLCE (an upper ontology of particulars), using a number of relevant ontology design patterns (also anchored in DOLCE) and described using OWL. In this OWL-based ontological model, Feature Types were conceived of as specialisations of "WebResource" concepts that in DOLCE-based ontology design patterns are 'information entities' representing 'real-world objects'. The ontological prototype model was populated with sample data to demonstrate the encoding of observation-centric Feature Types. The outcome was a conceptual model that was ontologically grounded and which met Antarctic community needs (RQ1.1.3).

In the absence of guidance from ISO and the OGC, prototyping activity was also undertaken to develop suitable access methods to Feature Catalogue content. This activity directly addressed RQ 1.1.4 – "*What methods are best suited to extract re-usable content from an ontologically-grounded Feature Catalogue ?".* REST-based services were established in preference to other available access options because it was argued that they presented the lowest barrier to uptake; provided the broadest access to Catalogue content (using current browser and internet technology) and were consistent with approaches being adopted for accessing content in ontology repositories more generally. Service development required creation of service query patterns such that each query would provide a unique URL for the Catalogue resources delivered. An enhanced version of SKOS (an informal ontology expressed using RDFS, usually used for classification systems) was trialled as one type of service output (the other types being HTML and XML).

Lessons learned along the investigative journey, mainly related to ontology development, were captured so that they could be used as practical guidance for communities wishing to embark on semantic annotation activities. Observations made will also be used next, in Chapter 5, to evaluate the work performed in this chapter, particularly in relation to community (and infrastructure) readiness to embrace semantic dataset annotation, facilitated by community-based, ontologically-grounded Feature Catalogues.

Chapter 5.

Feature Catalogue Evaluation and Discussion

The previous chapter described the process of developing a semantically-enabled Feature Catalogue (instantiated using OWL, RDBMS and SKOS data models). The design was grounded by practical requirements gathered from two large, multi-disciplinary scientific communities and a detailed review of their data modelling needs in data exchange scenarios. In conducting this research relevant standards were also evaluated in order to assess their suitability for supporting Feature Type Catalogue development. A number of standards deficiencies were noted.

The prototyping activity also involved an evaluation of two primary methods for accessing Catalogue resources: SPARQL and a REST-based interface. Since there are no existing implementation standards for creating REST interfaces to Feature Catalogue repositories, a model for how this can be achieved was demonstrated. Taken together, it is claimed here that the ISO 19110 enhanced, ontologically-grounded Feature Catalogue (FCAT) model and the associated REST-based access methods can provide a flexible set of tools for semantically annotating datasets that are typically exchanged within Antarctic-themed science.

At the start of the previous chapter, four factors were stated as being crucial considerations for development of the semantic Catalogue and its access interface. These high level considerations (or goals) were abstracted by the author from discussions with participating communities. Restated these considerations were that:

1. The Catalogue had to fit architecturally as a component within the evolving infrastructures of both Antarctic-connected communities participating in this study and incorporate (or interoperate with) any standards already adopted.
2. Any developed access interfaces to the Catalogue had to demonstrably support dataset annotation methods capable of being executed within current Antarctic scientific data infrastructure.
3. The Catalogue content model had to accommodate the types of data routinely captured and exchanged.
4. The content model also had to provide an appropriate semantic signature for included (Feature Type) concepts (i.e., the model should provide enough semantic content regarding a Feature Type to enable discrimination between datasets containing dissimilar/similar Feature Types for dataset integration purposes in the current SCAR and AODN infrastructure. Ideally,

this semantic signature should also be interoperable with existing semantic description frameworks already in use by more advanced scientific communities exchanging observation-centric data).

Data modelling exercises in the previous chapter were used to formatively evaluate goal ‘3’, and to a large extent goal ‘4’, using inputs and reflections from community focus groups. What has not yet been demonstrated and evaluated is whether goals ‘1’ and ‘2’ have been satisfied and the extent to which the semantic signature that is expressed through the Feature Type Catalogue ontological model is able to be used by other related semantically-enabled infrastructure (as stated in goal ‘4’).

In this chapter (in section 5.1) a desk-top evaluation of the SCAR and AODN infrastructure is conducted to determine whether the artefacts developed could be readily applied in these environments. This permits an evaluation of whether goals ‘1’, ‘2’ and parts of ‘4’ have been met by the work performed in Chapter 4 (conducted in order to answer RQ 1.1 *What characterises an ontologically-grounded Feature Catalogue that can support Antarctic science data publication through Web services*”). The reported facets of evaluation taken together build up to a descriptive, theoretical assessment of the ‘goodness of architectural fit’ between the artefacts developed in this research, existing Antarctic infrastructure and related connected systems (whose main focus is on exchanging observation-centric, spatially-enabled datasets).

‘Architectural fit’ is judged by the degree to which developed artefacts use, or conform to, various standards and design patterns used by the respective participating communities and by the functional interoperability of the developed artefacts with technological and content components of their scientific infrastructure. ‘Architectural fit’ is also evaluated with respect to artefact interoperability with a selection of the more semantically-advanced systems with which the SCAR and AODN communities intersect (i.e., connect).

No evaluation would be complete without a critical look at the limitations of the artefacts developed. Section 5.2 discusses how aspects of FCAT and its associated REST-based service access methods could be improved upon. This facet of the evaluation holds particular significance as guidance for communities wishing to take the prototyping activity in this thesis through to full scale implementation and areas requiring further research are flagged.

Section 5.3 of this chapter continues with the guidance theme and concludes with a discussion of ‘issues of practise’ that impact upon a community-sponsored Feature Catalogue build-process beyond prototyping, taking into consideration the ‘state of readiness’ of the communities concerned with respect to adopting semantic technologies.

5.1 'Architectural Fit' With Existing SCAR and AODN Infrastructure

The basic architectures of the SCAR and AODN infrastructures are depicted in Figures 5.1 and 5.2. Both are service-oriented architectures, although the SCAR community does not currently use the Global Change Master Directory (GCMD) as an OGC-CSW compliant service registry. The GCMD Directory has CSW (interface standard) capability but it is only providing this functionality at present when it acts as a node of the Global Earth Observation System (GEOSS (GEO, 2012)). The CSW-compliant service interface is therefore currently not accessible to SCAR members.

National SCAR participants mainly interact with the GCMD (Portal) in its capacity as an advanced dataset metadata system (rather than as a services registry). Services can, however, be registered with separate service metadata (called Services Entry Resource Format -SERF records (NASA, 2012b)) and a number of nations are now using this facility. Registered services are discovered within the GCMD Portal using SERFs and can then be linked to via URL. The services advertised in the GCMD must be directly related to the processing, viewing, analysis, archival, retrieval, production, interpretation, discovery, acquisition, modelling, visualizing, formatting, or indexing of Earth science data.

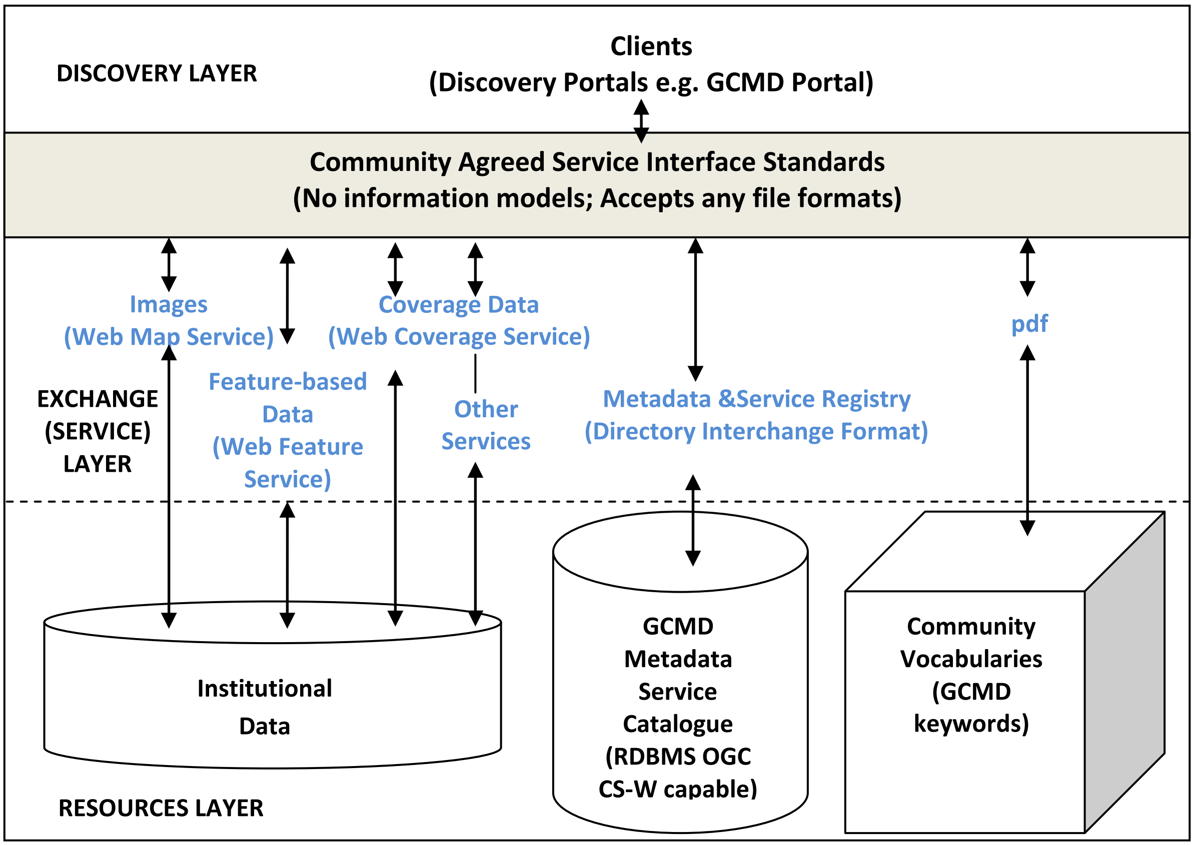

Figure 5.1 Basic Infrastructure Components Available to SCAR Community.

The scope and variety of service protocols allowable is therefore extremely broad. This contrasts with the AODN infrastructure, where the supported protocols are more restricted.

The GCMD Portal manages metadata in the Data Interchange Format (DIF (NASA, 2012c)), a format unique to the GCMD but which is mappable to the ISO 19115 metadata standard. A <Related_URL> field in the DIF enables the linking of datasets to the DIF metadata record (via URL). This field uses a set of sub-fields with controlled terms to describe the URL content. Many of these terms describe service types (e.g., 'Access Map Viewer'; 'GET Map Service (non OGC)'; 'GET Web Feature Service (WFS)'; 'GET Web Map Service (WMS)'). Each set of terms is accompanied by a plain text description. So, using this field it is also possible to link services directly to a dataset metadata record without first having created a SERF (i.e., service metadata record). Antarctic community members who use both the GCMD and AODN infrastructure use XSLT (Clark, 1999) translators to move between the DIF and ISO 19115 metadata standards.

The main vocabulary used by SCAR community members is the GCMD Science Keywords (Olsen *et al.,* 2007) This is a multi-level term vocabulary which is moderated across the first five hierarchy levels but which allows users to add their own specialised terms at the lower levels (i.e., in an uncontrolled manner).

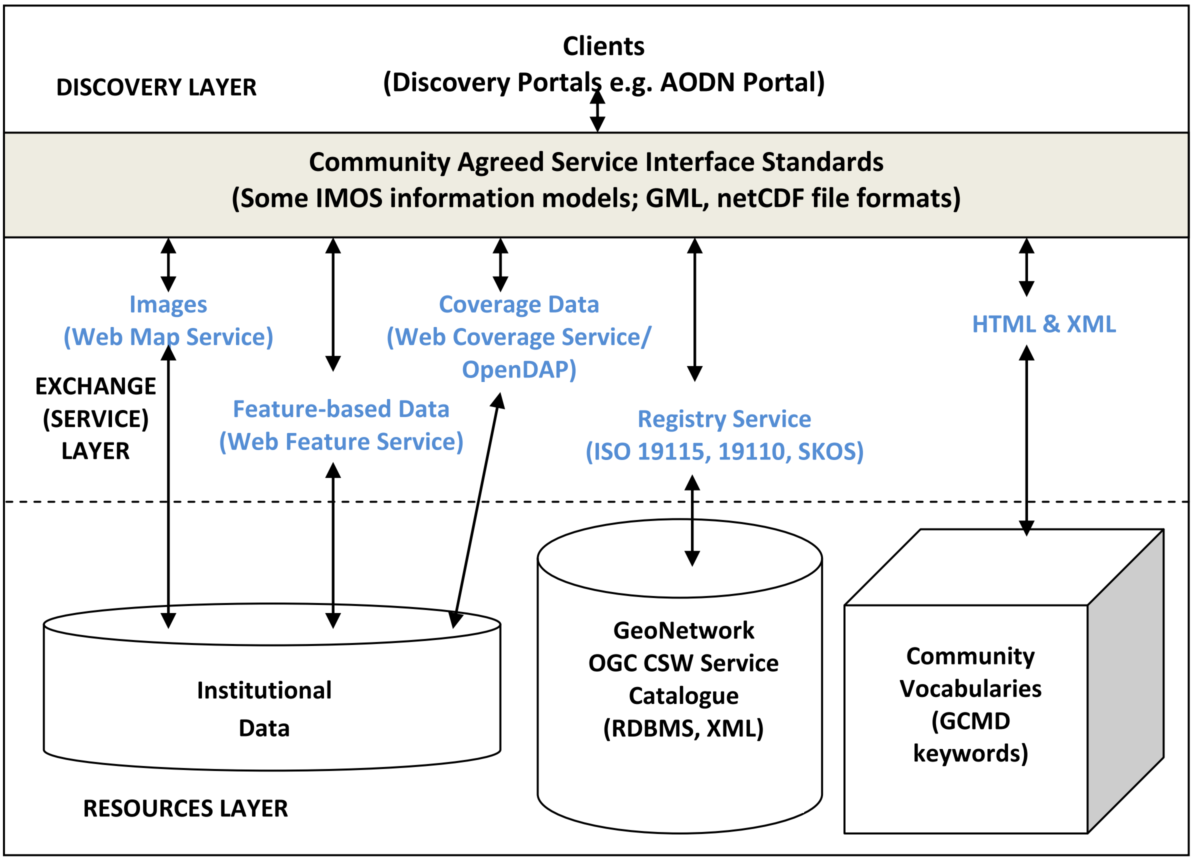

Figure 5.2 Basic Infrastructure Components Available to AODN Community.

These keywords are used within the GCMD metadata and for most GCMD Portal discovery services. They are a mixture of theme, parameter and material object type terms (without clear distinction).

The AODN infrastructure uses a customised version (i.e., the Metadata Entry and Search Tool - MEST) of GeoNetwork for its metadata registry. Unlike the GCMD registry, GeoNetwork is an open source product. GeoNetwork harnesses RDBMS (and XML) data stores (GeoNetwork, 2011). Some communities using GeoNetwork also harness the Sesame (semantic store (Aduna, 2012)) to persistently manage thesauri which can then be exported (using RDF files) and used within GeoNetwork as SKOS thesauri. It is also feasible to create local thesauri in GeoNetwork but the interface limits users to identifying a concept with an identifier, a preferred label, a definition and links to other concepts. Three types of links are used in GeoNetwork: SKOS related; broader and narrower terms. The AODN currently makes almost no use of this facility.

The AODN Portal is configured to discover and use OGC-compliant Web services and is also a DAP-enabled client. DAP (or Data Access Protocol (OPeNDAP, 2012)) is a protocol designed to hide the implementation of different collections of data using an interface based on the 'name-datatype value' conceptual model (Gallagher *et al.,* 2007). It is particularly popular in the climate and oceanographic sciences, but not often used in biological disciplines.

The AODN community has subscribed to a profile of ISO 19115 called the Marine Community Profile (MCP (AODCJF, 2008)). A 'Profile' of a standard is essentially an extension to that standard that is still conformant. Vocabularies used in this profile (within the ISO 19115 'MD_Keywords' element) are currently drawn from the GCMD Science Keywords. GeoNetwork provides a custom widget to expose and define internal links between dataset metadata and associated services. These internal associations (recorded in the underlying RDBMS) between dataset metadata and service metadata are achieved through a coupling of the dataset metadata Id and service Ids. In this widget a list of OGC service (layers) associated with dataset metadata is automatically generated if there is a GetCapabilities document URL advertised (in the metadata). If no GetCapabilities URL has been provided, a service (layer) name can be manually recorded against a given metadata record. Within the dataset metadata record there is also provision for recording a direct link to an 'online resource' which can be a link to a service using an ISO 19115 'CI_OnlineResource' metadata sub- element (see Figure 5.3 for a snippet of a metadata record pointing to a Web Map Service):

```
<gmd:onLine>
        <gmd:CI_OnlineResource>
                <gmd:linkage>
                        <gmd:URL>http://geonetwork3.aodn.org/ows/296</gmd:URL>
                </gmd:linkage>
                <gmd:protocol>
                <gco:CharacterString>OGC:WMS1.1.1httpgetmap</gco:CharacterString>
                </gmd:protocol>
        </gmd:CI_OnlineResource>
<gmd:onLine>
```

Figure 5.3 Online Link From Dataset Metadata To A WMS Service In GeoNetwork.

GeoNetwork also provides a facility to link a metadata dataset record to a cited Feature Catalogue (conformant with ISO 19110) which must be an instance of a Feature Catalogue stored in the same GeoNetwork instance as the metadata record. This is achieved via linking a Feature Catalogue Id (as allocated by GeoNetwork) to a Metadata Record Id (as generated by GeoNetwork) in the underlying RDBMS. Most recently this relationship can also be referenced (and recognised) within GeoNetwork by using sub-elements of an ‘MD_ContentInformation’ metadata element, within the dataset metadata record, using the pattern shown in Figure 5.4 (assuming that the ISO metadata profile being used supports the MD_ContentInformation’ element). As will be demonstrated later, this ‘within metadata’ linking mechanism does not harness the full capacity of the ‘MD_ContentInformation’ metadata element and has some limitations.

```
<gmd:contentInfo>
        <gmd:MD_FeatureCatalogueDescription>
                <gmd:includedWithDataset>
                        <gco:Boolean>false</gco:Boolean>
                </gmd:includedWithDataset>
                <gmd:featureCatalogueCitation uuidref="22a7c5a0-3bea-4e53-9356-
83d27d333cd8"
xlink:href="http://localhost:8080/geonetwork/srv/en/metadata.show?uuid=22a7c5a0-3bea-
4e53-9356-83d27d333cd8"/>
        </gmd:MD_FeatureCatalogueDescription>
</gmd:contentInfo>
```

Figure 5.4 GeoNetwork Online Link From Dataset Metadata To An Internally Referenced Feature Catalogue.

5.1.1 Potential Use Of Feature Catalogues In SCAR and AODN Infrastructure For Semantic Annotation

Armed with a basic overview of both the SCAR and AODN infrastructure it is now possible to examine how the Feature Catalogue application developed in this thesis, designed to facilitate semantic annotation can be readily applied in these contexts. Semantically annotating data exchanged via Web services usually involves attaching names, attributes, comments, or descriptions, to a document (resource), or to a selected part of a resource. These descriptors provide additional information (metadata) about an existing dataset (presented as a document) or item of data (within a document). Technically semantic annotation refers to instances where the information that is attached uses an ontology, but semantic annotation can also more loosely encompass the attachment of any type of descriptions to a resource (without the inclusion of an ontological definition). The more permissive definition of semantic annotation will be used here because it must be recognised that building capability within communities (and hence their infrastructure) is often achieved through incremental advancement, rather than by monumental shifts. Using formal ontologies and the tools that are required to create and interpret them can present barriers for domain practitioners not well versed in semantic technologies, whilst using informal or semi-formal semantics (e.g., SKOS or plain XML) may not.

Recall from Chapter 2 that in services complying with the OGC services suite there are several different levels at which resources can be annotated (Maue, 2009):

- In standalone metadata documents that are describing either data series, datasets or Feature Types.
- Within metadata elements used to describe service registry artefacts inside a registry.
- In (service) metadata documents specifically designed to describe services, for example within Capability Documents (where annotations could be placed in either the ‘ows:metadata’ field of the service description part of the Document, or in the ‘ows:ServiceIdentification’ keyword section of this Document. Note these two metadata elements are unique to the OGC service – and are not ISO 19115 elements).
- In the XML (or GML)-based schema document (i.e., the Application Schema level) for the data model.
- Within the XML (or GML) instance document that is actually serving the data (i.e., at the data instance level .

The OGC specification does not define the metadata content standard, nor the metadata elements that should be used in constructing metadata within the OGC technology stack. Since the communities under investigation in this study use the ISO 19139 encoding standard for metadata compliant to ISO 19115/19119 (or they map to this standard), there are three current possibilities that arise where annotations could be injected (or referenced) by using ISO 19115 metadata elements. It is feasible to include and use the following ISO 19115 elements in the various annotation levels discussed (particularly within dataset and registry metadata documents):

(i) 'MD_Keywords', or
(ii) 'MD_ContentInformation', or
(iii) 'MD_ApplicationSchemaInformation'.

Whilst semantic annotation does not have to be restricted to using these particular ISO 19115 metadata elements (and there may be better non metadata-related methods, e.g., using simple Xlinks within Application Schema and in data instance documents), it is a useful place to start, since most communities are using metadata to document their datasets. The alternate methods that do not harness these ISO 19115 metadata elements or the additional Capability Document service metadata elements (i.e., 'MetadataURL'; 'ows:ServiceIdentification' and 'modelReferences') will be discussed later.

Within their service registries, both communities use 'MD_Keywords' populated with controlled, or semi-controlled vocabularies to aid in dataset (and by default, linked service) discovery. Although both communities share a common vocabulary (i.e., GCMD Science Keywords), the ability to specialise the lower levels in the keyword hierarchy creates inconsistencies in how this vocabulary is applied by individual data providers. Semantically similar datasets within the AODN infrastructure may not be discovered as such, by using 'keyword'-based searches, because similar datasets may have been marked-up using keywords from these unaligned elements the vocabulary. The vocabularies used are also simple text lists.

At the dataset description level, the (ISO 19115) MCP metadata Profile (used within the AODN) does not currently include 'MD_ContentInformation' or 'MD_ApplicationSchemaInformation' elements. This is because the MCP was one of the first ISO 19115 compliant profiles to be developed internationally and at the time, the utility of these metadata elements were not appreciated. With the benefit of hindsight both of these elements should be incorporated into the MCP. Likewise, the GCMD does not currently have a dedicated placeholder for the types of information that are described in these aforementioned ISO 19115 metadata elements. Within the Capability Documents issued by SCAR and AODN data providers, there is an available option to include 'metadata' (through the 'ows:metadataURL') in the service's 'FeatureTypeList' description, but neither the SCAR, nor AODN communities currently issue guidance to data publishers about how any of these metadata elements should be completed. There is also a lack of guidance regarding completing the 'ows:ServiceIdentification' keyword section of the Capability Document.

Annotation Using MD_ContentInformation

Of the various metadata elements just mentioned, of most interest here is the potential use of 'MD_ContentInformation'. This ISO 19115 element contains, amongst other items a sub-element for identifying a Feature Catalogue (i.e., 'MD_FeatureCatalogueDescription'). This sub-element contains information about the languages supported by the Feature Catalogue, includes the option of recording the names of Catalogue-hosted Feature Types and requires a citation link to the Catalogue. It is this element that could be harnessed to anchor Feature Type references to datasets and services within registry metadata and within Capability Documents.

Maue (2009) has suggested adding another class (element) to the ISO 19115 metadata standard, i.e., an "MD_Reference" which could contain explicit pointers to resource or domain ontologies, but whilst an interesting idea, it is actually unnecessary if the Feature Catalogue developed in this thesis was used as the source of domain concepts. This is because the 'Feature Catalogue' (which also happens to be an ontological resource) is already an inbuilt component of the existing ISO 19115 specification and can utilise the 'MD_FeatureCatalogueDescription' element for semantic annotation purposes.

Parts of the 'MD_FeatureCatalogueDescription' have been implemented in GeoNetwork (see previous Figure 5.4), but not in the manner that provides for the best use of this element. From a review of the bug-tracking site for GeoNetwork, presently it does not seem possible:

- To link to a Catalogue external to GeoNetwork or name the Catalogue within GeoNetwork using anything other than a GeoNetwork generated Id, or
- To identify multiple Feature Types associated with the Catalogue, despite this appearing to be feasible from the 19110 schema included in GeoNetwork.

A demonstration Feature Catalogue description (part of the 'MD_ContentInformation' metadata element) is presented in Figure 5.5 showing ideally how the description could carry annotation links using appropriate URLs for FCAT's Feature Type individuals and the URL for the FCAT Feature Catalogue instance to which they relate. In this example the Feature Catalogue URL points to an external resource (as do the links to the individual Feature Types, which are part of this external catalogue).


```
<gmd:contentInfo>
    <gmd:MD_FeatureCatalogueDescription>
      <gmd:includedWithDataset>
        <gco:Boolean>false</gco:Boolean>
      </gmd:includedWithDataset>
      <gmd:featureTypes>
        <gco:LocalName xlink:href="http://authorityX/FeatureCatalogue1/FeatureType/1"> FeatureTypeA
</gco:LocalName>
      </gmd:featureTypes>
      <gmd:featureTypes>
       <gco:LocalName xlink:href="http://authorityX/FeatureCatalogue1/FeatureType/2"> FeatureTypeB
</gco:LocalName>
      </gmd:featureTypes>
      <gmd:featureCatalogueCitation>
        <gmd:CI_Citation>
          <gmd:title>
            <gco:CharacterString>AODN Feature Type Catalogue</gco:CharacterString>
          </gmd:title>
          <gmd:date>
            <gmd:CI_Date>
            .......
            </gmd:CI_Date>
          </gmd:date>
          <gmd:citedResponsibleParty>
            <gmd:CI_ResponsibleParty>
              <gmd:organisationName>
                <gco:CharacterString>AODN DO <gco:CharacterString>
              </gmd:organisationName>
              <gmd:positionName>
                <gco:CharacterString>AODN Catalogue Administrator</gco:CharacterString>
              </gmd:positionName>
              <gmd:contactInfo>
               .......
                 .......
                  <gmd:onlineResource>
                    <gmd:CI_OnlineResource>
                      <gmd:linkage>
                        <gmd:URL>http://authorityX/FeatureCatalogue/1/FORMAT=XML</gmd:URL>
                      </gmd:linkage>
                      <gmd:protocol>
                        <gco:CharacterString>WWW:LINK-1.0-http--link</gco:CharacterString>
                      </gmd:protocol>
                      ..........
               <gmd:contactInfo>
          </gmd:citedResponsibleParty>
        </gmd:CI_Citation>
      </gmd:featureCatalogueCitation>
    </gmd:MD_FeatureCatalogueDescription>
 </gmd:contentInfo>
```


Figure 5.5 Use of 'MD_ContentInformation' To Describe FCAT Feature Catalogue Content

A problem, however, using this basic method of appending Feature Type information to existing dataset (and within Capability Document) metadata (i.e., through the 'MD_ContentInformation' element), is that there is no provisioning in the supplied example (in Figure 5.5) for describing the format of the responses received from resolving either the Feature Type, or Feature Catalogue URLs. There is provision for stating the 'protocol' used but not format. Under these circumstances it would be best to provide FCAT responses rendered in XML (by default).

In order to use the injected annotation, via the method just demonstrated, a small number of infrastructure modifications need to take place, which harness existing technologies and standards and do not require the addition of semantic technologies.

For example, the AODN would need to import an FCAT (XML serialised) Profile of ISO 19110 into GeoNetwork and generate a populated, HTTP addressable FCAT Catalogue instance (RDBMS or XML-based) capable of producing resource documents that are conformant with this included Profile. The (ISO 19115) MCP Profile would also need to be updated to include missing ISO 19115 metadata elements such as "MD_ContentInformation". The (AODN) MEST registry version of GeoNetwork, underpinning functions of the AODN Portal, would require a 'branch' enhancement to remedy the two Feature Catalogue implementation deficiencies mentioned earlier (i.e., support for multiple Feature Types and improved linking between dataset metadata and Catalogue instances).

AODN data producers and data consumers could then use the MEST to manage Feature Catalogue content, link dataset and services metadata and Feature Catalogue content. AODN data consumers could use embedded links pointing to the Feature Catalogue to resolve semantic annotation for Feature Types identified in dataset and Capability Document metadata, without the use of any additional dedicated tooling (since this particular solution only incorporates serving plain XML responses).

Given that the GCMD DIF metadata records also don't include (DIF metadata element) place-holders for Feature Catalogue information there is also a requirement to make sure that any routine metadata mappings made between the two infrastructure metadata systems are still aligned (after modification to the MCP). Two suggestions to temporarily solve this issue are:

(a) In mapping from MCP metadata to DIF, use the DIF <Related_URL> to manage Feature Catalogue and Feature Type references, and

(b) Request that the GCMD add additional <URL_Content_Type> controlled terms, e.g., 'Get Feature Catalogue' and 'Get Feature Type Definition' to describe the resources being de-referenced by the URLs (Figure 5.6).

```
<Related_URL>
<URL_Content_Type> GET  Feature Catalogue </URL_Content_Type>
<URL_Description> Refers to a Feature Catalogue Resource. </URL_Description>
<URL> http://authorityX/FeatureCatalogue/1/ </URL>
</Related_URL>
                                AND
<Related_URL>
<URL_Content_Type> GET  Feature Type Description - XML </URL_Content_Type>
<URL_Description> Refers to a Feature Catalogue's Feature Type resource semantic description in
XML. Sub-types identify different resource representations (e.g., XML, RDF and OWL)
</URL_Description>
<URL> http://authorityX/FeatureCatalogue/1/FeatureType/1/FORMAT=XML </URL>
</Related_URL>
```

Figure 5.6 Suggested Additional GCMD 'URL_Content_Type' Vocabulary Terms

A better and longer-term solution for the GCMD is to include 'MD_ContentInformation' element type information in the DIF and then develop an appropriately matching ISO 19115 GCMD Profile.

Annotation Using ServiceIdentification Keyword

In addition to using ISO 19115 metadata elements in appropriate places in the Capability Document, there are other places in this Document also able to carry annotation. For example, it is possible to use the 'ows:MetadataUrl' element present in both the Web Feature Service (WFS) and Web Map Service (WMS) specifications which is normally used to provide a link to metadata documents describing FeatureTypes and MapLayers (see Figure 5.7). This 'ows:MetadataUrl' element also has a format attribute which defines the metadata format and its values are taken from a finite set, containing for example, terms such as 'ISO19115' or 'FGDC'. This set could possibly be augmented to include terms such as 'RDF', 'OWL', or 'WSML' (Maue, 2009) and in which case the annotation would be characterised as being capable of providing more formal semantic definitions.

```
<wfs:WFS_Capabilities version="1.1.0">
<FeatureTypeList>
   <FeatureType>
         <Name>rivers</Name>
          <MetadataURL type="0" format="rdf/xml">
         http://authorityX/FeatureCatalogue/1/FeatureType/1/FORMAT=SKOS
         </MetadataURL>
  </FeatureType>
  .......
```

Figure 5.7 Capability Document Snippet With Possible Feature Type Annotation Placement In A WFS Capability Document

Annotation Using 'modelReference'

At the Application Schema and data instance document levels it would appear that adding a SAWSDL "modelReference" attribute to XML (and by default GML) element type definitions is the preferred method (Lemmens *et al.,* 2007; Lefort, 2009; Maue, 2009; Janowicz *et al.,* 2009). Broring *et al.* (2009) have demonstrated an example of annotation at the data document instance level using a snippet from a Sensor Observation Service providing observations involving wind. By using an attribute to hold the SAWSDL "modelReference" the XML can point to either a concept 'WindDirection' or an individual 'HobartCity' drawn from a specific ontology. Their example has been adapted and presented here, as Figure 5.8. In this example, "ontologyX" is substituted for the path expression to a specific ontology that holds the concepts used in the example. Using this technique, AODN and SCAR data providers could use an FCATOWL ontology and REST-based access service to provide a reference to a sa:sampledFeature (Type) as well as a reference to the om:ObservedProperty (i.e., an FCATOWL:QualityAttribute) since both a Feature Type and its properties are ontologically expressed in FCATOWL.

```
<insertObservation service='SOS' version='1.0.0'
...
<om:Observation>
  ...
  ...
  <om:ObservedProperty sawsdl:modelReference='[ontologyX]:WindDirection'/>
  <om:featureOfInterest>
   <sa:SamplingPoint  gml:id='sampling01'>
     <sa:sampledFeature sawsdl:modelReference='[ontologyX]:HobartCity'/>
     ...
     ...
     </sa:SamplingPoint>
   </om:featureOfInterest>
...
```

Figure 5.8 SAWSDL "modelReference" in Annotation (adapted from Broring *et al.,* 2009 – Listing 1.1)

Using more formal semantics assumes, however, that the AODN and SCAR communities could usefully consume the resources supplied (using their own existing infrastructure). Whilst it is likely that semantics delivered in SKOS could be parsed and used with very little modification to existing tooling, the useful consumption of OWL would require more significant infrastructure modification to include Web Ontology Service (WOS) components (Lacasta *et al.,* 2007). Using OWL ontologies is particularly useful for semantic matching (Giunchiglia *et al.,* 2007) during data integration exercises (Vidal *et al.,* 2009), so for this and other reasons its use is a desirable end goal for the communities in this study. But the Feature Type descriptions generated from the FCATOWL ontology, would currently most benefit other communities with whom the AODN and SCAR communities intersect and who are already using semantically-enabled systems. These communities (e.g., Virtual Solar-Terrestrial Observatory –VSTO (VSTO, 2012); Sensor Network Group –SNG (W3C, 2012); Science Environment For Ecological Knowledge – SEEK (SEEK, 2008); MetOcean Domain Working Group (OGC, 2012c)) manage and maintain their own infrastructure but are usually willing to share services and tooling because it is a beneficial way in which to broaden and grow their own systems and user-base (Finney, 2007). It should be remembered, however, that SCAR and AODN data providers can also benefit from sharing semantically annotated data services with these allied communities, because its then feasible to leverage investments already made by these communities in developing semantic infrastructure components (which are not yet evident within SCAR or AODN systems), such as those required for semantic indexing, inferencing, semantic query and semantic mapping.

Resolving Linked Semantic Annotation

All of the discussion thus far has focussed on 'where' and 'how' to embed semantic annotation links. There is already a significant amount of literature about the syntactic methods of linking (DeRose *et al.,* 2001; Lefort, 2009; Lemmens *et al.,* 2007; Cox, 2010b; Schade *et al*., 2010) but there is less information available about the types of standard information that a user (machine or human) should receive when an annotation link is resolved.

Providing a 'modelReference' (i.e., an ontology URI and a direct reference to an entity (class, property, or individual)) in that particular ontology does not guarantee a thin client can adequately interpret the result. The simplest type of resolvable link would be the presentation of the complete ontology (e.g., as an OWL, RDF or SKOS file) at the end of a URL, with a pointer to an ontology member. The client is then left with the task of parsing the file and extracting from it any contextual information associated with that member that might be required. In a complex ontology, knowing what connected elements to extract is a difficult task for a machine client. This task could be executed, however, by loading the full ontology into a local ontological repository and then using SPARQL to extract information of interest (provided the ontology's knowledge model was well understood), or perhaps by creating an XML data model and using XQuery (Boag *et al.,* 2010), provided that no reasoning processes were anticipated. These approaches place the problem of locating all of the necessary ancillary information surrounding a linked resource and the burden of its interpretation on the client.

In developing the Feature Catalogue prototypes outlined in this research significant thought was put into how the FCATOWL, FCATSKOS or plain XML tagged data could be consumed by clients and what packages of information should be delivered in response to requests for information on resources. In the use-cases supporting the development of the ontologically-grounded Feature Catalogue, most use-cases did not involve the requirement to reason over the data. But even in these (non-reasoning) scenarios it was still considered mandatory to be able use the supplied descriptions to discriminate between different Feature Types, or determine if two Feature Types are essentially the same but have different characteristics, perhaps because their geographic representations or sampling (sensing) processes are different. The research in this thesis has argued that both the (Feature Type) concept and its context are therefore required in many situations. Broring *et al.* (2009) and Janowicz *et al.* (2009) outline similar issues for situations in which they discuss the challenges of correctly mapping sensor Web concepts in sensor Web "plug and play" scenarios. Janowicz *et al.* (2009) relegate the solution to this information granularity/context problem, in part, to the Web Observation Service (WOS) which in their particular design includes a "Lookup and Retrieval" function. In this function a "get model reference" service returns an appropriate ontology element ID for a given resource ID (e.g., a GML Feature ID). Then a "Retrieval" service executes semantic matchmaking between a query and (a) available web service advertisements and (b) Feature Type definitions. Precisely how this would all be achieved and what constitutes a "Feature Type definition" is not, however, explained.

As has been described earlier, in the Feature Catalogue access services developed in this thesis, Feature Type contextual information includes:

- A lexical definition of the Feature Type (its class type and location with respect to a DOLCE upper ontology in the case of the mooted FCATOWL service).
- A list and description of a Feature Type's attributes and in the case of those attributes that are considered observable properties, a description of their sampling process, units-of-measure and grounding datums.
- A Feature Type's spatio-temporal representation(s).
- A list of any known relationships that the Feature Type participates in with a link to the Feature Type(s) that are the range of that relationship.

In the case of a SPARQL end-point placed directly over FCATOWL this information could be provided directly through some standardised query templates. Ideally, however, a domain concept provider (such as FCATOWL and its associated access services) should serve information to clients in a number of forms (i.e., information bundles and formats) or supply scripts that perform any necessary data translations (the latter being difficult when trying to address issues of both format and information granularity).

The prevailing view of semantic Web service clients is that they work with data represented in RDF. In contrast, Web services usually expect (and emit) messages in XML. As has been canvassed already in Chapter 2, in order to enable a semantic client to communicate with published Web services, its semantic data must be “lowered” into the service provider’s expected input messages (e.g., expressed in XML), and the data coming back from the service in its output messages must be “lifted” up to the semantic level (e.g., expressed in RDF) for consumption by the client (Passant *et al.,* 2009). Sometimes it’s not the resource rendering, however, that is so important but also how much information should be delivered and of what type (i.e., information granularity).

In the Feature Catalogue application delivered in this research, the approach taken was to develop REST-based Web services, specialised for different use-cases that delivered Catalogue content at different levels of semantic granularity and in different formats. Lifting and lowering is masked and instead the problem is one where the Web client has to make a choice concerning the most suitable available resource format in which to receive data from the service. The data provider is in control of the granularity with which a resource is supplied because they are responsible for linking an appropriate URL (for Feature Type definitions) to a Capability Document or service description and can also use the FCAT services to directly reference other (associated) resource types (e.g., attributes aka observed properties) which are often separately described within an Application Schema or data instance document. Using the FCAT services, data providers have a choice about the level of information granularity that is supplied via given URLs. In framing the services trialled in Chapter 4, it was considered better to provide more contextual information in delivered payloads (for Feature Type definitions), rather than less, because it was reasoned that superfluous information can always be parsed out (or ignored).

Although time and resource constraints precluded demonstration of an OWL encoded output from the Rest-based Feature Catalogue services, versions of the semantic concept delivery services were developed for (XML, HTML and SKOS). By using any of these RESTful services, individual resources from within the relational Catalogue repository are accessible via simple URL reference (using URI templates; Gregorio *et al*., 2010) and it is possible to obtain various views of a Feature Type.

Since Rest-based services that deliver messages about Feature Type resources contain links to other resources all contextual information is resolvable, it is a question of how many “HTTP” traverses are required to get at the specific (contextual resource) information that is required. For this reason, in addition to the traditional iterative drill-down type of REST-based service (i.e., the sample IterativeCatalogueService developed in Chapter 4) an example of a more compact message form (i.e., a StreamLinedCatalogueService) was also provided. In the case of REST queries issued to this latter service the payload received by the client contains a full and expanded listing of all resources associated with the resource that was the focus of the initial query. In this way no extra HTTP traverses are required and all contextual data is provided in one hit.

The simplest and most straight-forward use of the Feature Catalogue and its REST interface for annotation purposes would involve the attachment of an appropriate URL (that describes a Catalogue resource) to a resource element in an Application Schema or instance document, chosen by the data provider so that the Catalogue resource URL delivers information in a specific form and format (e.g., a provider links a URL that resolves an IterativeCatalogueService for a named Feature Type that delivers information in XML). But using this approach, that is one of attaching a single URL, will result in a fixed response (and payload), even if there is another (REST) service available that could have served Catalogue content in a more suitable manifestation (for a particular Web client).

However, since it is possible to anchor to more than one URL using both 'Xlink' and SAWSDL constructs, more flexibility could be provided to take advantage of the full range of provided FCAT Feature Catalogue services. A Web client could be given the option of choosing the most suitable message response rendering for its purposes (in terms of both information granularity and format) if the Web client could be presented with a choice. This situation could be achieved by providing more than one URL as an option for the Web client through an "extended Xlink" (DeRose *et al.,* 2001), or by using a SAWSDL *modelReference* with multiple pointers. Unfortunately, in the SAWSDL case the nature of the relationship between such multiple presented links is relatively semantically opaque (for the Web client using currently available SAWSDL constructs). In an "extended Xlink", however, there is scope for establishing relationships amongst links and there are some descriptive attributes for links, but neither of these two approaches makes it particularly easy for a machine-client to interpret which service would be most appropriate. But using either approach provides all of the links and therefore all of the FCAT services that are available for exploitation.

Perhaps a better solution for providing a choice about semantic resource descriptions is already available under the HTTP protocol, at least with respect to formats. Cox (2010b) reminds OGC service users that under the HTTP protocol there is existing support for content negotiation between client and server to allow for the delivery of a resource representation that a client is capable of processing. The HTTP "Accept" and "Accept Language" headers can be used to both constrain the type of service offering a Web client is willing to take and can also be used to publicise the types of services that a provider can supply. Accept headers can also be used to indicate that the request is specifically limited to a small set of desired types. The Accept header gives a browser the chance to tell a Web server in which format it wants a resource to be rendered. Further, by giving a list of available options this content negotiation happens in a single request. Sample code for a typical HTTP Accept header has the form:

Accept: text/plain; q=0.5, text/html, text/x-dvi; q=0.8, text/x-c

Verbally, this would be interpreted by a server as "text/html" and "text/x-c" are the preferred media types, but if they do not exist, then send the text/x-dvi entity, and if that does not exist, send the "text/plain" entity. Understanding the nuances of the arcane header syntax is not required here, simply that the facility exists.

For example, the Feature Catalogue application in this thesis could advertise the following general URL that provides access to the Catalogue Resource with ID=1:

http://data.aad.gov.au/aadc/FeatureCatalogue/index2.cfm?/IterativeService/Catalogue/1

A GET request (from a Web client) to this URL could return different responses depending on whether the HTTP "Accept" value in the Web client HTTP header was *text/html*, *application/xml*, or *application/rdf+xml*. Cox (2010b) states the available variant responses in such a case are responses in their own right and they would be addressed as they are now in the REST-based Feature Catalogue application (i.e.):

http://data.aad.gov.au/aadc/FCAT/index2.cfm?/IterativeService/Catalogue/1/Format=HTML
http://data.aad.gov.au/aadc/FCAT/index2.cfm?/IterativeService/Catalogue/1/Format=XML
http://data.aad.gov.au/aadc/FCAT/index2.cfm?/IterativeService/Catalogue/1/Format=RDF

Which of these URLs would be resolved would depend on the Web client's stated preference for HTML, XML or RDF in the Accept header.

If no preference was selected by the issuing client, a specific format would be supplied by default (which is currently the case in the Feature Catalogue REST-based services if the "/Format" switch is left off of the service request). This HTTP header solution addresses the problem of response formats but not the problem of which service to use if information granularity is an issue. Solving this particular problem needs additional research.

With the potential for multiple annotations per service description, Maue *et al.* (2010) have developed and trialled a Semantic Annotations Proxy service (SAPR) that permits the injection of semantic annotation into service documents using a proxy registration process. The associated injected annotation procedure allows for a decoupling of the semantic annotations from the original service metadata. Different sets of annotations may be added to the service metadata (without changing anything in the original service document). This proxy-based method is interesting but relies on a service intermediary and because it is in early prototype phase it is uncertain if semantic constructs, in addition to those already mentioned (e.g., SAWSDL), will be supported. Whilst additional semantic constructs are not supported, this method offers an advantage only in the circumstances where existing services have been deployed by communities and additional or different annotation needs to be added post deployment (to that which has already been anchored).

Annotation Placement Preferences

Thus far four options have been provided in terms of where annotation can be placed within the SCAR and AODN infrastructure (i.e., Dataset/Registry metadata, Service Metadata, Application and instance document levels). Using any particular level in which to anchor the semantics has a different issue with respect to the 'find-ability' and 'usability' of the resource description (see Table 5.1, adapted from Appendix B; Maue, 2009).

Table 5.1 –Annotation Level Pros and Cons (adapted from Appendix B; Maue 2009)

Advantages, Drawbacks and Applications Of Using Annotation At Different Levels In The OGC Services Suite
Source: Dataset/Registry Metadata And Service Metadata In An OGC Capability Document **Advantages:** • Relatively easy to implement. • Easy for human users to understand (even for non-experts). • Helps to (quickly) identify what kind of data is served. • Needs no modification of the underlying data schema (because the annotation is not present in the schema). • Good recall when searching for datasets due to semantically-enabled query processing. • Is the only current solution for binary data like images (served by WCS or WMS). **Drawbacks:** • Semantically-enabled discovery resources require specialised interfaces that provide users with the ability to select the embedded concepts from within the metadata element in a record, or from within the Capability Document (if delivered in OWL). • Can impair readability of service metadata. • Potential consistency Problems: Changes in underlying data model will require changes to annotations in Capability Document or a metadata record. • Not possible to directly identify the parts (tagged elements) of the dataset that are actually represented by annotation since the annotation is anchored at a high level in the Capability

Document. **Suggested Applications:** • Basic Service Discovery. • Basic Service Evaluation.
Source: Application Schema **Advantages:** • Higher precision and better recall (of queries) since annotation is anchored to specific elements in the Schema. • Less problems with consistency, since changes to the data model have a direct impact on annotations in the Capability Document. **Drawbacks:** • Describing complex data models can be tedious, requires additional documentation. • Semantically-enabled discovery resources require specialised interfaces that provide users with the ability to select the embedded concepts from within the Schema (if delivered in OWL). • Annotations of processes (e.g., in a Web Processing Service) become rather complex when there is embedded schema semantics. **Suggested Applications:** • Service Discovery. • Service Evaluation. • Workflow Validation. • Merging (and Integration) of Datasets.
Source: Instance Schema **Advantages:** • A lightweight approach, if using a 'modelReference' attribute added to entities. • Flexible annotation of Features. • If annotation was also present at dataset/registry level a high level of semantic discrimination would be possible, with respect to finding and sub-setting data. • Good solution for data without explicit application schema. **Drawbacks:** • If no accompanying annotation at dataset/registry or service level, unlikely that resource discovery will be aided. • Would ideally need annotation tools in a client application (i.e., in Desktop GIS or Web mapping tools). • Increased data volume (verbosity). **Suggested Applications:** • Resource Discovery (with filtering). • Service Evaluation. • Quality Control for Resources. • Merging (and Integration) of Datasets.

It would seem most sensible to annotate at the dataset metadata, Capability Document and data instance level so that maximum utility is gained from the semantic annotation. This may be impractical and labour intensive, however, without modification to existing tooling (e.g., metadata editing tools and Web mapping servers) to enable user-friendly, and in some cases automated, linking of semantic annotation (to avoid repetition in annotation). For example, GeoNetwork currently automatically extracts layer information from Capability Documents and links this to dataset level metadata. A similar facility could automatically extract Feature Type definitions from these Capbility documents and use them to populate dataset metadata elements. As a first step, in particular to aid with dataset discovery, Feature Type information should, as a minimum, be recorded within dataset and registry level metadata elements.

Summary

In summary, it has been demonstrated that annotation using the FCAT semantic Feature Catalogue model can be achieved within the SCAR and AODN infrastructure in several ways:

(a) Provide links from a resource to content managed in external resources to allow local applications to access remote content for enrichment purposes, for example by using XLinks (DeRose *et al.,* 2001) as in Figure 5.5, or by using defined XML elements to carry references (e.g., 'gmd:URL' also demonstrated in Figure 5.5; or by using a 'metadataURL' as in Figure 5.7 for OGC WFS or WMS services; or by using a GCMD 'Related_URL').

(b) Provide a 'modelReference' to ontological descriptions (as in Figure 5.8). Note that often an accompanying lifting script is also required, i.e., a script used to capture the data transformation that may be needed for invocation and semantic data mapping purposes, which allows remote applications to lift content from a resource. Since the REST-based methods developed in this thesis provide resources in a variety of formats, lifting scripts are not required, provided that there are agreed mechanisms for a client to understand and then subsequently choose available formats.

(c) Provide embedded attributes to allow remote applications to lift content from a resource but through RDFa (Adida *et al.,* 2008). This particular approach was demonstrated earlier in Chapter 4 where examples were given for describing services with hREST (microformats) and RDFa (but without examples of embedded 'modelReferences').

Architecturally there is a demonstrably good fit between the artefacts developed in Chapter 4 and their immediate utility within the SCAR and AODN infrastructure, requiring relatively few and minor modifications to existing systems, standards and protocols. This is because as a source of semantic annotation, FCAT and its associated REST-based access interfaces and query patterns can leverage the types of annotation paradigms currently available to the SCAR and AODN communities. Since REST is the most basic protocol on the Web there is no reliance on OGC styled services or service query patterns that would subsequently limit the scope of the application of the FCAT Feature Catalogue services. All FCAT resources are identified by unique URLs that can be used in conjunction with any type of annotation service (be they OGC-based or where data is both expressed and annotated as RDF triples). Essentially FCAT and its services can provide semantic definitions (for Feature Types) that can be used anywhere a Linked Data type service is appropriate. This general approach (i.e., combining Semantic Web approaches with what is effectively Linked Data resources) has recently been proposed by a number of research groups (Phuoc and Hauswirth, 2009; Sequeda and Corcho, 2009; Janowicz *et al.,* 2010; Patni *et al.,* 2010). Despite differences in suggested methods all argue that features-of-interest and observations should be identified using URIs, looked up by dereferencing these URIs over HTTP, encoded in machine-readable knowledge representation languages such as RDF and OWL, and interlinked with other resources.

5.1.2 FCATOWL – Alignment With Other Observation-Centric Ontologies

Although it has been demonstrated that FCAT modelled resources can be successfully used within the system architectures of the SCAR and AODN communities, emphasis was placed on assessing syntactic and structural 'fit' rather than assessing ontological alignment. This is because the AODN and SCAR communities have no ontologies with which to compare and contrast the FCATOWL or FCATSKOS models. As previously mentioned there are, however, communities with which both the SCAR and AODN communities intersect, who have been early adopters of semantic technologies and who use observation-centric ontologies. Since the interoperability of the SCAR and AODN systems with those of intersecting communities will depend upon the degree of alignment between the ontologies used, it was decided to assess how FCATOWL aligned with these observation-based semantic models.

Ontology alignment is an area of considerable activity and research and it is used in any mediation process where it is needed to obtain interoperation between ontology models for any specific task, for example in: data integration (Calvanese *et al.,* 2001); merging (McGuiness *et al.,* 2000), transformation (Klein and Fensel, 2001) and translation (van Assem *et al.,* 2004).

Whilst it is not intended here to perform complex mappings between FCATOWL and each of the observation-centric ontologies that will be evaluated, a general assessment is made of the concordance between FCATOWL concepts and those of the compared ontologies. For each ontology the question asked is, "could an adequately populated FCATOWL ontology serve as a source of domain concepts for this ontology, for semantic annotation purposes" ?

A Recap of FCATOWL characteristics

A broad summary of FCATOWL (which was developed in the previous chapter) is provided as a necessary pre-cursor to subsequent comparative discussions.

FCATOWL (based on the FCAT enhanced model of the ISO-19110 *Feature Catalogue Methodology*) was designed to enable the semantic description of Feature Types, which can be considered to be both (DOLCE) 'perdurants' and 'endurants'. Endurants (also known as substances or material objects) are those entities that can be observed/perceived as a complete concept, at any point in time. Perdurants (also known as occurrents) are those entities for which only a part exists if we look at them at any given snapshot in time. The ontology is feature-centric and capable of modelling traditional GIS-type features (mainly endurants) and features which are considered observations (i.e., perdurants). FCATOWL is a generic, container ontology for capturing the type of semantic details about Feature Types which are most often required when trying to determine whether two or more datasets containing specific Feature Types are able to be meaningfully integrated, used together, or chained in a service workflow to suit a particular use-case. In these circumstances it is important to know:

- Whether two Feature Types are describing the same real-world material concept or occurrent.
- The type of spatio-temporal representation that the Feature Type has been encoded in.
- The individual properties of the Feature Type and where these properties are (DOLCE) qualities, what type of quality they are.
- The collection methods used in sensing the Feature Type properties.
- The units-of-measure used in the estimation of any of the Feature Type's properties and the datums or conventions used in grounding these units-of-measure.

Additionally, FCATOWL includes a mechanism to enable communities-of-interest to ascribe Feature Types to different sets of classification systems which ultimately helps to improve the ways in which a Feature Type might be discovered. In FCATOWL some governance facets are also built into the ontology by assigning Feature Types to profiles (i.e., collections of Features Types owned and managed by a particular user community).

FCATOWL Feature Types are an *IRE:WebResource* (a specialisation of a DOLCE: *information-realisation*) which means that an *FCATOWL:FeatureType* is a '*proxy'* for some real-world entity at any given time (t) about which you can make any number of assertions. A '*WebResource'* is a resource that is made available on the Web, hence is accessible through a Web protocol (e.g., a document or a Web service). The 'identity' of an FCATOWL Feature Type is therefore clearly defined as an information resource. This issue of identity could become significant in spatial data infrastructure since all HTTP (protocol) response codes should be used to indicate if a representation of an 'information' or 'non-information' resource is returned when dereferencing a URL (Cox, 2010b).

Finally, the FCATOWL ontology caters for the situation where a Feature Type may have different characteristics (attributes), depending on its spatio-temporal representation and/or its usage by a particular community-of-interest. It is therefore feasible for a given Feature Type to have two slightly different semantic descriptions (as expressed through its semantic constituent parts), but still be recognised as the same Feature Type. In all of these respects FCATOWL provides for a rich description of a Feature Type that can be applied in a wide range of use-cases.

It is important to note, however, that although FCATOWL is modelled on ISO 19110 it departs from the standard in several aspects. FCATOWL models a spatio-temporal footprint, includes 'collection' properties and does not include Feature Type 'operations'.

FCATOWL & The Semantic Sensor Network Ontology

Because FCATOWL draws heavily on the observation and measurement model it is highly complementary to ontologies recently developed for supporting Sensor Observation Services (Na *et al.,* 2007), which were introduced briefly in Chapter 2. Neuhaus and Compton (2009) have developed a Semantic Sensor Network Ontology (SSNO), which recently underwent review and subsequent generalisation through harnessing the idea of the Stimulus-Sensor-Observation Ontology Design Pattern (Janowicz and Compton, 2010). Figure 5.9 depicts the generalised version of SSNO, whilst Figure 5.10 describes elements of the Stimulus-Sensor-Observation (SSO) pattern to which SSNO now conforms. The SSO pattern was developed following the principle of minimal ontological commitments to make it reusable for a variety of application areas. A minimal set of classes and relations are used to express the notions of stimuli, sensor, and observations.

Based on the work of Quine (1995), the skeleton module of the SSNO ontology defines stimuli as the (only) link to the physical environment. In this pattern empirical science observes these stimuli using sensors to infer information about environmental properties and to construct features-of-interest. To tighten the semantics of the model and to improve interoperability, the enhanced version of SSNO was also cast in a DOLCE UltraLite framework (see Figure 5.11 for linkages to DOLCE).

Despite sharing a common upper ontology, there are notable differences between the class and property structures of the enhanced SSNO and FCATOWL in terms of how they have both been mapped into DOLCE, due to their separate development paths and emphases. For example, an *Observation* in SSNO is modelled as a DOLCE *Situation* and *Features-of-Interest* only have DOLCE *Qualities* as properties.

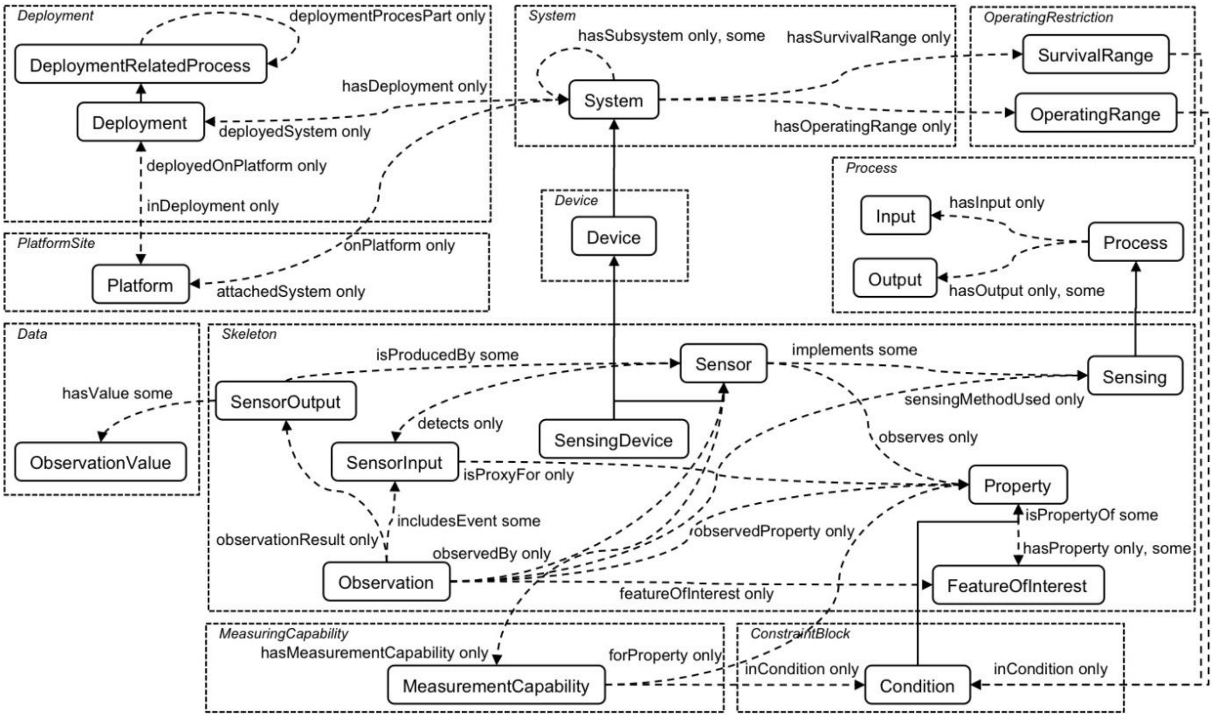

Figure 5.9 – Semantic Sensor Network Ontology Classes and Properties (Figure 2, Lefort *et al.,* 2011).

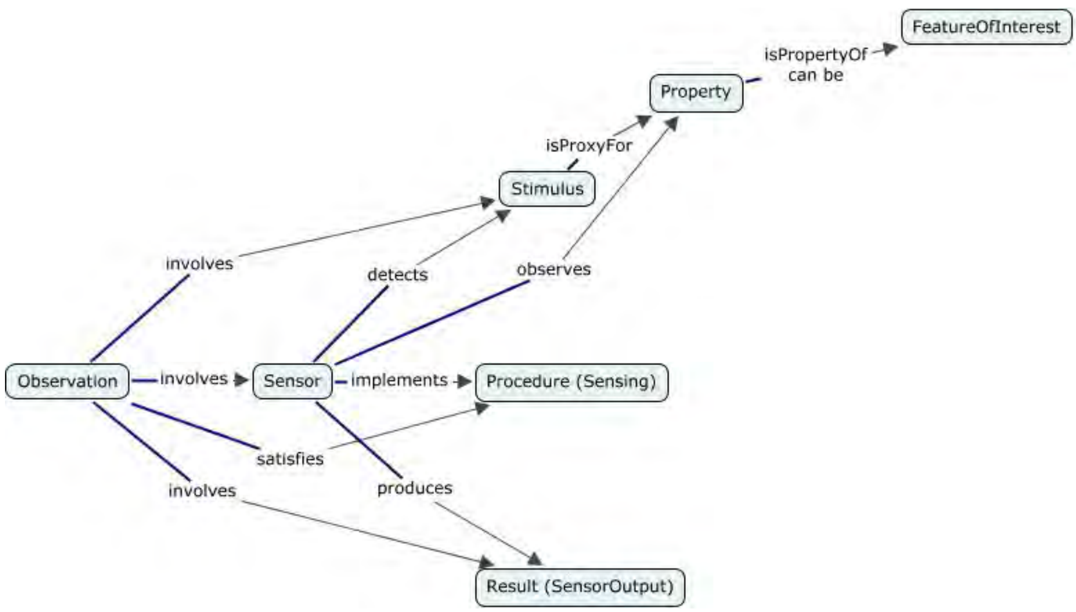

Figure 5.10 – Stimulus-Sensor-Observation (SSO) pattern (Figure 4, Lefort *et al.,* 2011).

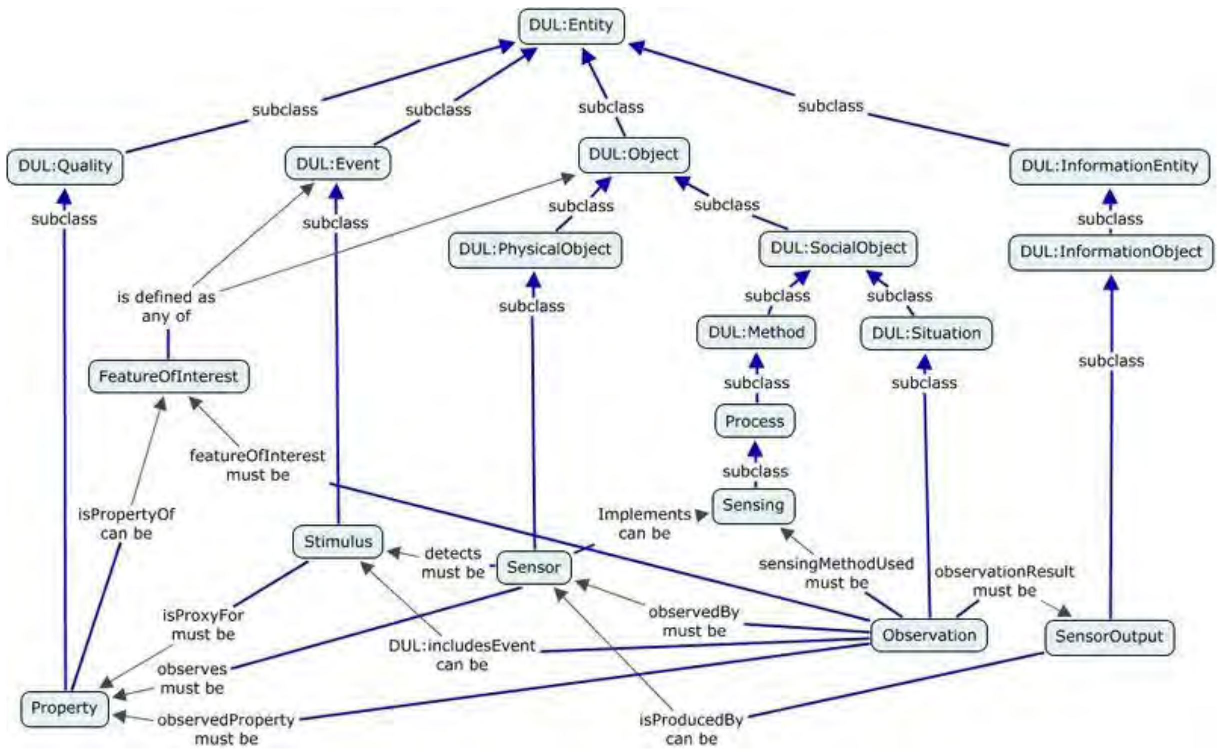

Figure 5.11 – Sensor Network Ontology & DOLCE (Lefort *et al.,* 2011).

SSNO does not attempt to model the domain of Feature Types but provides a hook (via the entity *Feature-Of-Interest*) to import any necessary ontological descriptions. Importantly, the two models (SSNO and FCATOWL) can be viewed as being synergistic and with sufficient commonality for FCATOWL to be used to seed domain features in SSNO. Semantically, FCATOWL defines a *Feature Type* to be the union of a DOLCE endurant and a perdurant (as does SSNO). FCATOWL also delineates feature *Attributes* into two types of DOLCE class - *Qualities* (equivalent to properties in SSNO terms) and *NonQualityAttributes*. So, if required, only those *Attributes* in FCATOWL that are considered *Qualities* could be referenced from within the SSNO ontology. Both FCATOWL and SSNO consider a (sensing) *Process* to be a DOLCE *Method*, so semantically at the core, there is concordance between the two ontologies.

Although not evident from the SSNO figures included, both FCATOWL and SSNO also represent the effective detectable results of an SSNO observation as a DOLCE *Region.* The relations and classes used to model perceivable changes in the environment which are directly or indirectly related to observable *Properties* differ, however, between the two ontologies. FCATOWL models the relationship between an observable *Property* (i.e., FCATOWL:*QualityAttribute*) and its *Region* (or *Quale*) simply and directly through a "hasRegion" relation to an FCATOWL: *Value* class. FCATOWL allows for the assignment of FCATOWL:*UnitsOfMeasure* to that *Value* class through a DOLCE "parameterises" relation. Using the Stimulus Sensing paradigm, SSNO more correctly semantically models the relationship between a *Result* and an observable *Property* (refer again to Figure 5.10). In SSNO a *Stimulus* is modelled as a "proxyFor" the observable *Property*, which is then "detected" by a *Sensor* which then links to the *Result* via a "produces" relation. Despite these differences it is still relatively easy to map from one to the other and these differences do not affect FCATOWL's use as a point of truth for serving up domain concepts for inclusion in SSNO, or SSNO derivatives.

Recognising that the SSNO does not take into account how observation data was created and doesn't address the problem of semantically identifying what types of features are being referred to, Devaraju and Kuhn (2010) have begun to explore the development of what they term a "process-centric domain ontology" to fill the gap. Presumably many such ontologies will be required as feed-stock for SSNO-based implementations given the breadth of domains in which sensors can and will be deployed. The example ontology described by Devaraju and Kuhn is taken from the Hydrology domain (see Figure 5.12 for a partial expose of an evaporation process ontology). They argue that a process-centric domain ontology first identifies geo-processes, entities (i.e., objects and matters), and their properties followed by the relations between them. The SSNO and process-oriented domain ontology are then combined into an application ontology, where the domain features and processes are mapped into *SSNO:Feature and SSNO:Observation* entities respectively and domain feature *Qualities* are mapped into the *SSNO:PhysicalProperty* class. Note that the mapping process uses the domain ontology to create subcategories (subclasses) of the aforementioned SSNO classes.

It is suggested that FCATOWL, once appropriately populated, would similarly be able to provide the necessary observable property-types and associated real-world entities that are applicable for establishing SSNO application-specific ontologies. This is because: FCATOWL characterises its inhering Feature Type *Attributes* as DOLCE *Qualities (*which are further categorised by *QualitySpaces*)*;* it identifies *FeatureTypes* that play an observation role and which are proxies for entities that are sub-categories of *DOLCE Processes* and establishes connections between an observation (*FeatureType*) and other types of features (e.g., those that play Features-Of-Interest and SampledFeatures within an observation) through the assignment of roles and semantic relationships (see Figure 5.13).

FCATOWL anchors its *FeatureTypes*, which are *InformationRealisations* to real-world entities (i.e., domain objects) through the "proxyFor" relation. Each *FeatureType* is expected to be characterised by either a DOLCE *Object* or a DOLCE *Event*. Because FCAT is a container ontology, FCATOWL *FeatureTypes* are OWL individuals, whereas SSNO *Features* are OWL classes. Therefore, using FCATOWL to seed the SSNO application ontology would first require the extraction of the FCATOWL individual *FeatureTypes* with their subsequent re-casting as classes in the application ontology.

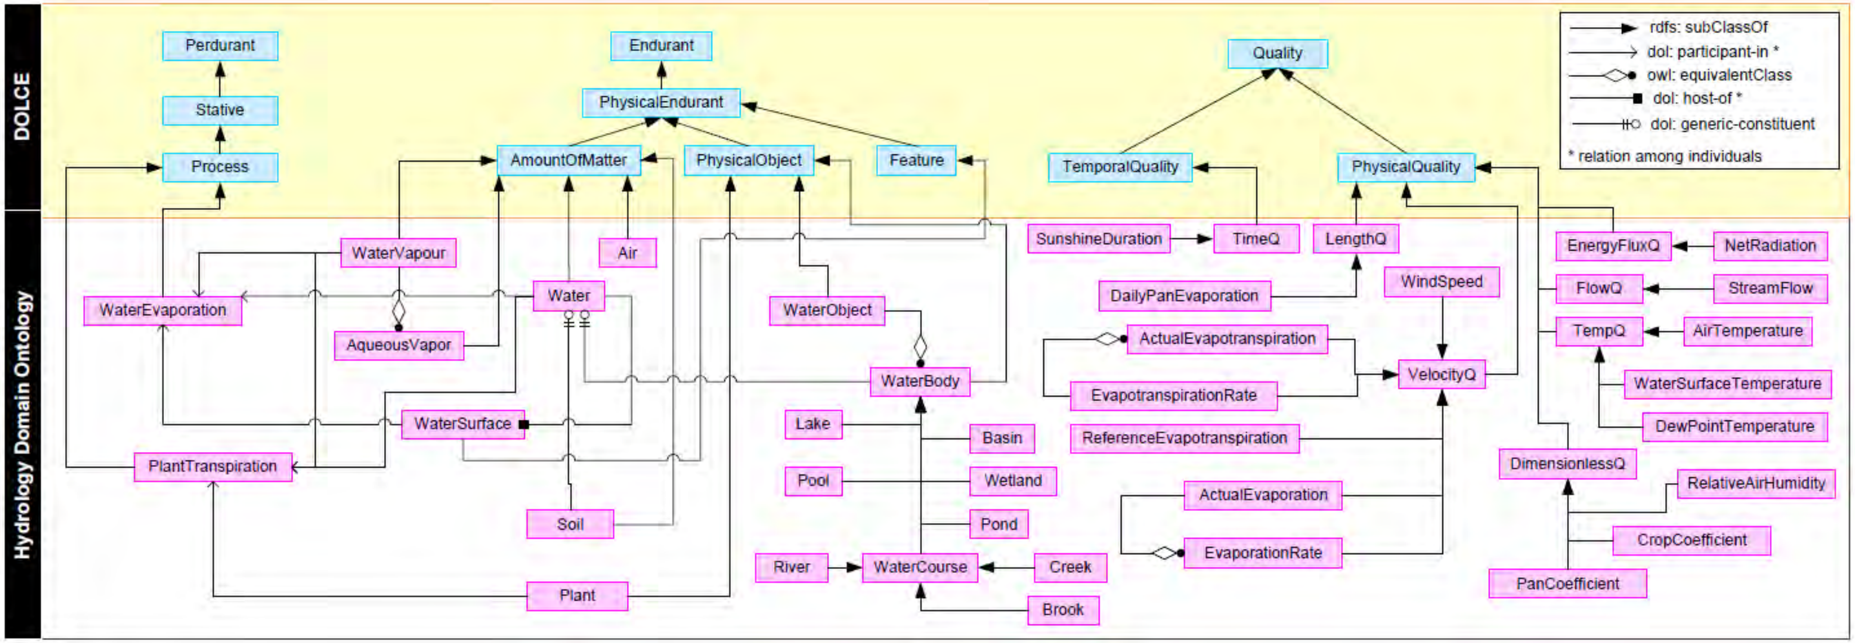

Figure 5.12 – Partial View of A Process-Oriented Ontology For An Evaporation Process (from Figure 3, Devaraju and Kuhn, 2010).

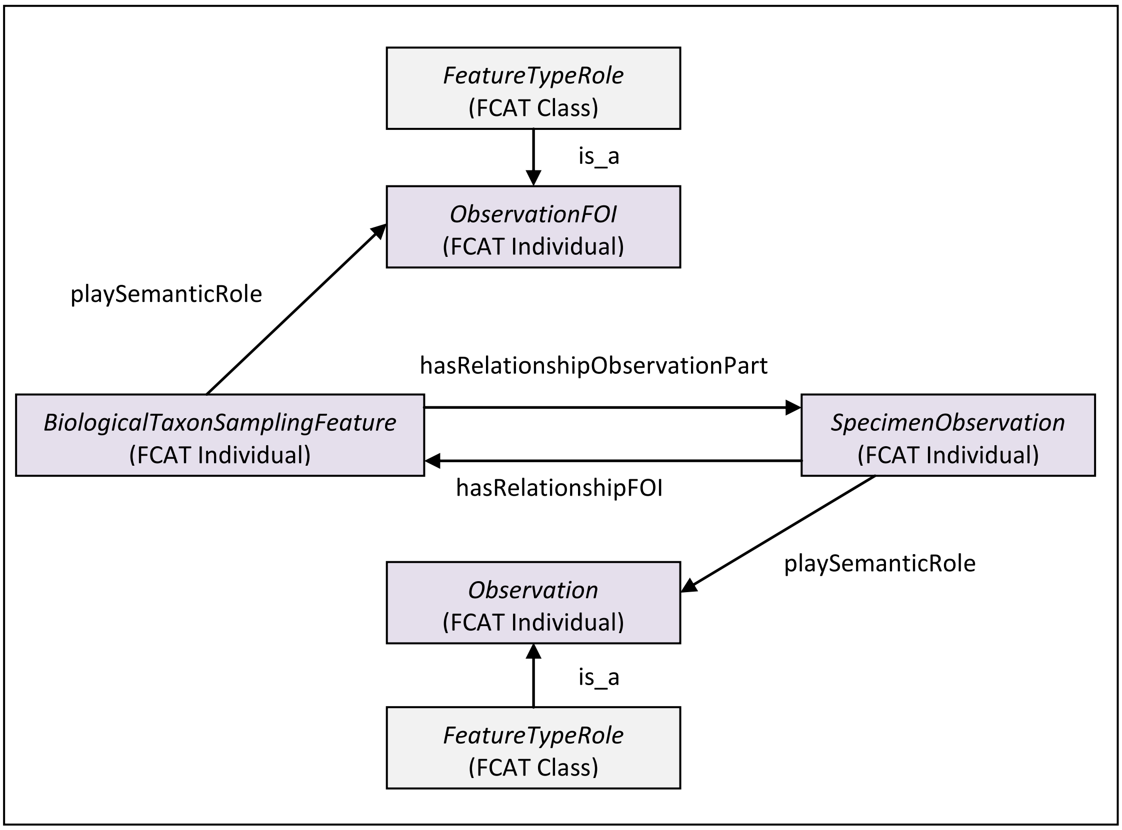

Figure 5.13 FCAT Semantic Relations That Link Roles To Feature Types

It is important to note that FCATOWL is capable of holding different (and legitimate) views of some of the entities used in constructing the container ontology. This is because OWL 2 DL allows punning, i.e., different uses of the same term whether it is an individual or a class (Hitzler *et al.,* 2009). For example, an *Instrument* class is specifically used in the FCATOWL model to describe the sensing method for a *QualityAttribute*. It is also possible that in populating FCATOWL, content providers will wish to create a Feature Type that is a "proxyFor" an *Instrument (*i.e., an *Instrument* will simultaneously be a type of DOLCE *PhysicalArtefact* (and a class in FCATOWL) and an FCATOWL *FeatureType* individual.

In SSNO, *Deployment* and *Platform* are separate and specific modules (and Classes) that have their own characteristics. The data model patterns suggested in section 4.2.4 (of this thesis), one of which is then serialised as a demonstration in section 4.4.2, also recommend that such entities be modelled as separate Feature Types in FCATOWL. It is possible, but doubtful that FCATOWL would be used to provide semantic descriptions for these types of entities to seed SSNO application ontologies. It is more likely that FCATOWL would be a source of domain concepts similar to those modelled by Devaraju and Kuhn (2010).

An aspect of FCATOWL which should be of interest to the Sensor Observation community is that FCATOWL categorizes *Qualities* (through their associated *QualitySpaces* and *ReferenceSpaces*) at a high level of discrimination, by leveraging the foundation work of Probst (2007). It is argued that in future scenarios, when operational (as opposed to experimental) sensor observation networks are more pervasive, the ability to more precisely identify *Quality* categories will aid discovery and integration of similar data drawn from different data sources.

Broring *et al.* (2009) have also identified the importance of the granularity of classification of observed properties for use-cases involving the automatic registration of "plug and play" Sensor Observation Services. They assert that the specified output of a sensor should produce appropriate values for the properties of a given features-of-interest. They cite the example of hypothetically registering a sensor for detecting a feature's property e.g., "wind direction", where the wind has been modelled as a 3-dimensional quality. They indicate that a robust sensor Web infrastructure would issue a warning if the sensor assigned to measure such a property was a 2D anemometer (because of the mismatch between the type of quality anticipated and that which the sensor is capable of producing).

FCATOWL, SERONTO & OBOE

The aim of developing the FCATOWL model was to make it as general as was possible to support a multi-disciplinary scientific base. Catapano *et al.* (2011) highlight that there is no harmonised model (across disciplines) for supporting the description of scientific observations. This is despite domain-oriented informatics experts perceiving the utility to be gained from adopting a core model for scientific observations and measurements that could be flexibly linked to domain-sanctioned controlled vocabularies. They assert that an observational abstraction would provide the underlying scaffolding for describing specific observations, and enable observations to be inter-related. They are aware of Cox's O&M abstraction (2006; 2010a) but did not see it as a suitable solution for observations involving biological data. Instead, SERONTO and OBOE are two ontologies mentioned by Catapano *et al.* (2011) that have emerged from two distinct biological communities that attempt to provide common scaffolding. In the next several paragraphs SERONTO will be reviewed, followed by OBOE, with respect to FCATOWL.

SERONTO (van der Werf *et al.,* 2009) is an ontology developed within ALTER-Net, a Long Term Biodiversity, Ecosystem, and Awareness Research Network funded by the European Union. The ALTER-Net system addresses major biodiversity issues at a European scale. To this end, SERONTO has been developed as a new approach to deal with the problem of integrating data from distributed data sources stored and collected at different locations within the European Union. It is marketed as an ontology of observations with theoretical roots in statistical methodology.

The core classes and properties of the SERONTO observation model are shown in Figure 5.14. This observation model can best be described as being of the form: a parameter “V” with scale “S” and unit “U” is obtained by the method “M” on the experimental unit “E” from the population “P” which has the values “h1” at time “t1”. A populated example of this model is given in Figure 5.15.

A *Parameter* does not correspond to the idea of an observable property (as expressed in the O&M model or its FCATOWL:*QualityAttribute* equivalent) but instead could be any of a ‘measurement’, a ‘treatment’ or a ‘classification’. An *Experimental Unit* (aka *PhysicalThing*) coincides with the idea of an SSNO feature-of-interest (although in SERONTO ontologically this seems to be restricted to endurants).

A SERONTO *ParameterMethod* is a type of convenience (n-ary) class that marries a *Parameter* to its *Method*, *Units-of-Measure* and *Measurements Scale. Method* is further sub-divided into three classes *ClassificationMethod, TreatmentMethod and MeasurementMethod*. *Values* are coupled to the *PhysicalObject* and a *Parameter* by means of the *ValueSet* class. *SelectionDescription* connects the sample with the population from which it is drawn (via the properties “hasSample” and “hasPopulation”).

Although not shown in Figure 5.14 a *SelectionDescription* also has a *SamplingMethod* which describes the type of method used e.g., ‘random sampling’, ‘convenient sampling’, ‘stratified sampling’, etc. Specifically how the sampling has been executed is described in another class called *SamplingProcedure* (which is also related to *SelectionDescription*). The last core classes that are included in SERONTO are the *ReferenceElement*, *ReferenceList* and *Reference* class. Apparently, during the development of SERONTO a need was expressed to be able to attach a *Reference* (which could be part of a *ReferenceList e.g., a species list*) to any SERONTO object. In addition a *ReferenceElement* can be associated with a *ValueSet* , via a special SERONTO property called “hasHelpObject”, which can then hold the requisite reference information. See Figure 5.16.

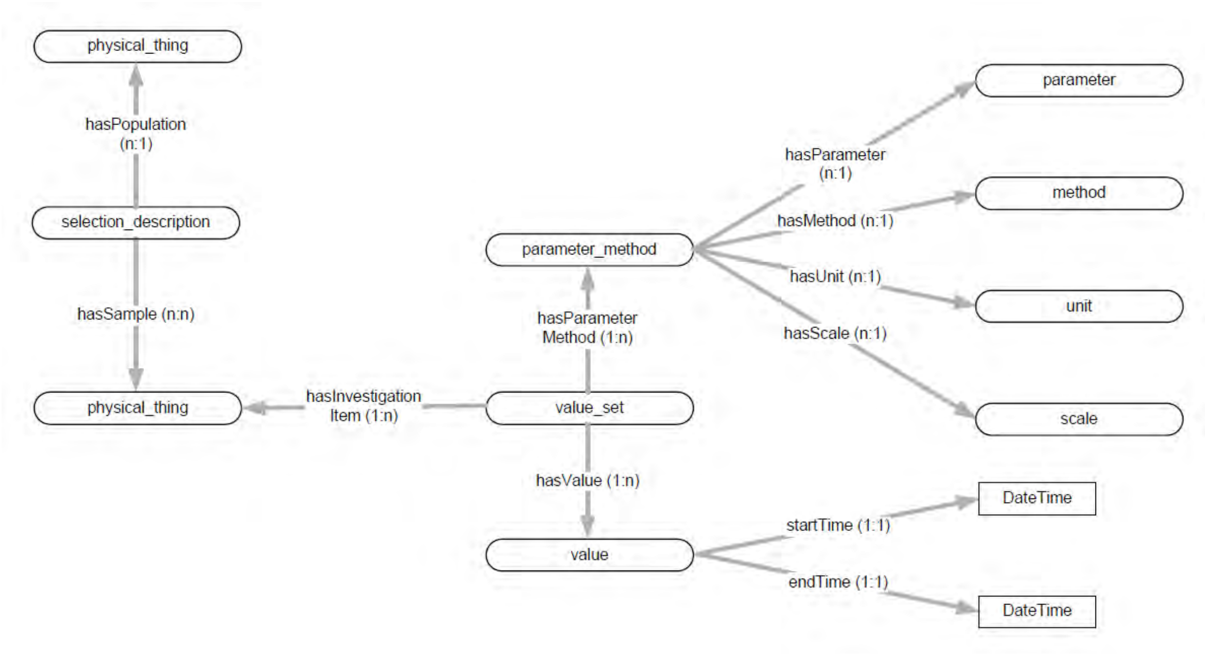

Figure 5.14 Basic Classes and Relations In SERONTO (from van der Werf *et al*., 2009; Figure 2).

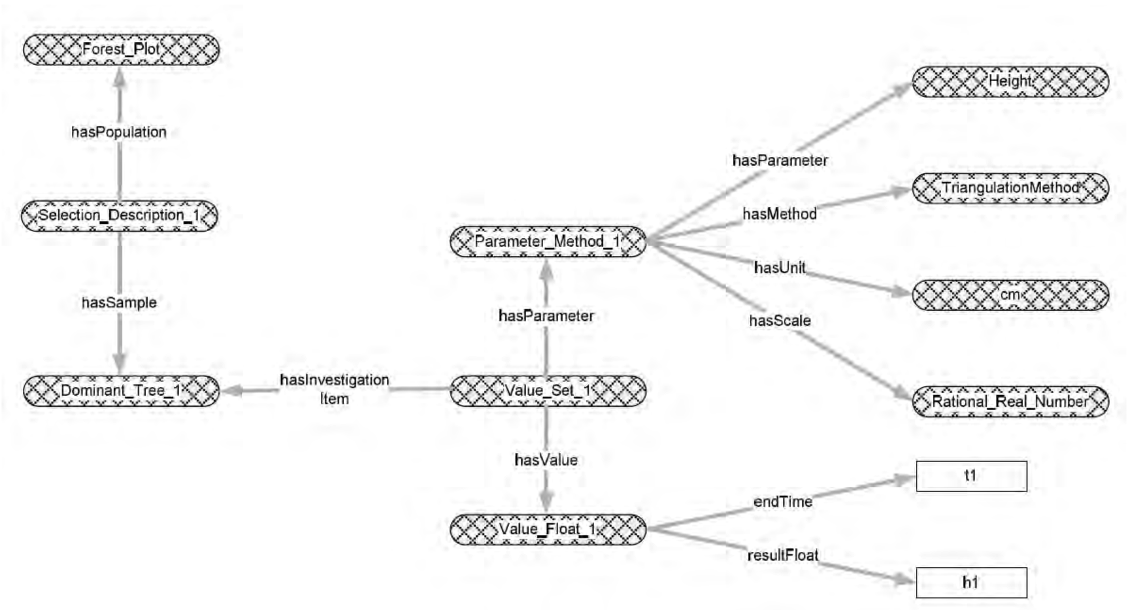

Figure 5.15 Representation of the Sample SERONTO Observation (from van der Werf *et al.,* 2009; Figure 3).

The extended schema of the SERONTO model focuses heavily on the constructs necessary to delineate the measurement scale associated with the observation and the characteristics of *Experimental Unit* selection (in addition to those classes already mentioned that are related to the *SelectionDescription* class). Measurements scale is considered an important aspect of a variable within SERONTO because it is from the scale that the possible statistical models can be inferred and this directs the assumptions that can be made in performing a particular analysis with SERONTO described data. Van der Werf *et al.* (2009) state that: “the scale can be divided into having string values (nominal_scale) or having numerical values (nominal_scale). The basic characteristic for the scale is the domain definition (i.e., which values can be chosen) and the interpretation of those values, which we call the linearity. The linearity has as possible instances (nominal, ordinal, interval circular and cyclic). The domain of a numerical scale can be defined by: the lower and upper mathematical boundary (this is not the boundary which can be measured by the device, but are the real domain boundaries); the mathematical base (decimal, binary, hexadecimal, etc); the numerical domain ($\mathbb{N}$: Natural numbers (including 0), $\mathbb{Z}$: Integer numbers, $\mathbb{R}$: Real numbers, $\mathbb{Q}$: Rational numbers (integer division)) and the modulo for cyclic or circular measurements. The domain for the nominal scale is a reference_list”. See Figure 5.17.

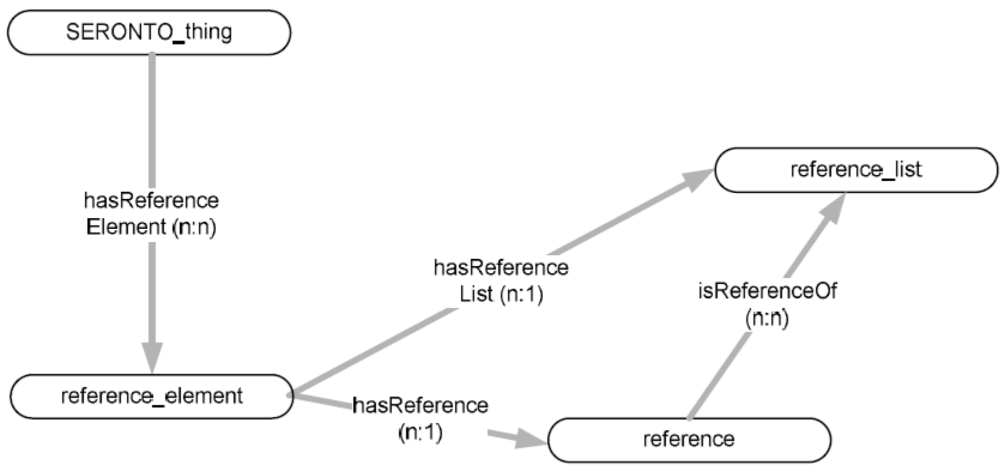

Figure 5.16 SERONTO Reference Classes (from van der Werf *et al.,* 2009).

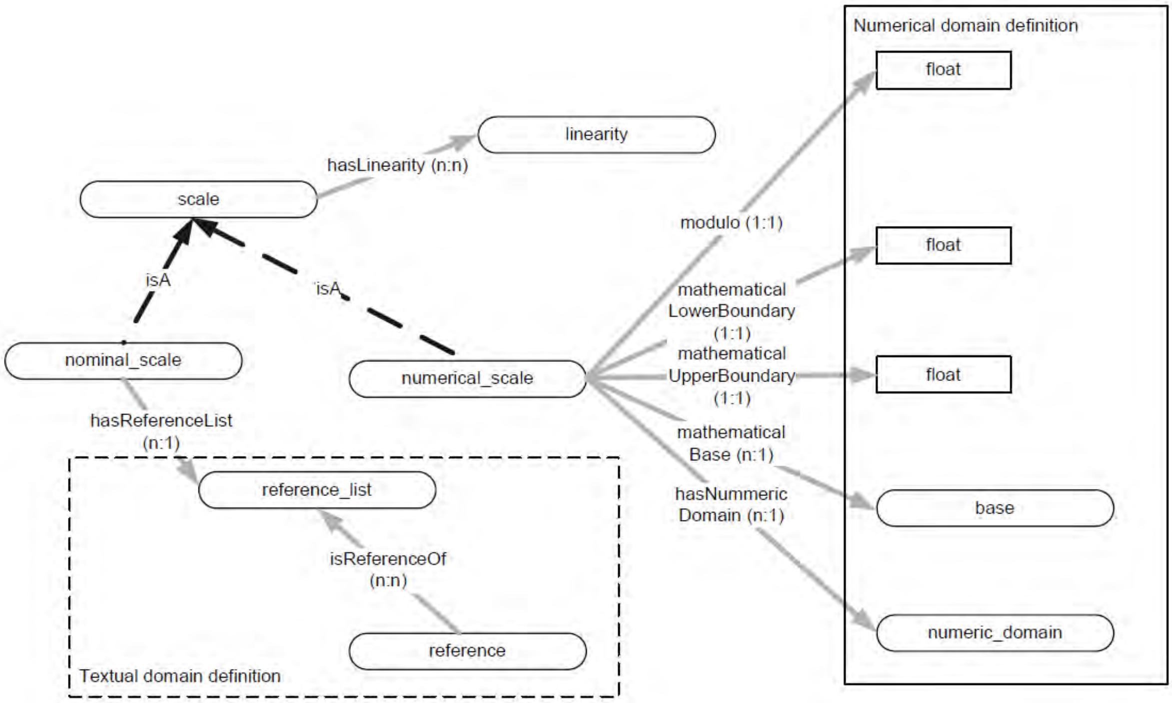

Figure 5.17 Classes For Describing The Scale Of A Parameter (from Figure 7, van der Werf, 2009).

SERONTO is not mapped to an upper ontology but derives its core classes from a concept called *SERONTOThing* (similar to OWL *Thing*). There are three major categories of class below *SERONTOThing (AbstractThing, PhysicalThing* and *ReferenceCatalogue)*.

OBOE, the other ontology mentioned by Catapano *et al.* (2011), is the Extensible Observation Ontology (Madin *et al.,* 2007) which was created within the Science Environment For Ecological Knowledge (SEEK) Project. SEEK is a collaborative project between many institutions striving to build a cyber infrastructure for ecological, environmental, and biodiversity research and to educate the ecological community about eco-informatics.

In OBOE the ontology separates *Observations* from the *Entity* being observed (See Figure 5.18). The *Observation* has a *Measurement* while the *Entity* has characteristics, and the *Measurement* is then of a specific characteristic. The *Entity* class is the extension point for including domain vocabularies (see Figure 5.19). *Observations* can occur 'within' a context, which in turn is also considered an *Observation*. Properties associating contextualised observations are transitive.

In comparison to OBOE, SERONTO appears to have additional value because of its use of *Reference Lists*, its superior ability to describe applied methods (through the chaining of methods and methods that can encompass other methods), the introduction of a *SelectionDescriptions* class explaining the origin of a research object, a time stamp that can be bound to every value and the provision of templates for specific domain use (Peterseil *et al.,* 2010).

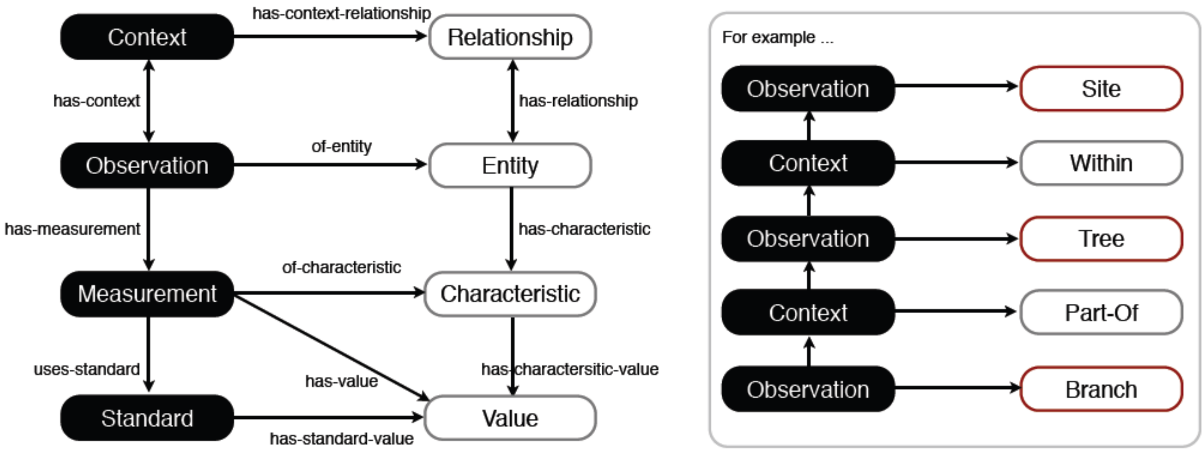

Figure 5.18. OBOE Core Structure (Berkley *et al.,* 2009)

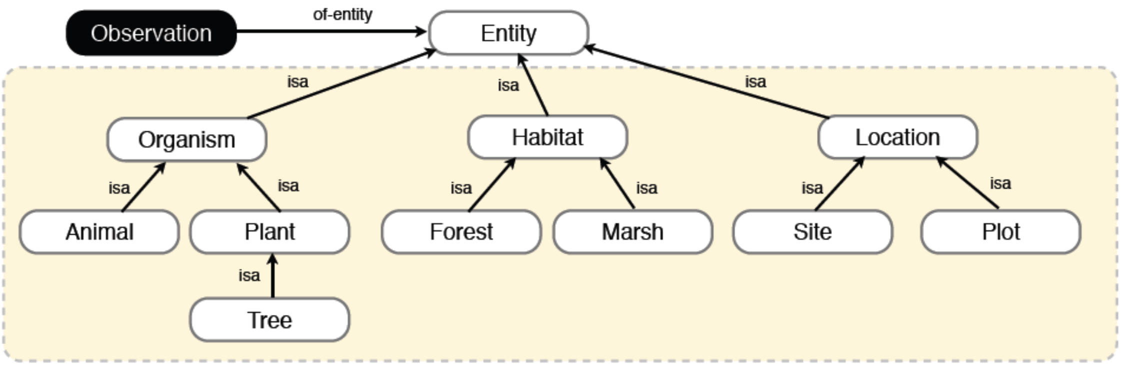

Figure 5.19 OBOE Domain Extension Point With Examples (Berkley *et al.,* 2009)

Both OBOE and SERONTO anticipate the injection of domain concepts through extension points, which could be supplied by FCATOWL. Additionally, because FCATOWL was designed firmly with biology data in mind it accommodates the inclusion of "scale" which has been shown to be of fundamental importance to biologists, as demonstrated through its prominence in SERONTO. FCATOWL has not, however, provided the level of detail associated with "scale" descriptions that are apparent in SERONTO. Scale is expressed through FCATOWL:*ReferenceSpaces* using a simplified taxonomy adapted from Probst (2007). The Probst (2007) taxonomy purposefully omits the "nominal scale" but he does not offer an alternate ontological classification. FCATOWL therefore includes nominals in its adapted taxonomy to appropriately accommodate the needs of the biological community. Like SERONTO, FCATOWL is also capable of holding reference-type information but this is achieved via OWL annotation properties. For example, using the individual FCATOWL:

*TaxonConceptName (*an individual of the *FCATOWL:NonQualityAttribute* class) for demonstration purposes, Figure 5.20 shows some annotation associated with this individual, which refers to a specific taxonomic reference list. Annotation can be ascribed to any FCATOWL class or individual.

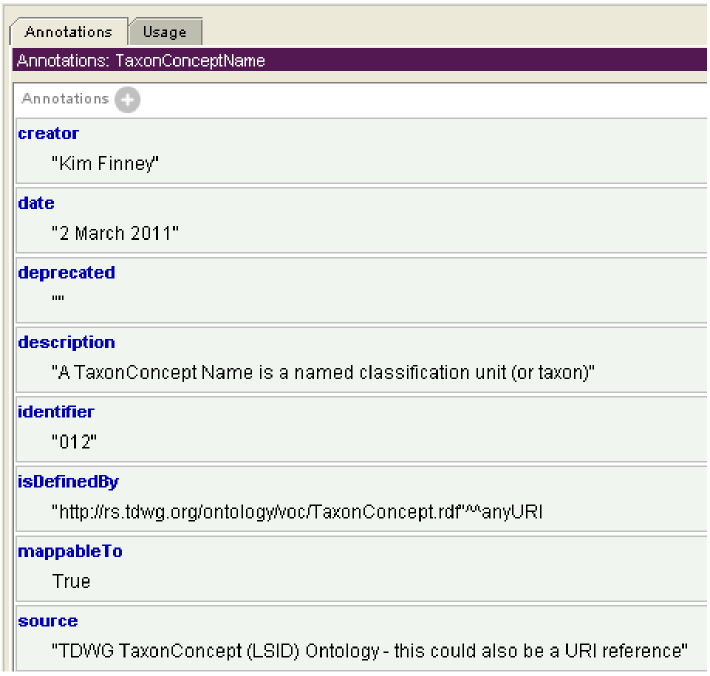

Figure 5.20 Annotation in FCAT Used For Reference (and Administration) Purposes

Summary of 'Architectural Fit' Using Ontology Alignment

In the absence of any ontological models in the SCAR and AODN communities with which to compare FCATOWL and because SCAR and AODN annotated data services should be interoperable with the semantic infrastructure of intersecting communities, a comparison of FCATOWL with three other observation-centric physical and biological ontology models, was undertaken. The assessment, demonstrated that *Feature Types*, as defined in FCATOWL, could be used as a ready source of domain concepts for these ontologies. This was particularly the case for the Sensor Observation related ontology, since FCATOWL and the Sensor Network Ontology (SSNO) share the upper-ontology DOLCE, even though some common model concepts had been treated slightly differently as would be expected from modelling *Feature Types* from different perspectives and against different use-cases.

Despite not sharing an upper ontology with the biology ontologies, from an information modelling perspective FCATOWL was compatible with both biological models as a potential source of domain concepts. For example, the inclusion of 'scale', 'sampling method' and 'units of measure' in FCATOWL's semantic signature enables the possibility to map from an FCATOWL *Feature Type* and an associated *QualityAttribute* to SERONTO's *Physical Thing* and its associated *Parameter* in a manner that preserves the intended semantic definition of the domain concept and its sensed property. In OBOE, FCATOWL *Feature Types* and *QualityAttributes* could be a source of OBOE *Entities and Characteristics*.

5.2 Current FCAT Limitations

It has been argued, and demonstrated throughout discussions in the previous section, that the Feature Catalogue content model (FCAT) and its various demonstrated implementations and associated REST-based query interfaces are collectively useful for creating semantic annotations consumable in the broader frameworks being established to create semantic Web services delivering scientific data.

One of the primary use-cases for which the Feature Catalogue (and associated services) was constructed was to facilitate the definition and publication of the key concepts that are used in the construction of AODN and SCAR data services in order to reduce unnecessary service heterogeneity (within those particular communities). The Catalogue is viewed as a source of re-usable terms that communities agree upon for use in certain scenarios, particularly during the construction of OGC data services. From a governance perspective the Feature Catalogue informally holds the details of the “contracts” that members of a community make with each other regarding the subjects (e.g., features and properties) of their data transactions. Defining and characterising these subjects informs the necessary syntactic structure and semantic content of service data models. Governance requirements also affect some aspects of how the Catalogue and its services should function. To completely fulfil its intended role, FCAT has some limitations which should be addressed before moving to an operational implementation. These limitations are now elaborated upon.

5.2.1 *NonQualityAttribute* Typing

In the FCATOWL semantic model, which attempts to stay as conformant as is possible to the Generalised Feature Model, Feature Type attributes are split into two types; *QualityAttributes* and *NonQualityAttributes*. *QualityAttributes* are further characterised according to their *QualitySpaces* which provides a rich source of semantic discrimination for these types of attributes. Since it is *QualityAttributes* that are the inhering qualities that are sensed (and hence observed or measured) these are often the dominant foci of scientific datasets. *NonQualityAttributes* by contrast in FCATOWL are a type of DOLCE Ultra-Lite *Object*. No further discrimination is currently applied. For example, a *PlatformName* (attribute) would be of the same class (i.e., DOLCE *Object*) as say a *Manufacturer* (attribute) in FCATOWL. Ontologically these are dissimilar concepts (e.g., a *Manufacturer* may in reality be typed as a *DOLCE Role* and *PlatformName* might be a type of *DOLCE*

Description). To cast these types of feature attributes in an enhanced ontological framework there would need to be some re-arrangement of the FCATOWL model. It would be possible to create a *NonQualityAttribute* as a DOLCE *Role*, rather than an *Object*, and still link it (i.e., the *NonQualityAttribute*) to a Feature Type. The *NonQualityAttribute* role would be something played by a DOLCE *Object*. In this scenario, attributes could be established as any sub-class of a DOLCE *Object* but be linked back to the *NonQualityAttribute* through a "playsRoleOf" type relation.

5.2.2 Attribute Enumeration

It is often the case that attributes (*Quality* and *NonQualityAttributes*) are drawn from a set of predefined enumerated values. FCATOWL, as is, has not modelled this scenario (although this is modelled in the relational database view). In future versions of FCATOWL this situation should be catered for by including an "hasEnumeratedValue" data property. The "hasEnumeratedValue" would have as its domain the union of an FCATOWL *NonQualityAttribute* and *QualityAttribute* and its range should be an OWL2 data range property (e.g., "DataOneOf"), which enumerates a set of literal alternatives. This approach would enable individuals of the classes *NonQualityAttribute* and *QualityAttribute* to link to enumerated lists, where required.

5.2.3 Properties and Access Service Descriptions

In their Feature Catalogue model, Stock *et al.* (2010) explicitly defined a range of 'property types' that could be used to express the relationship between Feature Types. FCATOWL has only modelled a few additional relationships for demonstration purposes (e.g., "synonymFor"; "hasRelationshipObservationFOI"; "hasRelationshipSampledFeature"; "hasRelationshipObservationPart"). There are also existing DOLCE relations that it is expected would be used quite frequently and these include:

(a) the "partOf" relation which would identify that one FCATOWL *Feature Type* was a host to another FCATOWL Feature Type, in which case the second Feature Type would need to be realised through its "proxyFor" relation as a type of DOLCE *Feature.* DOLCE *Features* are parasitic entities such as a 'lake's edge' or a 'mountain's peak', and
(b) the "constituentOf" relation that models scientific granularities, for example the relationship between a body, its organs and its cells,

Apart from additional property typing, the HTML output versions of the various prototype Catalogue service descriptions could also be enhanced by the inclusion of embedded RDFa to add semantic content (similar to that demonstrated for SKOS output). This would be best achieved by decoupling the HTML and XML hREST service descriptions.

5.2.4 Resource Identifiers, Stability of URLs and Resolution Methods

From a governance perspective service consumers require certainty around the persistence of the URIs used by providers to identify resources (Cryer *et al*., 2009; Cox, 2010b). More-over these identifiers also need to be actionable (e.g., resolvable). Actionable identifiers should contain information which locates an appropriate resolution service if presented to a suitable client. HTTP URIs are actionable by default. An HTTP URI necessarily begins with "http://" and thus is recognisable by its structure. The HTTP system provides mechanisms for clients to access a data object by its associated identifier without any need for further customisation. The character of an HTTP identified object might change, but it will always have the same identifier (for example a Feature Type might gain additional Attributes, but it is still the same Feature Type).

The FCAT services all use HTTP URIs and are therefore already "actionable", however, the first-order resources identified by the various REST-based services in this thesis (in Chapter 4) should be allocated a suitable persistent identifier. This was anticipated and current identifiers in the REST-based services should be considered placeholders only, pending a broader community discussion on the preferred syntax for such identifiers.

Currently the REST service syntax identifies a resource using a simple resource name type (e.g., FeatureType or Attribute), separated by a "/", followed by a number (or code). Information currently prefixed to the resource identification snippet includes:

> http://domain name/servicetype/servicecodehandle/servicename/.

Life Science (persistent) Identifiers (LSIDS (TDWG, 2011b)) used in biology and ecology often include a domain name, a naming authority and a namespace followed by an identifying term – usually a number. This might be a useful pattern to build upon. Registration of a domain name that is not directly tied to an organisational entity would aid in providing persistence (Cox, 2010b) and ideally, a single canonical URL should exist for each first order resource.

A postulated canonical URL could be the persistent global identifier for a resource which should normally be the URL that is book-marked to represent the Catalogue resource. Because there are different REST interface URL's pointing to the same Catalogue resource (which each supply different formats and descriptions at differing levels of granularity) by design in FCAT services, some additional work at the server end is required to make sure that all forms of the service are still accessible to Web clients (not just the canonical form).

A solution could be that the Catalogue host's server could issue an HTTP '303' redirection in response to any requests for this canonical URL (when anything other than the canonical service is required, recalling that a Web client can negotiate with a server about what types of service it is willing to accept). The server would redirect the client to an acceptable alternative URL, using a "Location HTTP header" that provides the URL of the alternative document (i.e., the appropriate REST query string for the service in question). Content negotiation should use "HTTP Accept headers" (as previously discussed) to re-direct a client to the most appropriate service (Sauermann and Cyganiak, 2008).

The main drawback, however, with this one-step redirection method is that there is insufficient information stored within an HTTP "Accept" header (using existing standard MIME types and sub-types) for a client to differentiate the different content views rendered by the different service types. Some non-standard types would be required to describe the various information view/format rendering combinations available. Alternatively, the '303' re-direction could present a semantic service description document (e.g., in SAWSDL) that the client could parse before choosing the appropriate URL (i.e., service rendering). The benefit of the one-step re-direction is its simplicity and it requires fewer steps, but the approach needs the client to understand the non-standard media type pairings. The alternate approach (using SAWSDL) permits more standardised terminology for client interpretation, at the possible expense of performance in accessing the resource. The client would also probably need to be thicker than a client used in the former approach (because it would need to be semantically enabled).

5.2.5 Summary Of Limitations & Further Research

Some of the limitations just mentioned (e.g., the need for a better semantic signature for FCAT *NonQualityAttributes,* facilities for ascribing enumerated values and enhancement of FCATOWL properties) can all be remedied through modifications to the ontological model. Other limitations mentioned are not so easily rectified.

Establishing meaningful URIs that can take account of the 'granularity' with which resources should be described and then made accessible (using a REST-based approach) is an open research question. Janowicz *et al.* (2011) describe a RESTful Proxy service that they had developed that complements a Sensor Observation Service in which they extract observation data using a REST-based approach. As in this research they explain the difficulties associated with constructing RESTful query strings, with appropriate resource sequencing, sufficient to expose representations of 'features' that encapsulate required aspects of semantic context (for a given situation). They conclude that there is no "context-free sensor data" and that "the selection, processing and publication of data using URIs will involve making certain decisions about features-of-interest" and will rely on more tightly defining the way in which "observation offerings" are made in Sensor Observation Services.

They also state that finding the right balance between crafting sufficiently expressive URI-based query strings and the extent to which navigation can be used by links that are embedded within provided resources needs to be determined by gaining more feedback from the user community. In this research similar problems were tackled by defining different services that 'chunked' information at different levels of aggregation. But it is acknowledged that this is an area requiring more investigation, especially since there are existing deficiencies with some of the general IT standards (i.e., lack of expressivity in HTTP *accept header* types) that limit the ability to label and hence select different types of (granular) service offerings.

5.3 Community Readiness For Semantic Annotation Approaches

So far in this chapter the artefacts developed in Chapter 4 have been evaluated, mainly with respect to the 'goodness of architectural' fit with existing community (and connected) infrastructure, from an information modelling, standards and technology implementation perspective. Along the way short-comings have also been noted. But little attention has been given to the 'readiness' or otherwise of the participating communities to embrace semantic annotation (and the artefacts developed in this research). This section examines how the research in this thesis might be taken forward cognisant of the current 'capabilities' of these participating communities.

In developing the FCAT Feature Catalogue and its various access mechanisms, two different content management paradigms were employed, one based on a triple-store approach and the other using relational technology linked to a set of RESTful services. The purpose of trialling the two approaches was to see which one offered the easiest implementation path and therefore the best prospect for a sustained development effort by the communities participating in this research.

5.3.1 Evaluation Of Implemented Approaches Given Community Capabilities

A comparison of these two approaches has therefore been made on the basis of several factors: (a) the availability and robustness of enabling tools; (b) software development effort required; (c) skill requirements and (d) flexibility of the chosen option to meet current and future needs. Table 5.2 summarises subjective assessments made by the author against each of these factors. These assessments were based on the author's knowledge of the capabilities of the key data providers in the AODCJF consortium and AODN Development Office and those peak institutions representing the 25 nations that are a party to SCADM, plus experiences gained in conducting this research. The terms "most" and "few" are used throughout Table 5.2 as relative terms since the data was not quantitatively derived and is qualitative in nature.

Table 5.2 – Summary Evaluation of The Two FCAT Implementation Methods

Factor	Ontology Development Environment and Triple-Store Issues	Relational Package Issues
Tools	• Many commercial and open source semantic development tools are now available (see http://www.w3.org/2001/sw/wiki/Tools and in particular http://www.mkbergman.com/sweet-tools/). The degree of tool robustness and level of ongoing maintenance is however highly variable. • Protege is a solid open source tool but many existing plug-ins do not work with the latest releases. Current visualisation tools (plug-ins) are limited, not functionally rich and are not particularly robust. • TopBraid™ was a solid development environment for use in conjunction with creating an Oracle™ triple store. This product costs $3000 (in 2008/9) for a single user license. Some communities are now opposed to using commercial products over open source approaches. • Different semantic reasoners may detect different taxonomic inconsistencies. Error messages produced by reasoners are often unhelpful in guiding an ontology developer to remedy inconsistencies and unsatisifiable conditions. Ontology debugging support is relatively poor. Methods are still developing for addressing this (see Wang *et al.*, 2005; Lehmann and Bühmann, 2010). • Most institutional community members represented in SCAR (SCADM) and in the AODCJF/AODN are not using triple store technologies. • Many standards associated with semantic web implementation are immature, still under refinement, not well-exercised (for example Terp).	• Very mature commercial and open source tools available for supporting relational database development and maintenance. • Mature and well exercised standards available for querying relational database content (e.g., SQL). • Wide availability of relational database connectors for many open source and commercial web development products and programming languages. • GIS web servers used by communities to assemble and serve out data as web services are currently mounted over relational database back-ends. • Most community members are not yet familiar with semantic environments like Jena that sit over relational systems to provide persistent storage for triples.
Development Effort	• With no knowledge of OWL, ontology editing tools, or triple-stores, community developers (even if they are already	• Providing HTML, XML, SKOS and OWL-based message responses through RESTful queries, by accessing FCAT data

	skilled in complementary areas such as relational database administration and programming) will need to devote time (e.g., 6 man-months) learning to work competently within these environments. • SPARQL or XQuery end-points are required for providing access to triples. Since the FCAT semantic data model will be opaque to users (clients), a range of template queries or query interfaces would need to be developed. An additional investment of time is required to learn these technologies. • SPARQL queries over OWL can be verbose and complex so more compact versions (of SPARQL) may need to be mastered (e.g., Terp). • Specific lifting and lowering scripts may also be required (particularly if the intent is to provide responses in plain XML, HTML and OWL without an associated relational store).	modelled in a relational data store, can be achieved using technologies that communities are very familiar with. • Data mappings, however, particularly as demonstrated for SKOS can be complex and changes to the underlying data model and/or message format can trigger the need for frequent (and perhaps not insignificant) code re-writes. • AADC programmers, unfamiliar with developing RESTful services and SKOS, found it difficult to undertake the programming task, despite having been given detailed requirements, standard sample response templates for every resource type (in each service), along with sample syntax for the REST queries and a populated prototype database. Manual acceptance testing of code operation was slow and tedious and eventually XPath was used to semi-automate testing and de-bugging
Skills	• Most current institutional members of the communities in SCAR and in the AODN, who are data providers, have almost nil experience in creating, populating and administering triple-store environments. (There are some notable exceptions but these institutional members predominantly align themselves with more specialised communities, for example those developing and deploying Sensor Observation Services and so they don't currently view themselves as part of the AODN or SCAR communities, even if institutionally they are). • Very few community members would identify as ontologist and most would not be familiar with developing ontologies in OWL.	• Most institutional members deploying data services are skilled in relational database administration, XML-based languages such as GML, and the administration and deployment of spatial data web servers (such as GeoServer and Mapserver and the use OpenLayer geospatial development libraries [http://openlayers.org/]). Many are also familiar with standards such as XSLT (Clark, 1999) and XPath (Clark and DeRose, 1999). • Some members would be familiar with the concepts behind RDF and SKOS and may have used both.
Flexibility	• The use of triple stores (native, or over relational database management systems) will become more pervasive and the requirement for semantic typing is growing. Back-end solutions such as Oracle, which provide both RDF and relational storage models in parallel with connections permitted between the two models, provide the flexibility to run hybrid solutions (Oracle, 2007). • Reasoning over data will require development of an OWL ontology. There will be an increasing demand for FCAT	• New applications such as D2R server (Bizer and Cyganiak, 2006) enable RDF and HTML browsers to navigate the content of relational databases and applications to query databases using the SPARQL query language. This is achieved through a D2RQ mapping configuration file. • However, if provided services are to include an OWL-based message response it might become increasingly difficult to maintain the FCATOWL model in a relational store without the

	content in OWL form (as a source of domain concepts). • Semantically-enabled demonstrator projects are now relatively common. There are many interesting test-beds (Bugel and Hilbring 2007; Klein 2007; Henson *et al*., 2009; Zhang *et al.,* 2010) addressing semantic problems faced in establishing data exchange networks. Implementations, however, would be characterised as bleeding and leading edge. Many institutions (and in fact nations) are still struggling to instantiate spatial data infrastructures – without semantics (Comert *et al*., 2010). Pervasive semantically enabled infrastructures are still some way off.	support of an overarching semantic tool layer. This is because each Feature Type in OWL obtains its formal semantic class through the "proxyFor" relation, where the range is any type of DOLCE *Event* or *Object*. New class information would need to be stored frequently and Feature Type individuals would then need to be asserted as being of those class types. This is achievable by creating new relational database entities to manage this type of content, but at some point an assessment would need to be made about the extent of the maintenance overhead being created by NOT using dedicated semantic management tools (to manage semantic content).

Having initially considered the two approaches trialled as an "either/or choice", as the two methods were unpacked in Chapter 4 it became apparent that elements of both methods could be harnessed to provide a good future-proof solution for the communities in question, particularly taking into account the factors listed in Table 5.2.

The communities participating in this study are still in the early phases of learning to create, administer and publish XML (GML) based Web services. Some institutions and data management hubs (such as the AODN Development Office) possess a locus of expertise that by exception has the capacity (given incentives and training) to implement semantic components that could serve the whole community. It is important, therefore, that these institutions be encouraged to take the lead in instantiating the technologies necessary for enabling semantic mark-up of community deployed services in order to improve the interoperability of institutionally-based data services. Where semantic enablement has been most effective is in situations where there is a balanced framework of easy-to-use tooling, good governance and demonstration, coupled with a real-world need (Smith *et al.,* 2007, Finney, 2007).

Short-to-Medium-Term Approaches

OWL is useful for reasoning and for tightly specifying semantics, but this level of formality is perhaps beyond where the AODN and SCAR data providers and consumers need to be at the present time. A phased approach from "less formal" to "formal" could be used to introduce the notion of semantic annotation. As a first pass the community could agree to establish a repository based on the Feature Catalogue relational data model, using traditional technologies (e.g., RDBMS and/or XML) with which the communities are already familiar. HTML, XML and SKOS or plain RDF RESTful services, designed initially to deliver simple payloads would be available to annotate community metadata records with Feature Type descriptions and these descriptions could also be linked (via XLink), as appropriate, in instances of GML services. This repository would additionally require a user-friendly front-end (Web-client) for facilitating the community governance and moderation of Catalogue content (the implementation details of which should be opaque to most community users).

There are, however, definite advantages to also embedding an FCAT (schema) model within GeoNetwork from an interoperability and software leveraging perspective. If using a GeoNetwork-assisted solution it is suggested that an FCAT ISO 19110 profile be serialised and imported into GeoNetwork that is capable of leveraging linked SKOS-based thesauri to populate various parts of the FCAT semantic model (e.g., Feature Type concept, Feature Type Quality Attributes aka parameters, units of measure, datum and collection methods (instrument) definitions)). This particular approach is suggested because GeoNetwork Version 2.8 onwards is likely to provide extended support for SKOS thesauri by utilising more of the SKOS vocabulary model. Current GeoNetwork interface functionality is too limited (Pigot, 2011, pers comm.) to support manipulation of the full FCATSKOS model developed in this thesis (particularly given the SKOS extensions). However, by populating a standalone basic FCAT RDBMS model (which would be replicated in GeoNetwork as a profile of ISO 19110), with concepts drawn from community-maintained SKOS thesauri, this would provide a pathway towards more formal (Feature Type) semantic annotation within the context of existing GeoNetwork tooling. This method would be more formal in the sense that RDF could begin to be used by the communities as opposed to just plain XML (see Figure 5.21).

Encouraging the community to work on developing appropriately granular and consistent thesauri has several advantages. These thesauri can be used to seed FCAT RDBMS content (as various pieces of contextual information for Feature Type definitions) as well as become controlled terms for 'MD_Keywords' and hence be leveraged as discovery vocabularies by community Portal technology. Fox *et al.* (2009) reported that "date-time", "parameter", "instrument" and "spatial extents" are core search criteria for almost all observational datasets (within the purview of the Virtual Solar-Terrestrial Observatory). This quartet, with the addition of "theme" of interest; "units of measure" in association with "parameters" and arguably "Feature Types" are the base search criteria for Antarctic-themed observational data. Community-sanctioned thesauri can also be mapped using SKOS (Lacasta, 2007) to other thesauri used by intersecting communities to aid with discovering services available from these communities. This would be a significant step forward for the AODN and SCAR community users who already share data services with each other and other allied communities (but using poorly structured and managed vocabularies).

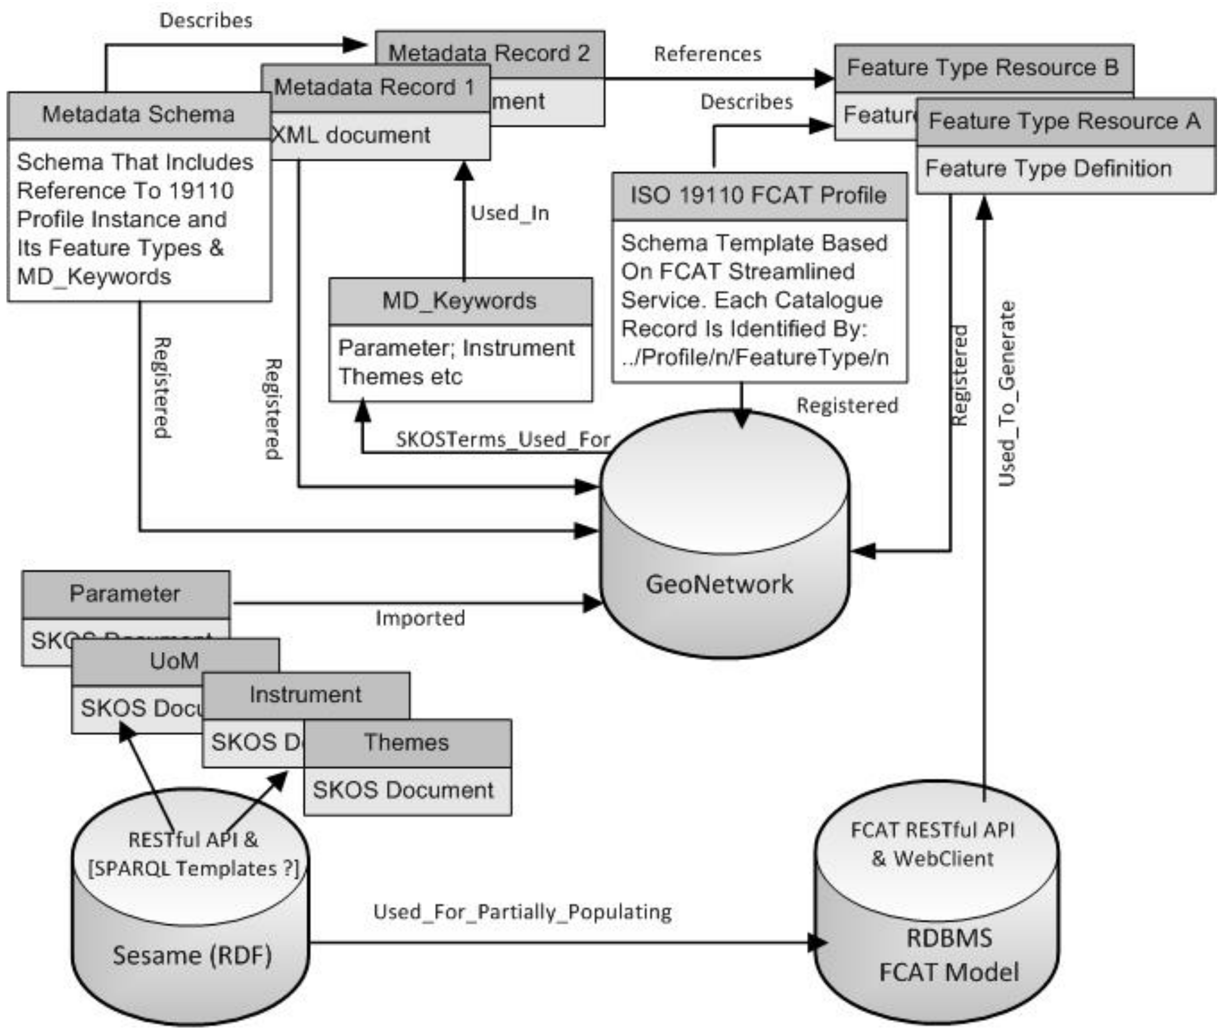

Figure 5.21 Short-term FCAT (Model) Implementation Option

The process of defining a SKOS-based version of the Feature Catalogue Model was initially trialled because it was believed to be a simpler model (than OWL) for communities to grasp. In hindsight, using SKOS in combination with REST lead the author towards an implementation that was completely sympathetic with the Linked Data paradigm, a paradigm that had not been adequately appreciated when this research commenced and which has only relatively recently risen to prominence in geo-spatial semantic research (Lehmann and Helmann, 2009; Barnaghi *et al.,* 2011; Lopez-Pellicer *et al.,* 2011). It is now believed that a Linked Data type approach, using interlinked vocabularies, anchored to an FCAT enhanced ISO 19110 Feature Catalogue model will be an extremely useful intermediate step towards full (and more formal) semantic enablement for the communities participating in this study. This is a very modular approach to using semantic content for annotation and accords well with advice from ontology experts interviewed in this thesis who rated ontology modularity highly when they were asked about what aspects of ontologies encouraged their re-use. Provided an ontology is modular it is easy to see how ontology snippets can ultimately be used in place of SKOS-based thesauri concepts in the Linked Data paradigm.

There must therefore be an easy route forward for ultimately linking Feature Catalogue content with OWL instantiations of the Feature Types and their associated contextual information. Because GeoNetwork can import RDF (SKOS model) formats (currently for thesauri manipulation) there is some immediate incentive for GeoNetwork users to configure a persistent semantic data store such as Sesame (in which community vocabularies could be easily managed with access available via Restful APIs and SPARQL templates). By starting in a well-bounded way, with simple knowledge models such as SKOS, infrastructure builders can become familiar with managing and administering RDF triple stores. There is then a potential future extension point for the later inclusion of OWL ontologies. OWLIM (Ontotext, 2011), an open source packaged storage and inference layer for Sesame that supports reasoning for OWL could, for example, be added later to provide a key part of a future semantic service layer.

Longer-Term Approaches

A two-phased approach to uptake of the research in this thesis is suggested for the participating communities. The first phase has just been described.

The second phase of development would incorporate an FCATOWL model repository. Ideally the OWL model should be managed via a triple store (such as Sesame or Oracle Semantic). This OWL repository should be seeded using the amassed content from the first repository (built in phase 1). A SPARQL end-point could be established over this repository and a number of SPARQL Query templates developed to extract Feature Type descriptions (along with RESTful APIs as developed in this thesis). These end-points, for exposing OWL descriptions of individual Feature Types, would then be made accessible from the Web-client established in phase one (alongside any existing services), so that any data providers wishing to anchor their annotations to a more formal semantic model, could do so (see Figure 5.22). The inclusion of an underlying ontological repository will not only provide for inference but will also provide for many new avenues to creatively visualise and search for semantic components (e.g., see jOWL (OntologyOnline, 2012); Stock *et al.,* 2009a; Paulheim and Probst, 2010; Larizgoitia and Toma, 2011).

However, an OWL-based encoding using triple-store technology and native SPARQL end-points (as opposed to templates) should not be attempted until the technology stack for using such approaches is more mature than it is at present. Ultimately the majority of community members, in the main, should not have to deal directly with the current crop of ontology editors and the often difficult processes of ontology creation. It is envisaged that new technologies will arise in the short-term to make the task of ontology establishment easier and that most community members should be re-using ontology fragments, rather than creating them from scratch.

In the two-phased approach, discussed above, it is believed that by becoming initially familiar with the tools required to support simple (but powerful and fully contextualised) semantic (but essentially non-ontological) annotations, the current barriers to making the next leap forward will be lowered for the communities in question. By using the content of the repository, built in phase one, to subsequently populate an OWL repository to be developed in phase two (based on FCATOWL), the community is less likely to be initially overwhelmed by new tools and techniques. The community's move from using less formal, to more formal types of expression should flow naturally from such a phased approach.

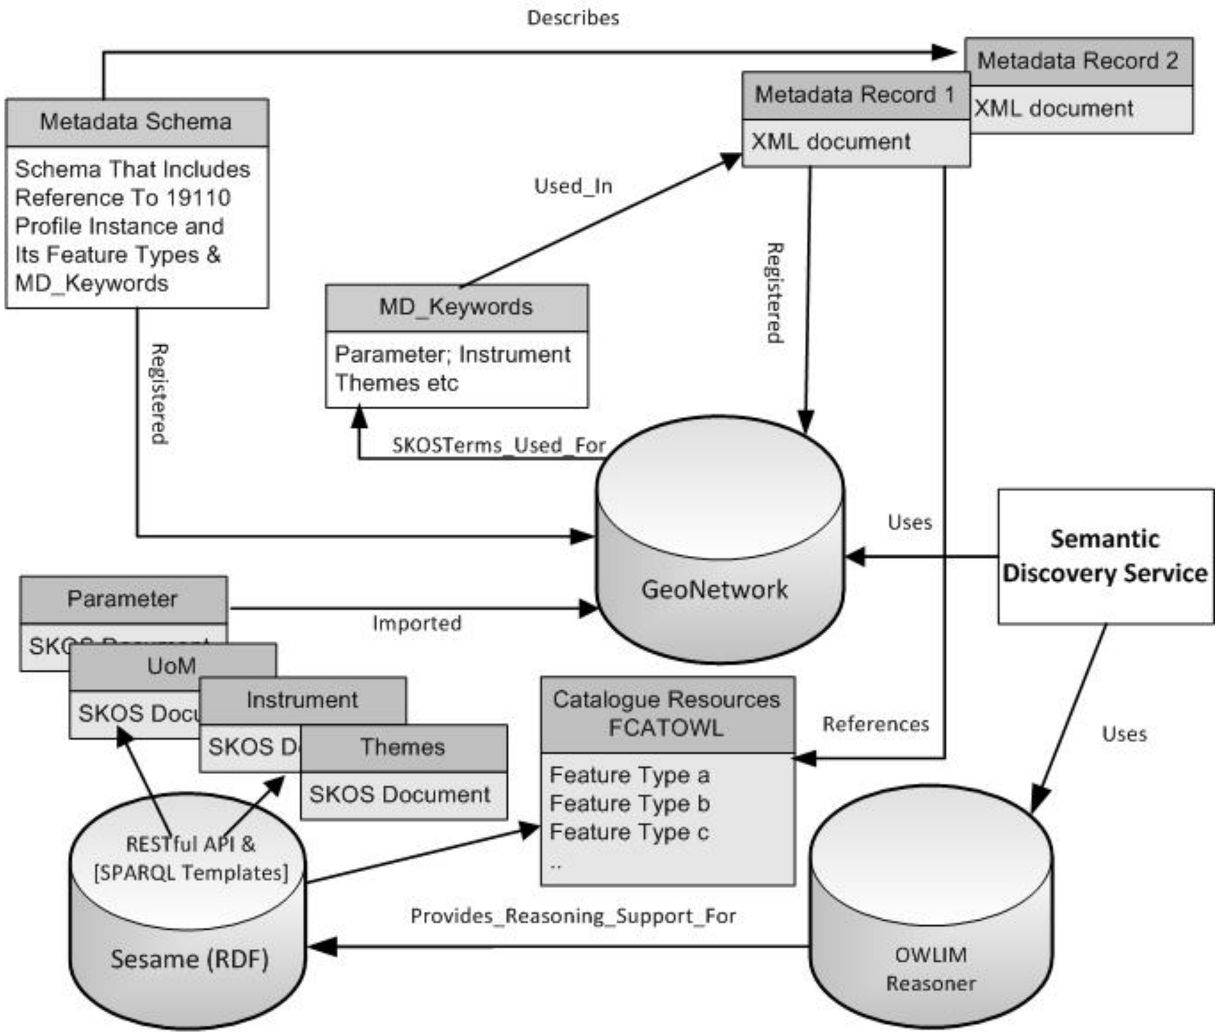

Figure 5.22 Longer-term FCAT (Model) Implementation Option

It is important to note that repository population (in phase one and particularly in phase two) should not be performed in isolation from other existing domain concept-building activities performed by other communities with similar goals and use-cases. Where-ever possible, to minimise effort and to increase interoperability, concepts identified and described by other communities (with similar or overlapping disciplines), should be mined and re-used to populate the FCAT Feature Catalogue. The next Chapter (6) provides the results of research into the factors that ontology experts consider are the most important for practically evaluating and selecting re-usable semantic content. It is hoped that with a blue-print for establishing a semantic repository and associated easy-to-access services, coupled with guidance (from Chapter 6) on how to most efficiently evaluate what content is already available, the communities covered in this study can make significant advances towards deploying semantically-enabled services.

5.3.2 Community Reaction and Feedback

Thus far, this research has demonstrated most of the components necessary to support semantic enablement of community data services (with the notable exception of a mock Web-client). It has also demonstrated that the approaches suggested are consistent with current trends in more broadly-based semantic-web activities and that a suitably populated FCAT Feature Catalogue would add value to the emerging semantic infrastructures of other external communities (as a source of domain concepts). In addition, the exploration of how to create simple and intuitive REST-based services for accessing Catalogue repository content has provided a greater depth of understanding about how to establish such services and how they can be used with existing technologies and standards to annotate scientific data services, particularly those conformant with OGC standards.

As this research developed, these key outcomes were continuously fed back to the AODN and SCAR communities (see meetings listed in Chapter 3). In some cases results were communicated more broadly to other allied communities through Conference and Workshop presentations.

In 2009 and in 2010 two detailed proposals were also written by the author, at the request of the AODC Joint Facility Board, to try to capitalise on Australian national data infrastructure funding made available by the Australian government via ANDS. Both of these proposals included a Feature Catalogue development component (designed to leverage aspects of this research). Whilst only one proposal was successful and funds were allocated for only one of the tasks outlined, each proposal was signed off in its entirety by all AODCJF and IMOS AODN Development Office (DO) representatives. This signalled an acceptance by AODN related agencies of the need for the types of semantic enablement presented in this thesis.

In September 2010, noting that funding from ANDS was unlikely (to pursue Catalogue development), the AODC JF Technical Committee convened a 2-day work-shop, to which the author was invited to lead discussions on how to prioritise activities associated with further establishing the AODN infrastructure (without additional funding). The minutes and outcomes of this meeting are presented in Appendix 17. As a result of this meeting, the Deputy Director of the AODN DO (and Australian GeoNetwork Development Coordinator) suggested that the FCAT extended ISO 19110 Feature Catalogue Model, developed in this research, be presented for discussion in the context of a possible ISO standards amendment and that the FCAT model be incorporated into GeoNetwork (see email snippet below).

```
Sent on 2 November 2010:

"One benefit of adding this into a release of the BlueNetMEST, or
Geonetwork, is that others could get to use it/comment on it.
However, if you prefer that we don't do this, thats fine - just let me know.

Paul raised the issue of the proposal going forward to ISO.
I've assumed that you have already been discussing these ideas with some of those
higher-level standards people, and have some thoughts on how to get these
suggestions adopted as part of revision to the ISO19110 standard?  If we can help
in the advocation process (by either supporting the good sense of the ideas, or by
declaring that the AODCJF is trialling it, or has trialled it, and finds it good),
then let us know....

...."
```

Community-related interactions to date, the general feedback garnered, current activities and agreed future community actions all demonstrate a shift in community perceptions about the need for semantic enablement that wasn't evident before this research commenced. In this regard it is considered that the goal of this research which was to "improve Web deployment of semantically described scientific datasets" has been facilitated (for participant communities). Through the conduct of this study it has been possible to answer RQ1.1 – *What characterise an ontologically-grounded Feature Catalogue that can support Antarctic science data publication through web services*" and to provide guidance on how such a Catalogue can be implemented. Prototyping these characteristics has confirmed for stakeholders, the utility of both informal and formal approaches to semantic information modelling and the value of REST-based services for annotation purposes.

5.4 Summary

In this chapter a summative evaluation of the research results reported in Chapter 4 was presented. It was demonstrated how the artefacts developed (i.e., the enhanced FCAT ISO-19110 based Feature Catalogue Model, the FCATOWL and RDBMS Feature Catalogue instantiations; and the REST-based access methods and query patterns) could be harnessed to provide semantic annotations for community-provided data services, using the existing infrastructure and standards of the AODN and SCAR communities. There were a range of possibilities for where annotations could be anchored. The alignment of the FCATOWL ontology with other observation-centric ontologies, used in semantically-enabled infrastructure of intersecting communities, was also evaluated. Effectively, what was being assessed in section 5.1 of this chapter was the level of 'architectural fit'.

The assessment was that the FCAT information model and its various instantiated derivatives could be used as a good source of domain entities, or features-of-interest in those systems that were reviewed. The semantic signature (created by contextualising a Feature Type) was highly compatible with other information models (when FCAT was considered as a source of domain entities). By richly contextualising the Feature Type entity in FCAT it was not difficult to see how FCAT Feature Type definitions, or the definitions of Feature Type QualityAttributes (or sensed parameters), could be discriminately mapped into existing observation-centric ontologies.

The REST-based access methods established for FCAT Catalogue content, in which unique HTTP URIs are provided for Catalogue resources, which themselves contain links to other resources, was not only an effective method for enabling semantic annotation within various parts of the OGC-services stack, it was also a transparent method for allowing any Web-based service to access FCAT Catalogue content. These REST-based access (and resource query patterns) broadened the utility of the Feature Catalogue as a source of semantic domain concepts.

Some limitations of the products that were developed were discussed and issues of resource identity and descriptive granularity were raised as significant issues requiring further research. This finding was corroborated by other recent independent research by Jancowiz *et al.* (2011).

Finally, this chapter provided "guidance" on approaches to semantic annotation for the communities that participated in the study, taking into consideration the current state of evolution of semantic technologies and community capabilities and practises. This guidance suggested a two-phase approach to semantic-enablement. Phase one involves leveraging existing infrastructure and tooling with which the communities are familiar, to commence the process of semantic annotation using informal semantic languages. Phase two, which would build upon achievements in phase one, and which would be executed when ontology tooling is more user-friendly and mature, involves utilising OWL and allied semantic technologies to realise a more sophisticated and formal ontological approach.

Chapter 6.

Practical Ontology Selection and Evaluation Data Analysis

In the previous chapter the results of the research related to (RQ1.1) characterising and prototyping an ontology-grounded Feature Catalogue and associated access services were evaluated and discussed. This chapter is concerned with presenting the research and data analysis conducted in order to gain deeper insight into how different scientific domains, who are semantically-enabling their data infrastructure, are approaching the task of ontology selection and evaluation. The key research question which is investigated is: *what typifies an expert-grounded ontology selection and evaluation framework that can support multi-disciplinary Antarctic science communities using Web services* (RQ1.2) ?

To assist communities like SCAR and the AODN to adopt ontologies as part of their data exchange protocols and infrastructure, it is necessary to establish practical guidance on how they could re-use existing (applicable) ontologies. The effort required to build good quality domain ontologies is significant (as evidenced from the literature outlined in Chapter 2 and from the author's personal experience in building FCATOWL and FCATSKOS). As an example, one of the most advanced ontology-building communities - OBO (Open Biological and Biomedical Ontologies), has been investing resources in ontology development since 2001 (Smith *et al*., 2007). Whilst OBO's focus does not overlap, to a large extent, with the observation-centric Earth Systems realms of the Antarctic communities in this study, the opportunity of capitalising on already sunk investments, of groups such as the OBO community, can only be realised if those domain experts building data exchange networks can understand how to apply the fruits of any existing ontology-building labour. To this end the research outlined next explores how ontology selection and evaluation is actually taking place in real-world projects, with a view to couching relevant findings in a selection and evaluation framework that could be applied by data network builders in communities like SCAR and the AODN.

As with the Feature Catalogue artefact development activities outlined in Chapter 4, the research presented in this chapter was grounded by interaction with experts. Experts who provided insight and information in support of the activities associated with this chapter were not, however, part of the SCAR or AODN communities, but were ontology practitioners, some of international repute. All of whom had been applying their skills in real-world ontology development projects (in scientific domains).

To address RQ 1.2 regarding ontology selection and evaluation practises, three sub-questions were posed to help focus the research (noting that one of the sub-questions - RQ1.2.2 has two additional nested sub-components). To recap (from Chapter 1, Figure 1.3), these questions were:

> RQ1.2.1: *What ontology selection and evaluation criteria are currently used across multi-disciplinary scientific communities (and are selection and evaluation methods consistent with those reported in the literature)*?,
>
> RQ1.2.2: *Is it feasible to derive a weighted evaluation criteria model in which criteria are rated according to importance* ? *If so, are evaluation criteria of equal weight or do some carry more importance than others (*RQ1.2.2.1)*?* AND *do relative levels of importance differ by application between scientific disciplines, or by any other discernible factor* (RQ1.2.2.2)?
>
> RQ1.2.3: *What evaluation measures can be used to assess evaluation criteria* ?

Figure 6.1 outlines the sequence of research that was conducted to address these questions. In this Figure, each step in the research process is pre-fixed by a superscript number (1 through 8) indicating that there were eight distinctive phases to data capture and analysis. Each of these phases is described in detail in this chapter, along with the results of data analysis. These sequence labels will be used later, as the data analyses unfold, to help anchor the reader back to this Figure (and hence the overall process that is being executed and the questions that are being addressed).

The research described in this chapter commences by developing and conducting a ‘Screening Survey’ which is primarily used to establish the suitability of experts as potential participants for this study. The survey, however, has another function which is to provide information that can later be used to help characterise (and stratify) experts to aid in subsequent data analysis phases.

Having selected and recruited suitable participants, each expert was then interviewed by the author and asked a series of ‘prompting’ questions designed to specifically uncover information and practises concerning:

- The type of ontology selection and evaluation ‘criteria’ used by the expert and/or communities in which they were anchored, when seeking ontologies that could be re-used to meet a community’s semantic requirements.

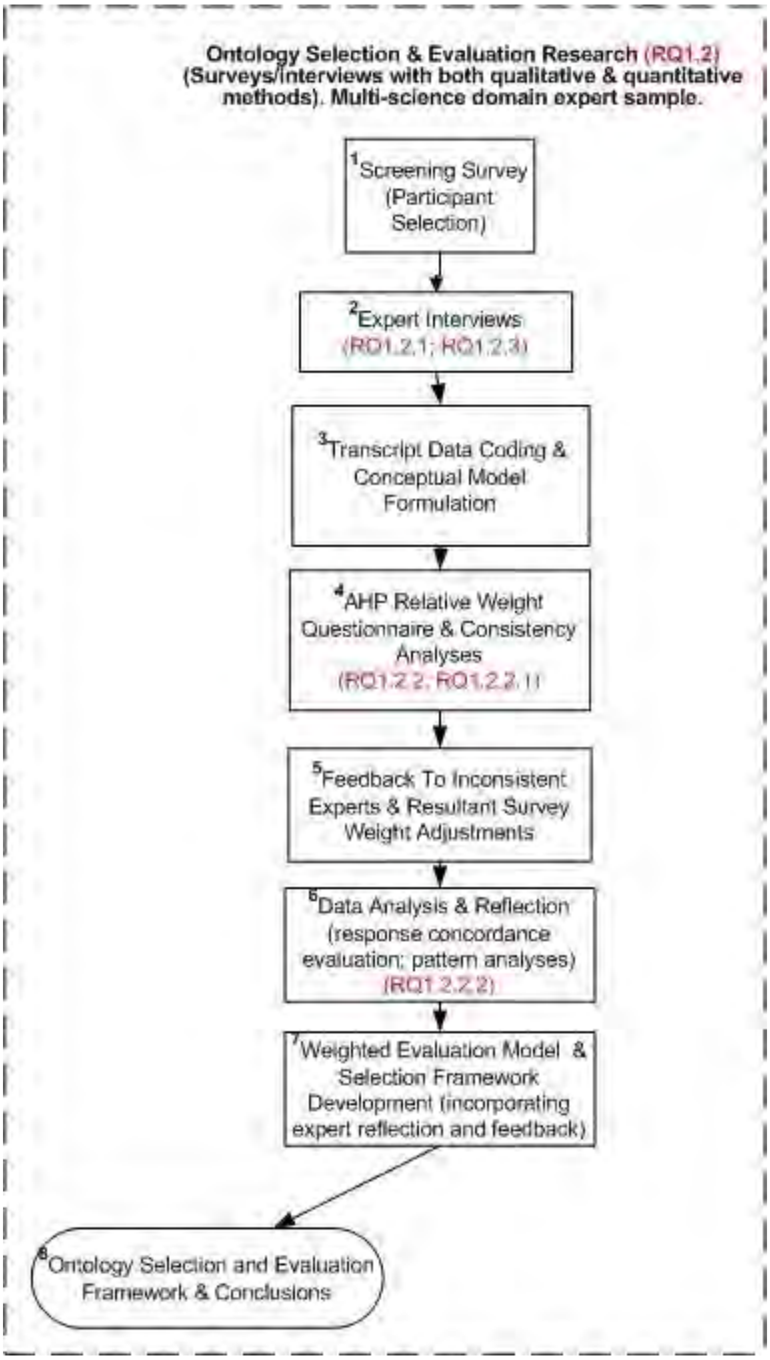

Figure 6.1 Sequence Of Data Analysis (excerpt from Figure 3.2 in Chapter 3)

- The types of ontology evaluation 'measures' (as opposed to criteria) being applied in ontology selection processes.
- The types of ontology selection and evaluation 'methodologies' used, and in particular whether these 'methods' were similar to those reported in the literature.
- The nature of ontology governance processes in communities developing and re-using semantic content (for the purposes of generating best-practise guidance).

- Ontology re-use practises in general that might hold 'guidance value' for multi-disciplinary Antarctic science communities wanting to semantically-enable their Web services.

These expert interviews had three functions. The first was to provide a thorough and comprehensive description of the ontology evaluation criteria being used in practise. This would enable the development of a hierarchical evaluation model, suitable as a decision-support framework, which could be used later for conducting an assessment of the relative 'importance' of ontology evaluation criteria, using the Analytical Hierarchical Processing (AHP) methodology (Saaty, 1980). The second function was to permit the generation of theories about practise relating to ontology selection and evaluation (incorporating issues of ontology governance). These theoretical conclusions, although generated as a consequence of analyses in this chapter, are presented and discussed in Chapter 7. The third function was to further build upon data from the screening survey to develop a comprehensive picture of the characteristics of each expert and the community(s) they represented, which could potentially be used in other phases of the research, particularly when addressing *whether relative levels of importance placed on ontology evaluation criteria differ by application between scientific disciplines, or by any other discernible factor* (RQ1.2.2.2).

To fulfil the first function, interview data was transcribed, analysed, coded and assigned categories (during phase 3 of the research, Figure 6.1) in order to create a hierarchical decision framework suitable for comparatively rating/evaluating ontologies (for re-use purposes). This result was achieved using purely qualitative approaches. Rather than use this evaluation framework to perform an ontology rating exercise, however, which was not the purpose of model development, experts were asked instead to complete a questionnaire that was structured so that a pair-wise rated comparison was made by each expert, between each element at each level in the evaluation model hierarchy (phase 4, in Figure 6.1). This revealed the relative level of importance that each expert was prepared to place on the various model evaluation elements that were derived from qualitative analysis of the interview data. The number of interviewed experts willing to participate in this phase of the research fell at this point to half the original participant number (i.e., '7').

The assignment of relative importance ratings, using AHP techniques, provided the raw quantitative material for a range of analyses subsequently performed in this thesis on these data. For example, internal checks were performed on the raw ratings allocated by each expert to determine whether the ratings that they had provided fell within an acceptable range of 'consistency' (i.e., consistency being a measure of the logical internal coherence between the various pairs of ratings provided). When internal coherence was outside of acceptable bounds, data was adjusted and fed back to the experts (phase 5, in Figure 6.1). Experts were asked if these altered data were acceptable, given the context provided in the feedback. Where there was agreement, expert provided raw data was subsequently adjusted.

In phase 6 (Figure 6.1) of the research process, further analyses were performed on the individual sets of ratings data to establish the degree of statistical ‘concordance’ between respondents (in terms of their ratings) and multi-dimensional-scaling was used to look for overall ‘patterns’ or ‘clustering’ in the responses from individual experts. In addition to these formal quantitative assessments, individual sets of data for each level in the model hierarchy, for each expert, were over-plotted and any discernible similarities in ratings between expert participants were analysed in detail through forensic inspection of the data. During this phase, considerable use was made of the ‘characterisation’ data collected during the Screening Survey and interviews, which was used to stratify experts into different groups, according to different sets of characterising dimensions. Any patterns or clusters in the ratings data between experts were assessed in light of these various types of stratification. Stratification was used as evidence to look for ‘justifiable’ reasons why patterns might exist. If patterns couldn’t be qualitatively justified, it had to be assumed that any ‘apparent’ patterns were no more than random, or that the data to hand wasn’t sufficient for elucidating why such patterns existed.

Although individual sets of ‘consistent’ importance (rating) data for each expert were available (at levels below CR <= 0.14), what was desired was a ‘group’ result for reported ‘importance’ ratings. Several measures of central tendency were therefore calculated to establish a ‘group’ result which was graphed and then fed back to each participant, along with the results of their individual assessments. Each expert was asked a number of questions, culminating in whether or not they were prepared to accept the group result.

In phase 7 (in Figure 6.1) of the research process a hierarchical, weighted evaluation criteria model was developed from a group result and provided to participants for review and feedback. Adjustments were subsequently made to the weighted model resulting from this feedback and a final model was developed. The weighted model was then augmented with specific evaluation measures, sourced from expert interviews and supplemented from the literature and bound together into an Ontology Selection and Evaluation Framework (in phase 8, in Figure 6.1).

Each of these different research phases are now explained in detail. Results in each phase are presented and discussed, where appropriate, in terms of the research questions that have driven the process.

6.1 Expert Screening Survey

The Screening Survey (step 1, in Figure 6.1) was a tool for evaluating an expert's suitability for participation in this study and a valuable means to help the author prepare for subsequent in-depth interviews (with the selected experts). The survey was created to target ontology developers and key personnel in communities known to be developing ontologies designed for application in either an OGC or TC 211 standards suite environment and/or for use in scientific data exchange scenarios. It was important to garner the participation of experts that reflected, to the best extent feasible, the disciplinary mix evident within Antarctic science, for example: oceanography, biology; ecology; meteorology; physics of the upper atmosphere, geology; glaciology; medicine and limnology. Importantly, the survey was also designed to:

- Introduce respondents (particularly those who would later go on to be interviewed) to the research questions, key terms and concepts that would recur throughout the research and to give these people an opportunity to start thinking about selection and evaluation issues in preparation for more in-depth questioning.
- Garner information about an individual, their associated community(s)-of-interest and their interpretation of key concepts to help begin to stratify and group experts in various ways which would be built upon later, as a result of in-depth interviews.

The survey itself consisted of a mixture of open and closed questions, conducted via a commercial Web-based survey instrument (QuestionPro), powered by SurveyConsole™. A survey template was built using wizards available in QuestionPro and where possible automatic field validation checks were included to minimise the chance of respondents only partially completing the survey.

As recommended by qualitative study experts, a pilot study was harnessed to pre-test the survey instrument (Baker, 1994; Neuman, 1999). Seven people (encompassing colleagues, friends and relatives) were contacted and invited to trial and then review the on-line screening survey. The aim of this exercise was to get volunteer reviewers to fine-tune the survey; review the email that would accompany the general invitation to take the survey and the associated participation consent forms (see Appendix 18 for the email sent to potential reviewers concerning the piloting of the survey).

Five of the seven people contacted tested the software and provided valuable feedback. Improvements made as a result of piloting the software included:

- Introduction of blank spacing to separate some questions to improve readability.

- Correction of a logical forking error (i.e., the software did not send the user to the right question at one of the forking points).
- Introduction of mandatory answers for many questions.
- Correction of minor typographical and spelling errors.
- Expansion of some boxes that were provided for user responses.
- Correction to an ambiguous question so that it's meaning was made much clearer.

The survey was then sent out widely via email to potential expert participants. Since these emails were sent to specific individuals, purposive sampling (i.e., sampling where the researcher chooses the sample based on who they think might be appropriate for the study) was used. In many cases recipients sent the email on to colleagues, so in these cases the author did not target the individuals involved, but in the main those who replied were part of the original cohort targeted by the author. For example, nine experts were referred as potential participants by some of the experts completing the initial survey. Six of these referrals appeared superficially suitable and were contacted separately and asked to take the survey. Three of these experts then went on to complete the survey. Fourteen respondents completed the survey in total. Of those that went on to participate in the study several experts were people who the author had worked with on committees, or with whom the author had been marginally affiliated due to professional work in the data management and data infrastructure-building sphere. None were close colleagues.

Thirteen people were ultimately assessed as suitable from a total of fourteen completed surveys. One expert had not practised re-use, nor had they developed ontologies for use in data exchange scenarios, so they were not invited to participate.

One expert ('JG') who did not claim to have ontology re-use experience, and instead stated that he did not understand the survey question on this topic, was included anyway (given his suitability was confirmed through a Web review of his work). Another expert ('JH') stated that his ontology use-case scenarios didn't involve scientific data exchange. Again, this expert was included anyway after verification of his suitability through a Web search of his published activities. Their suitability was tested and verified again during in-depth interview. Six of the responding experts claimed to have had direct experience with developing ontologies for use in ISO TC 211 standards environments (see Table 6.1).

Table 6.2 summarises basic information about those experts who responded and who were deemed suitable. Their sex, institutional affiliation and country of residence are not reported to preserve their anonymity. To be rated as suitable the respondent needed to be willing to participate in later phases of the research and needed to have experience (in a community setting) with ontology re-use, all within in a scientific data exchange environment. To preserve the privacy and confidentiality of experts, all of the results presented in this thesis are coded and any identifying material has been omitted from included material presented in appendices.

Table 6.1 Ontology Application

Expert Code	Use in scientific data exchange scenarios*	Use in Web service context using TC 211 standards
RH	1	0
WD	1	0
PF	1	0
JH	0	1
LL	1	1
DH	1	1
PM	1	0
CA	1	1
JG	1	0
RA	1	1
RL	1	0
SC	1	1
MH	1	0
Totals	**12**	**6**

*1= yes, 0 = no

Whilst eventual participant numbers are not high, the level of response received was not unexpected, since the field of expert ontology practitioners operating across science domains is not large (and certainly wasn't when this study commenced). Although the author would have preferred to work with a larger sample, Mason (2010) has asserted that in qualitative research, small sample sizes can be accommodated because one occurrence of a piece of data, or a code, is all that is necessary to ensure that it becomes part of the analysis framework. Frequencies (of data occurrence) are usually not as important as 'saturation' (i.e., when the collection of new data doesn't shed any new light on the issue under investigation (Glaser and Strauss, 1967). What was deemed important in this study was ensuring that those participating were diverse enough in background to enable the uncovering of the variety of views that are of interest and utility to this research. It is believed that this has been achieved through garnering participants from a range of scientific disciplines, countries, communities and institutions.

Table 6.2 Summary Of Experts Considered Suitable - Using Screening Survey Questionnaire

Expert Code	Primary Community Affiliation
RH	Taxonomic Database Working Group (TDWG)
WD	OBO

PF	Atmosphere/Meteorology
JH	Geoscience
LL	Hydrology
DH	Atlas of Living Australia
PM	OBO
CA	Reef Community
JG	Oceanography
RA	Hydrology
RL	Oceanography
SC	Geoscience
MH	OBO

6.1.1 Screening Survey-Broad Characterisaton Of Experts and Communities

Fifteen different types of communities were mentioned by responding experts (see Table 6.3). These communities were those with whom the expert was a part, or with whom they had worked in the past (on ontology development and reuse projects). A brief overview of these communities is provided next to anchor the reader in the contexts in which each expert has operated. The information presented demonstrates the broad representativeness, and breadth of experience collectively covered by the thirteen participants. Of note is the relatively high degree of overlap and the potential for cross-fertilisation of skills and practises between disciplines due to the number of experts involved in multiple communities. The survey specifically asked respondents to provide references to further information about communities with which they were associated. By consulting the literature and Web resources post Screening Survey receipt, any provided references regarding community affiliation were followed-up and this information has been included as part of the post survey data analysis. Most communities mentioned had international affiliations and had substantial memberships.

OBO Communities

Three communities that were cited ('Zebrafish', 'Phenoscape' and 'CARO') could be grouped under the one umbrella body, namely the OBO (Open Biological and Biomedical Ontology) Foundry (OBO Foundry, 2012a), which covers a range of developers of life-science ontologies.

The OBO Foundry communities started their ontology development in 2001 and are linked by the common principles that ontologies be open, orthogonal, instantiated in a well-specified syntax and designed to share a common space of identifiers. The main concept behind OBO is the issue of orthogonality, which if achieved would result in a single, non redundant system of ontologies. Each term would be defined in a single ontology and other ontologies, which need to use these terms, would ideally reference them (Smith *et al.*, 2007).

Table 6.3 Expert Community Affiliations - As Reported In Screening Survey Response

Expert Code	Community Identified	Disciplines Covered	Community Membership *	Community Governance Type +
RH	TDWG	Biodiversity Information Standards	IM	IG
WD	(a) Evolutionary Biology Community (OBO) (b) Model Organism Community (ZebraFish) (OBO)	Biology (Taxonomy) Biology (Taxonomy)	IM IM	IG NG
PF	(a) Earth Science Information Partners (ESIP) (b) National Aeronautics and Space Administration (NASA) (c) Aeronomy (d) Solar physics	Earth Science. Solid Earth, atmospheric, solar radiation. Middle and upper atmosphere. Solar atmosphere.	NM IM IM IM	IG IG IG IG
JH	Spatial Metadata Community & Geoscience	Management of spatial resources using metadata	NM	FG
LL	(a) Hydrology (Consortium of Universities for the Advancement of Hydrologic Science, Inc – CUAHSI) (b) Open GIS Consortium (OGC) (b) Energy	Hydrology Spatial data standards Distributed generation and biomass accounting.	IM	IG
DH	TDWG ALA	Bioinformatics	IM NM	FG FG
PM	(a) Ethosource/Ethosearch (b) Phenoscape TDWG	Animal Behaviour Fish Evolution Comparative Behaviour	IM NM	NG NG
CA	Coral Reef Ecology	Coral Reefs	NM	NG
JG	Marine Metadata Interoperability (MMI)	Marine observation parameters; platforms; devices and vocabularies for oceanographic science.	IM	IG
RA	IHO GeoSciML ANZLIC BoM, WMO	Marine XML, hydrography Geology Foundation models, Cadastre Meteorology, hydrology	IM IM NM IM	FG FG IG FG
RL	MMI Platforms. SeaDataNet SeaVox	Platforms and devices. Oceanographic parameters and metadata Oceanographic vocabularies	IM IM IM	IG FG FG

	CF Climate Science	Atmospheric science and oceanography	IM	FG
SC	Geoscience	Geology, geophysics, hydrogeology	IM	FG
	Water Resources	Hydrology and groundwater	IM	IG
	MetOcean	Climate science, oceanography	IM	IG
	Ecology	Taxonomy	IM	FG
MH	OBO	Anatomy, phenotype, ontology construction, interoperability and syntax	IM	FG
	Phenoscape (OBO)	Anatomy and evolutionary biology	IM	IG

* IM = International Membership; NM = National Membership

[+] FG = Formal Governance; IG = Informal Governance; NG = No Governance

Three experts (i.e., ‘MH’, ‘WD’, ‘PM’) are contributors to OBO ontology development activities. Some of their work was focussed on developing ontologies for different model organisms. Model organisms are species that have been widely studied usually because they are easy to maintain and to breed in a laboratory setting and they have particular experimental advantages. The **‘**ZebraFish’, for example, is a model organism**.** Model organisms are used to obtain (or infer) information about other species – including humans, who are more difficult to study directly (and ethically). Our collective understanding of human disease often relies on our ability to make reliable cross-species comparisons (using model organism data). A significant amount of model organism data is localized to anatomical structures. Some (OBO) ontologies represent structure, others represent function, others again represent stages of development, and some draw on combinations of these (Smith *et al.,* 2007).

To help improve cross-comparisons of data, two of the OBO experts had also participated in the development of a ‘framework’ anatomy ontology called the ‘Common Anatomy Reference Ontology’ (CARO). The function of this ontology is to better align cross species anatomical concept definitions and term relationships and to facilitate cross-species queries, for example when searching for similar phenotypes or gene expression. The OBO Foundry ontologies are all listed in the NCBIO BioPortal which covers ontologies in the fields of anatomy, phenotype, experimental conditions, imaging, chemistry, and health (Noy *et al.,* 2009).

TDWG

Another key biological ontology development community represented in survey is called the Taxonomic Database Working Group (TDWG), which covers biologists who are focussed on issues of ‘systematics’. Interestingly, despite the obvious (to the author) overlap in taxonomic systematics, it is a community quite distinct from the OBO Foundry communities (with no formally shared ontologies). Three experts stated that they had been involved with the TDWG community (i.e., 'DH', 'PM', 'RH').

TDWG is a not-for-profit scientific and educational association that is affiliated with the International Union of Biological Sciences. This group focuses on the development of standards for the exchange of biological and biodiversity data. In 2006 TDWG signed an MOU with the OGC in recognition that spatial information is an area of overlapping interest between the OGC and TDWG. The fact that TDWG has official affiliation with the OGC, but not with OBO either indicates both groups perceive little to be gained by cooperating, or they consider the fields they occupy to be very different. One expert ('PM'), however, stated that he has been involved in both TDWG and OBO.

Atlas Of Living Australia

Expert 'DH' is associated with TDWG and another biological community identified in this study, called the Atlas Of Living Australia (ALA (ALA, 2012)). The ALA is an AUS $50M funded Australian Government bioinformatics infrastructure project which is building a biodiversity information platform to provide scientists and the public with access to a wide range of biodiversity data. Linked data and ontological data description is of significant interest to this project.

Solar Terrestrial Physics

The Solar Terrestrial Physics Community was represented by one expert (i.e., 'PF'). This community is concerned with the interaction between the Earth and the Sun, with particular emphasis on the effect on the Earth of charged particles and energy from the Sun and of the interactions of solar and terrestrial magnetism. The ionosphere – the zone of charged particles around the Earth – is sustained by solar activity and is of particular importance in solar-terrestrial relations. Expert 'PF' is strongly associated with the Solar Terrestrial community's efforts to build a Virtual Solar Terrestrial Observatory (VSTO). This observatory is a unified semantic environment serving data from diverse data archives in the fields of solar, solar-terrestrial, and space physics.

Atmosphere/Meteorology

Experts 'SC', 'RA', 'RL' (and to some extent 'PF') have been, or are still currently active, in climate-science communities who are building grid computing facilities (BADC, 2012) to help model global climate. These communities are placing a heavy emphasis on methods required to better model the semantics of exchanged data that is ultimately used to run models and climate simulations. This category of communities could also be inclusive of sub-communities dealing specifically with Sensor Observation Services (SOS). Groups building SOS can apply their technologies and standards equally across all disciplines, since sensors are an interdisciplinary enabling technology. During interview (later in this study) 'JG', 'SC' and 'LL' specifically identified themselves with this (SOS) sub-community.

Oceanography (MMI, SeaDataNet)

In the discipline of oceanography two types of community were mentioned by experts ('JG', 'RL', 'PF' and 'SC'). All experts who had operated within the oceanography discipline had some affiliation with the Marine Metadata Interoperability (MMI (MMI, 2011)) project. This project was first funded by the National Science Foundation (NSF) in 2004 and is continuing to provide guidance, vocabularies and semantic services to the broad marine science and oceanographic communities. Its funding cycle is now at an end but the initiative is continuing with voluntary community support.

Experts 'RL' and 'JG' are also active participants in the SeaDataNet (SeaDataNet, 2011) federation. This is a virtual organisation of open digital repositories that manage, access and share data, information, products and knowledge originating from oceanographic fleets, new automatic observation systems and space sensors. It is a European consortium of National Oceanographic Data Centres and satellite data centres from 35 countries.

Hydrology (WIRADA, CUAHSI)

In the Hydrology sphere, three experts (i.e., 'SC', 'LL' and 'RA') explained that they were team members within the Water Information Research and Development Alliance (WIRADA (CSIRO, 2011). A partnership that has brought together the Australian Commonwealth Scientific Research Organisations (CSIRO)'s research and development expertise in water and information sciences and the Australian Bureau of Meteorology's operational role in hydrological analysis and prediction, in an effort to transform the way Australia manages its water resources. The main focus of activity for the experts in this study have been development of a water data transfer standard and assisting to develop methods and tools for managing and ensuring interoperability between different water information models.

There is a close relationship between the WIRADA project (and its constituent community) and a broader international community comprising of approximately 130 institutions world-wide (i.e., the Consortium of Universities for the Advancement of Hydrologic Science – CUAHSI (CUAHSI, 2011)). The US component of CUAHSI is partially supported by the NSF and its goal is to develop infrastructure and services for the advancement of hydrologic science. WIRADA and CUAHSI are key communities working together through an OGC Hydrology Domain Working Group that is seeking to evolve a hydrology mark-up language (i.e., WaterML) into an International Standard.

Geoscience (GeoSciML)

Four experts (i.e., 'JH', 'SC', 'RA' and 'PF') have been involved with communities in the geology discipline, with three ('SC', 'RA' and 'JH') specifically mentioning the community actively developing a semantic mark-up language for geology (i.e., GeoSciML (CGI, 2011)). GeoSciML is a project sponsored by the Commission for the Management and Application of Geoscience Information, which is a Commission of the International Union of Geological Sciences.

ANZLIC, Coral Reef, Animal Behaviour

Of the remainder of the communities, one is an Australia and New Zealand standards partnership focussed around spatial data standards, particularly spatial metadata (i.e., ANZLIC (ANZLIC, 2012), cited by expert 'JH') which is a relatively large community (in excess of 100 people). The remaining two cited communities, in comparison to all of the other communities that were mentioned, are quite small and possibly immature (i.e., a Coral Reef community identified by expert 'CA' and a community developing ontologies related to animal behaviour, cited by expert 'PM').

6.1.2 Expert Screening Survey– Potential Stratification Of Experts

Apart from providing data to assess a respondent's suitability for study participation, a goal stated earlier for the Screening Survey was for the survey to provide information that could be used to start looking for different ways in which experts might be stratified (or grouped) to assist with post interview data analyses and interpretation. Later in this study experts are asked to provide 'ratings' for the relative importance that they place on various ontology evaluation criteria, that have been nominated by experts (during interview) as being useful in selecting ontologies for reuse. The variability in how experts rate these evaluation criteria is specifically examined with respect to the stratification established here (and from additional stratification dimensions that are established post the in-depth interview process). By grouping experts it is possible to directly address (RQ1.2.2.2) – which asks if any relative differences in an expert's rating is reflective of '*disciplinary affiliation'* or '*other discernible factors'*.

The data captured from the survey offered several possible ways in which experts might be grouped. Apart from the obvious division of participants according to whether they have direct ontology development experience in ISO/OGC Standards environments (as shown in Table 6.1), potential groupings (based on the screening survey data) included clustering of experts by:

(i) The disciplines covered by expert community-centric activity.
(ii) The governance type of the communities that experts are involved in.
(iii) Expert perceptions of specific (and important) terms used in this study.

(iv) Expert ontology development experience.

(v) An expert's skill type.

The validity of these various group clusters were further evaluated during in-depth interviews and their significance, in light of information collected, was analysed post interview. Table 6.3 (presented earlier) has summarised returned survey data for community-of-interest affiliations, including the associated community governance classification, from which the following stratifications have been made.

(i) Stratification By Discipline

From the survey data it was found that there was a reasonable amount of cross-over in terms of community membership (e.g., 'JG' and 'RL' both nominated participation in the MMI community; 'MH', 'WD' and 'PM' all stated that they were active in the animal anatomy and evolutionary biology communities and 'SC' claimed to be active in hydrology, oceanography, atmospheric and geoscientific communities). Six main disciplines appeared to be covered by the experts: hydrology; biology (mainly animal); oceanography; geosciences, atmosphere (and solar physics)/meteorology and energy. Atmosphere, Solar Physics and Meteorology were grouped because many of the phenomena studied are closely interconnected. The energy discipline had only one representative expert and it appeared that the focus of this expert, within the energy discipline, was largely on bio accounting (a biological disciplinary focus). Water standards, marine vocabularies, and biodiversity (particularly taxonomy) were sub-disciplinary topics most often mentioned. Table 6.4 presents a possible grouping of experts based on the disciplines they cover (in their community ontology development work).

Table 6.4 Possible Expert Groupings of Respondents Based on Disciplines

Grouping	Expert Code
Hydrology	LL, RA, SC
Biology	RH, WD, DH, PM, CA, SC, MH, LL
Oceanography	JG, RL, SC
Geoscience	JH, RA, SC, PF
Atmosphere/Meteorology	PF, RA, RL, SC
Energy	LL

(ii) Stratification By Community Governance Type

Drawing from the survey responses, most communities that were listed were cited as having an international membership (ANZLIC, Phenoscape and ESIP communities were exceptions). Since ANZLIC is officially an Australia and New Zealand partnership, it was surprising to see that its membership was cited as "national". It has been reasoned, that in answering the (survey) questions about community membership, some experts may have confined themselves to responding about specific sub-groups, or specific activities within these larger groups, with which they had direct experience. This might account for some apparently contradictory survey responses (such as the one just mentioned). There was a relatively even split between communities concerning governance arrangements, with thirteen of the 24 communities mentioned having 'informal' governance and eleven cited as having 'formal' governance. Two experts mentioned that two of the communities that they were associated with had no governance. However, there appears to be a contradiction between the statements made by 'MH' and 'PM' with regards to the 'Phenoscape' community, with one respondent believing the community had "no governance arrangements" and the other stating it had "informal governance". This possibly indicates governance is very informal (if it exists at all). See Table 6.5 for potential stratification using community governance as a possible discriminating factor.

Table 6.5 Possible Groupings of Respondents Based on Community Governance Type

Grouping	Expert Code	Disciplines Represented
No Governance Group	WD, PM, CA	Biology
Informal Governance Group	RH, WD, PF, LL, JG, RA, RL, SC, MH	Biology, Atmosphere, Oceanography, Meteorology
Formal Governance Group	JH, DH, RA, RL, SC, MH	Biology, Geoscience, Meterology, Oceanography
People spanning both IG/FG	RA, RL, SC, MH	Biology, Atmosphere, Oceanography, Meteorology, Geoscience

(iii) Stratification By Expert Responses To Term Definitions

Another possible level of stratification investigated was whether expert interpretation of key terms used in this study, and which were provided for comment at the start of the questionnaire, differed amongst respondents. It seemed reasonable to assume that there could, for example, be differences of opinion about preferred ontology evaluation criteria if an expert took a broader interpretation of what the definition of an ontology was, than if they took a narrower, more formal view. In response to the definition questions posed in the survey, the first 4 definitions (for "ontology", "light-weight ontology", "formal ontology" and "ontology re-use") drew the most comment (see summary Table 6.6).

The comments supplied about the definition for "ontology" didn't really signal any major differences of view between the experts (see Table 6.7 for the actual comments). Differences of opinion emerged, however, when the terms, "light-weight" and "formal ontologies" were commented upon. One expert (at least) considered a "light-weight" ontology to be no more than a vocabulary and stated that these types of ontologies shouldn't be called ontologies at all (e.g., 'MH's' views in Table 6.7). It also transpired (during subsequent in-depth interview) that 'LL' shared 'MH's' view (see also 'LL's' comments for "ontology" in Table 6.7 which foreshadowed the view stated during interview). Other experts were prepared to accept a definition that gave a broader, more encompassing interpretation of ontologies, spanning 'informal' to 'formal'.

Table 6.6 Expert Agreement/Disagreement With Provided Definitions In Questionnaire

Expert Code	**ontology**	**Light-weight ontology**	**formal ontology**	**ontology re-use**	**encoding language**	**ontology selection**	**ontology evaluation**	**community-of-interest**	**Community governance**	**web services**
RH	0	1	0	0	1	0	1	1	0	1
WD	1	0	1	1	1	1	1	1	1	1
PF	0	0	0	1	1	1	1	1	1	1
JH	1	1	1	1	1	1	1	1	1	1
LL	0*	1	1	1	1	1	1	1	1	1
DH	1	1	1	1	1	1	1	1	1	1
PM	1	1	0	1	1	1	1	1	1	1
CA	1	1	1	1	1	1	1	1	1	1
JG	1	0	1	0	1	0	1	1	1	1
RA	0	1	0	1	1	1	1	1	1	1
RL	1	1	1	0	1	1	1	1	1	1
SC	1	1	1	1	1	1	1	1	1	0
MH	1	0	1	0	1	1	0	0	1	1
TOTAL (agree)	**9**	**9**	**9**	**9**	**13**	**11**	**12**	**11**	**12**	**12**

1 = agree; 0 = disagree

* LL's caveat: agrees with subsequent definitions for convenience but indicates disagreement with some of the definitions which follow, though doesn't say which ones. Implies ontology definition is too loose (not formal or unambiguous enough).

It is interesting to note here some inconsistencies in the responses provided by experts between the answers they gave for questions regarding the 'definitions of terms' and answers that they gave about their 'ontology development experience'. 'JH', 'RA', and 'SC' all agreed to the definition provided in the survey for a "Formal Ontology", which stated that a formal ontology was one capable of supporting reasoning. In response to the survey question asking experts about the ontologies that they had developed; 'JH', 'RA', and 'SC' all listed ontologies that they had developed as being "formal ontologies" despite these ontologies (apparently) being expressed in a language (XML and/or UML) that does not support reasoning. It would appear, that for these experts, their working definition of an ontology is perhaps broader (and looser) than the definition they actually signed off on (in the survey). In addition, both ‘PM’ and ‘RH’ questioned the definition of a formal ontology, ‘RH’ being prepared to include SKOS as a formal ontology despite its lack of reasoning power and ‘PM’ simply broadened the definition of ‘formal’ to include more forms of logic than First Order Logic (FOL), as did ‘PF’. However, both ‘PM’ and ‘RH’ then listed OWL ontologies that they had created as ‘light-weight’. The author can only conclude that perhaps this was because they had not exercised appropriate property and value restrictions when using this language or that they had used OWL Full (which does not lend itself to reasoning), or they just answered inconsistently.

Although there were embellishments of the provided definition for the “ontology re-use” term, in the main most people understood the term to mean a similar idea. The only point of significance that should be raised regarding these alternate definitions for “ontology re-use” is the idea, captured from expert commentary, that “re-use” can mean both the explicit referencing of existing ontological components, or borrowing components from an existing ontology and (possibly) re-working (or re-casting) them to fit within the framework of a new ontology.

Overall, there appeared to be a relatively good level of concordance regarding definitions for some of the key terms that are important in the context of this study (at least at a superficial level as measured by the responses to the survey questions, despite the apparently contradictory statements surrounding “formal”ontologies). The only stratification that could be realised, based on differences of views about the meanings of terms, would see ‘LL’ and ‘MH’ grouped in a cluster that believes ontologies by definition are “formal”. Apart from this grouping, there did not seem to be any other obvious stratification of experts. The expert commentary on the provided definitions was, however, considered to be useful for interpreting other information gathered later in the study.

Table 6.7 Alternate Definitions Provided By Experts For Survey Supplied Terms

Expert Code	**Term: Ontology**
	Definition Provided: An ontology makes explicit the types of concepts, or real-world objects used in a domain and generally includes information about concept properties, concept relationships and value restrictions on both properties and concepts.
RH	“The Gene Ontology wiki defines ontology as 'blah blah blah' - I find it hard to improve on this. http://wiki.geneontology.org/index.php/Glossary.”
PF	“I agree, except in the sense that 'type of concepts' is used.... it is usually, the types AND the concepts.”
LL	“In the context of this study, it is more convenient to use your definitions even if I don't totally agree with some of them (I belong to the 'ontologies are formal and unambiguous school) - So here I define the term ontology more loosely + any 'raw material' which can be converted into your definition of an ontology (this is the only way I can answer efficiently to the questions on how do you select what you reuse).”

RA	"A procedural rather than functional def :-): An ontology is a formalism of an epistomology, (which is the agreed concepts using in a domain of discourse, defined by that community)."
Expert Code	**Term: Light-weight Ontology**
	Definition Provided: A light-weight ontology is one which may lack formalism (inbuilt rules and syntax) required to support automatic reasoning.
WD	"NA - I haven't encountered this term before."
PF	"I would refer to the ontology spectrum (http://www.ksl.stanford.edu/people/dlm/papers/ontologies-come-of-age-mit-press-(with-citation).htm) and lightweight is often OWL-Lite. Prefer not to include 'lack' of formalism, prefer 'informal is a'."
JG	"A light-weight ontology is one which lacks significant detailed information (e.g., , , properties and value restrictions) while still adhering to the nominal form of an ontology."
MH	"I wouldn't call this an ontology at all, rather a controlled vocabulary."
Expert Code	**Term: Formal Ontology**
	Definition Provided: A formal ontology is one that is capable of being used for automatic reasoning through its expression in a language capable of supporting first-order logical inferences.
RH	"Many people would think of a SKOS ontology/vocabulary as being formal if it is agree and represents the domain but it needn't be good for inference."
PF	"FOL is too strict. DL is more than sufficient. Also formal does not always mean reasoning is used...."
PM	"A formal ontology is one that is capable of being used for automatic reasoning though its expression in a language capable of supporting logical inferences in description logic or more inclusive logical formalism."
RA	"would add governance constraints - someone has to do the fomalising, otherwise it's all form and no substance."
Expert Code	**Term: Ontology Re-use**
	Definition Provided: Ontology re-use is the process in which existing ontological components are used as input to generate new ontologies.
RH	"Ontology re-use is the process in which existing ontological components are used or referenced by other ontologies to avoid duplication of concepts."
JG	"Ontology re-use is the application of existing ontologies to create new applications or understanding (which do not necessarily ever get expressed as ontologies)."
RL	"Ontology re-use is the process in which existing ontological components are used as input to generate new ontologies or the process by which existing ontologies are adopted in full by communities not involved in their development."
MH	"I would consider re-use to be support for interoperability. I would not create new ontologies, but rather Xref them explicitly."
Expert Code	**Term: Ontology Selection**
	Definition Provided: Ontology selection is the process that is used to determine if one or more ontologies, or ontological components satisfy certain predefined criteria.
RH	"Ontology selection is a little like shopping. One doesn't go out just with a specific set of criterior. There is an element of discovery involved. One could go looking for shoes and come back with a complete outfit!"
JG	"Ontology selection is the process of determining the ontologies that will be most useful for a given application or purpose."
Expert Code	**Term: Ontology Evaluation**
	Definition Provided: Ontology evaluation is a task within the ontology selection process where an assessment against any given selection criteria takes place.
MH	"ontology evalulation can be both an initial task as part of ontology selection, but is also an ongoing task as part of ontology maintenance and collaboration."
Expert	**Term: Community of Interest**

Code	
	Definition Provided: A community-of-interest is a group of individuals that has a shared and identifiable set of interests. Communities are bound together informally through shared practices and contributions to common goals or more formally through the formation of an organisation.
MH	"A community-of-interest is a group of individuals that has a shared and identifiable set of interests and requirements. Communities are bound together informally through shared practices and contributions to common goals, use of common tools or ontologies, or more formally through the formation of an organisation."
Expert Code	**Term: Community Governance**
	Definition Provided: Community governance involves establishing structures, processes, policy and the norms that help the community function to achieve its goals.
RH	"Yes but... To be a single community their needs to be a single goal if not you have goal chaining. X may share goal A with Y and Y may share goal B with Z but X and Z don't share a goal. Are they in the same community-of-interest? If yes their is only one global community-of-interest if no then there are very few real community-of-interests out their."
Expert Code	**Term: Web Services**
	Definition Provided: Web services are software systems designed to support interoperable machine to machine interaction over a network. Three commonly supported architectures for such systems are Service-Oriented-Architectures (SOA), Representational State Transfer (REST) and Remote Procedure Calls (RPC).
SC	"Web services are software systems designed to support loosely-coupled machine to machine interaction over a network. Service-Oriented-Architectures (SOA) may be implemented using Representational State Transfer (REST) or Remote Procedure Calls (RPC)."

(iv) Stratification By An Expert's Ontology Development Experience

With respect to ontology development experience (summarised from survey responses and presented in Table 6.8), four of the thirteen experts had experience developing light-weight ontologies (six if 'SC' and 'JH' are both included) and the remainder (7) have experience with formal ontology development, or both. Two experts claimed to have used OWL-DL and another five claimed experience with OWL (without specifying a dialect). It is possible (in fact probable) that all five experts who listed OWL, had used OWL-DL.

Table 6.8 Expert Ontology Development Experience

Expert Code	Ontology Name	Community	Domain	Language	LW/Formal
RH	LsidVocs	TDWG	Biology	OWL (but says it should have been RDF/S)	Light-weight
WD	Teleost Anatomy Ontology	Icthyological Community	Biology (comparative phylogenetic data)	OBO	Formal
PF	VSTO	Aeronomy and Solar Physics	Aeronomy and Solar Physics – instruments and physical quantities.	OWL-DL	Formal
	SESDI	As above	Earth science, atmospheric physics,	OWL-DL	Formal

			solar radiation		
JH	Geoscience Australia ISO 19115 scope code	Metadata Community	Cross-domain	XML (GML) using gml.xsd	Formal
LL	Water Observation Sustainable Biomass and Bioenergy Ontology		Hydrology, sensor Energy and Biomass	OWL OWL	Formal Formal
DH	TDWG Ontology	TDWG Community	Biodiversity informatics	OWL	Light-weight
PM	ABO Core Teleost Taxonomy Ontology	 Phenoscape	Animal Behaviour Taxonomy of teleost and other fish	Protégé Frames (also has an OWL version that is not in use) OBO	Light-weight Light-weight
CA	Coral Reef Ontology	Coral Reef Ecologists UQ	Coral Reefs	OWL	Light-weight
JG	MMI Device Ontology MMI Platform Ontology		Marine science sensing and sampling devices Marine science observation platforms	OWL-DL OWL-DL	Formal Formal
RA	ANZLIC Harmonised Data Model		Cadastre, Native Title	ISO (19000) compliant UML models	Formal
RL	MMI Platforms Ontology NDG Vocabulary Server		Marine metadata Atmosphere and ocean sciences	OWL Relational database with export to RDF	Formal Formal
SC	GeoSciML WaterML	IUGS Committee for Geoscience Information	Geology, Geological maps Water resources observations	UML, XML Schema UML, XML Schema	Formal Formal
MH	ZebraFish Anatomy Ontology Common Anatomy Reference Ontology (CARO)	ZebraFish Model Organism database (ZFIN) OBO Template for anatomy ontologies	Zebrafish anatomy and developmental stages Anatomy ontology for an organism	OBO OBO	Formal Formal

(v) Stratification By An Expert's Skill Type

Seven of the thirteen respondents claimed to be a "domain expert with some ontological developments skills". Three rated themselves as "skilled ontological engineers and domain experts". Two claimed to be "skilled ontological engineers with some domain expertise". One stated that they were a "skilled ontological engineer with no domain expertise". The sample cohort is populated with a roughly equal split of people who are **domain experts** with **some form of ontological skills** and people who are **skilled ontological engineers** with **varying degrees of domain expertise (**see Table 6.9 for summary data). It would therefore be highly useful to partition analyses and interpretations on the basis of whether experts claimed to be ontological engineers or domain experts (or both).

Table 6.9 Skill Level Nominated By Experts

Expert Code	Skills	Disciplines Covered
RH	Domain expert with some ontological developments skills (DESO)	Biology
WD	Domain expert with some ontological developments skills (DESO)	Biology
PF	Skilled ontological engineer and domain expert (SOEDE)	Atmosphere/Meteorology; Geoscience
JH	Domain expert with some ontological developments skills (DESO)	Geoscience
LL	Skilled ontological engineer with some domain expertise (SOESDE)	Hydrology, Biology, Energy
DH	Domain expert with some ontological developments skills (DESO)	Biology
PM	Skilled ontological engineer and domain expert (SOEDE)	Biology
CA	Skilled ontological engineer with no domain expertise (SOENDE)	Biology
JG	Domain expert with some ontological developments skills (DESO)	Oceanography
RA	Skilled ontological engineer with some domain expertise (SOESDE)	Hydrology, Atmosphere/Meteorology, Geoscience
RL	Domain expert with some ontological developments skills (DESO)	Oceanography, Atmosphere/Meteorology
SC	Domain expert with some ontological developments skills (DESO)	Biology, Hydrology, Atmosphere/Meteorology, Geoscience, Oceanography
MH	Skilled ontological engineer and domain expert (SOEDE)	Biology

6.1.3 Re-use, Selection and Evaluation Methods (As Derived From Screening Survey)

A further goal of the screening survey was for the author to begin the process of gaining a better understanding of the issues that experts considered were of the most importance, in their contexts, for selecting and evaluating ontologies for re-use (prior to an in-depth examination via the interview process). All experts (bar 'JG', as explained earlier) claimed to have reuse experience which is why they were considered useful as study participants (see Table 6.10). Importantly in this study, there is also an interest in whether experts, who are engaged in reuse, are actively using selection and evaluation methods sourced from the academic literature (recall RQ 1.2.1 – *What ontology selection and evaluation criteria are currently used across multi-disciplinary scientific communities (and are selection and evaluation methods consistent with those reported in the literature ?*). Survey respondents were therefore asked if the methods they had used were based on methods reported in the literature; were based on methods used by colleagues in another community or domain; were based on methods developed in their own communities, or whether no evaluation method was used at all.

Table 6.11 shows that only two experts cited methods with academic origins (i.e., 'LL' and 'MH'), and two others (i.e., 'PF' and 'PM') attributed their methods to another community. All experts, however, stated that they used some form of evaluation method.

Table 6.10 An Expert's Experience With Ontology Reuse

Expert Code	Ontology Re-use	Community In Which Development Took Place
RH	Tried to re-use Dublin Core and VCARD	TDWG
WD	Re-used Zebrafish Anatomy Ontology	Icthyological Community
PF	Re-used SWEET Ontology	Solar and Aeronomy
JH	ISO TC 211 gmx_Codelists located at http://www.isotc211.org/2005/gml/gml.xsd	Metadata Community
LL	Yes (has experience but specific ontology not nominated)	Water Observation Community
DH	Re-used domain knowledge held in earlier XML schema based data models.	TDWG
PM	Re-used NCBI taxonomy ontology (also part of OBO) in the Teleost Taxonomy	OBO Community
CA	NASA SWEET, SPIRE and OBO ontologies	Coral Reef Ecology Community
JG	Didn't understand the question	
RA	ISO Geography Models	ANZLIC
RL	Yes (has experience but specific ontology not nominated)	IODE GETADE
SC	Yes (has experience but specific ontology not nominated)	Geoscience community
MH	For CARO project re-used components of the Foundational Model of Anatomy (FMA) and CARO itself is now re-used by groups developing their own ontologies.	OBO

Table 6.11 Source Of Evaluation Methods

Expert Code	Ontology Name	Method Reference	Method Source
RH	LsidVocs		Own community
WD	Teleost Anatomy		Own community
PF	SWEET 1.0	"http://sweet.jpl.nasa.gov"	Other community
PF	SWEET 2.0	"http://sweet.jpl.nasa/2.0"	Other community
JH	Geoscience Scope Codes	"Gained community feedback by ozmeta-l@erin.gov.au and asdi-l@lists.anzlic.org.au email lists plus emails directly to others interested in this field."	Own community
LL	Water Observations	"Unpublished paper on methodology – loosely connected to ontology quality work published by Gangemi et al and the ontology work published by Alanai et al."	Academic origin
DH	TDWG Ontology		Own community
PM	NCBI Ontology		Other community
CA	Coral Reef Ontology		Own community
JG	Device Ontology		Own community
JG	Platforms Ontology		Own community

RA	ANZLIC Harmonsied Data Model		Own community
RA	Lanslides Inventory Model	"Methodologies proposed in http://ijsdir.jrc.ec.europa.eu/index.php/ijsdir/article/viewFile/62/26"	Own community
RL	NDG Vocabulary Server		Own community
SC	GeoSciML		Own community
SC	WaterML		Own community
MH	CARO	"http://www.blackwell-synergy.com/doi/pdf/10.1111/j.0021-8782.2005.00376.x?cookieSet=1 , http://www.nature.com/nbt/journal/v25/n11/full/nbt1346.html, http://www.acsu.buffalo.edu/~md63/DonnellyAIMed05.pdf"	Academic origin
MH	ZebraFish Anatomy		Own community

Experts were also asked to nominate what issues guided their "selection" and "evaluation" processes so that it was possible for the author to start to become familiar with some of these issues prior to performing in-depth interviews. Table 6.12 summarises the selection and evaluation topics extracted from expert commentary. Responses for evaluation issues were more detailed and perhaps more specific than were provided for selection topics. Both sets of topics served as prompts for further elaboration during interviews. Appendix 19 provides expert responses in full, for selection and evaluation issues respectively.

Table 6.12 Summary of Reported Selection and Evaluation Topics - Survey Responses

Selection Topics	Evaluation Topics
Domain coverage	Easier to make it up yourself
Fit with "normalised skeleton"	Appropriate scope
Purity of entities	Modularity and ease of owl imports
Ontology focus (not vague)	Ease of transforming raw material into ontology
Availability of definitions	Compatibility with normalised skeleton
Relevance to problem	Size and coverage of domain
Compatible technology	Alignment with earlier work
Coherence	Currency within the community
Completeness	Following the structure of an ontology developed by an influential member of the community
Credibility	Extracted best parts from existing ontologies in terms of how well they fitted the domain
Community acceptance, formal specifications interoperability	Number of terms, organisation of terms and clarity of definitions
Avoiding overlaps	Well-scoped and not monolithic package and matches governance mandate of community
Potential impact locally or globally	Governance process transparent and suitable for local reuse constraints
	Suitable extensibility mechanisms

	Does it mix meta levels
	Is it documented
	Is it formalised in a way that compliance with the formal standard can be assured
	Subjective analysis
	Familiarity with existing ontology
	Available expertise
	Single inheritance

6.2 In-depth Interview Data Analyses

The screening survey provided useful background information on experts willing to participate in the study. All thirteen chosen experts were then interviewed by the author. A standard list of prompting questions (refer again to Appendix 4) was used to generate discussion during interviews (i.e., interviews were semi-structured in that specific information was being sought). Often respondents would need very little prompting to start talking about the topics and in these circumstances the question list was simply used as an aide memoir by the interviewer to make sure that all topics requiring investigation had been covered. All participants were also provided with a table of evaluation criteria sourced from Finney (2008). During the last stage of each interview, each expert was asked to make comment on the table's contents. Of interest was whether any of the listed criteria, if not already covered during the interview, had any resonance with an expert (positively or negatively). Some experts preferred to provide a written response on this question, post interview (because the allotted time for the interview had been exceeded at that point). The information returned was added to the interview transcripts for analysis.

Desired outcomes from these interviews were:

- The identification of key ontology selection and evaluation criteria and their categorical relationship (i.e., potential dimensionalisation), plus a further assessment of whether criteria and/or methods applied by experts were consistent with those described in the literature (RQ1.2.1).
- Any resulting dimensionalised model of the evaluation criteria would be used (in conjunction with the AHP method) to gain a deeper understanding of the importance placed on criteria by experts, when using these criteria in ontology selection processes (RQ1.2.2.1).
- An enhanced understanding of how communities of practise manage and govern ontology development processes (involving re-use) and the measures they use to assess ontology evaluation criteria (RQ1.2 & RQ1.2.3).

Table 6.13 summarises interview details. In five out of the thirteen interviews the author had the advantage of being able to react to the visual mannerisms of the experts as well as to auditory cues. Overall, however, the interviews conducted by phone provided for a significant amount of interactivity and only one expert ('PM') was possibly hampered in getting his messages across because of a poor phone connection.

Table 6.13 Interview Details

Expert Code	Interview Method	Interview Duration (hours:minutes)
RH	Phone	0:56
WD	Phone	0:40
PF	Phone	1:09
JH	Face-to-Face	1:06
LL	Face-to-Face	1:18
DH	Phone	1:02
PM	Phone	1:03
CA	Phone	0:37
JG	Phone	0:57
RA	Face-to-Face	1:17
RL	Face-to-Face	0:54
SC	Video Link	0:57
MH	Phone	0:54

Approximately 13 hours of expert commentary was captured and transcribed from audio into digital text. Each transcript was then loaded into Atlas.ti™ (i.e., qualitative analysis software) along with any other material supplied post interview by the expert. Tools in Atlas.ti™ provided a computerised means for establishing and managing codes and memos that were used to mark-up text and it was possible for each code to be assigned to a code family for classification purposes (see Figure 6.2 and 6.3 for screenshots of the Atlas.ti™ tool).

Each transcript was reviewed, in detail, sentence-by-sentence several (> 15) times each. Ultimately each pass through the transcript sought to extract information from a number of perspectives. Some of these included:

a. Gaining insight, from a holistic perspective, into the specific techniques and criteria used by experts to select and evaluate ontologies for re-use.
b. Differentiating between the stated importance placed on any cited selection and evaluation criteria.
c. Identifying where experts converged or diverged in the methods they used.

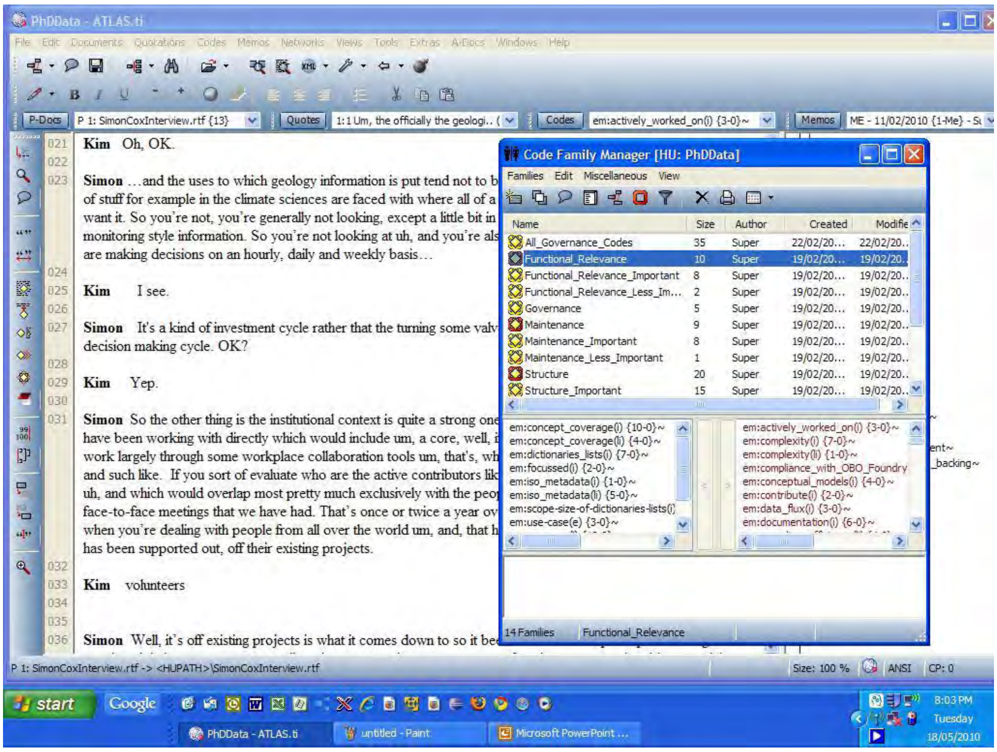

Figure 6.2 – Snapshot of Atlas.ti software showing code manager.

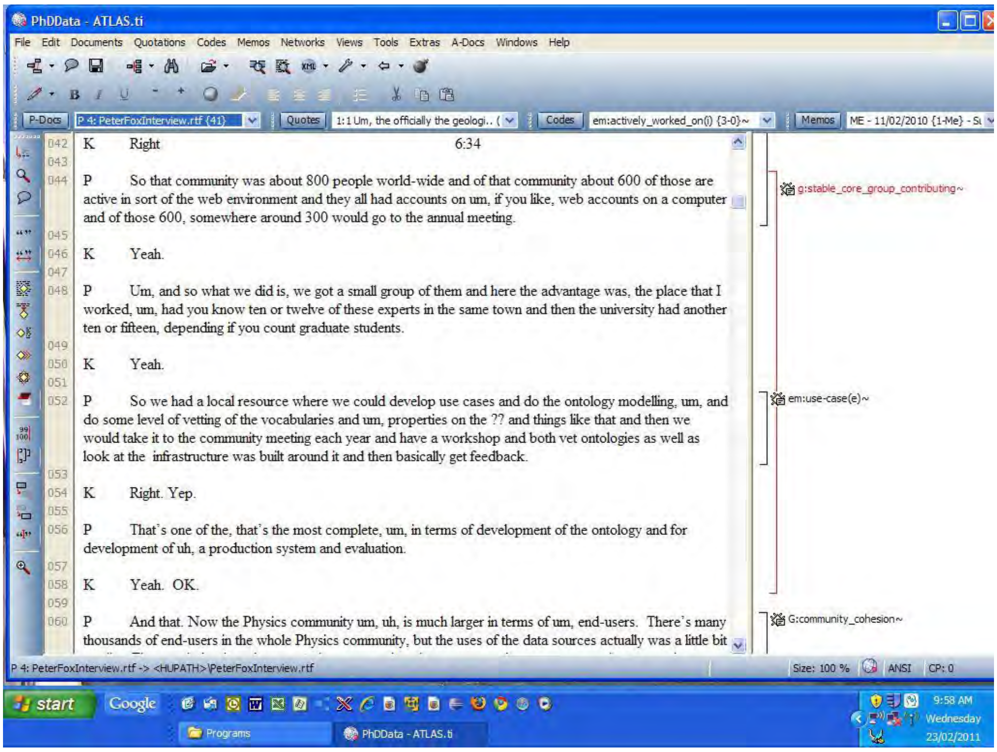

Figure 6.3 – Snapshot of Atlas.ti software showing part of a coded transcript.

d. Investigating any patterns or groupings in expert responses with respect to methods used, disciplines covered, communities experts were involved in, the governance arrangements evident in their communities of interest, and the types of use-cases the ontology development activities were aimed at satisfying.
e. Understanding the contexts in which experts were operating.
f. Differentiating between methods, criteria and techniques experts thought were successful vs those that they thought less effective.
g. Validating the expert's stated level of expertise (from the Screening Survey).

The interview material was also allowed to 'speak for itself' in that concepts, patterns, ideas and associations could emerge as part of the analysis, regardless of whether they informed the analytical perspectives mentioned above.

The descriptive information captured from experts was used directly to generate a dimensionalised hierarchical model of ontology evaluation criteria (as described in more detail later in 6.2.2), other information arising was used to further stratify experts, and some theoretical propositions were generated about the 'methodologies' used by communities to select and evaluate ontologies and how they govern ontology development and use (which is reported in Chapter 7).

According to Dey (1993) qualitative analysis involves 'breaking data down into bits and then beating the bits back together'. He offers that analysis is not just about describing the data, but is about describing the objects or events to which the data may refer. He goes on to elaborate that description is often not sufficient, what we often want to do is to 'interpret', 'understand' and 'explain'. In breaking the data down and then classifying it, the concepts we use in classification and the connections we make between them, provide the basis for a fresh description. Describing phenomena, classifying it and seeing how concepts interconnect is at the heart of qualitative analysis and these three tasks are what have been performed in this thesis on the data transcribed from the in-depth interviews (see Figure 6.4 for the circular and iterative process of qualitative analysis, which was performed on the data).

Deep descriptions (Denzin, 1978) i.e., descriptions which take account of context, actor intensions and processes in which actions or events are embedded, were formulated for various issues under investigation to support 'accounts' of phenomena of interest, such as the methodological approaches used by experts for ontology evaluation and issues of community-based ontology governance (described later in Chapter 7). These descriptions were built from data using a coding process involving the identification of concepts (Glaser and Strauss, 1967) that were evident within the transcribed text, which were then grouped into categories and themes (Richards, 2005). Themes integrate categories in a manner that sheds light on the research issues under investigation.

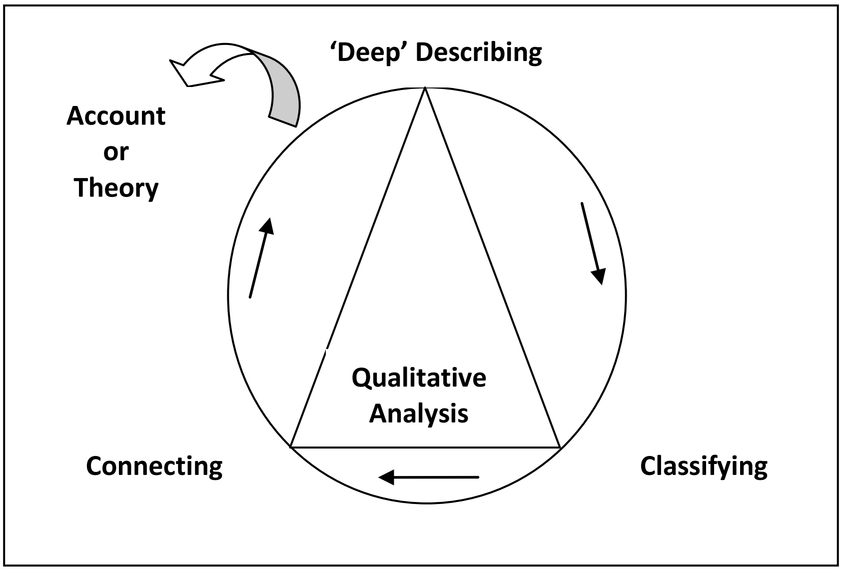

Figure 6.4 A Circular and Iterative Qualitative Analysis Method (adapted from Dey (1993) Fig. 3.1)

In terms of specific methodology, thematic analysis as described by Dey (1993) and Miles and Huberman (1994) was the approach used for coding and data analysis, harnessing the data-driven inductive approaches of Boyatzis (1998) but with the 'Template Analysis' method of Crabtree and Miller (1992), which involved the deductive a priori development of codes. Hierarchical coding is emphasised in Template Analysis; that is to say, broad themes encompass successively narrower, more specific ones. This method was chosen because the author had already undertaken some work on identifying selection and evaluation criteria and to ensure that this work did not bias outcomes, it was 'put on the table' up-front and the data was used, in part to put existing assumptions to the test. This was considered a 'more honest' approach than either (a) declaring concepts, categories and themes as 'emergent', when in reality some of these results would not be unexpected – particularly if they are already evident in the literature, or (b) pretending that there were no a priori ideas, but sub-consciously using the data to validate pre-conceptions through selective application of the data. Instead, a partially populated template was established up-front (particularly for selection and evaluation criteria and evaluation measures topics) and data either supported aspects of the template or caused its re-arrangement.

Thematic analysis was never-the-less highly inductive; themes were allowed to emerge from the data and were not imposed upon the transcripts by the author (even though, as explained some a priori codes/categories were conceived prior to the commencement of data analysis). Data collection and analysis took place simultaneously, particularly as background reading and references mentioned at interview formed part of the analysis process. Data from different experts were compared and contrasted as the author successively interviewed and transcribed material. Often as new transcripts were analysed this caused the author to re-visit already partially analysed data, looking to see if emergent concepts in new transcripts hadn't been missed in any previously analysed interview material.

The Template Analysis method that was used, as described by King (2004), harnessed codes and methods of coding that were not as prescriptive as those proposed in other coding methods such as Grounded Theory (Strauss and Corbin, 1998). Additionally, Template Analysis is not as wedded to Grounded Theory's realist methodology (Waring and Wainwright, 2008). King (2004) explains that it is permissible to develop some initial codes (perhaps aligned to the interview questions or the researcher's knowledge) a priori, which can be refined after exploration of the data (i.e., as further concepts emerge from data analysis). This is discouraged in Grounded Theory (Glaser, 1995).

6.2.1 Codes and Template Development

Several passes through all of the transcripts was necessary to really refine how to code and fully mark-up the transcripts. The first and second recorded interview transcripts were used initially, however, as training material not only to better understand the capabilities of the supporting software, but just to get used to the act of coding. Although coding proceeded as each interview was sequentially transcribed, it was important to arrive at an efficient technique to analyse and mark-up all of the data, since the sheer volume of material was daunting and the task was extremely labour and time intensive. Too many codes with only slight semantic or contextual differences appeared to make the process overly complicated and created obfuscation when trying to make sense of the data in-hand. This was because the concepts that were coded initially were so individualised, and there were so many of them that they defied meaningful processing using human cognition. Valuable lessons learnt whilst using the training data were that too many codes made meaningful interpretation impossible and too few codes meant potentially important concepts were missed.

Guidance was sought from Dey (1993) who suggested asking "who", "what", "when" "where" and "why" of the data and most significantly then asking "so what ?". This approach was explored but abandoned relatively quickly because it just didn't seem to fit comfortably with exploring the topics of interest. Keeping in mind the perspectives (stated earlier) which were pre-formulated to guide the analysis, it seemed more natural to base the coding on several concerns of the research as suggested by Bogdan and Biklen (1982) in Dey (1993), which included searching for: definitions of various constructs, criteria, processes, methods and communities; processes; settings ; strategies; relationships and attitudes. Reviewing the data from these perspectives provided for a rich set of concepts and connections.

As ideas, concepts and issues arose whilst reading through each transcript, a code was created and each code was given a description. If the idea, issue or concept recurred in the same transcript the code was applied again to the relevant segment (or quotation). If the same idea, issue or concept was encountered whilst reading through a different transcript an existing code was re-used and applied as text mark-up, providing it adequately covered the concept being flagged. For some purposes of evaluating the data it was reasoned that the more times a code occurred, the stronger the evidence was for it as being an issue of significance (within the cohort studied). This was obviously the case if the occurrences were spread between experts, rather than the topic being mentioned multiple times by an individual expert. Memos and notes (Urquhart, 2001) supplemented codes and were used to annotate ideas and often to make connections between data from one transcript to another.

As the importance of a 'concept' was of considerable interest to this research, particularly if it was an 'ontology evaluation criteria', sometimes codes covering the same issue were split into two. It was then clear from the coding, whether the expert had said that an issue was, or wasn't important, or whether the issue was just relatively less important (although experts were not always forthcoming about what a "more" or "less" important issue was being compared with).

The resultant variously categorised (template) data are listed in Appendix 20.

6.2.2 Hierarchical Ontology Evaluation Criteria Model

In this third phase of the research (refer again to Figure 6.1), in which coding and categorisation was performed, one goal was to create a hierarchical decision framework suitable for comparatively rating/evaluating ontologies that could address the criteria aspects of RQ1.2.1: *What ontology selection and evaluation criteria are currently used across multi-disciplinary scientific communities (and are selection and evaluation methods consistent with those reported in the literature)*? A hierarchical model was therefore developed through analysis of the interview data, by starting a priori with a template developed from evaluation measures in Finney (2008).

Five high level dimensions were identified in Finney (2008) that grouped "qualitative measures" deemed to fit within each of these dimensions. In Finney (2008) the term "measure" was used erroneously because what are being referred to both in Finney (2008) and by the experts in this study, are ontology ‘evaluation criteria’. A ‘measure’ is by definition a means by which criteria can be assessed. Criteria and measures are considered separate entities in this study. A description of the five dimensions developed in Finney (2008) is provided in Table 6.14.

The data captured from experts, through interview, resulted in the identification of forty-two criteria deemed as being of importance during ontology evaluation, as compared to the twenty-seven criteria listed in Finney (2008). A significant number of these additional criteria were ascribed to the dimension of “Governance” (e.g., Finney had 3 measures in this dimension, experts cited 10). Some of the criteria originally listed by Finney (2008) were also presented very differently by experts in that the criteria were framed in some cases much more generally. For example, in Finney (2008) a stated measure (which is a criterion in this study) appeared as follows: “Is Governance participatory for maintained dictionaries and lists”, which is very specific to the governance of just dictionaries and lists. Two interviewed experts suggested that this specific criterion should instead be framed as “Whether contribution to ongoing (ontology) development and maintenance is encouraged, open and facilitated”.

Table 6.14 Dimension Descriptions

Dimension	Description
Structure	Structural criteria which address matters primarily concerning engineering aspects of the ontology. This includes issues such as the world-view that the ontology models; the derivation of its concepts; the types of relations it permits; the axiomata it supports; its degree of modularity, consistency and extensibility.
Functional Relevance	Functional Relevance criteria which pertain to an ontology’s fitness for purpose (mainly from an engineering perspective). Criteria in this dimension are usually framed ‘relative’ to a particular community, use-case or application.
Usability	Usability criteria pertain to benchmarks set by an individual, or a group regarding the immediate, or longer-term ‘level of difficulty’ associated with using an ontology for the purposes at hand. Criteria in this dimension are also framed ‘relative’ to a particular community, use-case or application.
Maintenance	Maintenance criteria cover operationalised aspects of ontology management best practise (e.g., ontology versioning; availability of up-to-date ontology documentation).
Governance	Governance criteria are those that pertain to how a community organises itself to: set policies; establish protocols; issue guidance; monitor and moderate its own ontology development and manage its own activities.

Where there was mention of criteria that appeared in the (starter) Template, these criteria were retained, or existing criteria were edited to instead reflect expert description and opinion. New criteria were added and placed within an appropriate dimension (as determined by the author because experts raised issues and criteria but did not categorise them). Having placed all of the mentioned criteria into an existing dimension, the data were reviewed again to see if any further groupings or clusters were evident. This included another detailed assessment of the adequacy of the dimensions themselves. Despite a doubling of the data in the Template, the dimensions still appeared robust. Within each dimension it was possible, however, to discern an intermediate grouping of topics (i.e., dimension sub-categories), which have been described in Appendix 20 and depicted in Table 6.15, along with the evaluation criteria that they categorise.

Table 6.15 Hierarchical Classification of Evaluation Criteria (Post Interview Analysis)

Dimension	Sub-Category	Evaluation Criteria
Structure (S)	**Maturity** (SM)	(SM1) Maturity of the ontology (e.g., in terms of how many iterations it has been through).
	Structural Transparency (SS)	(SS1) Having access to documented conceptual models that summarise the scope and organisation of the ontology (e.g., UML, Concept Maps, Concept Dictionaries with URIs and schema descriptions), (SS2) Does the ontology already re-use other ontologies or ontology components and where it does are the dependencies clear ? (SS3) Whether the ontology includes instance sample data. (SS4) Extent to which ontological relationships define context and help give meaning (or definition) to concepts (e.g., discharge is a concept that can have multiple meanings but when coupled with m^3/s tells you it's a "flow"). (SS5) Whether concepts are well-defined both in terms of textual descriptions and logic.
	Conformity (SC)	(SC1) Whether the ontology conforms with language encoding principles and community development rules.
	Engineering (SE)	(SE1) Whether the ontology is extensible in terms of a user's ability to easily add new concepts or specialise existing ones. (SE2) Whether the ontology is sufficiently modular that it is easily able to be re-used. (SE3) Whether the ontology is logically consistent. (SE4) Encoding efficiency (ability to deliver complex and potentially voluminous data in compact structures). (SE5) Does the author have credibility (in developing ontologies and/or within the domain)? (SE6) Whether a visual inspection of the ontology gives confidence after looking at aspects such as: number of concepts, cohesion, tangledness, redundancy of concepts or properties.

Functional Relevance (FR)	**Application Relevance** (FRA)	(FRA1) Whether the ontology meet the use case goals or competency questions. (FRA2) The degree of concept coverage ? (FRA3) The applicability of included vocabularies (e.g., scope, type and size of dictionaries and lists). (FRA4) How focussed is the ontology for its stated purpose (e.g., does it focus on concepts and relations for the stated application or domain, or is it bloated with terms from many different spheres of interest) ?
	Standards Harmonisation (FRS)	(FRS1) Harmonisation with existing (formal or de jour) standards (e.g., with other accepted domain and application ontologies, upper ontologies, Dublin Core, FOAF, VCARD).
Usability (U)	**Ease of Application** (UE)	(UE1) Complexity of ontology (in terms of a user's ability to model instance data using the ontology and the level of expertise required) ? (UE2) Processing affordance (i.e., do the patterns chosen for constructing and encoding the concepts have any potential operating synergy with service software that can recognise these patterns therefore improving the scope to develop and associate reusable data manipulation software). (UE3) How accessible the ontology is (e.g., is it addressable via a URI, is it published in SKOS, XML, RDF or OWL)? (UE4) State of flux (i.e., how often are major revisions released) (UE5) Whether the ontology is interoperable with other ontologies (UE6) Are there easily accessible and mature manipulation tools associated with using or updating the ontology ?
	Sustainability (US)	(US1) The size and scope of the current user implementation base. (US2) Could we readily convince our community to use this ontology ? (US3) Estimation of whether this ontology will be used and sustained into the future.
Maintenance (M)	**Curation** (MC)	(MC1) The maintenance base of the ontology (in terms of number and skill of people maintaining the ontology). (MC2) Are there dedicated "gatekeepers" (or editors/curators) for the ontology? (MC3) Does the ontology have version control (e.g., are old versions maintained and accessible and is versioning apparent for the ontologies in use)?
	User Assistance (MU)	(MU1) Is there any type of help-desk associated with the ontology ? (MU2) Quality and availability of published documentation ?

Governance (G)	**Framework** (GF)	(GF1) Does the ontology development community have a published governance framework (e.g., Clear roles and responsibilities for participants, guiding principles, review processes, repositories, community portal or wiki) ? (GF2) Is there evidence that the ontology conforms with community governance policies and principles (e.g., naming conventions, scope rules, version control) ? (GF3) Does the community review or moderate individual ontology developments ?
	Community (GC)	(GC1) How mature the community is in terms of ontology development and its longevity and cohesion as a community of practise. (GC2) Whether the community is institutionally-backed (e.g., by standards bodies, government organisations, seed-funded start-ups). (GC3) Whether the community has a sustained core group supporting its ontology activities. (GC4) Whether the community's mandate is obvious and well-bounded.
	Behaviours (GB)	(GB1) Whether the community actively encourages open use of their ontologies. (GB2) Whether contribution to ongoing development and maintenance is encouraged, open and facilitated. (GB3) Whether the governance and the core group of developers engender trust and credibility.

The rationale behind the introduction of the sub-categories is that they break up the dimensions (which are at a high level of abstraction) into groups of criteria that pertain to different aspects of the dimension. Because an aim of this research is to gain a better understanding of which particular criteria, used by experts in the field, carry the most weight in ontology assessment and selection exercises, it was theorised that sub-categorising the data would provide a more manageable means for experts to focus on the different themes of the decision-making process. By providing chunks of information for them to process that are more detailed than dimensions, but less detailed than criteria, there is an intermediate level on which to focus expert judgements regarding the relative levels of importance they would place on issues in a decision-making process. Clustering of criteria at the dimension level (alone) was considered by the author to be too coarse for a weighting exercise (utilising AHP techniques). These sub-categories therefore provide for the intermediate classification of criteria. Brugha (1998) has shown that more 'structured' models can lead to higher consistency in expert responses.

The various considerations taken into account when coding and categorising these data (after Baker *et al.,* 2002) have already been summarised in the general methods chapter (i.e., Chapter 3, section

3.3.2). All dimensions, sub-categories and criteria were coded for the purposes of reference and subsequent analysis.

6.2.3 Additional Information Relevant To Expert Stratification And Subsequent Data Interpretation

Apart from using interview data to develop a hierarchical model of ontology evaluation criteria (which will be subjected to further analysis later within this chapter), it was also possible to derive additional information which:

- Was able to be used to validate data drawn from the screening survey (for example, sample development activities provided by expert 'PF' during interviews lead to a re-evaluation of his discipline groupings because it transpired that he had also worked closely with oceanographic communities).
- Permitted deeper investigation of expert responses with respect to cited ontology evaluation criteria.
- Provided for further stratification of experts (e.g., in terms of their community's maturity and focus of ontology development activity).

All of this information is drawn upon in subsequent phases of the research (particularly phase 6 shown in Figure 6.1).

Can The Level Of Expert Experience Be Inferred From Interview Data ?

Tables 6.16 and 6.17 show the disciplines that each expert has covered (from an ontology development perspective) and the colour used for each expert indicates the number of evaluation measures they raised during interview. Table 6.16 shows all evaluation criteria bar those mentioned in the 'governance' dimension and Table 6.17 shows all mentioned evaluation criteria, plus criteria and issues raised under governance codes. These results were separated in case the coded governance concepts, which also included governance issues, not necessarily just evaluation criteria, confounded results.

Whilst it is an extremely crude measure to equate the number of evaluation criteria mentioned, with an expert's level of sophistication and experience with ontology development and ontology evaluation, it never-the-less is considered a potential indicator that could be harnessed (judiciously). Used in combination with other indicators, it could help differentiate and interpret a particular expert's relative weight assignment for ontology evaluation criteria in the weighting exercise which will follow later in this chapter.

Table 6.16 Interviewees Discipline Coverage (Number of Evaluation Criteria Cited)

Expert Code*	Hydrology	Biology	Oceanography	Atmosphere Meteorology	Geoscience	Energy
SC						
RA						
LL						
PF						
JH						
JG						
RL						
RH						
CA						
WD						
DH						
MH						
PM						

*Legend	Number of evaluation measures cited.
	26
	17-20
	9-13
	6-7

Table 6.17 Interviewees Discipline Coverage (Evaluation Plus Governance Issues Cited)

Expert Code*	Hydrology	Biology	Oceanography	Atmosphere Meteorology	Geoscience	Energy
SC						
RA						
LL						
PF						
JH						
JG						
RL						
RH						
CA						
WD						
DH						
MH						
PM						

*Legend	Number of evaluation criteria plus governance issues cited.
	>=26
	<=25
	<=20
	<=15
	<= 10

The spread (diversity) of disciplinary experience was also thought to be another possible useful indicator of depth of experience. Unfortunately from the results, as collated, there appears to be little correlation (figuratively) between the number of evaluation criteria mentioned and the scope of the disciplines covered by the expert (see Table 6.16). Many of the experts who offered up a significant number of criteria (except for 'PF') had worked quite exclusively in the one discipline.

Likewise, an expert such as 'SC' who had worked across all five listed disciplines was not a particularly high scorer with regards to the number of non-governance related evaluation criteria mentioned. The lack of correlation between disciplinary association and the number of mentioned criteria isn't altered significantly even when governance criteria (and issues), mentioned by each participant, are added to the scores (see Table 6.17). The main difference is a relative change in rank for 'SC' in terms of the total number of criteria/issues cited. The data was not tested for statistical significance because the sample sizes were too small and the results were so visually obvious.

The number of evaluation criteria cited by individuals was also cross-checked against other data garnered from the screening survey (such as an expert's development experience with a particular language, experience with formal or informal ontologies and their level of expertise). There didn't appear to be any obvious connection between the number of criteria cited by an expert and any of these other types of expert stratifications

Although the self-assessed skill characterisations, provided through the screening survey, appeared in all but two cases to be relatively accurate (armed with the interview data) it was still difficult to gauge the relative depth of an individual's expertise, particularly in the area of ontology re-use. One self-characterisation (made by 'CA' during the Screening Survey) appeared perhaps to be an overestimate of their ontological skills and the other (by 'JG') an underestimate of theirs. 'CA' by their own admission at interview was still really developing his ontological skills.

Given that it has already been shown (via examination of information in Tables 6.18 and 6.19) that there was little overt evidence in the data to infer an expert's level of sophistication with respect to ontology development, using the range of disciplines they had covered as a surrogate measure, it was decided to look instead at the (ontological) maturity of the communities they had worked in as a possible clue to the depth of their expertise. Understanding the depth of each expert's skills and the range of their experience is important in terms of how an expert's opinion might be interpreted in subsequent phases of analysis in this study, particularly when examining why expert opinions diverge.

From Table 6.18, which summarises key facets of the communities who benefit from the assistance of the interviewed experts, several things are evident. In all but one expert's case (i.e., 'CA') the existing user-base for the developed ontological products is relatively large. In some cases the expert is either the sole ontological engineer for a community development effort, or they work with one other person (e.g., in cases involving 'PM', 'RH', 'RL', 'CA', 'JH'). In other cases the expert is part of a larger team (e.g., 'PF', 'SC', 'RA', 'MH', 'WD', 'PM', 'DH', 'LL'). It could be reasoned that those who work in larger teams have greater peer support and therefore may be exposed to a wider array of issues than those who work more predominantly on their own. 'PF', 'SC', and 'LL' are experts who have worked across a relatively large number of disciplines (and communities), have contributed a relatively high number of evaluation and governance issues during interview and have all worked in ontology development team environments.

Ontology Application

Another observation from Table 6.18 is that most community-based ontology development activities are less than 5 years old reflecting the immaturity of the activity in general (even in relatively cohesive and long-lived communities). Most community activity also seems to be targeting ontology development for the purposes of "improving data search functionality", or "enhancing community ability to integrate data from disparate sources" and/or "improving the interoperability of distributed systems". Many of the biology-centric communities have the additional application of "data annotation". Only a couple of the primary use-cases are mentioned as being application-specific (e.g., those mentioned by 'PF' and 'CA'), whilst the remainder of community activity appears to be striving for development of much broader-based (domain standards-setting) ontologies.

Table 6.18 Experts & Associated Community Characterisations (As Discussed In Interview)

Expert Code	Affiliated Community Type/Name	Approximate Size	Active (Ontology) Development Base	Maturity Estimate (with respect to ontology issues)	Primary Ontology Application
RH	TDWG	180 people at 2008 TDWG annual meeting.	1 person (for core ontology development) plus working groups for associated domain input.	Existed since 1985 but have only moved into ontologies in last 5 years.	Global ontology standard for interoperable exchange of biological data.
WD	Evolutionary biology and model organism (ZebraFish) (Phenoscape) (CARO)	80 people on mailing list. But interested user-base may be much larger.	6 ontology curators in total. But one curator per ontology. Appear to work as a team to some extent.	>5 years but have borrowed existing ontologies (with a longer development history) from	Ontologies for describing morphological characteristics (of fish). Includes anatomical reference

				the biomedical community (OBO).	ontologies. Ontologies are often applied in data annotation.
PF	Solar Terrestrial Community.	800 (600 have web accounts and 300 go to annual meetings).	12 experts at NCAR and another 10-15 experts at local Uni (including grad students).	Community together 20+ years – but ontology development only in last 5 years.	Application areas encompassed are searching, data integration, provenance, data mining, data product-type ontology - for describing data products. Often application-specific, except for topics such as UoM etc.
	Physics Community.	20 core research groups (>1000 people).	As above.	As above.	As above
	Oceanography Community.	Users of an Oceanographic Institution biological and chemical oceanography clearinghouse.	3 people	Community had worked together for "many years on vocabularies". < 1 year on ontologies.	As above
JH	Australian Spatial Data Standards (ANZLIC). Also ISO TC 211 participation.	> 100 people	2 people (one technically qualified within Geoscience Australia).	Spatial community has > 10 year history but only recent foray into ontologies (< 3 years ?)	Data searching, integration, interoperability. Standards setting.
LL	CSIRO Water For A Healthy Country Flagship (including activity with WIRADA, CUAHSI and GeoFabric).	CSIRO Flagship involves > 1000 people. CUAHSI is a multinational, multi-partner community (>100 formal members, but with a much larger user base).	> 5 people within CSIRO, possibly more within the Bureau of Meteorology. Unknown number working for CUAHSI.	WIRADA has been funded since 2007. Ontology work begun from inception. CUAHSI has been a community since 2005. Ontology work is only recent < 3 years.	International and national water data transfer standards. Encompassing data integration, interoperability, intelligent monitoring.
	W3C Incubator Group on Sensor Network Ontologies.	- unknown	-unknown	3 years but also using pre-existing ontologies.	Sensor and observation data integration, intelligent monitoring, sensor control.

DH	TDWG Atlas of Living Australia	As above. 17 Australian and State government agencies (> 1000 people).	As Above. > 20 people	As Above. Funded project that has been running for 3 years.	As Above. Biodiversity reference Portal for Australian flora and fauna. Key applications include searching, data integration, data annotation and data mining.
PM	Animal behaviour (a past project).		Mentioned two teams working on ontology and workshops	5 years	Putting animal behaviour terms into a gene ontology.
	Fishes Taxonomy Catalogue (Phenoscape) (Teleost)	20 on a mailing list and a further 20 who are interested.	1 person	3 years (but also borrowed existing OBO ontologies)	Fish systematic ontology
CA	Coral Reef Ecology Community	Limited as yet.	1 person	< 2 years old.	Specifically targeting Coral Reef observations and measurements. Application-specific.
JG	Marine Metadata Interoperability (MMI)	200-300 people engaged out of a community of >10,000.	2 principal experts (plus 5-7 domain experts assisting through teleconferences).	> 5 years old.	Platforms and Devices Ontologies. Applications include data integration, interoperability, searching.
	W3C Incubator Group on Sensor Network Ontologies.	As above.	As above.	As above.	As above.
RA	CSIRO Water Flagship	As Above.	As for LL Above.	As for LL	As Above.
	ANZLIC	As Above.	As for JH Above.	As for JH	As Above.
	International Hydrographic Organisation.	> 100 people	-unknown	-unknown	Interoperability, data integration, intelligent navigation.
	GeoSciML	International geological survey institutions plus CSIRO (8-10 organisations).	15-20 people (with 6 usually active at any one time)	6 years	Data interoperability, data integration.
RL	MMI	As above.	As above.	As above.	As above.
	Climate Science (based around NERC Data Grid vocabulary activity)	> 100 people spread over many UK and international institutions.	1 developer.	5 years (but drawing on 20+ years of domain activity).	Data search, integration and interoperability.

RL	SeaDataNet (incorporating SeaVox)	>100 people spread throughout Europe. Many (40) participating institutions.	1 developer.	3 years.	Data search, integration and interoperability.
SC	CSIRO Water Flagship	As above.	As above.	As above.	As above.
	GeoSciML	As above.	As above.	As above.	As above.
	Climate Science	-unkown	-unkown	-unkown	-unkown
	OGC standards Body	-unkown	-unkown	-unkown	-unkown
MH	Evolutionary biology and model organism (ZebraFish) (Phenoscape) (CARO)	As Above.	As Above.	As Above.	As Above.

Roles Played By Experts

Coding and analysis of the interview data also revealed that there were three main types of roles played by experts within their respective communities. The degree to which each role was played was dependent on the specifics of the community in question and on the period of an expert's community affiliation. The three roles identified included:

(a) Leader, Driver, Initiator

This role was characterised by the expert playing a vital role in a community's decision to develop semantic approaches for the description of community data. In many cases they (alone, or with the support of one or two others) developed incubator activities to seed the community's development efforts. Most experts who had played this type of role within a specific community, tended to remain affiliated with the group that they had helped to found and they may have continued to be a driver of that community's activity (usually with the benefit of an expanded interest-base).

(b) Maintainer

In other scenarios experts may no longer lead or drive activities within a community, but have continued to play a role which helps maintain community activity. This role could be technical, governance or marketing focussed.

The "maintenance" role, however, also encompassed the activities of experts who have not been associated with initiating a communty's ontology development activity, but who contribute on a regular basis to specific ontology development functions (e.g., as a specific ontology curator).

(c) Specialist (ontology and/or domain) Advisor

All experts had played the role of specialist technical advisor (which was one of the reasons that they were each selected to participate in this study). In this role their specific expertise in building and evaluating ontologies is used by a community to aid in a community's ontology development effort. The expert may be a continuing source of expertise for a community, or their role may have been short-lived.

Of the experts interviewed, nine ('PF', 'DH', 'LL', 'JG', 'RL', 'RA', 'RH', 'SC', 'MH') had played all three roles and the remainder had primarily played roles (b) and (c).

Frequency Of Evaluation Criteria Citation And Verbal Importance Ranking

Another view of an expert's data can be summarised as in Table 6.19, where it is possible to use the evaluation criteria as the focus for differentiation, rather than the experts themselves. In this Table we can see which experts mentioned particular (coded data) evaluation criteria (for all dimensions except for 'Governance'). In Table 6.19 if a code related to the same criteria and the only difference was whether it was considered important, or less important (i.e., suffixed with "i" or "li"), it was collapsed into one code because here (in this Table) we are interested in whether the issues is considered at all, not whether it is considered more or less important. Table 6.19 simply represents the presence or absence of mention of specific criteria, by expert.

Table 6.20 shows similar types of results but for data originally coded with a prefix of "g" (denoting governance-based criteria and issues). In both tables the criteria are coloured to help visually appreciate how many experts have mentioned each concept. This data serves to give some appreciation for the topicality of the coded concept within the expert cohort. More frequently mentioned concepts could be considered to be the more universally (and cross-discipline) considered issues. For a better understanding of the relative importance that is placed on a concept by experts, it is necessary to examine which of these tagged concepts were mentioned with qualification. Table 6.21 lists those concepts where experts who mentioned them, qualified the relative importance of the criteria or issue (usually loosely and without specific reference to other issues). During interviews if a concept was mentioned at all it was considered of some importance, so only those concepts where there was split opinion about the importance factor, or where the concept was mentioned because it was considered less important than other issues, have been listed in Table 6.21. These data will be important later, for validation and interpretation purposes, when looking at how expert's relatively rated individual ontology evaluation criteria in the hierarchical model weighting exercise.

Table 6.19 Evaluation Criteria Mentioned By Experts (Not Including Governance)

Evaluation Criteria	RH	WD	PF	JH	LL	DH	PM	CA	JG	RA	RL	SC	MH
em:actively_worked_on			1	1					1				
em:complexity	1	1	1			1		1					1
em:compliance_with_OBO_ Foundry principles							1						
em:concept_coverage		1	1	1	1	1	1	1	1		1		1
em:conceptual_models			1	1					1				1
em:contribute						1			1				
em:data_flux			1	1		1			1				
em:dictionaries_lists	1		1	1		1		1			1		
em:documentation			1	1					1				1
em:encoding_efficiency			1			1	1					1	
em:engineering_evaluation			1		1				1				
em:extensibility	1	1	1	1		1		1					1
em:focussed	1												1
em:gatekeepers_editors													1
em:general_accessibility					1						1		
em:help-desk													1
em:instance_samples			1	1	1	1			1				
em:interoperability		1											1
em:iso_metadata	1	1	1			1	1					1	
em:know_authors			1				1			1			1
em:language_encoding		1	1	1		1						1	1
em:logically-consistent													1
em:mainipulation_tool	1				1	1							
em:maintained_plus_applic ation			1										
em:maintenance_base			1			1			1	1			

em:major_revisions_a_detraction							1						
em:maturity_of_community							1						
em:maturity_of_ontology							1						
em:maturity_of_tools							1						
em:modularity	1	1	1	1	1	1	1	1				1	1
em:processing_affordance			1	1			1						
em:re-uses_ontologies			1										
em:relationships-that-define-context-and give-meaning					1								
em:saleability_to_community			1		1				1		1	1	
em:size-credibility-of-governance-body					1								
em:skilled_ontologists_involved							1						
em:survivorship			1	1		1	1						
em:transparent_dependency										1			
em:use-case	1	1	1	1		1	1	1	1	1			
em:user-base		1	1	1		1	1	1	1	1		1	1
em:versioning		1	1			1	1		1			1	1
em:well-defined				1	1					1	1		1
em:will-become-standard		1		1									
em:XSD	1		1	1	1	1							

***Legend**	**Expert Spread**
	Mentioned by <=10 experts
	Mentioned by <=8 experts
	Mentioned by <=6 experts
	Mentioned by <=4experts
	Mentioned by <=2 expert

Table 6.20 Governance Issues Mentioned By Experts

Governance Issues	RH	WD	PF	JH	LL	DH	PM	CA	JG	RA	RL	SC	MH
G:best_practise_guidance									1				1
G:clarity_of_roles				1						1			
G:collaboration_tools		1				1					1	1	
G:community_cohesion			1										
g:community_size			1				1		1			1	
G:community_stable_over_time												1	
G:consensus_decision-making	1								1			1	
G:facilitation			1										
G:formal_decision_making_processes			1			1							
G:framework					1	1	1		1				1
G:Fully_funded			1										
G:gatekeeprs		1									1		1
G:governance_bodies			1		1					1		1	
G:governance_types											1		
G:incentives													1
G:informal decision-making_processes			1				1		1			1	1
G:leveraged_investment	1		1			1						1	
G:light_governance						1	1						
g:limit_mandate	1					1						1	
G:long-term_investment_cycle			1									1	
G:majority_rules		1											
G:OGC_as_broker												1	
G:ontology_reviews													1
G:policies										1			
G:project-based_approach			1						1				
G:public_access_to_community_work							1					1	
g:re-using_and_forking												1	
G:rivalry_between_Communites							1		1	1		1	
G:small_development_group		1	1						1		1		1

g:stable_core_group_contributing		1	1			1						1	
G:strong_institutional_backing			1									1	
G:trust	1			1			1			1	1		
em:transparency_participatory		1		1	1	1	1		1	1	1	1	1

*Legend	Expert Spread
	Mentioned by =>5 experts
	Mentioned by 4 experts
	Mentioned by 3 experts
	Mentioned by 2 experts
	Mentioned by 1 expert

Table 6.21 Qualified Tagged Concepts (Spread Of Expert Opinion)

Tagged Concept	Number of Experts Considering The Concept Important	Number of Experts Considering The Concept NOT AS Important
em:complexity	5	1
em:concept_coverage	7	3
em:data_flux	3	1
em:encoding_efficiency	0	4
em:extensibility	6	1
em:instance_samples	4	1
em:iso_metadata	1	5
em:processing_affordance	2	1
em:survivorship	0	4
em:user_base	8	2
em:versioning	6	1
em:well_defined	5	1
em:XSD	3	2

6.2.4 Summary

In this section (6.2) the reader has been provided with an explanation of how the interview data was captured and analysed (i.e., coded, collated and in some cases interpreted). Of significance was the generation of a three-tiered, hierarchical ontology evaluation criteria model (Table 6.15). Whilst this model is used as the focus of analyses undertaken in the remaining research phases (i.e., 4 through 8, refer to Figure 6.1), most of the data presented in this section, in the form of summarised results, becomes input to analysis and interpretative activity which is performed primarily in phase 6. In the next section (i.e., 6.3) data that have been presented in this section are drawn upon, particularly the various stratifications, to assist in explaining patterns observed amongst expert with respect to their relative ratings of the importance of ontology evaluation criteria (i.e., those criteria listed in Table 6.15). Other data that was captured and analysed, which was not presented in this section is summarised and discussed in Chapter 7.

6.3 Quantitative Pair-wise Comparison Survey Data Analyses

In order to directly address RQ1.2.2. - *Is it feasible to derive a weighted evaluation criteria model in which criteria are rated according to importance ? If so, are evaluation criteria of equal weight or do some carry more importance than others,* it was necessary to issue another survey. This second text-based survey was emailed to the thirteen experts (refer Appendix 21). They were introduced to the hierarchical model shown in Table 6.17 and asked to answer approximately eighty questions that pair-wise compared each criterion, within each sub-category, within each dimension with respect to the level of importance that they would place on each criterion in a decision-making process involving ontology selection. These comparisons were performed in the context of the AHP (Saaty, 1980). A nine point rating scale recommended by Saaty (1980) was used for importance ratings, despite reservations about this scale which have been expressed by French (1988) and Goodwin and Right (1988), as already discussed in Chapter 3. Given that this study will be performing pair-wise consistency analysis, issues of concern regarding transitivity will be addressed. See Table 6.22 for the nine-point scale. These eighty questions also included pair-wise comparison of each sub-category and each dimension.

Only seven of the thirteen experts returned a survey despite many reminders and solicitations. Fortuitously, at this point in time, the author was provided with another expert referral. Given the relatively low number of survey responses (54%), despite already having commenced data analysis, the additional expert (denoted as ‘KS’) was contacted and asked if she would participate in the “weighting” exercise. ‘KS’ agreed and returned a pair-wise comparison survey, bringing the total to eight responses. The new ontology expert, ‘KS’, is active in the MMI community and works within biological and oceanographic communities. ‘KS’ considered herself to be a “domain expert with ontological skills”.

The answers provided by each expert (from completing the survey) were converted into reciprocal (skew-symmetric) matrices and placed in a Microsoft™ Excel spreadsheet for analysis using various AHP techniques (Saaty, 1980). For example, if an expert rated option A as being 9 times more important than option B, then it can be reasoned that option B is one ninth (1/9) as important as option A. These values are then used to populate a matrix and the normalised Eigenvector solution for the matrix is the calculated weights. At the same time as the eigenvector solution is calculated a consistency ratio is computed that provides an estimate of how well the data fit to a computationally ideal result.

Table 6.22 Importance Rating Scale Used In Survey

Intensity of Importance *	Definition	Explanation
1	Equal importance	Two elements contribute equally to the objective.
3	Moderate importance	Experience and judgement slightly favour one element over another.
5	Strong importance	Experience and judgement strongly favour one element over another.
7	Very strong importance	One element is favoured very strongly over another; its dominance is demonstrated in practise.
9	Extreme importance	The evidence favouring one element over another is of the highest possible order of affirmation.

*Intensities of 2, 4, 6, and 8 can be used to express intermediate values if an expert felt they were warranted.

Recall from an earlier explanation in Chapter 3 that data can be inconsistent due to ordinal inconsistencies (i.e., where the order of ranked options is logically inconsistent) or due to cardinal inconsistencies (i.e., where options are ascribed ratings that are logically inconsistent). The consistency ratio, however, doesn't discriminate between the type of inconsistency that exists it just tells us that the data is not consistent. Saaty (1999) recommends that the computed consistency should be < 0.10. Other practitioners of AHP disagree (Kauko, 2002; Kryvobokov, 2005) and allow higher inconsistency levels in the data.

For the purposes of this study 0.14, or less, was selected as the cut-off for inconsistent data. This level was determined on the basis of results depicted in upcoming Table 6.25 and the degree to which each expert's pair-wise comparison data had to be manipulated to improve consistency. As will be seen later, a balance was required between reaching consistency and changing expert data such that the data no longer represented an expert's view. Most inconsistent comparison data were able to be brought down to the 0.14 level of consistency with only a few manipulations of pair-wise comparison results (and the overall affect on the data was in most cases minimal, see upcoming section 6.3.2). However, to get inconsistent expert data below 0.14 and down to 0.1, or less, would have required many more data transformations that would have significantly changed patterns in an expert's data. This level of manipulation was considered undesirable and it was also felt that experts would not be willing to accept these more significant pattern changes to their responses.

6.3.1 Converting Raw Ratings To Weights

As an example of how the raw comparison data is converted to weights (using AHP Analysis), 'KS' provided the following data for the questions on the structural transparency sub-category from the comparative survey:

A. Structural Transparency Comparisons

Of the following pairs of structural transparency measures outlined below which one is more important in helping you make your ontology selection and by how much ?

Q1.1

SS1. Having access to documented conceptual models that summarise the scope and organisation of the ontology (e.g., UML, Concept Maps, XSDs), or

SS2. Does the ontology already re-use other ontologies or ontology components and where it does are the dependencies clear?

More Important Option	Importance Rating
SS1	5

Q1.2

SS1. Having access to documented conceptual models that summarise the scope and organisation of the ontology (e.g., UML, Concept Maps, XSDs), or

SS3. Whether the ontology includes instance sample data.

More Important Option	Importance Rating
SS1	3

Q1.3

SS1. Having access to documented conceptual models that summarise the scope and organisation of the ontology (e.g., UML, Concept Maps, XSDs), or

SS4. The extent to which ontological relationships define context and help give meaning (or definition) to concepts (e.g., discharge is a concept that can have multiple meanings but when coupled with m^3/s tells you it's a "flow").

More Important Option	Importance Rating
SS4	2

Q1.4

SS1. Having access to documented conceptual models that summarise the scope and organisation of the ontology (e.g., UML, Concept Maps, XSDs), or

SS5. Whether concepts are well-defined both in terms of textual descriptions and logic.

More Important Option	Importance Rating
SS5	6

Q1.5

SS2. Does the ontology already re-use other ontologies or ontology components and where it does are the dependencies clear ?, or

SS3. Whether the ontology includes instance sample data.

More Important Option	Importance Rating
SS3	3

Q1.6

SS2. Does the ontology already re-use other ontologies or ontology components and where it does are the dependencies clear ?, or

SS4. The extent to which ontological relationships define context and help give meaning (or definition) to concepts (e.g., discharge is a concept that can have multiple meanings but when coupled with m^3/s tells you it's a "flow").

More Important Option	Importance Rating
SS4	6

Q1.7

SS2. Does the ontology already re-use other ontologies or ontology components and where it does are the dependencies clear, or

SS5. Whether concepts are well-defined both in terms of textual descriptions and logic.

More Important Option	Importance Rating
SS5	8

Q1.8

SS3. Whether the ontology includes instance sample data, or

SS4. The extent to which ontological relationships define context and help give meaning (or definition) to concepts (e.g., discharge is a concept that can have multiple meanings but when coupled with m^3/s tells you it's a "flow").

More Important Option	Importance Rating
SS4	4

Q1.9

SS3. Whether the ontology includes instance sample data, or

SS5. Whether concepts are well-defined both in terms of textual descriptions and logic.

More Important Option	Importance Rating
SS5	7

Q1.10

SS4. The extent to which ontological relationships define context and help give meaning (or definition) to concepts (e.g., discharge is a concept that can have multiple meanings but when coupled with m^3/s tells you it's a "flow"), or

SS5. Whether concepts are well-defined both in terms of textual descriptions and logic.

More Important Option	Importance Rating
SS5	5

These comparison data for 'KS' were translated into a matrix, as shown in Table 6.23. The Eigenvector solution (described earlier in Chapter 3) for this matrix was calculated using a Microsoft™-Excel spreadsheet (using a macro listed in Appendix 22). The macro function iteratively raises the matrix to powers that are successively squared each time. Matrix row sums are calculated at each step and normalised. The function stops when the difference between the sums in two consecutive calculations is smaller than a prescribed value (i.e., 0.0001).

The data in Table 6.23 implies that 'KS' thinks criterion 'SS5' carries the most importance (57.5%) when making decisions based on "Structural Transparency" factors. Alternatively, it could be stated that she thinks criterion 'SS5' is about four times more important than criterion SS1 (from the ratio of SS5/SS1).

Table 6.23 Judgement Matrix For Structural Transparency Sub-Category (KS)

	SS1	SS2	SS3	SS4	SS5	Weights (eigenvector)
SS1	1	5	3	0.5	0.166	0.130
SS2	0.2	1	0.333	0.166	0.125	0.035
SS3	0.333	3	1	0.25	0.142	0.065
SS4	2	6	4	1	0.2	0.195
SS5	6	8	7	5	1	0.575

The same macro function also computed the consistency index (from which we can derive the consistency ratio). The Principal Eigenvalue (λ_{max}) is first required to compute the consistency index. An Eigenvalue approximation can be obtained from the summation of products between each element of the Eigenvector and the sum of columns of the reciprocal matrix:

$$\boldsymbol{\lambda_{max}}=[(1+0.2+0.33+2+6)X0.130]+[(5+1+3+6+8)X0.035]+[(3+0.33+1+4+7)X0.065][(0.5+0.16+0.25+1+5)X0.195]+[(0.16+0.125+0.14+0.2+1)X0.575]$$

$$\mathbf{= 5.32}$$

The consistency index is given by:

CI = $(\boldsymbol{\lambda_{max}} - n)/n - 1$

Where n = the size of the judgement matrix (in the example n = 5).

Therefore,

CI = (5.32 – 5)/4 **= 0.08**

The **CI** is then compared with a table of values from Saaty (1980), where he computed Random Consistency Indices (**RI**) for up to 500 matrices of various sizes (see Appendix 23 for a table of Saaty's consistency values). The **RI** for a matrix of n = 5 from this table is 1.11. The Consistency Ratio is defined as:

CR = CI/RI

Therefore,

CR = 0.08/1.11 = 0.072 *

*Note in the example above the author has rounded numbers and used approximation calculations for simplicity of demonstration purposes. The actual computed value of CR for 'KS', for this example using the MS-Excel macro is 0.075.

The judgement matrices with computed **weights**, **λ_{max}, CI and CR** for all comparisons, for one expert (as an example) are found below in Table 6.24.

Table 6.24 Matrix Results For Expert 'KS'

Structural Transparency:

	SS1	SS2	SS3	SS4	SS5		Weights
SS1	1	5	3	0.5	0.166667		0.130
SS2	0.2	1	0.333333	0.166667	0.125		0.035
SS3	0.333333	3	1	0.25	0.142857		0.065
SS4	2	6	4	1	0.2		0.195
SS5	6	8	7	5	1		0.575
						Lambda max	5.328
						Consistency index	0.082
						Consistency ratio	0.075

Structural Engineering:

	SE1	SE2	SE3	SE4	SE5	SE6		Weights
SE1	1	2	1	6	5	0.2		0.153
SE2	0.5	1	0.333333	5	4	0.142857		0.094
SE3	1	3	1	6	5	0.2		0.167
SE4	0.166667	0.2	0.166667	1	0.333333	0.125		0.028
SE5	0.2	0.25	0.2	3	1	0.125		0.043
SE6	5	7	5	8	8	1		0.515
							Lambda max	6.478
							Consistency index	0.096
							Consistency ratio	0.077

Structural Dimension:

	SM	SS	SC	SE		Weights
SM	1	0.142857	0.25	8		0.264
SS	7	1	5	0.333333		0.340
SC	4	0.2	1	0.166667		0.148
SE	0.125	3	6	1		0.248
				Lambda max		8.845
				Consistency index		1.615
				Consistency ratio		1.871

Application Relevance:

	FRA1	FRA2	FRA3	FRA4		Weights
FRA1	1	3	4	4		0.532
FRA2	0.333333	1	2	3		0.241
FRA3	0.25	0.5	1	1		0.118
FRA4	0.25	0.333333	1	1		0.108
				Lambda max		4.062
				Consistency index		0.021
				Consistency ratio		0.024

Functional Relevance:

	FRA	FRS		Weights
FRA	1	5		0.833
FRS	0.2	1		0.167
		Lambda max		2.000
		Consistency index		0.000
		Consistency ratio		0.000

Ease of Application:

	UE1	UE2	UE3	UE4	UE5	UE6		Weights
UE1	1	5	5	4	5	4		0.460
UE2	0.2	1	1	0.333333	1	0.333333		0.066
UE3	0.2	1	1	0.333333	1	0.333333		0.066
UE4	0.25	3	3	1	3	1		0.171
UE5	0.2	1	1	0.333333	1	0.333333		0.066
UE6	0.25	3	3	1	3	1		0.171
						Lambda max		6.130
						Consistency index		0.026
						Consistency ratio		0.021

Sustainability:

	US1	US2	US3		Weights
US1	1	0.166667	0.25		0.082
US2	6	1	4		0.682
US3	4	0.25	1		0.236
				Lambda max	3.108
				Consistency index	0.054
				Consistency ratio	0.094

Usability:

	UE	US		Weights
UE	1	0.25		0.200
US	4	1		0.800
		Lambda max		2.000
		Consistency index		0.000
		Consistency ratio		0.000

Curation:

	MC1	MC2	MC3		Weights
MC1	1	1	3		0.429
MC2	1	1	3		0.429
MC3	0.333333	0.333333	1		0.143
			Lambda max		3.000
			Consistency index		0.000
			Consistency ratio		0.000

User Assistance:

	MU1	MU2		Weights
MU1	1	0.25		0.200
MU2	4	1		0.800
		Lambda max		2.000
		Consistency index		0.000
		Consistency ratio		0.000

Maintenance:

	MC	MU		Weights
MC	1	3		0.750
MU	0.333333	1		0.250
		Lambda max		2.000
		Consistency index		0.000
		Consistency ratio		0.000

Framework:

	GF1	GF2	GF3		Weights
GF1	1	0.25	1		0.167
GF2	4	1	4		0.667
GF3	1	0.25	1		0.167
			Lambda max		3.000
			Consistency index		0.000
			Consistency ratio		0.000

Community:

	GC1	GC2	GC3	GC4		Weights
GC1	1	0.25	0.333333	1		0.105
GC2	4	1	3	4		0.528
GC3	3	0.333333	1	3		0.262
GC4	1	0.25	0.333333	1		0.105
				Lambda max		4.083
				Consistency index		0.028
				Consistency ratio		0.032

Behaviours:

	GB1	GB2	GB3		Weights
GB1	1	1	1		0.500
GB2	1	1	1		0.250
GB3	1	1	1		0.250
			Lambda max		3.000
			Consistency index		0.000
			Consistency ratio		0.000

Governance:

	GF	GC	GB		Weights
GF	1	0.333333	0.25		0.122
GC	3	1	0.5		0.320
GB	4	2	1		0.558
			Lambda max		3.018
			Consistency index		0.009
			Consistency ratio		0.016

All Dimensions:

	S	FR	U	M	G		Weights
S	1	1	3	5	7		0.352
FR	1	1	3	5	7		0.352
U	0.333333	0.333333	1	5	6		0.187
M	0.2	0.2	0.2	1	4		0.073
G	0.142857	0.142857	0.166667	0.25	1		0.035
					Lambda max		5.326
					Consistency index		0.081
					Consistency ratio		0.075

6.3.2 Improving Data Consistency

A large number of the expert matrices had inconsistency ratios > 0.10 (see Table 6.25). Several methods for improving consistency were investigated (and the pros and cons of using each are discussed later in Chapter 7). Ideally what is required is the smallest change possible that will bring the expert data to an appropriate level of consistency, whilst preserving the integrity of the respondent's original answers. The most obvious method is to simply ask each expert who has inconsistent data to take the survey again, but given no guidance as to why their answers were considered inconsistent in the first place, there is no guarantee that an expert will be consistent the next time around. A Delphi approach (Schmidt, 1997) was also considered, where feedback is given in the form of a group result and each expert can elect, or not, to change their respective ratings on the basis of how other experts have rated criteria. But survey respondent burden is high using these types of repeat survey approaches and the risk of losing more experts was a risk with a high likelihood of occurrence.

The method eventually chosen to help bring expert data to a greater level of consistency was to use an iterative process (facilitated by an 'R' open source custom-written software program) that changed matrix values. The changes were recorded and reflected back to experts, where their agreement was sought to the changes made to the original ratings. This request to modify the data was provided to experts by email. The original ratings data were included in the email along with a plot of old and new values so that it was visually easily for the expert to see what effect suggested changes to their ratings had on the calculated weights after the change (e.g., see Figures 6.5 to 6.10 for sample plots provided to expert 'DH'). Appendix 24 provides a copy of a typical email to an expert requesting that they consider a change to their data (this represented phase 5 in Figure 6.1).

Table 6.25 Summary of Consistency Ratios (> 0.10)

Sub-Category	RA	LL	MH	WD	DH	CA	KS	JH
Structural Transparency	0.866	0.172	0.141	0.281	0.168	0.163	0.075	0.254
Structural Engineering	0.389	0.223	0.138	0.097	0.151	0.121	0.077	0.128
Structural Dimension	0.242	NC	0.569	0.119	0.188	0.323	1.871	0.236
Application Relevance	0.166	0.229	0.005	0.017	0.088	0.098	0.024	0.102
Functional Relevance	0	0	0	0	0	0	0	0
Ease of Application	0.279	0.097	0.121	NC	0.392	0.268	0.021	0.143
Sustainability	0.490	0.025	0	0.118	0.118	0.118	0.094	0.047
Usability	0	0	0	0	0	0	0	0
Curation	0.011	0.011	0.006	0.006	0.160	0.034	0	0.003
User Assistance	0	0	0	0	0	0	0	0
Maintenance	0	0	0	0	0	0	0	0
Framework	0.118	0.025	0.025	0.118	0.034	0.034	0	0.034
Community	0.148	0.106	0.005	0.129	0.076	0.102	0.032	0.161
Behaviours	0.257	0.257	0.011	0.118	0.257	0.118	0	0.118
Governance	0.118	0.057	0.257	0.118	0.257	0.118	0.016	0.204
All Dimensions	0.285	0.082	0.107	0.115	0.112	0.045	0.075	0.210

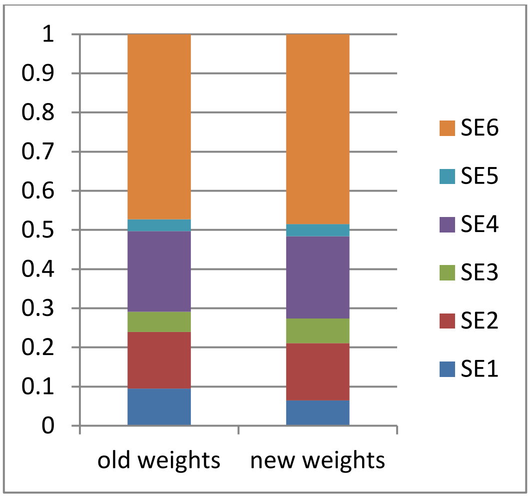

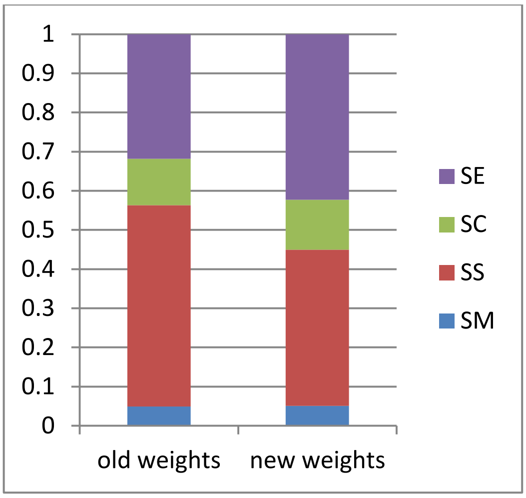

Figure 6.5 Expert DH (Structural Engineering)

Figure 6.6 Expert DH (Structural Dimension)

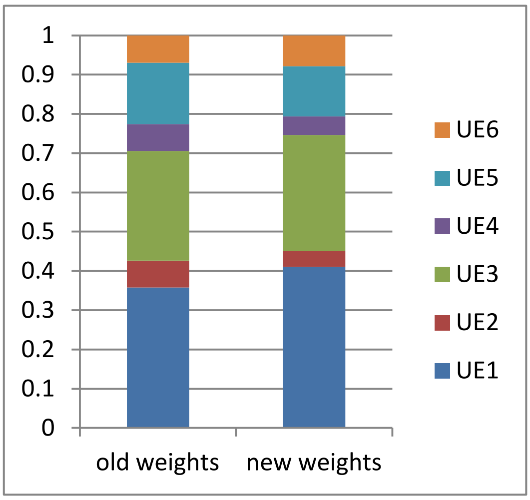

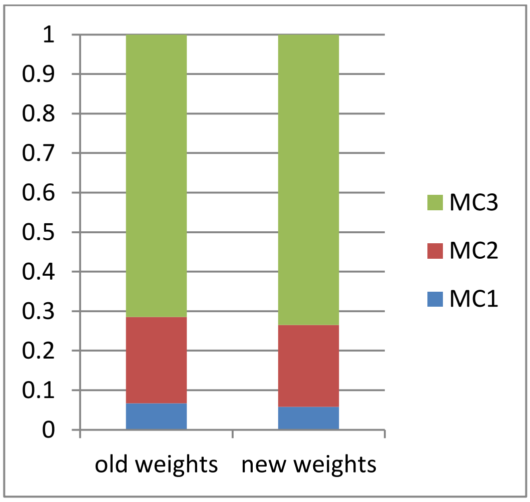

Figure 6.7 Expert DH (Ease of Application)

Figure 6.8 Expert DH (Curation)

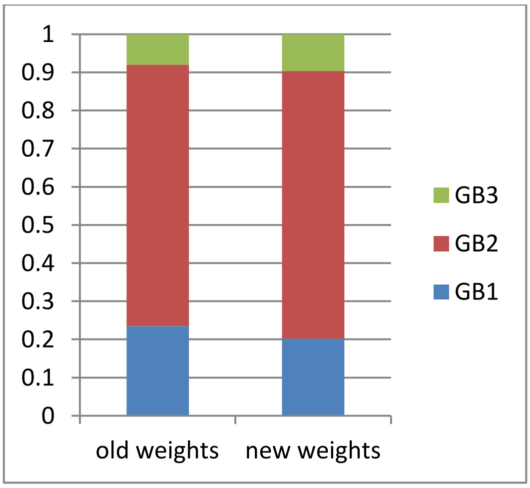

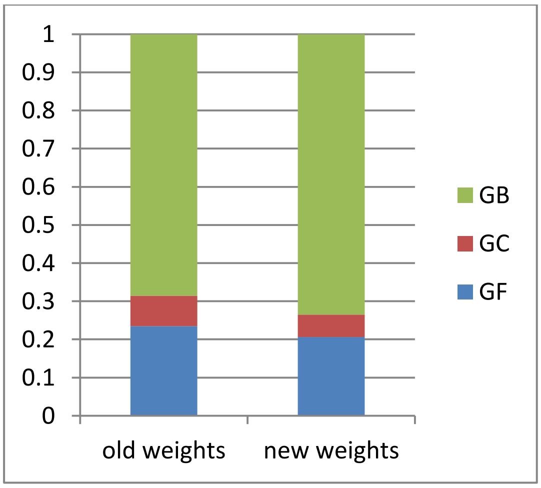

Figure 6.9 Expert DH (Behaviours)

Figure 6.10 Expert DH (Governance)

Changing matrix values to achieve consistency was performed using a structured approach and a script created in 'R' (see script in Appendix 25). The logic of the script was as follows: Assume a matrix "A" with elements A_{ij}; the value of the first matrix element (A_{11}) was changed (up or down, using values between 1 and 9); for each change to A_{11} a new consistency ratio (CR) was calculated; the values of A_{11} that produced the best three CRs were stored along with the new consistency ratios; the first matrix element (A_{11}) was then set back to its original value; the process repeated again for the second matrix element (A_{12}); again the best three CRs and the values that generated them were stored, i.e., the best three out of all calculations for matrix elements one (A_{11}) and two (A_{12}); the value of the second element (A_{12}) was set back to its original value and the script moves on to element three (A_{13}) and so on until all elements in the matrix have been processed. Only the best three CRs are kept at any one time by the script and the best three options are reported at the end of the script's execution, along with the matrix values that generated them. Figure 6.11 shows the 'R' script output using expert 'WD's' data matrices. Note only inconsistent matrices were eventually changed. No attempt was made to improve the consistency of matrices with a consistency value of < 0.14. The script used matrices values for all sub-category comparisons for ease of data structuring. Suggested new values for already consistent matrices were ignored.

Appendix 26 lists all of the plots, generated for each expert, showing all accepted changes to expert ratings data presented in terms of the effect these changes had on the calculated weights. It should be noted that when provided with suggested amendments, expert 'LL' made his own modifications instead of using those generated via the 'R' program and the author worked with this expert until all of his data fell within an acceptable level of consistency. All other experts accepted the changes suggested. The codes used in the plots are those shown previously in Table 6.15.

As is evidenced from inspecting the plots presented in Appendix 26, the relatively high number of changes required, indicates that all experts had some level of difficulty in keeping their answers sufficiently consistent. Of significance, most experts (7) had difficulty with comparing the criteria in the sub-category – "Structural Transparency" and four experts also had problems with assessing the criteria in the "Structural Engineering" sub-category. Not surprisingly then, seven experts also had difficulty with the pair-wise comparison of these sub-categories within the "Structural Dimension". Fifty percent of experts also had trouble being consistent when comparing criteria in the "Ease Of Application" sub-category.

After adjusting their ratings, expert matrix data were consistent to varying degrees, but all consistency ratios were subsequently less than 0.14 (see Table 6.26).

```
R is free software and comes with ABSOLUTELY NO WARRANTY.
You are welcome to redistribute it under certain conditions.
Type 'license()' or 'licence()' for distribution details.

  Natural language support but running in an English locale

R is a collaborative project with many contributors.
Type 'contributors()' for more information and
'citation()' on how to cite R or R packages in publications.

Type 'demo()' for some demos, 'help()' for on-line help, or
'help.start()' for an HTML browser interface to help.
Type 'q()' to quit R.

[Previously saved workspace restored]

> source('C:/Users/kim_fin/Documents/DocumentspostNovember2009/ahp.txt')
Data file: C:/Users/kim_fin/Documents/PhD/PhD_AHP_Weights_Survey1/
AHP_Analysis/WDforRanalysis.csv
Structural Transparency
 current CR of 0.2728
 CR of 0.1132 by setting element [1,5] to value 1 (current value 5)
 CR of 0.2180 by setting element [5,2] to value 1 (current value 5)
 CR of 0.2280 by setting element [4,2] to value 1 (current value 3)
Structural Engineering
 current CR of 0.0972
 CR of 0.0637 by setting element [1,6] to value 1 (current value 3)
 CR of 0.0655 by setting element [1,4] to value 3 (current value 2)
 CR of 0.0796 by setting element [6,4] to value 1 (current value 3)
Structural Dimension
 current CR of 0.1147
 CR of 0.0910 by setting element [3,1] to value 9 (current value 7)
 CR of 0.0584 by setting element [2,1] to value 2 (current value 7)
 CR of 0.0573 by setting element [4,2] to value 1 (current value 3)
Application Relevance
 current CR of 0.0161
 CR of 0.0015 by setting element [4,1] to value 2 (current value 3)
 CR of 0.0121 by setting element [2,1] to value 7 (current value 5)
 CR of 0.0121 by setting element [3,1] to value 7 (current value 5)
Functional Relevance
 current CR of NaN
Ease of Application
 current CR of 0.1014
 CR of 0.0734 by setting element [6,4] to value 9 (current value 3)
 CR of 0.0794 by setting element [2,1] to value 1 (current value 3)
 CR of 0.0832 by setting element [6,5] to value 1 (current value 3)
Sustainability
 current CR of 0.1169
 CR of 0.0032 by setting element [2,3] to value 2 (current value 5)
 CR of 0.0250 by setting element [2,1] to value 9 (current value 5)
Usability
 current CR of -Inf
Curation
 current CR of 0.0061
 CR of -0.0001 by setting element [3,1] to value 9 (current value 7)
 CR of 0.0023 by setting element [3,2] to value 2 (current value 3)
User Assistance
 current CR of NaN
Maintenance
 current CR of -Inf
Framework
 current CR of 0.1168
 CR of -0.0001 by setting element [1,2] to value 1 (current value 3)
 CR of -0.0001 by setting element [2,3] to value 1 (current value 3)
Community
 current CR of 0.1232
 CR of 0.0436 by setting element [2,4] to value 1 (current value 3)
 CR of 0.0436 by setting element [1,2] to value 1 (current value 3)
 CR of 0.0853 by setting element [3,4] to value 7 (current value 3)
Behaviours
 current CR of 0.1168
 CR of -0.0001 by setting element [2,1] to value 1 (current value 3)
 CR of -0.0001 by setting element [2,3] to value 9 (current value 3)
Governance
 current CR of 0.1168
 CR of -0.0001 by setting element [3,2] to value 1 (current value 3)
 CR of -0.0001 by setting element [1,2] to value 9 (current value 3)
All Dimensions
 current CR of 0.1113
 CR of 0.0780 by setting element [2,3] to value 1 (current value 3)
 CR of 0.0780 by setting element [5,1] to value 1 (current value 3)
 CR of 0.0780 by setting element [3,5] to value 1 (current value 3)
>
```

Figure 6.11 Sample 'R' script screen output for WD's matrix comparison data

Table 6.26 Summary of Consistency Ratios (< 0.14)

Consistency Ratio	RA	LL	MH	WD	DH	CA	KS	JH
Structural Transparency	0.1209	0.1008	0.1186	0.1132	0.1174	0.1035	0.0731	0.1338
Structural Engineering	0.1071	0.1251	0.1106	0.0972	0.1124	0.1208	0.0772	0.1277
Structural Dimension	0.1056	0.1179	0.0910	0.1147	0.1044	0.1312	0.1282	0.1312
Application Relevance	0.1372	0.0694	0.0053	0.0161	0.0845	0.0936	0.0231	0.0982
Functional Relevance	0	0	0	0	0	0	0	0
Ease of Application	0.1255	0.0963	0.1211	0.1014	0.118	0.1086	0.0210	0.1179
Sustainability	0.1169	0.0250	0	0.1169	0.1168	0.1169	0.0931	0.0462
Usability	0	0	0	0	0	0	0	0
Curation	0.0109	0.0109	0.0061	0.0061	0.1009	0.034	0.0001	0.0032
User Assistance	0	0	0	0	0	0	0	0
Maintenance	0	0	0	0	0	0	0	0
Framework	0.1169	0.0250	0.0250	0.1168	0.0332	0.0332	0	0.0332
Community	0.0834	0.1012	0.0053	0.1232	0.0733	0.0978	0.0308	0.0502
Behaviours	0.0061	0.0001	0.0109	0.1168	0.1169	0.1169	0	0.1168
Governance	0.1171	0.0560	0.0250	0.1168	0.1009	0.1169	0.0157	0.0692
All Dimensions	0.1349	0.0793	0.1043	0.113	0.1089	0.044	0.0727	0.0955

6.3.3 Analyses of Group Results

The primary purpose of the pair-wise comparison exercise was to gain a better understanding of the level of importance (or weight) that experts placed on certain ontology evaluation criteria within a decision-making process. At this point there are individual expert results that address RQ1.2.2. - *Is it feasible to derive a weighted evaluation criteria model in which criteria are rated according to importance ? If so, are evaluation criteria of equal weight or do some carry more importance than others.* However, a group result was preferred.

Because it was judged that the expert group would not sustain a Delphi approach, in order to improve consistency of responses or arrive at a group result, various normalised measures of central tendency were explored in the aggregated data (i.e., arithmetic mean (Abramowitz and Stegun, 1972), the geometric mean (Abramowitz and Stegun, 1972), median and Perth Measure (Kauko, 2002) as possible representative values for a group result. The formulae used for each measure are as follows:

Assume the sample space $\{x_1,....x_n\}$

(i) Arithmetic mean: $\frac{1}{n}\sum_{i=1}^{n}(x_i)$

(ii) Geometric mean: $n\sqrt{x_1 x_2 .. x_n}$

(iii) Median: ((n + 1)/2)th item in a sorted sample space

(iv) Perth Measure: $\frac{SUM((\text{Min } value \;\; of \;\; x)+(Max \;\; value \;\; of \;\; x)+ (Median \; *4))}{6}$

These formulae were applied in a Microsoft™Excel spreadsheet and grouped (expert) data for each sub-category and dimension were plotted for comparative purposes against the individual expert data (see Appendix 27). A sample plot for the 'Structural Engineering' sub-category is depicted in Figure 6.12. In this plot the label "Gmean" is the geometric mean; "NGmean" is the normalised geometric mean and "AdjMedian" is the normalised median. Expert names, dimensions, sub-categories and criteria are coded as in previous plots.

Because there is a high degree of variance in the group data none of the computed measures of central tendency (just mentioned) are considered ideally representative. The normalised geometric mean is probably the best of the central tendency measures using a visual inspection of the data. The geometric mean is often used in AHP to obtain an aggregated result (Ramanathan and Ganesh, 1994; Van Den Honert and Lootsma, 1996). The normalised geometric mean weights were taken as being indicative of a group outcome and these results were transcribed into the three tiers of the hierarchical model (see Figure 6.13). It can be seen from this Figure that three of the dimensions ('Functional Relevance', 'Structure' and 'Usability' are relatively similarly weighted (i.e., with weights of 29.5%; 28% and 21.8% respectively) and would appear to contribute a roughly equal amount each to the decision-making process in ontology selection tasks. The 'Maintenance' and 'Governance' dimensions, however, for this group of experts, would appear to contribute substantially less to the decision-making process (i.e., with weights of 11.5% and 9.2% respectively).

To test the actual level of concordance (or agreement) amongst the different experts across the full spectrum of pair-wise comparisons, Kendall's co-efficient of concordance was used at the 5 percent level of significance (0.05). If the test statistic *W* is equal to '1', then this means that all the experts have been unanimous, and each expert has assigned the same rating to the criteria in each sub-category and dimension. If *W* is equal to '0', then there is no overall trend of agreement among the experts, and their responses may be regarded as essentially random (Siegel and Castellan*,* 1988). Intermediate values of *W* indicate a greater or lesser degree of unanimity among the various responses. Kendall's co-efficient was calculated using the following, where: criterion *i* (in a sub-category, or a dimension) has a weight allocated to it by an expert *j*. This weight allocated by expert *j*, is subsequently converted to a rank $r_{i,j}$ (relative to other weights allocated by other experts for criterion *i*), where there are in total *n* criteria and *m* experts. Then the total rank R_i given to criterion *i*'s weight (across all experts) is:

$$R_i = \sum_{j=1}^{m} r_{i,j},$$

and the mean value of these total ranks is:

$$\bar{R} = \frac{1}{2}m(n+1).$$

The sum of squared deviations, *S*, is defined as:

$$S = \sum_{i=1}^{n}(R_i - \bar{R})^2,$$

and then Kendall's *W* (Kendall and Gibbons, 1990) is defined as:

$$W = \frac{12S}{m^2(n^3 - n)}.$$

Because $n > 20$ it is suggested (Legendre, 2010) that Chi-square ($\chi 2$) is then used to test for significance. Friedman's $\chi 2$ statistic is obtained from *W* using:

$$\chi 2 = m(n-1)W$$

This quantity is asymptotically distributed like chi-square with (n−1) degrees of freedom (*df*). The null hypothesis for Kendall's co-efficient of conformance is:

H_0: The expert rankings are not concordant (i.e., they disagree).

A web-based software program (Chang, 2012) was used to compute the co-efficient. First the program computed a table of ranks from the input weight data (i.e., all calculated ranks for weights across all categories for all experts) and then it used these ranks to calculate W and $\chi 2$. Since the computed value of $\chi 2$ (191.9855) was greater than the table value (Thompson, 1941) of $\chi 2$ (0.05; 56 *df*) = 43.19, H_0 was rejected, i.e., there is concordance amongst experts.

The computed value for W = 0.4285 and according to Schmidt (1997) that signifies a weak to moderate level of agreement. Schmidt (1997, Table 6) indicates confidence in the ranks should be 'fair' since the calculated co-efficient is greater than '0.3000' but less than '0.5000'.

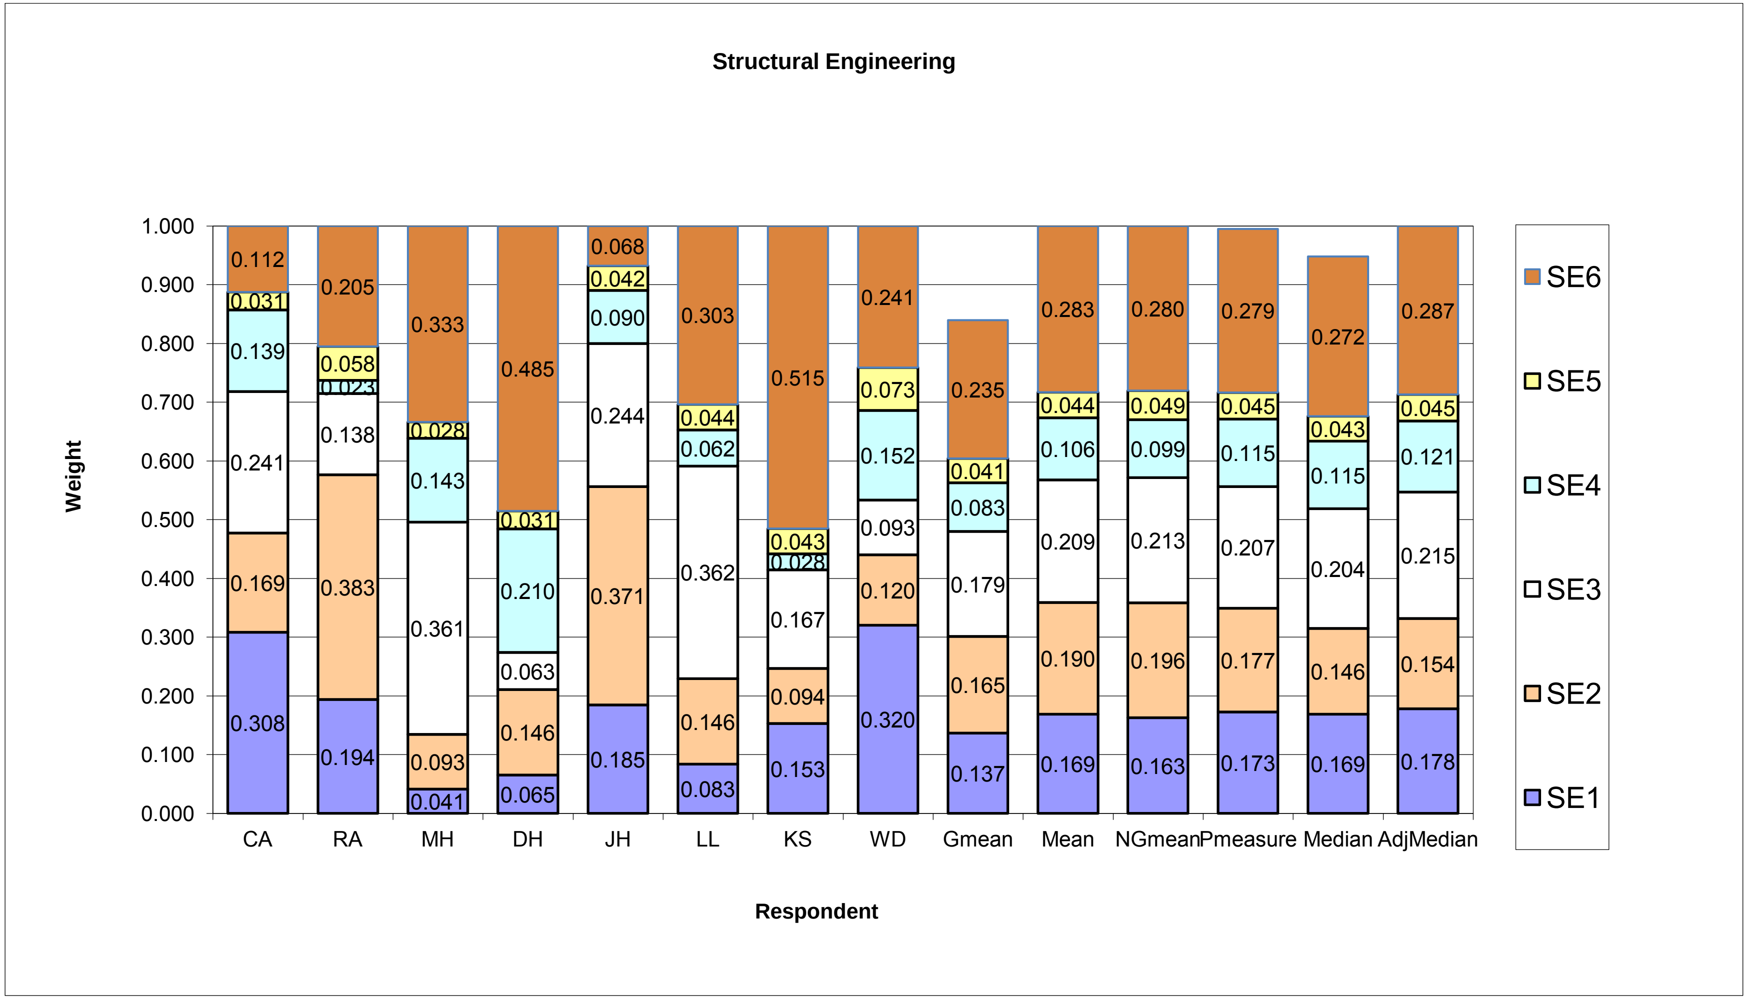

Figure 6.12 Structural Engineering Weight Data

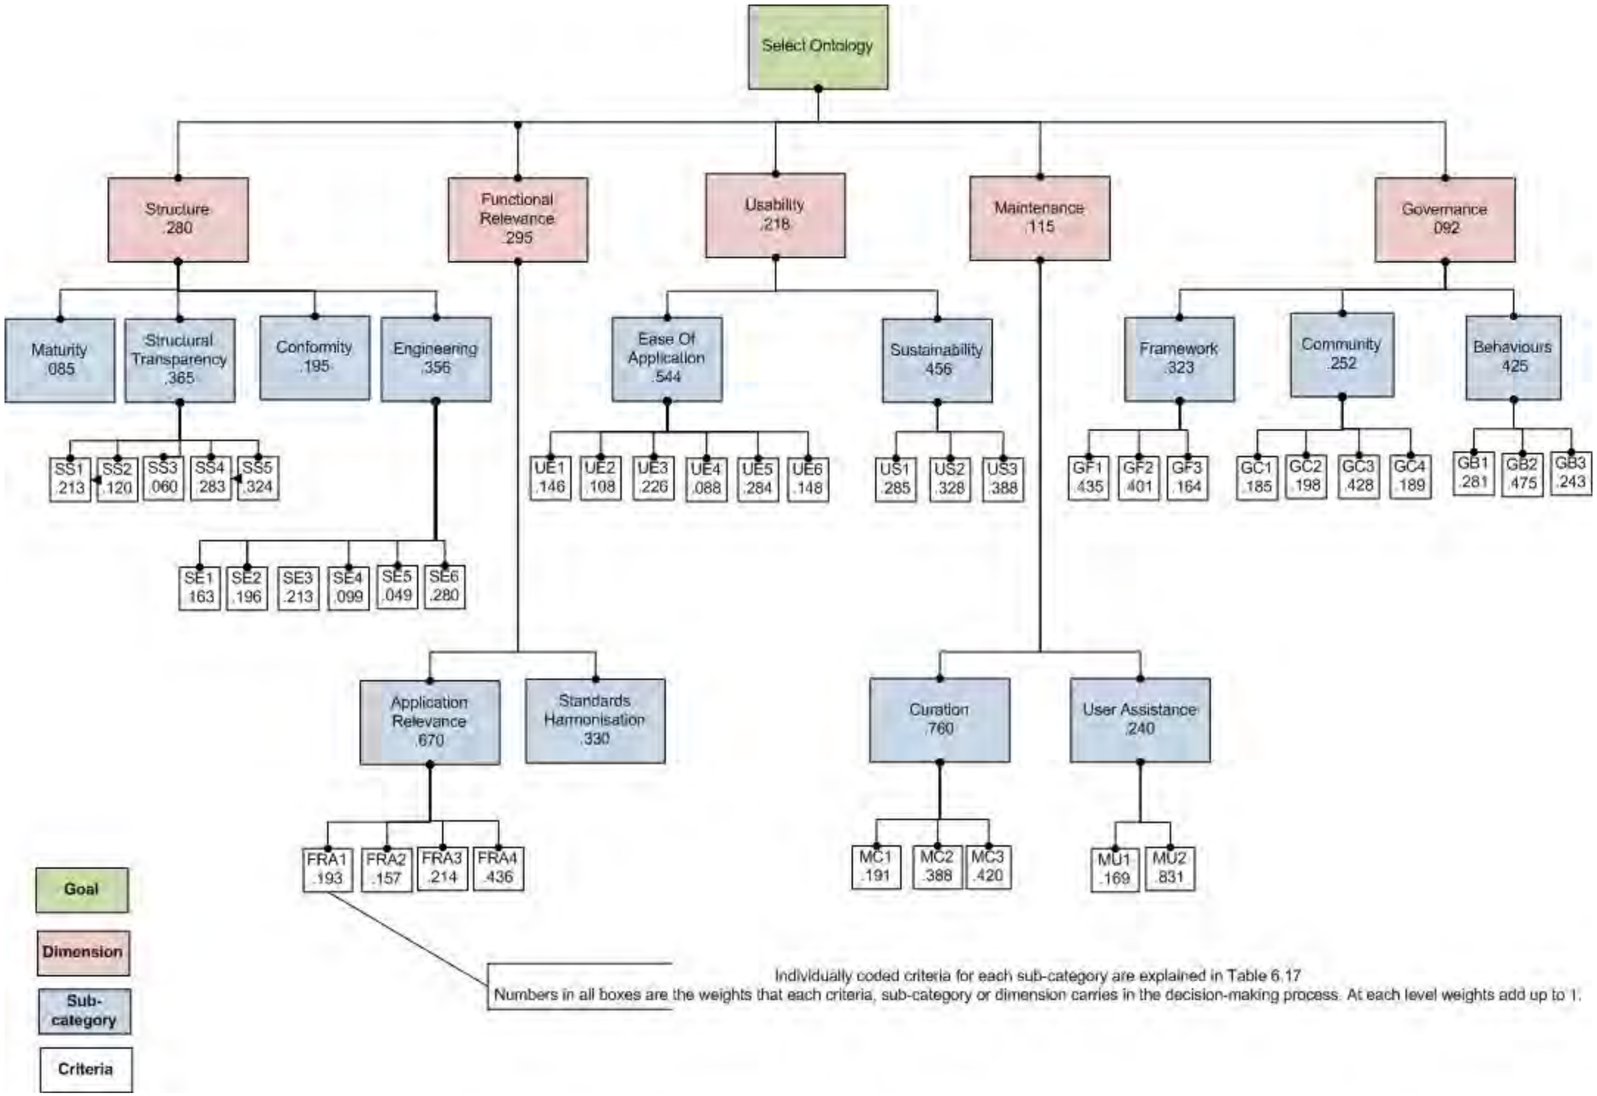

Figure 6.13 Hierarchical Weight Model (Using Normalised Geometric Mean Weight Values)

Whilst Kendall's W tells us something about the level of agreement between experts in their ranking of weights, it doesn't indicate anything about possible patterns or clusters in the data. To examine whether there were any similarities between experts in how they rated criterion, sub-categories and dimensions, multi-dimensional scaling (MDS) was used. This technique has previously been used by Chen *et al.* (2008) and Cox (2009).

A custom-made 'R' script (Appendix 28) was developed that took as input a matrix consisting of rows of experts and columns of weights (for all criteria, sub-categories and dimensions). Figure 6.14 shows the resultant MDS Plot.

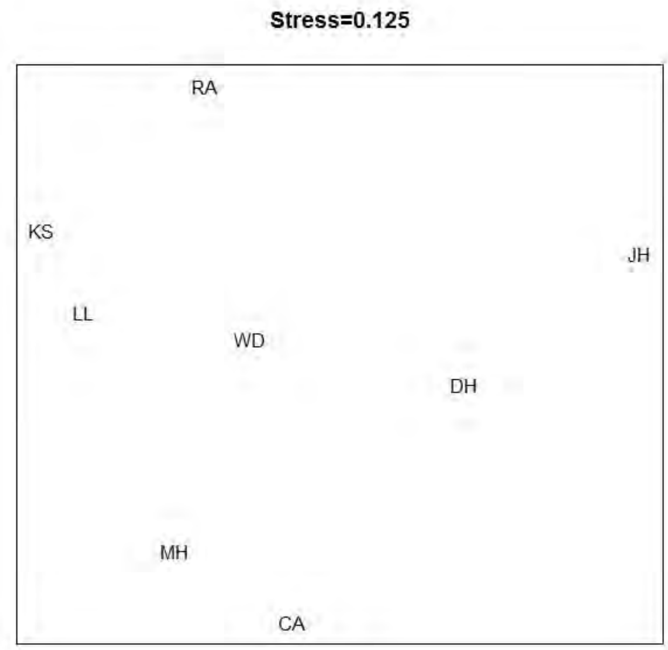

Figure 6.14 MDS Scatter Plot of Expert Weight Data

This plot tends to show dissimilarity between all experts (with 'KS' and 'LL' having more similarity than either has with anyone else, and 'MH' and 'CA' having more similarity with each other than they do with anyone else, but the similarities are weak). So, there are no obvious clusters of experts when the whole dataset is considered. In this example of MDS, experts' weights were rescaled to create two axes (a 2D configuration) such that the ordering of the individual pair-wise weights is preserved and data trends are enhanced. This was achieved by iteratively calculating a Stress function until this function reached a minimum value. The resulting configuration (Figure 6.14) represents the optimised clustering /trending of the data consistent with the original ordering of the data. A Stress value of '0' indicates that the data have been perfectly rescaled in two dimensions. Higher values approaching '1' indicate a failure to rescale in 2D (Steyvers, 2002). The Stress value measures the difference between a particular distance and its corresponding pseudo-distance ($d_{ij} - \widehat{d_{ij}}$) and indicates the “degree” of fit.

Raw Stress = $\sum (d_{ij} - \hat{d}_{ij})^2$

d_{ij} = distance between variables *i* and *j* in the configuration

$\hat{d}_{ij}$ = those values which minimize the stress, subject to the constraint that the d_{ij}'s have the same rank order as the input data

Raw Stress is generally normalized (between ‘0’ and ‘1’) so that it is possible to compare different configurations by making stress independent of the size or scale of the configuration.

The two most commonly used normalizing factors are:

Sum of squared distances: $\sum d_{ij}^2$ and

the sum of squared differences between the distances and their average:

$$\sum_i (d_{ij} - \bar{d})^2$$

Thus we have two measures of normalized stress:

Stress SQDIST $S_1 = \left[\sum (d_{ij} - \hat{d}_{ij})^2 / \sum d_{ij}^2 \right]^{\frac{1}{2}}$ and *Stress SQDEV* $S_2 = \left[\sum (d_i - \hat{d}_{ij})^2 / \sum (d_{ij} - \bar{d})^2 \right]^{\frac{1}{2}}$

where $\bar{d}$ = mean of all $d_{ij's}$

For 2D configurations, Stress (SQDIST) values between 0.10 and 0.20 (Kruskal and Wish, 1978) are considered to provide ‘fair ’rescaling of the resulting configuration. The computed Stress value (SQDIST) for the MDS plot in Figure 6.14 was 12.5% (or 0.125).

6.3.4 Weighted Criteria Model Data – Investigating Patterns In Expert Ratings

Given that there are no obvious clusters of experts with respect to how evaluation criteria (sub-categories and dimensions) were weighted, it is worthwhile forensically examining disaggregated weight data in terms of the potential expert stratification levels canvassed earlier in sections 6.1 and in 6.2. Recall that RQ1.2.2.2 asks “*do relative levels of* (ontology evaluation criteria) *importance differ by application between scientific disciplines, or by any other discernible factor*)”? Whilst it has been possible to derive a weighted evaluation criteria model, all analyses thus far demonstrates a relatively high degree of variation amongst experts as to the weights applied within the model. Here the group result is no longer of interest, but rather the variation in responses provided by each expert. From analyses in previous sections possible useful stratifications included the clustering of experts according to their:

(i) Experience in an ISO/OGC standards environment.

(ii) Discipline.

(iii) Community governance type.

(iv) Experience with formal vs light-weight ontologies.

(v) Ontology application area.

(vi) An expert's skill type.

(vii) Team based experience vs sole operator.

(viii) An expert's roles.

(IX) Overall level of expertise.

Table 6.27 summarises the expert clusters for each of (i) to (viii) above. It was not possible to objectively divide experts on the basis of (IX) because the proxies used previously, to gauge expertise, were considered too crude and of those experts ranked highest ('PF', 'SC' and 'LL'), only 'LL' participated in the pair-wise survey. Experts highlighted in yellow are those who took part in the pair-wise comparison survey and for whom pair-wise comparison data are available for analysis.

Table 6.27 Stratification of Experts

Stratification Type	Stratification Level	Expert Code
ISO/OGC	Experience	JH, LL, DH, CA, RA, SC
	No Experience	RH, WD, PF, PM, JG, RL, MH, KS
Discipline	Hydrology	SC, RA,LL,
	Biology	SC,LL,RH,CA,WD,DH,MH,PM,KS
	Oceanography	SC,PF,JG,RL
	Atmosphere/Meteorology	SC, RA,PF,RL
	Geoscience	SC, RA,PF,JH,
	Energy	LL
Governance Type	None	WD,PM,CA
	Informal Governance	RH,WD,PF,LL,JG,RA,RL,SC,MH, KS
	Formal Governance	JH,DH,RA,RL,SC,MH
	Experts spanning Informal/Formal	RA,RL,SC,MH

Ontology Type	Light-weight	DH,PM,CA, JH+,RA+,SC+
	Formal	WD,PF,LL,JG,MH,RH+,PM+ ,RL, KS
Expert Skill Type*	DESO	RH,WD,JH,DH, JG, RL,SC, KS
	SOESDE	LL, RA,
	SOENDE	CA
	SOEDE	PF, PM, MH
Team Experience	Sole operator (or two people)	PM,RH,RL,CA,JH
	Teamwork	PF, SC,RA,MH,WD,PM,DH,LL, KS
Ontology Application Area	Domain (including Domain Task)	RH, WD, JH, LL, SC, MH, RL, JG, DH, RA, PM, KS
	Application-specific	PF, CA
Expert Roles	Leader/Driver/Initiator	PF, DH, LL, JG, RL, RA, RH, SC, MH
	Maintainer	PF, DH, LL, JG, RL, RA, RH, SC, MH, JH, PM, CA, WD, KS,
	Specialist Advisor	PF, DH, LL, JG, RL, RA, RH, SC, MH, JH, PM, CA, WD, KS

* Codes are those used in Table 6.9.

+ placed in these stratification levels because of the encoding languages they listed in Table 6.8

For each dimension, i.e., the most abstract level of comparison within the devised hierarchical evaluation criteria weighted model, plots were produced showing the spread of expert data with respect to the normalised geometric mean (presented in each plot as a histogram). The normalised geometric mean is included to provide a central measure of tendency as a group bench-mark. The various levels of stratification in Table 6.27 were then consulted to assess whether there was any visual evidence to correlate patterns in expert responses with the various types of stratification listed.

From Figure 6.15 it is demonstrated that there is fairly good agreement between 'RA', 'MH', 'DH', 'LL' and 'KS' about the relative importance of the sub-categories in the "Structural" dimension. 'JH' and 'WD', although different to the previous group, are never-the-less similar to each other in terms of the level of importance they place on the various criteria of this dimension. 'CA' appears more of an outlier. 'CA' and 'JH' have both had less team development experience than the rest of the experts, so this was a possible differentiator. However, because 'WD's' data was similar in proportion to 'JH's' data and 'WD' have worked in team environments, "Team Experience" vs "Little Team Experience" is not a supportable rationale for the detected patterns. 'CA' may also differ from the rest of the experts given that his experience lies in developing "application-specific ontologies", rather than "domain ontologies" and it is possible that as a result, different weights might apply in developing these types of ontologies. But because 'CA' is the only expert who supplied weight data and who falls into this class (i.e., develops application-specific ontologies), it is almost impossible to attribute the difference to this factor alone. Had 'PF' (who is another member of the application-specific ontology development cluster) participated in the weighting exercise it might have been possible to infer more from the results.

'RA', "MH', 'LL' and 'DH' had all played the 'role' type – "Leader/Driver/Initiator", which may have been something this cluster of experts had in common, but the inclusion of 'KS' who had not played this role, discounted this stratification as being something which could be investigated further as the rationale for the pattern detected.

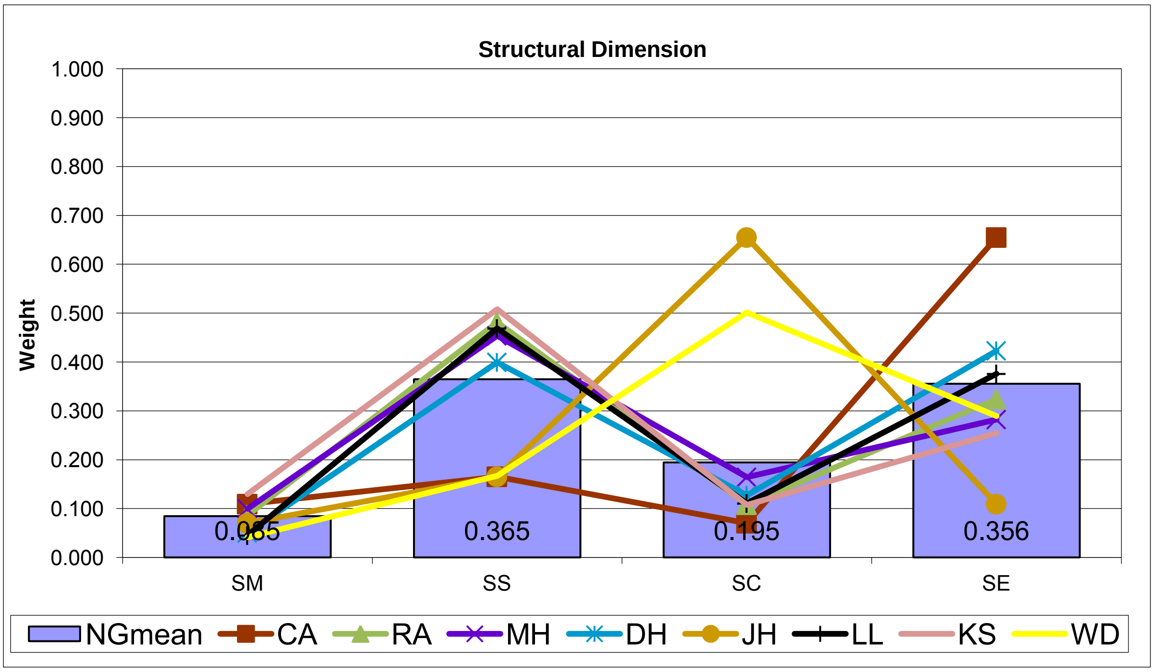

Figure 6.15 Expert Weight Data (Structural Dimension)

In terms of the "Functional Relevance" dimension (Figure 6.16) the data shows two very distinct, but opposite patterns. 'LL', 'KS', 'WD', 'RA', 'CA' and 'MH' fall into one group and 'DH' and 'JH' fall into another. There is no obvious reason from the stratification table (6.27) to explain these divergent clusters. It should be noted that 'JH' and 'DH' are particularly invested in standards setting groups and activities, which could have influenced their weight allocations. Importantly, however, six out of the eight experts had very similar views about the proportional importance of the two sub-categories within this dimension.

"Usability" (see Figure 6.17) was another polarising dimension, with two main groups, 'RA', 'LL' and 'KS' rating "Sustainability" factors as much more important than "Ease of Application" criteria. The remainder of experts disagreed and had the opposite perspective. No obvious stratification issues appeared to explain the pattern of separation within this dimension.

The "Maintenance" dimension (Figure 6.18) has a very significant level of agreement between all experts on the relative importance of "Curation" criteria over "User Assistance" criteria. Only 'MH' was the outlier.

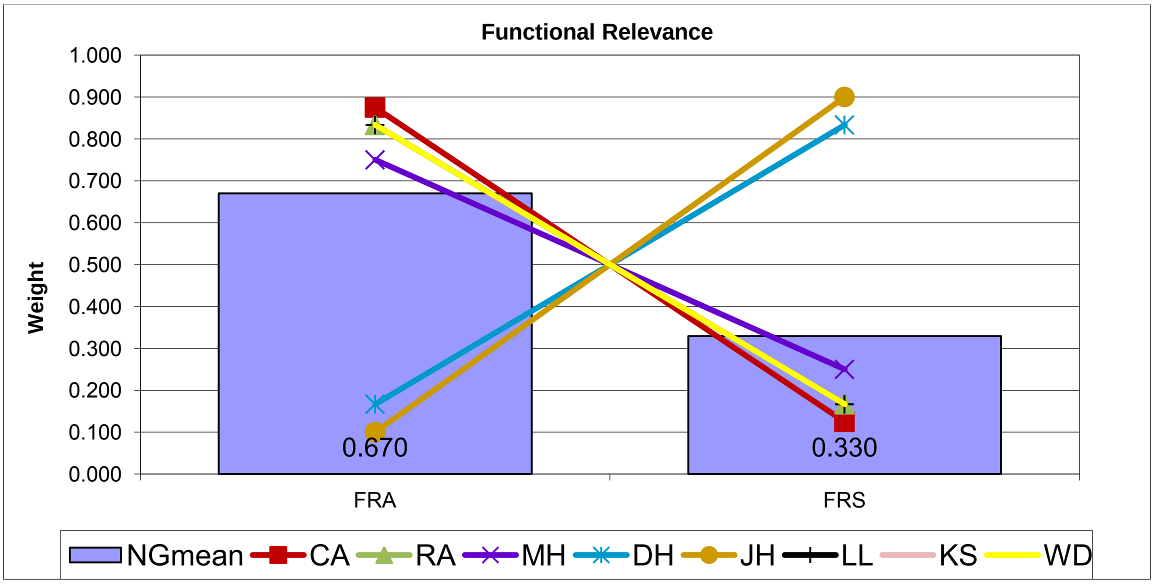

Figure 6.16 Expert Weight Data (Functional Relevance Dimension)

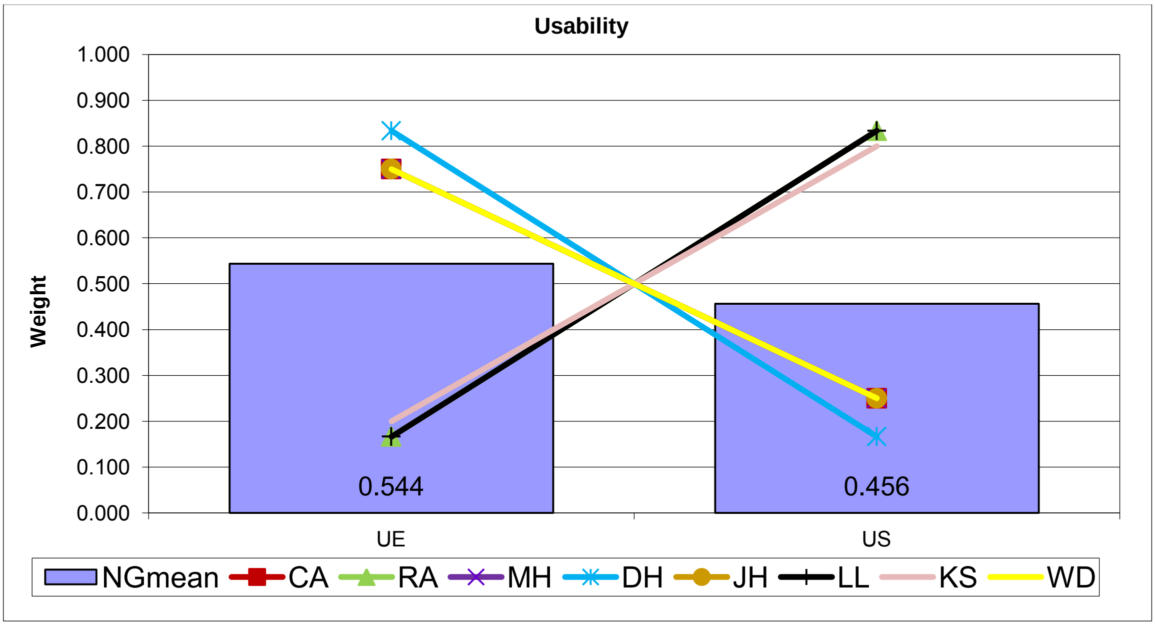

Figure 6.17 Expert Weight Data (Usability Dimension)

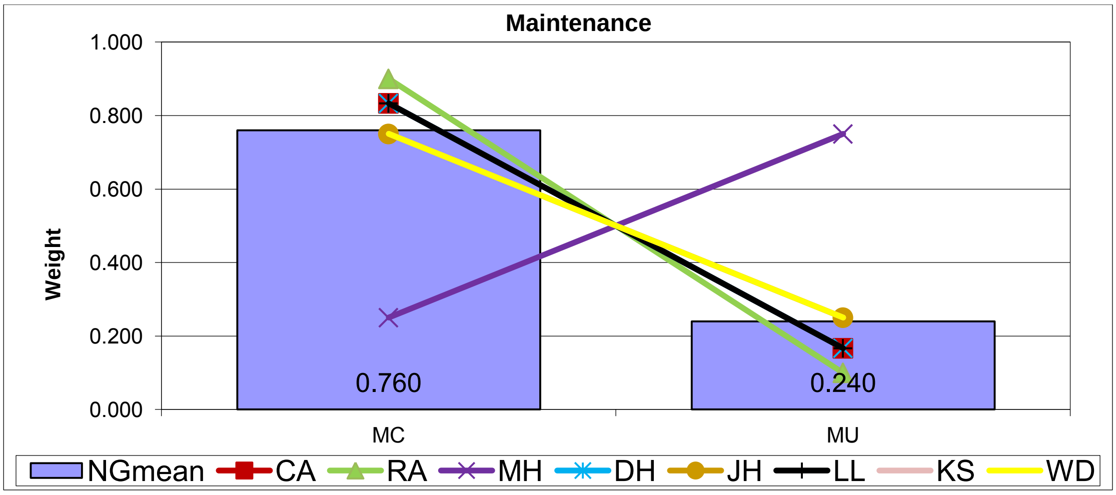

Figure 6.18 Expert Weight Data (Maintenance Dimension)

The patterns for the "Governance" dimension (Figure 6.19) were more complicated than for "Maintenance". There was strong agreement between 'RA' and 'JH' and to a lesser extent 'WD' that "Framework" factors were more important than "Community" and then "Behavioural" criteria, respectively. 'MH', 'LL' and 'KS' were also in relatively good agreement except that they had the order of preference reversed with "Behavioural" factors being most important. 'DH' and 'CA' were both outliers. The "Governance Type" stratification levels (from Table 6.27) were consulted and examined with respect to the distribution of experts in this plot to see if any patterns were reflected in the data. Even with the removal of the outliers, no rationale for the split was found as a result of referencing the stratification levels.

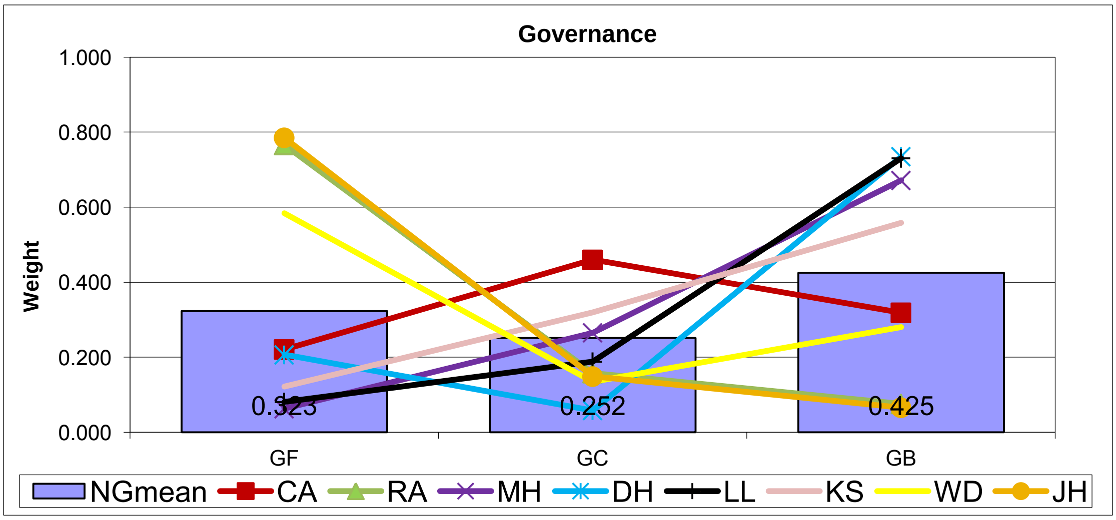

Figure 6.19 Expert Weight Data (Governance Dimension)

The overall picture for the "All" dimensions category (Figure 6.20) is relatively confused. There is some similarity of view between 'LL' and 'KS' and between 'DH' and 'JH' (although these latter couple differ in view to the former two). The rest of the data is highly mixed.

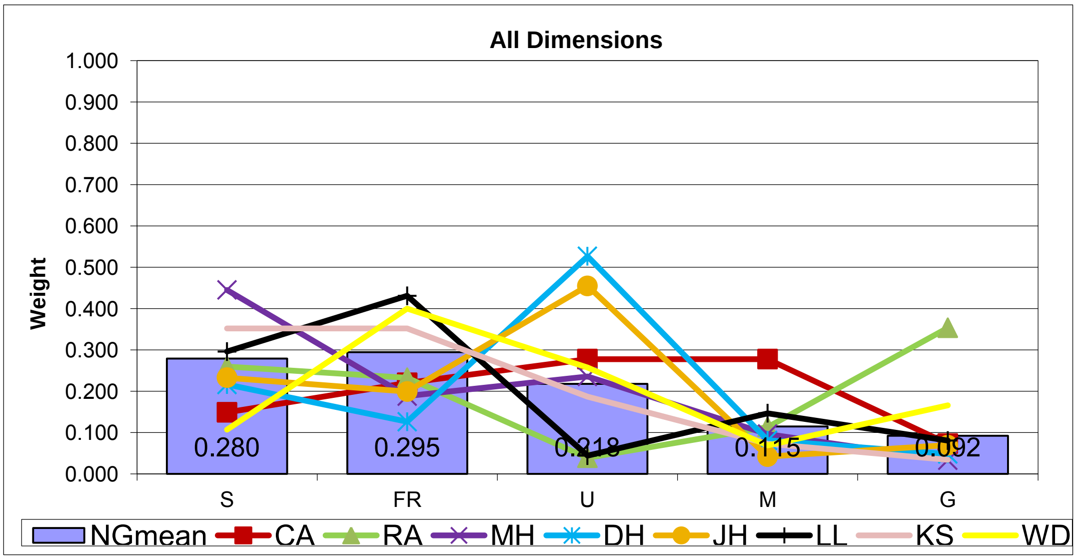

Figure 6.20 Expert Weight Data (All Dimensions)

Having assessed the data at the 'dimension' level of the hierarchical model, the 'sub-category' levels were then explored. As for the dimension level, at the sub-category tier there was also a clear lack of any apparent correlation between the spread of expert opinion and the stratification issues, as listed in Table 6.27. However, given that some patterns (i.e., clusters of experts) were apparent for the weights applied within the different sub-categories, the remaining data (below sub-categories) were examined, not only for evidence of stratification but also to see if the groupings of expert opinion that were evident at the dimension level, were also present within the more detailed levels of the data (i.e., between criteria weights within sub-categories).

Only those plots that demonstrated a continuation of a pattern already detected, or which depict other trends worthy of mention, are shown and discussed.

Within the "Structural Transparency" sub-category the data appears very mixed. The only readily discernible trend is that five out of the eight experts ranked criterion 'SS3' (i.e., Whether the ontology includes sample instance data) as having considerably lower importance than the other factors. Two of the three outliers, 'JH' and 'CA' agreed on placing 'SS5' (Whether concepts are well defined both in terms of textual description and logic) lower than 'SS3'. 'KS' placed 'SS2' above 'SS3'.

The weights applied in the "Structural Engineering" sub-category were also quite mixed. Significantly all experts agreed that 'SE5' (Does the author have credibility in developing ontologies and/or within the domain ?) was of least importance of the six criterion presented. There was some similarity of scoring by 'MH' and 'LL'. 'RA' and 'JH' showed some agreement (though different to 'MH' and 'LL'). Neither grouping reflected the patterns found earlier within the "Structural" dimension.

'DH' and 'JH' were relatively in synchronisation regarding weight allocation within the "Application Relevance" sub-category (Figure 6.21). It should be noted that these experts were also in alignment with respect to the weights they applied to the two sub-categories in the "Functional Relevance" dimension. Recall also that they both held a completely contrary opinion to the bulk of the other experts in ranking the "Functional Relevance" sub-categories. In this particular sub-category (as shown in Figure 6.21), six out of the eight experts did, however, consider 'FRA4' (How focussed is the ontology for its stated purpose ?) the most important criterion. The two outliers were 'KS' and 'WD' on this issue.

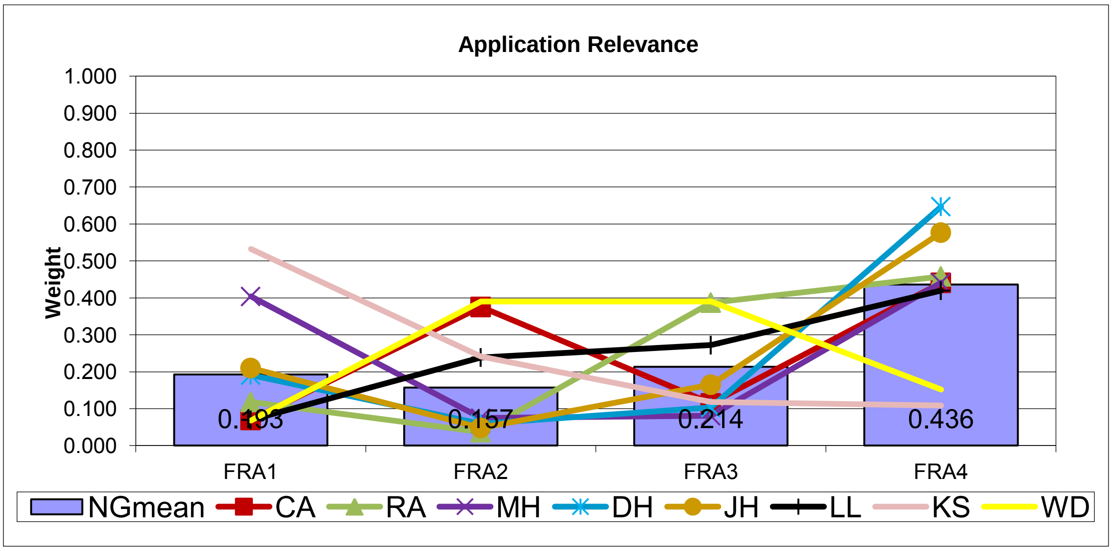

Figure 6.21 Expert Weight Data (Application Relevance)

Expert opinion was relatively confused with respect to the importance placed on the various criteria within the "Ease of Application" sub-category, so the plot has not been presented here.

The "Sustainability" sub-category data exhibited two groups that appeared to be in coherence (see Figure 6.22). 'LL', 'KS' and 'WD' rated US2 > US3 >US1 in very similar proportions, whereas 'DH' and 'JH' rated US3 > US1 > US2 (i.e., in a different order, but using similar proportions). The rest of the expert results were more mixed. 'LL' and 'KS' were also part of a group of three (including 'RA') who rated "Ease of Application" as more important than "Sustainability" in the "Usability" dimension comparisons.

The “Curation” sub-category comparison data showed some grouping of opinion. ‘WD’, ‘DH’ and ‘MH’ all rate MC3 > MC2 > MC1 in similar proportions. ‘RA’ and ‘LL’ were in close agreement and followed the ordination of the previous group but have ‘MC3’ and ‘MC2’ with almost equivalent weights.

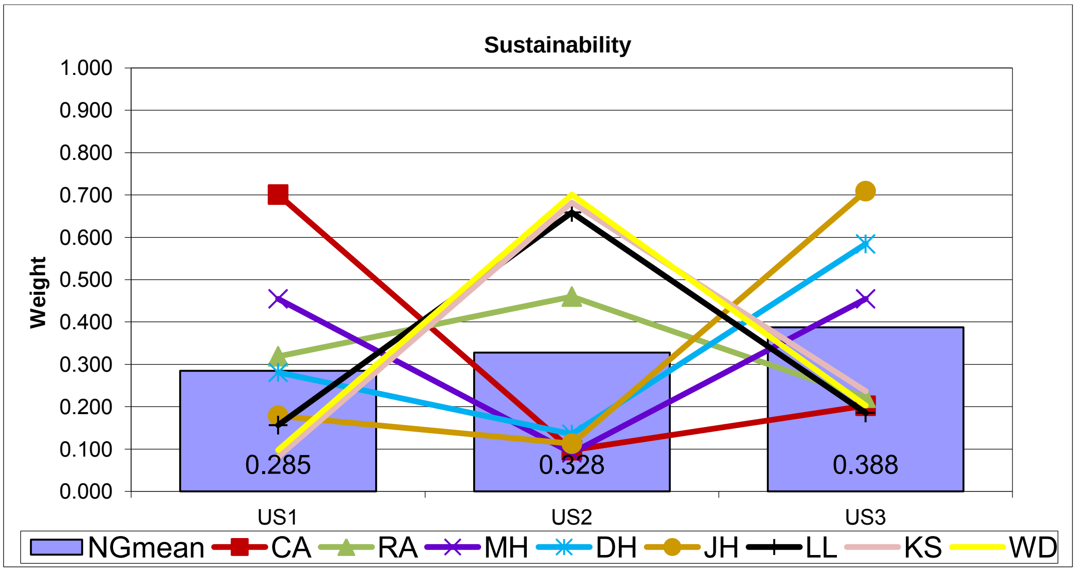

Figure 6.22 Expert Weight Data (Sustainability)

Figure 6.23 has been included only to show the rare occurrence in which all eight experts agreed on the ordination and the proportion of importance placed on the two criteria in the “User Assistance” sub-category. All considered ‘MU2’ (Quality and availability of published documentation) to be more important than ‘MU1’ (Is there any type of help-desk associated with the ontology ?).

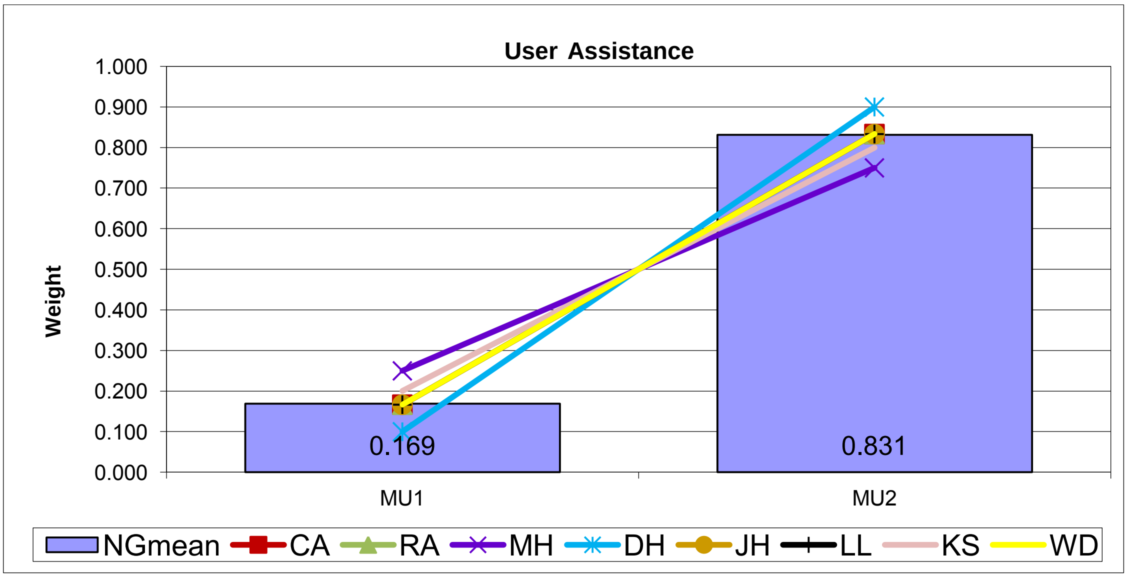

Figure 6.23 Expert Weight Data (User Assistance)

The “Framework” sub-category data (Figure 6.24) fell into two groups, with two experts considered possible outliers (i.e., ’KS’ and ‘CA’). ‘DH’, ‘JH’, ‘WD’ and ‘RA’ all ranked GF1 > GF2 >GF3. ‘LL’ and ‘MH’ rated GF2 > GF1 > GF3. Most experts (bar ‘CA’) considered ‘GF3’ (Does the community review or moderate individual ontology developments) to be of least importance.

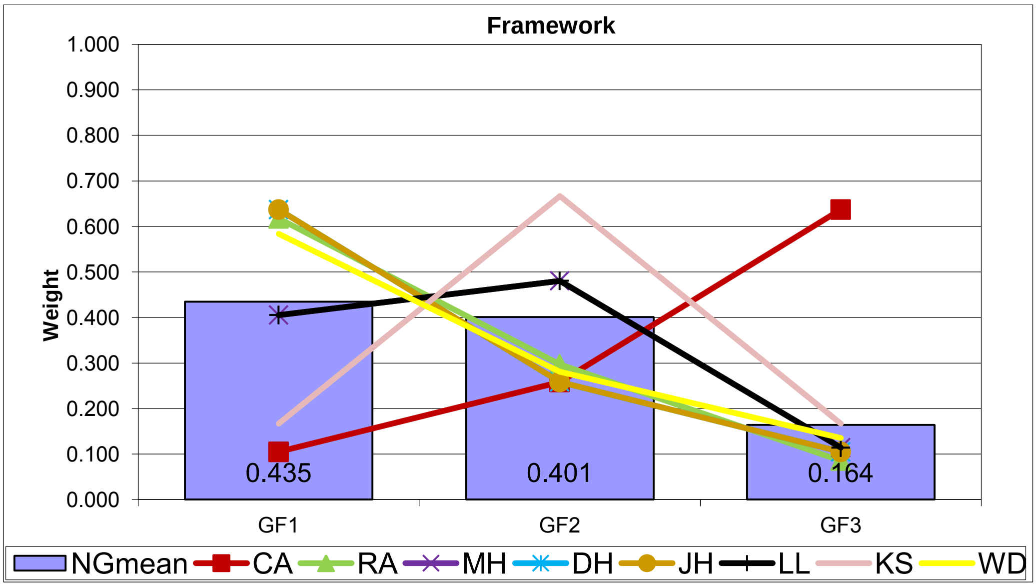

Figure 6.24 Expert Weight Data (Framework)

There is an obvious cluster of experts (‘LL’, ‘CA’, ‘WD’ and ‘DH’) that have similar ratings within the “Community” sub-category (see Figure 6.25). ‘MH’ follows the ordination rankings but places a heavier emphasis on criterion ‘GC1’ (How mature the community is in terms of ontology development and its longevity and cohesion as a community of practise) than the previous four experts. ‘JH’ and ‘KS’ could be considered outliers.

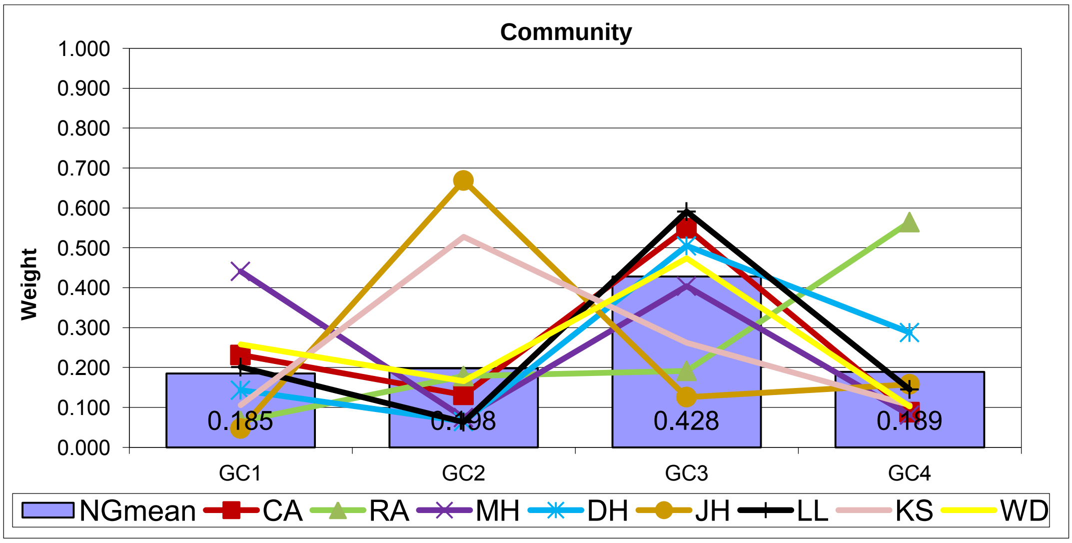

Figure 6.25 Expert Weight Data (Community)

Six out of the eight experts (Figure 6.26) rated 'GB2' (Whether contribution to ongoing development and maintenance is encouraged, open and facilitated) as the most important criteria in the "Behaviours" sub-category. Five of the experts ('DH', 'JH', 'LL', 'WD' and 'MH') had reasonably similar proportions in relation to 'GB2', but differed in their proportions for 'GB1' and 'GB3'.

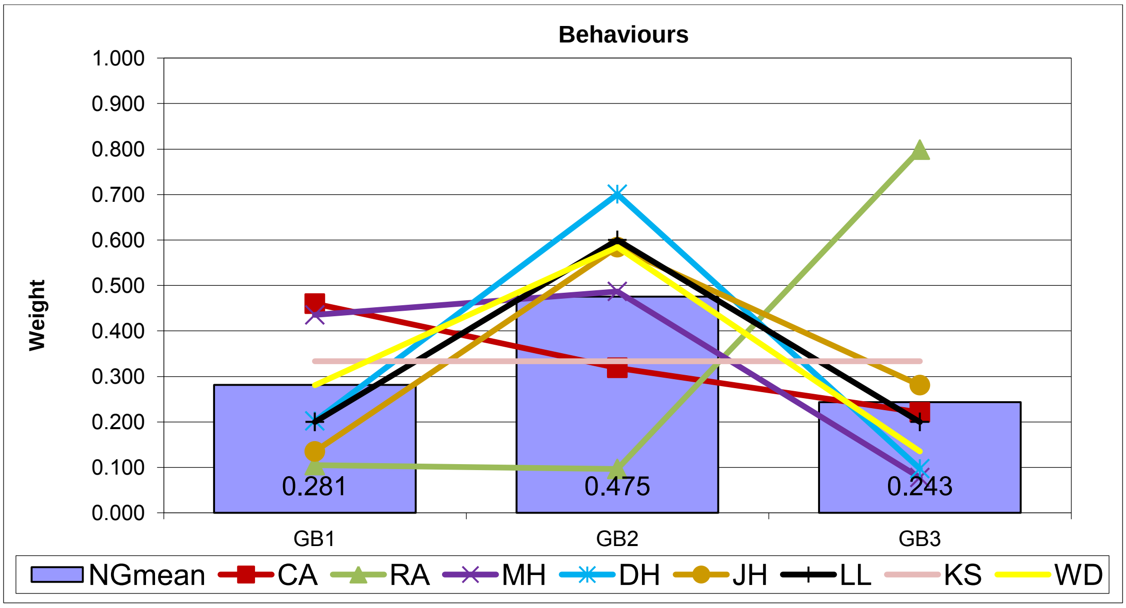

Figure 6.26 Expert Weight Data (Behaviours)

All of the previous analyses failed to find any explainable patterns in the way in which experts weighted evaluation criteria with respect to stratification performed on expert data derived from the screening survey and from interviews. In answering RQ1.2.2.2 it was therefore concluded that it wasn't possible to detect differences in expert weightings that could be attributed to membership of a particular scientific discipline, or any other discernible factor (using the data to hand).

6.3.5 A Comparison Of Criteria Importance Between Interviews and Pair-wise Survey

In addition to examining the data from the perspective of stratification, since some experts qualified their statements about evaluation criteria when they raised them during interview, it was possible to use this interview-based information to cross-check results available from the pair-wise comparison data. Table 6.28 is a slightly different view of Table 6.21 showing the individual experts who qualified a particular criterion, rather than just showing the total number of experts who qualified a criterion as being "important" or "less important". The first observation is that the data in Table 6.28 gives some comfort that the sub-sample of the expert cohort who participated in the pair-wise comparison survey, in the main, were experts covering both viewpoints on particular issues. If all pair-wise comparison survey participants were from a sub-group with similar views the sample could be viewed as relatively biased.

For each of the criteria listed in Table 6.28, the actual weights allocated by each expert were compared to the view expressed by an expert during interview about the relative importance of the criteria. Quite surprisingly, of the 35 weight values checked, nineteen (i.e., 54%) of the weights eventually allocated by experts during the pair-wise comparison exercise contradicted the general opinions stated at interview (i.e., a criterion was ranked in an opposite manner within its sub-category relative to the level of importance stated during interview).

On this basis alone it was concluded that there was no obvious evidence that experts (who later went on to participate in the pair-wise comparison survey) stacked the interview with their "pet" criteria and then voted them up in the ratings exercise (something considered possible by the author). The volatility of the swings in expert view is interesting in that when confronted with a "real" comparison (i.e., with a choice between alternative options) and where experts are asked to quantify relative importance, their initial informal ratings at interview were quite different from the views expressed later during the pair-wise comparison exercise.

Table 6.28 Evaluation Criteria Qualified During Interview By Expert

Criteria Code	Interviewed Experts Considering Criteria Important	Interviewed Experts Considering Criteria Less Important
UE1	RH, PF, DH, CA, MH*	WD
FRA2	JH, LL, PM, CA, JG, RL, MH	DH, WD, PF
UE4	JH, PF, JG	DH
SE4		DH, PM, SC, PF
SE1	DH, WD, CA,MH, RH, PF	JH
SS3	LL, JH, PF, JG	DH
UE2	JH, PF	PM
US1	PM, WD, MH, SC, RH, JH, RA, PF, JG	DH, CA,
MC3	DH, PM, WD, MH,PF, JG	SC
SS5	LL, MH, RA, RL	JH

*Experts who participated in the pair-wise comparison survey

Yet another way of cross-checking information provided by experts at interview and during the pair-wise comparative weighting exercise, is to inspect the ‘frequency’ with which ‘criteria’ (represented as concept codes) were mentioned at interview, and their eventual relative ‘rating’ by experts.

The evaluation criteria assembled in the hierarchical model (Table 6.15) represent those criteria mentioned by experts as being those used within communities for selecting ontologies for re-use purposes. Any criteria which were mentioned by experts were included in the model regardless of how frequently they were cited (across experts). However, during interviews a few concepts were mentioned by the majority of experts (i.e., ten out of the thirteen experts interviewed - refer to Table 6.19 - 20) and therefore it could be inferred that these might be some of the more universally important criteria. The most frequently cited concepts in Table 6.19 and 6.20 map to the following evaluation criteria (in Table 6.15):

FRA2: The degree of concept coverage

SE2: Whether the ontology is sufficiently modular that it is easily able to be reused.

US1: The size and scope of the current user implementation base.

GF1: Does the ontology development community have a published governance framework (e.g., Clear roles and responsibilities for participants, guiding principles, review processes, repositories, community portal or wiki).

FRA1: Whether the ontology meets the use-case goals or competency questions.

As will be evident from examining the data in the weighted hierarchical model diagrams presented next (in Figures 6.27 to 6.29), none of the concepts listed above were rated particularly highly during the pair-wise comparison exercise. ‘FRA1’ for example, ranks third out of a total of four siblings within the ‘Application Relevance’ sub-category; ‘GF1’ ranks second out of three siblings in the ‘Governance’ dimension; ‘SE2’ ranks third out of five siblings in the ‘Engineering’ sub-category and ‘US1’ is third out of three siblings in the “sustainability’ sub-category.

Whilst ‘frequency of mention’ is clearly not the same measure as ‘relative level of importance’ it was still surprising that the most frequently mentioned concepts, were not subsequently rated comparatively high in the level of importance exercise. Also ‘FRA4’, the most highly rated criterion in Figure 6.29, only drew mention by two experts during interview. It would appear that being potentially more universally applicable cannot be equated with level of importance, particularly given the most highly rated criterion was one of the least mentioned concepts (i.e., ‘em:focussed’ in Table 6.19).

6.4 Hierarchical Model Revision and Refinement

The previous section explained the survey methods used to capture criteria importance (or weight) data from experts and described the analytical techniques applied to evaluate those data (pertaining to individual experts and for the group as a whole). A weighted hierarchical model was derived as a result. Section 6.3 also exposed the reader to a wide variety of forensic data analyses that were conducted in order to seek explanations for any patterns that may have been evident between expert responses. One motivation for attempting to detect explainable patterns being that a clearer understanding of an expert's propensity, or inclination, to give a particular comparative rating would assist in the derivation of a better and more refined group representation of the model (or lead perhaps to the development of more than one model, tailoring each for a particular context). As it stands, the author was unable to link any evident level of expert stratification with the ratings provided when there were either similar or divergent views on criteria ratings.

Whilst there is no apparent (explainable) stratification, the (expert) sample size is small (although still considered representative) and there is considerable noise in the data which makes representing a group result, regarding the importance factors, problematic. It is openly acknowledged that a larger pool of experts, with more coherence in the data would add more credence to the modelled group result.

It is not possible in hind-cast to expand the expert cohort, but it is possible using the previous analyses to reduce some of the noise in the data so that (group) mean values better map rating opinions where there is apparent convergence amongst most experts and only a few experts are unexplainably divergent. This phase of the research (step 7 in Figure 6.1) is what is described next.

6.4.1 Expert Feedback and Hierarchical Model Revisions

All interviewed experts (plus 'KS') were provided with a copy of the "group" weights for each dimension, sub-category and criteria. The pair-wise comparison survey participants additionally received over-plots of all individual (expert) weights for each dimension and sub-category. All of these latter experts were asked to comment on how they felt about the adequacy of the normalised geometric mean as a representation of the group result (given that they had access to their own results in the context of those from other experts). The main question put to this sub-group of experts was:

> *"Do you believe that the normalised geometric mean values (selected to represent the 'group' result) are a good enough approximation (cognisant of the spread of expert*

values). If you don't, could you elaborate as to why (indicating how you would intuitively change either ranks, and/or actual histogram proportions in any graph, possibly in light of the weights given by others)"?

This group were also asked whether they had any comments to make on the model structure and if they could see value in applying the model in their community contexts.

The experts who did not participate in the pair-wise comparison survey were more simply asked:

"Would you intuitively change any of the weights (proportionately) and/or the rankings that indicate relative levels of criteria importance" ?

This sub-set of experts were also asked the same questions regarding model structure and community application, as those who participated in the pair-wise comparison survey.

Of the experts that replied (6 of 14), only one ('CA') was really comfortable with the mean as a good representation. Three experts ('RA', 'LL' and 'JH') had detected similar groupings to the author in a number of supplied plots and speculated (without the benefit of the expert stratification data) about why this might be evident. Two experts ('LL' and 'JH') were comfortable with a mean as a group result, but not if there were "explainable" patterns in the data that would be obscured by such a measure. 'RA', whilst also urging decomposition of the data to find useful patterns, had an interesting viewpoint where he stated "Given the failure of the common approaches I'd start with a qualitative judgement about the outliers, as the potential valuable insights ☺". He also made suggestions about why there might be divergences (which have already been assessed through the previous stratification exercise). These expert views supported the desirability of trying other approaches that might deliver a more representative group result.

After an exhaustive inspection of the data and any available supporting information, and contrary to the advice of 'RA', where there was a relatively definite "group trend" in the data, responses considered to be "outliers" were removed and the group mean re-calculated. This approach was particularly targeted at data in the middle and lower portions of the model (i.e., the weights applied to the sub-categories and criteria). A "trend" as defined in this context was demonstrated if five or more experts (out of 8) had quite similar views (in terms of both ordination and weight proportions). Because of the heterogeneity of the data at the "All dimension" level (with no obvious trend) it was reasoned that the existing group approximation should stand that gives roughly equal weight to three of the dimensions and a lesser, but also roughly equal weight to the other two dimensions.

Table 6.29 lists the sub-categories and/or the criteria that have had expert outliers removed. Figure 6.27 shows the revised model, with re-calculated normalised geometric mean weights after outlier removal. Figure 6.28 shows these weights as global percentages (i.e., the percent contribution that each level in the hierarchy makes to the decision making goal).

Table 6.29 Outliers Removed

Hierarchical Model Element Code	Experts Removed
Structural Dimension (S)	WD, JH, CA
Functional Relevance (FR)	JH, DH
Usability (U)	RA, LL, KS
Maintenance (M)	MH
Governance (G)	DH and CA
Behaviours (GB)	RA, CA, KS

The main changes in the model weights by removing outliers were:

(a) A slightly heavier emphasis on "Structural Transparency" criteria (0.468) than "Engineering" criteria (0.331) in the "Structural" dimension, where previously they had very similar values (0.365 and 0.356, respectively).

(b) A lower degree of importance now placed on "Standards Harmonisation" (0.171 down from 0.330) in the "Functional Relevance" dimension.

(c) More emphasis on "Ease of Application" (0.769 up from 0.544) in the "Usability" dimension.

(d) A slight flattening out of the values in the three "Governance" sub-categories by bringing down the value of the "Behaviours" sub-category (from 0.425 to 0.381).

(e) A relatively large increase in importance placed on 'GB2' (Whether contribution to ongoing development and maintenance is encouraged, open and facilitated) from a value of 0.475 to 0.611, at the expense mainly of the weight placed on 'GB3' (Whether the governance and the core group of developers engender trust and credibility).

In terms of expert feedback regarding the model structure three of the six experts who participated in the pair-wise comparison exercise (i.e., 'RA', 'WD', 'CA') responded positively about the value inherent in the model itself and felt that the model "covered" the issues well (e.g., 'WD' said "I think the structure of the model covers the relevant categories I would think important in ontology re-use").

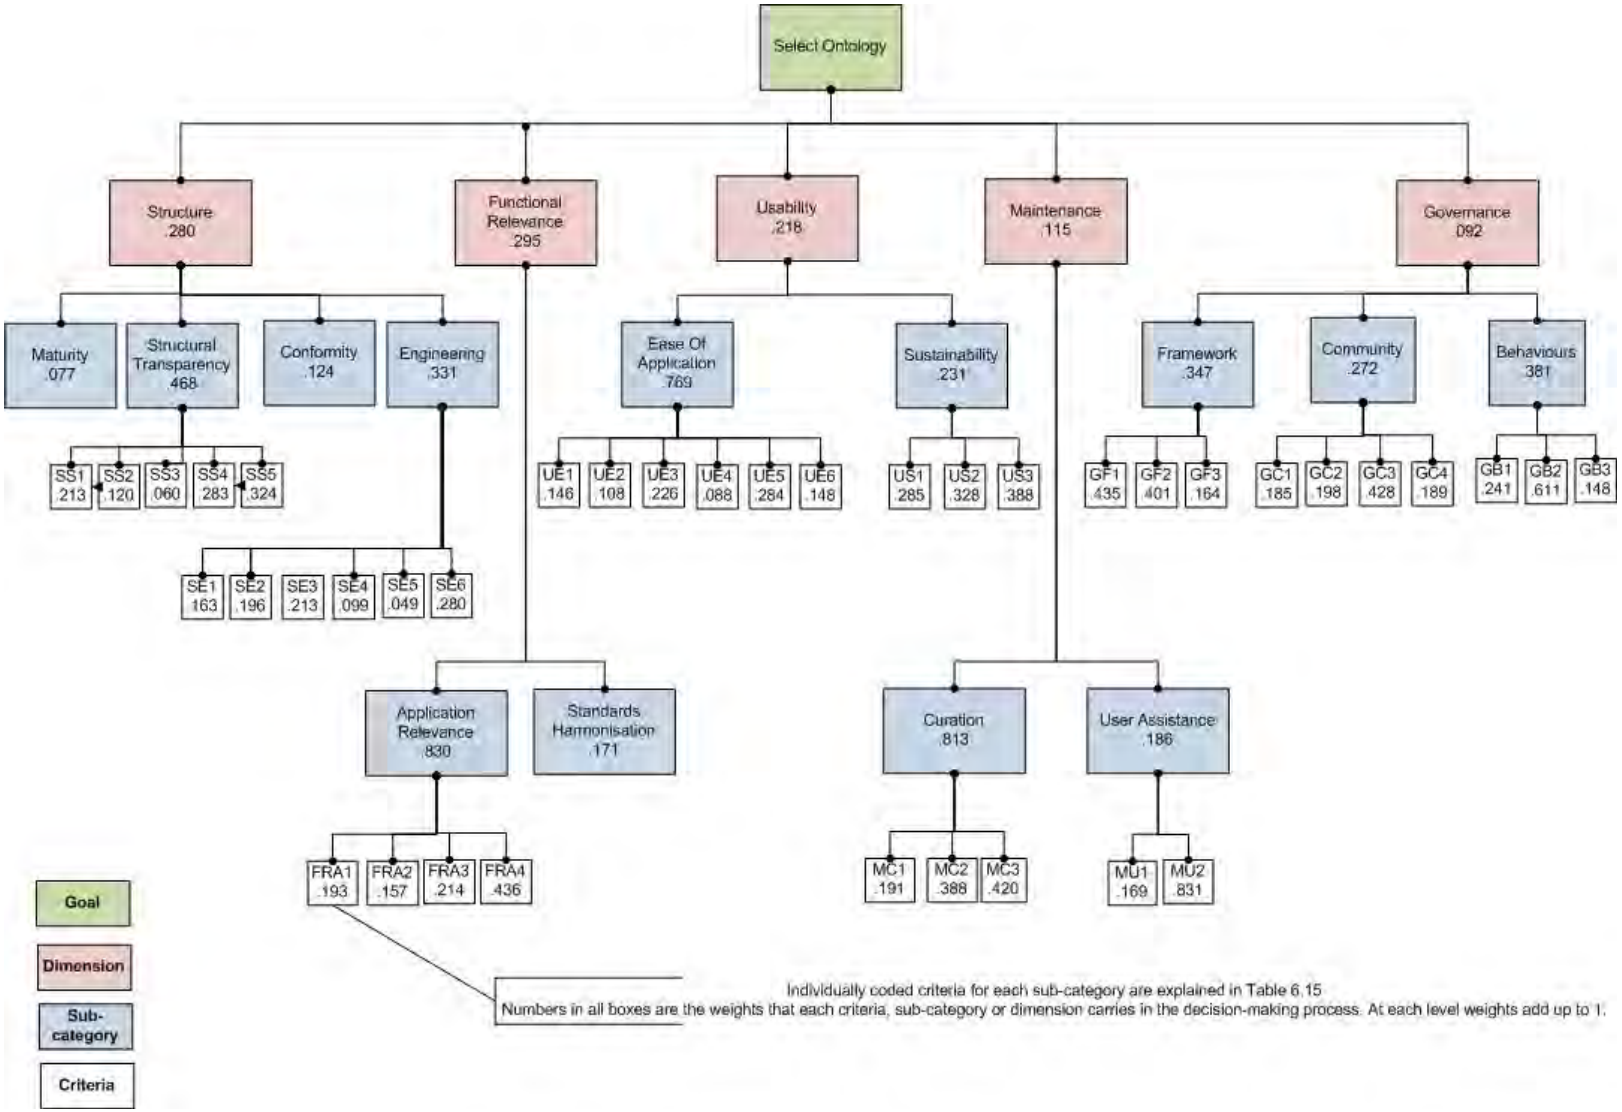

Figure 6.27 Revised Hierarchical Model (Weights)

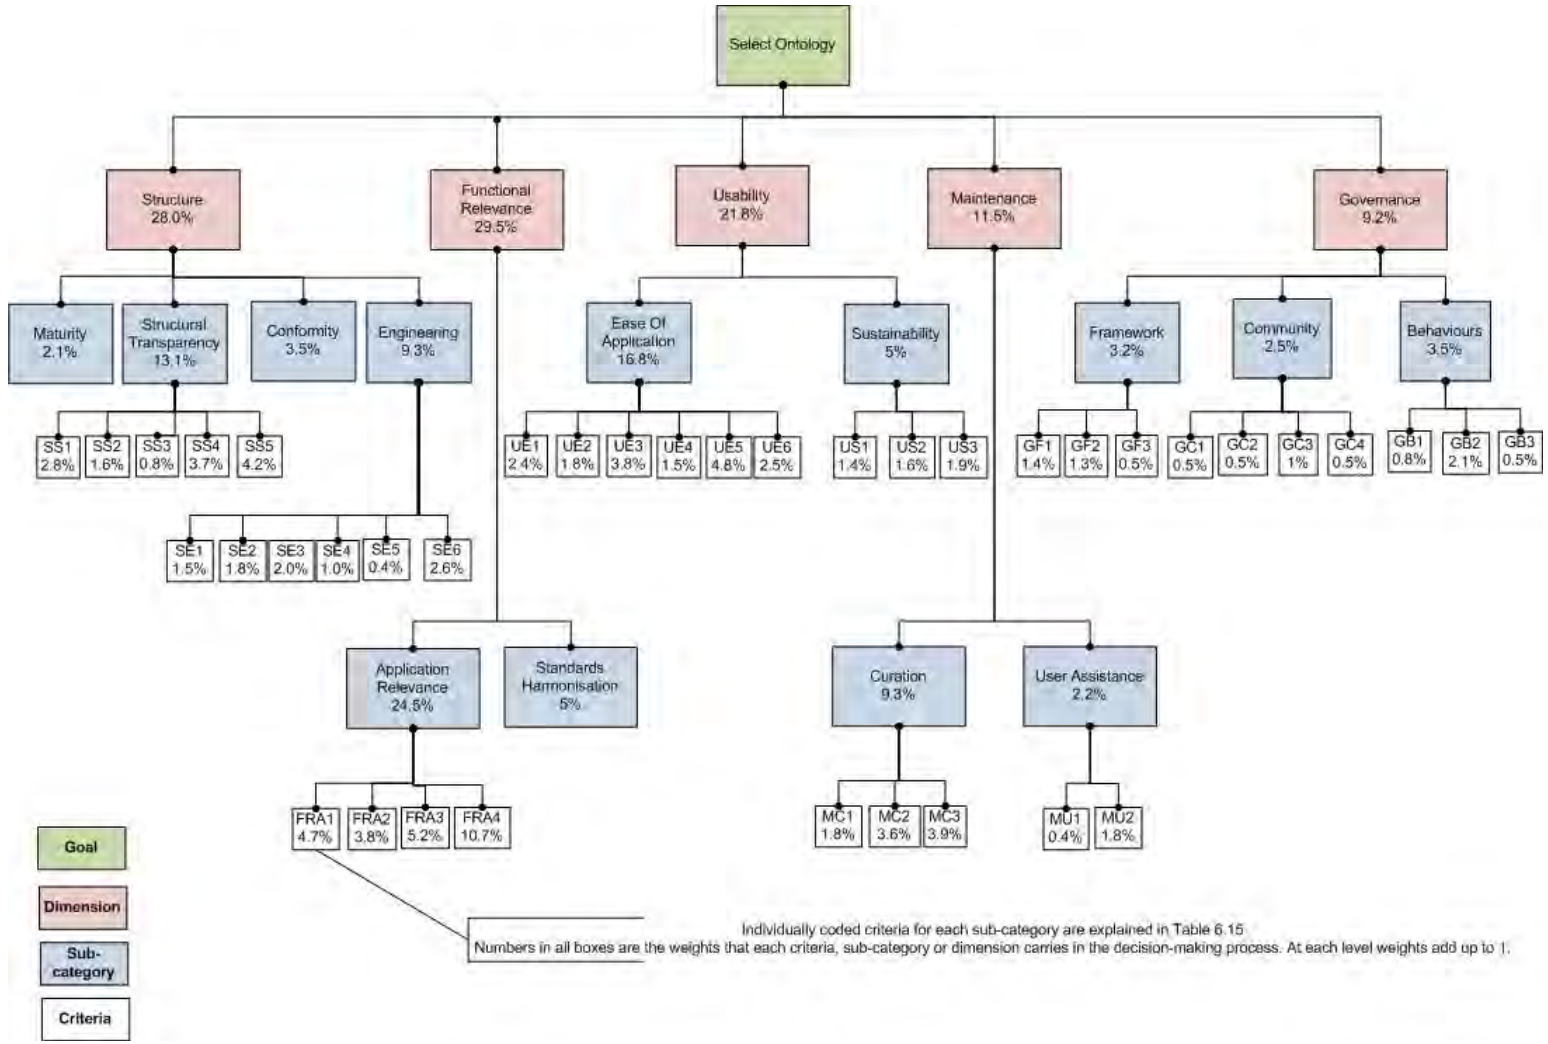

Figure 6.28 Revised Hierarchical Model (Global Percent Contributions)

Two experts ('JH' and 'DH') didn't comment and 'LL' had quite extensive feedback. It should be noted that in commenting on the structure, 'LL' was attempting to map this study's criteria to a questionnaire he had devised previously for a W3C Incubator ontology comparison activity and which he had begun to refine with more detail subsequent to receiving the expert cohort weight results and the hierarchical model, for comment. 'LL's' W3C Incubator Ontology survey question topics were originally at roughly the same level of abstraction as this study's dimension sub-categories. His interest in the model structure and the expert weights was therefore quite acute. His points of view, although very valuable, were assessed in light of underlying motivations (i.e., to improve the structure and applicability of his own emerging assessment model).

Two of the structural issues which 'LL' raised were "redundancy of criteria" and "criteria placement". 'LL' asserted that there was duplication between 'SC1' ("Whether the ontology conforms with language encoding principles and community development rules ?"), located in the "Conformity" sub-category and 'GF2' ("Is there evidence that the ontology conforms with community governance policies and principles (e.g., naming conventions, scope rules, version control ?"), located in the "Framework" sub-category.

On reflection, the author agrees with this point and it is possible that this duplication could have confused experts during their pair-wise comparisons (although 'LL' only picked this issue up after reviewing all of the data – not at the time of taking the survey).

In view of the duplication issue, 'SC1' could be remodelled to only focus on the language encoding principles (i.e., "Whether the ontology conforms with the chosen language's encoding principles?") and 'GF2' could be left to focus on the community-centric conventions. Alternatively 'GF2' could be removed. The first option wouldn't necessarily require any reconsideration of respective weights, although it is acknowledged that experts made their ratings based on the fact that community conformance policies were part of the 'SC1' sub-category (or criterion, since there is only one criterion in the category). This raises an issue of how to conserve, or re-apportion weights during this final stage of model refinement.

It must be remembered that the aim of this part of the study was to arrive at a model which could be generically useful and practical to apply (i.e., perhaps as a more rapid assessment technique than techniques already extant). The AHP process has served well as a structured (and quantitative) means to give an overall sense of the relative importance of different aspects of the decision-making process, from the viewpoints of multiple experts. During this final step in the model refinement process it is not plausible to stay true to the AHP process, in that it is not feasible to iteratively re-issue the model and get experts to re-take the comparative survey if it is subsequently decided (based on feedback, other qualitative data or because of practicality concerns) that some criteria should be omitted, combined or moved. However, the treatment of model weights in each of these circumstances should follow consistent rules.

In previous analyses of the model data it was possible to discern trends in expert responses for many of the sub-category weights and there was, as a result, more confidence in the relative weights in this layer of the model, than in either of the other two. Thus, any tweaking of the model at the criteria level should seek to preserve the weight ratios of the parent sub-categories. Where it is deemed appropriate to remove criteria from the model, the weight that the criteria carried should be equally apportioned to its siblings in the parent sub-category.

Therefore, in the case of the duplicated criteria ('SC1' and 'GF2') it was decided to remove criterion 'GF2', since it is essentially a conformance issue, and the weight carried by 'GF2' (1.3%) could then be apportioned (equally) to 'GF1' and 'GF3' (its local neighbours). However, before taking this step the model weight percentages were scanned to identify those criteria that provided the least "value" in terms of decision-making with a view to removing them altogether from the model. It was reasoned that the fewer criteria that an assessor must evaluate, the more rapid the assessment process (provided those issues that are left to be evaluated are of the most significance). Those criteria that comparatively add least to the decision making process, should be the ones to be omitted as the investment of effort required to evaluate them may not be worth the pay-back overall (if striving for a practical and rapid process).

Seven criteria ('GF3', 'GC1', 'GC2', 'GC4', 'GB3', 'SE5' and 'MU1') had values of 0.5%, or less. This was considered the arbitrary cut-off point for deleting evaluation criteria that contributed least to decision-making. Along with 'GF2', these seven criteria were removed from the model and their weights re-allocated. Of all of these criteria, the removal of 'GC2' (Whether the community is institutionally backed ?) was the only one of concern (to the author). Despite the low value allocated in the pair-wise comparison exercise, the significant advantages that institutional backing brings (through funding and resources) was specifically mentioned by two experts directly ('PF', 'SC'), neither of whom participated in the pair-wise comparison exercise, and by inference by two others ('DH' and 'RH'). Additionally, expert 'LL' felt that 'GC2' had been badly placed in the model and really should have been a "Framework" criterion. There is some sympathy for this view from the author, but it is uncertain what difference this would have made to its rating.

Expert ‘LL’ also noted in his feedback that ‘GF3’ was more of a behavioural issue than a “Framework” criterion, and had ‘GF3’ survived the cull, the author agreed it probably should have been moved into the “Behaviour” sub-category. In addition, there were several other criteria that ‘LL’ felt should be elsewhere in the model but in only two other circumstances did the author agree with his viewpoint. ‘LL’ considered ‘MU2’ (“Quality and availability of published documentation ?”) could have been better placed within the “Ease of Application” sub-category. This criterion could sit in that sub-category, but equally it could reside where it is, so no change to the model was made on the basis of this feedback. ‘LL ‘also considered that ‘UE5’ (“Whether the ontology is interoperable with other ontologies”) would be better placed in the “Functional Relevance” dimension. Criterion ‘UE5’ is focussed on how easy it is to apply the ontology given its engineering from an interoperability perspective. However, ‘LL’ has a point that this is in effect also something that could be evaluated under the sub-category “Application Relevance”. Since it could sit in either category, it is difficult to say what weight it would have drawn if placed in the “Application Relevance” sub-category. In its current placement it drew a rating of 4.8%, the highest rating in its local neighbourhood. How it would have rated in the other category is difficult to say in retrospect. Since it fits within either category it was left in situ.

‘LL’s’ remaining issues with criteria placement concerned: ‘UE4’ (“State of flux i.e., how often are major revisions released ?”) being better positioned in the “Curation” sub-category and ‘GB2’ (“Whether contribution to ongoing development and maintenance is encouraged, open and facilitated”?) also being better located in the “Curation” sub-category. Neither of these placements were deemed appropriate for the following reasons:

- The “State of Flux” (UE4) criterion is more about the perceived usability of an ontology if it is either updated too frequently, or perhaps not frequently enough. It is not about the act of versioning the ontology per se (which does belong under the “Curation” sub-category and is already there as a criterion), instead it is about an ontology’s usability given its state of flux.
- ‘GB2’ is a behavioural criterion and is about how (and whether) the community encourages and fosters contributions, not about whether the act of curation is actually taking place.

The feedback received from expert ‘LL’ indicates the inevitable subjectivity that can be associated with the process of categorisation. ‘LL’s’ detection of duplicates also highlights how different ‘perceptions’ can colour the way in which terms and expressions are applied. Both of these issues are discussed in more detail in Chapter 7.

6.4.2 Selection and Evaluation Framework

The final model (criteria and weights) resulting from all previous analyses is provided in Figure 6.29 with accompanying code descriptions as appear in Table 6.15. This last phase of the research is denoted as step 8 (in Figure 6.1).

There are several possibilities for how this model could be applied by the communities participating in this study, some of which will be outlined here as examples.

It could be assumed, for the purposes of evaluation, that an ontology that completely satisfies all of the listed evaluation criteria that are evident in the model in Figure 6.29 (and which scores top points for each evaluation criteria) could "conceptually" be considered a 'Gold Standard'. If this viewpoint is accepted then using the model's evaluation criteria permits both an evaluation of a single ontology (against a hypothetical Gold Standard) and also supports a comparative assessment between two or more ontologies.

A scenario for conducting an ontology evaluation exercise could encompass using the global percentages evident in the final model in two different ways. These (appropriately rounded) percentages could be considered to be the maximum score that a criterion could be given as a result of a rating activity (with individual criteria summing eventually to 100). But since there are many low value percentages, which don't really provide for a good degree of discrimination for comparative purposes, these global percentages could simply be used as multiplication factors. In this approach each criterion would be rated using a community- applied scoring system (similar to that demonstrated through the AHP process, where Saaty's (1980) nine-point rating scale was used). The weight factors would then be applied and total scores for each adjusted criterion then summed to give a final score.

If the purpose of the method was to assess a single ontology candidate, a threshold minimum (total) score would need to be set for the sum of criteria, prior to the evaluation exercise using expert judgement. An ontology falling below the threshold would be considered less than ideal. This may not of itself need to trigger rejection of the ontology, but it would indicate inadequacies with the ontology. Because the assessment is structured and targets evaluation of different facets (sub-categories and dimensions) of the ontology, it is possible using this technique, to quickly identify where the major inadequacies are.

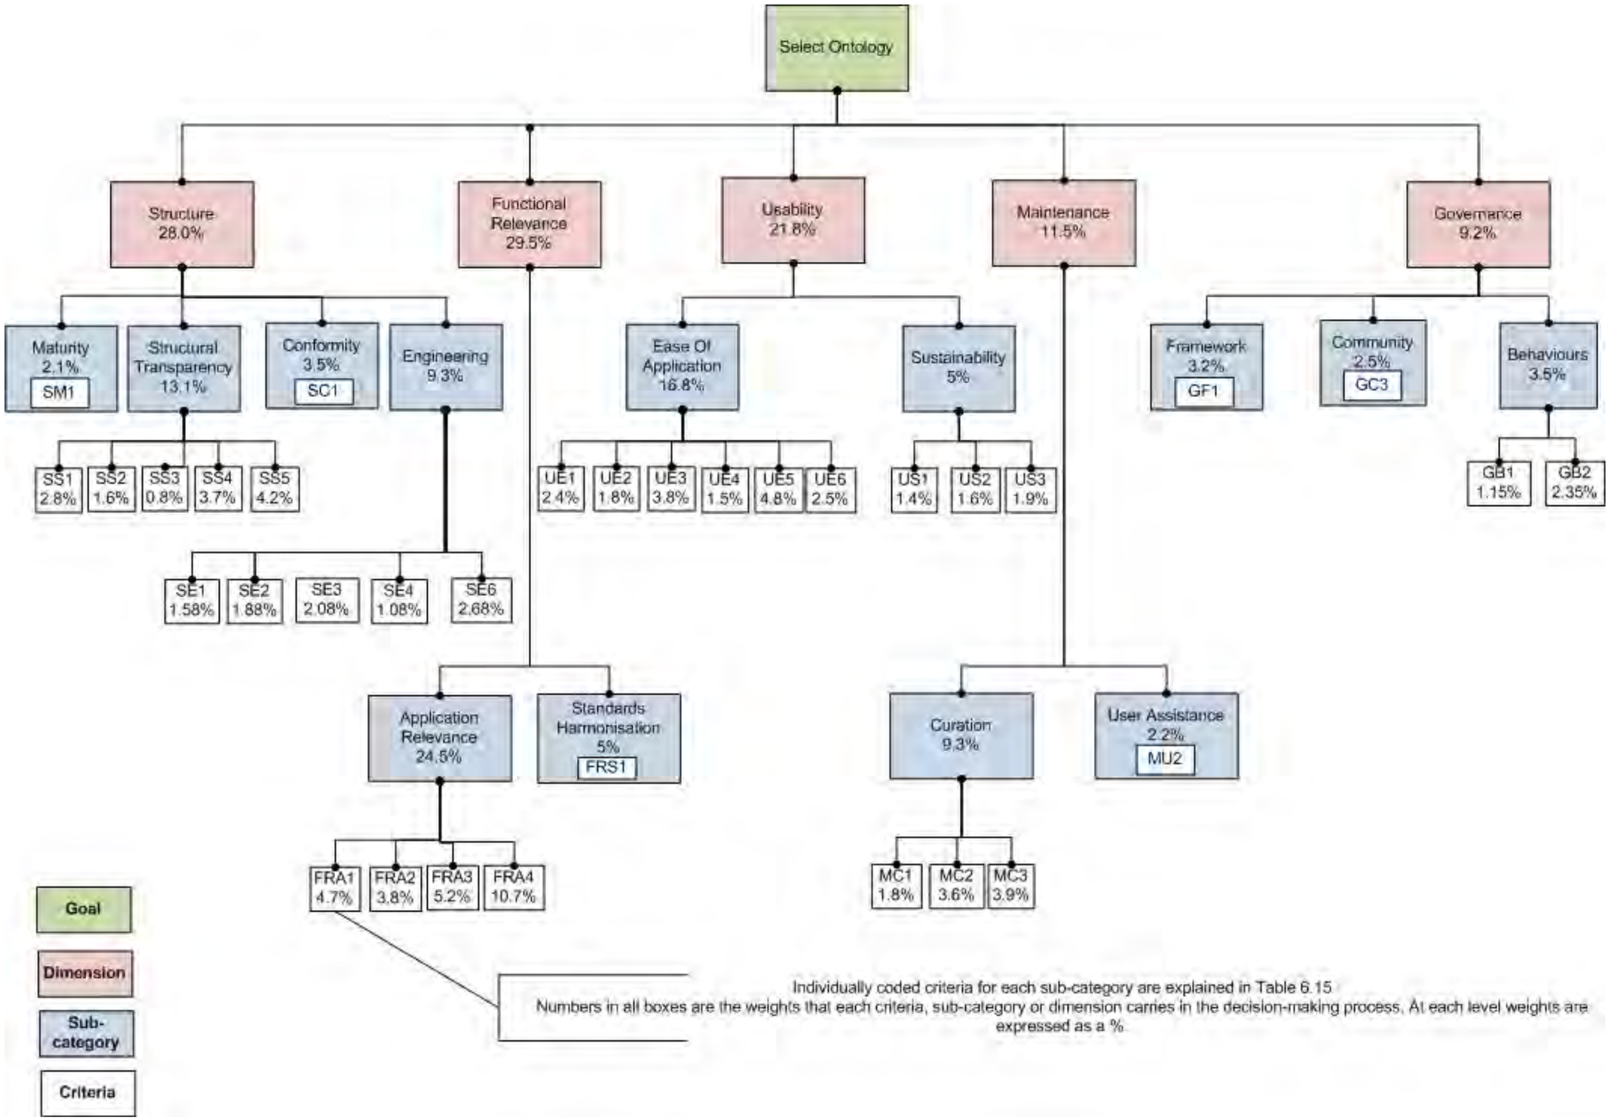

Figure 6.29 Final Model With Global Percentages

If the function of the evaluation was to compare two or more ontologies, the scoring system will indicate which ontology rates higher for the particular purposes in mind. Although the multiplication factors are considered to be indicative of the weights experts have collectively given to the level of importance placed on specific aspects of an ontology's utility and merit, there is still scope to modify the weights according to individual expert judgement and for specific project needs, prior to commencing an assessment exercise. Nothing more than a spreadsheet is required to manage the assessment task, although there are available proprietary (e.g., ExpertChoice (Expert Choice Academics, 2012)) and open source tools (e.g., SuperDecisions (SuperDecisions, 2012) that can be used to run AHP-style assessments.

Evaluation Criteria Measurement

Scoring aside, the criteria (represented in Figure 6.29) on their own do not lend themselves (in many cases) to direct evaluation. To make an assessment about whether a criterion is satisfied and to what extent, requires articulation of evaluation "measures". Measurement, according to the IEEE (1998) is "the act or process of assigning a number or category to an entity to describe an attribute of that entity". Table 6.30 lists a range of evaluation measures, organised by criteria, which were either (a) mentioned directly during expert interviews (refer to Appendix 20), or (b) extracted from expert-cited or supplied material (refer to Appendix 20), or (c) drawn from the general literature (and annotated as such in the Table). Measures sourced from the literature by the author, filled gaps where no measures were mentioned by experts. These latter measures were included because for the framework to have utility for the communities participating in this study, it needs to be complete. There were a number of cases in which criteria were mentioned by experts, but when questioned at interview about how these criteria were actually measured and evaluated, no examples could be provided. Table 6.30 summarises results with respect to RQ1.2.3 - "*What evaluation measures can be used to assess evaluation criteria* ?".

Yu (2008) has argued that without the analysis of ontology requirements for a given application, a criteria and measures-based approach alone may lack significance and relevance for ontology application. He reasons what one application (of an ontology) may consider to be an error, or an inconsistency, may not necessarily be so in another application. At interview, ten out of the thirteen experts mentioned that "meeting the use-case goals or competency questions" was an important criterion for an ontology to fulfil, lending credence to Yu's (2008) view. He further states that measures deemed important for one application may not necessarily be important for others.

The ontology types mentioned in this study included 'domain' and 'application or task' ontologies. Application ontology developments were less common having been developed by only two experts (i.e., 'PF' and 'CA'). Recall (from Table 6.17), however, that expert 'PF' was the most prodigious contributor of criteria. It should therefore be assumed that the evaluation criteria listed in the hierarchical model generally cover criteria relevant to the evaluation of both of these types of ontology. The main reasons cited by experts for developing and reusing ontologies included for the purposes of: "improving online data search functionality"; "enhancing community ability to integrate data from disparate sources"; "improving the interoperability of distributed systems" and "data annotation". Whilst these are broad areas of stated use, it is important to note that the criteria and measures reported should predominantly reflect these use-cases and presumably are generally applicable for the cited situations.

Table 6.30 Ontology Selection and Evaluation Criteria Measures

Sub-Category	Evaluation Criteria Code (from Table 6.17)	Suggested Evaluation Criteria Measure
Maturity	SM1	(1) Identify the number of ontology versions (iterations) that have been released (to form an impression of maturity from this perspective), and/or (2) Identify the length of time the ontology has been in use (to form an impression of maturity from this perspective).
Structural Transparency	SS1	(1) Availability of concept diagrams (or directed graphs) showing key concepts, concept relations and properties (could limit evaluation to "yes" or "no" and assess the completeness of these diagrams).
	SS2	(1) Number of "full" owl:imports or "partial" (MIREOT-like, i.e., selective use of classes from external ontologies) import statements (See Courtot *et al.,* 2011), and/or (2) Assess the impact of import statements on the scope of the original ontology with respect to its intended use.
	SS3	(1) Is instance data included (could limit evaluation to "yes" or "no")
	SS4	(1) Use Protege Metrics (Stanford University, 2011) or OntoQA (Tartir *et al.,* 2005) to gain an impression of the ontology design and the type and extent of the knowledge represented (e.g., examine class richness, ontology cohesion, class importance and inheritance).*
	SS5	(1) Assess the use of descriptive text, any inline annotation describing the origin of terms and the use of annotation.
Conformity	SC1	(1) Assess the availability of formalised term and property labelling conventions, and/or (2) Use commercial, open or community-based ontology editors to check for language-specific conformance, and/or (3) Perform random checks of the ontology against documented community encoding conventions (and any stated ontology design patterns).

Engineering	SE1	(1) Assess whether inherent ontology primitives enable appropriate sub-classing (i.e., specialisation) to suit needs.
	SE2	(1) Assess whether the ontology readily partitions*, and/or (2) Assess whether or not it is easy to extract concepts and dependent relations and axioms without loss of ontological consistency (when importing ontological fragments from elsewhere)*.
	SE3	(1) Perform consistency checks using a reasoner (including testing for unintended models through incorrect disjunction).
	SE4	(1) Qualitatively assess the verbosity of typically encoded data used in service transactions (in real-world use-case scenarios).
	SE6	(1) Perform a quick visual inspection of the ontology to assess whether the ontology instils a sense of confidence in its engineering after looking at aspects such as: the number of concepts, concept depth, cohesion, tangledness (i.e., number of direct subclass relation being present while simultaneously being derivable from a chain of other sub-class relations), any redundancy of concepts and properties, concepts not joined to anything else and its completeness.
Application Relevance	FRA1	(1) Assess the ontology scope (via an examination of competency questions and the use-cases it purports to satisfy) and the applicability of the implementation language, and (2) Test the ontology against a random sub-sample of appropriate competency questions (using available query engines).
	FRA2	(1) Inspect the ontology for missing concepts.
	FRA3	(1) Assess applicability of scope and type of included vocabularies.
	FRA4	(1) Qualitatively assess the ontology focus with respect to its intended use (through an examination of its concepts and properties relative to the use-cases it is intended to support). Does it include many "apparently" superfluous terms ?
Standards Harmonisation	FRS1	(1) Does the ontology derive from an upper ontology (could limit evaluation to "yes" or "no"), and (2) Can the ontology be used in conjunction with other (perhaps less ontologically formal standards) for example through use of semantic annotation of application schema or data instance documents, or through provided conversion tools that use both the ontology and other types of resources.
Ease Of Application	UE1	(1) Using known terms (not modelled as instances in the ontology) attempt to model the data and use the model to answer a number of competency questions. Assess (a) ease of modelling the data and (b) ease of querying the model to answer competency questions and (c) degree of precision in answering competency questions.
	UE2	(1) Does the ontology contain recognisable and re-usable ontology design patterns (evaluation could be limited to "yes" or "no" ?) and (2) Assess whether there are existing tools already capable of manipulating these design patterns.
	UE3	(1) Check compliance to Linked data deployment rules (see for example - http://www.w3.org/DesignIssues/LinkedData.html) using Vapour (http://validator.linkeddata.org/vapour) and/or a URI debugger.
	UE4	(1) Assess a Community-shared understanding of the: (a) conditions to be met for issuing a minor revision, and (b) conditions to be met for issuing a major revision, and (c) average or maximum period between minor and major revisions.
	UE5	(1) Look for any occurrences of ontology design patterns that would be likely to generate inconsistencies. The ideal situation is to find none.
	UE6	(1) Availability of software and User Interfaces exploiting the ontology

Sustainability	US1	(1) Number of users, and (2) Size and type of projects (and whether they are exploratory or production in nature), and (3) Size and commitment of governing organisations.
	US2	(1) Number of similar projects stopping, maintaining or starting similar developments: (a) using the same ontology language, and/or (b) using a different modelling language
	US3	(1) Number of projects using or intending to use the ontology: (a) using the same ontology language (b) using a different modelling language (2) Number of positive and/or negative reviews of the ontology.
Curation	MC1	(1) An assessment of the number and skill of people maintaining the ontology.
	MC2	(1) Number of ontology "committers".
	MC3	(1) An assessment of: (a) the type of version control system, and (b) the handling of versions through the URI scheme, and the availability and useability of older ontology versions.
User Assistance	MU2	(1) Availability and quality of: (a) Term dictionaries with term descriptions, (b) Descriptions of relations and axioms,
Framework	GF1	(1) Is the framework published (“yes” or “no”), and (2) An assessment of governance "strength" of the organisation supporting the development (e.g., is it a standards development organisation, research consortium, individual organisation with a research interest? Or with a commercial interest? Or with an operational interest?)
Community	GC3	(1) Number of persons in total (chairs, committers, followers), and (2) Number of hours per week allocated to the ontology development, and (3) Level of email traffic (or similar), and (4) Number of face-to-face meetings per year focusing on group work.
Behaviours	GB1	(1) Adequacy of licensing regime (open is preferred). (2) Community membership rules: are they sympathetic with your community’s participation ? (3) Positive mentions and referrals in academic and commercial conferences/workshops and in training events.
	GB2	(1) Is there a community web forum or wiki ? and (2) Are tools available to support contributions ?

*a measure taken from the literature (not mentioned by an expert)

Regardless, it is important in the application of the hierarchical evaluation model that criteria and measures are appropriately matched to requirements. Ontology requirements specify the competencies, capabilities, functionalities and qualities that are needed from a candidate ontology, for a given application.

Yu (2008) suggests articulating ‘questions’ for each stated requirement and then ensuring that these questions relate requirements to measures. Generally, in his evaluation methodology (i.e., ROMEO; Yu, 2008) the questions are based on aspects of an ontology evaluation criterion. Measures chosen to help evaluate a criterion should provide an answer to the question(s) posed. During interview several of this study's experts also mentioned the development of use-cases and inferred the use of competency questions as key parts of their overall ontology 'evaluation methodology', further supporting the need to incorporate a mapping between requirements, the criterion and measures to be used in any given situation. By employing this type of method only those criteria and measures applicable to the context at hand need to be used in a selection and evaluation process (drawing from the criteria and measures available in the Framework).

The weighted hierarchical model of evaluation criteria, whose utility has been grounded by derivation from expert experience, paired with relevant evaluation metrics (from Table 6.30), together with the described methods of application, is considered to be an 'expert- grounded ontology evaluation framework' suitable for application in Antarctic science. Since most of the experts interviewed were practising reuse and using their evaluation skills to build domain ontologies, the framework delivered is considered highly suitable for use in selecting ontologies that could be used to populate a Feature Catalogue (which is itself a repository for domain ontologies that can be used for dataset annotation purposes).

Comparison of Framework With Other Methods

In reviewing literature-based ontology evaluation methods, Yu (2008) observes that methods generally fall into two groups. The first type uses a life-cycle approach to ontology construction in which evaluation is performed iteratively, using pre-determined requirements, usually tested by competency questions, such as On-To-Knowledge (Sure *et al.,* 2003), Methontology (Lopez *et al.,* 1999) and the method proposed by Gruninger and Fox (1995). The other types specify criteria that are useful for evaluation, but generally can be independent of a build-cycle (e.g., most of the methods highlighted in Table 2.3, Chapter 2). The evaluation framework presented in this thesis falls into the latter category and is aimed at scenarios in which domain practitioners need to evaluate existing ontologies for the purposes of re-use. Re-use may be practised in building a completely new ontology, or is harnessed to help supplement an existing ontology in order to continuously improve its ability to support domain/task use-case goals (or its interoperability with other ontologies).

The framework presented in this thesis was inspired by the ONTOMETRIC approach of Lozano-Tello and Gomez-Perez (2004), because of ONTOMETRIC's use of AHP which was viewed as a good mechanism for generating a structured model of weighted criteria. Having used AHP to establish the bias that experts apply to certain types of evaluation criteria, its utility more broadly as decision-making tool, was reinforced, albeit with some noted drawbacks (discussed in the next chapter).

Despite the similarity of the underlying approach (i.e., the use of AHP), there are notable differences between ONTOMETRIC and the hierarchical framework that resulted from this research.

The most obvious difference is that ONTOMETRIC has ‘160’ characteristics (or evaluation criteria) against this hierarchical framework’s ‘42’. The main dimensions in ONTOMETRIC are categorised as “tools”; “language”; “content”; “methodology” and “cost”. Because ONTOMETRIC was developed in 2004, it is not surprising that the issue of ontology encoding language was considered significant for inclusion. In 2011, there has been a convergence in the use of languages towards OWL (mostly OWL-DL) for domain practitioners wishing to deploy ontologies on the Web. OBO Foundry ontologies are often still written in OBO Language but this has also now been mapped to OWL (Tirmizi *et al.,* 2011). Taxonomies and vocabularies in general are usually written in RDF. Presumably, as a result, ontology experts in this study didn’t consider language-specific issues to be something that required significant attention.

The “cost” dimension of ONTOMETRIC also didn’t appear to have resonance with this study’s experts. This is possibly because there has been a considerable surge, since ONTOMETRIC was developed, in the availability of open source tools for ontology development, management, search and visualisation (and even storage can be obtained for free). The costs incurred today are those involving the skilling-up of practitioners in the use of available tools and the costs associated with the “people-power” involved in ontology-building. As a result the issue of cost as an evaluation criteria, did not factor overtly in expert commentary. Rather the issue of sustainable resourcing for community development activities was an often raised topic. A number of experts indicated that part of their decision-making about whether to re-use a particular ontology, involved an assessment of the resources available, within a particular ontology source community, to maintain and sustain a source ontology into the future.

The example characteristics provided under the “cost” dimension in ONTOMETRIC relate more to material and capital costs such as those incurred in purchasing software, hardware, tools and ontologies. Most ontologies developed within the scientific domain, are now usually considered public goods, so characteristics such as ontology license costs would not typically be the type of issue, or criteria an expert would consider needed evaluation. The closest criteria in this study to ONTOMETRIC’s licensing (cost) characteristic was the “accessibility” of an ontology, but mainly from the perspectives of how well it is advertised and /or the ease with which it can, in whole, or in part be resolved on the web.

The other significant difference between the ONTOMETRIC approach and that given in this study was that this study provided one or more possible qualitative and/or quantitative measures for each of the criteria listed. Although ONTOMETRIC had a lexical-based scale to rate each ONTOMETRIC characteristic (which was converted to a number during AHP calculations) the actual measure being evaluated in most cases was not stated. For example, consider the ONTOMETRIC characteristics "Axioms_Infer_Knowledge" or "Importance_Of_Developed_Ontologies". It is not clear what aspects of these statements are to be evaluated. Independent reviewers might each assess different things for many of the characteristics listed (i.e., for those characteristics where the assessment is not simply numeric by virtue of its definition). Although the framework developed in this study leaves the assignment of a rating scale to the framework user, the scope of what is to be "measured" has been narrowed for each criterion provided in the hierarchical model.

It was not possible to perform a review of how the various characteristics in ONTOMETRIC (i.e., criteria in this study's hierarchical framework) were originally established because the methods explained in detail in Lozano-Tello (2002) were described in Spanish (with no apparent English translation available).

Apart from the commonality of approach with the ONTOMETRIC methodology, the framework derived in this study, also has some similarity with a framework developed by Fernandez-Breis *et al.* (2009). They have proposed a quality evaluation framework for bio-ontologies based on the ISO 9126 (ISO/IEC, 2004) standard for software quality. The Fernandez-Breis *et al.* (2009) augmented framework consists of seven dimensions (Structure, Functionality, Reliability, Usability, Efficiency, Maintainability and Quality-In-Use). Each of these dimensions has at least one, and in some cases two levels of sub-category. There are 38 quality metrics in total (i.e., statements considered to be at the same level as the evaluation criteria in the model in this thesis).

In applying the Fernandez-Breis *et al.* (2009) framework, ontology evaluators are expected to rate each quality metric with a value between 1 (worst) and 5 (best). The actual "metric" being used for each criteria is not evident from the description supplied in the publication, so it is assumed each evaluator could interpret the quality metric differently and therefore apply different measures. The method was trialled by Fernandez-Breis *et al.* (2009) using 8 MSc. students who were asked to compare and rate two ontologies. Each student was given 20 hours of training in designing ontologies prior to being asked to pilot the method. It was claimed that because no student reported problems with applying the method, coupled with the fact that there was some level of consensus amongst some students on ratings within dimensions, that anyone with some basic knowledge of ontology construction could easily apply the method. Some facets of evaluation e.g., 'tangledness', ‘cohesion’, ‘domain coverage’, ‘structural accuracy’, ‘reusability’ which were criteria in the Fernandez-Breis *et al.* (2009) applied ISO-derived framework require a relatively sophisticated understanding of ontology design (and the domain in question) and the author doubts that just 20 hours worth of introductory tutorial would prepare an individual to adequately assess such factors (unaided). Additionally, the experimental design of the pilot (using the 8 students) did not appear to permit some of the conclusions subsequently made about the ease of application of the method.

However, the goal of Fernandez-Breis *et al.* (2009) in developing their method, as was for the case in this study, was to establish a credible and easy-to-use methodology for ontology assessment. They founded their methodology in an existing software quality framework in contrast with this study which took a bottom-up approach of grounding criteria in expert experience. It is therefore interesting that the actual dimensions (sub-categories and criteria) listed by Fernandez-Breis *et al.* (2009) have a good degree of overlap with those listed in this study, which lends further credence to this study’s results, since the two sets of evaluation criteria were arrived at by different routes.

6.5 Practical Ontology Selection and Evaluation Summary

In this chapter the results of all research investigations into RQ1.2 – “*What typifies an expert-grounded ontology selection and evaluation framework that can support multi-disciplinary Antarctic science communities using Web services* and all of its sub-questions have been answered.

Through an expert Screening Survey, which helped provide information for selecting suitable study participants, thirteen experts familiar with ontology re-use were identified and recruited to take part in this research. These experts collectively covered six of the disciplines represented in Antarctic science and included individuals from national, as well as international institutions. There was a relatively high degree of overlap in the communities covered by the experts, with several experts active across a number of common communities. The coverage of disciplinary domains was considered good but the number of participants available and qualified to assist with the research was not optimal. As the study progressed some participants informally dropped out of the study. In following up responses to emails sent to experts it was clear that this was because almost all experts were in demand and were by their own admission time poor.

After analysis of the Screening Survey data, each participating expert was interviewed, their interviews recorded, transcribed and qualitatively coded (using Template Analysis techniques of Crabtree and Miller (1992) and King (2004)). An important outcome of this coding activity was the development of a three-tiered hierarchical ontology evaluation criteria model which could be used as input to a pair-wise comparative model element ratings exercise, harnessing AHP techniques (Saaty, 1980). The result of this exercise, after adjustment of any inconsistent expert provided data and by synthesising a group result, was a weighted hierarchical ontology evaluation criteria model. The weights in the model represented the relative level of importance that experts placed on the model's criteria, sub-categories and dimensions. Although the group result showed a 'fair' degree of concordance between expert ratings (as measured by Kendall's co-efficient of concordance), there was never-the-less considerable variability in weights allocated amongst respondents in a number of between model dimension and sub-category comparisons. Feedback obtained from experts indicated satisfaction that the model, in general terms, reflected the criteria used in ontology selection and evaluation. Opinion was divided, however, regarding the suitability of using the geometric mean as a measure of a 'group' result. Development of the weighted model in part, provided an outcome with respect to RQ1.2.2 – "*Is it feasible to derive a weighted evaluation criteria model in which criteria are rated according to importance ?*", but the results indicate that there was only a 'fair' degree of concordance, between experts regarding allocated weights. With only a 'fair' level of unanimity between experts, it was important to investigate whether there were any patterns, or clusters of experts, in terms of how elements of the model were rated, that might explain the variability and the less than ideal level of concordance. In doing so, RQ 1.2.2.1 and RQ 1.2.2.2, which both focus on unearthing reasons for any detected differences in how experts rated criteria, were also addressed.

The Screening Survey and the in-depth interviews permitted a relatively rich stratification of the experts according to various characteristics (e.g., disciplinary groupings; skill type; ontology development experience; ontology application area; community governance type; the roles played by experts within communities; community sizes/maturity and groupings based on expert perceptions of key terms). These stratifications were drawn upon, as possible explanations for patterns of similarity or dissimilarity between expert rating data. Where there were any discernible groupings in how experts rated specific criteria, sub-categories or dimensions, the various stratifications were consulted to see if they could explain the patterns in the results. In no cases could the stratifications be used to explain the patterns found.

There was also a relatively high degree of inconsistency found between ratings provided by some experts during the pair-wise comparison exercise and 'off-the-cuff" comments that they had made during interview, about the relative level of importance of particular evaluation criteria. So, the study revealed both variability between experts in how model elements were rated and an inconsistency within individual ratings made by experts. The necessity to adjust individual expert data due to relatively high levels of ordinal and transitive inconsistency during the pair-wise comparison exercise also demonstrates the degree of difficulty experienced by experts in cognitively rating ontology evaluation criteria.

The final hierarchical model delivered in this chapter was adjusted to remove very low rating criteria (i.e., criteria with ratings with a score of 0.5% or less) and the scores for these criteria were apportioned to siblings (i.e., other criteria within the same sub-category). In addressing RQ1.2.3 – *"What evaluation measures can be used to asses evaluation criteria ?",* evaluation measures mentioned at interview, or derived from material cited by experts, supplemented by other references that were sourced by the author from the literature (in cases where there were gaps), were matched with criteria from the hierarchical model. A description of how the model and measures could be applied by the communities mentioned earlier in this thesis was provided. Taken together, the model, the matched measures and application methods form the selection and evaluation framework developed from the experience of the experts that participated in all phases of the research described in this chapter.

Although a key outcome from qualitatively analysing the interview data was the formulation of a hierarchical evaluation criteria model, the data was also mined to uncover information concerning methods (as opposed to criteria) potentially used by experts, or communities to select and evaluate ontologies (RQ1.2.1). How communities manage and govern ontologies was also of interest. Methodologies for ontology selection and ontology governance form important guidance for the SCAR and AODN communities participating in this study. Data captured from the coding process pertaining to these issues is discussed next in Chapter 7 along with a more detailed dissection of some of the results reported in this chapter, particularly the variability and inconsistency in expert ratings of model elements.

Chapter 7.

Ontology Selection and Evaluation Discussion

The previous chapter presented the data analysis and results associated with investigations regarding ontology selection and evaluation criteria and methods, using data from ontology experts active in various (Antarctic-discipline relevant) scientific communities. An important aim, in Chapter 6, was to derive an expert-grounded evaluation framework that could be used by the SCAR and AODN communities in populating the Feature Catalogue artefact specified and prototyped in earlier chapters. The weighted hierarchical evaluation model that formed the basis of the eventuating framework was a reflection of expert practise and the weights applied to individual elements in the model were relative levels of importance. In developing the 'framework', results stemming from development of the hierarchical evaluation criteria model and the associated pair-wise expert analysis of hierarchical model elements (conducted to assign importance levels), were scrutinised from a number of perspectives (predominantly different stratifications of experts and community related data). The purpose of this scrutiny was to suggest reasons for any similarities or patterns in how experts weighted the evaluation model elements.

This chapter discusses possible reasons for why the patterns detected in the results could not be grounded in the various stratifications that were performed to group expert data in this study. This discussion of necessity encompasses a critical review of the methods and techniques used to solicit and analyse data and draws on the literature to rationalise observations of inconsistency in expert-provided information. Despite the inconsistencies found whilst generating the framework, the framework is still considered a 'practical' tool that can lead a novice (in ontological evaluation) to focus on evaluation criteria of importance and which permits the disambiguation of a decision-making activity into constituent parts, therefore making ontology assessment particularly transparent.

Data captured in Chapter 6, which is able to address a component of RQ1.2.1: *"....are selection and evaluation methods* [used by experts] *consistent with those reported in the literature* ?" which was put aside until now, will be exposed and discussed in this chapter. General theories and observations regarding community-based ontology governance, which emerged from in-depth interviews and subsequent data analysis, are also presented and the implications arising for the two communities-of-interest that participated in the Feature Catalogue development phase of this research are noted.

7.1 Inconsistencies in Expert Responses

There was a relatively high degree of inconsistency evident in the data captured from experts during this part of the study which was investigating issues of practise with respect to ontology selection and evaluation. This inconsistency occurred in three different areas, between:

a) Answers provided by individual experts during the pair-wise comparison ratings exercise (i.e., internal inconsistency). The AHP methodology is sensitive to inconsistencies made between pair-wise comparisons, and the level of inconsistency measured within each expert's original survey responses was relatively high (often CR > 0.10 with a maximum inconsistency of up to 0.866, which was later reduced to be <=0.14).
b) Experts in terms of their ratings of the various hierarchical evaluation model elements.
c) The general comments made by experts at interview regarding the overall level of importance of criteria and the subsequent rating of these criteria during the pair-wise comparison exercise.

Deeper insight into why these various forms of inconsistency might be manifest is provided here, which may aid further interpretation of the data and reveal limitations in the methods applied.

7.1.1 Methods and Expert Judgement

AHP is a problem-solving method and a systematic procedure for representing the elements of any problem (and their relative levels of importance, or 'preferences') which is why it was an attractive choice for use in this study to examine how ontology experts rate the various evaluation criteria that are being used in practise to select ontologies for reuse.

A moderate amount of inconsistency in AHP is inherent to the nature of human decision making, as each pair-wise comparison judgment will be approximate rather than exact (Saaty, 1980). However, high levels of inconsistency are indicative of problems with an expert's pair-wise preference judgments (or weights). An extreme example of inconsistency occurs if preference transitivity fails to hold. In general, an inconsistency index exceeding '0.2' suggests an inversion error (e.g., B is preferred to A, but A is really preferred to B by the expert) or a set of major inconsistent judgments. Values between '0.10' and '0.20' may indicate that a judgment was entered incorrectly, or that the expert really does have a high level of inconsistency (Williams *et al.,* 2007). Many results reported in Chapter 6 had weights > 0.10 Saaty's (1980) recommended upper CR limit (see Table 6.26).

AHP is based on three principles: problem decomposition, comparative judgement and synthesis of priorities (Saaty & Vargas, 1991). It is a theory of measurement for dealing with both quantifiable and intangible criteria (Saaty, 1980). As an analysis method it is predicated on a small set of assumptions or axioms. One essential assumption is that a decision-maker can quantify his or her preferences regarding the situation at hand (Keeney, 1982). This is generally taken to mean that relative preferences can be characterized by a single scalar number or point value. For example, if an expert can indicate that one outcome is “twice” or “five times” as preferable as another, this assumption is considered to hold. When this assumption holds, we can take the expert’s numeric judgments and apply a series of mathematical operations to them to obtain a solution (Hahn, 2003). A consequence of this quantification assumption is that error in judgments is typically considered to be non-existent or negligible. Expressed differently, judgments are typically taken to be certain and thus can be represented by scalar values (Hahn, 2003).

Causes Of Internal Inconsistency

In the traditional formulation of the AHP (as used in this research), expert judgments have been solicited as exact (or crisp) numbers. However, in many practical cases the human preference model is uncertain and decision-makers might be reluctant, or unable to assign the exact numerical values to the comparison judgments (Duran and Aguilo, 2008). Traditional AHP has an inability to adequately handle this potential inherent uncertainty and imprecision associated with the mapping of a decision-maker’s perception to exact numbers (Lefley and Sarkis, 1997; Deng 1999). So, although the use of the discrete scale of 1–9 has the advantage of simplicity (which was considered important in this study), the AHP technique (using this crisp scale) does not take into account this potential uncertainty associated with the mapping of one’s judgment to a number.

Recognising this inability to represent uncertainty, adaptations of AHP have been developed where the decision-maker’s comparison judgements are represented as fuzzy triangular numbers which allow for the incorporation of the vagueness of the human thinking (van Laarhoven and Pedrycz, 1983; Xu, 2000; Mikhailov and Tsvetinov, 2004; Sredjevic and Medeiros, 2008). This fuzzy number is often expressed as a triple (*a, b, c*), where ‘*b’, ‘a’* and ‘*c’* are the mean, the lower and the upper bounds, respectively. Saaty ‘s (1980) nine-point rating scale is often still used, but is transformed into triangular numbers.

Given the high degree of (ordinal and/or transitive) inconsistency evident within expert pair-wise comparison data in this study, perhaps an application of AHP using a fuzzy rating scale would have been preferable, since the inconsistency expressed may have been due to expert difficulty in rating hierarchical model elements that were considered very similar in terms of importance. However, it is not clear how fuzzy scales would improve situations where experts had genuine difficulty in assessing the relative importance of one element with respect to another because cognitively they did not have any sense of what was more or less important (i.e., they genuinely hadn't formed an opinion and couldn't do so for the purposes of the exercise). In such situations experts may have simply resorted to random choices, thus rendering their data internally inconsistent. It is feasible that this occurred within the expert cohort in this study. For example, one expert, who failed to complete the pair-wise comparison survey, and who withdrew from the exercise (after starting the survey), stated in part:

> *"things I was being asked to compare seemed semantically disjoint - like asking me do I prefer chalk to cheese!."*

Although application of AHP using fuzzy numbers is widespread, Stewart (2008) has expressed some concern with this approach with respect to the robustness of the method, in that he believes the ranges of imprecision should be made precise for the fuzzy number sets used. In effect he suggested that the limits defining a triangular fuzzy number should really be another fuzzy number (which presumably could go on *ad infinitum* in terms of boundary definition and would therefore significantly complicate AHP calculations).

Whilst the levels of inconsistency detected in this study require discussion, significant levels of inconsistency are not unusual in conducting AHP (Kryvobokov, 2005; Meixner, 2009). Williams *et al.* (2007) describe several types of error that experts make in undertaking pair-wise assessments (i.e., criteria, rank inversion, inconsistency, mechanical and omission errors). Criteria errors occur when the expert misjudges the importance of a criterion – either by interpreting an important criterion as unimportant, or the other way around. Rank inversion errors are human errors that occur when the expert mistakenly ranks a less attractive alternative higher than a more attractive alternative, when the two are very close in value. Omission errors are more likely when the expert skips from one aspect of the problem to another without considering any aspect in-depth, or focuses on one narrow aspect of the problem to the detriment of others.

Since experts in this study had to perform 80 pair-wise comparisons (in total) which is a relatively large number of comparisons to make, about a wide variety of issues, there was a propensity for these types of error to occur. It is widely recognised that the ability of humans to accurately express their knowledge decreases with increasing problem complexity. Thus, as the number of criteria in AHP increases, experts are likely to make inconsistent judgments during pair-wise comparison (Lin *et al.*, 2008). Brugha (1998) has demonstrated that when there are an excessive number of questions, there is a drop off in interest towards the end of the questioning process and there are consequently higher levels of reported inconsistencies.

Pair-Wise Comparison Survey Construction

In addition to the matters described above, survey structural issues could have contributed to inconsistent ratings/rankings. Although considerable care was taken in structuring the comparative survey, the way in which the model elements to be compared were framed, perhaps caused confusion. One expert (the same one who had difficulty comparatively rating model elements, discussed earlier), stated the following:

> *"I found the language in a lot of the questions difficult to understand. Sentences like ''Whether a visual inspection of the ontology gives confidence after looking at aspects such as: number of concepts, cohesion, tangledness, redundancy of concepts or properties.' means absolutely nothing to me - I'm a geologist turned oceanographic data manager, not a computer scientist. ....Sorry I can't be more help. A smaller number of questions with much more explanation and examples might be easier for the likes of me."*

However, no other expert who was provided with a survey indicated that they had any difficulty interpreting the criteria to be assessed. Other experts who failed to complete the survey did not attempt to commence it. In addition to the survey questionnaire, each survey participant was sent survey instructions which included sample marked up data; a table of the model dimensions, sub-categories and evaluation criteria and a graphical snapshot of the model showing the hierarchical arrangement of the questions, all to aid the respondent in completing the survey. On reflection, a glossary of terms should also have been included to provide experts with definitions for all terms used (as they were interpreted during interviews).

It was not possible to use a smaller number of questions, as suggested by the expert, because the very purpose of the survey was to get each expert to comparatively rate, in a pair-wise fashion, each model element.

Hierarchical Model Construction

It was stated in Chapter 3 that in constructing the hierarchical evaluation model particular attention was paid, during the template coding process, to avoid the introduction of redundant concepts. This seemed logically sensible but was also based on the recommendations of Baker *et al.* (2002) to optimise development of the decision-making model. Despite strict attention to this tenet, expert feedback still revealed possible duplication in two model element criteria (as described earlier in section 6.4.1, Chapter 6, where criteria 'GF2' and 'SC1' were considered inexact duplicates). This redundancy of criteria could have confused experts during their pair-wise assessments and thus contributed to detected levels of inconsistency. It should be noted though that there were '42' model element criteria and only two criteria were duplicates. The levels of detected inconsistency could not therefore be contributed to this factor alone.

In terms of model structure, expert 'LL' noted a number of criteria placements within sub-categories, where he did not agree with the categorisation of the criteria. Although expert 'LL' was the only expert who raised such concerns, his arguments for alternative placements could be supported by logic and in the main, most (but not all) of the alternate placements he raised could have been accommodated. This 'placement' issue was raised by expert 'LL' after he had received the summary of analysed weight data, and whilst he was comparing the hierarchical model in this study with his own developing evaluation model, not at the time of taking the survey. That aside, given it was feasible to place a small number of criteria (four) in an alternate sub-category, experts may have been subconsciously confused by the placement and therefore inconsistently rated the criteria. It has been shown elsewhere (Brugha, 1998; Belton and Stewart, 2002) that poor structuring of an AHP model can result in respondent confusion.

Despite the degree of rigour used to develop the model, the subjectivity of the act of classification is a weakness in the method that was used. Barsalou (1985) asserts that categorisation in the human mind is not a matter of something being completely in a category or completely out of it. Rather, an item may belong to a category by degrees. Additionally, Anderson and Pérez-Carballo (2001) write: "Variability [in categorisation] appears to be due to those subjective, cognitive, 'mentalist' processes going on in our minds, and the fact that the mind of every individual is different." A demonstration of this observation is an experimental study on manual indexing for Boolean information retrieval systems which showed that the degree of overlap in the keywords selected by two similarly trained people to represent the same document was, on average, no higher than 30% (Cleverdon, 1984).

Recognising that categorisation is subjective, a possible limitation of the study was that experts were provided with the hierarchical model at the time of being asked to complete the pair-wise comparison survey. A more robust approach would have been to introduce an intermediate step and provide the model to all experts for comment prior to this exercise in order to validate the author's arrangement of model elements (independent of the pair-wise comparison exercise). It is uncertain however, if this step would have picked up the possible alternate placements (because no expert taking the survey indicated at the time of studying the model that criteria should be placed elsewhere). Introducing this additional step would also have added yet another task for experts to complete, with the possibility of leading to an even higher drop-out rate than was experienced.

It should also be recalled that experts had already demonstrated a propensity to ‘rate’ criteria inconsistently (independent of the model’s structure), between in-depth interviews and the structured pair-wise comparison exercise. Fifty-four percent of the weights eventually allocated by experts during the pair-wise comparison exercise contradicted their general opinions about the level of importance of a criterion stated at interview. Interestingly, a (non-AHP) study by Einhorn (1974), investigating expert reliability (using *internal consistency* exhibited by pathologists as measured by the expert saying the same thing in similar situations) similarly found a low level of consistency. Einhorn (1974) (perhaps controversially and uncharitably) took this to suggest a low level of expert competence. However, others (e.g., Tversky and Kahneman, 1981) have shown that preferences can change according to the way a decision problem is framed and that the decision context (i.e., the nature of the set of alternatives) can also have an influence.

It is difficult to untangle the nature of the interplay between the ‘structure’ of the model and the various types of inconsistencies detected. Tversky and Kahneman, (1981) explain that rational choice requires that the preference between options should not reverse with changes of the problem ‘frame’. Because of imperfections of human perception and decision, however, they have observed changes of an expert’s perspective, by manipulating variations in the framing of acts, contingencies, or outcomes that has lead to reversals of the relative desirability of options. In this thesis there was an obvious difference between how criteria were discussed in the conversational style of an interview and the bald context of the comparative survey instrument. This difference in frame may have lead to the inconsistencies in ratings detected between interview and the survey exercise.

The order of questions can also apparently affect the answers in direct elicitation protocols such as surveys, as can how ‘hard’ or ‘easy’ the questions are perceived to be, or the technical substance of the questions (Burgman *et al.,* 2006). Additionally, the level of confidence assigned to a given judgment has been found to be affected by the confidence with which they rated previous judgments (Tormala and Petty, 2007). There is therefore a plethora of issues which may affect an expert’s preference rating and it is difficult, even with hindsight, to determine how to control for all of these issues.

Given the levels of inconsistency found (without justifiable explanation) and the lack of expert validation of the model structure (as an explicit step), an informative follow-up study would be to introduce the model hierarchy to a new group of (not so time-poor, and local) ontology experts; divide this group into two samples (randomly) and ask both groups if the model structure (a) addresses criteria of interest and (b) appropriately groups and nests issue of ‘similarity’ from a decision-making perspective. Model development could be performed as group consensus exercises but with recorded observation of each group's activity in developing the consensus model. Individuals within both groups would then be asked to pair-wise compare model elements (as in this study to ascertain model element weights). In this situation, of interest are any differences in model construction between the two groups and the resultant effects that this has on the estimation of local and global weights, between groups. The variability of within group responses could also be calibrated against this study's results. The similarity (or dissimilarity) between the pair of models could be described based on commonality in variables, for which a suite of similarity coefficients are available (Gower, 1985).

Since the outcomes of the AHP method are dependent on the 'structuring' of the decision hierarchy (model), a deeper analysis of any variability in how individual experts might differently frame the model structure would provide valuable insight into the decision-making process for the ontology selection problem.

7.1.2 Methods Trialled For Identifying & Improving Inconsistent Data

Since higher than acceptable levels of inconsistency were evident, and given concerns that additional surveys would lead to survey burden and the potential loss of further participants from the study, several methods for detecting and improving inconsistent ratings were reviewed and tested at the time of conducting the research explained in Chapter 6. These methods were expected to assist experts with reviewing their data during the analyses conducted.

Gower plots appeared to hold some promise for identifying which specific pair-wise comparisons in a respondent's matrix of data were inconsistent. Gower plots are diagrams that result from the singular value decomposition (SVD) of a square matrix and are similar to Biplots which are used widely in multivariate analysis and multi-dimensional scaling tasks (Gower and Hand, 1996). Li and Ma (2006) proposed a method to detect inconsistencies that involves drawing ordinal and cardinal Gower plots of a preference matrix (to recap cardinal errors are related to the values given to alternatives and ordinal errors are problems with the order of rated alternatives). The interpretation of Gower plots, which are constructed from skew-symmetric matrices, such as those found in AHP rely on the fact that when an expert's responses are cardinally inconsistent, the plot yields a set of points $P_1 = (u_1, v_1)....P_n = (u_n, v_n)$ on a straight line that does not cross the origin. Any lack of co-linearity provides a means of detecting "delinquent" comparisons, provided that the data had ordinal consistency to start with. Ordinal inconsistency can be detected by drawing a Gower plot of the tournament matrix of the skew-matrix. A tournament matrix, is a matrix reduced to the values '0', '1' and '-1' by converting a matrix element (r_{ij}) to '1' if $r_{ij} > 1$; to '0' if $r_{ij} = 1$ and to '-1' if $r_{ij} < 1$. In this case points P_1,...P_n should be equidistant from the origin, arranged in a counter clockwise order of preference, within a 180 degree arc. Points off the arc are ordinally inconsistent (Genest and Zhang, 1996).

Li and Ma (2006) then propose a mathematical method for helping experts to adjust preference values with a view to minimizing the change of preference and the degree of inconsistency. Despite this mathematical method being well-described, a test run through the examples presented in the paper, using the formulae supplied, did not yield the same Gower plots as depicted in the publication. Further investigations regarding SVD revealed that although the SVD and Eigenvalue decomposition are well-established and can be computed via state-of-the-art algorithms, it is not commonly mentioned that there is an intrinsic sign indeterminacy that can significantly impact the conclusions and interpretations drawn from their results (Bro *et al.,* 2008). SVD itself provides no means for assessing the sign of each singular vector. In actual algorithmic implementations of SVD, this indeterminacy is inherited so that the individual singular vectors have an 'arbitrary' sign. The actual sign is determined as a by-product of the computations that are used to ensure numerical stability. This determination of sign is essentially the same as assigning the sign randomly and hence the sign has no meaningful interpretation in terms of the data that the decomposition represents (Bro *et al.,* 2008). This is a significant problem for the production of Gower plots since ordinal inconsistency relies on the construction of a tournament matrix, whose construction is sign-dependent. A similar problem was found with another graphical method (Freeman, 1997), for determining inconsistency also harnessing a procedure by Gower (1977) called canonical analysis of asymmetry. Given the flaws of the approaches just mentioned, an ‘R’ software program was ultimately developed, as detailed in Chapter 6, to iteratively change individual pair-wise comparison results, until they reached a level of acceptability. These changes were reflected back to experts and their views sought as to whether these changes were acceptable, before they were used in subsequent analyses. All data in this study had values of CR <=0.14. CR values as high as 1.41 have been used in other studies (Kryvobokov, 2005).

7.1.3 Obtaining A Group Result

In this ontology selection and evaluation research the aim was to obtain a ‘group’ result. Moreno-Jimenez *et al.* (2002) describe three ways in which this is generally achieved using AHP techniques: (i) *Group Decision* where the individuals act jointly by looking for a common decision; (ii) *Negotiated Decision* where each actor solves the problem individually and then the agreement and disagreement zones are analysed in order to reach a consensus and (iii) *Systemic Decision* where each individual acts independently and a tolerance principle is used to look for a way of integrating all the positions**.**

Because it was not feasible to physically bring all experts together, nor could it be engineered so that people could be brought together 'virtually', all at the same time, option (i) above could not be used. Whilst the Delphi approach (Schmidt, 1997) was favoured to arrive at a 'negotiated decision', the additional effort that would be required from experts to participate in an iterative process to converge on a consensus, was judged to be too much of a burden. So, option (iii) was the approach eventually used.

According to Ramanatham and Ganesh (1994) and Forman and Peniwati (1998) the most often used method to obtain the preference of the group is by aggregation of individual priorities using a weighted geometric mean. A weighted geometric mean is used when group members are not considered to be equally important, and their ratings are therefore weighted accordingly. In this study, all experts were considered to carry the same level of importance so a (normalised) geometric mean was used without weights. Escobar *et al.* (2004) have demonstrated that the 'group' inconsistency is at least as good as the worst individual inconsistency for this aggregation approach. Since, no individual expert exceeded a CR of 0.14, the 'group' result should not exceed this value.

General Expert Disagreement

Because there was considerable heterogeneity in how experts rated particular model elements (see Chapter 6, section 6.3.4), explanations were sought to rationalise the more divergent opinions, regardless of the computed group result. Overall, all of the groupings of opinion that were evident surrounding the pair-wise comparison data for each dimension and each sub-category, were unexplainable on the basis of the various types of stratification discussed in Chapter 6. The only explanation left, given the data to hand, was that there were simple differences of opinion between the experts in relation to the weight that they would personally apply to criteria in a decision-making scenario.

Einhorn (1974) has argued that consensus (i.e., between-expert reliability) is a necessary condition for 'expertise'. His reasoning, however, seems at odds with the many reported situations in which experts do hold divergent views (whilst still being considered competent, as judged by their peers). Meta studies by Shanteau (2000) demonstrate such situations and he argues that experts more frequently diverge in views when the domain that they are working in is immature and dynamic. Semantic-enablement of scientific data infrastructure could be considered a very immature domain which could account for the variation in expert views expressed in this study.

Burgman *et al.* (2006) have produced a good summary of the literature about why expert estimates of quantities can often differ. Any of the issues that they list could conceivably apply to the experts in this study. These issues include:

- *Perception and memory*: Judgments and decisions are influenced by what we know and what we can remember about what we know. When we elicit a judgment or decision from an individual, and expect that decision to be informed by the prior knowledge, we are at the mercy of their memories. There are individual differences in memory function and there are neurological limits to capacity to store and retrieve information.
- *Framing*: A framing effect (discussed previously) occurs when a change in the presentation of a choice influences choice behaviour, even when the objective characteristics (outcomes or probabilities) are not changed (Tversky and Kahneman, 1981).
- *Heuristics and biases*: People naturally reduce the complexity of problems to make them more manageable. People try to simplify their judgments by applying 'rules of thumb', also known as heuristics. These are mental short cuts that reduce the cognitive burden associated with decision making (Shah & Oppenheimer, 2008). There are many types of problems associated with heuristic-based judgements, for example, often people, including experts, judge the probability of an event by how easy it is to recall an instance of such an event regardless of its true probability of occurrence. Cognitive biases are thinking patterns based on observations and generalizations that may lead to memory errors, inaccurate judgments, and faulty logic (Evans *et al.,* 1983).
- *Motivated reasoning*: Often, people distort interpretation of decision criteria and the evaluation of information to justify a predetermined choice, even when engaged in evaluating information in an attempt to be objective (Phillips 2002).
- *Values and attitudes*: Whether we are conscious of it or not, most decisions reflect value systems and attitudes. Values and attitudes are usually elicited using interviews, questionnaires and surveys. Gregory *et al.* (1997) asserts that the difficulty with survey techniques is that their design needs to closely mirror the complex cognitive and social processes shaping the participant's responses. The analyst needs to consider how the framing of each question affects responses, and how complexity and uncertainty of the issue at hand can cloud judgement.

Studies by Wilson and Schooler (1991) also demonstrate that it is possible that in some cases, experts can 'over-think' issues. Conventional wisdom indicates that we should think "hard" about our options when faced with a difficult decision. By devoting attention and conscious thought to a difficult decision, the prevailing view is that one can carefully consider and weigh the various options and choose the option that best matches one's goals. The work of Wilson and Schooler (1991) challenged this assumption with the results of studies that found when participants were asked to think about reasons for a decision that they had made, they made apparently worse decisions than participants who did not reflect on their reasons for a decision. In other experiments by Wilson and Schooler (1991) their conclusion was that thinking "hard" about reasons for judgments, or decisions may lead to judgments that depart from expert or consensual judgments. If the results found by Wilson and Schooler (1991) are generally applicable, this would draw into question many current techniques aimed at facilitating multi-criteria decision-making (not just AHP).

Patterns In The Data

Given that differences of opinion are likely to occur, in cases where one or two individuals had a different pattern of weights than the broader group (and if they were not constant outliers) these differences were not considered to be an issue and speculation about why these differences might exist would not be fruitful. The group mean for these types of situations was therefore considered a reasonable representation of "group-think", particularly in cases were outliers were removed before the geometric mean was calculated. But there were some situations where the group split into two and sometimes three groups, showing very distinctive trends. It is this type of variability that intuitively seems to warrant further examination and was particularly the case for responses provided for opinions expressed for sub-categories in the 'Usability' and 'Governance' dimensions.

The 'Usability' dimension had two sub-categories: 'Ease Of Application' (UE) and 'Sustainability' (US). In weighting these two sub-categories, experts were opposite in their responses. Experts 'RA', 'LL' and 'KS' rated 'Sustainability' factors much higher than 'Ease-of-Application' issues. The other five experts disagreed. It is the polarisation that is interesting here, in that there are not just two different trends but both groups felt that one set of factors was much more important than the other, yet there was divided opinion. The experts that chose 'Sustainability' over 'Ease Of Application' all come from very different backgrounds, have different self-described skills and have played differing roles in communities with respect to ontology development. Expert 'LL' and 'RA' had worked together on common projects, but not with expert 'KS', so the chance that all three had arrived at their opinions through cross-fertilisation and common experience, can be largely discounted.

Perhaps the difference in expert views here could be attributed to differences in an expert's cognitive style and personality traits. Burgman *et al.* (2006) note that this is a studied source of expert variance in opinion. It is plausible, for instance, that those experts who believed it more important that an ontology fit easily into existing applications, might also consider that by demonstration they could convince their community to adopt the ontology (because of its ease of application) and if the originating source of the borrowed ontology failed to maintain the source ontology into the future, it would be managed instead within the adopter's community for as long as it remained useful (i.e., they had a pragmatist viewpoint).

It certainly was the case that some experts stated at interview that it was often easier to take an existing ontology or ontology fragment, modify it if necessary to meet their own community's needs, without trying to influence its ongoing maintenance or its update at origin (hence forking from the originating source). For example, one expert stated: "*So, usually what I end up doing is picking up 3 or 4 of these things, taking what I think are the good bits out of them, merging them together and leaving the rest*". Other experts, 'LL' and 'RA' included, appeared more concerned with ensuring that re-used ontologies were identified as such and that their ongoing governance and use was reflective of this fact. This attitude is perhaps suggestive of a more utopian view of interoperating communities of practise and a set of behaviours commensurate with making this view a reality (which could be considered less pragmatic, but perhaps more visionary).

'Ease Of Application' ended up being the second highest (16.8%), weighted sub-category at the local (sub-category) level using the group results in the model, after 'Application Relevance' (24.5%). In contrast 'Sustainability' didn't rate particularly highly in a relative sense (5%). Clearly, therefore, most experts (in this study) thought it more important that an ontology be readily usable and compatible (or able to be integrated) with their current systems than whether the ontology they had chosen would be accepted by their community and sustained, or maintained into the future. In the main, most experts also thought that issues such as "the degree of concept coverage" and an ontology's ability to "meet use-case goals" (i.e., criteria in the sub-category 'Application Relevance') were of greater importance than either 'Ease Of Application' or 'Sustainability' factors. These views are inconsistent with other reported observations. For example, Obrst *et al.* (2007) have stated that "the ultimate evaluation of an ontology is in terms of its adoption and successful use, rather than its consistency or coverage. The Gene Ontology, while clearly impoverished in many representational aspects, is a fundamental success story".

The 'Sustainability' sub-category itself also had a polarising affect with respect to how experts rated the 'within sub-category', evaluation criteria. 'LL', 'KS' and 'WD' ranked (US2) -"whether they could convince their community to use an ontology" the most important criteria over (US3) – "Whether or not an ontology would be used and sustained into the future", followed by (US1) – "the size and scope of the current user-base”. Experts ‘DH’ and ‘JH’ ranked the options as US3 > US1 > US2 (as opposed to US2 > US3 > US1, the ordering by ‘LL’, ‘KS’ and ‘WD’). At this point a general observation (by the author) is that once you get down to the criteria level (i.e., the within sub-category comparisons), there are circumstances where it genuinely does become quite a difficult exercise to weight options comparatively, and this situation is one of them. Bearing this in mind there is some sympathy with the views of the expert cited earlier, who stated that they were being asked to compare “chalk to cheese”, in that the preference for one criterion over another may not be easily discernible. It is easy to see that expert views could be divergent on such issues. These ‘Usability’ model elements are mutually exclusive criteria but ones which mix technical and sociological matters, two of which are speculative and one of which is concrete. Interestingly, the most directly, immediately measurable (concrete) criterion (US1) is not the highest ranked by either group.

Despite these identified comparison difficulties, the AHP method, which mandates pair-wise comparisons from the leaf nodes of the hierarchy right through to the goal node, forces an expert to mentally weigh alternatives at differing conceptual levels of granularity, which ends up providing the expert with a much deeper understanding of their own rationale for decision-making (Brugha, 1998). But sometimes a preference for one thing or another may not be easy to weigh. Despite this observation, Williams *et al*. (2007) found that tools such as AHP can help decision makers develop a better understanding of the essence of a decision problem and therefore help them to reduce logical errors (especially when the information load is high).

The ‘Governance’ dimension also split the expert cohort in interesting ways. Expert ‘RA’, ‘JH’ and ‘WD’ rated ‘Framework’ factors ahead of ‘Community’ and ‘Behavioural’ factors. ‘Framework’ factors are those relating to how a community organises itself (for e.g., the OBO Foundry is a good example of an identifiable organising framework). ‘Community’ issues are those factors surrounding cohesion as a community-of-practise and other characterising features such as the community’s mandate and its level of backing. ‘Behaviours’ are more about the practises that a community exhibits (e.g., how it fosters trust, credibility etc).

‘MH’, ‘LL’ and ‘KS’ believed that behaviours were most important (‘DH’ and ‘CA’ were considered outliers, because they neither agreed with each other or either of the other two groups in terms of trend). This, again, is a difficult comparison for experts to make, since a community will thrive and survive based on the behaviours exhibited by its participants, but without an organising framework it can be more difficult for a community to make efficient use of its citizens to produce community goods. The interesting aspect of the results for governance is that experts ‘MH’ and ‘WD’ are both part of the OBO Foundry, and yet they were in opposite camps with respect to their weightings in this dimension. Being part of the same community and subject to similar governance (as well as holding similar roles within that community) doesn't appear to guarantee unity of opinion.

Notably 'LL' and 'KS' were of similar opinion again (i.e., as they were for rating the 'Usability' dimension). This could be more evidence for some similarity of cognitive style and personality traits between these two experts, or it could just be coincidence.

'Governance' issues were relatively frequently discussed by experts during interview but the 'Governance' dimension failed to rank highly in terms of importance during the pair-wise comparison exercise. Possible differences in perspective between experts who were interviewed but who did not go on to undertake the pair-wise rating activity and those who did are discussed in more detail later in section 7.3.1.

In summary considerable variation and inconsistency in expert judgement has been well noted in the literature (e.g., Dawes, 1971; Armstrong, 2001). Shanteau (2000) considers that this is an outcome of the contexts in which experts work. He claims that often single-point, optimal solutions do not exist in the complex environments in which most experts habit. Therefore experts operate "as if they have flat loss functions for deviations from optimality. They see small deviations as having minor consequences". He contrasts this with the research fraternity, where he says "in comparison, researchers often operate as if they have steep loss functions. That is, they view any deviation from optimality, no matter how slight, as having large consequences".

Shanteau (2000) also asserts that experts often work in realms where the basic science is still evolving and that is certainly true for the field of semantic enablement of data infrastructure. He argues that the "best answers" in emerging fields are soon obsolete and that new knowledge will likely provide better solutions tomorrow. "Once a domain has advanced to the point where all issues are resolved, there are few disagreements among experts because there is nothing to argue about. When a field has developed to that degree, however, the answers are known and agreed upon. Thus, total agreement among experts is an indication that there is no longer much of a role for experts to play in that domain." On the basis of the expert inconsistency detected in this study, if Shanteau (2000) is correct, the topic of ontology selection and evaluation has some time ahead of it before it could be considered a mature domain of activity.

7.2 Selection and Evaluation Methodologies

Apart from expert opinion about what evaluation criteria are of importance in selecting and evaluating an ontology for reuse, of interest in this thesis was the question of which 'methods', as opposed to ‘criteria’, were being employed by experts to perform the selection task. RQ1.2.1: asked *“....are selection and evaluation methods* [used by experts] *consistent with those reported in the literature* ?”. Information concerning this aspect of ontology selection and evaluation was derived both from the screening survey (outlined in Chapter 6, section 6.2, particularly Table 6.11) and from in-depth interviews.

In the ensuing discussion, the focus is on ‘methodologies’. In borrowing from the discipline of engineering Obrst *et al.* (2007) suggest that robust ontology evaluation methods should consist of:

- A formal, verifiable science base.
- Tested theories that allow prediction.
- Defined units of measure.
- Well-defined engineering practices.

Given these principles it is believed that many of the published methods, described earlier in Chapter 2, would not necessarily qualify as a robust method, mainly because they fail the second of the dot points of Obrst *et al.* (2007) and sometimes, as discussed in Yu (2008) and Vrandecic (2010), defined measures are absent. Few, if any, of the methods discussed in Chapter 2 offer the power of ‘prediction’ and if ontology evaluation methods are to be judged by such principles, at this immature stage of the ontology discipline’s evolution, no ‘selection and evaluation’ practise could be labelled as a ‘method’. A suggested alternate definition applied in this thesis, for an approach that would be considered an acceptably robust evaluation method, is where:

- There is a logically constructed and repeatable process.
- Process outputs (results) are documented.
- The ontology facets being evaluated are clearly identifiable.
- The facets being evaluated are measurable by a specific indicator.

7.2.1 Precis of Expert Data

Before moving to specifically address the issue of ‘methodologies’ a brief summary of the data concerning experts who participated in the selection and evaluation phases of the research is provided, to remind the reader about the main findings that emerged in Chapter 6. This will help put the methodology discussion in perspective and in context, and similarly for discussions in the subsequent section regarding governance.

Recall that the two main types of ‘expert’ that participated in the selection and evaluation component of this research were: people with ***domain expertise*** who had *acquired ontological skills* and ***ontologists*** with *varying degrees of domain expertise*. From an analysis of the data there didn't appear to be any obvious correlation between an expert's stated skill type (e.g., domain practitioner with acquired skills vs professional ontologist) and the number of mentioned evaluation criteria, or the number of governance-related issues raised by the expert. Whilst three of the professional ontologists ('PF', 'MH' and 'PM') were contributors of a relatively high number of criteria, two others ('RA' and 'CA') raised relatively few, and 'LL', the remaining professional ontologist raised an average number of criteria. This result was somewhat surprising since it was anticipated that the self-characterised, professional ontologists might individually raise more criteria and issues. It is worth noting, however, that 'PF' contributed the most criteria and remained the highest contributor, after governance issues (captured and coded separately) were added. 'PF' was also amongst three other experts ('RA', 'LL' and the domain-skilled 'SC') who were active in three or more science disciplinary areas. Fifty-percent of the professional ontologists (i.e., 'RA', 'LL' and 'PF') had applied their skills across multiple disciplines and the remainder were primarily single discipline (i.e., biologically) focussed.

Regardless of the number of evaluation criteria contributed during interview, the professional ontologists 'PF', 'MH', 'PM' and 'LL' all provided relatively deep insights across the range of questions raised at interview, and 'LL' was also highly active in providing valuable feedback on the structure of the hierarchical evaluation model, the weights applied by experts to the model elements and on issues of evaluation metrics. It is not asserted that these individuals were the only ones that provided "deep insights" but they consistently raised informative issues regarding "development methods", "governance" and/or "matters of practise", which demonstrated strong ontological capabilities exercised in community development settings.

The study had relatively few experts in total (14 inclusive of the late addition of expert 'KS'), but still achieved a good spread of disciplinary representation (i.e., three or more experts per discipline with the exception of the energy discipline, which had only one expert, i.e., expert 'LL'). This good disciplinary spread was mainly due to the fact that five experts had worked within more than one disciplinary area. The discipline most highly represented was biology. Disciplinary representation, however, was significantly reduced by the time that the evaluation criteria weighting exercise was conducted, due to the low (approximately fifty-percent) participation rate. For this exercise, the oceanography discipline was represented by 'KS', who joined the study after interviews were concluded. However, interviewed experts representing the oceanography discipline were provided with an opportunity to comment on the weights applied by other experts towards the end of the weighting exercise, so the oceanography sector was covered in the study to its end.

Most experts (9 out of 14) had worked in team situations, where teams were considered to be more than two people cooperating to achieve a common outcome. The norm in team situations was for smaller (< five people), rather than larger teams. ’PF’, ‘LL’, ‘RA’, ‘DH’ and ‘SC’ had all worked in projects that had drawn together larger teams (> 10 people) to work on ontology development. ‘LL’, ‘SC’ and ‘RA’ had all, however, worked on the same project, so only three different, large-scale projects which incorporated > 10 people were apparent from the data. Most (individual) projects cited did not have this level of people investment.

Of the experts interviewed nine (‘PF’, ‘DH’, ‘LL’, ‘JG’, ‘RL’, ‘RA’, ‘RH’, ‘SC’, ‘MH’) had played all three roles described in Chapter 6 (i.e., Leader/Driver/Initiator; Maintainer; Specialist (ontology and/or domain) Advisor)) and the remainder had primarily played the roles of ‘Maintainer’ and ‘Specialist Advisor’.

Fifteen different types of communities were mentioned by participating experts. Of note is the relatively high degree of overlap and the potential for cross-fertilisation of skills and practises between disciplines due to the number of experts involved in multiple communities.

Most communities mentioned had international affiliations and had substantial memberships. It was difficult to quantify community sizes since experts described their communities using different methods (e.g., some tried to estimate individual participants, others quoted size as a function of the number of institutions involved). The size of the membership, however, in most cases did not correlate well with the number of people actively working on ontologies within the community. Community membership is best simply interpreted as being the size of the broader community ready to be consumers of the ontological products developed. The number of people actually working on ontology development within communities appeared to be more of a reflection of either community maturity (i.e., age), or the presence of ‘funded/institutional backing’, than it had to do with actual community size. The relative immaturity (i.e., 6 years or less) of the semantic activities of most of the communities studied is also worth noting.

7.2.2 Reported Methods With Academic (Or External Community) Origins

It was reported by Paslaru Bontas-Simperl and Tempich (2006) that only a “small percentage” of practical ontology-related projects, that they had studied, followed any systematic approach to ontology-building and even less committed to a specific methodology. It was therefore speculated in this thesis that the ontology evaluation methods described in the literature, were possibly not those methods being used by practitioners actively developing ontologies in scientific domains, where the focus was on data exchange and data integration. Indeed, only two out of the thirteen experts (or

15%) who both completed a Screening Survey and who participated in interviews, provided citations to published ontology evaluation methods with academic origins and claimed to have used these methods.

However, in comparing the evaluation criteria listed in the expert-grounded evaluation framework with those criteria found in the literature, there is a high level of overlap (refer to Chapter 2 for a review of literature-based evaluation methods and associated evaluation criteria). A number of the criteria cited by experts, for example, were associated with evaluating an ontology's expandability; modularity; conciseness, clarity; coverage; organisational fitness and completeness. This suggests that even though experts claimed their evaluation 'methodologies' did not have origins in the academic literature, it is highly possible that individual evaluation criteria that have been 'picked up' and used by experts, did have their genesis in an academic context.

Vrandecic (2010) claimed that most evaluation criteria cited in the literature focussed on issues of semantics and context (refer again to discussions in Chapter 2) and less so on matters of vocabulary, syntax and structure. Whilst experts in this study did cover issues of semantics and context, there was also significant focus on vocabulary and syntax-related criteria. For example, see evaluation measures for criteria 'SS5'; 'SC1'; 'FRA3'; 'UE3'; 'MC3'; 'MU2'. This was an interesting outcome since Vrandecic (2010) claimed that there were relatively easy-to-perform types of checks for these latter-mentioned criteria. It would appear that these types of checks are being performed in practise despite not rating highly in the literature.

Of the two experts (i.e., expert 'LL' and 'MH') who did claim to be using published methods of academic origin, expert 'LL' cited two methods, one from Gangemi *et al.* (2005) and the other from Alani and Brewster (2005), both of which were also mentioned in the literature review in Chapter 2. Expert 'MH' provided three references, two of which were relatively general papers on OBO (Smith *et al.,* 2007) and the anatomical ontologies that are characteristic of model organisms (Bard, 2005) and the third was a paper by Donnelly *et al.* (2006) on a formal theory for spatial representation and reasoning in biomedical ontologies. This last paper presented a formal theory of "parthood" and "location" relations among ontology classes and explained how spatial information included in biomedical ontologies is often ambiguous and therefore the possibilities for implementing consistent automatic reasoning, within or across ontologies are limited.

The Donnelly *et al.* (2006) paper is an extremely detailed paper about 'mereology' (i.e., the relations of part to whole and the relations of 'part to part' within a whole) in the biomedical domain. The paper urges that different relational terms in a biomedical ontology be linked to a formal theory of spatial relations. Although this paper is not about a formal evaluation method, per se, it is possible to see how the issues raised in the paper, that are connected with mereology in biological ontologies, could be used as a source of evaluation criteria (particularly for assessing the adequacy of anatomical type ontologies). Of the three references cited by 'MH' this last paper was the only one with any distinguishable value with respect to defining elements pertaining to an evaluation method.

Two other experts (i.e., 'PF' and 'PM') indicated they had used methods devised by other communities. The reference provided by 'PF' to a "Semantic Web for Earth and Environmental Terminology (SWEET) ontology web site" (Raskin, 2011), did not hold any information on an ontology evaluation framework. A subsequent review of literature associated with the SWEET ontology, additionally did not uncover any published ontology evaluation methodologies. 'PM' referenced the National Center for Biotechnology Information (NCBI (NCBI, 2012)) as the source of an evaluation method, without a specific citation.

Even if we assume that the NCBI, which is a mature and credible source of a significant amount of research involving ontology development, had published a particular evaluation method that 'PM' had used, it still remains that 77% (10 out of 13) of the study's experts were using evaluation methodologies purportedly devised within their own community (and they did not consider these methods to stem from academic origins). The next section, therefore, examines what methods (if not those cited in the literature) were being applied by the experts in this study.

7.2.3 Descriptions Of Community-based Methods Used

All expert participants claimed to use some type of evaluation methodology and most of these methods purportedly had community-based origins. Despite this claim, few experts, citing community-centric methods could articulate the evaluation method used, or point to explanatory resources.

OBO Foundry Experts

For those experts who were part of the broader OBO Foundry community there was a relatively commonly stated approach (see Figure 7.1). In the main, evaluation occurred within two types of activity: within purposive workshops designed to address a new and significant (ontology) need or to address problems with existing (ontology) approaches; and as part of a continuing cycle of ontology maintenance and update (i.e., ongoing ontology curation). The various OBO ontologies cover different life-science sub-disciplines and as such involve separate sub-communities, for example various groups conducting research on different model organisms.

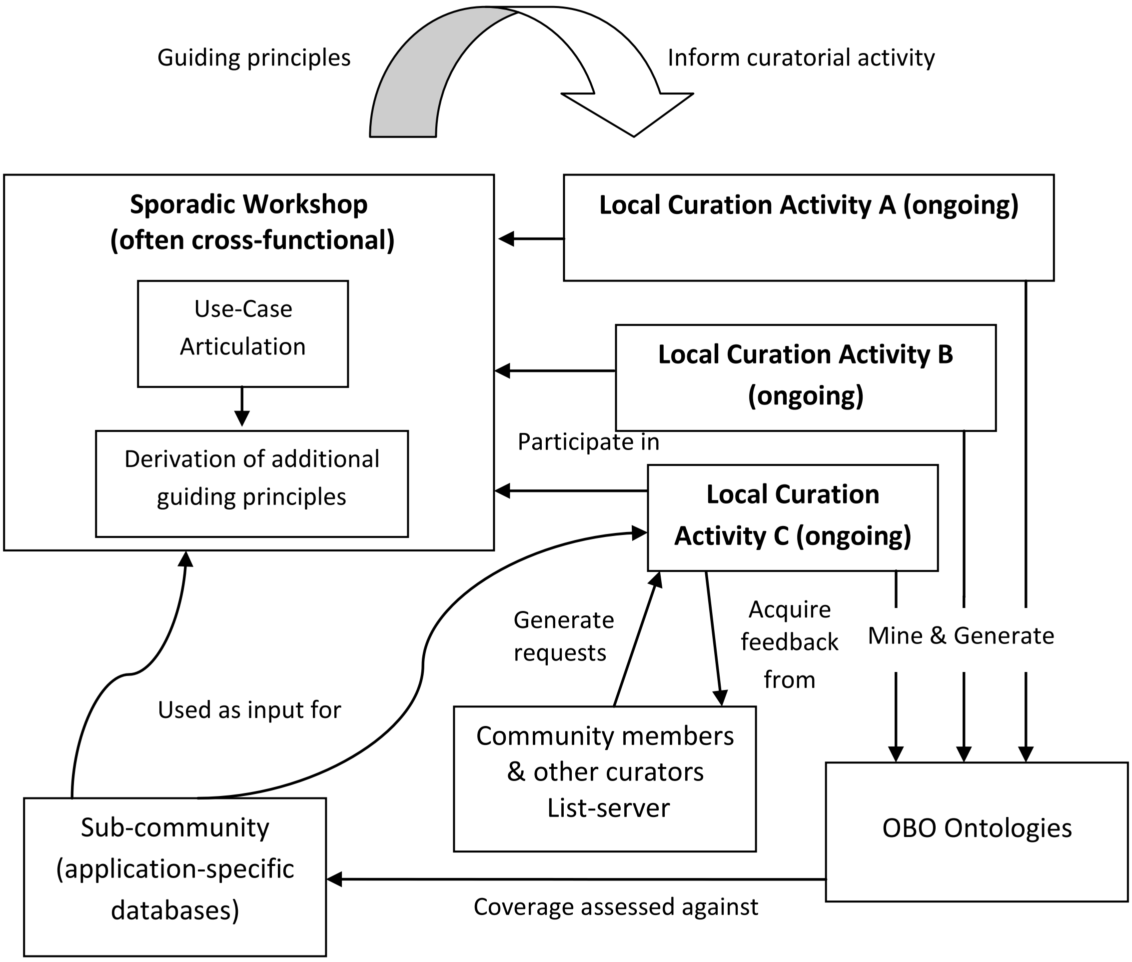

Figure 7.1 OBO Community Development and Evaluation Methodology Overview

When asked how they would first identify an ontology for possible re-use, OBO-affiliated experts said they would primarily use the OBO Foundry web site or the NCBO BioPortal to identify ontologies covering similar abstract, or domain concepts; harness their own collegiate networks; use community-based list-servers and as a last resort, trawl the Web. These OBO-centric specialists felt that they would trawl the Web rarely because the result was often thousands of hits for information of interest, with no really effective way of weeding out irrelevant material. Expert 'PM' stated:

> *"Lets see. I would look for ontologies that cover similar or an abstraction of the domain I'm trying to model. So for taxonomy I would look at other taxonomies and more or less follow their methodology.....*
>
> *So one thing to do in our case would be to go to the NCBO (BioPortal) which is a resource that I…..and I have looked for material because as part of the OBO community that's the place I would go to first. Their repository is basically a superset of what is in the OBO Foundry. There is a lot of material in there and it may not all be compliant but it is a source to look for biological ontologies that I would think of first because its large, its pretty well established, its not perfect .. sometimes our ontologies on there are not up to date but we*

have been able to work with them to get it up to date. They seem to be making an effort to make this resource work out for the community".

Expert 'PM' further explained:

"Yeah, and OBO (um) there's also a number of general OBO lists that serve for general discussion and so that could be another way that you could say 'does anybody know anything about how to represent this or write an ontology about that?' And so you almost always would get some response from somebody out there."

These OBO-centric experts primarily sought out potential sources of material (ontologies) that were already OBO compliant, or if not OBO compliant, at least built within the OBO community. The OBO community has put considerable effort into defining principles (OBO Foundry, 2012b) and guidance on a wide range of aspects related to ontology development, covering facets such as openness; naming conventions; format; relations; URIs; versioning; textual definitions and documentation to name just some of the principles. By staying within their own community of existing ontologies (for reuse purposes) the level of evaluation required is minimised because theoretically (at least) most ontology developers should be conforming to the same OBO guiding principles. The tooling developed within the OBO community is also another incentive to reuse existing OBO ontologies because the tools are geared to manipulate OBO conformant ontologies. Heavy reliance on OBO tooling to manipulate and manage OBO-linked ontologies was raised several times by the OBO experts.

A driver for the various OBO sub-communities to ensure conformance with OBO principles is the broader development of integrating ontologies such as the Ontology for Biomedical Investigations (OBI), which will describe biological and clinical investigations. This ontology will include a set of 'universal' terms that are applicable across various biological and technological domains as well as encompass some domain-specific terms that are relevant to given domains. OBI will eventually support the consistent annotation of biomedical investigations, regardless of the particular field of study and will also be able to represent the design of an investigation, the protocols, instrumentation and the material used, the data generated and the type of analysis performed on it. In order to use this ontology within the Foundry, individual sub-community developers will need to stay conformant with the OBO principles.

Because of all of these types of factors, the OBO experts considered it far easier to adopt, or adapt existing OBO ontologies to suit their purposes rather than re-use ontologies from elsewhere. Although expert 'PM' mentioned borrowing "taxonomic rank" concepts from the TDWG community (to use in an OBO ontology) it was primarily for mapping purposes, rather than because he believed that any additional intrinsic value was held by the TDWG ontology, from which the rank concepts were borrowed. It is worth noting again that there appeared to be no collaboration on ontology development between OBO sub-communities (sampled in this study) and those of other biological groups also sampled in this study (despite similar ground being covered).

In common with other experts in this study (e.g., ‘PF’, ‘RH’, ‘JH’, ‘DH’, ‘PM’, ‘CA’, ‘JG’, ‘RA’), the OBO-centric experts harness use-cases (and to some extent competency questions) as benchmarks for evaluating ontologies. In the OBO community these use-cases and competency questions are usually developed via community workshops (and much less frequently and less formally by individual ontology curators). Workshops appear to be scheduled when relatively large gaps in an ontology are uncovered and the structure of the ontology may need significant adjusting (not just simple amendments to the number, or type of included terms, which mostly occurs at the local curatorial level). See again, Figure 7.1.

Whilst none of the three OBO affiliated experts specifically mentioned the term “competency questions” they all implied the use of competency questions and a scan of OBO material on the web (Sansone, 2007; Smith, 2006) relating to individual ontology development and evaluation confirmed that competency questions and use cases are part of the OBO evaluation approach.

There was a strong impression resulting from interviews with the OBO experts that the majority of their evaluation approaches were highly inter-woven with their ongoing ontology development activities, which were incremental and driven by a constant need to gradually make a particular curated ontology better satisfy domain requirements. For example, ‘WD’ stated:

> “.... *And we have a group of about um, six curators. So these are people who are actually using the ontology and you know, they‘re really the people that are making the requests and I do that too. That’s also part of my job is curation. So there are probably six or seven of us that are actively using the ontology and we need a new term or we see that we need a new relationship, a new definition, then it’s requested. ....So the development really comes from the use, the use in this case.”*

Similary ‘MH explained:

> *“The second source is um our curators and we have nine curators here at...as they’re reading papers they spot terms that we don’t have or that have the wrong definition...*
>
> *...or have the wrong information on them and so then they’ll submit an internal request for that term ..”*

Evaluation also appears to involve a heavy emphasis on terminology. In order to assert what terms and definitions are required, the community of data providers and/or potential ontology users are often consulted to obtain lists that can then aid the evaluation process. Very often an ontology’s concept coverage is evaluated against databases that are already in use by model organism or biomedical communities. Using databases as a source of concepts against which an ontology's coverage is measured was also commonly used in other disciplines (particularly Oceanography, Solar Terrestrial Physics, Atmosphere and Meteorology disciplines). The more long-lived (mature) the community is, the greater the propensity there was for the existence of large databases of accumulated vocabularies. For example expert 'PF' (not an OBO expert) also indicated:

> *"Same methodology in their community: they've got a very similar situation where, um again, we worked on the use-cases. They had the extensive vocabulary because they have over many, many years evolved a data base schema and tables which represent their catalogue. ...So they had the basis of a vocabulary but of course, no structuring, no formal relationships, lots of duplication."*

In summary, the uniqueness of the approach reported by the OBO-centric experts lay in the scope of the ontologies that they were generally prepared to consider in their evaluation method (i.e., they confined themselves predominantly to existing OBO community ontologies for re-use purposes) and they had a strong existing set of criteria (which were well documented) against which to benchmark their ontologies. The drive to create orthogonal ontologies (i.e., ontologies that re-use term definitions that others have already created, rather than duplicate terms already in other ontologies) is also a significant differentiator.

The OBO community experts clearly undertook ontology evaluation, potentially against a wide range of benchmarks, but there wasn't evidence from interview of a "repeatable process", one that each expert regularly practised to evaluate ontologies. To some extent this observation is supported by recent work performed by Ghazvinian *et al.* (2011) which analysed the degree of re-use and term overlap between OBO ontologies. Ideally, term overlap (i.e., where two ontologies define similar terms independently), should be low given the goal of 'orthoganility', but results of the Ghazvinian *et al.* (2011) review reveals otherwise. In their study, among the 53 candidate ontologies that were analysed, only 2 ontologies were orthogonal within the OBO Foundry, having no overlap with other candidates. Additionally, only 30% of ontologies re-use terms, while 96% have terms overlapping with other ontologies. Ghazvinian *et al.* (2011) therefore concluded that these statistics indicated that the vast majority of the ontologies that currently overlap have not yet adopted any measure of term re-use. It is argued here that if a regular process of evaluation (aligned with the OBO principles) was occurring, term overlap would be much lower.

Experts In Solar Terrestrial Physics and Oceanography Communities

Although mentioned by the OBO experts, the application of use-cases was much more overtly mentioned for evaluation activities in the Solar Terrestrial Physics and Oceanography affiliated communities. However, still only two experts ('PF', 'JG') in these communities also mentioned, or implied the use of competency questions. A use-case-centred method promulgated by expert 'PF' involved gathering a small group of domain practitioners together and soliciting use-cases from them as the very first step in an ontology development and evaluation process (encompassing re-use scenarios). The exercise was structured through using MS-Word document templates. Expert 'PF' explained:

> *"We actually put a lot of effort into elaborating on the use-cases. We have a document that we use, a document format which is basically the Wikipedia .. if you have use cases on Wikipedia, it gives the sections. We put it in a WORD document, and we get it all filled in including the data sources and then all the actors and we elaborate on that. We use that as our working document. Because the use-case contains the semantics, the nouns and very phrases, we get them to pin them down very precisely about what they actually mean, including what assumptions have been made, um, you know, stages where there are, you know, implicit inferences in the use-cases ,Oh, if you want to do this, then you have to know this. But you didn't write that down in the use case. But that's okay. We can infer that for you and so it helps us in looking for where you need inference. Where is a simple query, all those things."*

Having worked with a small group, the information is taken away and ontologists (without domain practitioners) then mock-up a 'concept map' of the main classes and relationships that need to be modelled. At this point the ontologists also start to look for, and assess, other ontologies for re-use purposes. The concept maps are brought back to the small domain group and worked on again to check validity. The domain group are not initially exposed to any form of ontology mark-up. Ontology evaluation (for the purposes of reuse) is performed by ontologists, rather than domain practitioners. Expert 'PF' provides some insight into how potentially re-usable ontologies are initially evaluated:

> *"...So we um, we use the – we specifically throw them (ontologies) into Protégé or Swoop and run the validators. Uh, we run the inference checks over them, using Pellet or something like that. And then we have a series of tools just to make sure that they actually work, especially ones that have instances so we have some students that have instance evaluators, um who will typically run some queries across them, we'll and just actually do some engineering level routine checks on them.*
>
> *...Then, we'll actually pull them up into CMAP, we'll read them into CMAP and visually take a look at them. Um, and do a vetting session because it allows you in a much easier way to visualise what the ontology really expresses and it also gives you the level of completeness in the ontology, so for example, um, it shows up sort of where parts are disconnected where there are completely independent concepts that are not related to anything. It'll give you null basically. So at that engineering step, um, we're preparing for doing the evaluation, and we go down to that level."*

In the scenarios described by 'PF', if existing ontologies can be used to create the domain (or application-specific) ontology, ontologists will either adopt the whole ontology (if it meets most of their criteria) or create "mappings" through property assertions on namespaces between the new ontology and the one with re-usable components. The latter is usually the norm. Once an ontological model has been established by ontologists, a small group of domain practitioners is brought back into the activity again and the ontologists start to expose the small group to a developed (ontological) model. However, the model is exported back out of an ontology modelling tool such as Protege and a concept mapping tool (such as CMAP) is used instead for visualisation purposes so that domain practitioners do not have to understand the nuances of the ontology modelling language.

The preference in the approach mentioned by 'PF' appears to be to map to existing terms (and presumably properties) so that the new ontology is an extension of the old. There is also a commitment (at least by the initial facilitating group of ontologists) to provide feedback to the curators of the re-used ontologies (e.g., offering them editorial suggestions or extensions).

The development (and evaluation) methodology described during interview by expert 'PF' has been refined over time through use in multiple, cross-domain situations. Because it is a highly facilitated process and the skills required to "facilitate" are not easy to acquire, the method has been translated into a university course. Expert 'PF' explains:

> *"We taught last fall, called semantic E-science which is basically teaching this methodology and so we've spent a whole class - 3 hours - on developing use-cases, facilitating them, how you extract during the modelling, how you get the modelling going, how you help the knowledge modellers get the information they want, how you keep everyone engaged. It's a very social and dynamic exercise. We decided to teach a graduate course on it so we could figure out how we should write it down."*

Similar to the OBO-centric experts 'PF' did not find Google or Swoogle (a semantic search engine) of any particular value in finding ontologies for re-use, however, a web tool called Twine (Wissner and Spivack, 2009) was considered of some utility. Expert 'PF' and his teams were clearly more amenable to locating any type of ontology that might meet their need and from any source.

> *..Increasingly now we are using TWINE..So TWINE is I mean literally, it is meant to be.. threading related strands of information together to form twine as in strings form twine. And the semantically based uh...so it's a richer tagging environment where people can go and tag things according to more well-defined higher categories like semantic web, ontologies, Web 3.0, all of these things, and what you tend to find in there, is, among those broad categories, you tend to find uh, applications of semantics and the underlying ontologies to go with them."*

The issue of not being able to locate ontologies that could be valuable and which could be evaluated for re-use was a significantly recurring theme during conversations with all experts (no matter what their affiliation). For example, expert 'JG' stated:

"So finding ontologies is a pathetic state of affairs right now. There is... it's terrible."

The problems in locating ontologies identified by experts in this study closely match those problems reported by others (e.g., Zimmerman, 2010) for would-be ontology developers (wishing to practise reuse) which include:

- Ontologies defining the domain of interest simply don't exist.
- They exist but are difficult to find because they have been developed by small groups for experimentation, and they lack advertisement.
- They exist and can be found, but they are of poor quality, not complying with any standards or best practices.
- They exist and can be found but there are too many, of mixed quality, and it is difficult to assess which ones are appropriate for a specific use-case.

Clearly these issues point to the need for a significant improvement in the number and range of readily accessible ontology repositories, such as the ontologically-grounded Feature Catalogue developed earlier in this thesis, which itself should form part of a federated interoperability infrastructure for ontology metadata such as that being developed by the Open Ontology Repository (OOR) group (Baclawski and Schneider, 2009).

In trying to extract information from experts on the "evaluation method(s)" that they used, only expert 'PF''s description really fulfilled the criteria listed earlier (at the start of this section) for what constitutes a robust evaluation method. Expert 'PF' was able to explain "a logically constructed and repeatable process", which he was formally teaching to others. Interestingly, expert 'PF' emphasised the importance of three things for successful outcomes in seeking to build ontologies whilst practising re-use: (1) "keeping domain practitioners away from the details of ontology languages, particularly during the early stages of working to develop an appropriate ontological model"; (2) the need for "well-documented and refined use-cases" and (3) the need for a "small working group of domain practitioners in order to be able to manage and focus them".

The relevance of 'PF's advice was substantiated in two other interviews with experts: 'JG' (from an oceanography background) and 'RH' (from a biology discipline). Expert 'JG' explained the failure of an ontology activity that he had been involved in as follows:

"Um, so the interesting, um, it is clear that a major reason it has gone so long and so slowly and so poorly, is that we are consciously doing it in a very community centric way for better or for worse.

Author: Yeah.

JG: Um, it is very different than the way you develop an ontology if you've hired somebody, 'Here, go do this. Here are six people who know something about the domain. Produce an ontology'. uh uh, um, that said, a better ontologist would have moved this along quite a bit faster than I did. (laugh)

Author: But when you say that you were trying to do it in a community way, which is that, are you saying that's because you're having to involve the community in all of the steps and get a consensus perhaps ...

JG: Yes.

Author: ... that it is making the process very cumbersome and very slow as opposed to having a very tight team of skilled people, ontological engineers with some domain expertise and working over material that is handed to you...

JG: Right."

In relation to another ontology development activity 'JG' explained the failure of the process due to poor use-case construction:

"And it turned out what we had written in my mind was actually a whole bunch of use-cases that were functionally equivalent to search .... so it wasn't very well captured."

Expert 'RH' in his interview also explained why it was important to get people to articulate use-cases and then challenge the use-case creators (for clarification purposes to correctly ascertain what must be achieved) and also lamented that it was difficult to form small groups to work with because of fears community members had about being left out of development activities.

"What we should have been doing is saying look we have 32 use cases and they shouldn't go in the model. They are not to do with the ontology, they're things we need to be able to account of and some of the use cases … its not possible to model that because its not clear. 'What do you mean by it'. 'You can't specify it'. 'You've used those words because you don't know'. Does that make sense ?

Author: Yes it does. Its an engineering task and one done with a critical eye let's say rather than just assuming everything you've been told needs to be modelled. That's what you are saying ?

RH: Yes.

Author: Lets critically analyse what we are being told here and half of it is the way that it is because its not clear to the people anyway.

RH: Its also about what we are trying to achieve. Are we trying to create an ontology that represents the lexicon or..?

Author: Or are we trying to develop something specific and therefore need something new and clearer.

RH: Yep. Entirely. But that is a big cultural case because people want to be involved they see it as power, but it is like if you are not involved in the creation of these things.. you are frightened that someone is going to come along and say this is how you should do things."

Comparison Of Expert Reported Methods With The Literature

It was clear to the author in conducting interviews with many experts that there was considerable confusion surrounding the idea of an "evaluation methodology" and the related task of applying specific "evaluation criteria". Most experts during conversation dived into the detail of particular evaluation criteria when asked about the overall methodology they had applied, rather than describing an over-arching methodology. It was evident in a number of cases that an "evaluation methodology" and "evaluation criteria" were perceived to be the same thing.

The ontology development and evaluation methodology described by expert 'PF, which was the only method described in any detail, had many similarities with the method outlined in the literature by Annamalai & Sterling (2003). In the Annamalai & Sterling (2003) method they distinguish between domain ontologies, which they describe as only "loosely coupled to one another" and a purposive ontology as one "which explicitly defines the terms for supporting a specific purpose or use". They argue that a purposive ontology encodes specialised domain knowledge by composing various reusable domain ontologies and then affects the necessary application-specific extensions. Annamalai & Sterling (2003) advocate development of purposive ontologies whilst simultaneously pursuing the creation of reusable domain ontologies (see Figure 7.2).

In many senses, this was a similar technique to that espoused by expert 'PF' (who was one of only two experts claiming to develop 'Application' or 'Task' ontologies). The Annamalai & Sterling (2003) method incorporates a constant cycle of ontology evaluation.

Expert 'PF' highlighted that many ontology development tasks with which he is involved are "task or application-specific". In building these types of ontologies he has often had to work in a vacuum, where no relevant domain ontologies exist that he can re-use. His method often consequently develops re-usable domain components. Expert 'PF' explains this situation as follows:

"Now, in that context we looked at some existing ontologies and because of the science area that we focussed on and because we focussed on observed data initially, there was basically nothing in this area.

The SWEET Ontology did have some of the higher-level concepts like physical quantities and features and phenomena and things like that. But they didn't have any population at all for our environment. They didn't have any notion of the middle atmosphere, or upper atmosphere; it didn't have any notion of the sun or any of the things that we measured and it didn't have an instrument ontology at all, and at the time, an instrument ontology didn't exist. Um, what we did take from SWEET was certain spatial representations, and units, and the notion of time.

So we took the more foundational things, um, and imported those into the ontology and they were v.. quite stable.

But we developed our own. Now over time the author of SWEET, um we actually had a subsequent funded project together, and he has been moving some of our development back into SWEET."

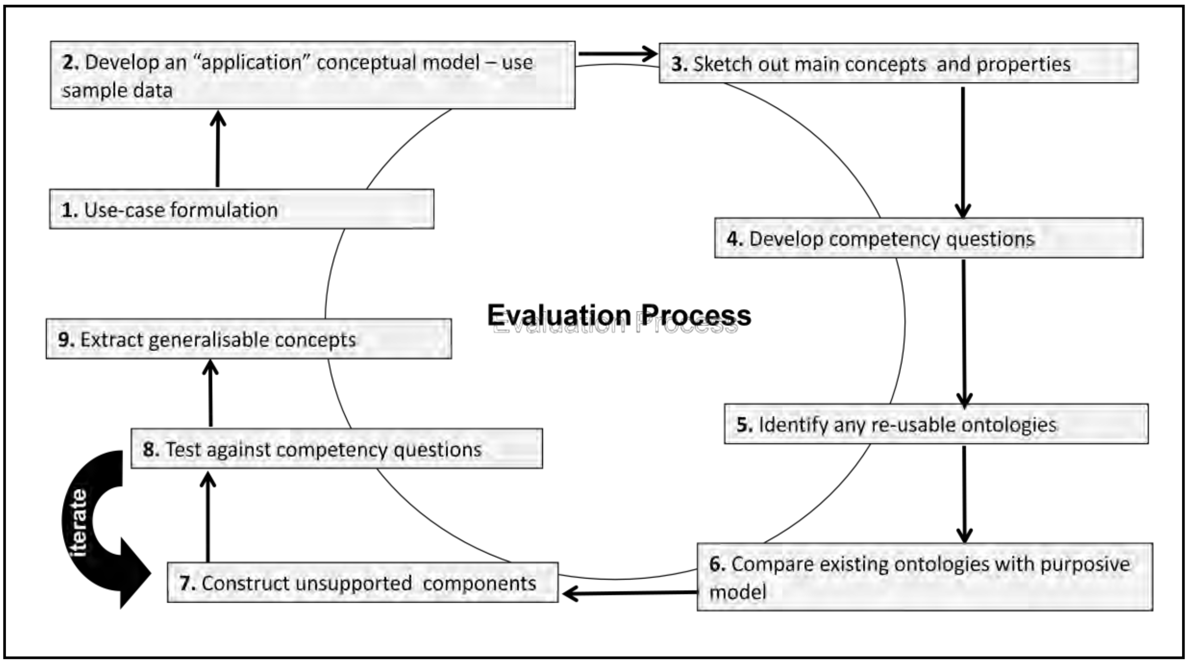

Figure 7.2 Visual Summary of the Ontology Development (and Evaluation) Method of Annamalai & Sterling (2003).

Expert 'PF' does not refer to "competency questions" per se but in his method, he structures and formulates queries (with the involvement of domain practitioners), which the ontology must be capable of answering and then tests the ontology model using these queries. These "queries" are essentially competency questions.

The additional elements in the evaluation method used by 'PF' and his teams (as opposed to that described by Annamalai & Sterling) are the use of simple visualisation tools such as CMAP which helps shield non-ontologists (initially) from issues of ontological detail and allows participants to focus on modelling the domain (or application) using simple, easily understandable constructs. There is also a very large emphasis on getting the detail and description of the use-cases right (sufficient for ontologists to make an initial model of the domain).

Whilst the evaluation methods described by experts in this study encompass elements of methods already presented from the literature in earlier chapters, the most obvious difference is the general lack of adoption of the more complex, algorithm-driven evaluation techniques, or techniques that appear to require a high level of ontological engineering skill (e.g., OntoClean (Guarino and Welty, 2004)). Where structural evaluation was performed it generally encompassed using:

- Visual inspections of the ontology model to assess completeness and to look for disconnected concepts.
- Using students (i.e., cheap labour) to run multiple queries over the ontology instance data to see if it met competency questions.
- Using a standard reasoner to check for obviously unsatisfiable concepts.
- Harnessing existing tools that were tuned for detecting inconsistencies in the ontological model (with respect to community-based rules, e.g., in the case of the OBO community).

Only one expert 'LL' reported using relatively complex structural measures, as reported in Gangemi *et al.* (2005) such as 'depth', 'breadth', 'tangledness', 'leaf and sibling distribution', 'density', 'modularity', 'consistency', 'complexity' and 'logical elements distribution', possibly measured automatically. For most other experts, inference from discussions was that many of the structural evaluation tasks performed were manual in nature. Even when there was mention of functional evaluation criteria such as examining concepts and concept coverage, the mining of existing databases for domain concepts did not appear to be automatic concept matching or automated natural language processing type data mining, but more manual interrogations of datasets or databases. This is probably most evident when inspecting the evaluation measures elicited from interview (as reported in Table 6.30).

7.3 Governance Issues

During interviews a relatively large number of governance issues were mentioned by experts, some of which could be, and were, framed as evaluation criteria. Governance matters featured quite extensively in all interviews and were either explicitly stated as being governance matters, or were mentioned indirectly and were noted by the author as issues considered governance related. There were ten governance criteria listed in the initial formulation of the ontology evaluation hierarchical model, spread across three sub-categories. After the pair-wise comparison exercise, the 'Governance' dimension accounted for 9.2% of the weight evident in the model which was the lowest contribution by any dimension. The fact that governance was a significant topic during interviews, but was later considered to be the least important theme of ontology evaluation when experts were asked to pair-wise rate model elements, is worth inspecting further.

7.3.1 Heavy vs Light Governance Perspectives

Two experts, 'RA' and 'WD', rated 'Governance' much more highly than their colleagues, with 'RA' ranking the 'Governance' dimension the most important theme of evaluation and 'WD' rating it third behind 'Functional Relevance' and 'Usability'. Standing these two experts aside, the interview notes were inspected once again to gauge how the other experts (who participated in the pair-wise comparison exercise) viewed governance issues before being asked to rate them in comparison to other factors.

Expert 'DH' had in fact professed that he did not believe in "heavy-handed governance". He preferred to see much more community (and semantic) freedom. He stated at one point:

> *"Obviously you need some governance in place, and some ability for uh, there to be a community process, which leads to, to, um, ratification, even beatification of certain properties and class of things as the preferred way to go……but at best, I think, the final products of these things have to be considered to be best practice recommendations rather than - obviously, at the level of a government agency or something like that is quite possible to require or ?? But at the level of inter-operability between, um, independent autonomous projects and activities, the best you can hope is to make it easier for people to go with the best practices than to bypass them. So I think the best hope for a solution in this area, is actually for a relatively light touch of the governance side."*

Expert 'DH''s view is that data providers (and data users) should be provided with "tooling" that lets them easily annotate data and information and he is not concerned if the information models that people conform to (for whatever purposes) changes (even in short time-frames). His approach is to be agile and adaptable. He states:

> *"In our case we're trying to architect what we're building, such that if in six months time we decide that the organisation of the information wasn't exactly the way that we'd like, in order to conform to, to some international models ……but it shouldn't be too hard either for us to try and repopulate a new instance, or else just add in extra properties that may be effectively synonyms of properties that we already have, but which give us the flexibility of moving forward."*

In reviewing the interview with 'DH' it was clear that he is a strong advocate for linking data through simple "RDF triple descriptions" and seems to subscribe to the notion of plurality of approach. His preference regarding the choice of semantic models in support of community activities is to let the market decide. Accordingly, the governance model that 'DH' would promote for community-based ontology development is one slanted towards that more suited to managing a 'bazaar' (as in the 'Cathedral and the Bazaar' essay by Raymond (1998)). A 'bazaar' is observed to babble with differing agendas and approaches, unlike the building of a 'cathedral', which is something carefully crafted by individual wizards, or small bands of mages, often working in splendid isolation. In a community ontology development governance scenario, aligned with the model of a 'bazaar', there would be few constraining rules, but plenty of tools to enable a high degree of semantic enablement. There would also be enough semantic craftsmen that you would quickly encode all of the possible permutations that you may want for a particular (ontological) concept that suits your purposes, albeit with a high degree of duplication and redundancy. This is in contrast to the 'cathedral' model of ontology governance, where solid guiding principles are set, more structured and consistent semantics are anticipated from development activities and there is less conceptual overlap. The community's assets in this model are actively checked for compliance with the guiding principles before they can gain the community's stamp of approval and community practitioners are encouraged to use community-sanctioned ontologies rather than develop alternate ontologies in parallel. The 'cathedral' model in this case would be similar to the governance approaches currently demonstrated by the OBO Foundry. Raymond (1998) was of course referring to software development communities and software debugging activities in his essay, but the analogy is a useful one regardless. Expert "DH" explains his position as follows:

> *"Um, it often, it, again, I suppose, this is one of the reasons why I think a light touch of governance is good. Um, it allows you to leave things so that it may be that your community ends up defining two completely parallel streams of activity that are sitting there right next to each other......and allowing those tussles not to be suppressed by an early (? inaudible) that, to trundle along for a while, may well help you to evaluate the real world benefits of both, probably get some prototyping going with both and also potentially end up with a community that keeps the two sides together so the questions of inter-operability between the two approaches are ones, um, yeah, that can be discussed rather than just be treated as a holy war."*

Other experts who participated in the pair-wise comparison activity didn't express any views that would explain their eventual rating of governance elements, one way or the other.

Those who didn't participate in the rating exercise also had their interviews revisited on this particular issue. Expert 'SC' was much more of a 'cathedral-builder'. He was actively seeking ways to formally bring institutions together under the auspices of a standards body, for the purpose of gaining agreement upon, and ratification of, vocabulary standards and semantic data models in hydrology. He stated:

> *"Um, but the idea is to, to be using that (OGC) as a venue to get some more broader scale agreement which would then be, and we're getting WMO to co-chair this. So, the idea is that OGC provide a convenient meeting place and then we would be developing agreements, which would then be forwarded to WMO for ratification. WMO has a commission for hydrology, which is the reason why that makes sense..... There's been a serious effort to make sure that these things aren't just the product of, of, of one academic shut in their office with the door closed thinking hard, and even though in many cases ultimately some of the key decisions do end up getting made by one person, they don't, they don't immediately get to say this is so. They're doing it, playing a particular role on behalf of the big group as being chair of some committee or, or, or something or other...."*

Plurality of approach is not something that expert 'SC' is comfortable with:

> *Now I last week discovered that, um, one of the better known ontology services in the natural sciences uh, the SWEET Ontologies from JPL, which you've probably come across those. Without any reference to us they've stood up a bunch of geo-ontologies.....*
>
> *And that's frustrating because we've been doing this. We would argue that we've been very, very public about letting people know that we've been doing it.*

Expert 'PF' was also more of an advocate for the 'cathedral-style' approach. In specific discussions surrounding governance models, he stated:

> *"Um, and so what we try and do, what I try and do before I leave an area, especially if I know I'm going to come back, I make sure they have some understanding of what governance means and if possible have a governance structure in place, but the reality is that it's a very unevenly dealt with, very much, uh, you know, I think there's an uncertainty in terms of semantics because in my mind because it's semantics, because it's science, it needs to be owned by the scientists and that means professional societies or scientific unions or other national entities and international entities. It's sort of the discussion we 're having at ICSU, is the, its needs to be the role that science plays. The reality is they don't really want to get involved with that on an ongoing basis. It shouldn't be a big burden but someone's got to do it. And when push comes to shove you want the international (scientific) union (ICSU) to say 'OK, you guys are going to use this particular um, ontology because we've looked at it and its all okay and the community is behind it'."*

Expert 'RL' also emphasised governance a number of times throughout his interview and at one point stated:

> *"Yes, I mean you'd want all those things in governance and you'd want governance to be effective. Weak governance is a big problem."*

Since ‘SC’, ‘PF’ and ‘RL’ did not participate in the pair-wise comparison exercise, this could explain why governance appeared to be an important factor initially, but was subsequently considered to have the least influence in a decision-making processes. It is impossible to say how these experts would have rated governance, had they availed themselves of the opportunity to do so, but from the views expressed they may have weighted the governance dimension more heavily (than it was).

7.3.2 Institutional Backing and Community Mandate

Some of the individual ‘Governance’ criteria in an early working version of the hierarchical evaluation model were considered so low in importance, in comparison to others, that the author omitted them from the final model (see Table 7.1 for omitted criteria). Note that ‘Governance’ criteria ‘GF2’ was eventually also removed (in addition to those listed), but this was because it was considered a duplication of another criteria, elsewhere in the model.

Table 7.1 Omitted Governance Criteria

Code	Criteria
GF3	Does the community review or moderate individual ontology developments.
GC1	How mature the community is in terms of ontology development and its longevity and cohesion as a community of practise.
GC2	Whether the community is institutionally backed.
GC4	Whether the community’s mandate is obvious and well-bounded.
GB3	Whether the governance and the core group of developers engender trust and credibility.

Although the criteria in Table 7.1 have been dropped from the final model it is worth emphasising that a couple of these criteria (‘GC2’ and ‘GC4’) were repeatedly extracted from interviews with experts who had played all three roles in a community (i.e., Leader/Driver/Initiator; Maintainer; Specialist Advisor) and who did not subsequently participate in the rating exercise. It is considered plausible that those who have played the role of “Leader/Driver/Initiator” are perhaps more attuned to some aspects of community governance that relate to how a community (ontology development activity) can best be made sustainable. These views may not have been adequately captured in the rating exercise because of the absence of these experts. Some common threads emerged from interviews with these types of experts that relate to the deleted criteria listed in Table 7.1.

Institutional backing, for example, can be manifest in a number of forms and it is important to understand how each form influences ontology creation and uptake. A government can decide to fully fund a specific initiative, in which case the funding is usually time-bound and the project can be relatively well-resourced. Whilst community productivity is initially high, if sustainable community practises are not built in from the start, when the source of funds is removed, community efforts may fall away. In rare cases, funding is ongoing as in the case of the National Centre For Biomedical Ontology (NCBO), which is closely allied to the OBO Foundry and which draws its base funding from the US National Institutes of Health. Sustainability in such cases is not entirely resource-dependent, but certainly growth in the user-base will be a function of the quality of the community products and the types of practises adopted by the community.

But even in situations where institutional backing involves community funding, one of the common issues that emerged (from interviews) was a perception that it was much easier to gain resources for "physical infrastructure-building" activities, often at the expense of activities related to the "information engineering" aspects of systems development. One expert stated:

> *"there is no research and development money.. it's all D and no R.."*

So getting the right mix of resourcing that spans all of a system's engineering viewpoints (i.e., enterprise, information, computational, engineering and technology) is vital to the success of a community-based semantic enablement project.

An alternate institutional model also evident from interviews is where specific institutions (government and non-government), with business requirements that can benefit from ontology development activity, support their staff to work on ontology development as part of their institution's project-based activities. Allocation of resources in this case will often be made overtly and a percentage of an individual's time can be devoted to community activities. Another allied type of institutional backing is where (sanctioned or unsanctioned) an individual may participate almost in a type of "skunk-works" mode in contributing to community activity. This latter type of model, not surprisingly is usually the least successful method for a community to both build ontologies and for it to attract users for its products.

Lastly, institutional backing can come from "institutions" who do not necessarily contribute resources, but who are seen as a credible peak body in a particular discipline or relative to a particular issue, and they lend their imprimatur to community activity. Community productivity (in ontology development) is directly proportional to the resources invested, so institutional backing that provides willing hands is crucial for the development of credible, useful and maintained products. But the broader adoption of a community's products can be influenced by the "seal of approval" given by a well-respected (disciplinary or interdisciplinary) body. At least two experts are currently engaged in liaising with, or lobbying, key peak bodies to play a role in domain ontology standards setting. This type of "backing" can also influence judgements made concerning evaluation criteria listed in Table 7.1 such as "Whether the governance and the core group of developers engender trust and credibility", because "seals of approval" and/or "recommendations" from trusted bodies provide a level of user assurance (and can be a potential indicator of quality).

In addition to governance issues related to institutional backing, another governance matter which was raised commonly by experts was a community's ability to adequately identify its mandate and the scope of its domain (and/or task) activity (i.e., criteria 'GC4'). There was a view that some communities stray beyond their remit, which appeared to create rivalries between groups. Experts felt that a community should only seek to "govern" those semantic components for which they have the resources, mandate and expertise. The issue of "mandate" and how to define it is possibly another reason why some experts are attempting to establish imprimaturs from peak bodies.

Conversely there was also a lament from a number of experts that there were many types of ontologies (or semantic concepts) that were more universally applicable and which should not be defined and duplicated, in a myriad of forms, by each community. Simple examples drawn from expert interviews included things like 'units-of-measure', 'notions of time', 'spatial relations', 'contact details' and' jurisdictional ontologies' such as concepts associated with 'countries'. In such cases communities were managing vocabularies and ontologies that they did not want to be the custodians of, or the point of truth for.

In yet other circumstances, experts alluded to the fact that their trust in the governance of a particular community (who were perceived to have a mandate in a particular area) was so low that they stood up versions of an ontology covering a particular semantic space, as an alternate point of truth for the ontology (or vocabulary). In other situations it wasn't the case that a community with a particular mandate wasn't trusted, it just had no means of sustaining an ontology development effort. In the most extreme cases experts had re-used elements of another community's ontology and deliberately forked so that they could maintain something credible for users to access.

There are a number of lessons to be drawn from these findings for communities such as SCAR and the AODN. Both are highly multi-disciplinary communities (both are organised thematically and geographically, with one focussing on the Antarctic and Southern Ocean and the other centred around marine disciplines of interest in Australia's jurisdiction). In developing their Feature Catalogue(s) both need to consider what semantic content they should endeavour to develop and govern as part of their own mandate and what semantic content should ideally be re-used from other sources. Of most significance is that for community semantic enablement activities to be sustainable, they need to be institutionally backed (ideally from both a resourcing and an imprimatur perspective). The data exchange infrastructure development activity which, for the AODN and SCAR, is largely carried out through a combination of institutional resourcing and project funds, should make explicit allowance for the information engineering aspects of systems development.

7.3.3 Governance Framework

A governance criterion that was favoured by experts and which was therefore retained, was ‘GF1’ – “Does the ontology development community have a published governance framework ?”. This particular criterion, which attracted an importance rating of 3.2%, was the highest governance factor (after removal of sibling governance criteria due to their relatively low overall ratings). Given ‘GF1’’s prominence (at the local level in the model), it is instructive to examine the various parts that make up a framework and review why this particular ‘Governance’ criteria was considered important by experts. Although the word “framework” was mentioned several times, it wasn’t clearly defined by experts. Instead, issues which were mentioned in a governance context and which appeared to the author to constitute “framework” components, are discussed here.

Inferences drawn from expert commentary were that a governance framework consisted of:

- **A set of guiding principles**: Guiding principles establish community norms in relation to both technical and behavioural matters. The OBO Foundry principles, for example, cover issues of format; naming conventions; term and property addressing methods, as well as covering behavioural expectations about ontology access, ontology maintenance and collaborative modes of development.
- **Declared roles** (e.g., gatekeepers/curators; committers; helpdesk; ontology police): A variety of roles were mentioned with respect to ontology-life-cycle management. The most commonly mentioned role was that of ‘curator’. ‘Gatekeeper’, ‘editor’ and ‘help-desk’ were also variously used as a synonym for this role. Experts indicated that the absence of identified ‘curators’ usually lead to poorer quality ontologies. Uncurated ontologies, that is, those without an identifiable “decision-making editor” who routinely responded to change requests, become gradually unusable or grow in an unsustainable and/or shallow manner. When there is no-one (designated) to amend core structural or design problems, the ontology will often get extended narrowly and in a specialised manner, along branches of the ontology that have few internal or external dependencies. In situations where orthoganility is a community guiding principle, a poorly curated ontology will often impact upon the development of other ontologies that need to re-use terms from the uncurated ontology. Curators are often assisted by ‘committers’, i.e., people who are “trusted” to make changes to ontologies but who are not necessarily viewed as curators. The importance of instituting

"change tracking" systems increases with the increasing number of 'committers'. The other key role mentioned was that of 'ontology reviewers', also referred to as the "ontology police". Reviewing can be a structured process, as in the OBO Foundry case-study, where reviewers are often "logicians", or reviewing can be a less formal process conducted by curators, or by users as a by-product of the normal cycle of "implementing" an ontology in a particular use-case.

- **Review mechanisms**: Most ontologies are in a constant state of evolution and their quality will reflect the level of scrutiny applied by users in assessing their fitness for the tasks for which they were designed. Smith (2008) presents two methods for undertaking ontology review, a peer review process analogous to 'peer review' in scientific publications, which is a rule-driven editorial process vs democratic ranking (by users). He concludes that both approaches have their benefits and detractions.
- **Ontology repository**: Most experts felt that finding suitable ontologies for re-use was a difficult task. An important framework element is therefore the registration of an ontology in a suitable repository so that it can be discovered, adequately described and managed. Hartmann *et al.* (2005a) have proposed an ontology metadata vocabulary (OMV) for describing the relevant properties of ontologies for supporting their reuse. Such metadata might form the basis of ancillary information registered in an ontology repository, however, a standard architecture for ontology repositories, their service interfaces and information models is still under active development within groups interested in the management of semantic content (e.g., Open Ontology Repository Initiative).
- **Community wiki**: To facilitate communication amongst the various framework role players and the ontology user community, mediums such as community portals, wikis and list-servers were often mentioned by experts. In some cases experts explained that the community wiki was also currently the community ontology repository.
- **Community tooling**: Development (or adoption) of "easy-to-use" tools can help a community conform to espoused norms as well as help grow the ontology developer and user-base. Tooling can include tailored ontology editors that are designed to operate with community ontologies and vocabularies and which make tasks like versioning, change-tracking and referencing of ontology modules (or terms) simple to achieve. Tooling can also encompass systems such as ontology repositories and ontology visualisation and graphing aids.

If a community can instantiate these "governance framework" components there appears to be a better chance that ensuing ontology development activities will meet community needs and be sustainable. In the communities discussed, the availability of people with ontological skills capable of playing the key roles mentioned was a reportedly limiting factor in nearly all cases. If the AODN and SCAR communities are to be successful in semantically-enabling their respective infrastructures, a partnership, or pooling of resources, with another community already exhibiting a governance framework (and skilled ontologists) would be the easiest route to success.

7.4 Summary

In this chapter inconsistencies found in the data with respect to expert opinion were discussed. Both 'within' and 'between' expert divergences were canvassed as well as any differences in expert opinions expressed between interviews and the pair-wise comparison exercise. It was speculated that some of the internal inconsistency (i.e., the 'within' expert inconsistency) may have been due to methodological limitations in the research techniques employed. A follow-up study was suggested to test for this possibility. The degree of detected inconsistency across the board, however, was strongly suggestive of expert-centric factors also at play, regardless of any perceived methodological limitations. A review of the literature revealed many studies in which experts (in general) had demonstrated considerable divergence of view. The situations in which these experts were studied were broad-based and not necessarily based on soliciting views using an AHP technique.

The remaining component of RQ1.2.1 was addressed in this chapter by discussing the data that emerged from analyses in Chapter 6, regarding the types of ontology evaluation methods used by the experts in this study. As was suggested at the start of this thesis, most (77%) experts were found to be using evaluation methods developed by their own communities, with only one expert able to fully substantiate a claim of using methods with academic origin. There was, however, overlap between the evaluation criteria cited by experts and those reported earlier in this thesis from the literature. Most experts appeared to confuse an 'evaluation method', with 'evaluation criteria'. The OBO-affiliated experts provided good insight into their broad methods of ontology development, encompassing elements of ontology evaluation, although there was insufficient detail to determine if what was described actually could be considered an 'evaluation method'. Outside of the methods with academic origin, most clarity regarding methodology emerged from the Solar Terrestrial Physics discipline. The methods used in this discipline closely paralleled that reported by Annamalai and Sterling (2003), although the method was not explicitly modelled on their technique.

Most experts were not using highly complex and difficult to apply evaluation methods, criteria or measures. In this sense it would appear that there is a disparity between what is outlined as 'best practise' in the literature and what is actually being performed in practise.

The chapter concluded by presenting a range of observations regarding community and ontology centric governance issues, including an outline of the elements of a derived 'governance framework', which emerged from information supplied by the participating experts. It was suggested that such a framework could be adopted by the SCAR and AODN communities to ensure good governance for their future semantic enablement activities.

Chapter 8.

Conclusions

The research in this thesis sought to address how Antarctic science communities could practically manage and select domain ontologies for use in semantically-enabled data exchange scenarios, given feature-centric Web service design patterns (RQ1). A semantically-enabled data exchange scenario is one in which the descriptive elements placed within data and their schema permit communication and cooperation between machines and between humans and machines. Ontologies provide the type of description necessary for machine to machine communication because they specify how a given scientific community models, defines, relates, interprets and encodes their knowledge. Semantic-enablement must therefore involve establishing a reference between the vocabularies used in the exchanged data and the ontologies developed to describe those vocabularies. The process of establishing these links is semantic annotation (Maue, 2009). Since most Antarctic-themed scientific resources are not yet deployed in a Semantic Web context (as argued in Chapter 1), the motivation for the investigations conducted in this thesis, was to explore how the process of semantic-enablement could be facilitated for current data infrastructure builders and scientific data providers, active in Antarctic science.

The benefits of semantic-enablement are manifold including much higher precision searches for existing online scientific data and services, greater possibilities for both human-assisted and automated data integration and by harnessing reasoning capabilities, many forms of service processing and chaining (i.e., linking the output of one service as input to another) can be automated.

The scientific communities who were the focus of this study are building and using service-oriented data exchange infrastructure designed to transmit and utilise four-dimensional observational and measurement data (i.e., data often with spatial and temporal components). These communities (i.e., the AODN and SCAR), who were used as case studies, had elected to adopt ISO spatial data standards and OGC Web service technologies as the foundation of their data exchange infrastructure. Like many other scientific communities of practise, these communities were novices with respect to semantic-enablement. Most of their efforts had been focussed on developing comprehensive lexical descriptions at the dataset level (i.e., dataset level metadata) and storing these documents in OGC-compliant Service Registries. Whilst these descriptions are useful for human users, and are a useful source of (non-ontological) search terms, they are not much more than uninterpretable verbage to machine clients.

The standards used by the SCAR and AODN communities direct that the semantic repository associated with the ISO and OGC technology suites to manage domain concepts, should be a Feature Type Catalogue (ISO, 2005b). In the ISO flavour of spatial data standards and the OGC Web service specifications, which are based on a Generalised Feature Model (ISO, 2002), a Feature Type (or its instance) is equivalent to an ontological concept. ISO 19110 (ISO, 2005b) is a conceptual specification for a Feature Type Catalogue model, expressed in UML and which is therefore not semantically grounded.

As was demonstrated in earlier chapters, ontologically-grounded Feature Type Catalogues are not being deployed within OGC-compliant infrastructure and it was reasoned that this was because there is no guidance (nor exemplars) to show how this can be achieved.

This lead to the postulation of a research sub-question (RQ1.1) which specifically asked "*what characterises an ontologically-grounded Feature Type Catalogue*" (see Figure 1.3, reproduced in this chapter, from Chapter 1). To investigate the topic adequately, four nested questions were also proposed as part of RQ1.1. The first question (RQ1.1.1) sought to address the types of use-cases and data models that a Feature Type Catalogue should be capable of supporting, particularly with respect to the contexts in which the two case study communities operated. Essentially this RQ was focussed on eliciting the full set of requirements that a Feature Catalogue should meet. The second of the questions (RQ 1.1.2) then asked whether the various ISO/OGC standards (including conceptual data models), relevant to the development of a Feature Type Catalogue, were appropriate to meet the needs of the studied communities. The third question (RQ1.1.3) sought to investigate how the ISO-based Feature Catalogue model (ISO 19110), or whatever enhanced model that might, of necessity have been derived from it, could be 'ontologically-grounded'. The last question (RQ1.1.4) pertaining to the characterisation of an ontologically grounded Feature Catalogue, asked what methods are best suited to access content from the Catalogue, cognisant that the methods used would need to be consistent with, and sympathetic to, interoperation with the existing components of the data infrastructure of the participant communities.

Given the dearth of semantic-enablement within the studied communities and their lack of familiarity, or experience with ontologies, semantic repositories and semantic annotation, it wasn't considered sufficient to just to focus on questions related to the development of technological artefacts (such as the Feature Type Catalogue).

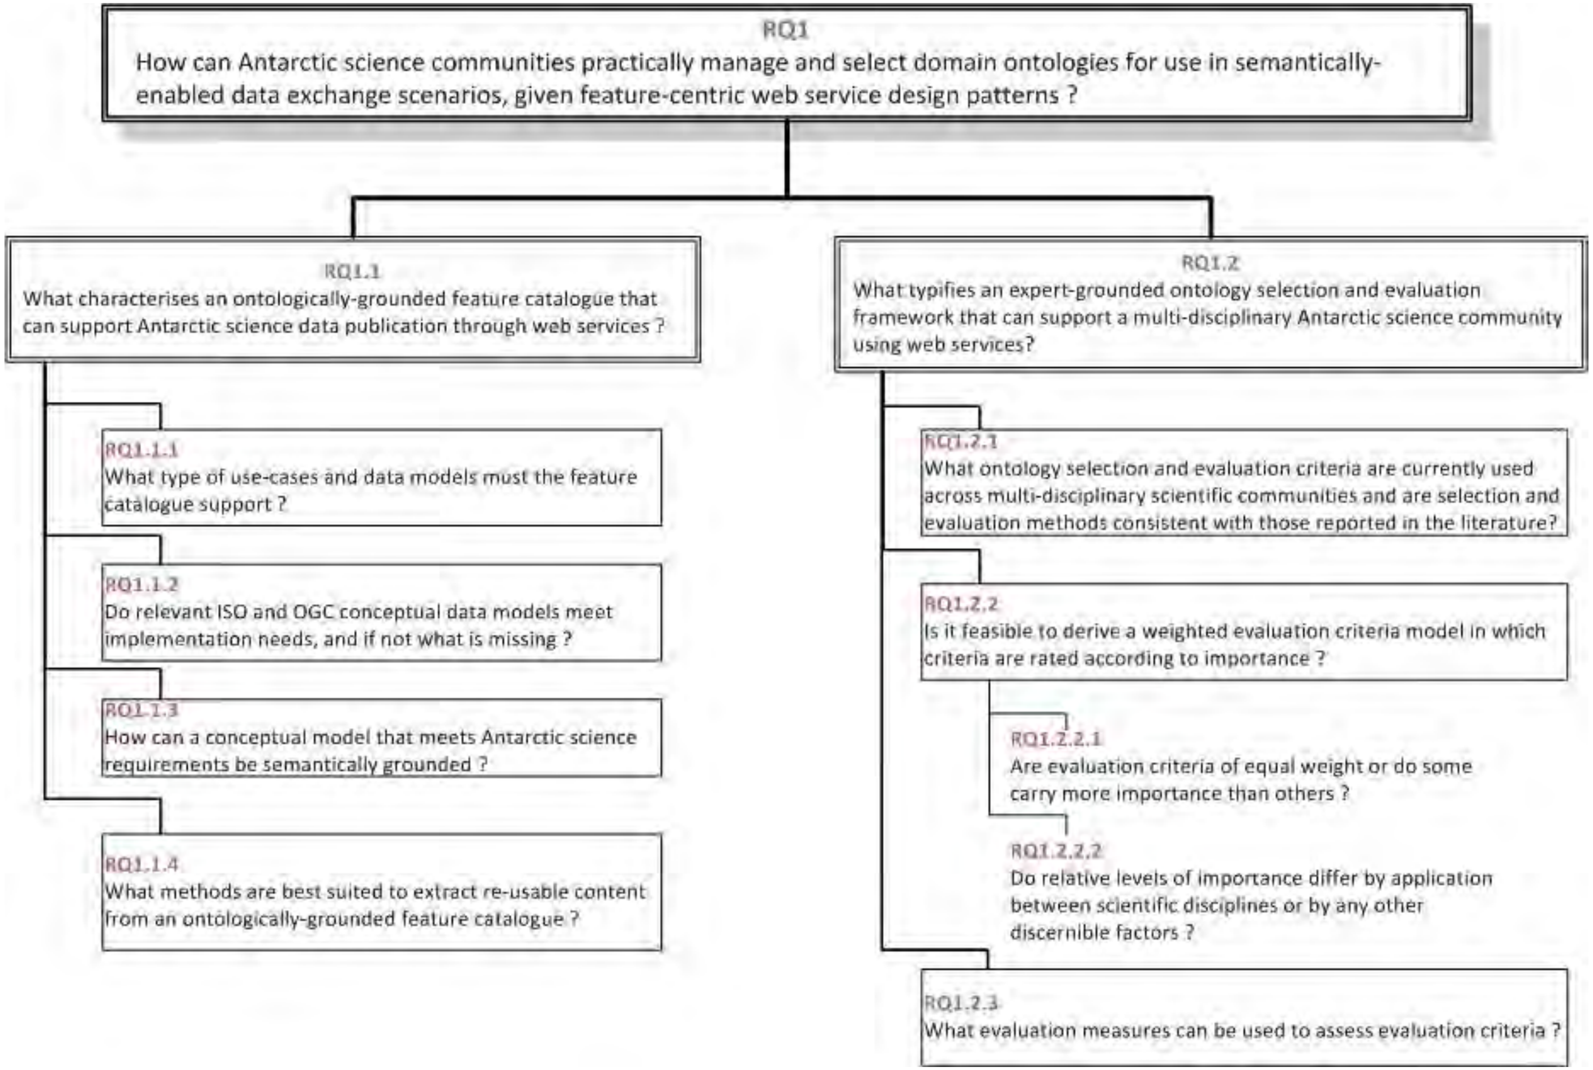

Figure 1.3. From Chapter 1

For the Feature Type Catalogue to be of practical utility to the groups in this study, they would need to be guided about how to populate the Catalogue with ontological content and then how to use the Catalogue for the purposes of semantic annotation. Whilst there are many reported methods in the literature for selecting and evaluating ontologies (e.g., Ontometric: Lozano-Tello and Gomez-Perez, 2004; OntoClean: Guarino and Welty, 2004; EvalExon: Spyns (2005)), most have been characterised as being resource intensive to apply and requiring a high level of ontological expertise (Hartmann *et al.,* 2005b; Kalfoglou & Hu, 2006; Blomqvist *et al.,* 2006). These issues were discussed in detail in Chapters 1 and 2. It was, therefore decided to investigate the methods and the criteria actually being used by scientific domains that had already embarked on the semantic-enablement journey.

The purpose of such investigations being to first identify the techniques that are used in infrastructure development activities and then to bring them together in the form of a guiding 'framework' that could then be applied by Antarctic communities to select and evaluate ontologies. These evaluated ontologies would be the feedstock to seed the Feature Type Catalogue and could consequently be used for semantic annotation. It was also considered that by using experts to compile such a framework that other general guidance, related to the task of ontology selection and evaluation, might emerge which could also be of practical significance.

The second major set of research questions accordingly centred on "*typifying an expert-grounded selection and evaluation framework, capable of supporting a multi-disciplinary Antarctic science community"* (RQ1.2). This research sub-question also had several nested questions (see Figure 1.3). Of significant interest were the evaluation criteria and the overall evaluation methodologies being used by experts already working within semantically-enabled scientific infrastructure (RQ1.2.1). Given that other studies (e.g., Paslaru Bontas-Simperl and Tempich (2006)) had found a low correlation between the methods cited in the literature and those practised by experts, it was anticipated that there might also be differences between theory and practise in this study. Because ontology evaluation can be a difficult and (inexact) exercise, even for skilled ontological engineers, in developing the guiding ontology evaluation framework it was considered pragmatic to only include those evaluation elements that were considered essential (i.e., of most importance in the selection decision-making process). To that end, this study was also interested in examining the various levels of importance (or weight) that experts placed on evaluation criteria (RQ 1.2.2). It was reasoned that highly weighted criteria were those that were the most influential in selecting ontologies for re-use. Recognising that 'criteria' must be measurable, also of interest was what evaluation 'measures' were being used to actually make assessments of individual ontologies, or which were being used to comparatively rate a number of ontologies in order to select the most suitable ontology for re-use (RQ1.2.3).

In reporting the results of the research designed to address all of these questions, this thesis is intended to be a "guide book" for those (Antarctic-related) scientific communities (using ISO and OGC standards) wishing to embark on the semantic-enablement exercise.

The two main contributions of this thesis have been the development of an ontologically-grounded Feature Type Catalogue (repository) based on ISO standards which can be used in OGC-compliant data infrastructure for the purpose of supporting semantic annotation of observation-centric scientific datasets. The second key contribution is the development of an expert-grounded ontology selection and evaluation framework that can be applied by communities to populate the Feature Type Catalogue with ontological content. To the author's knowledge the derivation of an ontologically-grounded Feature Catalogue is novel and has not been attempted before and the selection and evaluation framework, whilst not incorporating necessarily novel evaluation criteria or measures, is a reflection of current expert best-practise (rather than a re-statement of academic theory). As a methodology it is therefore more likely to be embraced than many of the methods reported earlier in this thesis that have arisen from academic exercises.

The remainder of this chapter summarises the main research findings and the various contributions that resulted from answering the research questions that drove this study. Where there were any perceived methodological limitations constraining conclusions that were reached, these are remarked upon. At various junctures, through the development of this thesis, additional research areas were mentioned which could provide enhanced validation of some of the conclusions reached, or which could help to improve techniques for semantically enabling aspects of scientific data exchange networks. These potential future research topics are restated and summarised in this chapter.

8.1 Feature Type Catalogue Related Contributions

In exploring RQ1.1 it was demonstrated through a variety of data modelling exercises, described in Chapter 4, that many of the Antarctic datasets exchanged by the communities participating in this study conform to a basic observation and measurement (O&M) pattern (Cox, 2006; 2010a). The basic O&M pattern, as presented by Cox (2006, 2010a), however, requires some augmentation and specialisation in order to cover the various attributes commonly found in Antarctic-themed datasets. This research has shown that biological datasets and their inherent (observation-centric) Feature Types present different modelling challenges to Feature Types that model purely physical phenomena or which generally encode GIS-centric datasets. Although the research in this thesis was deliberately designed to produce a Feature Type Catalogue artefact, some artefacts have also been opportunistically generated as a by-product of the exploratory investigations associated with data modelling. The suite of biological data model templates described in Chapter 4, are examples of such opportunistically emergent artefacts.

Emergent Artefacts

These generic biological models, based on the observation and measurement conceptualisations of Cox, (2010a), could be taken and systematically "specialised" by the SCAR and AODN communities to encode most types of biological data being captured through their observation and measurement research programs. To the author's knowledge neither community has yet developed biological encoding "patterns" that are fitted to the OGC observation and measurement standard, so this work provides a sound basis from which the communities can springboard to develop concrete implementations. Future research in this area, required to develop a set of canonical encodings, would involve development of standard O&M (Cox, 2010a) "result types" suitable for efficiently encoding specific types of data (conforming to the patterns identified in Appendix 7 of this thesis) that can be plugged into the "specialised" models. These O&M "result types" could be in the form of OGC 'Coverages', 'Feature Types' or simple record types. Before these models can be implemented, additional design work will also be required to specialise the 'Observation' entity in each of the template patterns, perhaps with some concurrent adjustments made to the associated O&M 'Sampling Feature' in order to reflect any needed variations in the temporal domain of the data (which is encoded within the 'Sampling Feature' entity).

The most significant conclusion arising from the research where the focus was to characterise an ontologically-grounded Feature Type Catalogue and which subsequently coloured the direction of the research outcomes, was that an observation-centric dataset contains Feature Types where the complete semantic signature requires reference to both its spatio-temporal footprint and its sampling context (if the purpose of the semantic definition is to establish Feature Type equivalence for use in data integration exercises). A detailed review of ISO 19110 (*Methodology for Feature Cataloguing*, (ISO, 2005b)) and how it related to companion standards (ISO 19109 (ISO, 2005c) and ISO 19126 (ISO, 2009a)) found ISO 19110 lacking, particularly in its exclusion of sampling context entities. Noting these drawbacks and armed with a relatively deep understanding about the community data to be transacted, coupled with use-cases surrounding the application of a Feature Catalogue, an enhanced ISO 19110 Feature Catalogue model was developed, that was definitively more able to meet community requirements. Development of this model was the culmination of the research into RQ 1.1.1 and 1.1.2.

The enhanced ISO model is also considered to be an 'opportunistic' emergent artefact from the research in this thesis. The ISO Generalised Feature Model is a meta-model that was conceived to define the elements and properties that constitute a feature. A feature is a typed object and was originally born out of conceptualisations involving data used within GIS and by GIS specialists. It is therefore geared to present a particular view of the world and it was adopted some time ago by the OGC as the basis for its data exchange and Web mapping services. The observation-centric view (Cox (2006, 2010a) of scientific data, albeit still cast in a Generalised Feature Model framework, came along much later and has only relatively recently been adopted as an OGC standard. Since the feature-centric model predates the observation-centric view a number of the ISO standards, such as the ISO 19110 *Methodology for Feature Cataloguing*, are still firmly rooted in a feature-centric paradigm.

In this research it has been demonstrated that an observation's sampling criteria are important components of a Feature Type's semantic definition, yet the ISO 19110 standard makes no provision for sampling criteria in defining a Feature Type. In fact "collection criteria" are deliberately omitted. Whilst this research indicated sampling criteria were important semantic components for defining observation-based data concepts, the value of including "operations" (i.e., functions that could be performed on a feature, or functions the feature could perform) were not considered important elements of a semantic definition.

The enhancements made to the ISO 19110 standard (in this thesis) to accommodate the requirements that emerged from investigating community datasets, use-cases and systems are thought to be those typically required by most multi-disciplinary scientific communities who are building infrastructure to exchange observed or measured data with spatio-temporal components. In this regard this study's enhanced Feature Catalogue model is thought to be highly generalisable and its development has identified gaps in the standard (when attempting to apply the Feature Catalogue Model to observation-centric data).

Artefacts By Design

Unlike the emergent artefacts above, by design, this research purposefully set out to develop a semantically-enabled Feature Catalogue artefact. This was achieved using an enhanced ISO 19110 model as a foundation and a Design Science research methodology. ISO 19110 provides descriptions for Catalogue resources, which are insufficient for either semi-formally (e.g., that can be attained by using RDFS) or formally (e.g., attained by using OWL) defining terms. The enhanced (ISO-based) Feature Catalogue conceptual model was therefore ontologically grounded (to attain the required level of formalism) by casting the conceptual model in terms of concepts and properties defined by DOLCE (an upper ontology of particulars (Masolo *et al.,* 2003)), using a number of relevant ontology design patterns (also anchored in DOLCE) and described using OWL-DL (stored natively and in Oracle™'s triple store). A second version of the model was also developed based on SKOS (which was generated from a relational data store, also in Oracle™).

In the OWL-based ontological model, Feature Types were conceived of as specialisations of "WebResource" concepts that in DOLCE-based ontology design patterns are 'information entities' representing 'real-world objects'. The outcome was a conceptual model that was ontologically grounded and which met Antarctic community needs, and thus satisfying RQ1.1.3.

In the absence of guidance from ISO and the OGC, prototyping activity was also undertaken to develop suitable access methods to Feature Type Catalogue content. Demonstration REST-based services, based on URI templates (Gregorio *et al.,* 2010) were established in preference to other available access options because it was argued that they presented the lowest barrier to uptake; provided the broadest access to Catalogue content (using current browser and internet technology) and were consistent with approaches being adopted for accessing content in ontology repositories more generally. Content access service development required creation of service query patterns such that each query would provide a unique URL for the Catalogue resources delivered. An enhanced version of SKOS was trialled as one type of service output (the other types being HTML and XML). These services are capable of delivering Catalogue content in different formats and at differing levels of information granularity.

Having established concrete implementations of the Feature Catalogue, research was also undertaken into how the Feature Catalogue could interoperate with the emerging infrastructures of the two case study communities and those of other communities, who use similar technologies and standards. Points at which the URI templates could be inserted into application schema documents, data service descriptions and service registries was discussed theoretically and demonstrated through a desk-top exercise, based on current ISO and OGC standards for developing application schema, metadata and registries, supplemented by reference to other IT standards (e.g., Xlink (DeRose *et al.,* 2001); SAWSDL (Kopecky *et al.,* 2007); RDFa (Adida *et al*., 2008) and Microformats (Kopecky *et al.,* 2008a)). These services and demonstrations of how the services could be coupled to existing infrastructure for semantic annotation purposes, provided concrete outcomes for RQ 1.1.4 –

"*What methods are best suited to extract re-usable content from an ontologically-grounded Feature Catalogue ?".*

It is important to note that the ontologically-grounded Feature Type Catalogue developed in this study was theoretically demonstrated to be a useful source of domain ontological content for other, already semantically-enabled scientific data exchange infrastructures (e.g., Sensor Networks and biological data exchange infrastructure). Its broad suitability was facilitated by the scope of the semantic signatures it could support in serving domain (Feature Type) concepts (e.g., associated details of scale, units-of-measure, datums and semantically significant sampling information). In some of the specific cases investigated there was also a relatively high degree of ontological homogeneity due to the use by other communities of the same upper level ontology (to that used in thesis, i.e., DOLCE).

The final contributions made with respect to development of the Feature Type Catalogue involved suggestions about how the SCAR and AODN communities could embrace the artefacts and learnings developed from this study and take them forward in a semantic-enablement activity (discussed in Chapter 5). It was suggested that these communities ease themselves into semantic-enablement using a staged approach by using less formal languages of expression (e.g., SKOS), existing tooling and semantic repositories based on technologies with which the communities are familiar (inclusive of REST-based query interfaces). An upgrade path would ultimately see these communities migrate towards the use of technologies such as OWL, SPARQL and triple-stores.

8.2 Ontology Selection and Evaluation Contributions

Another major artefact delivered by design in this thesis (through the harnessing of qualitative and quantitative methods), was an expert- grounded (and weighted) AHP-based (Saaty, 1980) hierarchical ontology evaluation model, which coupled with application guidance and associated evaluation metrics, formed a practical ontology evaluation framework suitable for supporting multi-disciplinary Antarctic science (RQ1.2). A Screening Survey was used to find suitable experts who could help address the research questions posed concerning ontology selection and evaluation practise. One expert joined the study after the conduct of the Screening Survey. The fourteen experts who were selected to participate in this part of the study were not drawn from the AODN and SCAR communities, but were ontology experts active in national and international groups already semantically-enabling their data infrastructure. These experts collectively covered six of the disciplines represented in Antarctic science. There was a relatively high degree of overlap in the communities covered by these experts, with several experts active across a number of common communities.

After analysis of the Screening Survey data, each participating expert was interviewed (using a semi-structured interview method), their interviews recorded, transcribed and qualitatively coded (primarily using the Template Analysis techniques of Crabtree and Miller (1992) and King (2004)) and a hierarchical model was developed. An AHP-based (Saaty, 1980) pair-wise comparison exercise was then conducted using individual expert responses to a questionnaire that elicited weights for each model element. These model weights were then framed as matrices of preferences and an Eigenvalue equation operating on these skewed matrices of comparisons was used to compute estimates of the relative importance of decision criteria (Genest and Zhang, 1996). The weighted hierarchical evaluation model that resulted is three-tiered and comprises of five dimensions at the top tier (i.e., 'Structure', 'Functional Relevance', 'Usability', 'Maintenance' and 'Governance'), which decompose into thirteen sub-categories in the second tier and forty-two individual evaluation criteria at the lowest level. The model's evaluation criteria were identified by interviewed experts as those being of most importance in selecting ontologies, or ontological components for re-use.

The weights in the final model delivered in this thesis represent the relative levels of importance that experts collectively place on the model's criteria, sub-categories and dimensions. This collective result was derived from the normalised geometric mean of all expert pair-wise comparison preference data. Although the computed group result showed a 'fair' degree of concordance between expert ratings (as measured by Kendall's co-efficient of concordance), there was never-the-less considerable variability in weights allocated amongst respondents in a number of between model dimension and sub-category comparisons. Post model development feedback obtained from experts remaining in the study to the end, indicated satisfaction that the model, in general terms, reflected the criteria used in ontology selection and evaluation. Opinion was divided, however, regarding the suitability of using the 'geometric mean' as a measure of the 'group' result. Development of the weighted model provided an outcome with respect to RQ1.2.2 – "*Is it feasible to derive a weighted evaluation criteria model in which criteria are rated according to importance ?*" (refer again to Figure 1.3), but the results indicated that there was only a 'fair' degree of concordance, between experts regarding allocated weights. With only a 'fair' level of unanimity between experts, patterns or clusters of experts, in terms of how elements of the model were rated, were investigated. In doing so, RQ 1.2.2.1 and RQ 1.2.2.2 which both focus on unearthing reasons for any detected differences in how experts rated criteria, were addressed (and reported in detail in Chapter 7).

In summary there was a high degree of 'within' and 'between' expert divergence and 'inconsistency' and it has been speculated that some of the detected internal inconsistency (i.e., the 'within' expert inconsistency) may have been due to methodological limitations in the research techniques employed. A follow-up study was suggested to test for this possibility. Possible limitations and further studies are summarised in subsequent sections of this chapter. The degree of detected inconsistency across the board, however, was strongly suggestive of expert-centric factors at play, regardless of any perceived methodological limitations. A review of the literature revealed many studies in which experts (in general) had demonstrated considerable divergence of view when asked to perform tasks involving preference elicitation using AHP and non-AHP based methods (Dawes, 1971; Einhorn 1974; Tversky and Kahneman, 1981; Cleverdon, 1984; Brugha 1998; Armstrong, 2001; Kryvobokov, 2005; Lin *et al.,* 2008; Meixner, 2009).

The methodology used in this thesis, to derive expert opinion on ontology selection and evaluation, was inspired by the ONTOMETRIC approach of Lozano-Tello and Gomez-Perez (2004) who also used an AHP approach to frame an ontology evaluation method. Despite the similarity of the underlying approach (i.e., the use of AHP), there are notable differences between ONTOMETRIC and the hierarchical evaluation framework that resulted from this research.

The most obvious difference is that ONTOMETRIC has '160' characteristics (or evaluation criteria) against this hierarchical model's '42'. It was argued earlier in Chapter 7 that as the number of criteria to be evaluated increases, experts are likely to make inconsistent judgments during pair-wise comparison (Lin *et al.,* 2008). Brugha (1998) also demonstrated that when there are an excessive number of questions, there is a drop off in interest towards the end of the questioning process and there are consequently higher levels of reported inconsistencies. Evidence from this study would suggest the same. A criticism of ONTOMETRIC is therefore that the number of criteria inherent in the approach is large and possibly difficult to apply in practice.

The main dimensions in ONTOMETRIC are categorised as "tools"; "language"; "content"; "methodology" and "cost", which contrasts significantly with the dimensions in this study's model. Issues related to ontology 'language' and material 'cost' are clearly not viewed by experts in 2011/12 as being of significant importance. This is because there has now been a convergence in the use of languages towards OWL (mostly OWL-DL) for deploying ontologies on the Web and open source tooling is now highly prevalent. Another significant difference between the ONTOMETRIC approach and that given in this study was that this study provided one or more possible qualitative and/or quantitative measures for each of the model criteria. Although ONTOMETRIC had a lexical-based scale to rate each ONTOMETRIC characteristic (which was converted to a number during AHP calculations) the actual measure being evaluated in most cases was not stated.

The last contributions made in this thesis relate to matters of practice and therefore are in the form of guidance, particularly for novices in ontology selection and evaluation. In addressing the remaining component of RQ1.2.1, regarding the types of ontology evaluation methods being used by the experts in this study, it was found that most (77%) experts were using methods developed by their own communities, rather than methods with an academic origin. There was, however, overlap between the evaluation criteria cited by experts and those reported earlier in this thesis from the literature. Most experts appeared to confuse an 'evaluation method', with 'evaluation criteria'. Outside of the methods cited with academic origin, most clarity regarding methodology emerged from the Solar Terrestrial Physics discipline. The methods used in this discipline closely paralleled that reported by Annamalai and Sterling (2003), although the method characterised was not explicitly modelled on their technique. Significantly, experts generally were not using highly complex and difficult to apply evaluation methods (e.g., OntoClean (Guarino and Welty, 2004)), criteria or measures. In this sense it would appear that there is a disparity between what is often outlined as 'best practise' in the literature and what is actually being performed in practice.

Data emerging from expert interviews also covered many matters associated with ontology and community governance. A recurring theme was that it was important, for example, that communities do not try to "govern" and hence be the "point of truth" for ontologies for which they neither have the mandate, nor resources to manage. A 'mandate' was considered something bestowed upon a community by a higher organising body (such as a scientific union or standards institution), or could be inherent by virtue of the business mandate of the community in question (e.g., institutions participating in delivering outcomes prescribed by water legislation would have a natural mandate to govern hydrologic cycle-related ontologies). "Institutional backing" with concomitant resources was a recurring characteristic of all successful semantic-enablement activity. Evidence of an appropriate "governance framework" with components as described in Chapter 7 (i.e., guiding principles; declared roles; review mechanisms; suitable ontology repository; community wiki and tooling) was considered an important factor by experts in this study in terms of both establishing credible semantic activities and for evaluating whether to re-use another community's ontologies.

8.3 Methodological Limitations

It has already been mentioned in this chapter that in pursuing the research in this thesis, operational constraints and other factors potentially lead to methodological limitations. These are now summarised.

Most of the work related to the development of the Feature Catalogue artefact and its associated REST-based resource access methods were prototypically implemented. A comprehensive implementation of operational systems was not performed and therefore evaluation of outcomes was confined, in a number of cases, to desk-top exercises that referenced the contexts in which the artefacts would be applied, in order to demonstrate their utility. Ideally these artefacts need to be executed in operational scenarios because in doing so, issues will come to light that were not evident in the confined, artificial environment of a prototyping activity. However, the prototypes do offer proof-of-concept, which was the primary aim of this research. A key area in which the Feature Type Catalogue needs further work is in terms of how it advertises its content. Recent work currently being performed as part of the Open Ontology Repository Initiative on this topic provides advances that could be now be incorporated into an operational version of the Feature Catalogue.

In conducting the selection and evaluation research with national and international experts, there was a relatively high (approximately 50%) drop-out rate at the time of conducting the pair-wise comparison exercise. This not only affected the degree of confidence that could be placed in the preference results (as being reflective of the original expert cohort) but it also affected the actual methods used in the study. Originally a Delphi approach (Schmidt, 1997) was favoured in order to both bring down the level of 'within' expert inconsistency and to draw the experts towards a consensus outcome for the derivation of preferences for model elements. This approach had to be abandoned in favour of a more mechanistic method, involving adjustment of preferences with concurrent feedback between the author and individual experts until inconsistency levels were at relatively acceptable levels (albeit still above the Saaty (1980) upper recommended CR level of 0.10). The added survey burden of the iterative Delphi approach was considered too risky in terms of causing the loss of more participants. Absent a group consensus, a computed (geometric mean) result was used instead to represent "group-think". Without the Delphi approach, the group result, although broadly representative of expert preferences, contained a relatively high degree of 'between expert' variance. For example, MDS (Steyvers, 2002) was applied to the entire dataset for all experts to assess similarities/dissimilarities between experts and there were no detectable similarity groupings between experts, with respect to their preference weights. The MDS stress value (0.125) indicated the computed outcome was a "fair" representation of the data (Kruskal and Wish, 1978).

Given the situation described above, more confidence could have been placed in the resultant preference data had expert inconsistencies been less than the (Saaty, 1980) CR 0.10 suggested upper limit and if the Delphi approach had been possible to apply, which by definition would have lowered the between expert heterogeneity. The loss of expertise during this part of the study considerably lowers confidence in the outcomes of the pair-wise comparison results.

It was also speculated whether experts might have arranged model elements differently to the author. Since model validation was not an explicit step, independent of providing the model at the time of pair-wise comparison, additional research was suggested to explore this issue.

8.4 Further Research

Given the levels of 'within expert' inconsistency found (without justifiable explanation) and the lack of expert validation of the model structure as an explicit step in this study (as described above), informative follow-up research would be to test the hierarchical evaluation model structure on a new group of experts. They could be asked if the model addresses criteria of interest and whether it appropriately groups and nests issue of 'similarity' from an expert's decision-making perspective. By splitting the group into two and then asking both groups to perform the same task, differences in model construction between the two groups and the resultant effects that this has on the estimation of local and global weights, between groups could be assessed. The variability of 'within' group responses could also be calibrated against this study's results. The similarity (or dissimilarity) between the pair of models could be described based on commonality in variables, using similarity coefficients (Gower, 1985). This would be a useful empirical test of the model's structure and would also add value to the body of literature on expert behaviour during multi-criteria decision making encompassing human conceptualisation and categorisation.

In investigating how and where domain ontological content, drawn from the Feature Type Catalogue could be injected to semantically-enable scientific data exchange infrastructure, some deficiencies were also noted that would benefit from additional research. In many cases links to ontological content were suggested through using a SAWSDL "modelReference" element within XML encodings, or by using an "extended Xlink". Currently included constructs within these standards are not particularly expressive and the author could not see how they could be used to provide an adequate, machine-interpretable description of the resource being referenced such that a machine client could select the best type of resource for its purpose. For example, if a resource could be provided in multiple formats and at differing levels of granularity (or expressiveness) and each of these options was linked and available, it is not clear how this type of information could be encoded within the "modelReference" or "extended Xlink" elements. Part of the problem can be addressed through the use of HTTP header "Accept Language" elements (as described in Chapter 5), but this does not solve all parts of the problem.

The issue of ontological content referencing and de-referencing is a topic that needs more research and exploration and will no doubt become very important as more communities attempt to make ontological content a component of their data networks.

8.5 Summary

As the pace of semantic-enablement advances there will be an increasing need for the development of easy-to-apply tools and reliable methods that arm the scientific domain expert with a ready means to quickly and efficiently mark-up his or her data with semantic annotation. These tools must fit the contexts in which scientific domain experts operate and be matched to the skills that these experts possess. People building today's scientific data infrastructure are primarily domain-based data managers and software engineers, not professional ontologists. It is unlikely that this situation will change significantly in the near future. Many of the standards that are, or have been used to build scientific infrastructure focus almost exclusively on the syntactic interoperability of Web services and associated software. Already sunk costs in these infrastructures, which are in many cases considerable, will mean that semantic-enablement must be achieved by 'bolting' components on to what already exists or by adapting the components already in use.

The research in this thesis, which in the case of the ontologically-grounded Feature Catalogue, has taken an existing component of the OGC-standards stack (i.e., the Feature Catalogue) and suggested ways in which it can be better configured to both suit and drive development of OGC-standards - based semantic infrastructure. By asking those relatively few experts who are already engaged in semantically-enabling scientific data infrastructure how they are choosing the ontological components that they are using and by compiling this information into a framework that can be re-used by others, it is hoped that the barriers to ontology use have been lowered.

Glossary Of Terms

Axiom: Axioms define and delimit the realm of analysis. An axiom is a logical statement that is assumed to be true. It is any mathematical statement that serves as a starting point from which other statements are logically derived.

Conceptual model: A conceptual model represents 'concepts' (entities) and the relationships between them.

Data Infrastructure: Is a digital infrastructure promoting data sharing and consumption.

Data model: A data model is an abstraction of the real world which incorporates only those properties thought to be relevant to the application at hand. A data model would normally define specific groups of entities, and their attributes and the relationships between these entities. A data model is independent of a computer system and its associated data structures.

Domain: Community context: A sphere of activity, concern, or function; a field of interest or discourse. Application context: The body of knowledge defining the range and scope of an application in terms of elements, rules and behaviours.

Earth system science: A scientific endeavour that seeks to integrate various fields of academic study to understand the Earth as a system. It considers interaction between the atmosphere, hydrosphere, lithosphere, biosphere and heliosphere.

Eigenvector: A special set of vectors associated with a linear system of equations (i.e., a matrix equation) that are sometimes also known as characteristic vectors, proper vectors, or latent vectors.

Encoding: The activity of converting data into code, such as converting spatial coordinates into XML.

Epistemic: Pertaining to knowledge or the conditions for acquiring it.

Evaluation Criteria: General qualities that an ontology ought to possess.

Evaluation Measure: Quantifies some aspect of an ontology.

Framework: An information architecture perspective: In terms of software design, a reusable software template, or skeleton, from which key enabling and supporting services can be selected, configured and integrated with application code. A conceptual perspective: is used in research to outline possible courses of action, or to present a preferred approach to an idea or thought.

GML: OGC`s XML-based language for describing and encoding geospatial information. GML is an application of XML, a specification developed by members of the Open GIS Consortium. http://www.opengis.org/techno/specs/00-029/GML.html ". GML is an XML encoding for spatial data. In a sense, it is a schema-writing language for spatial information.

Heuristic: Refers to experience-based techniques for problem solving, learning, and discovery.

Inference: Is the act or process of deriving logical conclusions from premises known or assumed to be true.

Inhere: To exist permanently and inseparably in, as a quality, attribute, or element.

Interface: A named set of operations that characterise the behaviour of an entity. For a given distributed computing technology, an interface is an implementation of one or more operations that include the syntax of the interaction between two functional entities. An interface is a shared communication boundary between two functional entities.

Interoperability: Capability to communicate, execute programs, or transfer data among various functional units in a manner that requires the user to have little or no knowledge of the unique characteristics of those units ISO 2382-1. "The ability for a system or components of a system to provide information portability and inter-application, cooperative process control. Interoperability, in the context of the OpenGIS Specification, is software components operating reciprocally (working with each other) to overcome tedious batch conversion tasks, import/export obstacles, and distributed resource access barriers imposed by heterogeneous processing environments and heterogeneous data.

Linked data: Is a method of publishing structured data so that it can be interlinked online. Linked Data principles include:
1. Use URIs as names for things,
2. Use HTTP URIs so that people can look up those names,
3. When someone looks up a URI, provide useful RDF information,
4. Include RDF statements that link to other URIs so that they can discover related things.

Matrix: Is a rectangular array of numbers, symbols, or expressions. The individual items in a matrix are called its elements or entries.

Metadata: Metadata (and its constituent elements) describes an information resource, or helps provide access to an information resource.

Multi-curve coverage type: Coverage characterized by a finite spatial domain consisting of curves. Often the curves represent features such as roads, railroads or streams.

Operation: Specification of a transformation or query that a service may be called to execute.

Semantic: The study of meaning.

Semantic enablement: Web Services are not built for their own sake but to encapsulate data or processing models. To exchange data between services, i.e., to make them interoperable, they have to share common schemas or translate between them. Semantic annotations linking feature types or instances to explicit and shared conceptualizations support the clarification of vocabularies used and negate ambiguities. Embedded reasoning engines, operating on embedded semantics can help automate search, retrieval and workflow processing. All of these components, established within a data infrastructure add up to semantic enablement.

Service: Distinct part of the functionality (as expressed in operations) that is provided by an entity through interfaces.

Service Catalogue: OGC Service Catalogues include indexed listings of feature collections (datasets), their contents, their coverages, and other metadata. They register the existence, location, and description of feature collections (wrapped in services) held by an Information Community. Service catalogues provide the capability to add and delete entries. At a minimum a service catalogue will include the name for the feature collection (datasets) and the location handle (service URL) that specifies where this data may be found.

Spatial data Infrastructure: A spatial data infrastructure (SDI) is a data infrastructure implementing a framework of geographic data, metadata, users and tools that are interactively connected in order to use spatial data in an efficient and flexible way.

Triangulation: Is used to indicate that more than two types of methods are used in a study with a view to double (or triple) checking results/conclusions.

UML: Unified Modeling Language is a standardized general-purpose modelling language in the field of object-orientated software engineering.

Validation: The process of testing an application, system or method to ensure that it conforms to a specification, set of rules or principles.

Web Service: The OGC follows a definition of a Web service as originally proposed by IBM, Motorola and others: Web Services are self-contained, self-describing, modular applications that can be published, located, and invoked across the Web. Web services perform functions, which can be anything from simple requests to complicated business processes. Once a Web service is deployed, other applications (and other Web services) can discover and invoke the deployed service.

Bibliography

AADC, 2012a. *SCAR Feature Catalogue.* [online] Available from: http://data.aad.gov.au/aadc/ftc [Accessed 16 March 2012].

AADC, 2012b. *Australian Antarctic Data Centre-Home Page.* [online] Available from: http://data.aad.gov.au/ [Accessed 11 February 2012].

AADC, 2012c. *Prototype REST Services For New Feature Catalogue.* [online] Available from: http://data.aad.gov.au/aadc/Featurecatalogue/samples.cfm [Accessed 10 January 2012].

AADC, 2012d. *SKOS Iterative Feature Catalogue Service* [online] Available from: http://data.aad.gov.au/aadc/FeatureCatalogue/index2.cfm?/IterativeCatalogueService/Catalogue/Format=skos [Accessed 10 January 2012].

Abramowitz, M. and Stegun, I.A., eds. 1972. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables,* 9th ed. New York: Dover.

Adida, B., Birbeck, M., McCarron, S., and Pemberton, S., 2008. *RDFa in XHTML: Syntax and Processing. A collection of attributes and processing rules for extending XHTML to support RDF.* W3C Recommendation 14 October 2008. [online] Available from: http://www.w3.org/TR/rdfa-syntax/ [Accessed 20 March 2009].

Aduna, 2012*. OpenRFD.org-home of Sesame Home Page.* [online] Available from: http://www.openrdf.org/ [Accessed 19 February 2012].

ALA, 2012. *Atlas of Living Australia-Home Page.* [online] Available from: http://www.ala.org.au/ [Accessed 17 March 2012].

Alani, H. and Brewster, C., 2005. Ontology Ranking Based on the Analysis of Concept Structures. *In: Third International Conference On Knowledge Capture (K-Cap*), 2-5 October 2005 Banff, Alberta, Canada. 51-8.

Allocca, C., d'Aquin, M., and Moaa, E., 2009. DOOR: Towards a formalization of ontology relations. *In: International conference on knowledge engineering and ontology development KEOD 2009,* 6-8 October 2009 Madeira, Portugal.

AMD, 2012. *Antarctic Master Directory-Home Page.* [online] Available from: http://gcmd.gsfc.nasa.gov/KeywordSearch/Home.do?Portal=amd&MetadataType=0 [Accessed 21 February 2012].

Anda,B. and Sjøberg, D.I.K., 2002. Towards an inspection technique for use case models, *In: Proceedings of the 14th international conference on Software engineering and knowledge engineering* , 15-19 July 2002 Ischia Island, Italy. 127-34.

Anderson, J. D. and Pérez-Carballo, J., 2001. The nature of indexing: how humans and machines analyse messages and texts for retrieval. Part I: Research, and the nature of human indexing. *Information Processing and Management* 37(2), 231–254.

ANDS, 2010. *ANDS Metadata Workshop 08-11-2010 Notes.* [online] Available from: http://www.ands.org.au/events/metadataworkshop08-11-2010/index.html [Accessed 24 October 2011].

Annamalai, M. and Sterling, L., 2003. Guidelines for Constructing Reusable Domain Ontologies. *In:* S. Cranefield, T. Finin, V. Tamma and S. Willmott, eds. *AAMAS03 Workshop on Ontologies in Agent Systems, CEUR Workshop Series, Vol. 73.* 15 July 2003 Melbourne, Australia. 71-4.

ANZLIC, 2012. *ANZLIC-Home Page.* [online] Available from: http://www.anzlic.org.au/ [Accessed 27 January 2012].

AODC, 2012. *Australian Oceanographic Data Centre.* [online] Available from: http://www.aodc.gov.au/ [Accessed 21 February 2012].

AODCJF, 2008. *Marine Community Profile Of ISO 19115 Version 1.4*. [online] Available from: http://www.aodc.gov.au/files/MarineCommunityProfilev1.4.pdf [Accessed 9 February 2009].

Armstrong, J. S., 2001. Judgemental bootstrapping: inferring experts rules for forecasting. *In: Principles of forecasting: a handbook for researchers and practitioners*. Boston: Kluwer Academic Publishers, 171–192.

Arora, N., Westenthaler, R., Behrendt, W., and Gangemi, A., 2006. Information Object Design Pattern for Modeling Domain Specific Knowledge. *In: 1st ECOOP Workshop on Domain-Specific Program Development (DSPD)in conjunction with ECOOP,* 3 July 2006 Nantes, France.

Baader, F., Calvanese, D., McGuinness, D., Nardi, D., and Patel-Schneider, P., eds. 2003. *The Description Logic Handbook Theory*, Cambridge : Cambridge University Press.

Baader, F., Horrocks, I., and Sattler, U., 2004. Description Logics. *In:* S. Staab and R. Studer, eds. *Handbook on Ontologies*, Springer. 3-28.

Baclawski, K. and Schneider, T., 2009. The open ontology repository initiative: Requirements and research challenges. *In: Proceedings of Workshop on Collaborative Construction, Management and Linking of Structured Knowledge at the ISWC 2009,* 25 October 2009 Washington DC, USA.

BADC, 2012. NERC Data Grid-Discovery Page. [online] Available from: http://ndgbeta.badc.rl.ac.uk/discovery [Accessed 27 March 2012].

Baker, T.L., 1994. *Doing Social Research.* 2nd ed. New York: McGraw-Hill Inc.

Baker, D., Bridges, D., Hunter, R., Johnson, G., Krupa, J., Murphy, J., and Sorenson, K., 2002. *Guidebook to Decision- Making Methods* [online], WSRC-IM-2002-00002, Department of Energy, USA. Available from: http://emi-web.inel.gov/Nissmg/Guidebook_2002.pdf [Accessed 18 February 2012].

Bakillah, M. and Mostafavi, M.A**.,** 2010**.** A real time semantic interoperability framework in ad hoc network of geospatial databases: disaster management context**.** *In: the 2010 Canadian Geomatics Conference and Symposium of Commission I, ISPRS Convergence in Geomatics – Shaping Canada's Competitive Landscape Volume XXXVIII - Part 1,* 15-18 June 2010 Calgary, Alberta, Canada.

Barbour, R. S., 1998. Mixing qualitative methods: Quality assurance or qualitative quagmire? *Qualitative Health Research*, 8(3), 352-361.

Bard, J.B.L., 2005. Anatomics: the intersection of anatomy and bioinformatics. *Journal of Anatomy* [online], 26(1), 1–16. Available from: http://onlinelibrary.wiley.com/doi/10.1111/j.0021-8782.2005.00376.x/abstract [Accessed 1 April 2011].

Barnaghi, P., Ganz, F., Abangar, H., Presser, M., and Moessner, K., 2011. Sense2Web: A Linked Data Platform for the Semantic Sensor Web, Submitted to *Semantic Web Journal* June 2011.

Barsalou, L. W., 1985. Ideals, central tendency, and frequency of instantiation as determinants of graded structure in categories, *Journal of Experimental Psychology: Learning, Memory, and Cognition.* 11(4), 629–654.

Batcheller, J.K., Reitsma F., and Gittings B.M., 2007. Implementing ISO-compliant Feature Metadata. *In: GISRUK 2007 Proceedings,* 11th - 13th April 2007 NUI Maynooth, Ireland.

Batcheller, J.K. and Reitsma, F.E., 2010. Implementing Feature Level Semantics For Spatial Data Discovery: supporting the reuse of legacy data using open source components. *Computers, Environment and Urban Systems,* 34(4), 333-344.

Baumann, P., 2010. *WCS 2.0 Overview: Core and Extensions.* Open Geospatial Consortium Best Practise Document 09-110r3. [online] Available from: http://portal.opengeospatial.org/files/?artifact_id=36266&version=2...pdf [Accessed 4 May 2011].

Bazerman, M.H. 1998. *Judgement in Managerial Decision Making*. 4th ed. New York: John Wiley.

Bechhofer, S., Carr, L., Goble, C.A., Kampa, S., and Board, T.M., 2002. The Semantics of Semantic Annotation. *In: On the Move to Meaningful Internet Systems, 2002 - DOA/CoopIS/ODBASE 2002 Confederated International Conferences DOA, CoopIS and ODBASE 2002* [online], 30 October – 1 November 2001 Irvine, California, USA. Available from: http://eprints.ecs.soton.ac.uk/6978/ [Accessed 23 December 2010].

Bechhofer, S., 2010. SKOS: Past, Present and Future and a little bit of history, architecture and engineering. *In: 7th Extended Semantic Web Conference*, 30 May – 3 June 2010 Heraklion, Greece.

Beckett, D., 2004. *RDF/XML syntax specification (revised).* W3C Recommendation REC-rdf-syntax-grammar-20040210, February 2004. [online] Available from: http://www.w3.org/TR/ rdf-syntax-grammar/ [Accessed 18 July 2008].

Belgian Biodiversity Platform, 2012. *RAMS Taxon Search.* [online] Available from: http://www.scarmarbin.be/rams.php [Accessed 20 January 2010].

Belton, V. and Gear, T., 1983. On a short-coming in Saaty's method of analytic hierarchies, *Omega*, 11, 228–30.

Belton, V. and Stewart, T. J., 2002. *Multiple Criteria Decision Analysis: An Integrated Approach*. Boston: Kluwer Academic Publishers.

Bergman, MK. and Giasson, F., 2012. *UMBEL Reference Concept Ontology and Vocabulary.* [online] Available from: http://www.umbel.org/ [Accessed 12 February 2012].

Berkley, C., Bowers, S., Jones, M.B., Madin, J.S., and Schildhauer, M., 2009. Improving Data Discovery in Metadata Repositories through Semantic Search. *In: International Conference on Complex, Intelligent and Software Intensive Systems - CISIS*, 16-19 March 2009 Fukuoka, Japan, 1152-9.

Bernard, L., Einspanier, U., Haubrock, S., Hubner, S., Klien,E., Kuhn, W., Lessing, R., Lutz, M., and Visser, U., 2004. Geotechnologien Science Report Ontology-Based Discovery and Retrieval of Geographic Information in Spatial Data Infrastructures. *In: Information Systems in Earth Management-Status Seminar* [online], 23-24 March 2004 RWTH Aachen University, Germany. Available from: http://www.delphiimm.de/meanings/index_eng.html. [Accessed 22 March 2011].

Berners-Lee, T., Hendler, J., and Lassila, O., 2001. The Semantic Web. *Scientific Am*, May 2001, 34–43.

Berners-Lee, T., 2006. *Linked Data-Design Issues.* W3C.[online] Available from: http://www.w3.org/DesignIssues/LinkedData.html. [Accessed 18 December 2010].

Bishr, Y., 1998. Overcoming the Semantic and other Barriers to GIS Interoperability. *International Journal of Geographical Information Science,* 12(4), 299-314.

Bizer, C. and Cyganiak, R., 2006. D2R Server – Publishing Relational Databases on the Semantic Web. *In: 5th International Semantic Web Conference (ISWC2006)* [online], 5-9 November 2006, Athens GA. Available from: http://richard.cyganiak.de/2008/papers/d2r-server-iswc2006.pdf [Accessed 13 July 2011].

Blomqvist, E., Öhgren, A., and Sandkuhl, K., 2006, Ontology Construction in an Enterprise context: Comparing and Evaluating two Approaches. *In: Proceedings of 8th International Conference on Enterprise Information Systems,* 23-27 May 2006 Paphos, Cyprus. 86-93.

Blue Ribbon Task Force on Sustainable Digital Preservation and Access and Smith Rumsey, A., 2010. *Sustainable economics for a digital planet: ensuring long-term access to digital information: final report of the Blue Ribbon Task Force on Sustainable Digital Preservation and Access*. [online] Available from: http://brtf.sdsc.edu/biblio/BRTF_Final_Report.pdf [Accessed 2 February 2011].

Boag*, S,* Chamberlin*, D,* Fernández*, M*F, Florescu, D, Robie, J and Siméon, J., 2010. *XQuery 1.0: An XML Query Language (Second Edition).* W3C Recommendation 14 December 2010. [online] Available from: http://www.w3.org/TR/xquery/ [Accessed 10 February 2011].

BODC, 2011. *NERC Data Grid.* [online] Available from: http://www.bodc.ac.uk/projects/uk/ndg/ [Accessed 20 March 2011].

Bogdan, R.C. and Biklen, S.K. 1982. *Qualitative Research for Education: An Introduction to Theory and Methods,* London: Allyn & Bacon Inc.

Booch, G., Rumbaugh, J., and Jacobson, I., 1999. *The unified modeling language user guide.* Reading, MA: Addison-Wesley.

Borges, JL., 1964*. Other inquisitions, 1937-1952*. Translated by Ruth L.C. Simms. Introd. By J.E. Irby, Austin, University of Texas Press.

Botts, M., 2007. *OpenGIS Sensor Model Language (SensorML) Implementation Specification*. OGC 07-000. Version: 1.0.0. [online] Available from: http//portal.opengeospatial.org/files/?artifact_id=21273 [Accessed 1 June 2009].

Boyatzis, R.E., 1998. *Transforming Qualitative Information: Thematic Analysis and Code Development.* Thousand Oaks, CA: Sage.

Bragge, J., Korhonen, P., Wallenius, H., and Wallenius, J., 2010. Bibliometric Analysis of Multiple Criteria Decision Making/Multiattribute Utility Theory. *In:* M. Ehrgott, B. Naujoks, T. Stewart, and J. Wallenius, eds. *Multiple Criteria decision making for sustainable energy and transportation systems.* Berlin, Germany: Springer-Verlag, 259–268.

Brank, J., Mladenic, D., and Grobelnik, M., 2006. Gold standard based ontology evaluation using instance assignment. *In: Proceedings of the 4th Workshop on Evaluating Ontologies for the Web (EON2006),* 22 May 2006 Edinburgh, United Kingdom.

Bray, T., Paoli, J., Sperberg-McQueen, C.M., Maler, E., and Yergeau F., 2008. *Extensible Markup Language (XML) 1.0 (5th Edition).* W3C Recommendation.

Brewer, J. and Hunter, A., 1989. *Multimethod research: A synthesis of styles*. Newbury Park, CA: Sage.

Brewster, C., Alani, H., Dasmahapatra, S.,and Wilks, Y., 2004. Data driven ontology evaluation. *In: Proc. of Intl. Conf. on Lang. Resources and Eval.*, 26-28 May 2004 Lisbon, Portugal.

Brickley, D. and Guha, R.V., 2004. *RDF vocabulary description language 1.0: RDF Schema.* W3C Recommendation 10 February 2004. [online] Available from:http://www.w3.org/TR/rdf-schema/ [Accessed 23 March 2009].

Bro, R., Acar, E., and Kolda, T., 2008. Resolving the Sign Ambiguity in the Singular Value Decomposition. *Journal of Chemometrics,* 22(2), 135-140.

Brodaric, B. and Probst, F., 2009. Enabling Cross-Disciplinary E-Science BY Integrating Geoscience Ontologies With DOLCE. *IEEE Intelligent Systems*, 24(1), 66-77

Brodeur, J., Bédard, Y., Edwards, G., and Moulin, B., 2003. Revisiting the Concept of Geospatial Data Interoperability within the Scope of Human Communication Process. *Transactions in GIS*, 7(2), 243-265.

Broring, A., Janowicz, K., Stasch, C., and Kuhn, W., 2009. Semantic Challenges for Sensor Plug and Play. *In: Proceedings of the 9th International Symposium on Web and Wireless Geographical Information Systems*, 7-8 December 2009 Maynooth, Ireland. Heidelberg:Springer-Verlag Berlin, 72-86.

Broring A, Maue P, Janowicz K, Nust D and Malewski C., 2011. Semantically-Enabled Sensor Plug & Play for the Sensor Web. *Sensors,* 11(8), 7568-7605.

Brugha, C., 1998. Structuring and Weighting Criteria in Multi Criteria Decision Making (MCDM), *In:* T.J. Stewart and R.C. Van den Honert, eds. *Trends in Multicriteria Decision Making: Proceedings of the 13th International Conference on Multiple Criteria Decision Making,* 6-10 January 1997 Cape Town, SA. Berlin: Springer-Verlag, 229-242.

Brumfiel, G., 2011. High-energy physics: Down the petabyte highway. *Nature,* 469, 282-283.

Bügel ,B. and Hilbring, D., 2007. Application of Semantic Services in ORCHESTRA **.** *In: International Symposium on Environmental Software Systems ISESS 2007.* 22-25 May 2007 Prague.

Buitelaar, P., Eigner, T., and Declerck, T., 2004. OntoSelect: A Dynamic Ontology Library with Support for Ontology Selection. *In: Proc. of the Demo Session at the International Semantic Web Conference,* Nov 2004 Hiroshima, Japan.

Burgman, M., Fidler, F., McBride, M., Walshe, T., and Wintle, B., 2006. *Eliciting Expert Judgments: Literature Review.* ACERA Project 0611 [online]. Available from: http://www.acera.unimelb.edu.au/materials/endorsed/0611.pdf [Accessed 3 May 2009].

Calvanese, D., Giacomo, G.D., and Lenzerini, M., 2001. A framework for ontology integration*. In: Proceedings Of The First Internationally Semantic Web Working Symposium (SWWS),* 30 July-1 August 2001, Stanford, CA.

Catapano, T., Hobern, D., Lapp, H., Morris, R.A., Morrison, N., Noy, N., Schildhauer, M., and Thau, D., 2011. *Recommendations for the Use of Knowledge Organisation Systems by GBIF*, Copenhagen Global Biodiversity Information Facility. [online]. Available from: http://links.gbif.org/gbif_kos_whitepaper_v1.pdf [Accessed 5 July 2011].

CCSDS, 2002. *Reference Model for an Open Archival Information System (OAIS). CCSDS 650.0-B-1, Blue Book.* Washington: NASA. [online] Available from: http://public.ccsds.org/publications/archive/650x0b1.PDF [Accessed 2 November 2011].

Celik, T. and Suda, B., 2012a. *hCard 1.0 Microformats Wiki.* [online] Available from: http://microformats.org/wiki/hcard [Accessed 4 January 2012].

Celik, T. and Suda, B., 2012b. *hCalendar 1.0 Microformats Wiki.* [online] Available from: http://microformats.org/wiki/hcalendar [Accessed 6 January 2012].

CGI, 2011. *GeoSciML Documentation and Download Page.* [online] Available from: http://www.cgi-iugs.org/tech_collaboration/geosciml.html [Accessed 2 October 2011].

Champion, M., Ferris, C., Newcomer, E., and Orchard, D., 2002. *Web Services Architecture*, W3C Working Draft. 14 November 2002. [online] Available from: http://www.w3.org/TR/2002/WD-ws-arch-20021114/ [Accessed 1 September 2008].

Chandrasekaran, B., Josephson, J.R., and Benjamins, V.R., 1998, *Ontology of Tasks and Methods.* Ohio-State University. [online] Available from: http://ksi.cpsc.ucalgary.ca/KAW/KAW98/chandra/index.html. [Accessed 12 October 2010].

Chang, A., 2012. *Kendall's W for Ranked Data Program.* [online] Available from: http://www.stattools.net/KendallW_Pgm.php [Accessed 12 January 2012].

Chen, C.C., Thakkar, S., Knoblock, C.A. and Shahabi. C., 2003. Automatically Annotating and Integrating Spatial Datasets. In: *Proceedings 8th International Symposium on Spatial and Temporal Databases (SSTD),* 24-27 July 2007 Santorini Island, Greece, Springer- Verlag, 469-488.

Chen, M.F., Tzeng, G.H., and Ding, C.D., 2008. Combining fuzzy AHP with MDS in identifying the preference similarity of alternatives. *Applied Soft Computing*, 8(1), 110–117.

Cherryholmes, C.C., 1992. Notes on Pragmatism and scientific realism. *Educational Researcher*, 21, 13-17.

Christensen, E., Curbera, F., Meredith, G., and Weerawarana, S., 2001. *Web Services Description Language (WSDL) 1.1,* W3C Note 15 March 2001. [online] Available from: http://www.w3.org/TR/2001/NOTE-wsdl-20010315 [Accessed 1 September 2008].

Clark, J., 1999. *XSL Transformations (XSLT) Version 1.0.* W3C Recommendation 16 November 1999. [online] Available from: http://www.w3.org/TR/xslt [Accessed 23 January 2008].

Clark, J. and DeRose, S., 1999. *XML Path Language (XPath) Version 1.0.* W3C Recommendation 16 November 1999. [online] Available from: http://www.w3.org/TR/xpath [Accessed 23 January 2008].

Clark and Parsia, 2011. *Pellet-OWL 2 Reasoner For Java.* [online] Available from: http://clarkparsia.com/pellet/ [Accessed 30 January 2012].

Cleverdon, C., 1984. Optimizing convenient online access to bibliographic databases. *Information Services and Use*, 4(1), 37–47.

CMAR, 2010. *Codes For Australian Aquatic Biota-Home Page.* [online] Available from: http://www.cmar.csiro.au/caab/ [Accessed 20 December 2010].

Cockburn, A., 1997. Structuring Use Cases with Goals, *Journal of Object-Oriented Programming*, Sep-Oct, 1997 and Nov-Dec, 1997.

Cockburn, A., 2000. *Writing Effective Use Cases.* Addison-Wesley. Available from: http://www.usecases.org/writing/books/writingeffectiveusecases/writingeffectiveucs.pdf [Accessed 16 January 2012].

Cömert, C., Ulutaş, D., Akıncı, H., and Kara, G., 2010. Semantic web services for implementing national spatial data infrastructures. *Scientific Research and Essays*, 5(7), 685–692.

Corcho, O., Gómez-Pérez, A., González-Cabero, R., and Suárez-Figueroa, M.C., 2004. ODEval: a Tool for Evaluating RDF(S), DAML+OIL,and OWL Concept Taxonomies. *In: IFIP WG12.6 -- First IFIP Conference on Artificial Intelligence Applications and Innovations (AIAI2004),* 22-27 August 2004 Toulouse, France. 369-382.

Corcho, O., Fernandez-Lopez, M., and Gomez-Perez, A., 2006. Ontological engineering: Principles, methods, tools and languages. *In:* C. Calero, F. Ruiz, and M. Piattini eds. *Ontologies for Software Engineering and Software Technology,* New York: Springer, 1-48.

COSMO, 2012. *COSMO Common Semantic Model.* [online] Available from: http://ontolog.cim3.net/cgi-bin/wiki.pl?COSMO [Accessed 15 March 2012].

Courtot, M., Gibson, F., Lister, A.L., Malone, J., Schober, D., Brinkman, R.R. and Ruttenberg, A., 2011. MIREOT: The minimum information to reference an external ontology term. *Appl. Ontol.,* 6(1), 23-33

Cox, M.A.A., 2009. Multidimensional Scaling as an Aid for the Analytic Network and Analytic Hierarchy Processes. *Journal of Data Science,* 7(3), 381-396.

Cox, S., 2006. *Observations and Measurements*. OGC® 05-087r4 Version: 0.14.7. [online] Available from: http://portal.opengeospatial.org/modules/admin/license_agreement.php?suppressHeaders=0&access_license_id=3&target=http://portal.opengeospatial.org/files/?artifact_id=17038 [Accessed 1 December 2009].

Cox, S., 2010a. *Geographic Information: Observations and Measurements OGC Abstract Specification Topic 20.* OGC 10-004r3 and ISO 19156 Version 2.0. [online] Available from: http://portal.opengeospatial.org/files/?artifact_id=41579 [Accessed 12 September 2010].

Cox, S., 2010b. *OGC Identifiers – the case for http URIs.* OGC White Paper. [online] Available from: http://portal.opengeospatial.org/files/?artifact_id=39467 [Accessed 12 September 2010].

Cox, S., 2011. *Observations and Measurements – XML Implementations.* OGC 10-025r1. Version 2.0. [online] Available from: http://portal.opengeospatial.org/files/?artifact_id=41510 [Accessed 15 June 2011].

Crabtree, B. and Miller, W., eds. 1992. *Doing Qualitative Research* London: Sage.

Craglia M., Kanellopoulos I., and Smits P., 2007. Metadata: where we are now, and where we should be going. *In: Proceedings of the 10th AGILE International Conference on Geographic Information Science 2007* [online], Aalborg University, Denmark. Available from: http://people.plan.aau.dk/~enc/AGILE2007/PDF/93_PDF.pdf [Accessed 12 April 2009].

Cranefield, J. and Yoong, P., 2007. To Whom Should Information Systems Research Be Relevant: The Case for an Ecological Perspective. *In:* H. Österle, J. Schelp and R.Winter, eds. *Proceedings of the 15th European Conference on Information Systems (ECIS2007),* 7-9 June 2007 St Gallen, Switzerland, 1313-1324.

Cross, V. and Pal, A., 2006. OntoCAT: An Ontology Consumer Analysis Tool and Its Use on Product Service Categorization Standards. *In: Proceedings of 1st Intl. Workshop on Applications and Business Aspects of the Semantic Web (SEBIZ 2006) at the 5th Intl. Semantic Web Conference (ISWC 2006*), 6 November 2006 Athens, Georgia, USA.

Cruz, I.F. and Sunna, W., 2008. Structural Alignment Methods with Applications to Geospatial Ontologies. *Transactions in GIS*, 12(6), 683–711.

Cryer, P., Hyam, R., Miller, C., Nicolson, N., Ó Tuama, E., Page, R., Rees, J., Riccardi, G., Richards, K., and White, R., 2009. *Adoption of Persistent Identifiers for Biodiversity Informatics.* Recommendations of the GBIF LSID GUID Task Group, 6 November 2009. [online] Available from: http://www2.gbif.org/Persistent-Identifier.pdf [Accessed 4 September 2010].

CSIRO, 2011. *Water Information Research and Development Alliance-Home Page.* [online] Available from: http://www.csiro.au/partnerships/WIRADA.html [Accessed 22 September 2011].

CUAHSI, 2011. *Consortium of Universities for the Advancement of Hydrologic Science-Home Page.* [online] Available from: http://www.cuahsi.org/ [Accessed 2 May 2011].

Cycorp, 2012. *Overview of OpenCyc.* [online] Available from: http://www.cyc.com/opencyc [Accessed 4 February 2012].

Dalkey, N.C. and Helmer, O., 1963. An experimental application of the Delphi method to the use of experts. *Management Science, 9*(3), 458-467.

d'Aquin, M. and Lewen H., 2009. Cupboard-A Place to Expose your Ontologies to Applications and the Community. *In: The Semantic Web: Research and Applications, 6th European Semantic Web Conference, ESWC 2009,* 31 May – 4 June 2009 Heraklion, Crete, Greece. 913-918.

Davidson, E.J., 2005. *Evaluation Methodology Basics.* London: Sage Publications.

Dawes, R.M., 1971. A case study of graduate admission: Application of 3 principles of human decision making. *American Psychologist*, 26(2), 180–188.

DCMI, 2012. *Dublin Core® Metadata Initiative-Home Page.* [online] Available from: http://dublincore.org/ [Accessed 17 March 2012]

de La Beaujardiere, J., 2006. *OpenGIS Web Map Server Implementation Specification.* Version 1.3.0., 15 March 2006. [online] Available from: http://www.opengeospatial.org/standards/wms [Accessed 23 July 2008].

Dean, M. and Schreiber, G., 2003. *Web Ontology Language (OWL) Reference Version 1.0.* Technical report, W3C. [online] Available from: http://www.w3.org/TR/2003/WD-owl-ref-20030221/ [Accessed 12 May 2008].

Deng, H., 1999. Multicriteria analysis with fuzzy pairwise comparisons, *International Journal of Approximate Reasoning,* 21(3), 215-231.

Denzin, K., 1978. *The Research Act.* New York: McGraw-Hill.

DeRose, S., Maler, E., and Orchard, D., 2001. *XML Linking Language (XLink) Version 1.0* W3C Recommendation 27 June 2001. [online] Available from: http://www.w3.org/TR/xlink/ [Accessed 18 August 2007].

Devaraju, A. and Kuhn, W., 2010. A Process-Centric Ontological Approach for Integrating Geo-Sensor Data*. In: 6th International Conference on Formal Ontology in Information Systems (FOIS 2010)*, 11-14 May 2010 Toronto, Canada. 199-212.

DeWalt, K. M. and DeWalt, B. R., 2002. *Participant observation: a guide for fieldworkers.* Walnut Creek, CA: AltaMira Press.

Dey, I., 1993. Qualitative Data Analysis A User Friendly Guide For Social Scientists. New York: Routledge.

Doerr, M., Lagoze, C., Hunter, J., and Baker, T., 2002. *Building Core Ontologies.* White Paper of the DELOS Working Group on Ontology Harmonization. [online] Available from: http://delos-noe.iei.pi.cnr.it/activities/standardizationforum/ontology/Ontologies.pdf [Accessed 10 February 2008].

Dogac, A., Kabak, Y., and Laleci, G.B., 2010. Enriching ebXML Registries with OWL Ontologies for Efficient Service Discovery*. In: Free and Open Source Software For Geospatial Conference (FOSS4G),* 6-9 September 2010 Barcelona, Spain.

Donnelly, M., Bittner, T., and Rosse, C., 2006. A formal theory for spatial representation and reasoning in biomedical ontologies. *Artificial Intelligence in Medicine,* 36(1), 1-27.

DNF, 2012. *Feature Catalogues.* [online] Available from: http://www.dnf.org/registry/feature_catalogues [Accessed 24 February 2012].

Duran, O. and Aguilo, J., 2008. Computer-aided machine- tool selection based on a Fuzzy-AHP approach. *Expert Systems with Applications,* 34 (3), 1787–1794.

Dyer, J.S., Fishburn, P.C., Steuer, R.E., Wallenius, J., and Zionts, S., 1992. Multiple Criteria Decision Making, Multi attribute Utility Theory: The Next Ten Years. *Management Science,* 38, 645-654.

Edwards, W., 1977. How to use multiattribute utility measurement for social decision making. *IEEE Trans. Systems, Man, Cybernetics* (SMC-7), 326–340.

Edwards, W. and Barron, F.H., 1994. SMARTS and SMARTER: Improved simple methods for multiattribute utility measurement. *Organ. Behavior Human Decision Processes,* 60(3), 306–325.

Einhorn, J., 1974. Expert judgment: Some necessary conditions and an example. *Journal of Applied Psychology*, 59(5), 562-571.

Encyclopedia Britannica Online, 2012. *Decision theory.* [online] Available from: http://www.britannica.com/EBchecked/topic/155155/decision-theory [Accessed 10 March 2012].

Escobar, M.T., Aguarón, J., and Moreno-Jiménez, J.M., 2004. A Note on AHP Group Consistency for the Row Geometric Mean Priorization Procedure, *European Journal of Operational Research*, 153(2), 318-322.

European Commission, 2010. *Riding the wave - How Europe can gain from the rising tide of scientific data - Final report of the High Level Expert Group on Scientific Data - October 2010.* Available from: http://cordis.europa.eu/fp7/ict/e-infrastructure/docs/hlg-sdi-report.pdf [Accessed 23 January 2011].

European Commission, 2012. *INSPIRE Geoportal.* [online] Available from: http://inspire-geoportal.ec.europa.eu/ [Accessed 20 March 2012].

Euzenat, J. and Shvaiko, P., 2007. *Ontology Matching*. Berlin: Springer-Verlag.

Evans, J.S.T., Barston, J.L., and Pollard, P., 1983. On the conflict between logic and belief in syllogistic reasoning. *Memory & Cognition,* 11(3), 295-306.

Expert Choice Academics, 2012. *Expert Choice Academics Limited Academic Licence Site.* [online] Available from: http://www.expertchoice.com/academic-program/free-trial [Accessed 13 January 2012].

Farlex, 2012a. *The Free Dictionary-Semantic.* [online] Available from: http://www.thefreedictionary.com/semantic [Accessed 9 April 2012].

Farlex, 2012b. *The Free Dictionary-Framework.* [online] Available from: http://www.thefreedictionary.com/framework [Accessed 15 February 2012].

Farlex, 2012c. *The Free Dictionary-template.* [online] Available from: http://www.thefreedictionary.com/template [Accessed 16 March 2012].

Farrell, J. and Lausen, H., 2007. *Semantic Annotations for WSDL and XML Schema.* W3C Recommendation 28 August 2007. [online] Available from: http://www.w3.org/TR/2007/REC-sawsdl-20070828/ [Accessed 28 November 2009].

Feeley, R., Seiler, P., Packard, A., and Frenklach, M., 2004. Consistency of a reaction dataset. *J. Phys. Chem. A*, 108(44), 9573–9583.

Fensel, D. and Bussler, C., 2002. The Web Service Modelling Framework WSMF, *Electronic Commerce Research and Applications*, 1(2), 113-137.

Fenton, N.E. and Melton, A., 1996. Measurement Theory and Software Measurement. *In:* A. Melton ed. *Software Measurement*, London: International Thomson Computer Press, 27-38.

Fernández, M., Cantador, I., and Castells, P., 2006. CORE: A Tool for Collaborative Ontology Reuse and Evaluation. *In: 4th International Workshop on Evaluation of Ontologies for the Web (EON 2006) at the 15th International World Wide Web Conference (WWW 2006),* May 2006 Edinburgh, UK.

Fernández-Breis, T. J., Aranguren, M. E., and Stevens, R., 2009. A Quality Evaluation Framework for Bio-Ontologies. *In: ICBO: International Conference on Biomedical Ontology*, 24-26 July 2009 Buffalo, NY.

Fielding, R.T., 2000. *Architectural Styles and the Design of Network-based Software Architectures,* [online], Thesis (PhD). University of California, Irvine. Available from: http://www.ics.uci.edu/~fielding/pubs/dissertation/top.htm [Accessed 3 February 2007].

Finney, K., 2007. A "Bottom up" Governance Framework for Developing Australia's Marine Spatial Data Infrastructure (SDI). *Data Science Journal* [online], 6, 64-90. Available from: http://www.informatik.uni-trier.de/~ley/db/journals/datascience/datascience6.html [Accessed 14 September 2008]

Finney, K.T., 2008. Deriving an Australian marine ontology from existing ontological models: a practical evaluation. *In: WALIS International Forum 08,* 12-14 March 2008 Perth W.A.

Finney, K. T., 2009. *Data and Information Management Strategy (DIMS), 2009-2013 (scar Report 34).* [online] Available from: http://www.scar.org/publications/reports/Report_34.pdf [Accessed 25 October 2010].

Finney, K.T. and Watts, D., 2011. REST-based Feature Catalogue Services. *International Journal of Geographical Information Science*, 25(9), 1507-1524.

Fitting, M., 2011. Intensional Logic, *In:* E. Zalta, ed. *The Stanford Encyclopedia of Philosophy (Winter 2011 Edition).* [online] Stanford:Stanford University Available from: http://plato.stanford.edu/archives/win2011/entries/logic-intensional [Accessed 10 January 2010].

Forman, E.H., 1990. Deriving Ratio Level Measures from Verbal Judgments, *George Washington University Working Paper*.

Forman, E. and Peniwati, K., 1998. Aggregating individual judgments and priorities with the Analytic Hierarchy Process. *European Journal of Operational Research*, 108(1), 165−169.

Forman, E.H. and Gass, S.I., 2001. The Analytic Hierarchy Process: an exposition. *Operations Research*, *49*, 469–486.

Fox, P. and Hendler, J., 2009. Semantic eScience: Encoding Meaning in Next-Generation Digitally Enhanced Science. *In:* T. Hey, S. Tansley and K. Tolle, eds. *The Fourth Paradigm – Data Intensive Scientific Discovery.* Washington: Microsoft Research Redmond.

Fox, P., McGuinness, D.L., Cinquini, L., West, P., Garcia, J., Benedict, J.L. and Middleton, D., 2009. Ontology-supported scientific data frameworks: The Virtual Solar-Terrestrial Observatory experience. *Computers and Geosciences* [online], 35(4), 724-738. Available from: http://dx.doi.org/10.1016/j.cageo.2007.12.019 [Accessed 10 April 2011].

Freeman, L.C., 1997. Uncovering Organisational Hierarchies. *Computational & Mathematical Organisation Theory,* 3(1), 5-18.

French, S., 1988. *Decision Theory: an Introduction to the Mathematics of Rationality*, Chichester: Ellis Horwood.

Gallagher, J., Potter, N., Sgouros, T., Hankin, S., and Flierl, G., 2007. *The Data Access Protocol—DAP 2.0. ESE-RFC-004.1.1*. [online] Available from: http://www.opendap.org/pdf/ESE-RFC-004v1.1.pdf [Accessed 2 February 2011].

Gangemi, A., Guarino, N., Masolo, C., Oltramari, A., and Schneider, L., 2002. Sweetening Ontologies with DOLCE. *In:* A. Gomez-Perez and V.R. Benjamins, eds. *Knowledge Engineering and Knowledge Management. Ontologies and the Semantic Web. Proceedings of the 13th International Conference (EKAW 2002),* 1-4 October 2002 Siguenza, Spain. Heidelberg: Springer, 166-181.

Gangemi, A., 2005. *DOLCE-Lite-Plus Ontology Container Information.* [online] Available from: http://www.w3.org/2001/sw/BestPractices/WNET/DLP3941_daml.html [Accessed 20 January 2009].

Gangemi, A., Catenacci, C., Ciaramita, M., and Lehmann, J., 2005. A theoretical framework for ontology evaluation and validation. *In: Semantic Web Applications and Perspectives (SWAP) –2nd Italian Semantic Web Workshop,* 15-16 December 2005 Trento, Italy.

Gangemi, A., 2007. *DOLCE UltraLite OWL Ontology.* [online] for Laboratory for Applied Ontology (ISTC-CNR) ed. Available from: http://www.loa-cnr.it/ontologies/DUL.owl. [Accessed 18 November 2010].

García-Castro, R. and Gómez-Pérez, A., 2011. Perspectives in Semantic Interoperability. *In: International Workshop on Semantic Interoperability IWSI 2011 In conjunction with ICAART 2011,* 20 January 2011 Rome, Italy.

Gass, S. I., 1998. Tournaments, transitivity and pairwise comparison matrices. *J. Oper. Res. Soc.* 49(6) 616–624.

Gass, S.I., 2005. Model World: The Great Debate—MAUT Versus AHP. *Interfaces*. 35(4), 308–312.

Genesereth, M.R. and Fikes, R.E., 1992. *Knowledge Interchange Format, Version 3.0 Reference Manual*. Technical Report Logic-92-1, Computer Science Department, Stanford University.

Genest, C. and Zhang, S.S., 1996. Graphical Analysis of Ratio-Scaled Paired Comparison Data, *Management Science* [online], 42(3), 335-349 Available from: http://www.jstor.org/stable/2634347 [Accessed 23 October 2007].

Génova, G., Llorens, J., Metz, P., Prieto-Díaz, R., and Astudillo, H., 2005. Open Issues in Industrial Use Case Modeling. *The Journal Of Object Technology* [online], 4(6), 7-14. Available from: http://www.jot.fm/issues/issue_2005_08/article1.pdf [Accessed 3 February 2012].

GEO, 2012. *What Is GEOSS: The Global Earth Observing System of Systems.* [online] Available from: http://www.earthobservations.org/geoss.shtml [Accessed 10 January 2012].

GeoNetwork, 2011. *GeoNetwork Developers Manual. Version 2.6.3. January 26 2011.* [online] Available from: http://geonetwork-opensource.org/manuals/2.6.3/developer/GeoNetworkDeveloperManual.pdf. [Accessed 8 February 2011].

GeoNetwork, 2012. *GeoNetwork Opensource-Home Page.* [online] http://geonetwork-opensource.org/ [Accessed 12 January 2012].

Ghazvinian, A., Noy, N.F., and Musen, M. A., 2011. How orthogonal are the OBO Foundry ontologies? *J Biomed Semantics,* 2(sup 2):S2.

Giguere, G., 2007. Collecting and analysing data in multidimensional scaling experiments: A guide for psychologists using SPSS*. Tutorials in Quantitative Methods for Psychology,* 2(1), 26-37.

Gil, Y., Ratnakar, V., and Deelman, E., 2006. Metadata Catalogs with Semantic Representations. *In: International Provenance and Annotation Workshop (IPAW'06),* 3-5 May 2006 Chicago, Illinois.

Giunchiglia, F., Shvaiko, P., and Yatskevich, M., 2007. Semantic Matching: Algorithms and Implementation. *Journal of Data Semantics (JoDS),* 9, 1-38.

Glaser, B.G. and Strauss, A.L., 1967. *The Discovery of Grounded Theory: Strategies for Qualitative Research.* New Jersey: AldineTransaction.

Glaser, B.G., ed., 1995. *Grounded Theory 1984-1994,* Vols 1-2. Mill Valley, CA: Sociology Press.

Glimm, B. and Krötzsch, M., 2010, SPARQL beyond Subgraph Matching. *In: International Semantic Web Conference 2010*, 9-11 November 2011 Shanghai, China. 241-256.

GNU Project, 2012. *The R Project For Statistical Computing-Home Page.* [online] Available from: http://www.r-project.org/ [Accessed 20 March 2012].

Goble, C. and De Roure, D., 2009. The Impact of Workflow Tools on Data-centric Research. *In:* T. Hey, S. Tansley and K. Tolle, eds. *The Fourth Paradigm – Data Intensive Scientific Discovery.* Washington: Microsoft Research Redmond.

Golafshani, N., 2003. Understanding Reliability and Validity in Qualitative Research. *The Qualitative Report* [online], 8(4), 597-607. Available from: http://www.nova.edu/ssss/QR/QR8-4/golafshani.pdf. [Accessed 6 February 2011].

Golbreich, C. and Wallace, E., 2009. *OWL 2 Web Ontology Language: New Features and Rationale.* W3C Recommendation. [online] Available from: http://www.w3.org/TR/owl2-new-features/ [Accessed 1 September 2010].

Gomez-Perez. A., 2001. Evaluation of ontologies. *Intl. Journal of Intelligent Systems*, 16, 391–409.

Gomez-Perez, A., 2004. Ontology evaluation. *In:* S. Staab and R. Studer, eds, *Handbook on Ontologies,* 1st ed. Springer. 251-274.

Gomez-Perez, A., Corcho, O., and Fernandez-Lopez, M., 2004. *Ontological Engineering*. Secaucus, NJ, USA: Springer-Verlag London.

Goodwin, P. and Wright, G., 1998. *Decision Analysis for Management Judgement*, 2nd ed. Chichester: John Wiley.

Gower, J.C., 1977. The Analysis of Asymmetry and Orthoganility, *In*: J. Barra, F. Brodeau and G. Romier eds. *Recent Developments in Statistics,* North Holland, Amsterdam.

Gower, J.C., 1985. Measures of similarity, dissimilarity and distance. *In:* S. Kotz and N.L. Johnson, eds. *Encyclopedia of Statistical Sciences*. New York : John Wiley & Sons.

Gower, J.C. and Hand, D.J., 1996. *Biplots*. London, UK: Chapman & Hall.

Grau, B.C., Horrocks, I., Kazakov Y., Sattler U., 2007. Extracting modules from ontologies: A logic-based approach. *In: Workshop on OWL Experiences and Directions OWLED (2009),* 23-24 October 2009 Chantilly, Virginia. Berlin Heidelberg: Springer, 159-186.

Green, J., Dolbear, C., Hart, G., Goodwin, J., and Engelbrecht, P., 2008. Creating a semantic integration system using spatial data. *In: International Semantic Web conference,* 26-30 October 2008 Karlsruhe, Germany.

Greene, J., Caracelli,V., and Graham, W., 1989. Toward a Conceptual Framework for Mixed-Methods. Evaluation Designs. *Educational Evaluation and Policy Analysis,* 11, 255-274.

Greene, J.C., 2007. *Mixed methods social inquiry.* San Francisco: Jossey-Bass.

Gregor, S. and Jones, D., 2007. The Anatomy of a Design Theory. *Journal For The Association Of Information Systems (AJIS),* 8(5), 312-335.

Gregorio, J., Fielding, R., Hadley, M., Nottingham, M., and Orchard, D., 2010. *URI Template Draft. Internet Engineering Task Force (IETF*). [online] Available from: http://tools.ietf.org/html/draft-gregorio-uritemplate-04 [Accessed 2 January 2011].

Gregory, R., Flynn, J., Johnson, S.M., Satterfield, T.A., Slovic, P., and Wagner, R., 1997. Decision-pathway surveys: A tool for resource managers. *Land Economics,* 73(2), 240-254.

Gregory, R., 2010. Design Science Research and the Grounded Theory Method: Characteristics, Differences, and Complementary Uses. *In: Proceedings of 18th European Conference on Information Systems (ECIS 2010)* [online], 7-9 June 2010 Pretoria. Available from: http://aisel.aisnet.org/ecis2010/44 [Accessed 2 February 2011].

Gruber, T.R., 1995. Towards principles for the design of ontologies used for knowledge sharing. *International Journal of Human-Computer Studies*, 43(5-6), 907-928.

Gruber, T., 2008. Ontology. *In:* L. Liu and M. Tamer Özsu., eds. *Encyclopedia of Database Systems*, Berlin: Springer-Verlag.

Gruninger, M. and Fox, M., 1995. Methodology for the design and evaluation of ontologies. *In: Workshop on Basic Ontological Issues in Knowledge Sharing, IJCAI'95,* 20-25 August 1995 Montreal, Canada.

Guangzhi, Z., 2009. *A Multidimensional and Visual Exploration Approach to Project Portfolio Management,* Thesis (PhD), Robinson College of Business, Georgia State University.

Guarino, N. and Giaretta, P., 1995. Ontologies and Knowledge Bases: Towards a Terminological Clarification. *In:* N. Mars ed. *Towards Very Large Knowledge Bases: Knowledge Building and Knowledge Sharing.* Amsterdam: IOS Press, 25-32.

Guarino, N., 1997. Understanding, building and using ontologies. *Int. Journal Human-Computer Studies,* 46(2), 293-310.

Guarino, N., 1998. Formal Ontology and Information Systems. *In:* N. Guarino ed. *Formal Ontology in Information Systems*. 6-8 June 1998 Trento, Italy. 3-15.

Guarino, N. and Welty, C., 2004. An Overview of OntoClean. *In*: S. Staab and R. Studer, eds. *The Handbook on Ontologies.* Berlin:Springer-Verlag. 151-172.

Guba, E.G. and Lincoln, Y.S., 1994. Competing Paradigms in Qualitative Research. *In:* N.K. Denzin and Y.S. Lincoln, eds. *Handbook of Qualitative Research*. Thousand Oaks : Sage Publications.

Gudgin, M., Hadley, M., Mendelsohn, N., Moreau, J.J., Nielsen, H.F., Karmarkar, A., and Lafon, Y., 2007. *Soap version 1.2 part 1: Messaging framework 2nd ed.* World Wide Web Consortium, W3C Recommendation 27 April 2007.

Hadley, M., 2009. *Web Application Description Language.* W3C Member Submission 31 August 2009. [online] Available from: http://www.w3.org/Submission/wadl/ [Accessed 4 December 2009].

Hahn, E.D., 2003. Decision Making with Uncertain Judgments: A Stochastic Formulation of the Analytic Hierarchy Process. *Decision Sciences*, 34(3), 443-466.

Hansen, T. and Miller, G., 2002. *A Framework for Prioritizing Use Cases Definition and Verification of Requirements Through Use Case Analysis and Early Prototyping.* [online] Available from: http://www.mindspring.com/~ggmiller/~ggmiller/hansen.pdf [Accessed 24 February 2012].

Hansson, S.O., 2005. *Decision Theory – An Introduction.* vol 19, Stockholm: Royal Institute of Technology.

Hartmann, J. and Sure, Y., 2004. An Infrastructure for Scalable, Reliable Semantic Portals. *IEEE Intelligent Systems,* 19(3), 58–65.

Hartmann, J., Palma, R., Sure, Y., Su´arez-Figueroa, M. C., Haase, P., G´omez- P´erez, A., and Studer, R., 2005a. Ontology metadata vocabulary and applications. *In: International Conference on Ontologies, Databases and Applications of Semantics. In: Workshop on Web Semantics (SWWS), LNCS 3762,* 31 October-4 November 2005 Agia Napa, Cyprus. Springer, 906–915.

Hartmann, J., Spyns, P., Giboin, A., Maynard, D., Cuel, R., Suarez-Figueroa, M.C. and Sure, Y., 2005b. *Methods for ontology evaluation.* KnowledgeWeb Deliverable D1.2.3 v1.3.

Hartmann, J., Palma, R., and Gomez-Perez, A., 2009. Ontology Repositories, *In:* S. Staab and R. Studer eds. *Handbook on ontologies*. Berlin: Springer. 551-571.

Havens, S., 2006. *OpenGIS Transducer Markup Language Implementation Specification OpenGIS Implementation Specification Version: 1.0.0 OGC® 06-010r2.*

Healy, M. and Perry, C., 2000. Comprehensive criteria to judge validity and reliability of qualitative research within the realism paradigm. *Qualitative Market Research*, 3(3), 118-126.

Henson, C.A., Pschorr, J.K., Sheth, A.P., and Thirunarayan, K., 2009. SemSOS: Semantic sensor Observation Service. *In: Proceedings of the 2009 International Symposium on Collaborative Technologies and Systems* (CTS '09), 18-22 May 2009 Washington, DC, USA, 44-53.

Hevner, A., March, S., Park, J., and Ram, S., 2004. Design Science in Information Systems Research. *MIS Quarterly,* 28(1), 75-105.

Hevner, A.R., and Chatterjee, S., 2010. Design Research in Information Systems. *In: Integrated Series in Information Systems,* Volume 22. Springer.

Higashino W, A., Felgar de Toledo M.B., Capretz M. A. M. 2009. REST and Resource-Oriented Architecture. *In: Alt, R., Fähnrich, K.-P., Franczyk, B., eds. Proceedings First International Symposium on Services Science ISSS'09,* Logos, Berlin, 2009.

Hitzler, P., Krötzsch, M., Parsia, B., Patel-Schneider, P.F., and Rudolph, S., eds. 2009. *OWL 2 Web Ontology Language: Primer*. W3C Recommendation 27 October 2009. [online] Available from: http://www.w3.org/TR/2009/REC-owl2-primer-20091027/ [Accessed 2 March 2010].

Hobona, G. and Brackin, R., 2011. *OGC® OWS-8 Cross Community Interoperability (CCI) Semantic Mediation Engineering Report. Public Engineering Report.* OGC 11-063r6.

Hoepfl, M. C., 1997. Choosing qualitative research: A primer for technology education researchers. *Journal of Technology Education* [online], 9(1), 47-63. Available from: http://scholar.lib.vt.edu/ejournals/JTE/v9n1/pdf/hoepfl.pdf [Accessed 3 January 2011].

Horridge, M., Drummond, N., Goodwin, J., Rector, A., and Wang, H.H., 2006. The Manchester OWL Syntax. *In: Proc.of the 2006 OWL Experiences and Directions Workshop (OWLED2006),* 10-11 November *2006* Athens, GA, USA.

Horridge, M., 2010. Ontogenesis. [online] January 22 2010. Available from: http://ontogenesis.knowledgeblog.org/88 [Accessed 22 January 2010].

Horridge, M., 2011. *Protege Owl Tutorial.* [online] The University of Manchester. Available from: http://owl.cs.manchester.ac.uk/tutorials/protegeowltutorial/ [Accessed 25 March 2011].

Horrocks, I., Fensel, D., Broekstra, J., Decker, S., Erdmann, M., Goble, C., Van Harmelen, F., Klein, M., Staab, S., and Studer, R., 2000. *The Ontology Interchange Language OIL.* Technical Report, Free University of Amsterdam. [online] Available from: http://www.ontoknowledge.org/oil/ [Accessed 18 May 2006].

Horrocks, I., 2002. DAML+OIL: a Description Logic for the Semantic Web. *Bull. of the IEEE Computer Society Technical Committee on Data Engineering*, 25(1), 4-9.

Hossein, M., Rajabifard, A., and Williamson, I.P., 2010. Development of an Interoperable Tool to Facilitate Spatial Data Integration in the Context of SDI. *International Journal of Geographical Information Science*, 24(4), 487 – 505.

Hovy, E., 2002. Comparing sets of semantic relations in ontologies. *In:* R. Green, C.A. Bean and S.H. Myaeng, eds*. The Semantics Of Relationships.* Dordrecht: Kluwer, 91-110.

Howe, K.R., 1988. Against the Quantitative-qualitative Incompatibility Thesis or Dogmas Die Hard. *Educational Researcher*, 17, 10-16.

HP, 2009. *Joseki –SPARQL Server For Jena-Home Page.* [online] Available from: http://www.joseki.org/ [Accessed 30 January 2011].

HP, 2012. *Jena-Home Page.* [online] Available from: http://incubator.apache.org/jena/ [Accessed 2 April 2012].

ICSM, 2008. *ICSM Harmonised Data Model Version 1.* [online] Available from: http://www.icsm.gov.au/icsm/harmonised_data_model/model1/hdm_V1.html [Accessed 17 October 2009].

ICSU, 2011. *Interim Report to the ICSU Committee on Scientific Planning and Review (April 2011).* Ad-hoc Strategic Coordinating Committee on Information and Data (SCCID Report).

IEEE, 1998a. *IEEE Standard for a Software Quality Metrics Methodology, revision. STD-1061-1998*. Piscataway, NJ: IEEE Standards Dept.

IEEE, 1998b. *IEEE Recommended Practice for Software Requirements Specifications, STD- 830-1998,* IEEE Computer Society.

IFOMIS, 2011. *Basic Formal Ontology*. [online] Available from: http://www.ifomis.org/bfo/ [Accessed 17 November 2011].

IMOS, 2012a. *Australian Ocean Data Network-Home Page.* [online] Available from: http://imos.org.au/aodn.html [Accessed 1 March 2012].

IMOS, 2012b. *Integrated Marine Observing System-Home Page.* [online] Available from: http://imos.org.au/ [Accessed 12 January 2012].

IMOS, 2012c. *IMOS Metadata Entry and Search Tool - Home Page.* [online] http://imosmest.aodn.org.au/geonetwork/srv/en/main.home [Accessed 12 February 2012].

Iivari, J., 2007. A Paradigmatic Analysis of Information Systems as a Design Science. *Scandinavian Journal of Information Systems,* 19(2), 39-63.

Isaac, A., Phipps, J., and Rubin, D., eds. 2009. *SKOS Use Case and Requirements.* W3C Working Group Note 18 August 2009. [online] Available from: http://www.w3.org/TR/skos-ucr/ [Accessed 4 December 2009].

Isaac, A. and Summers, E., 2009. *SKOS Simple Knowledge Organization System Primer*. W3C Working Group. [online] Available from: http://www.w3.org/TR/skos-primer/ [Accessed 21 July 2010].

ISO, 2002. ISO 19101:2002. *Geographic information – Reference Model.*

ISO, 2003. ISO 19115:2003. *Geographic information - Metadata.*

ISO, 2005a. ISO 19101:2005. *Geographic information. Reference model.*

ISO, 2005b. ISO 19110. *Geographic Information - Methodology for Feature Cataloguing*. Available from: http://www.iso.org/iso/iso_catalogue/catalogue_tc/catalogue_detail.htm?csnumber=39965 [Accessed on 21 April 2007].

ISO, 2005c. ISO 19109. *Geographic Information - Rules for Application Schema*. Available from: http://www.iso.org/iso/iso_catalogue/catalogue_tc/catalogue_detail.htm?csnumber=39891 [Accessed on 21 June 2007].

ISO, 2005d. ISO 11179-6. *Information Technology – Metadata Registries (MDR)--Part 6:Registration*. Available from: http://www.iso.org/iso/iso_catalogue/catalogue_tc/catalogue_detail.htm?csnumber=35348 [Accessed on 21 October 2007].

ISO, 2005e. ISO 19123. *Geographic information - Schema for coverage geometry and functions.* [online] Available from: http://www.iso.org/iso/catalogue_detail.htm?csnumber=40121.[Accessed 17 April 2010].

ISO, 2009a. ISO 19126 DR AS/NZS CP. *Geographic information - Feature concept dictionaries and registers**.***

ISO, 2009b. *Standards guide ISO/TC211 Geographic Information/Geomatics.*

ISO, 2012. ISO TC 211. *Geographic Information/Geomatics.* [online] Available from: http://www.isotc211.org/ [Accessed 2 February 2012].

ISO/IEC, 2004. ISO/IEC TR 9126-4:2004. *Software Engineering – Product Quality – Part 4: Quality in use metrics.* [online] Available from: http://www.iso.org/iso/iso_catalogue/catalogue_tc/catalogue_detail.htm?csnumber=39752 [Accessed 2 July 2009].

ISO/IEC, 2010. ISO/IEC 11179. *Information Technology – Metadata Registries.* [online] Available from: http://metadata-standards.org/11179/ [Accessed 23 September 2011].

Iyengar, S.S. and Lepper, M., 2000. When choice is demotivating: Can one desire too much of a good thing? *Journal of Personality and Social Psychology,* 79, 995-1006.

Jacobson, I., Christerson, M., Jonsson, P., and Oevergaard, G., 1992. *Object-Oriented Software Engineering: A Use Case Driven Approach,* Addison Wesley.

James, K., Batcheller, J.K., Reitsma F., and Gittings B.M., 2007. Implementing ISO-compliant Feature Metadata . *In: GISRUK 2007 Proceedings,* 11th - 13th April 2007 NUI Maynooth, Ireland.

Janowicz, K., Wilkes, M., and Lutz, M., 2008. Similarity-based Information Retrieval and its Role within Spatial Data Infrastructures. *In: 5th International Conference on Geographic Information Science,* 23-26 September 2008 Park City, Utah.

Janowicz, K., Schade, S., Bröring, A., Keßler, C., Stasch, C., Maue, P., and Diekhof, T., 2009. A Transparent Semantic Enabalement Layer for the Geospatial Web. *In: Terra Cognita 2009 Workshop In conjunction with the 8th International Semantic Web Conference (ISWC 2009),* 26 October 2009 Washington.

Janowicz, K., Broring, A., Stasch, C., and Everding, T., 2010. Towards meaningful URIs for linked sensor data*. In: Towards Digital Earth: Search, Discover and Share Geospatial Data, Workshop at Future Internet Symposium,* 20 September 2010, Berlin Germany.

Janowicz, K**.** and Compton, M., 2010. The Stimulus-Sensor-Observation Ontology Design Pattern and its Integration into the Semantic Sensor Network Ontology. *In:* K. Taylor, A. Ayyagari and D. De Roure eds. *The 3rd International workshop on Semantic Sensor Networks 2010 (SSN10) in conjunction with the 9th International Semantic Web Conference (ISWC 2010),* 7-11 November 2010 Shanghai,China.

Janowicz, K., Bröring, A., Stasch, C., Schade, S., Everding, T., and Llaves, A., 2011. A RESTful Proxy and Data Model for Linked Sensor Data. *International Journal of Digital Earth.*

Jansen, H., 2010. The Logic of Qualitative Survey Research and Its Position in the Field of Social Research Method. *Forum: Qualitative Social Research Sozialforschung (FQS),* 11(2).

JCADM, 2007. *JCADM-11-2007 Meeting-Annotated Agenda.* [online] Available from: http://scadm.scar.org/0files/jcadm_meetings/2007_jcadm11_agenda.pdf [Accessed 24 January 2010].

Jean, S., Pierra, G., and Aït Ameur, Y., 2006. Domain Ontologies: A Database-Oriented Analysis. *In: WEBIST (1) 2006,* 10-13 April 2006 Setubal, Portugal. 341-351.

Jirka, S., Bröring, A., and Stasch, C., 2009. Discovery mechanisms for the sensor web. *Sensors*, 9(4), 2661–2681.

Jupp, S., Bechhofer, S., and Stevens, R., 2008. SKOS with OWL: Don't be Full-ish! *In: OWLED 2008* [online], October 26-27 2008 Karlsruhe, Germany. Available from: http://www.webont.org/owled/2008/papers/owled2008eu_submission_22.pdf [Accessed 20 September 2009].

Kalfoglou, Y. and Schorlemmer, M., 2003. Ontology Mapping: the State of the Art. *The Knowledge Engineering Review,* 18(1), 1– 31.

Kalfoglou, Y. and Hu, B., 2006. Issues with Evaluating and Using Publicly Available Ontologies. *In: Proceedings of 4th International EON Workshop, Evaluating Ontologies for the Web,* May 2006 Edinburgh, UK.

Kalyanpur, A., Parsia, B., Sirin, E., and Hendler, J., 2005. Debugging Unsatisfiable Classes in OWL Ontologies. *Journal of Web Semantics*, 3(4), 268-293.

Kaner, C. and Bond, W.P., 2004. Software Engineering Metrics: What Do They Measure and How Do We Know? *In: 10th International Symposium On Software Metrics*, 14-16 September 2004 Chicago, Il.

Kauko, T. J., 2002. *Modelling the locational determinants of house prices: neural network and value tree approaches,* Thesis (PhD), Utrecht University, Utrech.

Kavouras, M., Kokla, M., and Tomai, E., 2005. Comparing categories among geographic ontologies. *Computers and Geosciences,* 31(2), 145-154.

Keefer, D., Kirkwood, C.W., and Corner, J., 2004. Perspective on Decision Analysis Applications. *Decision Analysis,* 1, 4-22.

Keeney, R. L. and Raiffa, H., 1976. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*, New York: John Wiley.

Keeney, R.L., 1982. Decision analysis: An overview. *Operations Research*, 30(5), 803-838.

Kendall, M.G. and Gibbons, J.D., 1990. *Rank correlation methods.* 5th ed. London: Edward Arnold.

Kendall, E., 2009. *OOR for public sector use: Metadata questions,* August 6 2009. [online] Available from: http://ontolog.cim3.net/file/work/OOR-Ontolog-Panel/2009-08-06_Ontology-Repository-Research-Issues/PublicSectorUse--ElisaKendall 20090806.pdf. [Accessed 21 November 2011].

Kennedy, J., Hyam, R., Kukla, R., and Paterson, T., 2006. Standard Data Model Representation for Taxonomic Information. *OMICS – A Journal of Integrative Biology,* 10(2), 220-230.

Kepler, F., Paz-Trillo, C., Riani, J., Ribeiro, M.M., Valdivia-Delgado, K., de Barros, L.N., and Wassermann, R., 2006. Classifying Ontologies. *In: 2nd Workshop on Ontologies and their Applications (WONTO'2006),* 23-27 October 2006 Ribeirão Preto, Brazil.

Kerherve, P., 2006. ISO 191xx series of geographic information standards*. In: Technical Conference on The WMO Information System* [online], 6-8 November 2006 Seoul, Republic of Korea. Available from: http://www.wmo.int/pages/prog/www/TECO-WIS/abstract_1-3-1_WMO-PKerherve_ISO-191xx-Geographic.doc [Accessed on 23 July 2008].

King, N., 2004. Using templates in the thematic analysis of text. In: C. Cassell and G. Symon, eds. *Essential Guide To Qualitative Methods in Organisational Research.* London: Sage.

King, R., 2012. *hResume Microformats Wiki.* [online] Available from: http://microformats.org/wiki/hresume [Accessed 4 January 2012].

Klein, M. and Fensel, D., 2001. Ontology versioning for the Semantic Web. *In: International Semantic Web Working Symposium (SWWS),* 30 July-1 August 2001, Stanford, CA.

Klien, E., Lutz, M., and Kuhn, W., 2006. Ontology-based discovery of geographic information services: An application in disaster management. *Computers, Environment and Urban Systems,* 3, 102–23.

Klein, E., ed. 2007. *D3.1 Ontologies in the SWING Application Requirement Specification Deliverable. SWING Semantic Web-Service Interoperability for Geospatial Decision Making FP6-26514.* [online]. Available from: http://members.sti2.at/~carmenb/swing-deliverables/D3.1.pdf [Accessed 14 September 2009].

Klien, E., Fitzner, D. I. and Maué, P., 2007. Baseline for Registering and Annotating Geodata in a Semantic Web Service Framework. *In: 10th AGILE International Conference on Geographic Information Science,* 8-11 May 2007 Aalborg, Denmark.

Kopecky, J., Vitvar, T., Bournez, C., and Farrell, J., 2007. SAWSDL: Semantic Annotations for WSDL and XML Schema, *IEEE Internet Computing*, 11(6), 60-67.

Kopecky, J., Gomadam, K., and Vitvar, T., 2008a. Hrests: An Html Microformat For Describing Restful Web Services. *In: Proceedings of the 2008 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology*, 9-12 December 2008 Sydney. 619-625.

Kopecky, J., Sheth, A., Vitvar, T., and Gomadam, K., 2008b. *hRests.* [online] Available from: http://knoesis.org/research/srl/projects/hRESTs/ [Accessed 6 January 2012].

Kopecky, J. and Vitvar, T., 2008. *D38v0.1 MicroWSMO.* WSMO Working Draft 19 February 2008. [online] Available from http://www.wsmo.org/TR/d38/v0.1/ [Accessed 15 June 2010].

Kopecky, J., Vitvar, T., Fensel, D., and Gomadam, K., 2009*. D12V0.1HRESTS & MICROWSMO.CMS WG Working Draft – March 10, 2009.* [online] Available from: http://cms-wg.sti2.org/doc/d12v01_20090310.pdf [Accessed 10 November 2010].

Kottmann, C. and Reed, C., 2009. *The OpenGIS Abstract Specification, Topic 5: Features, Version 5. Document 08-126.* [online] Wayland, Massachusetts, USA. Available from: http://portal.opengeospatial.org/files/?artifact_id=29536 [Accessed 17 August 2010].

Koutsomitropoulos, D.A., Domenech, R.B., and Solomou, G.D., 2011. A Structured Semantic Query Interface for Reasoning-Based Search and Retrieval. *In: ESWC 2011*, 29 May – 2 June 2011 Heraklion, Crete, Greece. 17-31.

Kresse, W. and Fadaie, K., 2004. *ISO Standards for Geographic Information*. Berlin: Springer.

KRR, 2012. *Hermit OWL Reasoner-Home Page.* [online] Available from: http://hermit-reasoner.com/ [Accessed 30 February 2012].

Kruskal, J.B. and Wish, M., 1978. *Multidimensional Scaling.* Sage University Paper series on Quantitative Applications in the Social Sciences, 07-011. Newbury Park, CA: Sage Publications.

Kryvobokov, M., 2005. Estimating the weights of location attributes with the Analytic Hierarchy Process in Donetsk, Ukraine. *Nordic Journal of Surveying and Real Estate Research,* 2(2), 5-29.

Kubik, T. and Iwaniak, A., 2010. Building and maintaining metadata repositories with the aid of ontology tools and technologies*. In: GSDI World Conference, Realising Spatially Enabled Societies* [online], 19-22 October 2010 Singapore. Available from: http://www.gsdi.org/gsdiconf/gsdi12/papers/143.pdf [Accessed 23 January 2012].

Kuhn, W., 2003. Semantic Reference Systems*. IJGIS,* 17(5), 405-409.

Kuhn, W., 2005. Geospatial Semantics: Why, of What, and How? *Journal on Data Semantics III*, *Special Issue on Semantic based Geographical Information Systems*, 3534, 1-24.

Kuhn, W., 2009. A Functional Ontology of Observation and Measurement*. In:* K. Janowicz, M. Raubal, and S. Levashkin, eds. *GeoSpatial Semantics · Third International Conference (GeoS 2009),* 3-4 December 2009 Mexico City. Springer-Verlag. 26–43.

Lacasta, J., Nogueras-Iso, J., Béjar, R., Muro-Medrano, P.R., and Zarazaga-Soria, J. F., 2007. A Web Ontology Service to facilitate interoperability within a Spatial Data Infrastructure: Applicability to discovery. *Data Knowl. Eng.,* 63(3), 947-971.

Lacy, L.W., 2005. *Owl: Representing Information Using the Web Ontology Language.* Victoria, BC, Canada: Trafford Publishing.

Lara, R., Roman, D., Polleres, A., and Fensel. D., 2004. A Conceptual Comparison of WSMO and OWL-S*. In: Proceedings of the European Conference on Web Services (ECOWS'04),* 27-30 September 2004 Erfurt, Germany.

Larizgoitia, K. and Toma, I., 2011. *ENVISION Deliverable 5.3 - Integration of the open source catalogue with the semantic discovery engine - Version 1.* ICT for Environmental Services and Climate Change Adaption. [online] Available from: http://www.envision-project.eu/wp-content/.../2011/06/D5.3-2.0_final.pdf [Accessed 2 February 2012].

Lassila, O. and McGuinness, D.L., 2001. *The Role of Frame-Based Representation on the Semantic Web*, Knowledge Systems Laboratory Report KSL-01-02, Stanford University, 2001.

Lassoued, Y., Wright, D.J., Bermudez, L., and Boucelma, O., 2008. Ontology-based Mediation of OGC Catalogue Service for the Web: A Virtual Solution for Integrating Coastal Web Atlases *In: 3rd International Conference on Data and Software Engineering ICSOFT 2008,* 5-8 July 2008 Porto, Portugal.

Lefley, F. and Sarkis, J., 1997. Short-termism and the appraisal of AMT capital projects in the US and UK. *International Journal of Production Research,* 35(2), 341–368.

Lefort, L., 2009., Review of semantic enablement techniques used in geospatial and semantic standards for legacy and opportunistic mashups. *In: Proc. of the 5th Australian Ontology Workshop (AOW 2009) 1 December 2009. Held in Conjunction with the 22nd Australian Joint Conference on Artificial Intelligence (AI'09) University of Melbourne*, 1-4 December 2009 Melbourne, Australia.

Lefort, L., Henson, C., and Taylor, K., 2011. *Semantic Sensor Network XG Final Repo*rt. [online] Available from: http://www.w3.org/2005/Incubator/ssn/XGR-ssn-20110628/ [Accessed 1 January 2012].

Legendre, P., 2010. Coefficient of concordance. *In: Encyclopedia of Research Design*. SAGE Publications.

Lehmann, J. and Hellmann, S., 2009. LinkedGeoData – Adding a Spatial Dimension to the Web of Data. In: 8th International Semantic Web Conference (ISWC 2009), 25 October 2009 Washington DC., USA. Springer. 731–746.

Lehmann, J. and Bühmann, B.L., 2010. ORE - A Tool for Repairing and Enriching Knowledge Bases*. In:* P. Patel-Schneider, Y. Pan, P. Hitzler, P. Mika, L. Zhang, J. Pan, I. Horrocks, et al., eds. *The Semantic Web ISWC 2010 9th International Semantic Web Conference* [online], 7-11 November 2010 Shanghai, China. 177-193. Available from: http://www.springerlink.com/index/L13L913624K60837.pdf [Accessed 3 June 2011].

Lemmens, R., Wytzisk, A., de By, R., Granell, C., Gould, M. and van Oosterom, P., 2006. Integrating Semantic and Syntactic Descriptions to Chain Geographic Services. *IEEE Internet Computing,* 10(5), 42-52.

Lemmens, R., de By, R., Gould, M., Wytzisk, A., Granell, C., and van Oosterom, P., 2007. Enhancing Geo-Service Chaining through Deep Service Descriptions, *Transactions in GIS*, 11(6), 849-871.

Lenat, D.B. and Guha, R.V., 1990. *Building Large Knowledge-Based Systems: Representation and Inference in the Cyc Project.* Addison-Wesley. Reading, Massachusetts.

Lesage, N., 2007. *OGC Cataloguing of ISO Metadata (CIM) Using the ebRIM profile of CSW.* OGC Discussion Paper. Version 0.1.7. OGC 07-038.

Li, H.L. and Ma, L.C., 2006. Adjusting Ordinal and Cardinal Inconsistencies in Decision Preferences Based on Gower Plots. *APJOR*, 23(3), 329-346.

Liao, H. and Hong, J., 2005. Analysis and Application of the Hierarchical Metadata Framework of the ISO 19115 Standard. *In: Proceedings of the Asian Association On Remote Sensing* [online], 10 November 2005 Hanoi, Vietnam. Available from: http://www.a-a-r-s.org/acrs/proceeding/ACRS2005/Papers/GDS1-4.pdf [Accessed 15 March 2008].

Lieberman, J., 2006. *Geospatial Semantic Web Interoperability Experiment Report*. OGC Discussion Paper. OGC 06-002r1.

Lin, C.C., Wang, W.C., and Yu, W.D., 2008. Improving AHP for construction with an adaptive AHP approach (A3), *Automation in construction*, 17(2), 180-187.

Lincoln, Y.S. and Guba, E.G., 1985. *Naturalistic inquiry.* Beverly Hills, CA: Sage.

Lincoln, Y.S. and Guba, E.G., 2003. Paradigmatic Controversies, Contradictions, and Emerging Confluences. *In:* N.K. Denzin and Y.S. Lincoln, eds. *The Landscape of Qualitative Design: Theories and Issues*. 2nd ed. Thousand Oaks : Sage Publications.

LOA, 2001. *DOLCE Lite Schema.* [online] Available from: http://www.loa.istc.cnr.it/ontologies/DOLCE-Lite.owl [Accessed 10 January 2012].

LOA, 2012. *Dolce: A Descriptive Ontology for Linguistic and Cognitive Engineering*. [online] Available from: http://www.loa-cnr.it/DOLCE.html [Accessed 17 January 2012].

Lopez, M. F., Gomez-Perez, A., Sierra, A. P., and Sierra, J. P., 1999. Building a chemical ontology using methontology and the ontology design environment. *IEEE Intelligent Systems*, 14 (5), 37-45.

López-Pellicer, F.J., Florczyk, A.J., Renteria-Agualimpia, W., Nogueras-Iso, J., Valiño, J., and Muro-Medrano, P.R., 2011. CSW2LD: a Linked Data frontend for CSW. *In: II Iberian Conference on Spatial Data Infrastructures (JIIDE'2011),* 9-11 November 2011 Barcelona.

Lozano-Tello, A., 2002. *Metrica De Idoneidad De Ontologias,* Thesis (PhD). University Extremadura, Spain.

Lozano-Tello, A. and Gomez-Perez, A., 2004. ONTOMETRIC: A method to choose the Appropriate Ontology. *Journal of Database Management*, 15(2), 1-18.

Lutz, M. and Klien, E., 2006. Ontology-based Retrieval of Geographic Information. *International Journal of Geographical Information Science*, 20(3), 233–260.

Lutz, M., 2007. Ontology-Based Descriptions for Semantic Discovery and Composition of Geoprocessing Services. *GeoInformatica*, 11, 1–36.

Madin, J., Bowers, S., Schildhauer, M., Krivov, S., Pennington, D., and Villa, F., 2007. An ontology for describing and synthesizing ecological observation data. *Ecological Informatics*, 2, 279-296.

Maedche, A. and Staab, S., 2002. Measuring similarity between ontologies. *In: Proceedings of the European Conference on Knowledge Acquisition and Management (EKAW-2002)*, 1-4 October 2004 Siguenza, Spain.

Malone, J. and Parkinson, H., 2010. Reference and Application Ontologies. Ontogenesis [online]. January 22, 2010. Available from: http://ontogenesis.knowledgeblog.org/295 [Accessed 23 January 2010]

Manning., M., Aggarwal., A. Gao, K.,Tucker-Kellog, G,. 2009. Scaling the walls of discovery: using semantic metadata for integrative problem solving. *Briefings in Bioinformatics*, 10(2), 164-176.

Manola, F. and Miller, E., 2004. *RDF primer.* W3C Recommendation. [online] Available from: http://www.w3.org/TR/rdf-primer/ [Accessed 2 June 2009].

March, S.T. and Smith, G.F., 1995. Design and natural science research on information technology. *Decision Support Systems,* 15, 251-266.

March, S. T. and Storey, V. C., 2008. Design science in the information systems discipline: An introduction to the special issue on design science research. *MIS Quarterly*, 32(4), 725–730.

Marks, E., and Bell, M., 2006. *Service Oriented Architecture – A Planning and Implementation Guide For Business and Technology.* Hoboken, New Jersey: John Wiley and Sons, Inc.

Markus, M. L., Majchrzak, A., and Gasser, L., 2002. A Design Theory for Systems that Support Emergent Knowledge Processes, *MIS Quarterly* 26(3), 179-212.

Mason, M., 2010. Sample Size and Saturation in PhD Studies Using Qualitative Interviews. *Forum: Qualitative Social Research Sozialforschung (FQS),* 11(3).

Martell R., ed. 2007. *CSW-ebRIM Registry Service - Part 1: ebRIM profile of CSW (Version: 1.0.0)*. OGC 07-110r2, Open Geospatial Consortium Inc. [online] Available from: http://portal.opengeospatial.org/files/?artifact_id=27092 [Accessed 3 July 2009].

Martin, D., Burstein, M., Hobbs, J., Lassila, O., McDermott, D., McIlraith, S., Narayanan, S., Paolucci, M., Parsia, B., Payne, T. R., Sirin, E., Srinivasan, N., and Sycara, K., 2004. *OWL-S: Semantic markup for web services,* Technical Report UNSPECIFIED, Member Submission, W3C. [online] Available from: http://www.w3.org/Submission/ OWL-S/ [Accessed 23 August 2009].

Mascardi, V., Cordì, V., and Rosso, P., 2007. *A Comparison of Upper Ontologies*, Technical Report DISI-TR-06-2, Genova, Italy.

Masolo, C., Borgo, S., Gangemi, A., Guarino, N., Oltramari, A., and Schneider, L., 2003. *WonderWeb Deliverable D17: ISTC-CNR.*

Maue, P., 2009. *Semantic annotations in OGC standards*. OGC Discussion Paper.

Maué, P. and Schade, S., 2009. Data Integration in the Geospatial Semantic Web*. Journal of Cases on Information Technology*, 11, 100–122.

Maué, P., Michels, H., and Roth, M., 2010. Injecting semantic annotations into (geospatial) Web service descriptions. Submitted to *Semantic Web Journal SWJ*, 0(0), 1-10 (in review).

Maxwell, J. A., 1992. Understanding and validity in qualitative research. *Harvard Educational Review, 62,* 279–300.

Maxwell, l. J., 2008. Designing a qualitative study. *In:* L. Bickman and D.J. Rog eds. *The handbook of applied social research methods.* 2nd ed. Thousand Oaks, CA: Sage Publications.

McDermott, D., Burstein, M., and Smith, D., 2001. Overcoming ontology mismatches in transactions with self-describing service agents. *In: Proc. Int'l Semantic Web Workshop*, 30 July-1 August 2001 Stanford University, CA. 285-302

McGuiness, D.L., Fikes, R., Rice, J., and Wilder, S., 2000. An environment for merging and testing large ontologies. *In: Proceedings of the Seventh International Conference on Principles Of Knowledge Representation and Reasoning (KR2000),* 12-15 April 2000 Breckenridge, Colorado.

McGuinness, D. and van Harmelen, F., 2004. *OWL web ontology language overview*.

McKay, J. and Marshall, P., 2005. *A Review of Design Science in Information Systems. In: Proceedings of the Australasian Conference on Information Systems,* 30 November-2 December 2005 Sydney.

McKay, J. and Marshall, P., 2008. *Foundations of Design Science In Information Systems. Research Into Information Systems In Organisations (RISO).* Faculty of Information and Communication Technologies, Swinburne University of Technology, Working Paper. Issue No 1/08.

Meixner, O., 2009. Fuzzy AHP Group Decision Analysis and its Application for the Evaluation of Energy Sources. *In: Proceedings of the 10th International Symposium on the Analytic Hierarchy/Network Process,* 29 July-1 August 2009 Pittsburtgh, PA, USA.

Merkin, B. G., 1979. *Group Choice*, NY: John Wiley & Sons.

Merriam, S. B., 1998. *Qualitative research and case study applications in education.* San Francisco: Jossey-Bass Publishers.

Microsoft, 2012. *Skype-Home Page.* [online] Available from: http://www.skype.com/intl/en/home [Accessed 20 February 2012].

Mikhailov, L. and Tsvetinov, P., 2004. Evaluation of Services using a Fuzzy Analytic Hierarchy Process. *Applied Soft Computing*, 5(1), 23–33.

Miles, M.B. and Huberman, A.M., 1994. *Qualitative data analysis: An expanded sourcebook*. Thousand Oaks, CA: Sage.

Miles, A. and Brickley, D., eds. 2005. *SKOS Core Guide.* W3C Public Working Draft, World Wide Web Consortium, November 2005. [online] Available from: http://www.w3.org/TR/swbp-skos-core-guide [Accessed 13 February 2009].

Milton, S.K. 2004. Top-Level Ontology: The Problem with Naturalism. *Formal Ontology in Information Systems*, IOS Press.

MINDLab, 2012. *SWOOP Semantic Web Ontology Editor-Home Page.* [online] Available from: http://code.google.com/p/swoop/ [Accessed 2 April 2012].

Minola, F. and Miller, E., 2004. *RDF Primer W3C Recommendation.* [online] Available from: http://www.w3.org/TR/2004/REC-rdf-primer-20040210/ [Accessed 23 September 2008].

MMI, 2011. *Marine Metadata Interoperability Project-Home Page.* [online] Available from: http://marinemetadata.org/ [Accessed 17 January 2012].

Moeller, J., 2010. Spatial Data Infrastructures (SDI)-Challenges and Opportunities. *Directions Magazine.* [online] Available from: http://www.directionsmag.com/articles/spatial-data-infrastructures-sdi-challenges-and-opportunities/122344 [Accessed 3 March 2011].

Moreno-Jiménez, J.M., Aguarón, J., and Escobar, M.T. 2002, Decisional Tools for Consensus Building in AHP-Group Decision Making. *In: 12th Mini Euro Conference*, 2-5 April 2002 Brussels, Belgium.

Morgenthal, J.P. and Evdemon, J., 2000. Eliminating Redundancy in XML Using ID/IDREF. *XML Journal* [online]. Available from: http://xml.sys-con.com/node/40075 [Accessed on August 18 2000].

Motik, B., Horrocks, I., Rosati, R., and Sattler, U., 2006. Can OWL and logic programming live together happily ever after. *In: Proceedings ISWC-2006,* 5-9 November 2006 Athens, Georgia. Springer, 501-514.

Motik, B., Patel-Schneider, P.F., and Grau, B.C., 2009. *OWL 2 Web Ontology Language: Direct semantics*. W3C Recommendation. [online] Available from: http://www.w3.org/TR/owl2-direct-semantics [Accessed 27 October 2010].

Na, A., Priest, M., Niedzwiadek, H., and Davidson, J., 2007. *OGC Implementation Specification 06-009r6: Sensor Observation Service,* Open Geospatial Consortium, Tech. Rep. [online] Available from: http://portal.opengeospatial.org/files/?artifact_id=26667 [Accessed 24 August 2009].

NASA, 2012a. *Global Change Master Directory*. [online] Available from: http://gcmd.nasa.gov/ [Accessed 16 February 2012].

NASA, 2012b. *Service Entry Resource Format (SERF) Writers Guide.* [online] Available from: http://gcmd.nasa.gov/User/serfguide/ [Accessed 19 February 2012].

NASA, 2012c. *Directory Interchange Format (DIF) Writers Guide.* [online] Available from: http://gcmd.nasa.gov/User/difguide/difman.html [Accessed 19 February 2012].

NCBI, 2012. *Welcome To NCBI.* [online] Available from: http://www.ncbi.nlm.nih.gov/ [Accessed 1 February 2012].

NCBO, 2012. *BioPortal.* [online] Available from: http://bioportal.bioontology.org/ [Accessed 16 January 2012].

Nebert, D. and Whiteside, A., eds., 2005. *OpenGIS Catalogue Services Specification, Version 2.0.0 with Corregendum OpenGIS® Implementation Specification OGC 04-021r3*.

Nebert, D., Whiteside, A., and Vretanos, P., eds. 2007. *OpenGIS - Catalogue Services Specification (version: 2.0.2)*. OGC 07-006r1, OGC. [online] Available from: http://portal.opengeospatial.org/files/?artifact_id=20555 [Accessed 27 October 2010].

NeOn, 2011. *Ontology Design Patterns-Home Page.* [online] Available from: http://www.ontologydesignpatterns.org [Accessed 10 December 2011].

Neuhaus, H. and Compton, M., 2009. The Semantic Sensor Network Ontology: A Generic Language to Describe Sensor Assets. *In: AGILE 2009 Pre-Conference Workshop Challenges in Geospatial Data Harmonisation,* 02 June 2009 Hannover, Germany.

Neuman, W.L., 1999. *Social Research Methods – Qualitative and Quantitative Approaches*. 4th ed. Boston: Allyn and Bacon.

Niles, I. and Pease, A., 2001. Toward a Standard Upper Ontology, *In:* Chris Welty and Barry Smith, eds. *Proceedings of the 2nd International Conference on Formal Ontology in Information Systems (FOIS-2001),* 17-19 October 2001 Ogunquit, Maine.

Noy, N. and Rector, A., 2006. *Defining N-ary Relations on the Semantic Web*. W3C Working Group Note 12 April 2006.

Noy, N.F., Shah, N.H., Whetzel, P.L., Dai, B., Dorf, M., Griffith, N., Jonquet, C., Rubin, D.L., Storey, M.A., Chute, C.G., and Musen, M.A., 2009. BioPortal: ontologies and integrated data resources at the click of a mouse. *Nucleic Acids Res 2009*, 37(Web Server issue), 170-173.

Nyulas, C., Noy, N.F., Dorf, M., Griffith, N., and Musen, M.A., 2009. Ontology-Driven Software: What We Learned From Using Ontologies As Infrastructure For Software Or How Does It Taste to Eat Our Own Dogfood*. In:* E.F. Kendall ed. *5th International Workshop on Semantic Web Enabled Software Engineering* **-** *SWESE 2009***,** 26 October 2009 Virginia, USA.

OASIS, 2002. *EbRIM: ebXML Registry Information Model v2.1*. [online] Available from: http://www.ebxml.org/specs/ebRIM.pdf. [Accessed 28 February 2010].

OBO Foundry, 2012a, *The Open Biological and Biomedical Ontologies-Home Page.* [online] Available from: http://www.obofoundry.org/ [Accessed 7 March 2012].

OBO Foundry, 2012b. *Category:Principles.* [online] Available from: http://www.obofoundry.org/wiki/index.php/Category:Principles [Accessed 1 February 2012].

O'Brien, J. and Marakas, G., 2007. *Introduction to Information Systems*. 13th ed. Boston: McGraw-Hill/Irwin.

Obrst , L., Ashpole , B., Ceusters, W., Mani, I., and Smith, B., 2007. The Evaluation of Ontologies. Toward Improved Semantic Interoperability. *In:* C. Baker and K. Cheung, eds. *Semantic Web: Revolutionizing Knowledge Discovery in the Life Sciences*. Secaucus, NJ, USA: Springer-Verlag New York, Inc., 139-158.

OGC, 2011. *OGC Reference Model, 2011-12-19. OGC 08-062r7, Version: 2.1.*

OGC, 2012a. *Open Geospatial Consortium.* Available from: http://www.opengeospatial.org/ [Accessed 25 January 2012].

OGC, 2012b. *OGC Standards and Supporting Documents.* [online] Available from: http://www.opengeospatial.org/standards [Accessed 10 April 2012].

OGC, 2012c. *MetOcean Domain Working Group-Home Page*. [online] Available from: http://external.opengis.org/twiki_public/MetOceanDWG [Accessed 2 February 2012].

Olsen, L.M., Major, G., Shein, K., Scialdone, Vogel,J.R., Leicester,S., Weir,H., Ritz, S., Stevens, T., Meaux, M., Solomon, C., Bilodeau, R., Holland, M., Northcutt, T., and Restrepo, R.A., 2007. *NASA/Global Change Master Directory (GCMD) Earth Science Keywords. Version 6.0.0.0.0* [online] Available from: http://gcmd.nasa.gov/Resources/valids/archives/keyword_list.html [Accessed 1 May 2009].

OMG, 2000. *OMG Unified Modelling Language Specification, Version 1.4,* Object Management Group, November 2000.

Ontology Online, 2012. *jOWL Home-Page.* [online] Available from: http://jowl.ontologyonline.org/ [Accessed 7 February 2012].

Ontotext, 2011. *OWLIM-Home Page.* [online] Available from: http://www.ontotext.com/owlim [Accessed 17 March 2012].

Onwuegbuzie, A. J. and Collins, K. M. T., 2007. A Typology of Mixed Methods Sampling Designs in Social Science Research*. The Qualitative Report* [online], 12(2), 281-316. Available from: http://www.nova.edu/ssss/QR/QR12-2/onwuegbuzie2.pdf [Accessed 23 March 2009].

OOR, 2012. *Open Ontology Repository.* [online] Available from: http://ontolog.cim3.net/cgi-bin/wiki.pl?OpenOntologyRepository [Accessed 11 March 2012].

OPeNDAP, 2012. *OPeNDAP Home Page.* [online] Available from: http://opendap.org/ [Accessed 19 February 2012].

OpenGIS, 2012. *OGC Coverage Type Schema.* [online] Available from: http://schemas.opengis.net/gmlcov/1.0/coverage.xsd [Accessed 20 March 2012].

Oracle™. 2007. *Semantic Data Integration for the Enterprise. An Oracle White Paper.* [online] Available from: http://www.lawebinteligente.net/c/document_library/get_file?p_l_id=10746andfolderId=10906&name=DLFE-4.pdf [Accessed 23 February 2008].

Oracle™, 2010a. *Jena Adapter For Oracle Database.*[online] Available from: http://download.oracle.com/docs/cd/E18283_01/appdev.112/e11828/sem_jena.htm#CBBBGFHI [Accessed 4 September 2011].

Oracle™, 2010b. *OTN Discussion Forum-Problems Setting Up Semantic Tech in Oracle 11G Running Linux.* [online] Available from: http://forums.oracle.com/forums/thread.jspa?messageID=4360682 [Accessed 4 September 2011].

Orme, A., Yao, H., and Etzkorn, L., 2006. Coupling metrics for ontology-based systems. *IEEE Software*, 23(2), 102–108.

Orme, A., Yao, H., and Etzkorn, L., 2007. Indicating ontology data quality, stability, and completeness throughout ontology evolution. *Software Maintenance and Evolution: Research and Practice*, 19(1), 49–75.

Osterle, H., Becker, J., Frank, U., Hess, T., Karagiannis, D., Krcmar, H., Loos, P., Mertens, P., Oberweis, A., and Sinz, E.J., 2011. Memorandum on design-oriented information systems research. *European Journal of Information Systems*, 20(1), 1–6.

Palma, R. and Haase, P., 2005. Oyster-sharing and re-using ontologies in a peer-to-peer community. *In: International Semantic Web Conference, volume 3729 of LNC,* 6-10 November 2005 Galway, Ireland. Springer, 1059-1062.

Parent, C., Spaccapietra, S., and Zimányi, E., 2006. *Conceptual Modeling for Traditional and Spatio-Temporal Applications, The MADS Approach*. Berlin: Springer-Verlag.

Parsons, M.A., de Bruin, T., Tomlinson, S., Campbell, H., Godoy, O., and LeClert, J., 2011. The State of Polar Data – The IPY Experience. *In:* I. Krupnik, et al., eds. *Understanding Earth's Polar Challenges: International Polar Year 2007–2008.* University of the Arctic, Rovaniemi, Finland /CCI Press (Printed Version), Edmonton, Alberta, Canada and ICSU/WMO Joint Committee for International Polar Year 2007–2008.

Paslaru Bontas Simperl, E., and Tempich, C., 2006. Ontology Engineering: A Reality Check. *In: 5th International Conference on Ontologies, Databases and Applications of Semantics ODBASE2006,* 29 October - 3 November 2006 Montpellier, France.

Passant, A., Kopecký, J., Corlosquet, S., Berrueta, D., Palmisano, D., and Polleres, A., 2009. *XSPARQL: Use cases*. W3C Member Submission 20 January 2009. [online] Available from: http://xsparql.deri.org/spec/usecases#sec:sawsdl_lowering [Accessed 12 September 2010].

Patel, C., Kaustubh, S., Yugyung, L., and Park, E.K., 2003. OntoKhoj: a semantic web portal for ontology searching, ranking and classification. *In: WIDM '03: Proceedings of the 5th ACM international workshop on Web information and data management,* 7-8 November 2003 New Orleans, Lousiana. 58-61.

Patel-Schneider, P.F., Hayes, P., and Horrocks, I., 2004. *OWL web ontology language semantics and abstract syntax.* W3C Recommendation. [online] Available from: http://www.w3.org/TR/owlsemantics/ [Accessed 16 September 2009].

Patni, H., Henson, C., and Sheth, A., 2010. Linked sensor data. *In: 2010 International Symposium on Collaborative Technologies and Systems,* 17-21 May 2010, Chicago.

Patton, M.Q., 1990. *Qualitative evaluation and research methods.* 2nd ed. Newbury Park, CA: Sage.

Paulheim, H. and Probst, F., 2010. Ontology-Enhanced User Interfaces: A Survey. *Int. J. Semantic Web Inf. Syst,* 6(2), 36-59.

Percivall, G., 2002. *The OpenGIS Abstract Specification, Topic 12: OpenGIS Service Architecture, Version 4.3.* OpenGIS Consortium Technical Report. [online] Available from: http://portal.opengeospatial.org/files/?artifact_id=1221 [Accessed 23 October 2009].

Percivall, G., 2003. *OGC Reference Model OGC 03-040.* OGC. [online] Available from: http://portal.opengeospatial.org/files/?artifact_id=3836 [Accessed 23 October 2009].

Peroni, S., Motta, E., and d'Aquin, M., 2008. Identifying Key Concepts in an Ontology Through the Integration of Cognitive Principles with Statistical and Topological Measures. *In: 3rd Asian Semantic Web Conference,* 8-11 December Bangkok, Thailand.

Peterseil, J., Blankman, D., Gaigalas, G., Magagna, B., Schentz, H., Vanacker, S., Vanhercke, L., and van der Werf, B., 2010. *European Biodiversity Observation Network: Design of a plan for an integrated biodiversity observing system in space and time.* EBONE-D7.1_1.0_TechnicalSpecification. [online] Available from: http://www.ebone.wur.nl/NR/rdonlyres/298E8A25-FEB4-4E0C-B7E3-1C725D3F7006/107708/EBONE_D71_10_TechnicalSpecification_20100331.pdf [Accessed 12 January 2011].

Phillips, F., 2002. The distortion of criteria after decision-making. *Organizational Behavior and Human Decision Processes,* 88(2), 769-784.

Phuoc, D.L. and Hauswirth, M., 2009. Linked open data in sensor data mashups. *In:* K. Taylor, A. Ayyagari and D. De Roure, eds., *Proceedings of the 2nd International Workshop on Semantic Sensor Networks (SSN09), CEUR, vol.522,* 25-29 October 2009 Washington, 1-16.

Pierce, C. S., 1931. *In:* C. Hartshorne and P. Weiss, eds. *Collected Papers of Charles Sanders Peirce.* Cambridge, MA: Harvard University Press.

Piirainen, K., Gonzalez, R.A., and Kolfschoten, G., 2010. Quo Vadis, Design Science ? – A Survey of Literature. *In:* R. Winter, J.L. Zhao and S. Aier, eds. *Global Perspectives on Design Science Research - 5th International Conference, DERIST 2010,* 4-5 June 2010 St Gallen, Switzerland.

Pisanelli, D., Gangemi, A., and Steve, G., 2003. An Ontology of Descriptions and Situations for Lyee's hypothetical World. *In: Proceedings of Lyee W03-The Second International Workshop on Lyee Methodology,* 24-26 September 2003 Stockholm, Sweden. 53-59.

Portele C., 2007. *Geography Markup Language (GML) Encoding Standard* version 3.2.1 OGC 07-036 Open Geospatial Consortium 2007-08-27. [online] Available from: http://portal.opengeospatial.org/files/?artifact_id=20509 [Accessed 6 February 2009].

Powell, J., 2011. The New Electronic Chart Product Specification S-101: An Overview. *International Journal on Marine Navigation and Safety of Sea Transportation (TransNav),* 5(2), 167-171.

Presutti, V. and Gangemi, A., 2008. Identity of resources and entities on the web. *International Journal of Semantic Web and Information Systems,* 4(2), 49-72.

Pries-Heje, J. and Baskerville, R., 2008. The Design Theory Nexus. *MIS Quarterly*, 32(4), 731-755.

Princeton University, 2012. *WordNet.* [online] Available from: http://wordnet.princeton.edu/ [Accessed 27 February 2012].

Probst, F., 2006. An Ontological Analysis of Observations and Measurements. *In: Proc. of the 4th International Conference on Geographic Information Science GIScience2006,* 20-23 September 2006 Munster, Germany. 304-320.

Probst, F., 2007. *Semantic Reference Systems for Observations and Measurements.* [online] Thesis (PhD), Institute for Geoinformatics, University of Munster, Germany. Available from: http://ifgi.uni-muenster.de//~probsfl/publications/PROBST-Thesis-SemanticReferenceSystemsForObservationsAndMeasurements.pdf [Accessed 18 October 2009].

Prud'hommeaux, E. and Seaborne, A., 2008. *SPARQL Query Language for RDF.* W3C Recommendation. [online] Available from: http://www.w3.org/TR/rdf-sparql-query/ [Accessed 12 October 2010].

Pschorr, J., Henson, C., Patni, H., and Sheth, A.P., 2010. Sensor Discovery on Linked Data. *In: Proceedings of the 7th Extended Semantic Web Conference,* 30 May – 3 June 2010 Heraklion, Crete, Greece.

Purchase, H. C., Andrienko, N., Jankun-Kelly, T. J., and Ward, M., 2008. Theoretical foundations of information visualization. *In:* A. Kerren, J.Stasko, J Fekete and D. Smith eds. *Information Visualization: Human-Centered Issues and Perspectives*, *Notes In Computer Science*, 4950. Springer, 46-64.

QuestionPro, 2012. *QuestionPro Online Survey Software-Home Page*. [online] Available from: http://www.questionpro.com/home [Accessed 12 February 2012].

Quine, W.V.O., 1995. *From Stimulus to Science.* Cambridge, MA: Harvard University Press.

Ramanathan, R. and Ganesh, L.S., 1994, Group preference aggregation methods employed in AHP: an evaluation and an intrinsic process for deriving members' weightages. *European Journal of Operational Research,* 79(20), 249-265.

Raskin, R., 2011. *Semantic Web for Earth and Environmental Terminology (SWEET).* [online] Available from: http://sweet.jpl.nasa.gov/ [Accessed 10 Feb 2012].

Raymond, E.S., 1998. *The Cathedral and the Bazaar.* [online] Available from: http://www.jus.uio.no/sisu/the_cathedral_and_the_bazaar.eric_s_raymond/_4.html [Accessed 2 January 2010].

Read, D., 2004. Utility theory from Jeremy Bentham to Daniel Kahneman. Working Paper – LSEOR 04-64. *Thinking Reasoning (16/05/2005).* Available from: http://www.lse.ac.uk/collections/operationalResearch/Default.htm [Accessed 2 June 2009].

Reed**,** C.**,** 2010. The Open Geospatial Consortium and Web Services Standards. *In:* P Zhao and L Di, eds. *Geospatial Web Services: Advances in Information Interoperability*. Hershey, Pa: IGI Global.

Reitsma, F. and Albrecht, J., 2005. Modeling with the Semantic Web in the Geosciences. *IEEE Intelligent Systems,* 20 (2), 86–88.

Renear, A.H., Sacchi, S., and Wickett, K.M., 2010. Definitions of dataset in the scientific and technical literature. *In: Proceedings of the 73rd ASIS\&T Annual Meeting on Navigating Streams in an Information Ecosystem - Volume 47,* 22-27 October 2010 Pittsburg, PA.

Richards, L., 2005. *Handling qualitative data.* London: Sage.

Rodriguez, A., 2008. *RESTful Web Services: The basics.* IBM Technical Paper. [online] Available from: http://www.ibm.com/developerworks/webservices/library/ws-restful/index.html [Accessed on 11 December 2009].

Roman, D., Keller, U., Lausen, H., Oren, E., Bussler, C., Kifer, M., and Fensel, D., 2004. *Web Service Modeling Ontology (WSMO)*, version 1.0. [online] Available from: http://www.wsmo.org/2004/d2/v1.0/ [Accessed 23 August 2009].

Roussey, K., 2005. *Guidelines to build ontologies: A Bibliographic study.* [online] Available from: http://www.towntology.net/Documents/guidelines.pdf [Accessed 9 November 2010].

Saaty, T., 1980, *The Analytic Hierarchy Process: Planning, Priority Setting, Resource Allocation*, New York: McGraw-Hill.

Saaty, T. L., 1987. Rank generation, preservation, and reversal in the analytic hierarchy decision process. *Decision Sci.,* 18(2), 157–177.

Saaty, T. L., 1990. An exposition of the AHP in reply to the paper- Remarks on the analytic hierarchy process, *Management Science.* 36(3) 259–268.

Saaty, T.L. and Vargas, L.G., 1991. *Prediction, projection and forecasting*. Boston: Kluwer Academic Publishers.

Saaty, T.L., 1994. *Fundamentals of Decision Making*, Pittsburgh, PA: RWS Publications.

Saaty, T.L. 1999*. Decision Making for Leaders: The Analytic Hierarchy Process for Decisions in a Complex World*. Pittsburgh, Pennsylvania: RWS Publications.

Sabou, M., Lopez, V., Motta, E., and Uren, V., 2006. Ontology Selection: Ontology Evaluation on the Real Semantic Web. *In: WWW2006,* 22–26 May 2006 Edinburgh, UK.

Sansone, S., 2007. *Ontology for Biomedical Investigation (OBI), the OBO Foundry and related synergistic standards initiatives.* [online] Available from: http://www.casimir.org.uk/storyfiles/42.0.wp4_sansone.pdf [Accessed 1 Feb 2012].

Sauermann, L. and Cyganiak, R., 2008. *Cool URIs for the Semantic Web,* W3C Interest Group Note 3 December 2008. [online] Available from: http://www.w3.org/TR/2008/NOTE-cooluris-20081203 [Accessed 17 November 2009].

SCADM, 2009. *SCADM Amsterdam 2009 Meeting-Home Page.* [online] Available from: http://scadm.scar.org/scadm/mr_ams2009.html [Accessed 24 March 2010].

SCADM, 2010. *SCADM Buenos Aires 2010 Meeting-Home Page.* [online] Available from: http://scadm.scar.org/scadm/mr_ba2010.html [Accessed 1 November 2011].

SCADM, 2011. *Status of the Antarctic Master Directory.* [online] Available from: http://scadm.scar.org/0files/scadm_sep2011/AMDPresentation2011.pdf [Accessed 8 November 2011].

SCADM, 2012. *SCADM-Home Page.* [online] Available from: http://scadm.scar.org/index.php [Accessed 1 February 2012].

SCAR, 2012. *Scientific Committee on Antarctic Research-Home Page.* [online] Available from: http://www.scar.org/ [Accessed 16 March 2012].

Schade, S., Sahlmann, A., Lutz, M., Probst, F., and Kuhn, W., 2004. Comparing approaches for semantic service description and matchmaking*. In:* R. Meersman and Z. Tari eds. *On the Move to Meaningful Internet Systems 2004, CoopIS, DOA, and ODBASE,* 25-29 October 2004 Agia Napa, Cyprus. Berlin: Springer, 1062–1079.

Schade, S., Granell, C., and Díaz, L., 2010. Augmenting SDI with Linked Data*. In: Workshop On Linked Spatiotemporal Data 2010 (LSTD 2010),* 14-17 September 2010 Zurich, Switzerland.

Schensul, S. L., Schensul, J. J. and LeCompte, M.D., 1999. *Essential ethnographic methods: observations, interviews, and questionnaires (Book 2 in Ethnographer's Toolkit).* Walnut Creek, CA: AltaMira Press.

Schmidt, R.C., 1997. Managing Delphi Surveys Using Nonparametric Statistical Techniques. *Decision Sciences, 28*(3), 763-774.

Schneider, D., 2009. *Quantitative Ecology: Spatial and Temporal Scaling*. 2nd ed. Academic Press.

Schneider, M., 2010. SPARQLAS Implementing SPARQL Queries With OWL Syntax*. In: TWOMDE 2010* [online], 28 June-2 July 2010 Malaga, Spain. Available from: http://ceur-ws.org/Vol-604/paper10.pdf [Accessed 13 January 2011].

Schut, P. and Whiteside, A., ed., 2005. *OpenGIS® Web Processing Service.* OGC 05-007r4. [online] Available from: http://portal.opengeospatial.org/files/?artifact_id=12184 [Accessed 20 March 2010].

Schut, P., 2007. *OGC Implementation Specification 05-007r7: OpenGIS Web Processing Service.* [online] Available from: http://portal.opengeospatial.org/files/?artifact_id=40310 [Accessed 23 March 2010].

Schuurman N. and Leszczynski, A., 2006. Ontology-based Metadata. *Transactions in GIS*, 10(5), 709-726.

Schwering, A., 2008. Approaches to Semantic Similarity Measurement for Geo-Spatial Data: A Survey. *Transactions in GIS,* 12(1), 5-29.

Scott, J. and Marshall, G., 2005. *A Dictionary of Sociology.* Oxford: Oxford University Press.

Scriven, M., 1991. *Evaluation Thesaurus,* 4th ed. Newbury Park, CA: Sage Publications.

SeaDataNet, 2011. *SeaDataNet2 Context and Objectives.* [online] Available from: http://www.seadatanet.org/Overview/Context-and-Objectives [Accessed 12 September 2011].

SEEK, 2008. *Science Environment For Ecological Knowledge (SEEK)-Home Page.* [online] Available from: http://seek.ecoinformatics.org./ [Accessed 2 February 2010].

Senkler, K., Voges, U., and Remke, A., 2004. An ISO 19115/19119 Profile For OGC Catalogue Services CSW 2.0. *In: 10th EC GI & GIS Workshop, ESDI State of the Art*, 23-25 June 2004 Warsaw, Poland.

Sequeda, J. and Corcho, O., 2009. Linked stream data: A position paper . *In:* K. Taylor, A. Ayyagari and D. De Roure, eds., *Proceedings of the 2nd International Workshop on Semantic Sensor Networks (SSN09), CEUR, vol.522,* 25-29 October 2009 Washington, 148-157.

Shah, A.K. and Oppenheimer, D.M., 2008. Heuristics made easy: An effort-reduction framework. *Psychological Bulletin, 134*(2), 207-222.

Shanteau, J., 2000. Why Do Experts Disagree?. *In:* B. Green, R. Cressy, F. Delmar, T. Eisenberg, B. Howcroft, M. Lewis, D. Schoenmaker, J. Shanteau, and R. Vivian, eds. *Risk behaviour and risk management in business life*. Dordrecht, The Netherlands: Kluwer Academic Press, 186-196.

Sholler, D., 2008. *Gartner SOA User Survey: Adoption Trends and Characteristics.* [online] Available from: http://www.gartner.com/DisplayDocument?ref=g_search&id=765720&subref=simplesearch [Accessed 14 January 2012].

Sidi, R., 2010. Data is at the heart of new science ecosystem. *Research Information* [online], Feb/March 2010. Available from: www.researchinformation.info [Accessed 2 February 2011].

Siegel, S. and Castellan, N., 1988. *Nonparametric Statistics for the Behavioral Sciences*, 2nd edition. McGraw-Hill.

Simon, H. A., 1996. *The Sciences of the Artificial.* 3rd ed. Cambridge, MA: MIT Press.

Simonis, I. and Echterhoff, J., eds. 2011. *OGC Sensor Planning Service Implementation Standard. Version: 2.0 OGC 09-000.*

Simperl, E., 2009. Reusing ontologies on the Semantic Web: A feasibility study*. Data & Knowledge Engineering*, 68, 905–925.

Sinnig, D., Chalin, P., and Khendek, F., 2009. LTS semantics for use case models. *In: ACM Symposium on Applied Computing,* 9-12 March 2009 Hawaii, USA, 365-370.

Sirin, E., Bulka, B., and Smith, M., 2010. Terp: Syntax for OWL-friendly SPARQL Queries*. In: 7th OWL Experiences and Directions Workshop.* 21-22 June 2010. San Francisco, California.

Smith, B., 2006. *Common Anatomy Reference Ontology Workshop Minutes September 8-9, 2006 in Seattle*. [online] Available from: http://www.bioontology.org/wiki/images/6/67/CARO-minutes.doc [Accessed 1 February 2012].

Smith, B., Ashburner, M., Rosse, C., Bard, J., Bug, W., Ceusters, W., Goldberg, L.J., Eilbeck, K., Ireland, A., Mungall, C.J., The OBI Consortium, Leontis, N., Rocca-Serra, P., Ruttenberg, A., Sansone, S.A., Scheuermann, R.H., Shah, N., Whetzel, P.L., and Lewis, S., 2007. The OBO Foundry: coordinated evolution of ontologies to support biomedical data integration. *Nature Biotechnology* [online]*,* 25, 1251 – 1255. Available from: http://www.nature.com/nbt/journal/v25/n11/full/nbt1346.html [Accessed 28 April 2009].

Smith, B., 2008. The Evaluation of Ontologies: Editorial Review vs. Democratic Ranking. *In: Proceedings of InterOntology 2008,* 26-27 February 2008 Tokyo, Japan. 127-138.

Smith, N.J. and Gannon, M., 2008. Political Risk In Light Rail Transit PPP Projects. *In: Proceedings of the Institution of Civil Engineers Management, Procurement and Law,* November 2008. 179–185.

Sowa, JF., 2010. *Ontology.* [online] Available from: http://www.jfsowa.com/ontology/ [Accessed 19 September 2011].

Spyns, P., 2005. *EvaLexon: Assessing triples mined from texts.* Technical Report 09, STAR Lab, Brussels, Belgium.

Sredjevic, B. and Medeiros, Y.D.P., 2008. Fuzzy AHP Assessment of Water Management Plans. *Water Resources Management,* 22(7), 877-894.

Staab, S. and Maedche, A., 2000. Axioms are Objects, too - Ontology Engineering beyond the modelling of Concepts and Relations*. In: Proceedings of the ECAI 2000 Workshop on Ontologies and Problem Solving Methods,* 21-22 August 2000 Berlin, Germany. 1-16.

Stanford University, 2011. *Protégé Wiki-Ontology Metrics.* [online] Available from: http://protegewiki.stanford.edu/wiki/Pr4_UG_rv_On_Ontology_metrics [Accessed 13 November 2011].

Stanford University, 2012a. *Protege-Home Page.* [online] Available from: http://protege.stanford.edu [Accessed 21 March 2012].

Stanford University, 2012b. *Dublin Core OWL DL Schema.* [online] Available from: http://protege.stanford.edu/plugins/owl/dc/protege-dc.owl [Accessed 30 January 2012].

Staub, P., Gnägi, H.R., and Morf A., 2008. Semantic Interoperability through the Definition of Conceptual Model Transformations. *Transactions in GIS*, 12(2), 193-207.

Stecher, R. and Niederée , C., 2005**.** Ontology Fitness - Supporting Ontology Quality beyond Logical Consistency**.** *In: Formal Ontologies Meet Industry Workshop*, 9-10 June 2005 Castelnuovo del Garda, Italy.

Stevens, S.S., 1946. On the Theory of Scales of Measurement. *Science,* 103, 677-680.

Stevens, S.S., 1951. *Mathematics, Measurement and Psychophysics.* New York: John Wiley.

Stewart, T.J., 2008. Robustness Analysis and MCDA. *Newsletter of EWG MCDA* [online], 3(18). Available from: http://www.inescc.pt/~ewgmcda/ForStewart.pdf [Accessed 2 June 2010].

Steyvers, M., 2002. Multidimensional Scaling*. In: Encyclopedia of Cognitive Science.* London: Nature Publishing Group.

Stock, K., 2008*. OGC Catalogue Service Specification 2.0.2 Feature Type Catalogue Extension Package for ebRIM (ISO/TS15000-3) Profile of CSW 2.0.* OGC Discussion Paper. Version 1. OGC-07-172r1. 2008-02-19.

Stock, K., Robertson, A., Bishr, M., Stojanovic, T., Ortmann, J., Reitsma, F., and Medyckyj-Scott, D., 2009a. eScience for Sea Science: A Semantic Knowledge Infrastructure for Marine Scientists. *In: 5th IEEE International Conference on e-Science,* 9-11 December 2009. Oxford, UK.

Stock, K., Small, M., Ou, Y., and Reitsma, F., 2009b. *OWL Application Profile of CSW.* OGC Discussion Paper 09-010. Technical report.

Stock, K., Atkinson, R., Higgins, C., Small, M., Woolf, A., Millard, K., and Arctur, D., 2010. A semantic registry using a Feature Type Catalogue instead of ontologies to support spatial data infrastructures. *International Journal of Geographical Information Science,* 24(2), 231 – 252.

Stock, K., Robertson, A., and Small, M., 2011. Representing OGC Geospatial Web Services in OWL-S Web Service Ontologies. Submitted to *International Journal of Spatial Data Infrastructures Research*. (in review).

Stojanovic, L., Stojanovic, N., Gonzalez, J., and Studer, R., 2003. OntoManager – A System for the Usage-Based Ontology Management. *In: On The Move to Meaningful Internet Systems 2003 CoopIS, DOA, and ODBASE,* November 2003 Catania, Sicily, Italy. Berlin, Heidelberg : Springer. 858-875.

Strauss, A. and Corbin, J., 1998. *Basics of qualitative research.* 2nd ed. Thousand Oaks, CA: Sage.

Stuckenschmidt, H., 2008. Debugging OWL Ontologies - A Reality Check. *In: EON 2008,* 1-2 June 2008 Tenerife, Spain.

Super Decisions, 2012. *Super Decisions Software for Decision-Making- Home Page.* [online] Available from: http://www.superdecisions.com/ [Accessed 13 February 2012].

Sure, Y., Staab, S., and Studer, R., 2003. On-To-Knowledge Methodology. *In:* S. Staab and R. Studer, eds, *Handbook on Ontologies,* 1st ed. Springer. 117-132.

Takase, T., Makino, S., Kawanaka, S., Ueno, K., Ferris, C., and Ryman, A., 2008. *Definition languages for RESTful Web services: WADL vs. WSDL 2.0.* IBM Report. [online] Available from: http://download.boulder.ibm.com/ibmdl/pub/software/dw/specs/ws-wadlwsdl/WADLWSDLpaper20080621.pdf [Accessed on 11 December 2009].

Tartir, S., Arpinar, I.B., Moore, M., Sheth, A.P. and Aleman-Meza, B., 2005. OntoQA: Metric-Based Ontology Quality Analysis. *In: IEEE Workshop on Knowledge Acquisition from Distributed, Autonomous, Semantically heterogeneous Data and Knowledge Sources,* 27 November 2005 Houston, TX, USA.

Tashakkori, A. and Teddlie, C.D., 1998. *Mixed Methodology: Combining Qualitative and Quantitative Approaches*. Thousand Oaks: Sage Publications.

TDWG, 2007a. *Ontology (LSID) Vocabularies – Globals.* [online] Available from: http://rs.tdwg.org/ontology/voc/Common.rdf [Accessed July 2007].

TDWG, 2007b. *Taxon Concept (LSID) Ontology.* [online] Available from: http://rs.tdwg.org/ontology/voc/TaxonConcept.rdf [Accessed July 2007].

TDWG, 2007c. *Taxon Rank (LSID) Ontology*. [online] Available from: http://rs.tdwg.org/ontology/voc/TaxonRank [Accessed July 2007].

TDWG, 2011a. *TDWG-Home Page.* [online] Available from: http://www.tdwg.org/ [Accessed 10 September 2011].

TDWG, 2011b. *Life Science Identifiers (LSIDs).* [online] Available from: http://wiki.tdwg.org/twiki/bin/view/GUID/LSID [Accessed 16 March 2012].

Teddlie, C. and Tashakkori, A., 2009. *Foundations of mixed methods research: Integrating quantitative and qualitative techniques in the social and behavioral sciences.* Thousand Oaks, CA: Sage.

The Opengroup, 2009. *Navigating the SOA Open Standards Landscape Around Architecture. OASIS and OMG White Paper.* [online] Available from: http://www.adobe.com/devnet/livecycle/pdfs/soa_standards.pdf [Accessed 11 December 2009].

Thomas, E., Alani, H., Sleeman, D., and Brewster, C., 2005. Searching and Ranking Ontologies on The Semantic Web. *In: Workshop On Ontology Management: Searching, Selection, Ranking and Segmentation. 3rd K-CAP 2005*. 2 October 2005 Banff, Canada.

Thompson, C.M., 1941. Table Of Percentage Points of The Chi-Square Distribution. *Biometrika*, 32, 188-189.

Tirmizi, S. H., Aitken, S., Moreira, D. A., Mungall. C., Sequeda, J., Shah, N. H., and Miranker, D. P., 2011. Mapping between the OBO and OWL ontology languages. *J Biomed Semantics* [online], 2(supp.1):S3. Available from: http://www.jbiomedsem.com/content/2/S1/S3 [Accessed 7 July 2011].

Tormala, Z.L. and R.E. Petty. 2007. Contextual contrast and perceived knowledge: Exploring the implications for persuasion. *Journal of Experimental Social Psychology,* 43(1), 17-30.

Trochim, W.M., 2006. *The Research Methods Knowledge Base*, 2nd Edition. [online] Available from: http://www.socialresearchmethods.net/kb/ [Accessed 2 February 2011].

Tsarkov, D. and Horrocks, I., 2012. *FaCT++ Site.* [online] Available from: http://owl.man.ac.uk/factplusplus/ [Accessed 30 February 2012].

Tversky, A. and Kahneman, D., 1981 The Framing of Decisions and the Psychology of Choice. *Science New Series*, 211(4481), 453-458.

Ulrich, W., 2002. Critical Systems Heuristics. *In:* H.G. Daellenbach and R.L. Flood, eds. *The Informed Student Guide to Management Science.* Thomson Learning: London.

Uren, V., Cimiano, P., Iria, J., Handschuh, S., Vargas-Vera, M., Motta, E., and Ciravegna, F., 2006. Semantic annotation for knowledge management: Requirements and a survey of the state of the art, *Web Semantics: Science, Services and Agents on the World Wide Web*, 4(1), 14-28.

Urquhart, C., 2001. An Encounter with Grounded Theory: Tackling the Practical and Philosophical Issues. *In:* E. Trauth, ed. *Qualitative Research in IS: Issues and Trends*. Hershey: Idea Group Publishing, 104-140.

Uschold, M. and Grüninger, M., 2004. Ontologies and semantics for seamless connectivity. *ACM SIGMOD Record*, 33(4), 58-64.

Uschold, M., Healy, M., Williamson, K., Clark, P., and Woods, S., 1998, Ontology reuse and application. *In:* N. Guarino, ed. *Proc of the Int Conf on Formal Ontology in Information Systems - FOIS'98 (Frontiers in AI and Applications v46),* 6-8 June 1998 Trento, Italy. Amsterdam: IOS Press, 179-192.

UTAS, 2012. *UTAS Human Research Ethics.* [online] Available from: http://www.research.utas.edu.au/human_ethics/social_science_forms.htm [Accessed 16 February 2012].

Vaccari, L., Schvaiko, P., and Marchese, M., 2009. A Geo-service Semantic Integration in Spatial Data Infrastructures. *International Journal of Spatial Data Infrastructures Research*, 4, 24-51.

Vaishnavi, V. and Kuechler, W., 2004. *Design Research in Information Systems.* [online] January 20, 2004, last updated September 30, 2011. Available from: http://desrist.org/desrist. [Accessed 15 February 2012].

Valente, A. and Breuke, J., 1996. Towards principled core ontologies. *In*: B.R. Gaines and M. Mussen, eds. *Proceedings of the KAW-96*, 9-14 November 1996 Banff, Canada.

Van Assem, M., Menken, M.R., Schreiber, G., Wielmaker, J., and Wileinga, B., 2004. A method for converting thesauri to RDF/OWL. *In:* S.A. McIlraith, D. Plexousakis, and F. van Harmelan, eds., *Proceedings Of The Third International Semantic Web Conference (ISWC 2004),* 7-11 November 2004 Hiroshima, Japan.

Van Den Honert, R.C. and Lootsma, F.A., 1996, Group preference aggregation in the multiplicative AHP: the model of the group decision process and Pareto optimality, *European Journal of Operational Research,* 96(2), 363-370.

Van der Werf, B., Adamescu, M., Ayromlou, M., Bertrand, N., Borovec, J., Boussard, H., Cazacu, C., van Daele, T., Datcu, S., Frenzel, M., Hammen, V., Karasti, H., Kertesz, M., Kuitunen, P., Lane, M., Lieskovsky, J., Magagna, B., Peterseil, J., Rennie, S., Schentz, H., Schleidt, K., and Tuominen, L., 2009. *SERONTO A Socio-Ecological Research and Observation Ontology: the core ontology.* Version 2009-04-09, Project GOCE-CT-2003-505298.

van Heijst, G., Schreiber, A.T., and Wielinga, B.J. 1997. Using Explicit Ontologies in KBS Development. *Int. Journal of Human-Computer Studies,* 46, 293-310.

van Laarhoven, P.J.M. and Pedrycz, W., 1983. Fuzzy extension for Saaty's priority theory, *Fuzzy Sets and Systems,* 11(3), 229-241.

Van Rees, R., 2003. Clarity in the usage of the terms ontology, taxonomy and classification. *In: Proceedings of the CIB W78's 20th International Conference on Construction IT, Construction IT Bridging the Distance*, 23-25 April 2003 Waiheke, New Zealand. 432-440.

Vanopstal, K., Vander Stichele, R., Laureys, G., and Buysschaert, J., 2009. Vocabularies and Retrieval Tools in Biomedicine: Disentangling the Terminological Knot. *Journal of Medical Systems,* 99.

Venable, J., 2009. Identifying and Addressing Stakeholder Interests in Design Science Research: An Analysis Using Critical Systems Heuristics*. In: Information Systems – Creativity and Innovation in Small and Medium-Sized Enterprises.* 301, 93–112

Verschuren, P. and Hartog, R., 2005. Evaluation in Design-Oriented Research. *Quality & Quantity,* 39, 733-762.

Vidal, V.M.P., Sacremento, E.R., de Macedo, J.A.F., and Casanova, M.A., 2009. An Ontology-Based Framework For Geographic Data Integration*. In: ER '09 Proceedings of the ER 2009 Workshops on Advances In Conceptual Modelling Challenging Perspectives, Lecture Notes In Computer Science, Volume 5833.* 9-12 November 2009 Gramado, Brazil, Heidelberg : Springer-Verlag, 337.

Viljanen, K., Tuominen, J., Salonoja, M., and Hyvonen, E., 2010. Linked Open Ontology Services. *In: Proceedings of the 1st Workshop on Ontology Repositories and Editors for the Semantic Web (ORES-2010),* 31 May 2010 Hersonissos, Crete, Greece.

Vinoski, S., 2007. REST Eye for the SOA Guy**.** *IEEE Internet Computing*, 11(1), 82-84.

Voges, U. and Senkler, K., 2007. *OpenGIS Catalogue Services Specification 2.0.2 - ISO Metadata Application Profile*. OGC 07-045, Open Geospatial Consortium Inc. [online] Available from: http://portal.opengeospatial.org/files/?artifact_id=21460 [Accessed 8 June 2008].

Vrandečić, D., 2010. Ontology Evaluation for the Web-PhD Proposal. *In:* J. Diederich, E. Motta, and E. Bontas, eds. *Proceedings of the KnowledgeWeb PhD Symposium KWEPSY 2006,* 17 June 2006 Budva, Montenegro.

Vretanos, P., 2005. *OGC Implementation Specification 04-094: Web Feature Service Implementation Specification.* [online] Available from: http://portal.opengeospatial.org/files/?artifact_id=8339 [Accessed 23 March 2010].

VSTO, 2012. *Virtual Solar Terrestrial Observatory.* [online] Available from: http://www.vsto.org/ [Accessed 6 February 2012].

Wallenius, J., Dyer, J.S., Fishburn, P.C., Steuer, R.E., Zionts, S., and Deb, K., 2008. Multiple Criteria Decision Making, Multiattribute Utility Theory: Recent Accomplishments and What Lies Ahead. *Management Science,* 54(7), 1336-1349.

Walter, V. and Fritsch, D., 1999. Matching Spatial Data Sets: a Statistical Approach. *International Journal of Geographic Information Sciences,* 13(5), 445-473.

Wang, H., Horridge, M., Rector, A.L., Drummond, N., and Seidenberg, J., 2005. Debugging OWL-DL Ontologies: A Heuristic Approach. *In: International Semantic Web Conference 2005*, 6-10 November 2005 Galway, Ireland. 745-757.

Ware, J., 2009. *A Systematic Analysis To Identify, Mitigate, Quantify, And Measure Risk Factors Contributing To Falls.* Thesis (PhD), University of Central Florida, Orlando, Florida.

Waring, T. and Wainwright, D., 2008. Issues and Challenges in the Use of Template Analysis: Two Comparative Case Studies from the Field. *The Electronic Journal Of Business Research Methods.* 6(1), 85-94.

Warren, L., 2004. *Uncertainties in the Analytic Hierarchy Process.* DSTO Information Sciences Laboratory, Australian Government, Dept. Of Defence, Defence Science and Technology Organisation. Scientific Publication. Report Number: DSTO-TN-0597.

Washington, N. and Lewis, S., 2008. Ontologies: Scientific Data Sharing Made Easy. *Nature Education*, 1(3).

Wehrmann, T., Gebhardt, S., Klinger, V., and Künzer, C., 2011. Data processing using Web Processing Service orchestration within a Spatial Data Infrastructure. *In: 34th International Symposium for remote Sensing of the Environment,* 10-15 April 2011 Sydney, Australia.

Whiteside, A., ed., 2005. *OGC Web Services Common Specification Version 1.0.* Wayland, Massachusetts.

Whiteside, A. and Greenwood, J., eds. 2010*. OGC Web Services Common Standard v2.0.0.* Number OGC 06-121r9. [online] Available from: http://portal.opengeospatial.org/files/?artifact_id=38867 [Accessed 23 June 2011].

Whyte, L.L., 1969. *Hierarchical Structures* . American Elsevier, New York, N.Y.

Wikipedia, 2012a. *Wikipedia-Semantics.* [online] Available from: http://en.wikipedia.org/wiki/Semantics [Accessed 22 February 2012].

Wikipedia, 2012b. *Wikipedia-Interoperability.* [online] Available from: http://en.wikipedia.org/wiki/Interoperability [Accessed 3 February 2012].

Wikipedia, 2012c *Wikipedia-Notation3.* [online] Available from: http://en.wikipedia.org/wiki/Notation3 [Accessed 4 January 2012].

Williams, M., Dennis, A., Stam, A., and Aronson J., 2007. The impact of DSS use and information load on errors and decision quality, *European Journal of Operational Research,* 176(1), 468-81.

Williams, D., 2011. The Planet at Their Fingertips: Climate Modeling Data Heats Up. *SciDAC Review* [online], 9. Available from: http://www.scidacreview.org/0902/html/esg.html [Accessed 2 February 2011].

Wilson, T.D. and Schooler, J.W., 1991. Thinking too much: Introspection can reduce the quality of preferences and decisions. *Journal of Personality and Social Psychology,* 60(2), 181–192.

Wissner, J. and Spivack, N., 2009. *Semantic Web Use Cases and Case Studies-Case Study Twine.* [online] Available from: http://www.w3.org/2001/sw/sweo/public/UseCases/Twine/ [Accessed 1 February 2012].

Woolf, A., 2007*. Climate Science Modelling Language Version 2 User Manual.* [online] British Atmospheric Data Centre. Available from: http://csml.badc.rl.ac.uk/index.php?option=com_content&view=article&id=53%3Auserguide&catid=34%3Adocuments&Itemid=55 [Accessed 2 June 2011].

Worboys, M., 1994. Object-oriented approaches to Georefeneced information, *IJGIS*, 8(4), 385-399.

W3C, 2001. *N-Triples W3C RDF Core WG Internal Working Draft.* [online] Available from: http://www.w3.org/2001/sw/RDFCore/ntriples/ [Accessed 4 December 2009].

W3C, 2007. *W3C RDF Powered Validation Service-Home Page.* [online] Available from: http://www.w3.org/RDF/Validator/ [Accessed 24 November 2009].

W3C, 2011. *OWL-DL compatible version of SKOS.* [online] Available from: http://www.w3.org/2009/08/skos-reference/skos-owl1-dl.rdf [Accessed 30 December 2011].

W3C, 2012. *Semantic Sensor Networks Community Group-Home Page.* [online] Available from: http://www.w3.org/community/ssn-cg/ [Accessed 16 March 2012].

Xu, R., 2000. Fuzzy least square priority method in the analytic hierarchy process. *Fuzzy Sets and Systems,* 112(3), 395-404.

Yang, C., Raskin, R., Goodchild, M., and Gahegan, M., 2010. Geospatial Cyberinfrastructure: Past, present and Future, *Computers, Environment and Urban Systems,* 34, 264–277.

Yin, R.K., 1994**.** Case Study Research: Design and Methods. 2nd edn. Newbury Park: Sage.

Yoo, H., Moon, A., Par, Y., and Kim, S., 2010. *Framework of Semantic Web Services and Service Mediation*. [online] Available from: http://www.iiis.org/CDs2010/CD2010SCI/SCI_2010/PapersPdf/SA487TV.pdf [Accessed 19 July 2011].

Yu, J., Thom, J., and Tam, A., 2005. Evaluating an ontology with OntoClean. *In: Proceedings of 10th Australasian Document Computing Symposium (ADCS)*, 12 December 2005 University of Sydney.

Yu, J., Thom, J.A. and Tam. A., 2007. Ontology evaluation using Wikipedia categories for browsing. *In: Proc. 16th ACM Conference on Information and Knowledge Management,* 6-9 November 2007 Lisboa, Portugal.

Yu**,** J., 2008*. Requirements-oriented methodology for evaluating* ontologies*,* Thesis (PhD). School of Computer Science and Information Technology. RMIT. Melbourne, Australia.

Yu, J., Thom, J.A., and Tam, A., 2009. Requirements-oriented methodology for evaluating ontologies. *Information Systems*, 34(8), 766–791.

Yue, P., Gong, J., Di, L., He, L., and Wei, Y., 2010. Semantic Provenance Registration and Discovery using Geospatial Catalogue Service. *In:* S.S. Sahoo, J. Zhao, P. Missier and J.M. Gomez-Perez, eds. *Proceedings of the Second International Workshop on the role of Semantic Web in Provenance Management (SWPM 2010) Workshop at the 9th International Semantic Web Conference (ISWC2010) - ISWC 2010 Workshops Volume VII,* 7 November 2010 Shanghai, China.

Zhang, X., Cheng, G., and Qu, Y., 2007. Ontology Summarization Based on RDF Sentence Graph*. In: 16th International World Wide Web Conference,* 8-12 May 2007 Banff, Alberta, Canada.

Zhang, X., Cheng, G., Ge, W., and Qu, Y., 2009. Summarizing Vocabularies in the Global Semantic Web. *Journal of Computer Science and Technology,* 24(1), 165-174.

Zhang, C., Zhao, T., Li, W., and Osleeb, J.P, 2010. Towards logic-based geospatial feature discovery and integration using web feature service and geospatial semantic web. *International Journal of Geographical Information Science,* 24(6), 903-923.

Zimmermann, A., 2010. Ontology Recommendation for the Data Publishers. *In: Proceedings of the 1st Workshop on Ontology Repositories and Editors for the Semantic Web*, 31 May 2010 Hersonissos, Crete.

Ontology Management and Selection In Re-Use Scenarios

APPENDICES

Kim Finney (B.Sc, M.Sc)

Submitted in fulfilment of the requirements for the degree of Doctor of Philosophy

University of Tasmania

School of Computing and Information Systems

October 2012

List Of Appendices

Appendix 1 – On-line Screening Survey Questions

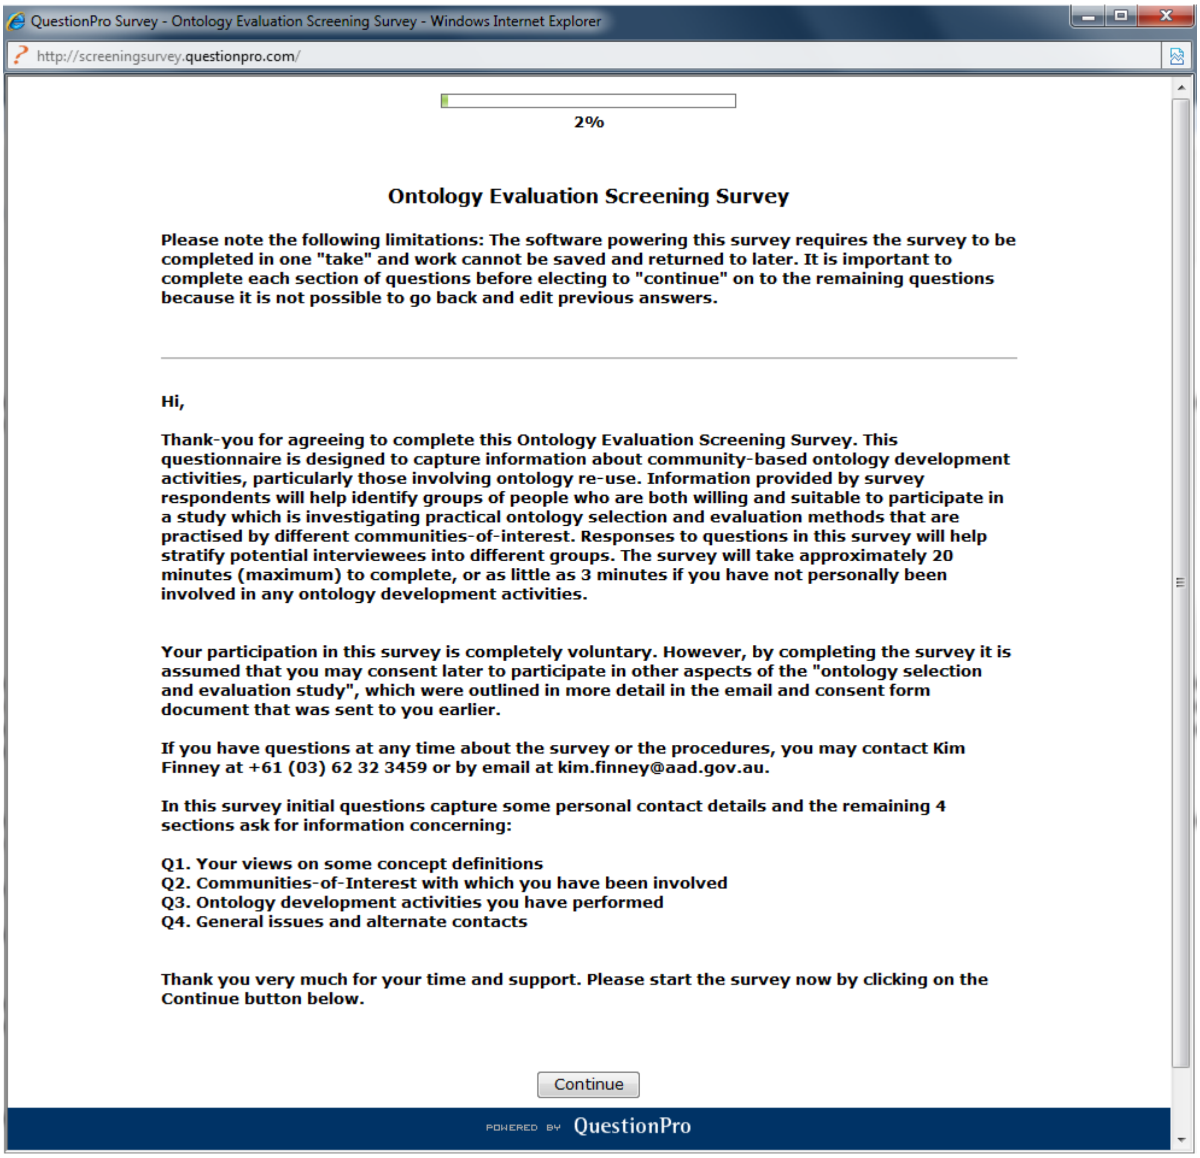

QuestionPro Survey - Ontology Evaluation Screening Survey - Windows Internet Explorer

http://screeningsurvey.questionpro.com/

2%

Ontology Evaluation Screening Survey

Please note the following limitations: The software powering this survey requires the survey to be completed in one "take" and work cannot be saved and returned to later. It is important to complete each section of questions before electing to "continue" on to the remaining questions because it is not possible to go back and edit previous answers.

Hi,

Thank-you for agreeing to complete this Ontology Evaluation Screening Survey. This questionnaire is designed to capture information about community-based ontology development activities, particularly those involving ontology re-use. Information provided by survey respondents will help identify groups of people who are both willing and suitable to participate in a study which is investigating practical ontology selection and evaluation methods that are practised by different communities-of-interest. Responses to questions in this survey will help stratify potential interviewees into different groups. The survey will take approximately 20 minutes (maximum) to complete, or as little as 3 minutes if you have not personally been involved in any ontology development activities.

Your participation in this survey is completely voluntary. However, by completing the survey it is assumed that you may consent later to participate in other aspects of the "ontology selection and evaluation study", which were outlined in more detail in the email and consent form document that was sent to you earlier.

If you have questions at any time about the survey or the procedures, you may contact Kim Finney at +61 (03) 62 32 3459 or by email at kim.finney@aad.gov.au.

In this survey initial questions capture some personal contact details and the remaining 4 sections ask for information concerning:

Q1. Your views on some concept definitions
Q2. Communities-of-Interest with which you have been involved
Q3. Ontology development activities you have performed
Q4. General issues and alternate contacts

Thank you very much for your time and support. Please start the survey now by clicking on the Continue button below.

Continue

POWERED BY QuestionPro

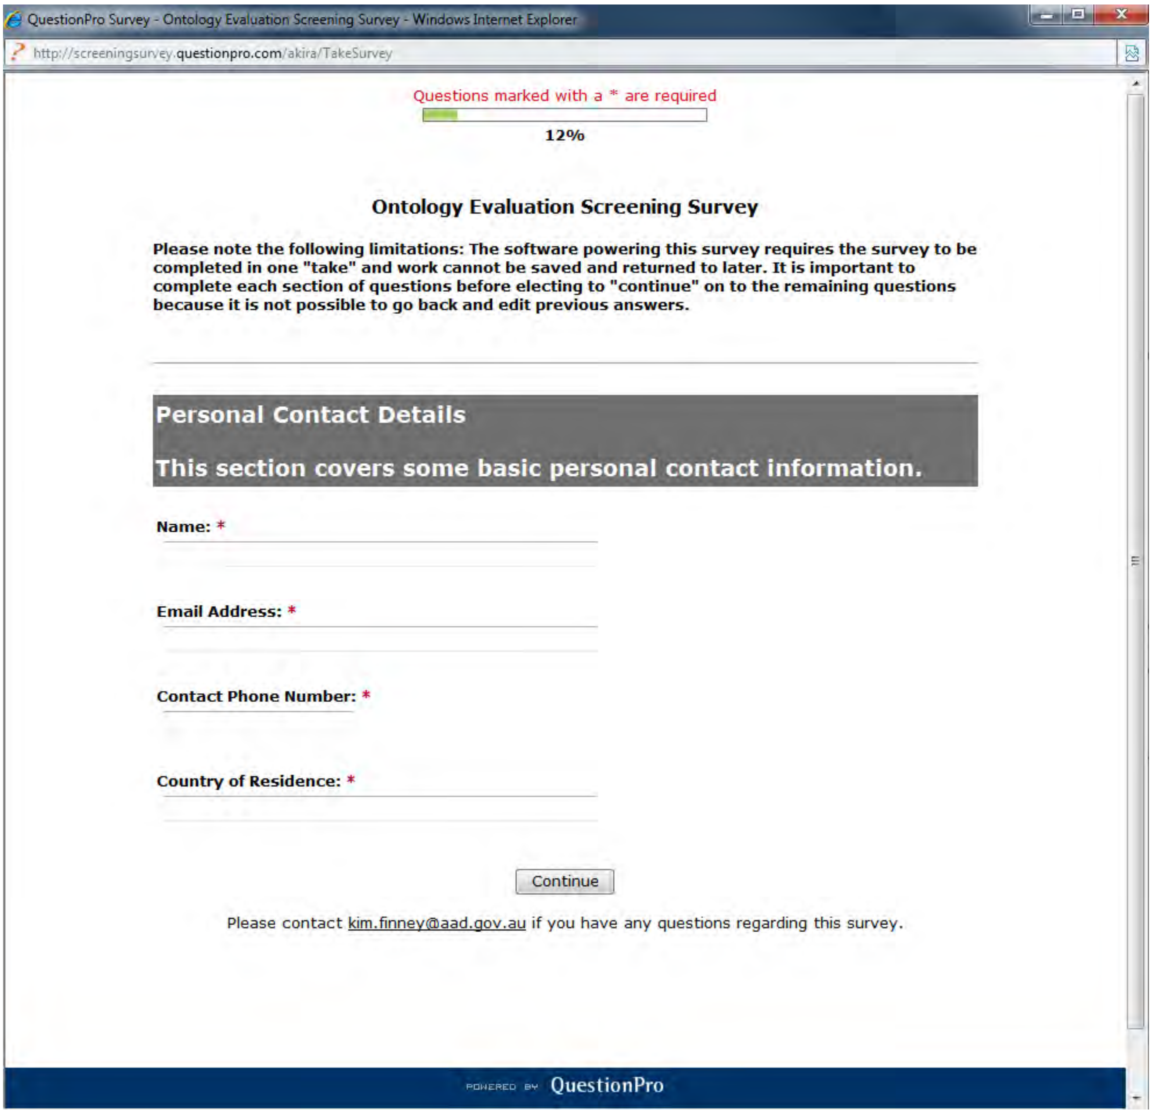

QuestionPro Survey - Ontology Evaluation Screening Survey - Windows Internet Explorer

http://screeningsurvey.questionpro.com/akira/TakeSurvey

Questions marked with a * are required

12%

Ontology Evaluation Screening Survey

Please note the following limitations: The software powering this survey requires the survey to be completed in one "take" and work cannot be saved and returned to later. It is important to complete each section of questions before electing to "continue" on to the remaining questions because it is not possible to go back and edit previous answers.

Personal Contact Details

This section covers some basic personal contact information.

Name: *

Email Address: *

Contact Phone Number: *

Country of Residence: *

Continue

Please contact kim.finney@aad.gov.au if you have any questions regarding this survey.

POWERED BY QuestionPro

QuestionPro Survey - Ontology Evaluation Screening Survey - Windows Internet Explorer

http://screeningsurvey.questionpro.com/akira/TakeSurvey

Questions marked with a * are required

34%

Ontology Evaluation Screening Survey

Please note the following limitations: The software powering this survey requires the survey to be completed in one "take" and work cannot be saved and returned to later. It is important to complete each section of questions before electing to "continue" on to the remaining questions because it is not possible to go back and edit previous answers.

Q1 Definitions

In order to place the responses that you provide later in the survey in an appropriate context, this first set of questions asks you to indicate whether you agree with the definitions provided for terms used in this survey. If you don't agree with any of the definitions presented please provide your own interpretation of the term, as you would apply it in your everyday work.

Q1.1 Do you agree with the following definition of "ontology" ? If "no", please supply your own definition.

An ontology makes explicit the types of concepts, or real-world objects used in a domain and generally includes information about concept properties, concept relationships and value restrictions on both properties and concepts.

*

- Yes
- No, an alternate definition

is:

Q1.2 Do you agree with the following definition of "light-weight ontology" ? If "no", please supply your own definition.

A light-weight ontology is one which may lack formalism (inbuilt rules and syntax) required to support automatic reasoning.

*

- Yes
- No, an alternate definition

is:

Q1.3 Do you agree with the following definition of a "formal ontology" ? If "no", please supply your own definition.

A formal ontology is one that is capable of being used for automatic reasoning through its expression in a language capable of supporting first-order logical inferences.

*

- Yes
- No, an alternate definition

is:

QuestionPro Survey - Ontology Evaluation Screening Survey - Windows Internet Explorer

http://screeningsurvey.questionpro.com/akira/TakeSurvey

Q1.4 Do you agree with the following definition of "ontology re-use" ? If "no", please supply your own definition.

Ontology re-use is the process in which existing ontological components are used as input to generate new ontologies.

*

- Yes
- No, an alternate definition is:

Q1.5 Do you agree with the following definition of "ontology encoding language" ? If "no", please supply your own definition.

An ontology encoding language is one chosen to represent a conceptual knowledge model. Languages may differ in terms of their expressiveness, rigour and ability to code unambiguous semantics.

*

- Yes
- No, an alternate definition is:

Q1.6 Do you agree with the following definition of "ontology selection" ? If "no", please supply your own definition.

Ontology selection is the process that is used to determine if one or more ontologies, or ontological components satisfy certain predefined criteria.

*

- Yes
- No, an alternate definition is:

Q1.7 Do you agree with the following definition of "ontology evaluation" ? If "no", please supply your own definition.

Ontology evaluation is a task within the ontology selection process where an assessment against any given selection criteria takes place.

*

- Yes
- No, an alternate definition is:

POWERED BY QuestionPro

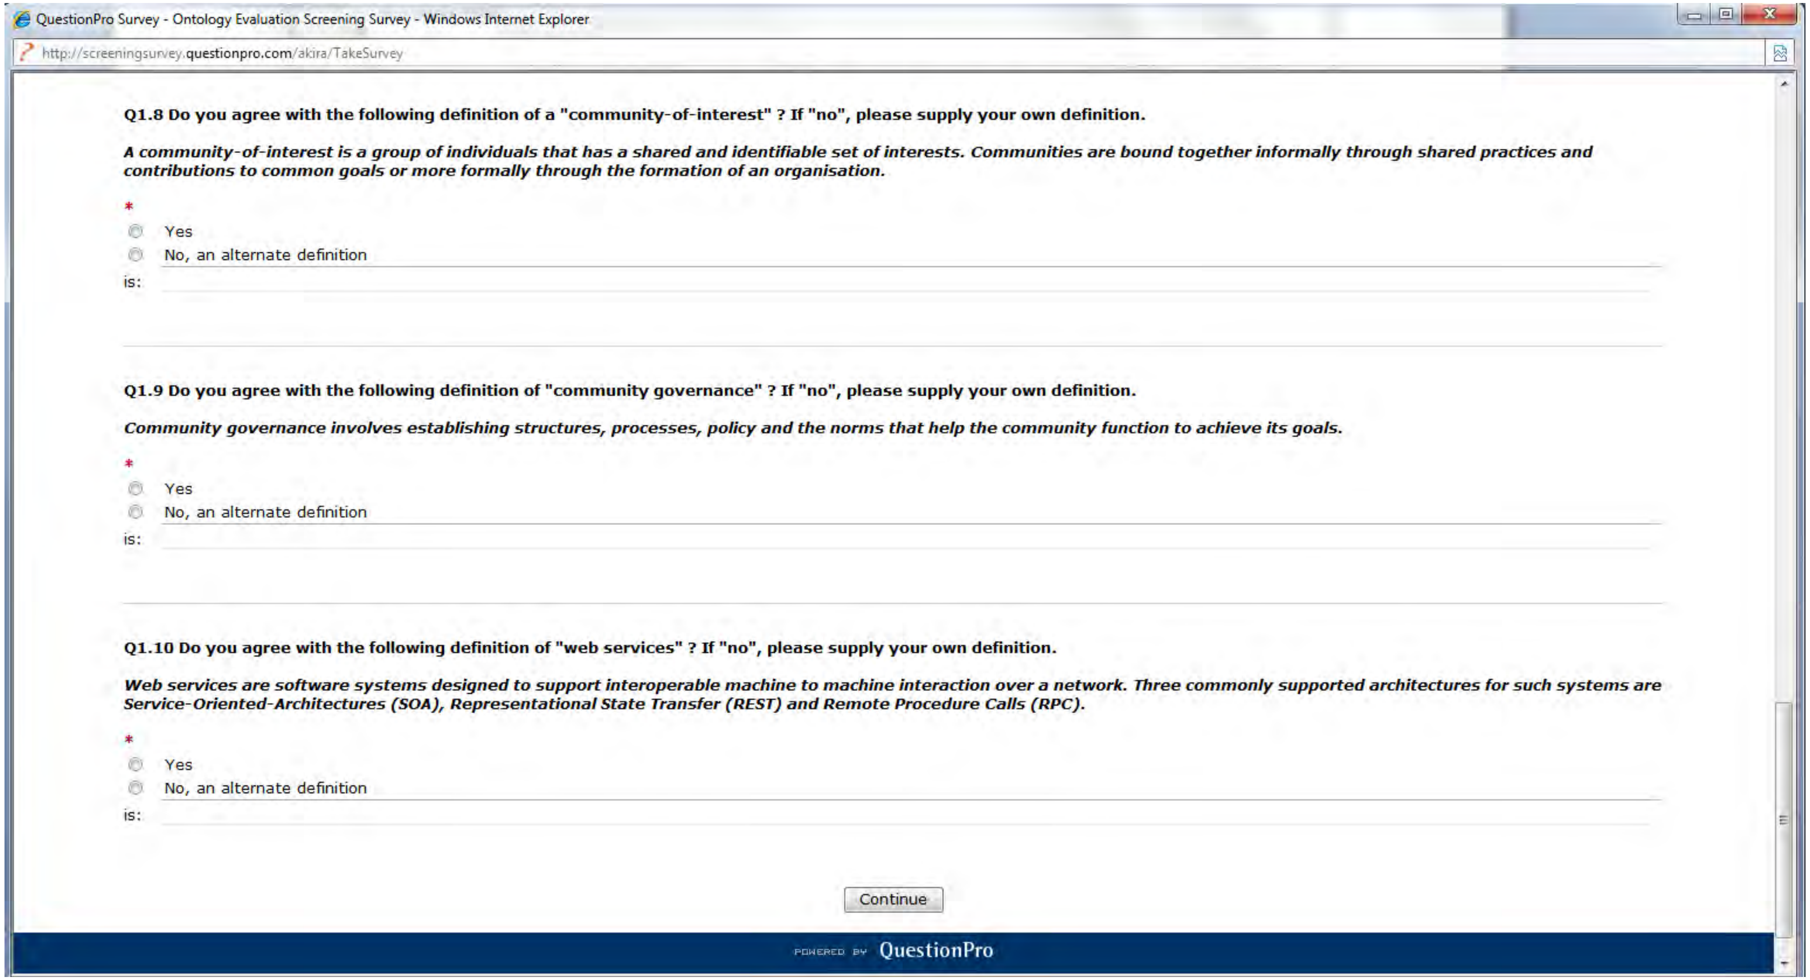
QuestionPro Survey - Ontology Evaluation Screening Survey - Windows Internet Explorer

http://screeningsurvey.questionpro.com/akira/TakeSurvey

Q1.8 Do you agree with the following definition of a "community-of-interest" ? If "no", please supply your own definition.

A community-of-interest is a group of individuals that has a shared and identifiable set of interests. Communities are bound together informally through shared practices and contributions to common goals or more formally through the formation of an organisation.

*

- Yes
- No, an alternate definition is:

Q1.9 Do you agree with the following definition of "community governance" ? If "no", please supply your own definition.

Community governance involves establishing structures, processes, policy and the norms that help the community function to achieve its goals.

*

- Yes
- No, an alternate definition is:

Q1.10 Do you agree with the following definition of "web services" ? If "no", please supply your own definition.

Web services are software systems designed to support interoperable machine to machine interaction over a network. Three commonly supported architectures for such systems are Service-Oriented-Architectures (SOA), Representational State Transfer (REST) and Remote Procedure Calls (RPC).

*

- Yes
- No, an alternate definition is:

Continue

POWERED BY QuestionPro

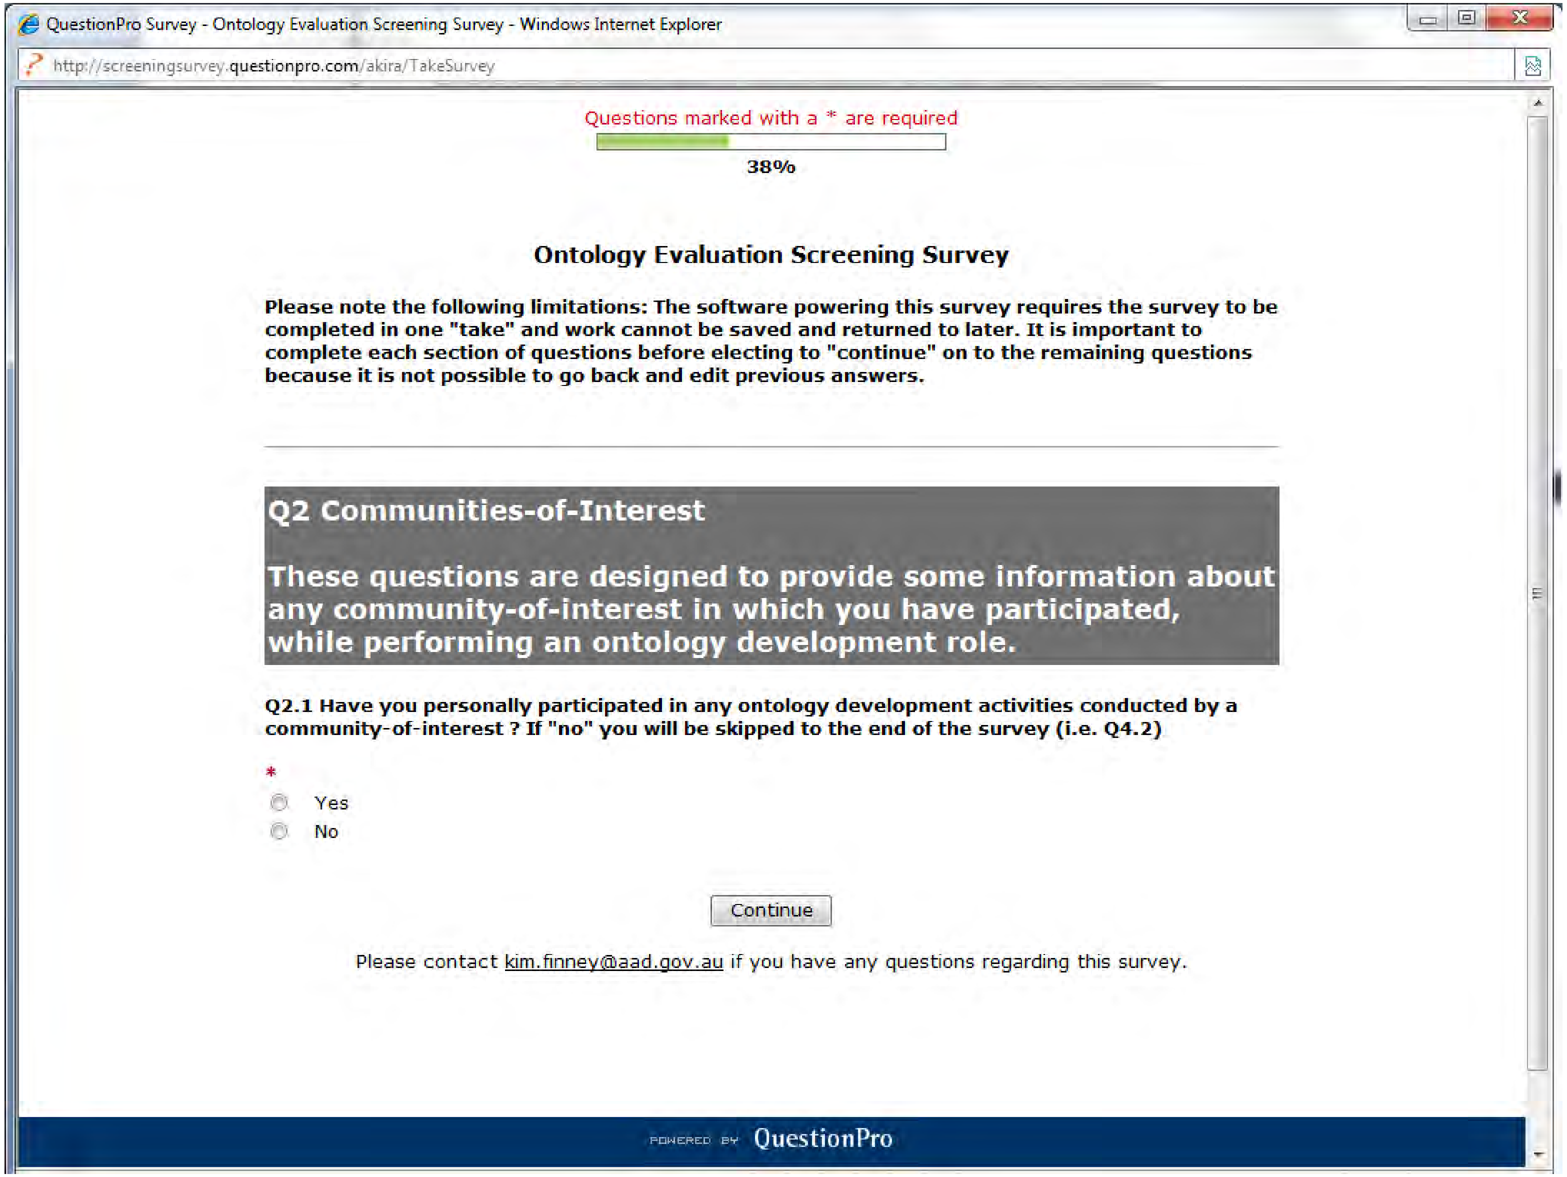
QuestionPro Survey - Ontology Evaluation Screening Survey - Windows Internet Explorer

http://screeningsurvey.questionpro.com/akira/TakeSurvey

Questions marked with a * are required

38%

Ontology Evaluation Screening Survey

Please note the following limitations: The software powering this survey requires the survey to be completed in one "take" and work cannot be saved and returned to later. It is important to complete each section of questions before electing to "continue" on to the remaining questions because it is not possible to go back and edit previous answers.

Q2 Communities-of-Interest

These questions are designed to provide some information about any community-of-interest in which you have participated, while performing an ontology development role.

Q2.1 Have you personally participated in any ontology development activities conducted by a community-of-interest ? If "no" you will be skipped to the end of the survey (i.e. Q4.2)

*

- Yes
- No

Continue

Please contact kim.finney@aad.gov.au if you have any questions regarding this survey.

POWERED BY QuestionPro

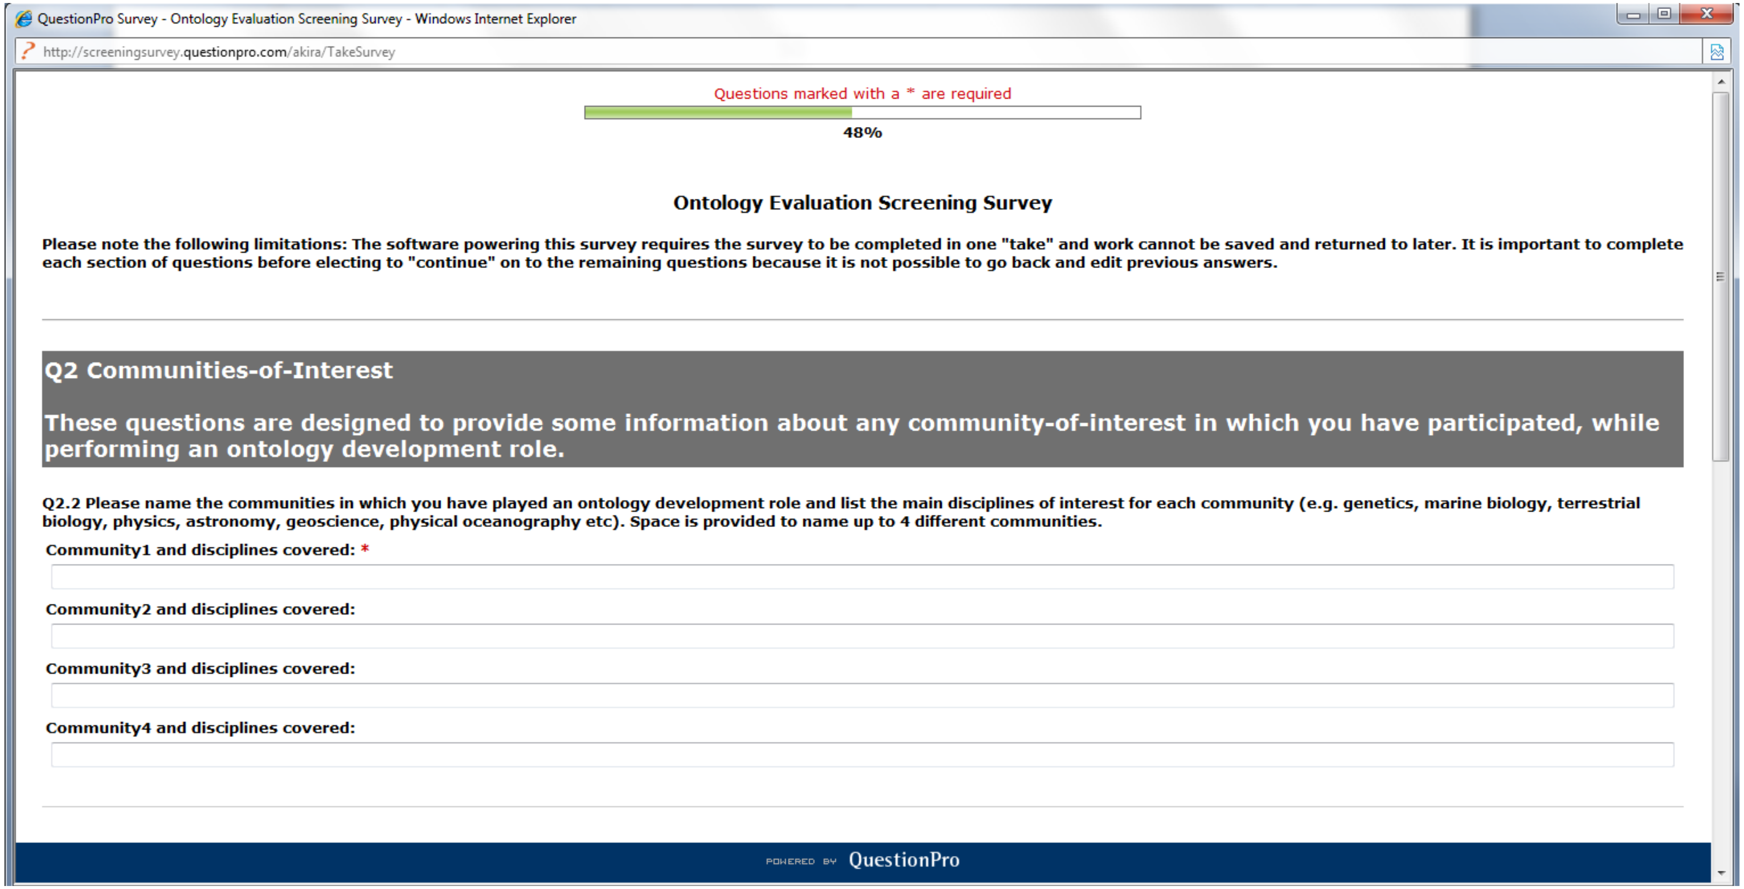

QuestionPro Survey - Ontology Evaluation Screening Survey - Windows Internet Explorer

http://screeningsurvey.questionpro.com/akira/TakeSurvey

Questions marked with a * are required

48%

Ontology Evaluation Screening Survey

Please note the following limitations: The software powering this survey requires the survey to be completed in one "take" and work cannot be saved and returned to later. It is important to complete each section of questions before electing to "continue" on to the remaining questions because it is not possible to go back and edit previous answers.

Q2 Communities-of-Interest

These questions are designed to provide some information about any community-of-interest in which you have participated, while performing an ontology development role.

Q2.2 Please name the communities in which you have played an ontology development role and list the main disciplines of interest for each community (e.g. genetics, marine biology, terrestrial biology, physics, astronomy, geoscience, physical oceanography etc). Space is provided to name up to 4 different communities.

Community1 and disciplines covered: *

Community2 and disciplines covered:

Community3 and disciplines covered:

Community4 and disciplines covered:

POWERED BY QuestionPro

QuestionPro Survey - Ontology Evaluation Screening Survey - Windows Internet Explorer

http://screeningsurvey.questionpro.com/akira/TakeSurvey

Q2.3 Please indicate if the communities that you listed above have national or international membership, or both, by checking the appropriate boxes. Please answer the question using the communities listed in the same order as you described them in Q2.2 above.

	National Membership	International Membership
Community1	☐	☐
Community2	☐	☐
Community3	☐	☐
Community4	☐	☐

Q2.4 Please indicate if the communities that you listed above at Q2.2 have an explicitly formalised governance mechanism, an informal governance mechanism or no obvious governance mechanism, by checking the appropriate radio button. Please answer the question using the communities listed in the same order as you described them in Q2.2 above.

	Formal Governance...	Informal Governance...	No Governance
Community1	○	○	○
Community2	○	○	○
Community3	○	○	○
Community4	○	○	○

Q2.5 For the communities that you listed at Q2.2 please indicate if there are any on-line resources, web sites or wikis that describe the community and/or its activities. Please answer the question using the communities listed in the same order as you described them in Q2.2 above.

Community1: Resources

Community2: Resources

Community3: Resources

Community4: Resources

Continue

Please contact kim.finney@aad.gov.au if you have any questions regarding this survey.

POWERED BY QuestionPro

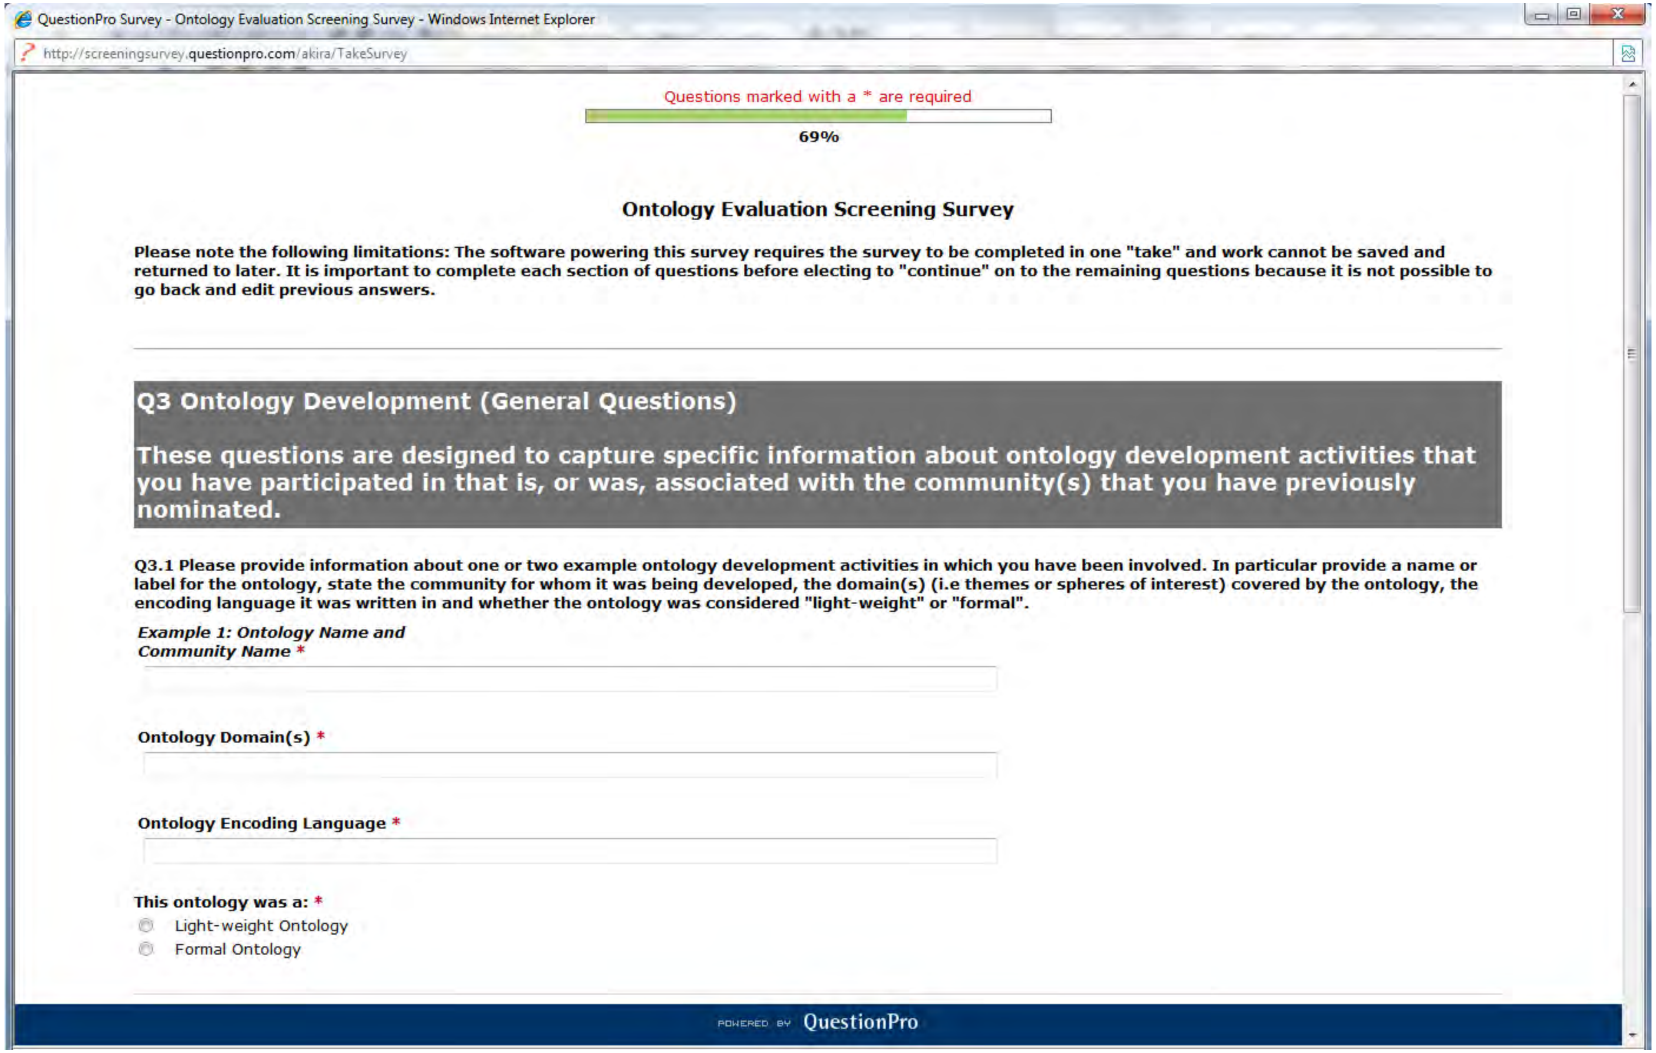

QuestionPro Survey - Ontology Evaluation Screening Survey - Windows Internet Explorer

http://screeningsurvey.questionpro.com/akira/TakeSurvey

Questions marked with a * are required

69%

Ontology Evaluation Screening Survey

Please note the following limitations: The software powering this survey requires the survey to be completed in one "take" and work cannot be saved and returned to later. It is important to complete each section of questions before electing to "continue" on to the remaining questions because it is not possible to go back and edit previous answers.

Q3 Ontology Development (General Questions)

These questions are designed to capture specific information about ontology development activities that you have participated in that is, or was, associated with the community(s) that you have previously nominated.

Q3.1 Please provide information about one or two example ontology development activities in which you have been involved. In particular provide a name or label for the ontology, state the community for whom it was being developed, the domain(s) (i.e themes or spheres of interest) covered by the ontology, the encoding language it was written in and whether the ontology was considered "light-weight" or "formal".

***Example 1: Ontology Name and Community Name* ***

Ontology Domain(s) *

Ontology Encoding Language *

This ontology was a: *

- Light-weight Ontology
- Formal Ontology

POWERED BY QuestionPro

QuestionPro Survey - Ontology Evaluation Screening Survey - Windows Internet Explorer

http://screeningsurvey.questionpro.com/akira/TakeSurvey

This ontology was a: *

- Light-weight Ontology
- Formal Ontology

Example 2: Ontology Name and Community Name

Ontology Domain(s)

Ontology Encoding Language

This ontology was a:

- Light-weight Ontology
- Formal Ontology

Q3.2 In the ontology development exercises that you were involved in did you re-use any existing ontologies or ontological components ? If "yes" please name a community in which this ontology re-use took place. If "no" you will be skipped to Q4.1

*

- No
- Yes. Name community

Continue

Please contact kim.finney@aad.gov.au if you have any questions regarding this survey.

POWERED BY QuestionPro

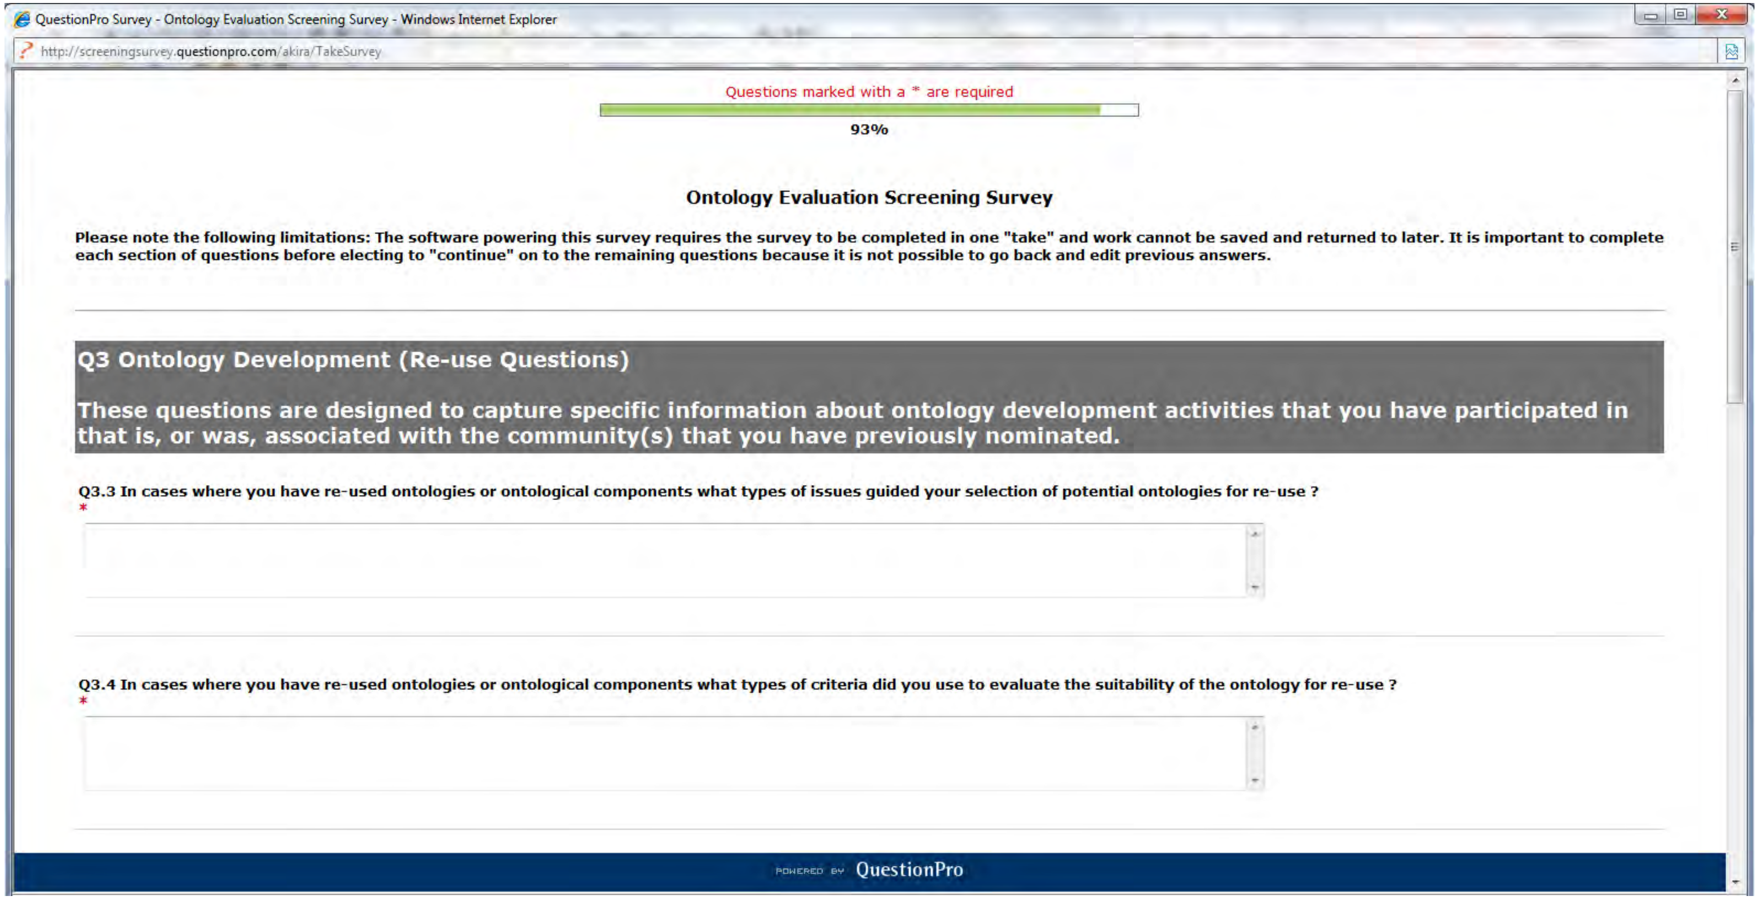

QuestionPro Survey - Ontology Evaluation Screening Survey - Windows Internet Explorer

http://screeningsurvey.questionpro.com/akira/TakeSurvey

Questions marked with a * are required

93%

Ontology Evaluation Screening Survey

Please note the following limitations: The software powering this survey requires the survey to be completed in one "take" and work cannot be saved and returned to later. It is important to complete each section of questions before electing to "continue" on to the remaining questions because it is not possible to go back and edit previous answers.

Q3 Ontology Development (Re-use Questions)

These questions are designed to capture specific information about ontology development activities that you have participated in that is, or was, associated with the community(s) that you have previously nominated.

Q3.3 In cases where you have re-used ontologies or ontological components what types of issues guided your selection of potential ontologies for re-use ?
*

Q3.4 In cases where you have re-used ontologies or ontological components what types of criteria did you use to evaluate the suitability of the ontology for re-use ?
*

POWERED BY QuestionPro

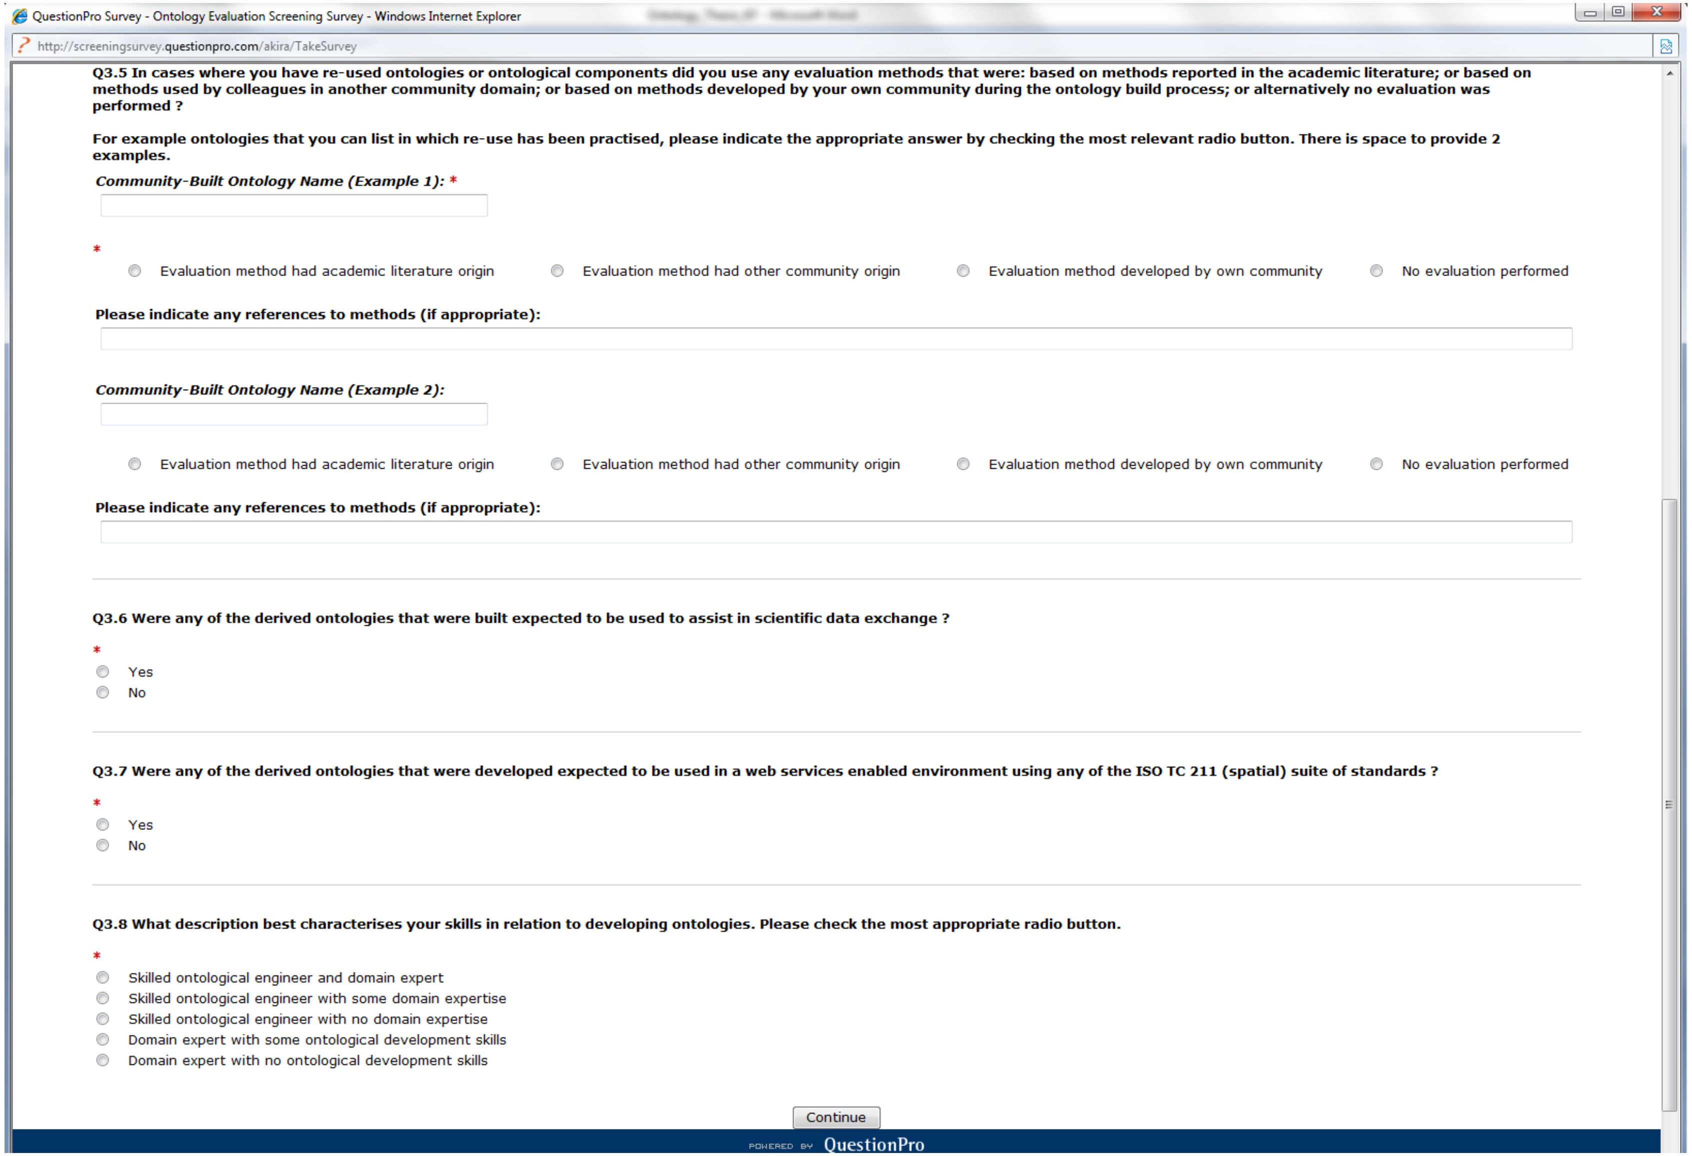

QuestionPro Survey - Ontology Evaluation Screening Survey - Windows Internet Explorer

http://screeningsurvey.questionpro.com/akira/TakeSurvey

Q3.5 In cases where you have re-used ontologies or ontological components did you use any evaluation methods that were: based on methods reported in the academic literature; or based on methods used by colleagues in another community domain; or based on methods developed by your own community during the ontology build process; or alternatively no evaluation was performed ?

For example ontologies that you can list in which re-use has been practised, please indicate the appropriate answer by checking the most relevant radio button. There is space to provide 2 examples.

Community-Built Ontology Name (Example 1): *

*

- Evaluation method had academic literature origin
- Evaluation method had other community origin
- Evaluation method developed by own community
- No evaluation performed

Please indicate any references to methods (if appropriate):

Community-Built Ontology Name (Example 2):

- Evaluation method had academic literature origin
- Evaluation method had other community origin
- Evaluation method developed by own community
- No evaluation performed

Please indicate any references to methods (if appropriate):

Q3.6 Were any of the derived ontologies that were built expected to be used to assist in scientific data exchange ?

*

- Yes
- No

Q3.7 Were any of the derived ontologies that were developed expected to be used in a web services enabled environment using any of the ISO TC 211 (spatial) suite of standards ?

*

- Yes
- No

Q3.8 What description best characterises your skills in relation to developing ontologies. Please check the most appropriate radio button.

*

- Skilled ontological engineer and domain expert
- Skilled ontological engineer with some domain expertise
- Skilled ontological engineer with no domain expertise
- Domain expert with some ontological development skills
- Domain expert with no ontological development skills

Continue

POWERED BY QuestionPro

QuestionPro Survey - Ontology Evaluation Screening Survey - Windows Internet Explorer

http://screeningsurvey.questionpro.com/akira/TakeSurvey

Questions marked with a * are required

100%

Ontology Evaluation Screening Survey

Please note the following limitations: The software powering this survey requires the survey to be completed in one "take" and work cannot be saved and returned to later. It is important to complete each section of questions before electing to "continue" on to the remaining questions because it is not possible to go back and edit previous answers.

Q4 General

Q4.1 If the community(s) you nominated in answer to previous questions did not re-use ontologies in any of their ontology development activities, was it an explicit decision not to do so ?

*

- No, we just didn't re-use any.
- Yes, and please give any reasons for the decision

Q4.2 Can you provide contact details for any other person that you know of who has been involved in community-based ontology development where ontology re-use has been practised ?

*

- No
- Yes and contact details are:

Continue

Please contact kim.finney@aad.gov.au if you have any questions regarding this survey.

POWERED BY QuestionPro

Appendix 2 – Sample Screening Survey Email

Dear Colleague,

I am seeking your assistance in a research project that I am undertaking as part of my PhD studies on the topic of "Selection and evaluation of ontologies for re-use". I am also an active participant in the marine and Antarctic data management communities and therefore have a professional, as well as personal interest in the type of information that you may be able to provide. I am investigating whether the practise of re-using ontologies during community-based ontology development exercises can be facilitated, and therefore increased, if there were practical and rapid methods available for selecting and evaluating existing ontologies.

I am trying to find people who have been involved in specific types of community-based ontology building activities that would be willing to complete an on-line survey and then potentially participate in either a face-to-face interview, and/or a later survey designed to validate any findings emerging from an analysis of the interview material. You have been selected because I am personally aware that you are a recognised and respected contributor in community activities that may involve ontology re-use, or a colleague of yours has referred me to you because of your expertise.

The initial Screening Survey that is phase 1 of this research should take less than 30 minutes to complete. If you are later selected for interview (phase 2), the expectation is for the interview to take between 1 and 1.5 hours in duration and I would make an appointment to meet with you at your convenience. If you reside outside of Australia I would organise either a phone hook-up or an interview by video. If you are selected for the validation component of the research (phase 3) this would involve a two-part survey, with each part possibly taking 30 mins to complete. The worst-case scenario is that you would need to be involved in all three phases, totalling a maximum of 3 hours of your time, over a period of 6 months.

I am anticipating that the available pool of people active in the area of community ontology-building, under the conditions of interest, is relatively small and therefore your agreement to participate in this study would be invaluable. I am hopeful that an outcome from this study will be a more practical method of ontology evaluation, ultimately leading to increased ontology re-use and hence an improvement in the interoperability of on-line scientific systems.

The initial Screening Survey can be found on-line at
http://www.questionpro.com/akira/TakeSurvey?id=957173.

I have attached an information sheet of the research aims and the general methodology that will be followed for further clarification. I have also attached a Consent Form, which if you choose to participate in the study must be completed and mailed or emailed back to me. The deadline for survey submission is XX Month 2008.

Thank-you in advance for taking the time to read this email and for considering taking part.

Yours sincerely, Kim Finney PhD student (University of Tasmania)

Appendix 3 - Participant Information Sheet & Consent Form

Practical Ontology Selection and Evaluation Methods Study

Invitation

You are invited to participate in a research study that I am undertaking on the topic of "Selection and evaluation of ontologies for re-use". I am investigating whether the practise of re-using ontologies during community-based ontology development exercises can be facilitated, and therefore increased, if there were practical and rapid methods available for selecting and evaluating existing ontologies.

The study is being conducted by me (Kim Finney) in my capacity as a part-time PhD candidate, enrolled at the University of Tasmania. I am also Manager of the Australian Antarctic Data Centre, located in Kingston, Tasmania.

1. 'What is the purpose of this study?'

From a scientific community perspective one of the main impediments to realising the Semantic Web vision (as espoused by Tim Berners-Lee) is that most scientific data, even those data deployed on the web, are not generally expressed or encoded in an unambiguously defined, machine-interpretable manner. To achieve semantic interoperability scientific communities must deploy their data using ontologies. Building an ontology is a highly resource-intensive task so many knowledge and web engineers extol the virtues of re-using existing ontologies, or their ontological components. However, cost-effective re-use of ontologies necessarily implies that there also exists efficient and well-developed methods for selecting and evaluating those ontologies that are candidates for re-use. Unfortunately there isn't a comprehensive and global approach to the selection and evaluation problem, despite a range of literature on the topic.

Of those selection and evaluation methods that are reported in the academic literature, many are not considered to be particularly useful or applicable in real-world ontology development scenarios. This research aims to explore whether "academic techniques" are being used in practical development activities and if they are not, what methods are being employed. Further, it is proposed that there may be a set of practical measures which can be used by different scientific communities to select and evaluate ontologies, regardless of the community domain, albeit under set conditions. Of specific interest are communities who are seeking to exchange and manipulate four-dimensional (3D, & time) scientific datasets in a web service environment. Ultimately, this research seeks to investigate if it is possible to identify a set of rapid ontology assessment techniques that can be practically applied in real-world ontology development exercises involving ontology re-use.

2. 'Why have I been invited to participate in this study?'

I am approaching people who have been involved in specific types of community-based ontology building activities that would be willing to complete an on-line survey and then potentially participate in either a face-to-face interview, and/or a later survey designed to validate any findings emerging from an analysis of the interview material. You have been selected because I am personally aware that you are a recognised and respected contributor in community activities that may involve ontology re-use, or a colleague of yours has referred me to you because of your expertise.

3. 'What does this study involve?'

Building on work already performed by Finney (2008) this study will primarily use survey and interview methodologies to investigate how ontology selection and evaluation is being conducted by different scientific communities. Ideally, an evaluation model will then be developed post discussions with the various community experts and the significance of the nominated evaluation criteria will be quantitatively assessed. The ability to postulate a rapid ontology assessment technique will be dependent upon the outcomes from interviews and analysis of the interview data.

Initially people will be asked to complete a screening survey so that potential participants can be assessed for suitability to participate in other stages of the study. This initial stage also permits classification of respondents into different community domains. If a respondent has had suitable ontology development experience, they will be asked to participate in a face-to-face interview which is designed to elicit information on their experience in selecting, evaluating and re-using ontologies in ontology development exercises. A practical ontology selection and evaluation model will then be postulated, derived from an analysis of interview data. This model will then be validated with survey participants.

Reference to paper cited above:

Finney, K.T. (2008). Deriving an Australian Marine Ontology from Existing Ontological Models: a practical evaluation. In *proceedings of the WALIS Conference, March 2008,* Perth Convention Centre, Western Australia. Retrievable from the WWW at http://www.walis.wa.gov.au/forum/peer-review.

It is important that you understand that your involvement is this study is voluntary. While I welcome your participation, I respect your right to decline. There will be no consequences to you if you decide not to participate. If you decide to discontinue participation at any time, you may do so without providing an explanation. All information will be treated in a confidential manner, and your name will not be used in any publication arising out of the research. All of the research will be kept in a locked cabinet in my office at the AAD Kingston Head Office and on a secure server on the Divisional storage area network.

4. Are there any possible benefits from participation in this study?

I am hopeful that an outcome from this study will be a more practical method of ontology selection and evaluation, ultimately leading to increased ontology re-use and hence an improvement in the interoperability of on-line scientific systems.

5. Are there any possible risks from participation in this study?

There are no specific risks anticipated with participation in this study.

6. What if I have questions about this research?

If you would like to discuss any aspect of this study please feel free to contact me on ph 03 62 32 3459 on by email at kim.finney@aad.gov.au. Once I have analysed the information I will email you a summary of my findings. You are welcome to contact me at that time to discuss any issue relating to this research study.

This study has been approved by the University of Tasmania, Social Sciences Human Research Ethics Committee. If you have concerns or complaints about the conduct of this study should contact the Executive Officer of the HREC (Tasmania) Network on (03) 6226 7479 or email human.ethics@utas.edu.au. The Executive Officer is the person nominated to receive complaints from research participants. You will need to quote [*HREC project number*].

Thank you for taking the time to consider this study.

If you wish to take part, please sign the attached consent form.

This information sheet is for you to keep.

CONSENT FORM FOR PARTICIPANTS

Institution Letterhead

CONSENT FORM

Title of Project: Practical Ontology Selection and Evaluation Methods In Re-Use Scenarios

---

1. I have read and understood the 'Information Sheet' for this project.
2. The nature and possible effects of the study have been explained to me.
3. I understand that the study involves completing a screening survey questionairre (less than 20 minutes), which may then lead to participation in a face-to-face interview (approximately 1.5 hours) and a quantitative survey (approximately 30 minutes X 2 surveys). The purpose of these interactions is to ascertain how different communities of practise select and evaluate ontologies.
4. I understand that participation involves no risks.
5. I understand that all research data will be securely stored on the Australian Antarctic Division premises for at least five years, and will be destroyed when no longer required.
6. Any questions that I have asked have been answered to my satisfaction.
7. I agree that research data gathered from me for the study may be published provided that I cannot be identified as a participant.
8. I understand that the researcher will maintain my identity as confidential and that any information I supply to the researcher will be used only for the purposes of the research.
9. I agree to participate in this investigation and understand that I may withdraw at any time without any effect, and if I so wish, may request that any data I have supplied to date be withdrawn from the research.

Name of Participant:

Signature: Date:

Statement by Investigator

☐ I have explained the project & the implications of participation in it to this volunteer and I believe that the consent is informed and that he/she understands the implications of participation

If the Investigator has not had an opportunity to talk to participants prior to them participating, the following must be ticked.

☐ The participant has received the Information Sheet where my details have been provided so participants have the opportunity to contact me prior to consenting to participate in this project.

Name of investigator ____________________

Signature of investigator ____________________ Date ____________

Appendix 4 – Follow-up Interview Questions

These interviews are centred primarily on development tasks where the interviewee has indicated that they have re-used ontologies. Questions will be framed to get the interviewee to focus on those experiences where re-use was practiced.

1.0 Opportunity For Clarification

1.1 Review interviewee's response to survey and use some-time early in the interview to clarify anything that was ambiguous or unclear from the interviewee's survey responses.

2.0 Community-centric Questions

2.1 Can you elaborate a little about the communities in which you have performed ontology development work, for example:

- How many participants would you say are in the community ?
- Is there a peak group responsible for ontology development work or is it a more diffuse sort of activity ?
- How long has the community been actively developing ontologies ?
- How does the community manage itself and its activities in general ?

2.2 If you were part of an ontology development team, how big was it and how was it specifically governed ?

3.0 Individual Experience

3.1 How would you describe your ontology experience vs your domain experience with respect to the ontology development tasks that you've mentioned ? Areas of interest are:

- Length of experience ?
- History of involvement – what roles were played ?

4.0 Ontology Development

4.1 How have you gone about selecting ontolgies for re-use ? i.e

- How did you know that they existed ?
- How did you know whether they were worth pursuing or not, what types of issues were of importance in making the selections ?

4.2 Could you provide an overview of a specific ontology development process in general terms, i.e. the typical tasks that you performed to create the ontology ?

- What use-cases or purpose was the ontology meant to address ?
- How would you classify the ontology – application or domain ?
- Did you re-use an entire ontology or a subset ?
- What were the steps involved (use-case formulation, competency questions, conceptual modelling, encoding, evaluation/trial process ?)
- How long did it take (elapsed time) and how many man-hours involved collectively if done by a team ?

4.3 What contexts was the ontology meant to operate in ?

- Which domains ?
- What type of IT environments – any using ISO TC 211 or OGC standards?

- What type of applications ?

5.0 Ontology Evaluation

5.1 Could we focus in on the evaluation process used to assess the suitability of an ontology or ontological component being considered for re-use and specifically talk about the criteria and measures that you used in the evaluation ? i.e.

- Did you use a documented method ? Was this method developed by the group or borrowed from elsewhere ? If so where ?
- Were the criteria predominantly qualitative, quantitative or a mixture of both ?
- What were the criteria and the measures ?
- Out of the criteria and measures mentioned which were given the most weight in coming to a decision about whether to include re-usable components or not ?
- Were there things that you would have liked to use by way of criteria or measures but didn't ? and why didn't you use them ?
- Were there any limitations of the evaluation method itself that you used ?
- Would you use different criteria if you were to do another re-use exercise and if so, why ?
- How long did evaluation take relative to the whole development task ?
- If any specific dimensions out of structure, function, maintenance, governance or usability are not mentioned prompt to see why they may not have been considered important.

5.2 Did you consider that the derived ontology that you developed was ultimately fit for the task ?

- If not could its failure be attributed in any way to the evaluation method used ? If yes could you elaborate.

Appendix 5 – MarineBiotaStatistics GML Object


```xml
<!-- edited using XMLSPY v6 sp2 (http://www.xmlspy.com) by Kim Finney(AADC) -->
<schema xmlns:gml="http://www.opengis.net/gml" xmlns="http://www.w3.org/2001/XMLSchema"
targetNamespace="http://www.w3.org/2001/XMLSchema" elementFormDefault="qualified"
attributeFormDefault="unqualified" version="1.0">
    <annotation>
      <documentation>Schema for CPR Biological Coverage Values</documentation>
    </annotation>
    <import namespace=http://www.opengis.net/gml
    schemaLocation="http://schemas.opengis.net/gml3.1.1/base/gml.xsd"/>
<!--======================================================================= -->
<!--==== IMOS MarineBiotaStatistic Value Type  substitutable in Value Arrays===========-->
<!--======================================================================= -->
    <element name="MarineBiotaStatistic" type="imos:MarineBiotaStatisticType"/>
    <complexType name="MarineBiotaStatisticType">
      <complexContent>
        <restriction base="gml:CompositeValueType">
          <sequence>
            <element ref="gml:metaDataProperty" minOccurs="0" maxOccurs="unbounded"/>
            <element ref="gml:description" minOccurs="0"/>
            <element ref="gml:name" minOccurs="0" maxOccurs="unbounded"/>
            <element ref="marineBiotaStatisticValueComponent" minOccurs="0" maxOccurs="unbounded"/>
            <element ref="marineBiotaStatisticValueComponents" minOccurs="0"/>
          </sequence>
        </restriction>
      </complexContent>
    </complexType>
<!--======================================================================= -->
<!--==== Now defining the type as a property=====================================-->
<!--======================================================================= -->
    <element name="marineBiotaStatisticValueComponent" type="MarineBiotaStatisticPropertyType"/>
    <complexType name="marineBiotaStatisticValuePropertyType">
      <sequence minOccurs="0">
        <element ref="_MarineBiotaStatisticValue"/>
      </sequence>
      <attributeGroup ref="gml:AssociationAttributeGroup"/>
    </complexType>
    <element name="marineBiotaStatisticValueComponents" type="MarineBiotaStatisticValueArrayPropertyType"/>
    <complexType name="MarineBiotaStatisticValueArrayPropertyType">
      <sequence>
        <element ref="_MarineBiotaStatisticValue" maxOccurs="unbounded"/>
      </sequence>
    </complexType>
<!--======================================================================= -->
<!--==== Now defining the abstract _MarineBiotaStatisticValue=====================-->
<!--======================================================================= -->
    <element name="_MarineBiotaStatisticValue" abstract="true" substitutionGroup="gml:_Value"/>
<!--===================================================================================== -->
<!--=Now must define the individual elements that can substitute for the abstract MarineBiotaStatisticValue=-->
<!--===================================================================================== -->
    <element name="TaxonConceptName" type="gml:CategoryOrNullListType"
    substitutionGroup="_MarineBiotaStatisticValue"/>
    <complexType name="TaxonConceptNameType">
      <complexContent>
        <restriction base="gml:CategoryPropertyType">
          <attribute name="codespace" type="anyURI" use="optional"/>
        </restriction>
      </complexContent>
    </complexType>
    <element name="TaxonRank" type="gml:CategoryOrNullListType" substitutionGroup="_MarineBiotaStatisticValue"/>
    <complexType name="TaxonRankType">
```

```xml
      <complexContent>
        <restriction base="gml:CategoryPropertyType">
          <attribute name="codespace" type="anyURI" use="optional"/>
        </restriction>
      </complexContent>
    </complexType>
      <element name="TaxonConceptLSID" type="gml:CategoryOrNullListType"
      substitutionGroup="_MarineBiotaStatisticValue"/>
    <complexType name="TaxonConceptLSIDType">
      <complexContent>
        <restriction base="gml:CategoryPropertyType">
          <attribute name="codespace" type="anyURI" use="optional"/>
        </restriction>
      </complexContent>
    </complexType>
      <element name="TaxonPlacementFormal" type="gml:CategoryOrNullListType"
      substitutionGroup="_MarineBiotaStatisticValue"/>
    <complexType name="TaxonPlacementFormalType">
      <complexContent>
        <restriction base="gml:CategoryPropertyType">
          <attribute name="codespace" type="anyURI" use="optional"/>
        </restriction>
      </complexContent>
    </complexType>
    <element name="TaxonCount" type="gml:CountType" substitutionGroup="_MarineBiotaStatisticValue"/>
    <complexType name="TaxonCountType">
      <complexContent>
        <restriction base="gml:CountPropertyType"/>
      </complexContent>
    </complexType>
      <element name="TaxonMaturation" type="gml:CategoryOrNullListType"
      substitutionGroup="_MarineBiotaStatisticValue"/>
    <complexType name="TaxonMaturationType">
      <complexContent>
        <restriction base="gml:CategoryPropertyType">
          <attribute name="codespace" type="anyURI" use="optional"/>
        </restriction>
      </complexContent>
    </complexType>
</schema>
```

Appendix 6 – CPR GML Instance Document including internal reference to composite phenomena (highlighted)


```xml
<?xml version="1.0" encoding="UTF-8"?>
<!--================================================================-->
<!-- Test CPR Tow Dataset From Imaginery Deployment on Aurora Australis======-->
<!--================================================================-->
<imos:Dataset xmlns:imos="http://aadc-maps.aad.gov.au/imos" xsi:schemaLocation="http://aadc-maps.aad.gov.au/imos
http://aadc-maps.aad.gov.au/imosDatasetWrapper.xsd" xmlns:csml="http://www.ndg.nerc.ac.uk/csml"
xmlns:om="http://www.opengis.net.com/om" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xmlns:xlink="http://www.w3.org/1999/xlink" xmlns:gco="http://www.isotc211.otg/2005/gco"
xmlns:gmd="http://www.isotc211.org/2005/gmd" xmlns:gml="http://www.opengis.net/gml"
xmlns:swe="http://www.opengis.net/swe" xmlns:aodcjf="http//www.aodc.gov.au/metadata"
gml:id="CPRTestDatasetNo1">
<!--============================================================================-->
<!-- Start of Observation Collection For Tow. A Tow = Continuous recording during a Voyage.====-->
<!--============================================================================-->
    <imos:imosObservationCollection gml:id="TowNo1Aurora">
      <gml:name>CPR Tow No 1 </gml:name>
      <gml:description>CPR instrument deployment from leg 2 of Aurora Australis voyage</gml:description>
      <gml:boundedBy>
          <gml:Envelope srsName="urn:ogc:def:crs:EPSG::4326">
            <gml:pos>-73.333 120.546 </gml:pos>
            <gml:pos>-53.111 122.000</gml:pos>
          </gml:Envelope>
      </gml:boundedBy>
      <om:time>
          <gml:TimePeriod gml:id="TowCollTime1">
            <gml:startPosition>2005-01-11T18:22:25.00</gml:startPosition>
            <gml:endPosition>2005-01-11T19:22:45.00</gml:endPosition>
          </gml:TimePeriod>
<!--================================================================-->
<!-- Included a linestring location to outline voyage track =====================-->
<!--================================================================-->
      <om:location gml:id="CPRVoyageTrack">
            <gml:Linestring srsName="urn:ogc:def:crs:EPSG::4326">
              <gml:posList dimension="4">53.0 10.0 53.0 11.0 53.0 11.5 53.0 12.0</gml:posList>
            </gml:Linestring>
      </om:location>
      </om:time>
<!--=====Metadata and pointer to metadata record, attached at the collection = tow level===-->
      <om:metaDataProperty>
<!--=============Dataset citation information including name of dataset creator====-->
      <aodcjf:MD_Identification>
          <aodcjf:citation>
            <aodcjf:CI_Citation>
              <aodcjf:title>Continuous Plankton Recorder Data (2005) Captured By Dr Graham Hosie</aodcjf:title>
              <aodcjf:date>12/07/05</aodcjf:date>
              <aodcjf:citedResponsibleParty>
                  <aodcjf:CI_ResponsibleParty>
                    <aodcjf:individualName>Dr Graham Hosie</aodcjf:individualName>
                    <aodcjf:organisationName>AADC</aodcjf:organisationName>
                    <aodcjf:role>
                      <aodcjf:CI_RoleCode>prinicipalInvestigator</aodcjf:CI_RoleCode>
                    </aodcjf:role>
                  </aodcjf:CI_ResponsibleParty>
              </aodcjf:citedResponsibleParty>
            </aodcjf:CI_Citation>
          </aodcjf:citation>
          <aodcjf:abstract>AADC Continuous Plankton Recorder Data</aodcjf:abstract>
      </aodcjf:MD_Identification>
<!--=============Metadata Identifier and name of metadata creator==============-->
```

```xml
      <aodcjf:MD_Metadata>
        <aodcjf:fileIdentifier xlink:href="http://www.imos.gov.au/CAASM/record#1"/>
        <aodcjf:contact>
          <aodcjf:CI_ResponsibleParty>
              <aodcjf:individualName>Dave Connell</aodcjf:individualName>
              <aodcjf:organisationName>AADC</aodcjf:organisationName>
              <aodcjf:role>
                <aodcjf:CI_RoleCode>custodian</aodcjf:CI_RoleCode>
              </aodcjf:role>
          </aodcjf:CI_ResponsibleParty>
        </aodcjf:contact>
      </aodcjf:MD_Metadata>
      <aodcjf:MD_SecurityConstraints>
        <aodcjf:classification>unclassified</aodcjf:classification>
      </aodcjf:MD_SecurityConstraints>
    </om:metaDataProperty>
<!--=================================================================-->
<!-- First Observation Member from Tow. Observation = Tow (Voyage) Segment==-->
<!--=================================================================-->
    <imos:Observation gml:id="01">
      <imos:ownedBy gml:id="proj1">IMOS Ships Of Opportunity</imos:ownedBy>
      <imos:hostedBy gml:id="plat1">
        <imos:platformType>Ship</imos:platformType>
        <imos:platformName>Aurora Australis</imos:platformName>
      </imos:hostedBy>
      <om:time>
        <gml:TimePeriod gml:id="CPRObservationTime1">
          <gml:startPosition>2005-01-11T18:22:25.00</gml:startPosition>
          <gml:endPosition>2005-01-11T18:22:45.00</gml:endPosition>
        </gml:TimePeriod>
      </om:time>
      <imos:InstrumentProcedure gml:id="01">
        <imos:instrumentType xlink:href="urn:x-imos:def:instruments:GCMD:CPR">CPR</imos:instrumentType>
        <imos:instrumentSerialNo>123456</imos:instrumentSerialNo>
        <imos:processingMethod>
          <gml:name xlink:href="urn:x-
          imos:def:processingprocedures:IMOS:CPRProcessing">CPRProcessingMethod</gml:name>
          <gml:description xlink:href="http:www.aadc.gov.au/procedures/CPR.doc"/>
        </imos:processingMethod>
      </imos:InstrumentProcedure>
      <imos:observedProperty>
        <swe:CompositePhenomenon gml:id="CPRTaxon" Dimension="6">
          <gml:identifier codespace="http://aadc-
          maps.aad.gov.au/imos/phenonemon_dictionary">imos:phenomena:MarineBiotaStatistic</gml:identifier>
          <swe:component xlink:href="imos:phenomena:taxonConceptName"/>
          <swe:component xlink:href="imos:phenomena:taxonConceptLSID"/>
          <swe:component xlink:href="imos:phenomena:taxonPlacementFormal"/>
          <swe:component xlink:href="imos:phenomena:taxonRank"/>
          <swe:component xlink:href="imos:phenomena:taxonCount"/>
          <swe:component xlink:href="imos:phenomena:taxonMaturationStage"/>
        </swe:CompositePhenomenon>
      </imos:observedProperty>
<!--=====================================================-->
<!-- Feature of Interest (CPRSegment)========================-->
<!--=====================================================-->
      <imos:featureOfInterest gml:id="CPRSegmentFeatureTestNo1">
        <om:location>
          <gml:LinesString srsName="urn:ogc:def:crs:EPSG::4326">
              <gml:posList dimension="2">-64.622 120.546 -64.623 120.545</gml:posList>
          </gml:LinesString>
        </om:location>
        <imos:locationQcFlag codeSpace="http://www.imos.gov.au/QCFlags/location">-1</imos:locationQcFlag>
        <om:time>
```

```xml
          <gml:TimePeriod gml:id="CPRSegmentTime1">
            <gml:startPosition>2005-01-11T18:22:25.00</gml:startPosition>
            <gml:endPosition>2005-01-11T18:22:45.00</gml:endPosition>
          </gml:TimePeriod>
        </om:time>
        <imos:timeQcFlag codeSpace="http://www.imos.gov.au/QCFlags/time">0</imos:timeQcFlag>
        <imos:segmentId>1</imos:segmentId>
        <imos:segmentLength uom= "http://www.imos.gov.au/UnitsDictionary.xml#nmile">5</imos:segmentLength>
        <imos:surveyName>VoyageAA010007</imos:surveyName>
        <imos:fluorometerMeasure
        uom="http://www.imos.gov.au/UnitsDictionary.xml#flur">34</imos:fluorometerMeasure>
        <imos:salinityMeasure uom="http://www.imos.gov.au/UnitsDictionary.xml#sal">36</imos:salinityMeasure>
        <imos:temperatureMeasure
        uom="http://www.imos.gov.au/UnitsDictionary.xml#degC">18</imos:temperatureMeasure>
        <imos:licorMeasure uom="http://www.imos.gov.au/UnitsDictionary.xml#lux">300</imos:licorMeasure>
        <imos:totalAbundance
        uom="http://www.imos.gov.au/UnitsDictionary.xml#individuals">150</imos:totalAbundance>
      </imos:featureOfInterest>
<!--=================================================================-->
<!-- Value (CPRMulticurveCoverageData)=================================-->
<!--=================================================================-->
      <imos:result gml:id="CPRCoverageDataTestNo1">
        <gml:multiCurveDomain>
          <gml:MultiCurve srsName="urn:EPSG:GeographicCRS:4326">
            <gml:curveMember>
              <gml:Linestring gml:id="CPRSegmentTrack">
<!--=================================================================-->
<!-- poslist can be used to define segment track in as much detail as required=====-->
<!--=================================================================-->
                <gml:posList dimension="3">-64.622 120.536 -64.644 120.546 -64.623 120.545</gml:posList>
              </gml:Linestring>
            </gml:curveMember>
          </gml:MultiCurve>
        </gml:multiCurveDomain>
        <gml:rangeSet>
<!--=====================================================================================-->
<!-- First ValueArray conisists of a composite value type holding all taxon related info for first taxon===========-->
<!--=====================================================================================-->

          <gml:ValueArray gml:id="CPRCoverageDataTaxonType1">
            <imos:marineBiotaStatisticValueComponents>
              <gml:CompositeValue>
                <imos:marineBiotaStatisticValueComponent>
                  <imos:TaxonConceptName>
                    <gml:CategoryList codeLIst=
                    "http://www.imos.gov.au/TaxonNameServer/Name/#201">foraminifera</gml:CategoryList>
                  </imos:TaxonConceptName>
                </imos:marineBiotaStatisticValueComponent>
                <imos:marineBiotaStatisticValueComponent>
                  <imos:TaxonConceptLSID>
                    <gml:CategoryList codeLIst=
                    "http://www.imos.gov.au/TaxonNameServer/LSID/#201">LSID4567</gml:CategoryList>
                  </imos:TaxonConceptLSID>
                </imos:marineBiotaStatisticValueComponent>
                <imos:marineBiotaStatisticValueComponent>
                  <imos:TaxonRank>
                    <gml:CategoryList codeLIst=
                    "http://www.imos.gov.au/TaxonNameServer/Rank/#201">class</gml:CategoryList>
                  </imos:TaxonRank>
                </imos:marineBiotaStatisticValueComponent>
                <imos:marineBiotaStatisticValueComponent>
                  <imos:TaxonPlacementFormal>
```

```xml
                <gml:CategoryList codeLIst= "http://www.imos.gov.au/TaxonNameServer/#201">Eucaryota,
                Protozoa, Biciliata, Rhizarai, Foraminfera</gml:CategoryList>
              </imos:TaxonPlacementFormal>
            </imos:marineBiotaStatisticValueComponent>
            <imos:marineBiotaStatisticValueComponent>
              <imos:TaxonCount>
                <gml:Count>123</gml:Count>
              </imos:TaxonCount>
            </imos:marineBiotaStatisticValueComponent>
            <imos:marineBiotaStatisticValueComponent>
              <imos:TaxonMaturationStage>
                <gml:CategoryList
                codeLIst="http://www.imos.gov.au/LifeStageDictionary/#l40">juveniles</gml:CategoryList>
              </imos:TaxonMaturationStage>
            </imos:marineBiotaStatisticValueComponent>
          </gml:CompositeValue>
        </imos:marineBiotaStatisticValueComponents>
      </gml:ValueArray>
<!--=========================================================================-->
<!-- The 2nd ValueArray containing all data for second taxon===========================-->
<!--=========================================================================-->
      <gml:ValueArray gml:id="CPRCoverageDataTaxonType2">
        <imos:marineBiotaStatisticValueComponents>
          <gml:CompositeValue>
            <imos:marineBiotaStatisticValueComponent>
              <imos:TaxonConceptName>
                <gml:CategoryList
                codeLIst="http://www.imos.gov.au/TaxonNameServer/Name/#401">Allogromiida</gml:Categor
                yList>
              </imos:TaxonConceptName>
            </imos:marineBiotaStatisticValueComponent>
            <imos:marineBiotaStatisticValueComponent>
              <imos:TaxonConceptLSID>
                <gml:CategoryList codeLIst=
                "http://www.imos.gov.au/TaxonNameServer/LSID/#401">LSID4569</gml:CategoryList>
              </imos:TaxonConceptLSID>
            </imos:marineBiotaStatisticValueComponent>
            <imos:marineBiotaStatisticValueComponent>
              <imos:TaxonRank>
                <gml:CategoryList codeLIst=
                "http://www.imos.gov.au/TaxonNameServer/Rank/#401">order</gml:CategoryList>
              </imos:TaxonRank>
            </imos:marineBiotaStatisticValueComponent>
            <imos:marineBiotaStatisticValueComponent>
              <imos:TaxonPlacementFormal>
                <gml:CategoryList codeLIst="http://www.imos.gov.au/TaxonNameServer/#401">Eucaryota,
                Protozoa, Biciliata, Rhizarai, Foraminfera, Allogromiida</gml:CategoryList>
              </imos:TaxonPlacementFormal>
            </imos:marineBiotaStatisticValueComponent>
            <imos:marineBiotaStatisticValueComponent>
              <imos:TaxonCount>
                <gml:Count>12000</gml:Count>
              </imos:TaxonCount>
            </imos:marineBiotaStatisticValueComponent>
            <imos:marineBiotaStatisticValueComponent>
              <imos:TaxonMaturationStage>
                <gml:CategoryList
                codeLIst="http://www.imos.gov.au/LifeStageDictionary/#l50">adults</gml:CategoryList>
              </imos:TaxonMaturationStage>
            </imos:marineBiotaStatisticValueComponent>
          </gml:CompositeValue>
        </imos:marineBiotaStatisticValueComponents>
      </gml:ValueArray>
```

```xml
<!--=========================================================================-->
<!-- The 3rd ValueArray consisting of all data for 3rd taxon============================-->
<!--=========================================================================-->
              <gml:ValueArray gml:id="CPRCoverageDataTaxonType3">
                 <imos:marineBiotaStatisticValueComponents>
                   <gml:CompositeValue>
                     <imos:marineBiotaStatisticValueComponent>
                         <imos:TaxonConceptName>
                           <gml:CategoryList codeLIst=
                           "http://www.imos.gov.au/TaxonNameServer/Name/#301">Lagynidae</gml:CategoryList>
                         </imos:TaxonConceptName>
                     </imos:marineBiotaStatisticValueComponent>
                     <imos:marineBiotaStatisticValueComponent>
                         <imos:TaxonConceptLSID>
                           <gml:CategoryList codeLIst=
                           "http://www.imos.gov.au/TaxonNameServer/LSID/#301">LSID4578</gml:CategoryList>
                         </imos:TaxonConceptLSID>
                     </imos:marineBiotaStatisticValueComponent>
                     <imos:marineBiotaStatisticValueComponent>
                         <imos:TaxonRank>
                           <gml:CategoryList codeLIst=
                           "http://www.imos.gov.au/TaxonNameServer/Rank/#301">family</gml:CategoryList>
                         </imos:TaxonRank>
                     </imos:marineBiotaStatisticValueComponent>
                     <imos:marineBiotaStatisticValueComponent>
                         <imos:TaxonPlacementFormal>
                           <gml:CategoryList codeLIst= "http://www.imos.gov.au/TaxonNameServer/#301">Eucaryota,
                           Protozoa, Biciliata, Rhizarai, Foraminfera, Allogromiida,Lagynidae</gml:CategoryList>
                         </imos:TaxonPlacementFormal>
                     </imos:marineBiotaStatisticValueComponent>
                     <imos:marineBiotaStatisticValueComponent>
                         <imos:TaxonCount>
                           <gml:Count>10</gml:Count>
                         </imos:TaxonCount>
                     </imos:marineBiotaStatisticValueComponent>
                     <imos:marineBiotaStatisticValueComponent>
                         <imos:TaxonMaturationStage>
                           <gml:CategoryList codeLIst="http://www.imos.gov.au/LifeStageDictionary/#l60">adults and
                           juveniles</gml:CategoryList>
                         </imos:TaxonMaturationStage>
                     </imos:marineBiotaStatisticValueComponent>
                   </gml:CompositeValue>
                 </imos:marineBiotaStatisticValueComponents>
              </gml:ValueArray>
            </gml:rangeSet>
          </imos:result>
      </imos:Observation>
<!--===========================================================================================-->
<!-- 2nd Observation Member from Tow. Observation = Tow (Voyage) Segment 2 - pattern now repeats==-->
<!--===========================================================================================-->
      <imos:Observation gml:id="02">
          <imos:ownedBy gml:id="proj1">IMOS Ships Of Opportunity</imos:ownedBy>
          <imos:hostedBy gml:id="plat1">
            <imos:platformType>Ship</imos:platformType>
            <imos:platformName>Aurora Australis</imos:platformName>
          </imos:hostedBy>
          <om:time>
            <gml:TimePeriod gml:id="CPRTime2">
              <gml:startPosition>2005-01-11T18:22:25.00</gml:startPosition>
              <gml:endPosition>2005-01-11T18:22:45.00</gml:endPosition>
            </gml:TimePeriod>
          </om:time>
          <imos:InstrumentProcedure gml:id="02">
```

```xml
        <imos:instrumentType xlink:href="urn:x-imos:def:instruments:GCMD:CPR">CPR</imos:instrumentType>
        <imos:instrumentSerialNo>123456</imos:instrumentSerialNo>
        <imos:processingMethod>
          <gml:name xlink:href="urn:x-
          imos:def:processingprocedures:IMOS:CPRProcessing">CPRProcessingMethod</gml:name>
          <gml:description xlink:href="http:www.aadc.gov.au/procedures/CPR.doc"/>
        </imos:processingMethod>
      </imos:InstrumentProcedure>
      <imos:observedProperty>
        <swe:CompositePhenomenon gml:id="CPRTaxon" Dimension="6">
          <gml:identifier codespace="http://aadc-
          maps.aad.gov.au/imos/phenonemon_dictionary">imos:phenomena:MarineBiotaStatistic</gml:identifier>
          <swe:component xlink:href="imos:phenomena:taxonConceptName"/>
          <swe:component xlink:href="imos:phenomena:taxonConceptLSID"/>
          <swe:component xlink:href="imos:phenomena:taxonPlacementFormal"/>
          <swe:component xlink:href="imos:phenomena:taxonRank"/>
          <swe:component xlink:href="imos:phenomena:taxonCount"/>
          <swe:component xlink:href="imos:phenomena:taxonMaturationStage"/>
        </swe:CompositePhenomenon>
      </imos:observedProperty>
<!--===================================================================-->
<!--============= Feature of Interest (CPRSegment)========================-->
<!--===================================================================-->
      <imos:featureOfInterest gml:id="CPRSegmentFeatureTestNo2">
        <om:location>
          <gml:LinesString srsName="urn:ogc:def:crs:EPSG::4326">
              <gml:posList dimension="2">-64.622 120.546 -64.623 120.545</gml:posList>
          </gml:LinesString>
        </om:location>
        <imos:locationQcFlag codeSpace="http://www.imos.gov.au/QCFlags/location">-1</imos:locationQcFlag>
        <om:time>
          <gml:TimePeriod gml:id="CPRTime4">
              <gml:startPosition>2005-01-11T18:22:25.00</gml:startPosition>
              <gml:endPosition>2005-01-11T18:22:45.00</gml:endPosition>
          </gml:TimePeriod>
        </om:time>
        <imos:timeQcFlag codeSpace="http://www.imos.gov.au/QCFlags/time">0</imos:timeQcFlag>
        <imos:segmentId>1</imos:segmentId>
        <imos:segmentLength uom= "http://www.imos.gov.au/UnitsDictionary.xml#nmile">5</imos:segmentLength>
        <imos:surveyName>VoyageAA010007</imos:surveyName>
        <imos:fluorometerMeasure
        uom="http://www.imos.gov.au/UnitsDictionary.xml#flur">44</imos:fluorometerMeasure>
        <imos:salinityMeasure uom="http://www.imos.gov.au/UnitsDictionary.xml#sal">21</imos:salinityMeasure>
        <imos:temperatureMeasure
        uom="http://www.imos.gov.au/UnitsDictionary.xml#degC">17</imos:temperatureMeasure>
        <imos:licorMeasure uom="http://www.imos.gov.au/UnitsDictionary.xml#lux">302</imos:licorMeasure>
        <imos:totalAbundance
        uom="http://www.imos.gov.au/UnitsDictionary.xml#individuals">1230</imos:totalAbundance>
      </imos:featureOfInterest>
<!--===================================================================-->
<!-- The 1st ValueArray for 1st Taxon sampled on second segment=============-->
<!--===================================================================-->
      <imos:result gml:id="CPRCoverageDataTestNo2">
        <gml:multiCurveDomain>
          <gml:MultiCurve srsName="urn:EPSG:GeographicCRS:4326">
              <gml:curveMember>
                <gml:Linestring>
                  <gml:posList dimension="3">-64.622 120.546 -64.644 120.566 -64.623 120.545</gml:posList>
                </gml:Linestring>
              </gml:curveMember>
          </gml:MultiCurve>
        </gml:multiCurveDomain>
        <gml:rangeSet>
```

```xml
            <gml:ValueArray gml:id="CPRCoverageDataTaxonSeg2Type1">
                <imos:marineBiotaStatisticValueComponents>
                  <gml:CompositeValue>
                    <imos:marineBiotaStatisticValueComponent>
                        <imos:TaxonConceptName>
                          <gml:CategoryList codeLIst=
                          "http://www.imos.gov.au/TaxonNameServer/Name/#201">foraminifera</gml:CategoryList>
                        </imos:TaxonConceptName>
                    </imos:marineBiotaStatisticValueComponent>
                    <imos:marineBiotaStatisticValueComponent>
                        <imos:TaxonConceptLSID>
                          <gml:CategoryList       codeLIst=
                          "http://www.imos.gov.au/TaxonNameServer/LSID/#201">LSID4567</gml:CategoryList>
                        </imos:TaxonConceptLSID>
                    </imos:marineBiotaStatisticValueComponent>
                    <imos:marineBiotaStatisticValueComponent>
                        <imos:TaxonRank>
                          <gml:CategoryList codeLIst=
                          "http://www.imos.gov.au/TaxonNameServer/Rank/#201">class</gml:CategoryList>
                        </imos:TaxonRank>
                    </imos:marineBiotaStatisticValueComponent>
                    <imos:marineBiotaStatisticValueComponent>
                        <imos:TaxonPlacementFormal>
                          <gml:CategoryList codeLIst= "http://www.imos.gov.au/TaxonNameServer/#201">Eucaryota,
                          Protozoa, Biciliata, Rhizarai, Foraminfera</gml:CategoryList>
                        </imos:TaxonPlacementFormal>
                    </imos:marineBiotaStatisticValueComponent>
                    <imos:marineBiotaStatisticValueComponent>
                        <imos:TaxonCount>
                          <gml:Count>500</gml:Count>
                        </imos:TaxonCount>
                    </imos:marineBiotaStatisticValueComponent>
                    <imos:marineBiotaStatisticValueComponent>
                        <imos:TaxonMaturationStage>
                          <gml:CategoryList
                          codeLIst="http://www.imos.gov.au/LifeStageDictionary/#l40">juveniles</gml:CategoryList>
                        </imos:TaxonMaturationStage>
                    </imos:marineBiotaStatisticValueComponent>
                  </gml:CompositeValue>
                </imos:marineBiotaStatisticValueComponents>
            </gml:ValueArray>
<!--===================================================================================-->
<!-- The 2nd ValueArray containing all data for second taxon sampled on second segment===========-->
<!--===================================================================================-->
            <gml:ValueArray gml:id="CPRCoverageDataTaxonType2">
                <imos:marineBiotaStatisticValueComponents>
                  <gml:CompositeValue>
                    <imos:marineBiotaStatisticValueComponent>
                        <imos:TaxonConceptName>
                          <gml:CategoryList codeLIst=
                          "http://www.imos.gov.au/TaxonNameServer/Name/#401">Allogromiida</gml:CategoryList>
                        </imos:TaxonConceptName>
                    </imos:marineBiotaStatisticValueComponent>
                    <imos:marineBiotaStatisticValueComponent>
                        <imos:TaxonConceptLSID>
                          <gml:CategoryList codeLIst=
                          "http://www.imos.gov.au/TaxonNameServer/LSID/#401">LSID4569</gml:CategoryList>
                        </imos:TaxonConceptLSID>
                    </imos:marineBiotaStatisticValueComponent>
                    <imos:marineBiotaStatisticValueComponent>
                        <imos:TaxonRank>
                          <gml:CategoryList codeLIst=
                          "http://www.imos.gov.au/TaxonNameServer/Rank/#401">order</gml:CategoryList>
```

```xml
                    </imos:TaxonRank>
                  </imos:marineBiotaStatisticValueComponent>
                  <imos:marineBiotaStatisticValueComponent>
                    <imos:TaxonPlacementFormal>
                      <gml:CategoryList codeLIst= "http://www.imos.gov.au/TaxonNameServer/#401">Eucaryota,
                      Protozoa, Biciliata, Rhizarai, Foraminfera, Allogromiida</gml:CategoryList>
                    </imos:TaxonPlacementFormal>
                  </imos:marineBiotaStatisticValueComponent>
                  <imos:marineBiotaStatisticValueComponent>
                    <imos:TaxonCount>
                      <gml:Count>5</gml:Count>
                    </imos:TaxonCount>
                  </imos:marineBiotaStatisticValueComponent>
                  <imos:marineBiotaStatisticValueComponent>
                    <imos:TaxonMaturationStage>
                      <gml:CategoryList
                      codeLIst="http://www.imos.gov.au/LifeStageDictionary/#l50">adults</gml:CategoryList>
                    </imos:TaxonMaturationStage>
                  </imos:marineBiotaStatisticValueComponent>
                </gml:CompositeValue>
              </imos:marineBiotaStatisticValueComponents>
            </gml:ValueArray>
          </gml:rangeSet>
        </imos:result>
      </imos:Observation>
    </imos:imosObservationCollection>
</imos:Dataset>
```

Appendix 7 – Characteristics Of Biological Data

Focus of Sampling	Typical Sampling Method	Spatio-Temporal Sampling Dimension	Proxy For	Minimum Value Set Members	Examples
Single taxa	In-situ extractive sampling of host medium, take/catch specimen, photograph	1D Location (point) @ fixed time (t)	taxa population in that region	Location (L) @ t, Taxa name, measured parameter(s)	Trapping an animal at a particular location and time and recording details (taxa, age, sex) about the animal. Taking a soil sample at a particular location and time and analysing it for the presence of a particular type of taxa, then recording how many taxa were present. A taxonomic specimen collected at a location at a specific time is dissected and various parameters are recorded (e.g. weight, length, blood type, age, sex).
Single taxa	Single quadrat, transect, multiple quadrats in gridded sampling (random, non-random, stratified), photograph, satellite image,	2D Location (polygon, grid or line) @ fixed time (t)	taxa population in that region	Location (L), time (t), Taxa name, measured parameter(s), [possibly repeated for different locations (Ln) at a fixed time (t) i.e. (Ln @ t)]	Sampling along a transect using an aerial photograph. Discrete sample points are identified along the transect at which an estimation of the % cover of a taxa is made. Even though locations vary, the time and taxa are constant. A grid is placed over a satellite image and for each grid within the scene % cover of a taxa is estimated. A processed satellite image (e.g. Landsat) where each pixel in an image has a value which can be correlated with the presence or absence of a particular species.
Single taxa	Single quadrat, transect, multiple quadrats in gridded sampling	2D Location (polygon, grid or line) @ varying time	taxa population in that region	Location (L), time (t), Taxa name, measured	A long transect is established and a person traverses the transect and at intervals stops and checks for the presence or absence of a

	(random, non-random, stratified), photograph, satellite image,	(t1..n).		parameter(s), [possibly repeated for L2 @ t2, Ln @ tn]	particular taxa. Each sampling station is sampled at a different time. So location and time vary but taxa remains constant.
Single taxa	Tagged	2 or 3D Location (line with height or depth) @ varying time (t1..n).	taxa population in that region OR Physical environment of region	Location (L), time (t), Taxa name, measured parameter(s), [repeated for different locations (Ln) at different times (tn) i.e. Ln@tn]	An animal fitted with a tracking device and other sensors will have its location tracked over time as well as having other environmental parameters sampled. So taxa is constant but location and time varies.
Multiple taxa	In-situ extractive sampling of host medium, take/catch specimens, photograph	1D Location (point) @ fixed time (t)	Community assemblage(s) in that region	Location (L) @ t, Taxa name, measured parameter(s) [repeated for each taxa 1..n]	Surface water sample at a particular location and time which is subsequently analysed to determine all plankton species present in the sample along with an estimate of the numbers of each type identified.
Multiple taxa	Single quadrat, transect, multiple quadrats in gridded sampling (random, non-random, stratified), photograph, satellite image, integrating towed devices e.g. fishing net.	2 or 3D Location (polygon, grid or line) @ fixed time (t)	Community assemblage(s) in that region	Location (L), time (t), Taxa name, measured parameter(s), [possibly repeated for different locations (Ln) at a fixed time (t) for each taxa i.e. Ln@t]. Depth or height may be important domain parameters.	Satellite image of a particular habitat overlain with a grid, where for each grid cell, individual vegetation types are identified and an estimate of % cover is made. Type of taxa and location vary but time is constant. Still image from an underwater camera is used to identify the various taxa identifiable on the seafloor. Location and time are constant but taxa varies. Depth may be an important domain parameter (along with location) which would classify the observation as being spatially 3D. A fishing net opens at time t and closes at time n. The net is emptied and fish are sorted, weighed and sexed. The location is a ship's track (possibly with depth) and the time is an

					interval. All identified taxa are linked to the track and the interval. So taxa varies but all taxa are assigned to a track and a time interval. A processed satellite image (e.g. Landsat) where each pixel in an image has a value which can be correlated with the presence or multiple species.
Multiple taxa	Single quadrat, transect, multiple quadrats in gridded sampling (random, non-random, stratified), photograph, satellite image, discrete sampling towed devices e.g. CPR recorder.	2 or 3D Location (polygon, grid or line) @ varying time (t1..n).	Community assemblage(s) in that region	Location (L), time (t), Taxa name, measured parameter(s), [possibly repeated for different locations (Ln) at different times (tn) i.e. Ln@tn for each taxa] Depth or height may be important domain parameters.	A continuous plankton recorder is towed behind a ship. At the end of the voyage the silk onto which plankton has stuck is cut up into discrete samples. Each sample is given a 3D location and time along the ship's track. The silk segment is the analysed for different taxa and the counts for each taxa. Location, time and taxa varies. Videos taken a long a ship's track on the seafloor which is subsequently analysed frame by frame for taxa present. Location, time and taxa varies.

Appendix 8 – Feature Catalogue Use-Case Vignettes

The reader should note that the words 'Feature Type' and 'attribute' are used loosely within the use-cases (since they are, to a large degree framed consistent with how the community has communicated their needs to the author, which at times could include ambiguities). It is by examining the various ways in which a Catalogue user wants to be able to interact with the system that gives additional clues as to what elements should constitute a Feature Type vs what should be considered a Feature Type's attribute. The following sections present the use-case vignettes for each class of interaction.

Discovery and Search Scenarios

Use Case 1
A marine (program) community member queries a Feature Catalogue, through a web client and elects to see all features listed in the catalogue. She then elects to see only those Feature Types that have been nominated for use by the IMOS community. She wants to browse through the attribute information for each of the retrieved Features Types.

Use Case 2
A community member wants to search the catalogue for Feature Type information through a web client and rather than browse via Feature Types belonging to her community of interest, or by simple alphabetical Feature Type listings, she wants to browse via a classification scheme that differentiates Feature Types on the basis of their sampling method.

Use Case 3
A community member is unsure of the provenance of several Feature Types in use by her community. She uses a Feature Catalogue web client to search for the Feature Types in question and accesses maintenance information about the Feature Types such as change history notes, who requested the feature be created, who created the feature, when it was created and the source of the definition.

Use Case 4
A community member wants to search the catalogue for Feature Type information through a web client. Once she has found a Feature Type that she is interested in she wants to see what its lexical definition is, whether it is a synonym for any other features (inside the catalogue) or if it maps to any other concepts outside of the catalogue (in use by other communities of interest). If there are such relationships with external concepts that are online, she wants to be able to navigate directly to those resources to inspect them.

Use Case 5
A SCAR community member wishes to examine all Feature Types that are considered to be components of other features (e.g. they have a partOf or componentOf type transitive relationship). The catalogue is queried to display all features that are deemed to have other features as their components.

Use Case 6
The catalogue administrator is interested in which features are most frequently accessed by catalogue users. She uses the catalogue's administration console to evaluate routinely captured statistics on catalogue operations.

Use Case 7
A SCAR community member is interested in using a particular Feature Type defined by marine (AODN) community colleagues. They use the catalogue to request information on a particular Feature Type and inspect all of the attributes of that feature. They notice that one of the attributes draws its permitted values from a list of terms. They are curious about the types of terms in this list and so request to see the available options in the list. They are satisfied that it covers their domain of interest and so they decide to use the feature as it is.

Use Case 8
A community member is interested in a specific observation Feature Type since they wish to build a data service delivering that particular type of observation, conformant with the agreed community description for such Feature Types. They query the catalogue to retrieve the description of the (observation) Feature Type. From that description they are able to link to and then inspect the observation's feature-of-interest and the associated observed properties.

Use Case 9
An AODN community member is building a web client to provide access to a wide variety of datasets that it will source from remote data servers within their community. They want to provide an intuitive, multi-faceted discovery interface and so query the Feature Catalogue first to see a description of all of the features available in the catalogue and then descriptions of all of the properties (attributes) that are listed in the catalogue. Since measured properties are further described by their estimation method and if appropriate by the component used to make the measurement, they are able to build their discovery interface using the facets of Feature Type, Observed Property and Sampling Method (obtained from the Catalogue).

Use Case 10
An AODN community member is building a data service and wants to make sure that they assign a symbol to a Feature Type that is consistent with how other community members will be depicting their Feature Types through portrayal services. They query the Catalogue for the Feature Type they are interested in and there is a textual description of the symbol noting any constraints in use and a URL to the graphic depiction of the symbol and a URL to a symbology catalogue entry.

Use Case 11
A community member is building a data service and wants to make sure that they use the correct Feature Type definitions within their service, consistent with community practise. They access the specific Feature Type they are interested in and then link out to one of the sample encoding schemas in which the Feature Type participates.

Add and Moderation Scenarios

Use Case 12
A marine community data service provider wants to edit attribute details for a given Feature Type in the catalogue because informal community discussions have indicated that the feature is no longer described appropriately. She logs into the catalogue and edits the feature in question. The catalogue automatically emails nominated community contacts and notifies them of the intended changes. The original feature is unaltered and still in operation, whilst the community decides whether or not to accept the changes. Once the community decides that the changes to the properties are to be committed, the changed feature is instantiated by a catalogue administrator and a change note is added to the Feature Type.

Use Case 13
A marine community member wants to set up a community profile for Feature Types. The new community project starting up wants a way of defining, and standardising on the features that will be used in exchanged datasets. She uses a catalogue web client to either define or select elements that constitute Feature Type definitions for the use-case in question. The profile is submitted for moderation by the community and after approval is given, the catalogue administrator commits the new profile and any new constituent elements for public use within the catalogue.

Use Case 14
A marine community member wants to add a new Feature Type classification scheme to the catalogue. The community member logs into the catalogue and elects to add a new scheme. When the community member is satisfied that they have created the requisite scheme they submit the new scheme for community consideration. After the catalogue performs a number of simple integrity and validation checks an email is automatically sent to nominated community members notifying them that a new scheme has been submitted for their consideration. When all nominated members have reviewed the scheme details and are satisfied that it should be a recognised community scheme for use in a community standard profile, the catalogue administrator commits the scheme to the catalogue. Until the scheme is committed it is considered as being "under review" and is only accessible using a subset of the catalogue's functionality.

Use Case 15
A marine community moderator wants to change a feature attribute of a Feature Type from being a mandatory attribute to one which is optional. She logs into the system, changes the designation of the attribute and submits the modification to the community for review.

Download Scenarios

Use Case 16
A metadata service provider decides to use the Feature Catalogue as a source of Feature Type information to help its users mark-up metadata records with Feature Type information in order to start standardising on the way features are described in metadata. The catalogue has an accessible service end-point that allows the metadata provider to dynamically link to the catalogue from its metadata application to extract Feature Type information. Alternatively,

periodic extraction of the desired information can be undertaken in bulk by the metadata application administrator for local caching.

Use Case 17
A marine community data service provider is going to build data services and wants to make sure that they will be conformant with how the community has agreed to exchange certain dataset Feature Types. She issues a service call to the catalogue to download a description of the Feature Types and properties used by her community. Armed with this description and aware of existing community application schema that the Feature Type participates in she builds an appropriate data delivery service.

Use Case 18
A marine community data service provider has built a data service and wants to link dynamically to a feature type definition for a specific feature type, using the service she has created. She uses a catalogue service that accesses a feature’s definition details. She notices that there are several available formats in which to receive the Feature Type details. She chooses the one that is most suitable for resolution from within her service.

Use Case 19
A SCAR community data service provider has built a data service and wants to link (via an xlink in the data service’s output) to a Feature Type definition. She includes links to the catalogue service that accesses a feature’s definition details from appropriate points in her data delivery schema.

Use Case 20
A sophisticated marine community data service user is interested in accessing some very specific data services listed in a service registry. Since this particular registry does not have the type of classification schemes that would assist our sophisticated data user, they use the community Feature Catalogue to help them shape their service registry query. To populate their web service query (targeting the service registry) they use a Feature Catalogue service that identifies all of the Feature Types used by their community that have a specific property (e.g. Sea Surface Temperature). The results returned from the Feature Catalogue, which lists all Feature Types matching this criterion, are then used to construct a registry query in which all of these Feature Types are enumerated.

Use Case 21
A sophisticated marine community data service user and utility provider is interested in developing some useful software tools for community colleagues. They want to build some automated (on the fly) data integration software that facilitates aggregation of datasets. This software will retrieve and then combine datasets delivered by data services based on whether their Feature Types have sufficiently common attributes types, units of measure and estimation methods. To determine dataset compatibility, the data integration software uses a Feature Catalogue service to download Feature Type descriptions and parses this information using a rule-base to produce a list of candidate Feature Types that can be used for data acquisition and aggregation purposes.

Management and Administration Scenarios

Use Case 22
At specified periods the catalogue administrator (and host) determines that it is time to do a bulk update of catalogue contents. A new catalogue version is released. All features that are still in use are transferred to the new catalogue version. Some features are deprecated and they are not transferred to the new catalogue but remain visible as part of the old catalogue (version). At a suitable time in the future the old catalogue version is removed from public access, but for a period it remains accessible in tandem with the new one.

Use Case 23
At specified periods the catalogue administrator (and host) determines that it is time to update Catalogue functionality. The user community is consulted via an online discussion forum about changes that should be made. All software change requests and completed updates are publicly logged.

Use Case 24
Community moderators (elected by the community) and community members interested in being involved in moderation processes are listed on an administered register, which is used by the Catalogue system to automatically contact community members when moderation processes are required.

Appendix 9 – Encapsulating Observation Features Example

Mapping '*SpecimenObservation*' Into A Catalogue Complex Feature Type.

Feature Type:{
  id:*001*
  name:*SpecimenObservation*
  definition: *An observation about a biological specimen*
  definitionSource: *AODC JF Technical Committee*
    SpatioTemporalGeometry:{
      id:*0001*
      nameValue:*1D-Point (@fixed t)*
      description: *A single cartesian coordinate at a single time(t)*
      shape: *GM_Point*
      definitionSource: *AODC JF Technical Committee* }
    Attribute:{
      id:*007*
      name: *observerName*
      definition: *Name of observer*
      definitionSource: *AODC JF Technical Committee* }
    hasRelationshipFOI: {*BiologicalTaxonSamplingFeature}*
    hasRelationshipSampledFeature: {*FishPopulationFeature}*
    ownedBy: {BiologicalProject}
    hostedBy: {Ship}
    inTheme: {*Taxa Observations}}*
Feature Type:{
  id:*099*
  name: *Biological Project*
  definition: *A project involving the study of biological entities.*
  definitionSource: *AODC JF Technical Committee*
    Attribute: {
      id:*017*
      name: *projectName*
      definition: *Project or activity identifier name/code*
      definitionSource: *AODC JF Technical Committee* }
    Attribute: {
      id:*018*
      name: *projectPeriod*
      scale: *Ratio*
      definition: *Period of time covered by project (start date to end date).*
      definitionSource: *AODC JF Technical Committee* }
      characterisedByUoM:{
        id:*028*
        name: *Date*
        definition: *A calendar date is fully specified by the year (numbered by some scheme beyond the scope of the calendar itself), the month (identified by name or number), and the day of the month (numbered sequentially starting at 1).*
        definitionSource: *AODC technical Committee*
        symbol: *dd/mm/yy-dd/mm/yy*

characterisedByDatum:{
id: *077*
name: *GregorianCalendar*
definition: ......
definitionSource: *Wikipedia*}}
Owns: {*SpecimenObservation*}}
inTheme: {*Project Categories}*
Feature Type:{
id:*033*
name: *Ship*
definition: *A type of floating marine platform.*
definitionSource: *AODC JF Technical Committee*
Attribute: {
id:*022*
name: *platformName*
definition: *Platform identifier code*
definitionSource: *AODC JF Technical Committee* }
Attribute: {
id:*024*
name: *platformType*
definition: *The type of Platform (e.g. ship, satellite)*
definitionSource: *AODC JF Technical Committee* }
hosts: {*SpecimenObservation*}}
inTheme: {*Observation Platforms}*
Feature Type:{
id:*004*
name: *BiologicalTaxonSamplingFeature*
definition: *A sampling feature for capturing biological taxon information and associated properties*
definitionSource: *AODC JF Technical Committee*
SpatioTemporalGeometry:{
id:*0001*
name: *1D-Point (@fixed t)*
description: *A single cartesian coordinate at a single time(t)*
shape: *GM_Point}*
Attribute:{
id:*036*
name: *Latitude*
scale: *Ratio*
definition: *The angular distance of a location south or north of the equator.*
definitionSource: *AODC JF Technical Committee*
characterisedByUoM:{
id:111
name: degrees, minutes and seconds of latitude
definition: *some text*
definitionSource: *Wikipedia (http://en.wikipedia.org/wiki/Latitude)*
symbol: ^{0}*:’:”*
characterisedByDatum:{
id:*010*

name: *World Geodetic System of 1984 (WGS84)*
definition: *Standard coordinate frame for earth*
definitionSource: *Wikipedia*
http://en.wikipedia.org/wiki/World_Geodetic_System }}}

Attribute:{
id:*037*
name: *Longitude*
scale: *Ratio*
definition: *The east-west position of a point on the Earth's surface.*
definitionSource: *AODC JF Technical Committee*
characterisedByUoM:{
id:*112*
name: *degrees, minutes and seconds of longitude*
definition: *some text*
definitionSource: *Wikipedia (http://en.wikipedia.org/wiki/Longitude)*
symbol: *^{0}:':"*
characterisedByDatum:{
id:*010*
name: *World Geodetic System of 1984 (WGS84)*
definition: *Standard coordinate frame for earth*
definitionSource: *Wikipedia*
http://en.wikipedia.org/wiki/World_Geodetic_System }}}

Attribute:{
id:*012*
name: *taxonConceptName*
scale: *Nominal*
definition: *is a named classification unit (or taxon)*
definitionSource: *TDWG Taxon Concept (LSID) Ontology*
onlineOntologyFlag: *True*
OntologyURL: *http://rs.tdwg.org/ontology/voc/TaxonConcept.rdf*}

Attribute: {
id:*033*
name: *samplingTime*
scale: *Ratio*
definition: *Time when sample was collected*
definitionSource: *AODC JF Technical Committee*
characterisedByUoM:{
id:*009*
name: *24 hour clock*
definition: *convention of time keeping in which the day runs from midnight to midnight and is divided into 24 hours*
definitionSource: *Wikipedia*
symbol: *hr:min:seconds*
characterisedByDatum:{
id: *015*
name: *Coordinated Universal Time (UTC)*

definition: *Standard based on international Atomic Time*
definitionSource: *Wikipedia*}}}
Attribute: {
id:*044*
name: *weight*
scale: *Ratio*
definition: *Force on an object due to gravity*
definitionSource: *Wikipedia*
characterisedByUoM:{
id:*016*
name: *kilogram*
definition: *The base system of mass*
definitionSource: *International System of Units*
symbol: *kg*}
characterisedBySamplingMethod:{
id: *109*
processName: *Weigh*
processDescription: *The process for estimating mass*
InstrumentName: *Doran 50kg Fish Scales*
InstrumentType: *Digital Weighing Scales*}}
hasRelationshipObservation: {*SpecimenObservation}*
hasRelationshipSampledFeature: {*FishPopulationFeature}*
inTheme: {*Taxa Observations}*
Feature Type:{
id:*007*
name: *FishPopulationFeature*
definition: *A domain feature representing a fish population*
definitionSource: *AODC JF Technical Committee*
SpatioTemporalGeometry:{
id:*0007*
name: *3D-Volume*
scale: *Ratio*
description: *A Volume with no specific time constraints.*
shape: *GF_ Solid}*
Attribute:{
id:*012*
name: *taxonConceptName*
definition: *is a named classification unit (or taxon)*
definitionSource: *TDWG Taxon Concept (LSID) Ontology*
onlineOntologyFlag: *True*
OntologyURL: *http://rs.tdwg.org/ontology/voc/TaxonConcept.rdf*}
Attribute:{*...plus any other attributes that are specific to this feature type*}
hasRelationshipObservation: *SpecimenObservation*
hasRelationshipFOI: *BiologicalTaxonSamplingFeature*
inTheme: {*Taxa Observations}*

In the example provided above there are three rules which guide the mapping process from a schema encoding into the enhanced Feature Catalogue Model:

1. Observed properties that would normally be recorded in “ObservedProperty” become direct attributes of the sampling “feature-of-interest”.
2. If the sampling “feature-of-interest” is a feature collection, assign properties to the member feature type, and
3. The sampling processes described in “ObservationProcess” become attributes of the observed properties (which are themselves attributes). This mapping is similar in concept to the association “attributeOfAttribute” (from ISO 19109) which is used to create connections between attributes.

It should also be noted that the specialised Observation and the feature-of-interest will normally have the same ‘SpatioTemporalGeometry’. In cases where the specialised Observation is a collection this will not be the case.

Appendix 10 - FCATOWL Ontology Schema (on enclosed DVD in back cover)

Appendix 11 – FCAT Relational Model Data Dictionary

Table and Column Names	Description	Constraints
ATTRIBUTE	A characteristic of a Feature Type	
Attribute_Code	The unique Id for the Attribute	PK
Attribute_Definition	The lexical description for the Attribute	
Attribute_Name	The local name for the Attribute	
Attribute_Type	A classification of the attribute according to One of Probst's (2007) PhD Thesis Quality Space Taxa (from look-up table). Excludes "Shape" which is included in this Feature Catalogue Schema as "Spatio_Temporal_Geometry".	
Attribute_Value_Data_Type	The type of Attribute drawn fom a nominated listed namespace.	
Enumerated_Value_Domain	A binary flag (0, or 1) indicating whether the value of the attribute is selected from an enumerated list.	
Attribute_Value_Domain	The actual value of the attribute if from an enumerated list	
Scale_Type	Scale (e.g. interval, ratio etc drawn from a code list).	
Definition_Source_Code	A unique Id for a Definition Source.	FK
Deprecated_Flag	A binary flag indicating whether the Attribute is no longer in use and has been deprecated.	
Is_Mandatory	A binary flag denoting field is mandatory.	
ATTRIBUTE_METHOD	Links a Sampling Method to an Attribute	
Attribute_Code	The unique Id for the Attribute	PK,FK
Method_Code	The unique Id for the Sampling Method	PK,FK
ATTRIBUTE_TYPE	A classification of the attribute according to One of Probst's (2007) PhD Thesis Quality Space Taxa (from look-up table). Excludes "Shape" which is included in this Feature Catalogue Schema as "Spatio_Temporal_Geometry".	
Attribute_Type_Code	A unique Id for an Attribute_Type.	PK
Attribute_Type_Definition	A description of the Attribute Type.	
Attribute_Type_Name	Name of the Attribute_Type (e.g. 1D Spatial Location Quality, Volume). Covers all of Probst's spatial and non-spatial quality types except for "Shape".	
Definition_Source_Code	A unique Id for a Definition Source.	FK
Deprecated_Flag	A binary flag indicating whether the Attribute_Type_Code is no longer in use and has been deprecated.	
ATTRIBUTE_UOM	Links an Attribute to a Unit Of Measure and a Datum.	
Attribute_Code	The unique Id for the Attribute	PK,FK
Datum_Code	The unique Id for a Datum	FK
Deprecated_Flag	A binary flag indicating whether the Attribute UoM is no longer in use and has been deprecated.	
Uom_Code	The unique Id for a UoM	PK
ATTRIBUTE_VALUE_DOMAIN	Value for an enumerated Attribute domain.	

	including its codes and interpretation.	
Attribute_Label	Label given to Attribute (for an attribute of colour this could be red or green or blue)	
Attribute_Name	Name of the Attribute that the Attribute_Value is a pick list for.	
Attribute_Value_Code	A unique Id for an Attribute value	PK
Attribute_Value_Definition	A description of the Attribute value	
Definition_Source_Code	A unique Id for a Definition Source.	FK
Deprecated_Flag	A binary flag indicating whether the Attribute Attribute Value is no longer in use and has been deprecated.	
CHANGE_LOG	Audit log of changes to the schema content.	
Change_Date	The date on which the edit was made.	
Change_Description	A general description of the edit and why it was made.	
Change_Log_Id	A unique Id for the Change_Log entry.	PK
Fcat_Column_Name	The name of the Column in which the edit has taken place.	
Fcat_Table_Name	The name of the Table in which the edit has taken place.	
Feature_Catalogue_Code	Id of the Feature Catalogue in which this edit took place.	FK
Feature_Catalogue_Version _No	Versionof the Feature Catalogue in which this edit took place.	
New_Value	The value of the column after it was changed.	
Old_Value	The value of the column before it was changed.	
CLASSIFICATION_SCHEME	A Classification Scheme in which a Feature Type might participate.	
Classification_Scheme_Code	A unique Id for the Classification_Scheme.	PK
Classification_Scheme_C_Date	Date on which the Classification Scheme was created in the Feature_Catalogue.	
Classification_Scheme_Name	The name of the Classification_Scheme (e.g. SCAR TOPO Mapping Scheme).	
Definition_Source_Code	A unique Id for a Definition Source.	
Deprecated_Flag	A binary flag indicating whether the Classification-Scheme is no longer in use and has been deprecated.	
DATUM	A Datum which anchors an Attributes UoM.	
Datum_Code	A unique Id for a Datum.	PK
Datum_Definition	A description of the Datum	
Datum_Name	Name of the Datum (e.g. WGS 84)	
Definition_Source_Code	A unique Id for a Definition Source.	FK
Deprecated_Flag	A binary flag indicating whether the Datum is no longer in use and has been deprecated.	
DEFINITION_SOURCE	A description of the source of a definition which can contain links to online resources.	
Definition_Authority	Citation to the information identifying the source material for the definition (particularly if there is no Source URN to refer to).	
Definition_Sourced_Date	Date on which information was sourced.	
Definition_Source_Code	A unique Id for a Definition Source.	PK
Definition_Source_Urn	A unique Id conforming to standard for	

	naming URNs.	
Online_Ontology_Flag	A binary flag indicating if there is an online ontology which provides a definition for the term (in which case the URN should be a URL to the ontological definition).	
FEATURE_ CATALOGUE	A Feature Catalogue contains its identification and contact information and a broad definition of some of the Feature Types (as per ISO 19110).	
Deprecated_Flag	A binary flag indicating whether the Feature Catalogue is no longer in use and has been deprecated.	
Feature_Catalogue_Code	A unique Id for the Feature Catalogue.	PK
Feature_Catalogue_Description	A description of the Feature Catalogue.	
Feature_Catalogue_Name	The name given to the Feature Catalogue.	
Feature_Catalogue_Producer	The contact details of the Catalogue owner.	
Feature_Catalogue_Version_No	The version number of the Catalogue	
Feature_Cat_Creation_Date	The date on which the Catalogue was created.	
FEATURE_TYPE	A class of real-world phenomena (or object) with common properties.	
Feature_Type_Code	A unique Id for the Feature Type	FK
Deprecated_Flag	A binary flag indicating whether the Feature Type is no longer in use and has been deprecated.	
Has_Relationship	A flag indicating whether this Feature_Type participates in any Feature_Type to Feature_Type associations (e.g. Feature_Type A is part_of Feature_Type B). Actual relationships are found in the Relationship Table.	PK
Feature_Type_Definition	A description of the Feature Type	
Feature_Type_Name	A Feature Type Name	
Feature_Type_Urn	A URN for the Feature Type	
Definition_Source_Code	A unique Id for a Definition Source.	
FEATURE_TYPE_ATTRIBUTE	Links a Feature Type to an Attribute and a Feature Type Geometry and a Uinit of Measure.	
Attribute_Code	A unique Id for an Attribute.	PK,FK
Deprecated_Flag	A binary flag indicating whether the Feature Type Attribute is no longer in use and has been deprecated.	
Feature_Type_Code	A unique Id for a Feature Type.	PK,FK
Geometry_Code	A unique Id for a Feature Type Geometry.	
Uom_Code	A unique Id for a Unit of Measure.	
FEATURE_TYPE_ATTRIBUTE_PROFILE	Links several entities together (UoM, Attribute, Datum, Feature Type, Feature Type Profile, Spatio Temporal Geometry, Sampling Method).	
Added_By_Name	A unique Id for the Feature Type	
Attribute_Code	A unique Id for an Attribute.	PK,FK
Date_Added_To_Profile	The date the Feature Type and associated components were added to the Profile.	
Datum_Code	A unique Id for a Datum.	FK
Deprecated_Flag	A binary flag indicating whether the Feature	

	Type Attribute Profile is no longer in use and has been deprecated.	
Feature_Catalogue_Code	A unique Id for a Feature Catalogue.	PK,FK
Feature_Type_Code	A unique Id for a Feature Type.	PK,FK
Feature_Type_Profile_Id	A unique Id for a Feature Type Profile.	PK,FK
Geometry_Code	A unique Id for a Spatio-Temporal Geometry.	PK,FK
Uom_Code	A unique Id for a Unit of Measure.	
Method_Code	A unique Id for a Sampling Method.	FK
FEATURE_TYPE_CATALOGUE	Links a Feature Catalogue and a Feature Type.	
Deprecated_Flag	A binary flag indicating whether the Feature Type Catalogue is no longer in use and has been deprecated.	
Feature_Catalogue_Code	A unique Id for a Feature Catalogue.	PK,FK
Feature_Type_Code	A unique Id for a Feature Type.	PK,FK
FEATURE_TYPE_GEOMETRY	Links a Feature Type to a Spatio-Temporal Geometry.	
Deprecated_Flag	A binary flag indicating whether the Feature Type Geometry is no longer in use and has been deprecated.	
Feature_Type_Code	A unique Id for a Feature Type	PK,FK
Geometry_Code	A unique Id for a Spatio-Temporal Geometry	PK,FK
FEATURE_TYPE_PROFILE	A collection of Feature Types and their properties which belong to a specified user community.	
Deprecated_Flag	A binary flag indicating whether the Feature Type Profile is no longer in use and has been deprecated.	
Feature_Type_Profile_ID	A unique Id for the Profile	PK
Feature_Type_Profile_C_Date	The date on which the profile was created.	
Feature_Type_Profile_Description	General description of the profile (main purposes for which it is currently used)	
Feature_Type_Profile_Mod_Email	Contact email of the person maintaining and moderating the profile.	
Feature_Type_Profile_Mod_Name	The name of the person maninating and moderating the profile.	
Feature_Type_Profile_Name	Nominal name given to profile for human consumption.	
Profile_User_Code	A unique Id for a given User Profile.	FK
FEATURE_TYPE_SYMBOLOGY	Link between a Symbol, a Feature Type, a Feature Type Profile and a Spatio-Temporal Geometry.	
Deprecated_Flag	A binary flag indicating whether the Feature Type Symbology is no longer in use and has been deprecated.	
Feature_Type_Code	A unique Id for a Feature Type	PK,FK
Feature_Type_Profile_Id	A unique Id for a Feature Type Profile	FK
Geometry_Code	A unique Id for a Spatio-Temporal Geomtery.	FK
Symbol_Code	A unique Id for a symbol.	PK
FEATURE_TYPE_THEME_ CLASSIFICATION	Links a Feature Type, Theme, Classification Scheme, Feature Type Profile and a Symbol.	
Classification_Scheme_Code	A unique Id for a Classification Scheme.	PK,FK
Date_Added_To_Scheme	Dateon which the Feature Type was added	

	to the scheme.	
Deprecated_Flag	A binary flag indicating whether the Feature Type Theme Classification is no longer in use and has been deprecated.	
Feature_Type_Code	A unique Id for a Feature Type	PK,FK
Feature_Type_Profile_Id	A unique Id for a Feature Type Profile	PK,FK
Symbol_Code	A unique Id for a Symbol	FK
Theme_Code	A unique Id for a Theme.	PK,FK
PROFILE_USER	A class of user that owns a Feature Type Profile. The Profile User could be a named community, an application or a community project.	
Deprecated_Flag	A binary flag indicating whether theProfile User is no longer in use and has been deprecated.	
Profile_User_Code	A unique Id for a Profile User.	PK
Profile_User_Contact_Email	Contact email for the Profile User.	
Profile_User_Contact_Name	Point of contact for the User Profile.	
Profile_User_Description	A description of the Profile User.	
Profile_User_Name	Local name given to the Profile User.	
Profile_User_Url	URL to an online application.	
RELATIONSHIP	A named association between two Feature Types.	
Deprecated_Flag	A binary flag indicating whether the Relationship is no longer in use and has been deprecated.	
Object_Feature_Type	The unique Id of the FeatureType which is the object of the relationhsip (e.g. Subject_Feature_Type has_relationship Object_Feature_Type).	
Relationship_Class_Name	Additional information about a relationship property (e.g. Inverse, Functional, Transitive, Symmetric etc - all available from the Relationship Class Look-up Table).	
Relationship_Code	Unique Id for a relationship between two Feature Types.	PK
Relationship_Type_Name	Name of the relationship between two Feature_Types (e.g. part_of, synonym_for).	
Subject_Feature_Type	The unique_id of the Feature_Type which is the object of the relationhsip (e.g. Subject_Feature_Type has_relationship Object_Feature_Type).	
RELATIONSHIP_CLASS	The type of 'ontological' relationship drawn from a code list (e.g. Equivalent, Symmetric, Transitive).	
Definition_Source_Code	A unique Id for a Definition Source.	FK
Relationship_Class_Code	A unique Id for the Relationship Class.	PK
Relationship_Class_Description	A description of the Relationship Class.	
Relationship_Class_Name	A name for the Realtionship Class.	
RELATIONSHIP_TYPE	The type of realtion between two Feature Types (e.g., 'partOf').	
Definition_Source_Code	A unique Id for a Definition Source.	FK
Relationship_Type_Code	A unique Id for the Realtionship Type.	PK
Relationship_Type_Description	A description of the Relationship Type.	
Relationship_Type_Name	A name for the Realationship Type.	

SAMPLING_METHOD	A class containing attributes that describe the method by which a property has been derived, generated or sampled.	
Instrument_Name	A name for an instrument.	
Instrument_Type	A type of instrument drawn from a code list.	
Method_Code	A unique Id for the Sampling Method.	PK
Method_Description	A description of the Sampling Method.	
Method_Name	A name for the Sampling Method.	
SCALE	A scale code table	
Definition_Source_Code	A unique Id for a Definition Source.	FK
Deprecated_Flag	A binary flag indicating whether the Scale is no longer in use and has been deprecated.	
Scale_Code	A unique Id for a Scale.	PK
Scale_Definition	A description of the Scale.	
Scale_Type	The Scale type.	
SCHEMA	A reference to an online Application Schema.	
Schema_Code	A unique Id for the Application Schema.	PK
Schema_URI	The URL linking to the Schema.	
SCHEMA_ASSOCIATION	Links an Application Schema to a Feature Type and Feature Type Profile.	
Feature_Type_Code	A unque Id for a Feature Type.	PK,FK
Feature_Type_Profile_Id	A unique Id for a Feature Type Profile	PK,FK
Schema_Code	A unique Id for an Application Schema.	PK,FK
SCHEME_THEME	Links a Classification Schem to a Theme.	
Classification_Scheme_Code	A unique Id for a Classification Scheme.	PK,FK
Deprecated_Flag	A binary flag indicating whether the Scheme Theme is no longer in use and has been deprecated.	
Heirarchy_Level	Indicates if the Theme is at the base level in the classification (e.g. Level 1) or at a lower level in the heirachy (e.g. Level 2 or 3 or 4).	
Parent_Theme_Code	A unique Id for the Parent Theme.	
Parent_Theme_Name	If Heirarchy Level is 2 or higher the Parent_Theme_Name is the Theme Name immediately above this Theme.	
Theme_Code	A unique Id for the Theme.	PK,FK
SPATIO_TEMPORAL_GEOMETRY	A lexical description and code for the domain coordinates (temporal and spatial) of a Feature Type. The descriptions are drawn from a code list of predefined types (e.g. CSML Feature Types or O&M Sampling Features). The domain coordinates set the framework for a range of properties associated with an observation result type.	
Definition_Source_Code	A unique Id for a Definition Source.	FK
Deprecated_Flag	A binary flag indicating whether the Spatio-Temporal Geometry is no longer in use and has been deprecated.	
Geometry_Code	A unique Id for the spatio-temporal geometry.	PK
Geometry_Description	A description of the spatio-temporal geometry (e.g. 2D-polygon; 2D-Profile Series).	
Geometry_Name	A specific geometry type (e.g. point, line,	

	polygon, and the various sampling feature geometries drawn from O&M or community developed schema).	
Geometry_Shape	Description of the actual geometry	
SYMBOLOGY	A description of the attributes that depict how a Feature Type will be displayed in an online visualisation system.	
Symbol_Code	A unique Id for the symbol.	PK
Symbol_Display_Scale_Range	Information about display scales for the symbol.	
Symbol_Name	A name for the symbol.	
Symbol_URL	A URL to a sample of the symbology or to the symbol entry in a symbology catalogue.	
THEME	A user-defined group to which one or more Feature Types belong. They are components of a Classification Scheme.	
Definition_Source_Code	A unique Id for a Definition Source.	FK
Deprecated_Flag	A binary flag indicating whether the Theme is no longer in use and has been deprecated.	
Theme_Code	A unique Id for a Theme.	PK
Theme_Description	A description of the Theme.	
Theme_Name	A name for the Theme.	
UNITS_OF_MEASURE	Units Of Measure for an Attribute.	
Definition_Source_Code	A unique Id for a Definition Source.	FK
Deprecated_Flag	A binary flag indicating whether the Units Of Measure is no longer in use and has been deprecated.	
Uom_Code	A unique Id for the UoM	PK
Uom_Name	A name for the measurement	
Uom_Units	Units for the measurements scale	

Appendix 12 – hREST Capability Documents

IterativeFeatureCatalogueService

HTML Rendering

`1. GET Profile n in Catalogue n`

To make a query to retrieve individual profiles from a specific Feature Catalogue using the http method GET, use the request
`http://data.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/{id}/Profile/{id}`

The parameter: `id` - (a positive integer) is the identifier of a particular Catalogue and Profile respectively.

The output format for the request provides profile information encompassing links to any Feature Types associated with this profile. The response is supplied according to an XML schema document - Feature_Profile.xsd

`2. GET ALL Profiles in Catalogue n`

To make a query to retrieve all profiles from a specific Feature Catalogue using the http method GET, use the request
`http://data.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/{id}/Profile/`

The parameter: `id` - (a positive integer) is the identifier of a particular Profile and Catalogue respectively. A `/Profiles/` with no parameter following indicates the return of all profiles.

The output format for the request provides profile information for each profile stored in the Feature Catalogue encompassing links to any Feature Types associated with each individual profile. The response is supplied according to an XML schema document - Feature_Profile.xsd

`3. GET Feature Type n, from Profile n in Catalogue n`

To make a query to retrieve a specific Feature Type, from a specific Profile within a specific Catalogue using the http method GET, use the request
`http://data.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/{id}/Profile/{id}/FeatureType/{id}/`

The parameter: `id` (a positive integer) - identifies the id of the Catalogue, id of the Profile and the id of the Feature Type respectively.

The output format for the request provides information for that feature type with links to associated geometries, symbology, feature type attributes, themes that the feature type may belong to and associated classification schemes as well as any feature type relationships. The response is supplied according to an XML schema document - Feature_Types.xsd

`4. GET All Feature Types for Profile n in Catalogue n`

To make a query to retrieve all feature types, for a specific Profile from a specific Feature Catalogue using the http method GET, use the request

```
http://data.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/{id}/Profile/{id}/FeatureType/
```

The parameter: `id` - (a positive integer) identifies the id of the Catalogue and Profile respectively and a `/` with no parameter following the FeatureType indicates all features are returned.

The output format for the request provides information for all feature types for profile n where each feature type has links to associated geometries, symbology, feature type attributes, themes that the feature type may belong to and associated classification schemes as well as any feature type relationships. The response is supplied according to an XML schema document - Feature_Types.xsd

`5. GET Attribute n, for Feature Type n, in Profile n in Catalogue n`

To make a query to retrieve Attribute n for Feature Type n, for Profile n from a specific Feature Catalogue using the http method GET, use the request

```
http://data.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/{id}/Profile/{id}/FeatureType/{id}/Attribute/{id}
```

The parameter: `id` - (positive integer) identifies the id of the Catalogue, Profile, Feature Type and Attribute respectively.

The output format for the request provides information for the attribute and has links to associated units of measure. The response is supplied according to an XML schema document - Attributes.xsd

`6. GET All Attributes for Feature Type n, in Profile n of Catalogue n`

To make a query to retrieve all attributes for feature type n, for profile n from the Feature Catalogue using the http method GET, use the request

```
http://data.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/{id}/Profile/{id}/FeatureType/{id}/Attribute/
```

The parameter: `id` - (positive integer) identifies the id of the Catalogue, Profile, Feature Type and Attribute respectively and a `/` with no parameter following the Attribute indicates all attributes are returned.

The output format for the request provides information for all attributes and has links to associated units of measure. The response is supplied according to an XML schema document - Attributes.xsd

`7. GET Geometry n, for Feature Type n, in Profile n in Catalogue n`

To make a query to retrieve Geometry n for Feature Type n, for Profile n from a specific Feature Catalogue using the http method GET, use the request

```
http://data.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/{id}/Profile/{id}/FeatureType/{id}/Geometry/{id}
```

The parameter: `id` - (positive integer) identifies the id of the Catalogue, Profile, Feature Type and Geometry respectively.

The output format for the request provides information for geometries. The response is supplied according to an XML schema document - Geometries.xsd

`8. GET All Geometries for Feature Type n, in Profile n of Catalogue n`

To make a query to retrieve all geometries for feature type n, for profile n from the Feature Catalogue using the http method GET, use the request
`http://data.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/{id}/Profile/{id}/FeatureType/{id}/Geometry/`

The parameter: `id` - (positive integer) identifies the id of the Catalogue, Profile, Feature Type and Geometry respectively and a `/` with no parameter following the Geometry indicates all geometries are returned.

The output format for the request provides information for all geometries. The response is supplied according to an XML schema document - Geometries.xsd

`9. GET Unit-of-Measure n, for Attribute n, in Feature Type n, in Profile n in Catalogue n`

To make a query to retrieve Unit-of-Measure n, for Attribute n in Feature Type n, for Profile n from a specific Feature Catalogue using the http method GET, use the request
`http://data.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/{id}/Profile/{id}/FeatureType/{id}/Attribute/{id}/UoM/{id}`

The parameter: `id` - (positive integer) identifies the id of the Catalogue, Profile, Feature Type, Attribute and UoM respectively.

The output format for the request provides information for units-of-measure. The response is supplied according to an XML schema document - UoMs.xsd

`10. GET All Units-of-Measure, for Attribute n, in Feature Type n, in Profile n in Catalogue n`

To make a query to retrieve all Units-of-Measure, for Attribute n in Feature Type n, for Profile n from a specific Feature Catalogue using the http method GET, use the request
`http://data.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/{id}/Profile/{id}/FeatureType/{id}/Attribute/{id}/UoM/`

The parameter: `id` - (positive integer) identifies the id of the Catalogue, Profile, Feature Type, Attribute and UoM respectively and a `/` with no parameter following the UoM indicates all UoMs are returned.

The output format for the request provides information for units-of-measure. The response is supplied according to an XML schema document - UoMs.xsd

`11. GET Responses Using Formatting Option`

To make a query to retrieve Feature Catalogue content, in a specific format, using the http method GET, use the request
`http://data.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/{id}/Profile/{id}/Format={format}`

The parameters: `id` - (positive integer) identifies the id of the Catalogue and Profile respectively and `format` - (a string) identifies the format of the output file. The format option is specified at the end of the service request. Currently XML and html are options. Specifying no format provides XML output by default.

The output format for the sample request provides profile information. The response is supplied according to an XML schema document - Feature_Profile.xsd

Associated HTML Mark-Up

```html
<html>
<body>
<div class="service" id="svc">
<span class="label"><strong>IterativeFeatureCatalogueService</strong></span>
<div class="operation" id="1"><br/>
<code class="label"> 1. GET Profile n in Catalogue n</code><br/>
To make a query to retrieve individual profiles from a specific Feature Catalogue using the http method <span
  class="method">GET</span>, use the request <code class="address">
  http://data.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/{id}/Profile/{id}</code>
  <br/> </p>
<p><span class="input"> The parameter: <code>id</code> - (a positive integer) is the identifier of a particular
  Catalogue and Profile respectively.</span><br/></p>
<p><span class="output">The output format for the request provides profile information encompassing links to
  any Feature Types associated with this profile. The response is supplied according to an XML schema document
  - <a href  ="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/Feature_profile.xsd">Feature_Profile.xsd
  </a></span></p>
</div>
<div class="operation" id="2">
<code class="label"> 2. GET ALL Profiles in Catalogue n</code><br/>
<p> To make a query to retrieve all profiles from a specific Feature Catalogue using the http method <span
  class="method">GET</span>, use the request <code class="address">
  http://data.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/{id}/Profile/</code>
  <br/> </p>
<p><span class="input"> The parameter: <code>id</code> - (a positive integer) is the identifier of a particular
  Profile and Catalogue respectively. A <code>/Profiles/</code> with no parameter following indicates the return
  of all profiles.</span><br/></p>
<p><span class="output">The output format for the request provides profile information for each profile stored
  in the Feature Catalogue encompassing links to any Feature Types associated with each individual profile. The
  response is supplied according to an XML schema document - <a href
  ="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/Feature_profile.xsd">Feature_Profile.xsd
  </a></span></p>
</div>
<div class="operation" id="3">
<code class="label"> 3. GET Feature Type n, from Profile n in Catalogue n</code><br/>
<p> To make a query to retrieve a specific Feature Type, from a specific Profile within a specific Catalogue using
  the http method <span class="method">GET</span>, use the request <code class="address">
  http://data.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/{id}/Profile/{id}/Featur
  eType/{id}/</code> <br/> </p>
<p><span class="input"> The parameter: <code>id</code> (a positive integer) - identifies the id of the
  Catalogue, id of the Profile and the id of the Feature Type respectively.</span><br/></p>
<p><span class="output">The output format for the request provides information for that feature type with links
  to associated geometries, symbology, feature type attributes, themes that the feature type may belong to and
  associated classification schemes as well as any feature type relationships. The response is supplied according
  to an XML schema document - <a href
  ="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/Feature_Types.xsd">Feature_Types.xsd
  </a></span></p>
</div>
<div class="operation" id="4">
<code class="label"> 4. GET All Feature Types for Profile n in Catalogue n</code><br/>
<p> To make a query to retrieve all feature types, for a specific Profile from a specific Feature Catalogue using
  the http method <span class="method">GET</span>, use the request <code class="address">
```

```
http://data.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/{id}/Profile/{id}/Featur
eType/</code> <br/> </p>
<p><span class="input"> The parameter: <code>id</code> - (a positive integer) identifies the id of the Catalogue
and Profile respectively and a <code>/</code> with no parameter following the FeatureType indicates all
features are returned.</span><br/></p>
<p><span class="output">The output format for the request provides information for all feature types for profile
n where each feature type has links to associated geometries, symbology, feature type attributes, themes that
the feature type may belong to and associated classification schemes as well as any feature type relationships.
The response is supplied according to an XML schema document - <a href
="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/Feature_Types.xsd">Feature_Types.xsd
</a></span></p>
</div>
<div class="operation" id="5">
<code class="label"> 5. GET Attribute n, for Feature Type n, in Profile n in Catalogue n</code><br/>
<p> To make a query to retrieve Attribute n for Feature Type n, for Profile n from a specific Feature Catalogue
using the http method <span class="method">GET</span>, use the request <code class="address">
http://data.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/{id}/Profile/{id}/Featur
eType/{id}/Attribute/{id}</code> <br/> </p>
<p><span class="input"> The parameter: <code>id</code> - (positive integer) identifies the id of the Catalogue,
Profile, Feature Type and Attribute respectively.</span><br/></p>
<p><span class="output">The output format for the request provides information for the attribute and has links
to associated units of measure. The response is supplied according to an XML schema document - <a href
="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/Attributes.xsd">Attributes.xsd</a></span></p>
</div>
</div>
<div class="operation" id="6">
<code class="label"> 6. GET All Attributes for Feature Type n, in Profile n of Catalogue n</code><br/>
<p> To make a query to retrieve all attributes for feature type n, for profile n from the Feature Catalogue using
the http method <span class="method">GET</span>, use the request <code class="address">
http://data.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/{id}/Profile/{id}/Featur
eType/{id}/Attribute/</code> <br/> </p>
<p><span class="input"> The parameter: <code>id</code> - (positive integer) identifies the id of the Catalogue,
Profile, Feature Type and Attribute respectively and a <code>/</code> with no parameter following the
Attribute indicates all attributes are returned.</span><br/></p>
<p><span class="output">The output format for the request provides information for all attributes and has links
to associated units of measure. The response is supplied according to an XML schema document - <a href
="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/Attributes.xsd">Attributes.xsd</a></span></p>
</div>
</div>
<div class="operation" id="7">
<code class="label"> 7. GET Geometry n, for Feature Type n, in Profile n in Catalogue n</code><br/>
<p> To make a query to retrieve Geometry n for Feature Type n, for Profile n from a specific Feature Catalogue
using the http method <span class="method">GET</span>, use the request <code class="address">
http://data.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/{id}/Profile/{id}/Featur
eType/{id}/Geometry/{id}</code> <br/> </p>
<p><span class="input"> The parameter: <code>id</code> - (positive integer) identifies the id of the Catalogue,
Profile, Feature Type and Geometry respectively.</span><br/></p>
<p><span class="output">The output format for the request provides information for geometries. The response
is supplied according to an XML schema document - <a href
="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/Geometries.xsd">Geometries.xsd</a></span></p>
</div>
<div class="operation" id="8">
<code class="label"> 8. GET All Geometries for Feature Type n, in Profile n of Catalogue n</code><br/>
<p> To make a query to retrieve all geometries for feature type n, for profile n from the Feature Catalogue using
the http method <span class="method">GET</span>, use the request <code class="address">
http://data.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/{id}/Profile/{id}/Featur
eType/{id}/Geometry/</code> <br/> </p>
```

```
<p><span class="input"> The parameter: <code>id</code> - (positive integer) identifies the id of the Catalogue,
  Profile, Feature Type and Geometry respectively and a <code>/</code> with no parameter following the
  Geometry indicates all geometries are returned.</span><br/></p>
<p><span class="output">The output format for the request provides information for all geometries. The
  response is supplied according to an XML schema document - <a href
  ="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/Geometries.xsd">Geometries.xsd</a></span></p>
</div>
<div class="operation" id="9">
<code class="label"> 9. GET Unit-of-Measure n, for Attribute n, in Feature Type n, in Profile n in Catalogue
  n</code><br/>
<p> To make a query to retrieve Unit-of-Measure n, for Attribute n in Feature Type n, for Profile n from a specific
  Feature Catalogue using the http method <span class="method">GET</span>, use the request <code
  class="address">
  http://data.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/{id}/Profile/{id}/Featur
  eType/{id}/Attribute/{id}/UoM/{id}</code> <br/> </p>
<p><span class="input"> The parameter: <code>id</code> - (positive integer) identifies the id of the Catalogue,
  Profile, Feature Type, Attribute and UoM respectively.</span><br/></p>
<p><span class="output">The output format for the request provides information for units-of-measure. The
  response is supplied according to an XML schema document - <a href
  ="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/UoMs.xsd">UoMs.xsd</a></span></p>
</div>
<div class="operation" id="10">
<code class="label"> 10. GET All Units-of-Measure, for Attribute n, in Feature Type n, in Profile n in Catalogue
  n</code><br/>
<p> To make a query to retrieve all Units-of-Measure, for Attribute n in Feature Type n, for Profile n from a
  specific Feature Catalogue using the http method <span class="method">GET</span>, use the request <code
  class="address">
  http://data.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/{id}/Profile/{id}/Featur
  eType/{id}/Attribute/{id}/UoM/</code> <br/> </p>
<p><span class="input"> The parameter: <code>id</code> - (positive integer) identifies the id of the Catalogue,
  Profile, Feature Type, Attribute and UoM respectively and a <code>/</code> with no parameter following the
  UoM indicates all UoMs are returned.</span><br/></p>
<p><span class="output">The output format for the request provides information for units-of-measure. The
  response is supplied according to an XML schema document - <a href
  ="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/UoMs.xsd">UoMs.xsd</a></span></p>
</div>
<div class="operation" id="11">
<code class="label"> 11. GET Responses Using Formatting Option </code><br/>
<p> To make a query to retrieve Feature Catalogue content, in a specific format, using the http method <span
  class="method">GET</span>, use the request <code class="address">
  http://data.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/{id}/Profile/{id}/Format
  ={format}</code> <br/> </p>
<p><span class="input"> The parameters: <code>id</code> - (positive integer) identifies the id of the Catalogue
  and Profile respectively and <code>format</code> - (a string) identifies the format of the output file. The
  format option is specified at the end of the service request. Currently XML and html are options. Specifying no
  format provides XML output by default.</span><br/></p>
<p><span class="output">The output format for the sample request provides profile information. The response
  is supplied according to an XML schema document - <a href
  ="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/Feature_Profile.xsd">Feature_Profile.xsd
  </a></span></p>
</div>
</div>
</body>
</html>
```

StreamlinedFeatureCatalogueService

`1. GET Profile n in Catalogue n`

To make a query to retrieve an individual profile from a specific Feature Catalogue using the http method GET, use the request
`http://data.aad.gov.au/aadc/FeatureCatalogue/?/StreamlinedCatalogueService/Catalogue/{id}/Profile/{id}`

The parameter: `id` - (a positive integer) is the identifier of a particular catalogue and profile respectively.

The output format for the request provides profile information encompassing Catalogue, Feature Types, Feature Type Attributes, UoMs, Datums, Feature Type Relationships and the Classification Schemes and Themes that the profile's Feature Types belong to. The response is supplied according to an XML schema document - Profile_Package.xsd

`2. GET ALL Profiles in Catalogue n`

To make a query to retrieve all profiles from a specific Feature Catalogue using the http method GET, use the request
`http://data.aad.gov.au/aadc/FeatureCatalogue/?/StreamlinedCatalogueService/Catalogue/{id}/Profile/`

The parameter: `id` - (a positive integer) is the identifier of a particular catalogue. No parameter following the `/Profiles/` indicates all profiles are returned.

The output format for the request provides profile information encompassing Catalogue, Feature Types, Feature Type Attributes, UoMs, Datums, Feature Type Relationships and the Classification Schemes and Themes that the profile's Feature Types belong to. The response is supplied according to an XML schema document - Profile_Package.xsd

`3. GET Feature Type n, from Profile n in Catalogue n`

To make a query to retrieve a specific Feature Type, from a specific profile from the Feature Catalogue using the http method GET, use the request
`http://data.aad.gov.au/aadc/FeatureCatalogue/?/StreamlinedCatalogueService/Catalogue/{id}/Profile/{id}/FeatureType/{id}/`

The parameter: `id` - (a positive integer) identifies the id of the Catalogue, the id of the Profile and the id of the Feature Type respectively.

The output format for the request provides profile information encompassing Catalogue, Feature Types, Feature Type Attributes, UoMs, Datums, Feature Type Relationships and the Classification Schemes and Themes that the profile's Feature Types belong to. The response is supplied according to an XML schema document - Profile_Package.xsd

`4. GET Responses Using Formatting Option`

To make a query to retrieve Feature Catalogue content, in a specific format, using the http method GET, use the request
`http://data.aad.gov.au/aadc/FeatureCatalogue/?/StreamlinedCatalogueService/Catalogue/{id}/Profile/{id}/Format={format}`

The parameters: `id` - (a positive integer) identifies the id of the Catalogue and the id of Profile respectively and `format` - (a string) identifies the format of the output file. The format option is specified at the end of the service request. Currently XML and html are options. Specifying no format provides XML output by default.

The output format for the request provides profile information. The response is supplied according to an XML schema document - Profile_Package.xsd

Associated HTML Mark-Up

```
<html>
<body>
<div class="service" id="svc">
<span class="label"><strong>StreamlinedFeatureCatalogueService</strong></span>
<div class="operation" id="1"><br/>
<code class="label"> 1. GET Profile n in Catalogue n</code><br/>
To make a query to retrieve an individual profile from a specific Feature Catalogue using the http method <span
  class="method">GET</span>, use the request <code class="address">
  http://data.aad.gov.au/aadc/FeatureCatalogue/?/StreamlinedCatalogueService/Catalogue/{id}/Profile/{id}</co
  de> <br/> </p>
<p><span class="input"> The parameter: <code>id</code> - (a positive integer) is the identifier of a particular
  catalogue and profile respectively.</span><br/></p>
<p><span class="output">The output format for the request provides profile information encompassing
  Catalogue, Feature Types, Feature Type Attributes, UoMs, Datums, Feature Type Relationships and the
  Classification Schemes and Themes that the profile's Feature Types belong to. The response is supplied
  according to an XML schema document - <a href
  ="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/Profile_Package.xsd">Profile_Package.xsd
  </a></span></p>
</div>
<div class="operation" id="2">
<code class="label"> 2. GET ALL Profiles in Catalogue n</code><br/>
<p> To make a query to retrieve all profiles from a specific Feature Catalogue using the http method <span
  class="method">GET</span>, use the request <code class="address">
  http://data.aad.gov.au/aadc/FeatureCatalogue/?/StreamlinedCatalogueService/Catalogue/{id}/Profile/</code>
  <br/> </p>
<p><span class="input"> The parameter: <code>id</code> - (a positive integer) is the identifier of a particular
  catalogue. No parameter following the <code>/Profiles/</code> indicates all profiles are
  returned.</span><br/></p>
<p><span class="output">The output format for the request provides profile information encompassing
  Catalogue, Feature Types, Feature Type Attributes, UoMs, Datums, Feature Type Relationships and the
  Classification Schemes and Themes that the profile's Feature Types belong to. The response is supplied
  according to an XML schema document - <a href
  ="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/Profile_Package.xsd">Profile_Package.xsd
  </a></span></p>
</div>
<div class="operation" id="3">
<code class="label"> 3. GET Feature Type n, from Profile n in Catalogue n</code><br/>
<p> To make a query to retrieve a specific Feature Type, from a specific profile from the Feature Catalogue using
  the http method <span class="method">GET</span>, use the request <code class="address">
  http://data.aad.gov.au/aadc/FeatureCatalogue/?/StreamlinedCatalogueService/Catalogue/{id}/Profile/{id}/Fea
  tureType/{id}/</code> <br/> </p>
<p><span class="input"> The parameter: <code>id</code> - (a positive integer) identifies the id of the
  Catalogue, the id of the Profile and the id of the Feature Type respectively.</span><br/></p>
<p><span class="output">The output format for the request provides profile information encompassing
  Catalogue, Feature Types, Feature Type Attributes, UoMs, Datums, Feature Type Relationships and the
  Classification Schemes and Themes that the profile's Feature Types belong to. The response is supplied
```

```
  according to an XML schema document - <a href
  ="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/Profile_Package.xsd">Profile_Package.xsd
  </a></span></p>
</div>
<div class="operation" id="4">
<code class="label"> 4. GET Responses Using Formatting Option</code><br/>
<p> To make a query to retrieve Feature Catalogue content, in a specific format, using the http method <span
  class="method">GET</span>, use the request <code class="address">
  http://data.aad.gov.au/aadc/FeatureCatalogue/?/StreamlinedCatalogueService/Catalogue/{id}/Profile/{id}/For
  mat={format}</code> <br/> </p>
<p><span class="input"> The parameters: <code>id</code> - (a positive integer) identifies the id of the
  Catalogue and the id of Profile respectively and <code>format</code> - (a string) identifies the format of the
  output file. The format option is specified at the end of the service request. Currently XML and html are
  options. Specifying no format provides XML output by default.</span><br/></p>
<p><span class="output">The output format for the request provides profile information. The response is
  supplied according to an XML schema document - <a href
  ="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/Profile_Package.xsd">Profile_Package.xsd
  </a></span></p>
</div>
</div>
</body>
</html>
```

ComponentFeatureCatalogueService

`1. GET Profile n in Catalogue n`

To make a query to retrieve individual Profiles from a specific Feature Catalogue using the http method GET, use the request

```
http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Ca
talogue/{id}/Profile/{id}
```

The parameter: `id` - (a positive integer) is the identifier of a particular Catalogue and Profile respectively.

The output format for the request provides profile information encompassing links to any Feature Types associated with this profile. The response is supplied according to an XML schema document - Feature_Profile.xsd

`2. GET ALL Profiles in Catalogue n`

To make a query to retrieve all Profiles from a specific Feature Catalogue using the http method GET, use the request

```
http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Ca
talogue/{id}/Profile/
```

The parameter: `id` - (a positive integer) is the identifier of a particular Catalogue. A `/` with no parameter following indicates the return of all profiles.

The output format for the request provides profile information for each profile stored in a specific Feature Catalogue encompassing links to any Feature Types associated with each individual profile. The response is supplied according to an XML schema document - Feature_Profile.xsd

`3. GET Feature Type n in Catalogue n`

To make a query to retrieve a specific Feature Type from a specific Feature Catalogue using the http method GET, use the request
`http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/{id}/FeatureType/{id}/`

The parameter: `id` (a positive integer) - identifies the id of the Catalogue and Feature Type respectively.

The output format for the request provides information for that feature type with links to associated geometries, symbology, feature type attributes, themes that the feature type may belong to and associated classification schemes as well as any feature type relationships. The response is supplied according to an XML schema document - Feature_Types.xsd

`4. GET All Feature Types in Catalogue n`

To make a query to retrieve all feature types from the Feature Catalogue using the http method GET, use the request
`http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/{id}/FeatureType/`

The parameter: `id` (a positive integer) - identifies the id of the Catalogue. A `/` with no parameter following the FeatureType indicates all features are returned.

The output format for the request provides information for all feature types where each feature type has links to associated geometries, symbology, feature type attributes, themes that the feature type may belong to and associated classification schemes as well as any feature type relationships. The response is supplied according to an XML schema document - Feature_Types.xsd

`5. GET Catalogue n`

To make a query to retrieve Catalogue n using the http method GET, use the request
`http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/{id}`

The parameter: `id` - (positive integer) identifies the id of the Catalogue.

The output format for the request provides information for the catalogue. The response is supplied according to an XML schema document - Catalogues.xsd

`6. GET All Catalogues`

To make a query to retrieve all catalogues using the http method GET, use the request
`http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/`

A `/` with no parameter following the Catalogue indicates all catalogues are returned.

The output format for the request provides information for all catalogues. The response is supplied according to an XML schema document - Catalogues.xsd

`7. GET Attribute n in Catalogue n`

To make a query to retrieve Attribute n from a specific Feature Catalogue using the http method GET, use the request `http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/{id}/Attribute/{id}`

The parameter: `id` - (positive integer) identifies the id of the Catalogue and Attribute respectively.

The output format for the request provides information for the attribute and has links to associated units of measure. The response is supplied according to an XML schema document - [Attributes.xsd]

`8. GET All Attributes in Catalogue n`

To make a query to retrieve all Attributes from a specific Feature Catalogue using the http method GET, use the request
`http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/{id}/Attribute/`

The parameter: `id` - (positive integer) identifies the id of the Catalogue. A `/` with no parameter following the Attribute indicates all attributes are returned.

The output format for the request provides information for all attributes and has links to associated units of measure. The response is supplied according to an XML schema document - [Attributes.xsd]

`9. GET Geometry n in Catalogue n`

To make a query to retrieve Geometry n from a specific Feature Catalogue using the http method GET, use the request
`http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/{id}/Geometry/{id}`

The parameter: `id` - (positive integer) identifies the id of the Catalogue and Geometry respectively.

The output format for the request provides information for the geometry. The response is supplied according to an XML schema document - [Geometries.xsd]

`10. GET All Geometries in Catalogue n`

To make a query to retrieve all geometries from the Feature Catalogue using the http method GET, use the request
`http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/{id}/Geometry/`

The parameter: `id` - (positive integer) identifies the id of the Catalogue. A `/` with no parameter following the Geometry indicates all geometries are returned.

The output format for the request provides information for all geometries. The response is supplied according to an XML schema document - [Geometries.xsd]

`11. GET Unit-of-Measure n in Catalogue n`

To make a query to retrieve Unit-of-Measure n from a specific Feature Catalogue using the http method GET, use the request
`http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/{id}/UoM/{id}`

The parameter: `id` - (positive integer) identifies the id of the Catalogue and UoM respectively.

The output format for the request provides information for unit-of-measure. The response is supplied according to an XML schema document - UoMs.xsd

`12. GET All Units-of-Measure in Catalogue n`

To make a query to retrieve all Units-of-Measure from a specific Feature Catalogue using the http method GET, use the request
`http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/{id}/UoM/`

The parameter: `id` - (positive integer) identifies the id of a particular Catalogue. A `/` with no parameter following the UoM indicates all UoMs are returned.

The output format for the request provides information for units-of-measure. The response is supplied according to an XML schema document - UoMs.xsd

`13. GET Relationship n in Catalogue n`

To make a query to retrieve Relationship n from a specific Feature Catalogue using the http method GET, use the request
`http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/{id}/Relationship/{id}`

The parameter: `id` - (positive integer) identifies the id of the Catalogue and Relationship respectively.

The output format for the request provides information for the relationship. The response is supplied according to an XML schema document - Relationships.xsd

`14. GET All Relationships in Catalogue n`

To make a query to retrieve all Relationships from a specific Feature Catalogue using the http method GET, use the request
`http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/{id}/Relationship/`

The parameter: `id` - (positive integer) identifies the id of a particular Catalogue. A `/` with no parameter following the Relationship indicates all relations are returned.

The output format for the request provides information for relationships. The response is supplied according to an XML schema document - Relationships.xsd

`15. GET Classification Scheme n in Catalogue n`

To make a query to retrieve Classification Scheme n from a specific Feature Catalogue using the http method GET, use the request
`http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/{id}/Scheme/{id}`

The parameter: `id` - (positive integer) identifies the id of the Catalogue and Scheme respectively.

The output format for the request provides information for the scheme. The response is supplied according to an XML schema document - Scheme_Package.xsd

`16. GET All Classification Schemes in Catalogue n`

To make a query to retrieve all Classification Schemes from a specific Feature Catalogue using the http method GET, use the request
`http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/{id}/Scheme/`

The parameter: `id` - (positive integer) identifies the id of a particular Catalogue. A `/` with no parameter following the Scheme indicates all schemes are returned.

The output format for the request provides information for schemes. The response is supplied according to an XML schema document - [Scheme_Package.xsd]()

`17. GET Responses Using Formatting Option`

To make a query to retrieve Feature Catalogue content, in a specific format, using the http method GET, use the request
`http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/{id}/Profile/{id}/Format={format}`

The parameters: `id` (positive integer) identifies the id of the Profile and `format` - (a string) identifies the format of the output file. The format option is specified at the end of the service request. Currently XML and html are options. Specifying no format provides XML output by default.

The output format for the sample request provides profile information. The response is supplied according to an XML schema document - [Feature_Profile.xsd]()

Associated HTML Mark-Up

```
<html>
<body>
<div class="service" id="svc">
<span class="label"><strong>ComponentFeatureCatalogueService</strong></span>
<div class="operation" id="1"><br/>
<code class="label"> 1. GET Profile n in Catalogue n</code><br/>
To make a query to retrieve individual Profiles from a specific Feature Catalogue using the http method <span
  class="method">GET</span>, use the request <code class="address">
  http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/{id}/Profile/{id}</co
  de> <br/> </p>
<p><span class="input"> The parameter: <code>id</code> - (a positive integer) is the identifier of a particular
  Catalogue and Profile respectively.</span><br/></p>
<p><span class="output">The output format for the request provides profile information encompassing links to
  any Feature Types associated with this profile. The response is supplied according to an XML schema document
  - <a href ="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/Feature_profile.xsd">Feature_Profile.xsd
  </a></span></p>
</div>
<div class="operation" id="2">
<code class="label"> 2. GET ALL Profiles in Catalogue n</code><br/>
<p> To make a query to retrieve all Profiles from a specific Feature Catalogue using the http method <span
  class="method">GET</span>, use the request <code class="address">
  http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/{id}/Profile/</code>
  <br/> </p>
<p><span class="input"> The parameter: <code>id</code> - (a positive integer) is the identifier of a particular
  Catalogue. A <code>/</code> with no parameter following indicates the return of all
  profiles.</span><br/></p>
<p><span class="output">The output format for the request provides profile information for each profile stored
  in a specific Feature Catalogue encompassing links to any Feature Types associated with each individual profile.
```

```html
  The response is supplied according to an XML schema document - <a href
  ="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/Feature_profile.xsd">Feature_Profile.xsd
  </a></span></p>
</div>
<div class="operation" id="3">
<code class="label"> 3. GET Feature Type n in Catalogue n</code><br/>
<p> To make a query to retrieve a specific Feature Type from a specific Feature Catalogue using the http method
  <span class="method">GET</span>, use the request <code class="address">
  http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/{id}/FeatureType/{i
  d}/</code> <br/> </p>
<p><span class="input"> The parameter: <code>id</code> (a positive integer) - identifies the id of the Catalogue
  and Feature Type respectively.</span><br/></p>
<p><span class="output">The output format for the request provides information for that feature type with links
  to associated geometries, symbology, feature type attributes, themes that the feature type may belong to and
  associated classification schemes as well as any feature type relationships. The response is supplied according
  to an XML schema document - <a href
  ="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/Feature_Types.xsd">Feature_Types.xsd
  </a></span></p>
</div>
<div class="operation" id="4">
<code class="label"> 4. GET All Feature Types in Catalogue n</code><br/>
<p> To make a query to retrieve all feature types from the Feature Catalogue using the http method <span
  class="method">GET</span>, use the request <code class="address">
  http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/{id}/FeatureType/</
  code> <br/> </p>
<p><span class="input"> The parameter: <code>id</code> (a positive integer) - identifies the id of the
  Catalogue. A <code>/</code> with no parameter following the FeatureType indicates all features are
  returned.</span><br/></p>
<p><span class="output">The output format for the request provides information for all feature types where
  each feature type has links to associated geometries, symbology, feature type attributes, themes that the
  feature type may belong to and associated classification schemes as well as any feature type relationships. The
  response is supplied according to an XML schema document - <a href
  ="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/Feature_Types.xsd">Feature_Types.xsd
  </a></span></p>
</div>
<div class="operation" id="5">
<code class="label"> 5. GET Catalogue n </code><br/>
<p> To make a query to retrieve Catalogue n  using the http method <span class="method">GET</span>, use the
  request <code class="address">
  http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/{id}</code> <br/>
  </p>
<p><span class="input"> The parameter: <code>id</code> - (positive integer) identifies the id of the
  Catalogue.</span><br/></p>
<p><span class="output">The output format for the request provides information for the catalogue. The
  response is supplied according to an XML schema document - <a href
  ="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/Catalogues.xsd">Catalogues.xsd</a></span></p>
</div>
<div class="operation" id="6">
<code class="label"> 6. GET All Catalogues </code><br/>
<p> To make a query to retrieve all catalogues  using the http method <span class="method">GET</span>, use
  the request <code class="address">
  http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/</code> <br/> </p>
<p><span class="input"> A <code>/</code> with no parameter following the Catalogue indicates all catalogues
  are returned.</span><br/></p>
<p><span class="output">The output format for the request provides information for all catalogues. The
  response is supplied according to an XML schema document - <a href
  ="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/Catalogues.xsd">Catalogues.xsd</a></span></p>
```

```
</div>
<div class="operation" id="7">
<code class="label"> 7. GET Attribute n in Catalogue n</code><br/>
<p> To make a query to retrieve Attribute n from a specific Feature Catalogue using the http method <span class="method">GET</span>, use the request <code class="address"> http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/{id}/Attribute/{id}</code> <br/> </p>
<p><span class="input"> The parameter: <code>id</code> - (positive integer) identifies the id of the Catalogue and Attribute respectively.</span><br/></p>
<p><span class="output">The output format for the request provides information for the attribute and has links to associated units of measure. The response is supplied according to an XML schema document - <a href ="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/Attributes.xsd">Attributes.xsd</a></span></p>
</div>
<div class="operation" id="8">
<code class="label"> 8. GET All Attributes in Catalogue n </code><br/>
<p> To make a query to retrieve all Attributes from a specific Feature Catalogue using the http method <span class="method">GET</span>, use the request <code class="address"> http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/{id}/Attribute/</code> <br/> </p>
<p><span class="input"> The parameter: <code>id</code> - (positive integer) identifies the id of the Catalogue. A <code>/</code> with no parameter following the Attribute indicates all attributes are returned.</span><br/></p>
<p><span class="output">The output format for the request provides information for all attributes and has links to associated units of measure. The response is supplied according to an XML schema document - <a href ="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/Attributes.xsd">Attributes.xsd</a></span></p>
</div>
<div class="operation" id="9">
<code class="label"> 9. GET Geometry n  in Catalogue n</code><br/>
<p> To make a query to retrieve Geometry n from a specific Feature Catalogue using the http method <span class="method">GET</span>, use the request <code class="address"> http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/{id}/Geometry/{id}</code> <br/> </p>
<p><span class="input"> The parameter: <code>id</code> - (positive integer) identifies the id of the Catalogue and Geometry respectively.</span><br/></p>
<p><span class="output">The output format for the request provides information for the geometry. The response is supplied according to an XML schema document - <a href ="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/Geometries.xsd">Geometries.xsd</a></span></p>
</div>
<div class="operation" id="10">
<code class="label"> 10. GET All Geometries  in Catalogue n</code><br/>
<p> To make a query to retrieve all geometries from the Feature Catalogue using the http method <span class="method">GET</span>, use the request <code class="address"> http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/{id}/Geometry/</code> <br/> </p>
<p><span class="input"> The parameter: <code>id</code> - (positive integer) identifies the id of the Catalogue. A <code>/</code> with no parameter following the Geometry indicates all geometries are returned.</span><br/></p>
<p><span class="output">The output format for the request provides information for all geometries. The response is supplied according to an XML schema document - <a href ="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/Geometries.xsd">Geometries.xsd</a></span></p>
</div>
<div class="operation" id="11">
<code class="label"> 11. GET Unit-of-Measure n  in Catalogue n</code><br/>
<p> To make a query to retrieve Unit-of-Measure n from a specific Feature Catalogue using the http method <span class="method">GET</span>, use the request <code class="address"> http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/{id}/UoM/{id}</code> <br/> </p>
```

```
<p><span class="input"> The parameter: <code>id</code> - (positive integer) identifies the id of the Catalogue
  and UoM respectively.</span><br/></p>
<p><span class="output">The output format for the request provides information for unit-of-measure. The
  response is supplied according to an XML schema document - <a href
  ="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/UoMs.xsd">UoMs.xsd</a></span></p>
</div>
<div class="operation" id="12">
<code class="label"> 12. GET All Units-of-Measure  in Catalogue n</code><br/>
<p> To make a query to retrieve all Units-of-Measure from a specific Feature Catalogue using the http method
  <span class="method">GET</span>, use the request <code class="address">
  http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/{id}/UoM/</code>
  <br/> </p>
<p><span class="input"> The parameter: <code>id</code> - (positive integer) identifies the id of a particular
  Catalogue. A <code>/</code> with no parameter following the UoM indicates all UoMs are
  returned.</span><br/></p>
<p><span class="output">The output format for the request provides information for units-of-measure. The
  response is supplied according to an XML schema document - <a href
  ="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/UoMs.xsd">UoMs.xsd</a></span></p>
</div>
<div class="operation" id="13">
<code class="label"> 13. GET Relationship n  in Catalogue n</code><br/>
<p> To make a query to retrieve Relationship n from a specific Feature Catalogue using the http method <span
  class="method">GET</span>, use the request <code class="address">
  http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/{id}/Relationship/{id
  }</code> <br/> </p>
<p><span class="input"> The parameter: <code>id</code> - (positive integer) identifies the id of the Catalogue
  and Relationship respectively.</span><br/></p>
<p><span class="output">The output format for the request provides information for the relationship. The
  response is supplied according to an XML schema document - <a href
  ="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/Relationships.xsd">Relationships.xsd</a></span></p>
</div>
<div class="operation" id="14">
<code class="label"> 14. GET All Relationships  in Catalogue n</code><br/>
<p> To make a query to retrieve all Relationships from a specific Feature Catalogue using the http method <span
  class="method">GET</span>, use the request <code class="address">
  http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/{id}/Relationship/</
  code> <br/> </p>
<p><span class="input"> The parameter: <code>id</code> - (positive integer) identifies the id of a particular
  Catalogue. A <code>/</code> with no parameter following the Relationship indicates all relations are
  returned.</span><br/></p>
<p><span class="output">The output format for the request provides information for relationships. The
  response is supplied according to an XML schema document - <a href
  ="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/Relationships.xsd">Relationships.xsd</a></span></p>
</div>
<div class="operation" id="15">
<code class="label"> 15. GET Classification Scheme n  in Catalogue n</code><br/>
<p> To make a query to retrieve Classification Scheme n from a specific Feature Catalogue using the http method
  <span class="method">GET</span>, use the request <code class="address">
  http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/{id}/Scheme/{id}</c
  ode> <br/> </p>
<p><span class="input"> The parameter: <code>id</code> - (positive integer) identifies the id of the Catalogue
  and Scheme respectively.</span><br/></p>
<p><span class="output">The output format for the request provides information for the scheme. The response
  is supplied according to an XML schema document - <a href
  ="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/Scheme_Package.xsd">Scheme_Package.xsd</a></span
  ></p>
</div>
```

```
<div class="operation" id="16">
<code class="label"> 16. GET All Classification Schemes  in Catalogue n</code><br/>
<p> To make a query to retrieve all Classification Schemes from a specific Feature Catalogue using the http
  method <span class="method">GET</span>, use the request <code class="address">
  http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/{id}/Scheme/</code
  > <br/> </p>
<p><span class="input"> The parameter: <code>id</code> - (positive integer) identifies the id of a particular
  Catalogue. A <code>/</code> with no parameter following the Scheme indicates all schemes are
  returned.</span><br/></p>
<p><span class="output">The output format for the request provides information for schemes. The response is
  supplied according to an XML schema document - <a href
  ="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/Scheme_Package.xsd">Scheme_Package.xsd</a></span
  ></p>
</div>
<div class="operation" id="17">
<code class="label"> 17. GET Responses Using Formatting Option </code><br/>
<p> To make a query to retrieve Feature Catalogue content, in a specific format, using the http method <span
  class="method">GET</span>, use the request <code class="address">
  http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/{id}/Profile/{id}/For
  mat={format}</code> <br/> </p>
<p><span class="input"> The parameters: <code>id</code> (positive integer) identifies the id of the Profile and
  <code>format</code> - (a string) identifies the format of the output file. The format option is specified at the
  end of the service request. Currently XML and html are options. Specifying no format provides XML output by
  default.</span><br/></p>
<p><span class="output">The output format for the sample request provides profile information. The response
  is supplied according to an XML schema document - <a href
  ="http://data.aad.gov.au/aadc/FeatureCatalogue/xsd/Feature_Profile.xsd">Feature_Profile.xsd
  </a></span></p>
</div>
</div>
</body>
</html>
```

Appendix 13 – Feature_Type_Attribute_Collection (Collection Format)


```xml
<?xml version="1.0" encoding="UTF-8"?>
<rdf:RDF
   xmlns:fcat="http://data.aad.gov.au/aadc/FeatureCatalogue/fcatskos-extension#"
   xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
   xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
   xmlns:dc="http://purl.org/dc/elements/1.1/"
   xmlns:skos="http://www.w3.org/2004/02/skos/core#">

<!-- Declaration of a Feature_Type_Attribute_Collection (i.e. A feature type and its spatio_temporal_geometry. This is a convenience concept established so that I can distinguish between Feature Types that have the same name but different collections of attributes for a particular geometry)-->
    <fcat:Feature_Type_Attribute_Collection rdf:about="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/1/Format=skos">
      <skos:externalID>1</skos:externalID>
      <skos:member>
        <fcat:Feature_Type rdf:about="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/1/FeatureType/5/Format=skos">
        </fcat:Feature_Type>
      </skos:member>
      <skos:member>
        <fcat:Geometry rdf:about="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/1/FeatureType/5/Geometry/4/Format=skos">
          <skos:prefLabel>Polygon</skos:prefLabel>
          <dc:coverage>2D</dc:coverage>
          <skos:externalID>5</skos:externalID>
        </fcat:Geometry>
      </skos:member>
      <skos:member>
        <fcat:Symbol rdf:about="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1Feature_Type_Attribute_Collection/1/FeatureType/5/Symbol/4/Format=skos"/>
      </skos:member>
      <skos:member>
        <fcat:Attribute rdf:about="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/1/FeatureType/5/Attribute/3/Format=skos"/>
      </skos:member>
      <skos:member>
        <fcat:Attribute rdf:about="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/1/FeatureType/5/Attribute/4/Format=skos"/>
      </skos:member>
      <skos:member>
        <fcat:Attribute rdf:about="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/1/FeatureType/5/Attribute/5/Format=skos"/>
      </skos:member>
      <skos:member>
        <fcat:Attribute rdf:about="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/1/FeatureType/5/Attribute/12/Format=skos"/>
      </skos:member>
      <skos:member>
        <fcat:Attribute rdf:about="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/1/FeatureType/5/Attribute/27/Format=skos"/>
      </skos:member>
```

<skos:member>
<fcat:Theme rdf:about="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/1/FeatureType/5/Theme/3/Format=skos"/>
</skos:member>
<fcat:memberOf rdf:resource="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Format=skos"/>
</fcat:Feature_Type_Attribute_Collection>
</rdf:RDF>

Appendix 14 – Feature Type (Concept Format)


```xml
<?xml version="1.0" encoding="UTF-8"?>
<rdf:RDF
   xmlns:fcat="http://data.aad.gov.au/aadc/FeatureCatalogue/fcatskos-extension#"
   xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
   xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
   xmlns:dc="http://purl.org/dc/elements/1.1/"
   xmlns:skos="http://www.w3.org/2004/02/skos/core#">

<!-- Declaration of a Feature_Type (i.e. A Feature Type and its spatio_temporal_geometry, all of its attributes with respect to that geometry type and any relations that are linked to this specific Feature Type)-->
  <fcat:Feature_Type rdf:about="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/1/FeatureType/5/Format=skos">
    <skos:prefLabel>Ice_Field</skos:prefLabel>
    <skos:externalID>5</skos:externalID>
    <skos:definition>A big pack of compacted snow (ice)</skos:definition>
<!-- Declaration that the Feature Type has relationships with Attributes, Theme, Geometry and relationship concepts -->
    <fcat:hasGeometry rdf:resource="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/1/FeatureType/5/Geometry/4/Format=skos"/>
    <fcat:hasSymbol rdf:resource="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1Feature_Type_Attribute_Collection/1/FeatureType/5/Symbol/4/Format=skos"/>
    <fcat:hasAttribute rdf:resource="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/1/FeatureType/5/Attribute/3/Format=skos"/>
    <fcat:hasAttribute rdf:resource="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/1/FeatureType/5/Attribute/4/Format=skos"/>
    <fcat:hasAttribute rdf:resource="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/1/FeatureType/5/Attribute/5/Format=skos"/>
    <fcat:hasAttribute rdf:resource="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/1/FeatureType/5/Attribute/12/Format=skos"/>
    <fcat:hasAttribute rdf:resource="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/1/FeatureType/5/Attribute/27/Format=skos"/>
    <fcat:hasTheme  rdf:resource="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/1/FeatureType/5/Theme/3/Format=skos"/>
<!-- Declaration that the Feature Type has an "exactMatch" relationship with another concept (i.e. Feature Type 7) ->
     <skos:exactMatch rdf:resource="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/1/FeatureType/7/Format=skos"/>
    <fcat:memberOf rdf:resource="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Format=skos"/
</fcat:Feature_Type>
</rdf:RDF>
```

Appendix 15 – SKOS Streamlined Service Pattern

```xml
<?xml version="1.0" encoding="UTF-8"?>
<rdf:RDF
   xmlns:fcat="http://data.aad.gov.au/aadc/FeatureCatalogue/fcatskos-extension#"
   xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
   xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
   xmlns:dc="http://purl.org/dc/elements/1.1/"
   xmlns:skos="http://www.w3.org/2004/02/skos/core#"
   xmlns:v="http://www.w3.org/2006/vcard/ns#">

<!-- Declaration of the UoMs and Datums used by the Feature_Type_Attribute_Collections below-->
    <fcat:Attribute_Type rdf:about="http://data-
    dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Co
    llection/5_1/FeatureType/5/Attribute/3/Attribute_Type/4/Format=skos">
      <skos:prefLabel>Basic_Spatial_Location_Quality</skos:prefLabel>
      <skos:externalID>4</skos:externalID>
  </fcat:Attribute_Type>
    <fcat:Datum rdf:about="http://data-
    dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Co
    llection/5_1/FeatureType/5/Attribute/5/UoM/4/Datum/1/Format=skos">
    <skos:prefLabel>WGS 84</skos:prefLabel>
    <skos:externalID>1</skos:externalID>
    <skos:definition>A spheroid for describing the earth</skos:definition>
    <dc:publisher>Kim</dc:publisher>
<!-- With dc:source this is how you would do it if the source was an ontology concept-->
<!-- If it was a textual source you would do the following pattern: -->
<!-- <dc:source>G.G. Tester. (2009) "Snow and Ice Atlas". Wiley Press, London UK.<dc:source> -->
    <dc:source rdf:resource="http://www.someontology/concept/#"/>
  </fcat:Datum>
  <fcat:UoM rdf:about="http://data-
  dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Coll
  ection/5_1/FeatureType/5/Attribute/5/UoM/4/Format=skos">
    <skos:prefLabel>Decimal Degrees</skos:prefLabel>
    <skos:altLabel>degrees</skos:altLabel>
    <skos:externalID>4</skos:externalID>
    <dc:publisher>Kim</dc:publisher>
<!-- With dc:source this is how you would do it if the source was an ontology concept-->
<!-- If it was a textual source you would do the following pattern: -->
<!-- <dc:source>G.G. Tester. (2009) "Snow and Ice Atlas". Wiley Press, London UK.<dc:source> -->
    <dc:source rdf:resource="http://www.someontology/concept/#"/>
    <fcat:hasDatum rdf:resource = "http://data-
    dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Co
    llection/5_1/FeatureType/5/Attribute/5/UoM/4/Datum/1/Format=skos"/>
  </fcat:UoM>
  <fcat:UoM rdf:about="http://data-
  dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Coll
  ection/5_1/FeatureType/5/Attribute/4/UoM/6/Format=skos">
    <skos:prefLabel>Metres</skos:prefLabel>
    <skos:altLabel>m</skos:altLabel>
    <skos:externalID>6</skos:externalID>
    <dc:publisher>Kim</dc:publisher>
<!-- With dc:source this is how you would do it if the source was an ontology concept-->
<!-- If it was a textual source you would do the following pattern: -->
<!-- <dc:source>G.G. Tester. (2009) "Snow and Ice Atlas". Wiley Press, London UK.<dc:source> -->
    <dc:source rdf:resource="http://www.someontology/concept/#"/>
  </fcat:UoM>
<!-- Declaration of the Feature Catalogue and its contents (the declarations above are bits and pieces that need to be
declared before they are assigned as resources to Attributes within concepts or collections appearing below)-->
  <fcat:Feature_Catalogue rdf:about="http://data-
  dev.aad.gov.au/aadc/FeatureCatalogue/?/StreamlinedCatalogueService/Catalogue/1/Format=skos">
  <dc:title>AADC Feature Catalogue</dc:title>
```

```xml
  <skos:externalID>1</skos:externalID>
  <dc:description>The best Feature Catalogue in Australia managed by the AADC</dc:description>
  <dc:identifier>Verison 1</dc:identifier>
  <dc:date>17 June 2009</dc:date>
  <dc:creator>Australian Antarctic Data Centre</dc:creator>
<!-- Declaration of the catalogue Feature_Type_Profile members-->
  <skos:member>
<!-- Declaration of the Feature_Type_Profile (collection)-->
    <fcat:Feature_Type_Profile rdf:about="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/StreamlinedCatalogueService/Catalogue/1/Profile/1/Format=skos">
    <skos:prefLabel>Test_profile_No1</skos:prefLabel>
    <skos:externalID>1</skos:externalID>
    <dc:description>A test profile numbered 1</dc:description>
    <dc:contributer>kim Finney</dc:contributer>
    <v:email>kim.finney@aad.gov.au</v:email>
<!-- Declaration of the Profile_User member-->
    <skos:member>
      <fcat:Profile_User rdf:about="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Profile_User/2/Format=skos">
        <skos:prefLabel>AADC</skos:prefLabel>
        <skos:externalID>2</skos:externalID>
        <dc:description>AADC represents a broad range of Antarctic data users</dc:description>
        <v:organization>Australian Antarctic Data Centre</v:organization>
        <v:fn>Ursula Harriss</v:fn>
        <v:email>Ursula.Harriss@aad.gov.au</v:email>
        </fcat:Profile_User>
      </skos:member>
      <skos:member>
<!-- Declaration of the Feature_Type_Profile_Attribute_Collection member (which is itself also a collection of members)-->

      <fcat:Feature_Type_Attribute_Collection rdf:about="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/StreamlinedCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/5_1/Format=skos">
        <skos:externalID>1</skos:externalID>
        <skos:member>
          <fcat:Geometry rdf:about="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/5_1/FeatureType/5/Geometry/1/Format=skos">
          <skos:prefLabel>Polygon</skos:prefLabel>
          <dc:coverage>2D</dc:coverage>
          <skos:externalID>1</skos:externalID>
          </fcat:Geometry>
        </skos:member>
        <skos:member>
          <fcat:Symbol rdf:about="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/5_1/FeatureType/5/Symbol/4/Format=skos">
            <skos:prefLabel>star</skos:prefLabel>
            <skos:externalID>4</skos:externalID>
            <dc:description>A blue star with a red circular centre</dc:description>
            <skos:prefSymbol rdf:resource="http:/example.com/symbols/0001.jpg"/>
            <skos:altSymbol rdf:resource="http:/example.com/symbols/0001a.jpg"/>
          </fcat:Symbol>
        </skos:member>
        <skos:member>
          <fcat:Attribute rdf:about="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/5_1/FeatureType/5/Attribute/3/Format=skos">
            <skos:prefLabel>Name</skos:prefLabel>
            <skos:externalID>3</skos:externalID>
```

```
      <fcat:hasAttributeType rdf:resource="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/5_1/FeatureType/5/Attribute/3/Attribute_Type/4/Format=skos"/>
    </fcat:Attribute>
  </skos:member>
  <skos:member>
    <fcat:Attribute rdf:about="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/5_1/FeatureType/5/Attribute/4/Format=skos">
      <skos:prefLabel>Latitude</skos:prefLabel>
      <skos:externalID>4</skos:externalID>
      <fcat:hasAttributeType rdf:resource="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/5_1/FeatureType/5/Attribute/4/Attribute_Type/4/Format=skos"/>
      <fcat:hasUoM rdf:resource="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/5_1/FeatureType/5/Attribute/4/UoM/4/Format=skos"/>
  </fcat:Attribute>
  </skos:member>
  <skos:member>
    <fcat:Attribute rdf:about="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/5_1/FeatureType/5/Attribute/5/Format=skos">
      <skos:prefLabel>Longitude</skos:prefLabel>
      <skos:externalID>5</skos:externalID>
      <fcat:hasAttributeType rdf:resource="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/5_1/FeatureType/5/Attribute/5/Attribute_Type/4/Format=skos"/>
      <fcat:hasUoM rdf:resource="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/5_1/FeatureType/5/Attribute/5/UoM/4/Format=skos"/>
    </fcat:Attribute>
  </skos:member>
  <skos:member>
    <fcat:Attribute rdf:about="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/5_1/FeatureType/5/Attribute/12/Format=skos">
      <skos:prefLabel>Ice_Field_Type</skos:prefLabel>
      <skos:externalID>12</skos:externalID>
      <fcat:hasAttributeType rdf:resource="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/5_1/FeatureType/5/Attribute/12/Attribute_Type/4/Format=skos"/>
    </fcat:Attribute>
  </skos:member>
  <skos:member>
    <fcat:Attribute rdf:about="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/5_1/FeatureType/5/Attribute/27/Format=skos">
      <skos:prefLabel>Depth_Value</skos:prefLabel>
      <skos:externalID>27</skos:externalID>
      <fcat:hasAttributeType rdf:resource="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/5_1/FeatureType/5/Attribute/27/Attribute_Type/4/Format=skos"/>
      <fcat:hasUoM rdf:resource="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/5_1/FeatureType/5/Attribute/27/UoM/6/Format=skos"/>
    </fcat:Attribute>
  </skos:member>
  <skos:member>
    <fcat:Theme rdf:about="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Classification_Scheme/1/Theme/3">
        <skos:prefLabel>Geomorphology</skos:prefLabel>
```

<skos:externalID>3</skos:externalID>
<fcat:memberOf rdf:resource="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Classification_Scheme/1/Format=skos"/>
</fcat:Theme>
</skos:member>
<skos:member>
<fcat:Feature_Type rdf:about="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/StreamlinedCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/5_1/FeatureType/5/Format=skos">
<skos:prefLabel>Ice_Field</skos:prefLabel>
<skos:externalID>5</skos:externalID>
<skos:definition>A big pack of compacted snow (ice)</skos:definition>
<dc:date>17 June 2009</dc:date>
<dc:publisher>Kim</dc:publisher>
<!-- With dc:source this is how you would do it if the source was an ontology concept-->
<!-- If it was a textual source you would do the following pattern: -->
<!-- <dc:source>G.G. Tester. (2009) "Snow and Ice Atlas". Wiley Press, London UK.<dc:source> -->
<dc:source rdf:resource="http://www.someontology/concept/#"/>
<!-- Declaration that the Feature Type has relationships with Attributes, Theme and Geometry concepts -->
<fcat:hasGeometry rdf:resource="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/5_1/FeatureType/5/Geometry/4/Format=skos"/>
<fcat:hasSymbol rdf:resource="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/5_1/FeatureType/5/Symbol/4/Format=skos"/>
<fcat:hasAttribute rdf:resource="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/5_1/FeatureType/5/Attribute/3/Format=skos"/>
<fcat:hasAttribute rdf:resource="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/5_1/FeatureType/5/Attribute/4/Format=skos"/>
<fcat:hasAttribute rdf:resource="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/5_1/FeatureType/5/Attribute/5/Format=skos"/>
<fcat:hasAttribute rdf:resource="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/5_1/FeatureType/5/Attribute/12/Format=skos"/>
<fcat:hasAttribute rdf:resource="http://data-dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Feature_Type_Attribute_Collection/5_1/FeatureType/5/Attribute/27/Format=skos"/>
<fcat:hasTheme rdf:resource="http://data.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Profile/1/Classification_Scheme/1/Theme/3"/>
<!-- Declaration that the Feature Type has a relationship with another concept (i.e. Feature Type 7)-->

<skos:exactMatch rdf:resource="http://data.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalague/1/Profile/1/Feature_Type_Attribute_Collection/7_2/Feature_Type/7"/>
</fcat:Feature_Type>
</skos:member>
</fcat:Feature_Type_Attribute_Collection>
</skos:member>
</fcat:Feature_Type_Profile>
</skos:member>
<!-- If there were more profiles they would be declared and inserted here-->
</fcat:Feature_Catalogue>
<!-- If there was more than one catalogue the next one would be declared and inserted here-->
</rdf:RDF>

Appendix 16 – SKOS Component Service Patterns

Catalogue


```xml
<?xml version="1.0" encoding="UTF-8"?>
<rdf:RDF
   xmlns:fcat="http://data.aad.gov.au/aadc/FeatureCatalogue/fcatskos-extension#"
   xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
   xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
   xmlns:skos="http://www.w3.org/2004/02/skos/core#"
   xmlns:dc="http://purl.org/dc/elements/1.1/"
   xmlns:v="http://www.w3.org/2006/vcard/ns#">

<!-- Declaration of the Feature_Catalogue Scheme-->
  <fcat:Feature_Catalogue rdf:about="http://data-
  dev.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/1/Format=skos">
    <dc:title>AADC Feature Catalogue</dc:title>
    <skos:externalID>1</skos:externalID>
    <dc:description>The best Feature Catalogue in Australia managed by the AADC</dc:description>
    <dc:identifier>Verison 1</dc:identifier>
    <dc:date>17 June 2009</dc:date>
    <dc:creator>Australian Antarctic Data Centre</dc:creator>
  </fcat:Feature_Catalogue>
<!-- If no number is provided after the "/" then the above Catalogue declaration would be repeated a number of times
depending on how many catalogues were in the database-->
</rdf:RDF>
```


Profile


```xml
<?xml version="1.0" encoding="UTF-8"?>
<rdf:RDF
   xmlns:fcat="http://data.aad.gov.au/aadc/FeatureCatalogue/fcatskos-extension#"
   xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
   xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
   xmlns:skos="http://www.w3.org/2004/02/skos/core#"
   xmlns:dc="http://purl.org/dc/elements/1.1/"
   xmlns:v="http://www.w3.org/2006/vcard/ns#">

<!-- Declaration of the Feature_Type_Profile Collection (i.e. a number of Feature_Type_Profiles)-->
  <fcat:Feature_Type_Profile rdf:about="http://data-
  dev.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/1/Profile/1/Format=skos">
      <skos:prefLabel>Test_profile_No1</skos:prefLabel>
      <skos:externalID>1</skos:externalID>
      <dc:description>A test profile numbered 1</dc:description>
      <dc:contributer>kim Finney</dc:contributer>
      <v:email>kim.finney@aad.gov.au</v:email>
      <fcat:memberOf rdf:resource="http://data-
      dev.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/1/Format=skos"/>
<!-- Declaration of the members of this Feature_Type_Profile Collection (i.e. a Profile_User and one or more
Feature_Types)-->
      <skos:member>
        <fcat:Profile_User rdf:about="http://data-
        dev.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/1/Profile/1/Profile_User/2/Form
        at=skos">
            <skos:prefLabel>AADC</skos:prefLabel>
            <skos:externalID>2</skos:externalID>
            <dc:description>AADC represents a broad range of Antarctic data users</dc:description>
            <v:organization>Australian Antarctic Data Centre</v:organization>
            <v:fn>Ursula Harriss</v:fn>
            <v:email>Ursula.Harriss@aad.gov.au</v:email>
        </fcat:Profile_User>
      </skos:member>
```

```
      <skos:member>
        <fcat:Feature_Type rdf:about="http://data-
        dev.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/1/Feature_Type_Attribute_Colle
        ction/1/FeatureType/1/Format=skos"/>
      </skos:member>
<!--This is where any additional Feature_Types would be placed if the profile contained more than one Feature_Type: as
per the following example:
  <skos:member> Note in the example below its the same Feature Type but with a different geometry and attributes. This
  is evident because the Feature_Type_Attribute_Collection number is different but the Feature Type number is the same)
    <fcat:Feature_Type rdf:resource="http://data-
    dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Feature_Type_Attribute_Collection/2
    /FeatureType/1/Format=skos"/>
  </skos:member>
  <skos:member> In this example the Feature Type itself is different.
    <fcat:Feature_Type rdf:resource="http://data-
    dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Feature_Type_Attribute_Collection/3
    /FeatureType/2/Format=skos"/>
  </skos:member> -->
  </fcat:Feature_Type_Profile>
<!--The pattern above for Feature_Type_Profile would be repeated if there was more than one profile. Additional profiles
would be inserted here-->
</rdf:RDF>
```


Feature Type


```
<?xml version="1.0" encoding="UTF-8"?>
<rdf:RDF
   xmlns:fcat="http://data.aad.gov.au/aadc/FeatureCatalogue/fcatskos-extension#"
   xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
   xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
   xmlns:dc="http://purl.org/dc/elements/1.1/"
   xmlns:skos="http://www.w3.org/2004/02/skos/core#">

<!-- Declaration of a Feature_Type (i.e. A Feature Type and its spatio_temporal_geometry, all of its attributes with respect
to that geometry type and any relations that are linked to this specific Feature Type)-->
  <fcat:Feature_Type rdf:about="http://data-
  dev.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/1/Feature_Type_Attribute_Collection/
  1/FeatureType/5/Format=skos">
    <skos:prefLabel>Ice_Field</skos:prefLabel>
    <skos:externalID>5</skos:externalID>
    <skos:definition>A big pack of compacted snow (ice)</skos:definition>
<!-- Declaration that the Feature Type has relationships with Attributes, Theme, Geometry and relationship concepts -->
    <fcat:hasGeometry rdf:resource="http://data-
    dev.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/1/Feature_Type_Attribute_Collectio
    n/1/FeatureType/5/Geometry/4/Format=skos"/>
    <fcat:hasSymbol rdf:resource="http://data-
    dev.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/1/Feature_Type_Attribute_Collectio
    n/1/FeatureType/5/Symbol/4/Format=skos"/>
    <fcat:hasAttribute rdf:resource="http://data-
    dev.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/1/Feature_Type_Attribute_Collectio
    n/1/FeatureType/5/Attribute/3/Format=skos"/>
    <fcat:hasAttribute rdf:resource="http://data-
    dev.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/1/Feature_Type_Attribute_Collectio
    n/1/FeatureType/5/Attribute/4/Format=skos"/>
    <fcat:hasAttribute rdf:resource="http://data-
    dev.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/1/Feature_Type_Attribute_Collectio
    n/1/FeatureType/5/Attribute/5/Format=skos"/>
    <fcat:hasAttribute rdf:resource="http://data-
    dev.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/1/Feature_Type_Attribute_Collectio
    n/1/FeatureType/5/Attribute/12/Format=skos"/>
    <fcat:hasAttribute rdf:resource="http://data-
    dev.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/1/Feature_Type_Attribute_Collectio
    n/1/FeatureType/5/Attribute/27/Format=skos"/>
```

```
      <fcat:hasTheme  rdf:resource="http://data-
      dev.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/1/Feature_Type_Attribute_Collectio
      n/1/FeatureType/5/Theme/3/Format=skos"/>
<!-- Declaration that the Feature Type has an "exactMatch" relationship with another concept (i.e. Feature Type 7)-->

      <skos:exactMatch rdf:resource="http://data-
      dev.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/1/Feature_Type_Attribute_Collectio
      n/1/FeatureType/7/Format=skos"/>
      <fcat:memberOf rdf:resource="http://data-
      dev.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/1/Format=skos"/>

      </fcat:Feature_Type>
</rdf:RDF>
```

Geometry

```
<?xml version="1.0" encoding="UTF-8"?>
<rdf:RDF
   xmlns:fcat="http://data.aad.gov.au/aadc/FeatureCatalogue/fcatskos-extension#"
   xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
   xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
   xmlns:skos="http://www.w3.org/2004/02/skos/core#"
   xmlns:dc="http://purl.org/dc/elements/1.1/">

  <fcat:Geometry rdf:about="http://data-
  dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Geometry/4/Format=skos">
      <skos:prefLabel>Polygon</skos:prefLabel>
      <dc:coverage>2D</dc:coverage>
      <skos:externalID>5</skos:externalID>
      <skos:definition>A shape bounded by connecting arcs.</skos:definition>
      <dc:publisher>Kim</dc:publisher>
<!-- With dc:source this is how you would do it if the source was an ontology concept-->
<!-- If it was a textual source you would do the following pattern: -->
<!-- <dc:source>G.G. Tester. (2009) "Snow and Ice Atlas". Wiley Press, London UK.<dc:source> -->
      <dc:source rdf:resource="http://www.someontology/concept/#"/>
      <fcat:memberOf rdf:resource="http://data-
      dev.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/1/Format=skos"/>
  </fcat:Geometry>
</rdf:RDF>
```

Attribute

```
<?xml version="1.0" encoding="UTF-8"?>
<rdf:RDF
   xmlns:fcat="http://data.aad.gov.au/aadc/FeatureCatalogue/fcatskos-extension#"
   xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
   xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
   xmlns:dc="http://purl.org/dc/elements/1.1/"
   xmlns:skos="http://www.w3.org/2004/02/skos/core#">

  <rdfs:Class rdf:ID="Attribute_Type4">
      <skos:externalID>4</skos:externalID>
      <skos:prefLabel>Basic_Spatial_Location_Quality</skos:prefLabel>
      <rdf:type rdf:resource="http://data.aad.gov.au/aadc/FeatureCatalogue/fcatskos-extension#Attribute_Type"/>
  </rdfs:Class>
  <fcat:Attribute rdf:about="http://data-
  dev.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/1/Attribute/5/Format=skos">
      <skos:prefLabel>Longitude</skos:prefLabel>
      <skos:externalID>5</skos:externalID>
      <skos:definition>A geographic coordinate (i.e. the east west component)</skos:definition>
      <dc:publisher>Kim</dc:publisher>
<!-- With dc:source this is how you would do it if the source was an ontology concept-->
<!-- If it was a textual source you would do the following pattern: -->
```

```
<!-- <dc:source>G.G. Tester. (2009) "Snow and Ice Atlas". Wiley Press, London UK.<dc:source> -->
    <dc:source rdf:resource="http://www.someontology/concept/#"/>
    <fcat:memberOf rdf:resource="http://data-
    dev.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/1/Format=skos"/>
    <fcat:hasAttributeType rdf:resource="#Attribute_Type4"/>
    <fcat:hasUoM rdf:resource="http://data-
    dev.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/1/UoM/2/Format=skos"/>
  </fcat:Attribute>
</rdf:RDF>
```

UoM

```
<?xml version="1.0" encoding="UTF-8"?>
<rdf:RDF
   xmlns:fcat="http://data.aad.gov.au/aadc/FeatureCatalogue/fcatskos-extension#"
   xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
   xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
   xmlns:dc="http://purl.org/dc/elements/1.1/"
   xmlns:skos="http://www.w3.org/2004/02/skos/core#">

  <fcat:UoM rdf:about="http://data-
  dev.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/1/UoM/2/Format=skos">
    <skos:prefLabel>Kilogram</skos:prefLabel>
    <skos:altLabel>kg</skos:altLabel>
    <skos:externalID>1</skos:externalID>
    <dc:publisher>Kim</dc:publisher>
<!-- With dc:source this is how you would do it if the source was an ontology concept-->
<!-- If it was a textual source you would do the following pattern: -->
<!-- <dc:source>G.G. Tester. (2009) "Snow and Ice Atlas". Wiley Press, London UK.<dc:source> -->
    <dc:source rdf:resource="http://www.someontology/concept/#"/>
    <fcat:hasDatum rdf:resource="http://data-
    dev.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/1/Datum/1/Format=skos"/>
    <fcat:memberOf rdf:resource="http://data-
    dev.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/1/Format=skos"/>
  </fcat:UoM>
</rdf:RDF>
```

Datum

```
<?xml version="1.0" encoding="UTF-8"?>
<rdf:RDF
   xmlns:fcat="http://data.aad.gov.au/aadc/FeatureCatalogue/fcatskos-extension#"
   xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
   xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
   xmlns:dc="http://purl.org/dc/elements/1.1/"
   xmlns:skos="http://www.w3.org/2004/02/skos/core#">
  <fcat:Datum rdf:about="http://data-
  dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1/Datum/1/Format=skos">
    <skos:prefLabel>WGS 84</skos:prefLabel>
    <skos:externalID>1</skos:externalID>
    <skos:definition>A spheroid for describing the earth</skos:definition>
    <dc:publisher>Kim</dc:publisher>
<!-- With dc:source this is how you would do it if the source was an ontology concept-->
<!-- If it was a textual source you would do the following pattern: -->
<!-- <dc:source>G.G. Tester. (2009) "Snow and Ice Atlas". Wiley Press, London UK.<dc:source> -->
    <dc:source rdf:resource="http://www.someontology/concept/#"/>
    <fcat:memberOf rdf:resource="http://data-
    dev.aad.gov.au/aadc/FeatureCatalogue/?/IterativeCatalogueService/Catalogue/1"/>
  </fcat:Datum>
</rdf:RDF>
```

ClassificationScheme

```
<?xml version="1.0" encoding="UTF-8"?>
<rdf:RDF
    xmlns:fcat="http://data.aad.gov.au/aadc/FeatureCatalogue/fcatskos-extension#"
    xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
    xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
    xmlns:dc="http://purl.org/dc/elements/1.1/"
    xmlns:skos="http://www.w3.org/2004/02/skos/core#">
  <fcat:Classification_Scheme rdf:about="http://data-
  dev.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/1/Classification_Scheme/1/Format=sk
  os">
    <skos:prefLabel>TestScheme 1</skos:prefLabel>
    <skos:externalID>1</skos:externalID>
    <skos:definition>A scheme I'm using to test</skos:definition>
    <dc:publisher>Kim</dc:publisher>
<!-- With dc:source this is how you would do it if the source was an ontology concept-->
<!-- If it was a textual source you would do the following pattern: -->
<!-- <dc:source>G.G. Tester. (2009) "Snow and Ice Atlas". Wiley Press, London UK.<dc:source> -->
    <dc:source rdf:resource="http://www.someontology/concept/#"/>
    <fcat:memberOf rdf:resource="http://data-
    dev.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/1/Format=skos"/>
    <skos:member>
      <fcat:Theme rdf:about=
      "http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/1/Classification_Scheme/
      1/Theme/3/Format=skos">
        <skos:prefLabel>Geomorphology</skos:prefLabel>
        <skos:externalID>3</skos:externalID>
        <skos:definition>A type of geological structure</skos:definition>
        <dc:publisher>Kim</dc:publisher>
<!-- With dc:source this is how you would do it if the source was an ontology concept-->
<!-- If it was a textual source you would do the following pattern: -->
<!-- <dc:source>G.G. Tester. (2009) "Snow and Ice Atlas". Wiley Press, London UK.<dc:source> -->
        <dc:source rdf:resource="http://www.someontology/concept/#"/>
        <skos:broader rdf:resource=
        "http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/1/Classification_Schem
        e/1Theme/4/Format=skos"/>
        <fcat:memberOf rdf:resource="http://data-
        dev.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/1/Classification_Scheme/1/For
        mat=skos"/>
      </fcat:Theme>
    </skos:member>
    <skos:member>
      <fcat:Theme rdf:about=
      "http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/1/Classification_Scheme/
      1Theme/4/Format=skos">
        <skos:prefLabel>Sedimentary Geomorphology</skos:prefLabel>
        <skos:externalID>4</skos:externalID>
        <skos:definition>A type of sedimentary geological structure</skos:definition>
        <dc:publisher>Kim</dc:publisher>
<!-- With dc:source this is how you would do it if the source was an ontology concept-->
<!-- If it was a textual source you would do the following pattern: -->
<!-- <dc:source>G.G. Tester. (2009) "Snow and Ice Atlas". Wiley Press, London UK.<dc:source> -->
        <dc:source rdf:resource="http://www.someontology/concept/#"/>
        <skos:narrower rdf:resource=
        "http://data.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/1/Classification_Schem
        e/1Theme/3/Format=skos"/>
        <fcat:memberOf rdf:resource="http://data-
        dev.aad.gov.au/aadc/FeatureCatalogue/?/ComponentCatalogueService/Catalogue/1/Classification_Scheme/1/For
        mat=skos"/>
      </fcat:Theme>
    </skos:member>
  </fcat:Classification_Scheme>
</rdf:RDF>
```

Appendix 17 - AODC JF/AODN Workshop (28-29/09/2010)

Actions Arising Discussion Paper

1.0 Background

The AODC JF Technical Committee met for two days on 28 and 29th September to workshop a range of technical issues associated with further developing the AODN infrastructure and provider services. A range of tasks were identified at the conclusion of the meeting that require organised action in order to systematically address ongoing infrastructure development. This paper, as agreed at the meeting, has been prepared to highlight and further characterise the tasks identified so that they can be adequately prioritised and resourced by both the AODN DO and the AODC JF.

2.0 Overall Infrastructure

The general infrastructure components discussed included:

- An AODN Portal
- AODN central MEST and agency MEST's
- Provider Services (WMS and WFS with the possibility of these also acting as wrappers for other types of data access services e.g. CSV, shapefiles)
- Feature Catalogue and Feature Catalogue Service
- Service Metadata Catalogue
- Symbology Catalogue and associated services

(a) Portal

It was decided that the general search paradigm to be delivered by the AODN Portal would be based around Feature Type and Parameter selection. For performance reasons portal functionality was conceived of as being divided conceptually into two parts – stage 1 and stage 2. Stage 1 functionality provided mainly for user selection of feature types and parameters with the option for spatio-temporal refinement. Stage 1 functionality would be supported mainly by content sourced from a Feature Catalogue, (service) layer metadata and provider-based WMS. Stage 2 functionality would provide user support for more detailed feature type, or parameter specific filters and data download. This stage would need to be supported by portal access to provider-based WFS (for the same feature types and parameters accessed in stage 1).

(b) Provider Services

It was agreed that these services (both WMS and WFS) would need to conform to an agreed community application schema (including the need for agreement on the format of any wrapped data payloads such as CSV etc). It was unclear what range of data payload formats should be supported (outside of plain WMS and WFS) and how they should be provided. However Geoserver 2.0.1 currently can output in GML 3.1.1, CSV, Shapefile and JSON fomats.

(c) Service Metadata Catalogue

Two (service) layer metadata records would be required for every data service provided (one for a WMS and one for a WFS). However one may be found sufficient with both service URL's embedded. A data service would be defined as delivering a single Feature Type. So two (service) metadata records are required for every feature type a provider wishes to supply. It was identified that there are currently some shortcomings with the metadata catalogue with respect to the capture and storage of metadata. The content of (service) layer metadata needs to be better defined, agreed and published (as a guide for users). Andrew Walsh has done some work on this already. It was agreed that the metadata should include details of the Catalogue used and the Feature Types and parameters that appear in the data (layer) service.

(d) Feature Catalogue and Service

An appropriately configured Feature Catalogue is required to manage the community agreed Feature Types and parameters. Ideally the catalogue should provide a service interface. Both the Portal and the community of data providers would need to access the catalogue content. The Portal would use a combination of (service) layer metadata and feature catalogue content to populate navigation menus.

(e) Symbology Catalogue and associated services

It was agreed that the symbols used by providers would be standardised per Feature Type.

3.0 Current Infrastructure Components and Practises

The following dot-points summarise components or practises that need to be changed if the community is to follow the infrastructure design identified above.

- AODN portal navigation currently uses facilities/data sources as the primary data access paradigm.
- AODN portal configuration and navigation is not automatically built according to service metadata or feature catalogue content.
- Data provider services use ad-hoc symbology.
- Data provider services use a controlled but unsupervised parameter vocabulary and do not adhere to any community application schema.
- Data is primarily delivered via WMS not WFS.
- There is no standardisation/agreement on what constitutes a dataset.
- The MCP does not cater for service metadata that includes Feature Catalogue identification and Feature Type content.

- The MEST does not include the latest version of the GCMD keywords (see workshop paper by Walsh) and the keyword widget does not currently cater for a 5 level nested vocab hierarchy.

- MEST does not include any facility to control entry of 'Data Parameters' so vocabularies used to date are uncontrolled.

4.0 Tasks/Actions Required

The tasks summarised below are either required, or strongly recommended in order to deliver the capability identified in 2.0 above. It was agreed that it was highly desirable for as many community members as is possible to be involved in infrastructure development. Engagement can be accomplished by active participation on working groups that are addressing tasks and/or by reviewing and commenting upon the work that groups are producing. This can only be achieved if people are made of aware of the context in which their piece of the puzzle fits and if working groups and developers regularly summarise the work that has been done and distribute it for comment. Regular reviews of the entire development activity should also be conducted and progress in general terms on all fronts should be communicated frequently to the Technical Committee/AODN DO (**i.e. this should be managed and run like a large project**).

(a) An AODN Portal

- Document the general design, functionality and content flows to be used by the portal and distribute to the community so that the logic of the design can be further evaluated, commented upon and verified as meeting community/user requirements. (*The activity we have done to date was a whiteboard brainstorming exercise which should now be done with additional rigour*).

- Potentially generate dummy components and services to build and test the portal and replace intermittently with real components/services as they become available.

(b) Provider Services (WMS and WFS with the possibility of these also acting as wrappers for other types of data access services e.g. CSV, shapefiles).

  - Scope out a general framework for defining AODN Community Feature Types.

  - Use data streams captured by CTD instruments as an exemplar:

    - Define a Feature Type or Feature Types for this stream of data and put it into the AADC Catalogue.

    - Agree on a simple symbol for the Feature Type(s).

    - Agree on payload types, content and formats (e.g. straight WFS or WFS wrapping something else)

    - Generate one or more application schema (instance documents plus schema) to provide access to CTD data.

    - Test Geoserver capability to generate the schema and serve sample CTD data.

(c) Feature Catalogue and Feature Catalogue Service.

- Evaluate the current Feature Catalogue capability within GeoNetwork/MEST. Establish its current fitness for use in the AODN context. Determine whether to support an enhancement of the current AADC Feature Catalogue (as demonstrated at the workshop) or support development, use and population of the GeoNetwork version.

- Populate the chosen catalogue.

(d) Service Metadata Catalogue

- Determine and publish the form, content and granularity of the (service) layer metadata that will be used by the portal to build menus and access services.

- Ensure that the MEST can hold the appropriate content and if possible harvest the appropriate information automatically from a providers OGC compliant service.

(e) Symbology Catalogue and associated Services

- Review community needs for symbology creation, use, management and publication.

- Assess existing symbology development and management tools (including the AADC Symbology Library and Editor). Make recommendations on use of most suitable tools/and or any development required to existing facilities.

(f) MEST Capability

- If required enhance the existing feature catalogue capability in the MEST to cater for extensions to ISO19115

- Upgrade the GCMD keyword picker to V6.0.0.0

- Include a ‘Data Parameter’ picker in the MEST for ISO19115-MCP metadata using CF standard or BODC P011 lists to control vocabulary.

Appendix 18 – Email To Colleagues Regarding Piloting The Screening Survey

Colleagues, Friends, Relatives,

I am just about to commence the data gathering phase of my PhD research and wondered if you would be willing to assist me by being friendly guinea pigs during a pilot of my initial screening survey ? I want to make sure before I send out invitations for participation in the survey that I haven't got any survey software glitches, typos, ambiguous or unintelligible questions, errors in attached documents etc. I realise that you may not be able to answer some of the questions because the survey will be targeting specific people who I hope can address them, but your general review will be nonetheless invaluable. If you have time it would be great if you could take the survey a couple of times as there are some obvious branching questions that, depending on your answer, should take you to different questions in the survey. It would be good to know if this branching is working correctly in the on-line survey. I automatically receive a copy of a completed survey. Please feel free to type rubbish into some of the required responses. This will help me test things like whether there is sufficient space for respondents to supply answers, whether entered text gets stored and copied to me etc.

One limitation of the commercial software I am using for the on-line survey is that you can't go back and review/change previous answers once a screen of data is committed, or return to a partially completed and saved copy of the survey. I can purchase this as an additional feature but would be very interested in your views about whether this functionality is indeed required. Would there be a significant improvement in your survey experience if this functionality was added, or would you consider any improvement in this area as marginal only ?

I hope to send the "real" surveys out by mid next week so if you are able to assist by reviewing the attached participant information and by subsequently taking the survey before this time I would be extremely grateful. Any feedback no matter how small, or seemingly insignificant is welcome (including on any aspect of the material attached). The text below will be the email sent to targeted participants.

Many thanks if you can help.

Regards, Kim

"Dear Colleague,

I am seeking your assistance in a research project that I am undertaking as part of my PhD studies on the topic of “Selection and evaluation of ontologies for re-use”. I am also an active participant in the marine and Antarctic data management communities and therefore have a professional, as well as personal interest in the type of information that you may be able to provide. I am investigating whether the practise of re-using ontologies during community-based ontology development exercises can be facilitated, and therefore increased, if there were practical and rapid methods available for selecting and evaluating existing ontologies.

I am trying to find people who have been involved in specific types of community-based ontology building activities that would be willing to complete an on-line survey and then potentially participate in either a face-to-face interview, and/or a later survey designed to validate any findings emerging from an analysis of the interview material. You have been selected because I am personally aware that you are a recognised and respected contributor in community activities that may involve ontology development, or a colleague of yours has referred me to you because of your expertise.

The initial Screening Survey that is phase 1 of this research should take less than 30 minutes to complete. If you are later selected for interview (phase 2), the expectation is for the interview to take between 1 and 1.5 hours in duration and I would make an appointment to meet with you at your convenience. If you reside outside of Australia I would organise either a phone hook-up or an interview by video. If you are selected for the validation component of the research (phase 3) this would involve a two-part survey, with each part possibly taking 30 mins to complete. The worst-case scenario is that you would need to be involved in all three phases, totalling a maximum of 3 hours of your time, over a period of 6 months.

I am anticipating that the available pool of people active in the area of community ontology-building, under the conditions of interest, is relatively small and therefore your agreement to participate in this study would be invaluable. I am hopeful that an outcome from this study will be a more practical method of ontology evaluation, ultimately leading to increased ontology re-use and hence an improvement in the interoperability of on-line scientific systems.

The initial Screening Survey can be found on-line at http://www.questionpro.com/akira/TakeSurvey?id=957173.

I have attached an information sheet of the research aims and the general methodology that will be followed for further clarification. I have also attached a Consent Form, which if you choose to participate in the study must be completed (by inserting your name and date) and mailed, faxed or emailed back to me prior to taking the on-line survey as part of the ethics process as designated by the University of Tasmania. Signatures on the forms will not be required if you email the form back with your name typed into the signature block. I will take receipt of the completed form as your consent.

The deadline for survey submission is XX October 2008.

Thank-you in advance for taking the time to read this email and its attachments and for considering taking part.

Yours sincerely,

Kim Finney

PhD student (University of Tasmania)"

Attachments (Participant Information Sheet Screening Survey.doc- is a little more background material; Consent Form Screening Survey.doc - is a consent form required for University ethics purposes).

Appendix 19-Selection and Evaluation Responses To Screening Survey

Selection Issues Raised By Experts

Expert Code	Selection Issues
RH	"We are biologists so we don't want to invent properties for peoples phone numbers. In fact we don't want to invent constructs to represent people."
WD	"We wanted to connect to the model organism databases in biology; therefore, we looked within the OBO community for an appropriate ontology to re-use."
PF	"Domain coverage, and service ontology descriptions."
JH	"Need to add extra hierarchy Levels to meet the growing use of ISO 19115 metadata standard for resources other than the traditional Geographic Information community."
LL	"Definition of vocabularies (externally-managed code list) in XML schemas Coverage of domain Compatibility with normalised skeleton (borrowed from another project) Potential impact (globally, in Australia)."
DH	"The TDWG ontologies sought to reuse structures, definitions and terms from other TDWG data standards where applicable and to reuse elements from simple models such as FOAF, Open GIS Consortium standards, etc. Note that the TDWG ontologies are still rather loose models, intended primarily to support data exchange and integration."
PM	"Acceptance of our ontology within the OBO repository."
CA	"The generalised domain of knowledge and concepts i was modelling."
JG	"availability; coherence/credibility; completeness"
RA	"Have been looking at the re-use of models for a while. The issue is whether a dependency simplifies or complicates the task at hand - re-definition complicates and is more work, but importing something which is vague or whose lifecycle is unclear is next to impossible. Competing options for re-use also complicates matters, and overlaps in otherwise suitable ontologies are very difficult to resolve consistently."
RL	"Suitability of terms, availability of definitions for terms, entity purity (e.g. avoid lists where 'helicopter' is included as a ship name'"
SC	"Relevance to problem. Use of compatible technology."
MH	"formal specifications, applicability, interoperability."

Evaluation Issues Raised By Experts

Expert Code	Evaluation Criteria
RH	"If some one else has done it it would be nice not for us to do it but this is very difficult. Far easier to make it up yourself."
WD	"We selected the ontology for re-use based on taxonomic and anatomical scope (e.g., we are building a skeletal ontology for fishes, and so wanted to re-use an ontology within that scope)."
PF	"Modularity and ease of owl:imports"
JH	"The formal development of these ontologies need to relate to a consider the existing ISO TC 211 code lists."
LL	"Ease of process into transformation chain (transition from 'raw material' to

	'ontology') Compatibility with normalised skeleton (borrowed from another project) Size and coverage of domain Existence of information enabling future tasks e.g. alignement with earlier work”
DH	“Currency within the community and more generally”
PM	“Following the structure of an ontology developed by an influential member of the OBO community, which our group (Phenoscape) wished to affiliate with”
CA	“If and how well they fitted the domain as well as the generalised concepts and relationships. I reviewed existing ontologies and took the best parts to make my ontology.”
JG	“number of terms; coverage of terms; organization of terms; clarity of definitions.”
RA	“Is it well scoped - i.e. does it mix concerns from different domains into a single monolithic package, or does its scope well match the governance mandate of the defining community? Is the governance process transparent and suitable for the local (re)usage constraints? Are there suitable extensibility mechanisms, both technically and governance-wise? Is the ontology/model well designed? does it mix meta-levels? Is it fully documented? is it formalised in a way that compliance with the formalism standards can be assured?”
RL	“Subjective analysis.”
SC	“Relevance to problem. Use of compatible technology. Familiarity with existing ontology and available expertise.”
MH	“Use of single inheritance, definitions, standard relations.”

Appendix 20 - Interview Analysis Categorised Codes

The original eight categories that were created to describe important facets of the data included:

(i) **Governance**: which encompassed all types of community governance issues,

(ii) **Structure**: covering quotes about ontology Structural Concepts, mainly from an evaluation perspective,

(iii) **Maintenance**: covering concepts raised about ontology maintenance, mainly from an evaluation perspective,

(iv) **Usability**: which addressed quotes surrounding user-centric issues, mainly from an evaluation perspective,

(v) **Functional Relevance**: covering issues which were not really structural but which pertained to an ontology's fitness for use, mainly from an evaluation perspective,

(vi) **Selection:** covering matters pertaining to ontology selection techniques,

(vii) **Communities:** encompassing material about the size, scope and capabilities of the communities covered by expert activity,

(viii) **Application**: a single class of code that was used to tag expert information covering the applications to which expert ontology development activities pertained.

Governance Code Descriptions

Code Tag	Description
em:maturity_of_community	The maturity of the community (time it has been working together, cohesion around goals and norms, maturity of products).
em:size-credibility-of-governance-body	Is the governance instituted by a single active member or is it more broadly based and pervasive ?
em:transparency_participatory(i)	Is the governance transparent and participatory for base concepts, lists and dictionaries (i:denotes considered important) ?
G:best_practise_guidance	Lots of guides and best practise examples developed to help community members. Ingredient for success.
G:clarity_of_roles	The different roles (and obligations) involved in developing, managing and governing ontologies should be clear and in evidence.
G:collaboration_tools	Use of collaboration tools, wikis subversion code

	repositories etc to bind the community together.
G:community_cohesion	The degree to which a community is cohesive and meets regularly and shares common focussed goals.
g:community_size	Size of the community having an impact on its function and success.
G:community_stable_over_time	Success based on stability of community over a long period of time and the technical material they are dealing with doesn't change quickly over time (exception being new instrumentation)
G:consensus_decision-making	Decisions are made in the community according to consensus
G:facilitation	Facilitation is key to developing good ontologies. Specific skills are required.
G:formal_decision_making_processes	Formal ratification processes are required surrounding ontology development and maintenance. Ingredient for success.
G:framework	Having a framework (of some kind) ingredient for success.
G:Fully_funded	Importance of funding for either development or governance activity. Ingredient for success.
G:gatekeeprs	Moderators for activity surrounding an ontology's development and maintenance. Ingredient for success.
G:governance_bodies	What or who is an appropriate governance body ? Important to get this right. Speculative statement (but also an ingredient for success).
G:governance_types	Distinction between different types of governance (e.g. technical vs content governance). Content determines terms and changes to terms, technical - serves the terms, versions them etc.
G:incentives	Rewards for participating in community ontology build and maintenance processes (e.g publications written up as a result of working group activities). Ingredient for success.
G:informal decision-making_processes	Where community decision-making processes are stated as being relatively informal by experts.
G:leveraged_investment	Has been able to leverage investment in other projects to subsidise work on ontologies.
G:light_governance	Formal but light governance that has maximum flexibility - (was ill defined so not sure what this exactly is).
g:limit_mandate	Set boundaries around what the community has within its mandate cognisant of the mandate of others. Ingredient for success.
G:long-term_investment_cycle	Choose to model data used in long-term studies and to support legislation, regulation with long term investment cycle. As opposed to short term bursts of investment and bursts of demand for the data. Ingredient for success.
G:majority_rules	Decision-making processes in community is by majority rule and views of the invested and active (as opposed to consensus).
G:OGC_as_broker	Use of a well-known standards player as a broker for domain harmonisation issues.
G:ontology_reviews	A community-based governance activity that vets submitted ontologies. Ingredient for success.

G:policies	Bind ontology (usage) policies as a contract with the delivery of ontologies – particularly useful for provenance issues. Ingredient for success.
G:project-based_approach	Treat ontology buid tasks more as project-based approach rather than a broad community based approach. Ingredient for success.
G:public_access_to_community_work	Speculative issue about whether more open, less moderated approaches to ontology build/update processes result in a larger user base than a more tightly controlled community governance approach.
g:re-using_and_forking	Forking and managing derivative because of poor governance of original ontology. Negative outcome of poor governance.
G:rivalry_between_Communites	Apparent rivalry between communities in the same or similar domain. Negative impact on productivity and success.
G:role_leader	Taken a lead role
G:role_mainatiner	Played role of maintainer
G:role_policeman	Ontology policy. Polices rules_principles
G:role_editor	Edito role like a curator
G:role_curator	Curator role- type of maintainer -editor
G:role_advisor	Advises – specialist expertise
G:role_gatekeeper	Is a curator and/or a committer
G:role_initiator	Doesn't just lead but also initiated community or activity
G:role_committer	Commits edits, updates and adds to ontology base. Trusted.
G:small_development_group	Form small development groups to undertake ontology build tasks. Ingredient for success.
g:stable_core_group_contributing	Community has a stable core group contributing to governance and development. Ingredient for success.
G:strong_institutional_backing	Have one or more institutional backers. Ingredient for success.
G:trust	Trust is important. Trust in competence of developers and in community membership (to do the right thing). Ingredients for success.

Structure Code Descriptions

Code Tag	Description
em:compliance_with_OBO _Foundry principles	For ontologies that are developed under the OBO framework – are they compliant ?
em:conceptual_models(i)	Quality of access to conceptual models (considered important).
em:encoding_efficiency(li)	Encoding efficiency - ability to deliver complex and potentially voluminous data in compact structures (considered less important).
em:engineering_evaluation(i)	Importing the ontology into a tool and visually inspect the ontology (is considered important).
em:extensibility(i)	Extensibility of an ontology in terms of ability to easily add new concepts or specialise existing ones (considered important).
em:extensibility(li)	Extensibility of ontology in terms of ability to easily add new concepts or specialise existing ones (considered less important)

em:instance_samples(i)	Quality of access to instance examples (considered important).
em:instance_samples(li)	Quality of access to instance examples (considered less important).
em:know_authors(i)	Better to know the ontology author and their reputation and whether or not it was developed in a team environment.
em:language_encoding(i)	Conformance with language encoding principles and rules (considered important).
em:logically-consistent(i)	Is the ontology logically consistent (e.g. is_a completeness, use of multiple inheritance, use of plurals, use of conjunctions).
em:maturity_of_ontology	How mature the ontology is in terms of the number of revisions it has gone through (linked possibly to size of user base and fitness for purpose ?).
em:modularity(i)	Modularity of ontology for the purposes of re-usability (considered important).
em:re-uses_ontologies(i)	If an ontology already re-uses other ontologies (is considered important)
em:relationships-that-define-context-and give-meaning	Where the relationships between concepts actually help provide semantic meaning (through the relationship context).
em:transparent_dependency(i)	Is it obvious whether the ontology has any critical dependencies on other ontologies (considered important) ?
em:well-defined(i)	Does the ontology contain well-defined concepts (considered important) ?
em:well-defined(li)	Does the ontology contain well-defined concepts ? (considered less important)
em:XSD(i)	Quality of access to XSDs (considered important).
em:XSD(li)	Quality of access to XSDs (considered less important).

Maintenance Code Descriptions

Code Tag	Description
em:actively_worked_on(i)	Is the ontology actively being worked on (is considered important) ?
em:documentation(i)	Quality of access to documentation (considered important).
em:gatekeepers_editors(i)	Lack of ontology gatekeepers or editors with the mandate to manage and moderate ontology editing leads to poorly structured ontology (considered important).
em:help-desk(i)	Is the ontology supported by an active help-desk team ?
em:maintained_plus_application(i)	Where an ontology is actively maintained but also has an application built around it (is considered important). Possibly more of an ingredient for success (?).
em:maintenance_base(i)	What is the maintenance base i.e in terms of people (skills, numbers) maintaining the ontology (considered important).
em:versioning(e)	Example of how poor versioning affects ontology quality.
em:versioning(li)	How often does the custodian release new versions and are old versions maintained and accessible ? (considered less important).
em:versions(i)	How often does the custodian release new versions and are old versions maintained and accessible ? (considered important).

Usability Code Descriptions

Code Tag	Description
em:complexity(i)	Complexity in terms of a user's ability to fit their own instance data into the ontological model, with particular concerns about the level of expertise required (considered important).
em:complexity(li)	Complexity in terms of a user's ability to fit their own instance data into the ontological model, with particular concerns about the level of expertise required (considered less important).
em:data_flux(i)	What is the state of flux i.e how often are there major revisions that would affect your use of the ontology (considered important).
em:flux(li)	What is the state of flux i.e how often are there major revisions that would affect your use of the ontology (considered less important).
em:general_accessibility(i)	Is the ontology easily accessible (open, linkable, downloadable, usable) ?
em:interoperability(i)	The level of "compatibility" in the sense that you can easily take an ontology (or a component) and use it in your own (existing) ontology i.e. its degree of interoperability.
em:mainpulation_tool(i)	Assertion that (bundled ?) ontology manipulation tools help ontology uptake (considered important).
em:major_revisions_a_detraction	More than flux – actually stating that too many revisions are a negative.
em:processing_affordance (i)	Processing affordance i.e. do the patterns chosen for constructing and encoding the concept have any potential operating synergy with service software that can recognise these patterns therefore improving the scope to develop and associate re-usable data manipulation software (is considered important).
em:processing_affordance (li)	Processing affordance i.e. do the patterns chosen for constructing and encoding the concept have any potential operating synergy with service software that can recognise these patterns therefore improving the scope to develop and associate re-usable data manipulation software (is considered less important).
em:saleability_to_community(i)	Could we readily convince our community user base to use this ontology ? (considered important).
em:skilled_ontologists_involved	Skilled ontologist's involvement considered important, particularly in early development phases.
em:survivorship(li)	Will the ontology be a survivor or a one minute wonder ? (considered less important – often because it's almost impossible to judge in advance. Some caveats can however be applied)
em:user-base(i)	What is the current user implementation base (considered less important).
em:user_base (li)	Potential for the ontology to become a standard (considered less important)
em:will-become-standard(i)	Potential for the ontology to become a standard (considered important)

Functional Relevance Code Descriptions

Code Tag	Description
em:concept_coverage(i)	Actual domain concept coverage - what fraction actually exists, even though assumption is that coverage will be incomplete (considered important).
em:concept_coverage(li)	Actual domain concept coverage - what fraction actually exists, even though assumption is that coverage will be incomplete (considered less important).
em:dictionaries_lists(i)	Applicability of included dictionaries and lists (considered important).
em:focussed(i)	Is the ontology focussed on the domain and the use-case it intentionally represents ? bloated with other superfluous stuff ?
em:iso_metadata(i)	Harmonisation with ISO 19115 metadata standard (i.e what overlaps exists and is there scope to include ISO 19115 elements (considered important).
em:iso_metadata(li)	Harmonisation with ISO 19115 metadata standard (i.e what overlaps exists and is there scope to include ISO 19115 elements (considered less important).
em:scope-size-of-dictionaries-lists(i)	
em:use-case(e)	Example of the usefulness of developing use-cases that an ontology must meet.
em:use-case(i)	Can the ontology readily meet the use-case goals (considered important).
em:use-case(ic)	Can the ontology readily meet the use-case goals (a statement about important caveats regarding use-cases).

Selection Code Descriptions

Code Tag	Description
S:build_process	A quotation about the process used to select and build an ontology with a focus on the build.
S:finding_ontologies	A quotation specifically about finding existing ontologies.
S:foundational_ontologies	A quotation about the use of foundation ontologies, or an opinion about them.
S:selection_processes	A quotation about the selection process itself – may also encompass some reference to evaluation.

Communities Code Descriptions

Code Tag	Description
C:community_capability	A quotation about a community's capability to either build, maintain or govern an ontology re-use effort.
C:size_scope	A quotation about the size and scope of a particular community of interest.
C:power_by_involvement	A quotation about people wanting to only use that which they have had a hand in creating.

Sub-category Descriptions

Code Tag	Description
Maturity	Maturity of the ontology
Structural Transparency	Groups criteria that pertain to issues that make the structure of the ontology transparent.
Conformity	Conformance of the ontology with language encoding principles.
Engineering	Groups criteria pertaining to evaluating how well the ontology has been engineered.
Application Relevance	Groups criteria that evaluate how relevant the ontology is for the use-case application.
Standards Harmonisation	Is the ontology compatible with other 'standard' or 'agreed' ontologies.
Ease of Application	Groups criteria that evaluate how easy it is for a user to harness the ontology for re-use.
Sustainability	Groups criteria that evaluate the 'sustainability' or the potential longevity of the ontology.
Curation	Groups criteria that evaluate how ell the ontology is curated.
User Assistance	Groups criteria that evaluate the degree of 'assistance' available to users in using the ontology.
Framework	Groups critera that pertain to governance framework issues.
Community	Groups criteria that characetrises the ontology 'owning' community.
Behaviours	Groups criteria that evaluate how a community behaves in governing an ontology.

Methodology Code Descriptions

Code Tag	Description
Met_use-case_good	Good experiences because of well appled use-case solicitation.
Met_community	Community-centric method.
Met_academic	Academic origin method.
Met_course	Developed courses on method.
Met_review	A review process as part of method.
Met_police	Having reviewers of practise.
Met_tools	Tools - a key part of method.
Met_use_case_bad	Bad experience because of poor use-case development.
Met_students	Using students as part of method.
Met_listservers	Using listeservers for method.
Met_workshops	Workshops as key part of method.
Met_ontology_experts	Experts vs domain practitioners – using both (but for different things)
Met_abstractawaydetail	Hiding ontology language detail from domain people.
Met_no_concepts	Can't find useful ontologies easily.
Met_poor_discovery	Search tools not good.
Met_own_group	Use own group most in finding ontos and tools.
Met_confusion	Think evaluation and overall methodology are the same thing.
Met-group_size	Small groups better – more focussed.

Original Evaluation Criteria Starter Template

Dimension	Qualitative Measure
Structure	1. Conformance with language encoding principles and rules. 2. Encoding efficiency (ability to deliver complex and potentially voluminous data in compact structures). 3. Extensibility of ontology in terms of ability to easily add new concepts or specialise existing ones. 4. Modularity of ontology for re-usability.
Functional Relevance	1. Can the ontology readily meet the use-case goals? 2. Actual domain concept coverage (what fraction exists, even though assumption is that the coverage will be incomplete ?) 3. Harmonisation with ISO 19115 Metadata standard (i.e. what overlaps exist and is there scope to include ISO 19115 elements). 4. Applicability of included dictionaries and lists.
Usability	1. Complexity (in terms of user ability to model instance data using the ontology, i.e. level of expertise required). 2. Processing affordance (i.e. do the patterns chosen for constructing and encoding the concept have any potential operating synergy with service software that can recognise these patterns therefore improving scope to develop and associate re-usable data manipulation software ?). 3. Will the ontology be a “survivor” or a “one-minute-wonder”? 4. What is its current user implementation base? 5. Is it actively being worked on? 6. What is its state of flux (i.e. how often are there major revisions)? 7. Is it likely to become a standard, or parts of it be elevated to a domain ontology? 8. Could we readily convince our community-user base to use this ontology? 9. Assessment of any included ontology manipulation tools.
Maintenance	1. What is the maintenance base (i.e. in terms of people maintaining the ontology)? 2. Quality of access to XSDs. 3. Quality of access to instance samples. 4. Quality of access to conceptual models. 5. Quality of, and, access to documentation. 6. How often does the custodian release new versions? 7. Are old versions maintained and accessible?
Governance	1. Do the custodians encourage adopters to contribute to the base? 2. Is Governance transparent and participatory for the base concepts? 3. Is Governance participatory for maintained dictionaries and lists?

Ontology Evaluation Measures Code Table

Measure Code	Decsription
SM1M	The number of ontology versions that have been released to form an impression of maturity.
SM1M	Length of time the ontology has been in use
SS1M	Availability of concept diagrams showing key concepts, concept relations and properties.
SS2M	Number of “full” owl imports or (MIREOT-like, i.e., selective use of classes from external ontologies) import statements (refer Courtot et al 2011)
SS2M	The impact of import statements on the original ontology
SS3M	Is instance data available
SS5M	Assess the use of descriptive text, any inline annotation describing the origin of terms and the use of annotation.
SC1M	Assess the availability of term and property labelling conventions
SC1M	Use commercial, open or community-based ontology editors to check for

	language-specific conformance.
SC1M	Random checks of the ontology against documented community encoding conventions and any stated ontology design patterns.
SE1M	Whether inherent ontology primitives enable appropriate sub-classing to suit needs.
SE3M	Consistency checks using a reasoner including testing for unintended models through incorrect disjunction.
SE4M	Qualitatively assess the verbosity of typically encoded data used in service transactions.
SE6M	Perform a quick visual inspection of the ontology to assess whether the ontology instils a sense of confidence in its engineering after looking at aspects such as: the number of concepts, concept depth, cohesion, tangledness (i.e., number of direct subclass relation being present while simultaneously being derivable from a chain of other sub-class relations), and any redundancy of concepts and properties.
FRA1M	Assess the ontology scope via an examination of competency questions and the use-cases it purports to satisfy and the applicability of the implementation language.
FRA1M	Test the ontology against random competency questions using available query engines.
FRA2M	Inspect the ontology for missing concepts.
FRA3M	Assess applicability of scope and type of included vocabularies.
FRA4M	Qualitatively assess the ontology focus with respect to its intended use through an examination of its concepts and properties relative to the use-cases it is intended to support.
FRS1M	Does the ontology derive from an upper ontology.
FRS1M	Can the ontology be used in conjunction with other perhaps less ontologically formal standards for example through use of semantic annotation of application schema or data instance documents, or through provided conversion tools that use both the ontology and other types of resources.
UE1M	Using known terms not modelled as instances in the ontology attempt to model the data and use the model to answer a number of competency questions. Assess ease of modelling the data and ease of querying the model to answer competency questions and degree of precision in answering competency questions.
UE2M	Does the ontology contain recognisable and re-usable ontology design patterns.
UE2M	Assess whether there are existing tools already capable of manipulating these design patterns.
UE3M	Check compliance to Linked data deployment rules using Vapour and/or a URI debugger.
UE4M	Assess a Community-shared understanding of the conditions to be met for issuing a minor revision, and conditions to be met for issuing a major revision, andaverage or maximum period between minor and major revisions.
UE5M	Look for any occurrences of ontology design patterns that would be likely to generate inconsistencies. The ideal situation is to find none.
UE6M	Availability of software and User Interfaces exploiting the ontology.
US1M	Number of users.
US1M	Size and type of projects and whether they are exploratory or production in nature.
US1M	Size and commitment of governing organisations.
US2M	Number of similar projects stopping, maintaining or starting similar

	developments: (a) using the same ontology language, and/or (b) using a different modelling language
US3M	Number of projects using or intending to use the ontology using the same ontology language and using a different modelling language
US3M	Number of positive and/or negative reviews of the ontology.
MC1M	The number and skill of people maintaining the ontology.
MC2M	Number of ontology committers.
MC3M	The type of version control system, and the handling of versions through the URI scheme, and the availability and useability of older ontology versions.
MU2M	Availability and quality of term dictionaries with term descriptions, Descriptions of relations and axioms.
GF1M	Is the framework published
GF1M	An assessment of governance strength of the organisation supporting the development (e.g., is it a standards development organisation, research consortium, individual organisation with a research interest? Or with a commercial interest? Or with an operational interest?)
GC3M	Number of persons in total – chairs committers followers.
GC3M	Number of hours per week allocated to the ontology development
GC3M	Level of email traffic
GC3M	Number of face-to-face meetings per year focusing on group work.
GB1M	licensing regime
GB1M	Community membership rules: are they sympathetic with your community's participation ?
GB1M	Positive mentions and referrals in academic and commercial conferences/workshops and in training events.
GB2M	community web forum or wiki ?
GB2M	tools available to support contributions ?

Governance Framework

Category	Description
Forums	Need collaborative spaces for community to work together and for publishing material.
Roles	Certain types of roles are required to initiatie, drive maintain activity and for developing trust.
Mandate	Community should govern with a mandate (self-acclaimed or appointed).
Rules	Governing rules; behaviours; principles; practises.
People	Adequate resourcing. With different roles.
Review/Moderation	Need to manage/moderate activity for conformance/compliance and provide material to check independently for compliance.
Tools	Tooling v. important to development, update, conformance, checking.

Appendix 21-Post Interview Ontology Evaluation Measures Ranking Survey

Thank-you for agreeing to complete this survey. This questionnaire is designed to capture the relative weight you would give to a range of ontology selection and evaluation measures, grouped under the five dimensions of “Structure”; “Functional Relevance”; “Usability”; “Maintenance” and “Governance”. Each dimension has also been further divided into sub-categories. The selection and evaluation measures that I am asking you to rate have been distilled from interviewing 13 ontology/domain experts, of which you may have been one. The purpose of this portion of my research is to investigate whether there are specific evaluation measures that are routinely applied by practitioners (as opposed to academics) in selecting ontologies for re-use (in whole or in part). Further, if there are commonly applied criteria, which of those criteria carry the most weight in decision-making. I hope to be able to develop a “rapid assessment process” for ontology selection based on the outcomes of this research. The answers that you provide in this survey will allow me to analyse the data using a quantitative technique called “Analytical Hierarchical Processing” (AHP). For this process to reach an acceptably robust outcome it may be necessary for me to run the same survey past you a second time, providing you with an opportunity to modify your answers in light of the responses provided by others (i.e. the Delphi Approach). If this becomes necessary I would provide you with the range and modal score for each question in the survey (incorporated for your convenience next to each question) and a copy of your previous responses. You would only be required to enter values in the second survey for questions where you have changed your mind about the answer.

There are 87 questions in total, spread over 6 categories of comparison. Each question asks you to consider the relative importance of the pair of statements made. A pilot run through the survey took on average 30 minutes, so whilst there appears to be a large number of questions, each one can potentially be dealt with quickly.

I sincerely appreciate your participation in this survey, and possibly a follow-up version which would give you the opportunity to change your ratings. This and any necessary follow-up survey will draw my PhD research to a close and I hope to have the Thesis completed by the end of 2010 at which time I will provide you with access to the results.

If you have any questions about the survey, my research or the procedures used you can contact Kim Finney on +61 3 62 95 1004 or by email on kim.finney@aad.gov.au.

Survey Instructions

The survey walks you through a pair-wise comparison of the relative importance that you place on ontology evaluation measures when making a **decision to select an existing ontology (in whole or part) for re-use**. The AHP technique requires pair-wise comparisons of elements at each level in a constructed hierarchy. The hierarchical model constructed from information provided through the expert interviews is attached (AHPWeight_Analysis.jpg document). The individual evaluation measures associated with the model are also attached (Evaluation_Measures_post_interview_for_AHP_V2.doc).

At each question you are asked to provide a rating taken from the 9-point rating scale depicted in Table 1.

Table 1. Rating Scale*

Intensity of Importance	Definition	Explanation
1	Equal importance	Two elements contribute equally to the objective.
3	Moderate importance	Experience and judgement slightly favour one element over another.
5	Strong importance	Experience and judgement strongly favour one element over another.
7	Very strong importance	One element is favoured very strongly over another; its dominance is demonstrated in practise.
9	Extreme importance	The evidence favouring one element over another is of the highest possible order of affirmation.

*Intensities of 2, 4, 6, and 8 can be used to express intermediate values if you feel they are warranted.

Sample questions

A. Of the two structural transparency evaluation measures outlined below which one is more important in helping you make your ontology selection and by how much ?

SS1. Having access to documented conceptual models that summarise the scope and organisation of the ontology (e.g. UML, Concept Maps, XSDs), or

SS3. Whether the ontology includes instance sample data.

More Important Option	Importance Rating
SS3	5

*The answers provided in blue can be interpreted as follows: Of the two evaluation measures presented, the option coded "**SS3**" (i.e. "*Whether the ontology includes instance sample data*") is considered strongly more important in the decision-making process (with a rating of **5**) than the option coded "SS1".*

B. Overall, of the two dimensions outlined below which one is more important in helping you make your ontology selection and by how much ?

S. Structural evaluation measures (as outlined in Table 1 of the attachment)

G. Governance evaluation measures (as outlined in Table 1 of the attachment)

More Important Option	Importance Rating
S	7

*The answers provided in blue can be interpreted as follows: Of the two dimensions (or major categories) of evaluation measures, the option coded "**S**" (i.e. "*Structural evaluation measures*") is considered very strongly more important in the decision-making process (and its dominance is demonstrated in practise with a rating of **7**) over Governance measures, coded as "G".*

Questionnaire

Q1. Comparisons Within The Structural Dimension

A. Structural Transparency Comparisons

Of the following pairs of structural transparency measures outlined below which one is more important in helping you make your ontology selection and by how much ?

Q1.1

SS1. Having access to documented conceptual models that summarise the scope and organisation of the ontology (e.g. UML, Concept Maps, XSDs), or

SS2. Does the ontology already re-use other ontologies or ontology components and where it does are the dependencies clear?

More Important Option	Importance Rating

Q1.2

SS1. Having access to documented conceptual models that summarise the scope and organisation of the ontology (e.g. UML, Concept Maps, XSDs), or

SS3. Whether the ontology includes instance sample data.

More Important Option	Importance Rating

Q1.3

SS1. Having access to documented conceptual models that summarise the scope and organisation of the ontology (e.g. UML, Concept Maps, XSDs), or

SS4. The extent to which ontological relationships define context and help give meaning (or definition) to concepts (e.g. discharge is a concept that can have multiple meanings but when coupled with m^3/s tells you it's a "flow").

More Important Option	Importance Rating

Q1.4

SS1. Having access to documented conceptual models that summarise the scope and organisation of the ontology (e.g. UML, Concept Maps, XSDs), or

SS5. Whether concepts are well-defined both in terms of textual descriptions and logic.

More Important Option	Importance Rating

Q1.5

SS2. Does the ontology already re-use other ontologies or ontology components and where it does are the dependencies clear ?, or

SS3. Whether the ontology includes instance sample data.

More Important Option	Importance Rating

Q1.6

SS2. Does the ontology already re-use other ontologies or ontology components and where it does are the dependencies clear ?, or

SS4. The extent to which ontological relationships define context and help give meaning (or definition) to concepts (e.g. discharge is a concept that can have multiple meanings but when coupled with m^3/s tells you it's a "flow").

More Important Option	Importance Rating

Q1.7

SS2. Does the ontology already re-use other ontologies or ontology components and where it does are the dependencies clear, or

SS5. Whether concepts are well-defined both in terms of textual descriptions and logic.

More Important Option	Importance Rating

Q1.8

SS3. Whether the ontology includes instance sample data, or

SS4. The extent to which ontological relationships define context and help give meaning (or definition) to concepts (e.g. discharge is a concept that can have multiple meanings but when coupled with m^3/s tells you it's a "flow").

More Important Option	Importance Rating

Q1.9

SS3. Whether the ontology includes instance sample data, or

SS5. Whether concepts are well-defined both in terms of textual descriptions and logic.

More Important Option	Importance Rating

Q1.10

SS4. The extent to which ontological relationships define context and help give meaning (or definition) to concepts (e.g. discharge is a concept that can have multiple meanings but when coupled with m^3/s tells you it's a "flow"), or

SS5. Whether concepts are well-defined both in terms of textual descriptions and logic.

More Important Option	Importance Rating

B. Structural Engineering Comparisons

Of the following pairs of structural engineering measures outlined below which one is more important in helping you make your ontology selection and by how much ?

Q1.11

SE1. Whether the ontology is extensible in terms of a user's ability to easily add new concepts or specialise existing ones, or

SE2. Whether the ontology is sufficiently modular that it is easily able to be re-used.

More Important Option	Importance Rating

Q1.12

SE1. Whether the ontology is extensible in terms of a user's ability to easily add new concepts or specialise existing ones, or

SE3. Whether the ontology is logically consistent.

More Important Option	Importance Rating

Q1.13

SE1. Whether the ontology is extensible in terms of a user’s ability to easily add new concepts or specialise existing ones, or

SE4. The encoding efficiency of the ontology (i.e. its ability to deliver complex and potentially voluminous data in compact structures).

More Important Option	Importance Rating

Q1.14

SE1. Whether the ontology is extensible in terms of a user’s ability to easily add new concepts or specialise existing ones,

SE5. Does the author have credibility (in developing ontologies and/or within the domain)?

More Important Option	Importance Rating

Q1.15

SE1. Whether the ontology is extensible in terms of a user’s ability to easily add new concepts or specialise existing ones,,or

SE6. Whether a visual inspection of the ontology gives confidence after looking at aspects such as: number of concepts, cohesion, tangledness, redundancy of concepts or properties.

More Important Option	Importance Rating

Q1.16

SE2. Whether the ontology is sufficiently modular that it is easily able to be re-used, or

SE3. Whether the ontology is logically consistent.

More Important Option	Importance Rating

Q1.17

SE2. Whether the ontology is sufficiently modular that it is easily able to be re-used, or

SE4. The encoding efficiency of the ontology (i.e. its ability to deliver complex and potentially voluminous data in compact structures).

More Important Option	Importance Rating

Q1.18

SE2. Whether the ontology is sufficiently modular that it is easily able to be re-used, or

SE5. Does the author have credibility (in developing ontologies and/or within the domain)?

More Important Option	Importance Rating

Q1.19

SE2. Whether the ontology is sufficiently modular that it is easily able to be re-used, or

SE6. Whether a visual inspection of the ontology gives confidence after looking at aspects such as: number of concepts, cohesion, tangledness, redundancy of concepts or properties.

More Important Option	Importance Rating

Q1.20

SE3. Whether the ontology is logically consistent, or

SE4. The encoding efficiency of the ontology (i.e. its ability to deliver complex and potentially voluminous data in compact structures).

More Important Option	Importance Rating

Q1.21

SE3. Whether the ontology is logically consistent, or

SE5. Does the author have credibility (in developing ontologies and/or within the domain)?Does the author have credibility (in developing ontologies and/or within the domain)?

More Important Option	Importance Rating

Q1.22

SE3. Whether the ontology is logically consistent, or

SE6. Whether a visual inspection of the ontology gives confidence after looking at aspects such as: number of concepts, cohesion, tangledness, redundancy of concepts or properties.

More Important Option	Importance Rating

Q1.23

SE4. The encoding efficiency of the ontology (i.e. its ability to deliver complex and potentially voluminous data in compact structures), or

SE5. Does the author have credibility (in developing ontologies and/or within the domain)?Does the author have credibility (in developing ontologies and/or within the domain)?

More Important Option	Importance Rating

Q1.24

SE4. The encoding efficiency of the ontology (i.e. its ability to deliver complex and potentially voluminous data in compact structures), or

SE6. Whether a visual inspection of the ontology gives confidence after looking at aspects such as: number of concepts, cohesion, tangledness, redundancy of concepts or properties.

More Important Option	Importance Rating

Q1.25

SE5. Does the author have credibility (in developing ontologies and/or within the domain)?Does the author have credibility (in developing ontologies and/or within the domain)?, or

SE6. Whether a visual inspection of the ontology gives confidence after looking at aspects such as: number of concepts, cohesion, tangledness, redundancy of concepts or properties.

More Important Option	Importance Rating

C. Structural Dimension Sub-category Comparisons

Of the following pairs of structural dimension sub-categories outlined below which one is more important in helping you make your ontology selection and by how much ?

Q1.26

SM. Maturity of the ontology (e.g. in terms of how many iterations/versions it has been through), or

SS. The structural transparency of the ontology (as outlined in Table 1 of the attachment).

More Important Option	Importance Rating

Q1.27

SM. Maturity of the ontology (e.g. in terms of how many iterations/versions it has been through), or

SC. Whether the ontology conforms with language encoding principles and community development rules.

More Important Option	Importance Rating

Q1.28

SM. Maturity of the ontology (e.g. in terms of how many iterations/versions it has been through), or

SE. The structural engineering of the ontology (as outlined in Table 1 of the attachment).

More Important Option	Importance Rating

Q1.29

SS. The structural transparency of the ontology (as outlined in Table 1 of the attachment), or

SC. Whether the ontology conforms with language encoding principles and community development rules.

More Important Option	Importance Rating

Q1.30

SS. The structural transparency of the ontology (as outlined in Table 1 of the attachment), or

SE. The structural engineering of the ontology (as outlined in Table 1 of the attachment).

More Important Option	Importance Rating

Q1.31

SC. Whether the ontology conforms with language encoding principles and community development rules, or

SE. The structural engineering of the ontology (as outlined in Table 1 of the attachment).

More Important Option	Importance Rating

Q2. Comparisons Within The Functional Relevance Dimension

A. Functional Relevance Comparisons

Of the following pairs of functional relevance measures outlined below which one is more important in helping you make your ontology selection and by how much ?

Q2.1

FRA1. Whether the ontology meet the use case goals or competency questions, or

FRA2. The degree of concept coverage.

More Important Option	Importance Rating

Q2.2

FRA1. Whether the ontology meet the use case goals or competency questions, or

FRA3. The applicability of included vocabularies (e.g. scope, type and size of dictionaries and lists).

More Important Option	Importance Rating

Q2.3

FRA1. Whether the ontology meet the use case goals or competency questions, or

FRA4. How focussed the ontology is for its stated purpose (e.g. does it focus on concepts and relations for the stated application or domain, or is it bloated with terms from many different spheres of interest).

More Important Option	Importance Rating

Q2.4

FRA2. The degree of concept coverage, or

FRA3. The applicability of included vocabularies (e.g. scope, type and size of dictionaries and lists).

More Important Option	Importance Rating

Q2.5

FRA2. The degree of concept coverage, or

FRA4. How focussed the ontology is for its stated purpose (e.g. does it focus on concepts and relations for the stated application or domain, or is it bloated with terms from many different spheres of interest).

More Important Option	Importance Rating

Q2.6

FRA3. The applicability of included vocabularies (e.g. scope, type and size of dictionaries and lists).

FRA4. How focussed the ontology is for its stated purpose (e.g. does it focus on concepts and relations for the stated application or domain, or is it bloated with terms from many different spheres of interest).

More Important Option	Importance Rating

B. Functional Relevance Dimension Sub-category Comparisons

Of the following pair of functional relevance dimension sub-categories outlined below which one is more important in helping you make your ontology selection and by how much ?

Q2.7

FRA. The application relevance of the ontology (as outlined in Table 1 of the attachment), or

FRS. Harmonisation with existing (formal or de jour) standards (e.g. ISO 19115 Metadata, Dublin Core, FOAF, VCARD).

More Important Option	Importance Rating

Q3. Comparisons Within The Usability Dimension

A. Ease of Application Comparisons

Of the following pairs of ease of application measures outlined below which one is more important in helping you make your ontology selection and by how much ?

Q3.1

UE1. Complexity of ontology (in terms of a user's ability to model instance data using the ontology and the level of expertise required), or

UE2. Processing affordance (i.e. do the patterns chosen for constructing and encoding the concepts have any potential operating synergy with service software that can recognise these patterns therefore improving the scope to develop and associate reusable data manipulation software).

More Important Option	Importance Rating

Q3.2

UE1. Complexity of ontology (in terms of a user's ability to model instance data using the ontology and the level of expertise required), or

UE3. How accessible the ontology is (e.g. is it addressable via a URI, is it published in SKOS, XML, RDF or OWL)?

More Important Option	Importance Rating

Q3.3

UE1. Complexity of ontology (in terms of a user's ability to model instance data using the ontology and the level of expertise required), or

UE4. State of flux (i.e. how often are major revisions released ?)

More Important Option	Importance Rating

Q3.4

UE1. Complexity of ontology (in terms of a user's ability to model instance data using the ontology and the level of expertise required), or

UE5. Whether the ontology is interoperable with other ontologies ?

More Important Option	Importance Rating

Q3.5

UE1. Complexity of ontology (in terms of a user's ability to model instance data using the ontology and the level of expertise required), or

UE6. Are there easily accessible and mature manipulation tools associated with using or updating the ontology ?

More Important Option	Importance Rating

Q3.6

UE2. Processing affordance (i.e. do the patterns chosen for constructing and encoding the concepts have any potential operating synergy with service software that can recognise these patterns therefore improving the scope to develop and associate reusable data manipulation software), or

UE3. How accessible the ontology is (e.g. is it addressable via a URI, is it published in SKOS, XML, RDF or OWL)?

More Important Option	Importance Rating

Q3.7

UE2. Processing affordance (i.e. do the patterns chosen for constructing and encoding the concepts have any potential operating synergy with service software that can recognise these patterns therefore improving the scope to develop and associate reusable data manipulation software), or

UE4. State of flux (i.e. how often are major revisions released ?

Q3.8

More Important Option	Importance Rating

UE2. Processing affordance (i.e. do the patterns chosen for constructing and encoding the concepts have any potential operating synergy with service software that can recognise these patterns therefore improving the scope to develop and associate reusable data manipulation software), or

UE5. Whether the ontology is interoperable with other ontologies ?

More Important Option	Importance Rating

Q3.9

UE2. Processing affordance (i.e. do the patterns chosen for constructing and encoding the concepts have any potential operating synergy with service software that can recognise these patterns therefore improving the scope to develop and associate reusable data manipulation software), or

UE6. Are there easily accessible and mature manipulation tools associated with using or updating the ontology ?

More Important Option	Importance Rating

Q3.10

UE3. How accessible the ontology is (e.g. is it addressable via a URI, is it published in SKOS, XML, RDF or OWL), or

UE4. State of flux (i.e. how often are major revisions released ?

More Important Option	Importance Rating

Q3.11

UE3. How accessible the ontology is (e.g. is it addressable via a URI, is it published in SKOS, XML, RDF or OWL), or

UE5. Whether the ontology is interoperable with other ontologies ?

More Important Option	Importance Rating

Q3.12

UE3. How accessible the ontology is (e.g. is it addressable via a URI, is it published in SKOS, XML, RDF or OWL), or

UE6. Are there easily accessible and mature manipulation tools associated with using or updating the ontology ?

More Important Option	Importance Rating

Q3.13

UE4. State of flux (i.e. how often are major revisions released, or

UE5. Whether the ontology is interoperable with other ontologies ?

More Important Option	Importance Rating

Q3.14

UE4. State of flux (i.e. how often are major revisions released, or

UE6. Are there easily accessible and mature manipulation tools associated with using or updating the ontology ?

More Important Option	Importance Rating

Q3.15

UE5. Whether the ontology is interoperable with other ontologies, or

UE6. Are there easily accessible and mature manipulation tools associated with using or updating the ontology ?

More Important Option	Importance Rating

B. Sustainability Comparisons

Of the following pairs of sustainability measures outlined below which one is more important in helping you make your ontology selection and by how much ?

Q3.15

US1. The size and scope of the current user implementation base, or

US2. Could we readily convince our community to use this ontology ?

More Important Option	Importance Rating

Q3.16

US1. The size and scope of the current user implementation base, or

US3. Estimation of whether this ontology will be used and sustained into the future.

More Important Option	Importance Rating

Q3.17

US2. Could we readily convince our community to use this ontology ?

US3. Estimation of whether this ontology will be used and sustained into the future.

More Important Option	Importance Rating

C. Usability Dimension Sub-category Comparisons

Of the following pair of usability dimension sub-categories outlined below which one is more important in helping you make your ontology selection and by how much ?

Q3.18

UE. The ease of application of the ontology (as outlined in Table 1 of the attachment), or

US. The sustainability of the ontology (as outlined in Table 1 of the attachment),

More Important Option	Importance Rating

Q4. Comparisons Within The Maintenance Dimension

A. Curation Comparisons

Of the following pairs of curation measures outlined below which one is more important in helping you make your ontology selection and by how much ?

Q4.1

MC1. The maintenance base of the ontology (in terms of number and skill of people maintaining the ontology), or

MC2. Are there dedicated "gatekeepers" (or editors/curators) for the ontology?

More Important Option	Importance Rating

Q4.2

MC1. The maintenance base of the ontology (in terms of number and skill of people maintaining the ontology), or

MC3. Does the ontology have version control (e.g. are old versions maintained and accessible and is versioning apparent for the ontologies in use)?

More Important Option	Importance Rating

Q4.3

MC2. Are there dedicated "gatekeepers" (or editors/curators) for the ontology?

MC3. Does the ontology have version control (e.g. are old versions maintained and accessible and is versioning apparent for the ontologies in use)?

More Important Option	Importance Rating

B. User Assistance Comparison

Q4.4

MU1. Is there any type of help-desk associated with the ontology ?

MU2. Quality and availability of published documentation ?

More Important Option	Importance Rating

C. Maintenance Dimension Sub-category Comparison

Of the following pair of maintenance dimension sub-categories outlined below which one is more important in helping you make your ontology selection and by how much ?

Q4.5

MC. The curation of the ontology (as outlined in Table 1 of the attachment), or

MU. The user assistance for the ontology (as outlined in Table 1 of the attachment),

More Important Option	Importance Rating

Q5. Comparisons Within The Governance Dimension

A. Framework Comparisons

Of the following pairs of framework measures outlined below which one is more important in helping you make your ontology selection and by how much ?

Q5.1

GF1. Does the ontology development community have a published governance framework (e.g. Clear roles and responsibilities for participants, guiding principles, review processes, repositories, community portal or wiki), or

GF2. Is there evidence that the ontology conforms with community governance policies and principles (e.g. naming conventions, scope rules, version control) ?

More Important Option	Importance Rating

Q5.2

GF1. Does the ontology development community have a published governance framework (e.g. Clear roles and responsibilities for participants, guiding principles, review processes, repositories, community portal or wiki), or

GF3. Does the community review or moderate individual ontology developments ?

More Important Option	Importance Rating

Q5.3

GF2. Is there evidence that the ontology conforms with community governance policies and principles (e.g. naming conventions, scope rules, version control), or

GF3. Does the community review or moderate individual ontology developments ?

More Important Option	Importance Rating

B. Community Comparisons

Of the following pairs of community measures outlined below which one is more important in helping you make your ontology selection and by how much ?

Q5.4

GC1. How mature the community is in terms of ontology development and its longevity and cohesion as a community of practise, or

GC2. Whether the community is institutionally-backed (e.g. by standards bodies, government organisations, seed-funded start-ups) ?

More Important Option	Importance Rating

Q5.5

GC1. How mature the community is in terms of ontology development and its longevity and cohesion as a community of practise, or

GC3. Whether the community has a sustained core group supporting its ontology activities ?

More Important Option	Importance Rating

Q5.6

GC1. How mature the community is in terms of ontology development and its longevity and cohesion as a community of practise, or

GC4. Whether the community's mandate is obvious and well-bounded ?

More Important Option	Importance Rating

Q5.7

GC2. Whether the community is institutionally-backed (e.g. by standards bodies, government organisations, seed-funded start-ups), or

GC3. Whether the community has a sustained core group supporting its ontology activities ?

More Important Option	Importance Rating

Q5.8

GC2. Whether the community is institutionally-backed (e.g. by standards bodies, government organisations, seed-funded start-ups), or

GC4. Whether the community's mandate is obvious and well-bounded ?

More Important Option	Importance Rating

Q5.9

GC3. Whether the community has a sustained core group supporting its ontology activities, or

GC4. Whether the community's mandate is obvious and well-bounded ?

More Important Option	Importance Rating

C. Behaviours Comparisons

Of the following pairs of behaviour measures outlined below which one is more important in helping you make your ontology selection and by how much ?

Q5.10

GB1. Whether the community actively encourages open use of their ontologies, or

GB2. Whether contribution to ongoing development and maintenance is encouraged, open and facilitated ?

More Important Option	Importance Rating

Q5.11

GB1. Whether the community actively encourages open use of their ontologies, or

GB3. Whether the governance and the core group of developers engender trust and credibility.

More Important Option	Importance Rating

Q5.12

GB2. Whether contribution to ongoing development and maintenance is encouraged, open and facilitated, or

GB3. Whether the governance and the core group of developers engender trust and credibility.

More Important Option	Importance Rating

D. Governance Dimension Sub-category Comparison

Of the following pairs of governance dimension sub-categories outlined below which one is more important in helping you make your ontology selection and by how much ?

Q5.13

GF. A published governance framework for the ontology (as outlined in Table 1 of the attachment), or

GC. Community centric issues (as outlined in Table 1 of the attachment),

More Important Option	Importance Rating

Q5.14

GF. A published governance framework for the ontology (as outlined in Table 1 of the attachment), or

GB. Community behaviours (as outlined in Table 1 of the attachment),

More Important Option	Importance Rating

Q5.15

GC. Community centric issues (as outlined in Table 1 of the attachment),

GB. Community behaviours (as outlined in Table 1 of the attachment),

More Important Option	Importance Rating

Q6. Comparisons Between ALL Dimensions

Of the following pairs of dimensions (that group evaluation measures) outlined below which one is more important in helping you make your ontology selection and by how much ?

Q6.1

S. Structural evaluation measures (as outlined in Table 1 of the attachment), or

FR. Functional relevance evaluation measures (as outlined in Table 1 of the attachment),

More Important Option	Importance Rating

Q6.2

S. Structural evaluation measures (as outlined in Table 1 of the attachment), or

U. Usability evaluation measures (as outlined in Table 1 of the attachment),

More Important Option	Importance Rating

Q6.3

S. Structural evaluation measures (as outlined in Table 1 of the attachment), or

M. Maintenance evaluation measures (as outlined in Table 1 of the attachment),

More Important Option	Importance Rating

Q6.4

S. Structural evaluation measures (as outlined in Table 1 of the attachment), or

G. Governance evaluation measures (as outlined in Table 1 of the attachment),

More Important Option	Importance Rating

Q6.5

FR. Functional relevance evaluation measures (as outlined in Table 1 of the attachment),

U. Usability evaluation measures (as outlined in Table 1 of the attachment),

More Important Option	Importance Rating

Q6.6

FR. Functional relevance evaluation measures (as outlined in Table 1 of the attachment),

M. Maintenance evaluation measures (as outlined in Table 1 of the attachment),

More Important Option	Importance Rating

Q6.7

FR. Functional relevance evaluation measures (as outlined in Table 1 of the attachment),

G. Governance evaluation measures (as outlined in Table 1 of the attachment),

More Important Option	Importance Rating

Q6.8

U. Usability evaluation measures (as outlined in Table 1 of the attachment),

M. Maintenance evaluation measures (as outlined in Table 1 of the attachment),

More Important Option	Importance Rating

Q6.9

U. Usability evaluation measures (as outlined in Table 1 of the attachment),

G. Governance evaluation measures (as outlined in Table 1 of the attachment),

More Important Option	Importance Rating

Q6.10

M. Maintenance evaluation measures (as outlined in Table 1 of the attachment),

G. Governance evaluation measures (as outlined in Table 1 of the attachment),

More Important Option	Importance Rating

Appendix 22 - Ms Excel Spreadsheet Macro For Calculating Eigenvectors and Saaty Consistency Ratios

```
Sub ahp()
Dim i As Integer, n As Integer, r As Integer, c As Integer
  Dim cw As Integer, cn As Integer, ci As Integer, ce As Integer
  Dim sr As String, sn1 As String, sc As String, SN As String
  Dim lambda As Double, consistency As Double, ratio As Double
  n = Selection.Rows.Count
  m = Selection.Columns.Count
  If n <> m Then
    MsgBox ("AHP: Your selected array must be square.")
  Else
    r = Selection.Row
    c = Selection.Column
    sr = r
    sc = c
    SN = r + n
    sn1 = r + n - 1
    cw = c + n
    cn = cw + 1
    ci = cn + 1
    ce = ci + 1
    cf = ce + 1
    For i = 0 To (n - 1)
      Cells(r + i, c + i).Value = 2
    Next i
    Cells(r, cw).FormulaR1C1 = "=sum(RC" & sc & ":RC[-1])"
    Cells(r + n, cw).FormulaR1C1 = "=sum(R" & sr & "C:R[-1]C)"
    Cells(r, cn).FormulaR1C1 = "=RC[-1]/R" & SN & "C[-1]"
    Cells(r, ci).FormulaR1C1 = "=mmult(RC" & sc & ":RC[-3],R" & sr & "C[-1]:R" & sn1 & "C[-1])"
    Cells(r, ce).FormulaR1C1 = "=abs(RC[-1]-RC[-3])"
    Cells(r + n, ce).FormulaR1C1 = "=sum(R" & sr & "C:R[-1]C)"
    Range(Cells(r, cw), Cells(r + n - 1, ce)).Select
    Selection.FillDown
    i = 0
    Do While (i < 100) And (Cells(r + n, ce).Value > 0.0001)
      i = i + 1
      Range(Cells(r, ci), Cells(r + n - 1, ci)).Select
      Selection.Copy
      Range(Cells(r, cw), Cells(r + n - 1, cw)).Select
      Selection.PasteSpecial Paste:=xlValues, Operation:=xlNone, SkipBlanks:=False, Transpose:=False
    Loop
    For i = 0 To (n - 1)
      Cells(r + i, c + i).Value = 1
    Next i
    Cells(r, cf).FormulaR1C1 = "=RC[-2]/RC[-3]"
    Range(Cells(r, cf), Cells(r + n - 1, cf)).Select
    Selection.FillDown
    Cells(r + n, cf).FormulaR1C1 = "=sum(R" & sr & "C:R[-1]C)"
    lambda = Cells(r + n, cf).Value / n
    SN = n
    Cells(r + n + 1, cn - 2).Value = "Lambda max"
    Cells(r + n + 1, cn).Value = lambda
    Cells(r + n + 2, cn - 2).Value = "Consistency index"
    Cells(r + n + 2, cn).Value = (lambda - n) / (n - 1)
```

```
        Cells(r + n + 3, cn - 2).Value = "Consistency ratio"
        Cells(r + n + 3, cn).FormulaR1C1 = "=r[-1]c/(0.35+(" & SN & "-2)^2/((" & SN & "-2)^2 + 6) + 0.08 * sqrt(" &
SN & "- 2 ))"
        If r > 1 Then
            Cells(r - 1, cn).Value = "Weights"
        End If
        Range(Cells(r, cn), Cells(r + n - 1, cn)).Select
        Selection.Copy
        Selection.PasteSpecial Paste:=xlValues, Operation:=xlNone, SkipBlanks:=False, Transpose:=False
        Range(Cells(r, cw), Cells(r + n, cw)).Select
        Selection.Clear
        Range(Cells(r, ci), Cells(r + n, cf)).Select
        Selection.Clear
        Range(Cells(r, cn), Cells(r + n + 3, cn)).Select
        Selection.NumberFormat = "0.000"
    End If
End Sub
```

Appendix 23 – Saaty Table of Consistency Values From Saaty (1980)

S. No.	Size of Matrix (n)	Random Consistency Index (RI)
1	1	0
2	2	0
3	3	0.52
4	4	0.89
5	5	1.11
6	6	1.25
7	7	1.35
8	8	1.40
9	9	1.45
10	10	1.49

Appendix 24 – Example Email To Expert Requesting They Consider A Change To Their Ratings Data

Hi WD,

I've taken leave to complete my PhD and have been doing the final analysis of my survey data in prep for getting some preliminary results back out to participants and to enable commencement of the write-up.

The type of analyses I've been doing on the survey data is based on the Analytical Hierarchical Processing method (by Saaty, 1980). Part of the method is to look at the consistency (or more normally the inconsistency) in the pair-wise responses provided by participants. It is recommended that a calculated consistency ratio should not really exceed 10%. The consistency ratio cumulatively measures the degree to which respondents might inadvertently place illogical preference weights between pair-wise comparisons.

All experts who responded to my survey had varying degrees of unacceptable inconsistency (which is not unusual). So I have taken the data I was supplied with and iteratively (through a controlled process) manipulated each expert's pair-wise comparison data until the consistency ratio becomes acceptable. Because I have used a method that attempts to preserve as much of the original matrix of data as possible (for each expert) I have, in some cases, accepted a higher than 10% (in)consistency. The ultimate purpose of the pair-wise comparisons (you will recall) is to arrive at a set of weights for each criteria that experts have identified (through interview) as being important for the selection and evaluation of ontologies (in re-use scenarios).

The purpose of my email is to reflect back to you the changes that I have made with respect to your answers and to see if you are willing to accept these amendments (which brings your comparison data within an acceptable tolerance and will allow me to use it in my study). I'll assume **no response by 31st January 2011** to be a positive one (i.e. you are agreeing to the modifications suggested so that I can use the amended data in any forward analyses). I am extremely grateful for the time you have invested thus far and appreciate how busy each of the experts in this study are, which is why I have taken the approach that I have to bring the data within tolerance, rather than asking all experts to re-take portions of the survey (which would be relatively time consuming).

Of the 80 odd pair-wise comparisons I asked you to make in the survey (your original responses are attached) I had to adjust 1 of the pair-wise comparisons, trying as best I could to preserve the pattern of your responses, whilst bringing the data into tolerance. Below I list the changes I have made. I've also included a document which graphs a comparison of the weights calculated for specific criteria, both before and after (pair-wise) data adjustment. This will give you a quick visual feel for what affect the changes I have made ultimately have on the calculation of weights, as derived from your responses (which is the goal of this particular exercise).

Adjusted pair-wise comparisons:

Structural Transparency
SS1 and SS5, current value 5 adjusted to 1

Best regards, Kim

Kim Finney
Manager, Australian Antarctic Data Centre
Australian Antarctic Division
Department of Sustainability, Environment, Water, Population and Communities
+61 3 62 32 3459
http://data.aad.gov.au/

	old weights	new weights
SS1	0.400	0.265
SS2	0.127	0.127
SS3	0.053	0.058
SS4	0.147	0.171
SS5	0.273	0.378

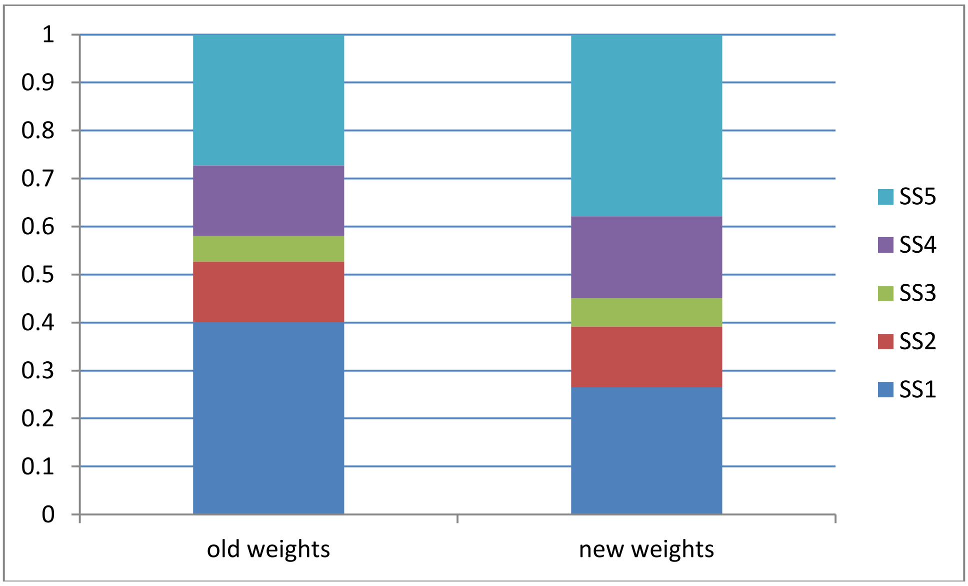

Appendix 25 – R Script For Improving Matrix Consistency Ratios

```
# run this script in R with:

# source('C:/Users/kim_fin/Documents/DocumentspostNovember2009/ahp.txt') ##adjust path if necessary

# reference CI values from random data

RI=c(0,0,0.5800,0.9000,1.1200,1.2400,1.3200,1.4100,1.4500,1.4900)

# file to read in

datafile='C:/Users/kim_fin/Documents/PhD/PhD_AHP_Weights_Survey1/AHP_Analysis/AHP_Adjusted_Compa
risons/FirstAdjustments/2ndAdjustments/LaufortforRanalysisnewasperLefort.csv' ##adjust

cat(sprintf("Data file: %s\n",datafile))

alldata=read.table(datafile,sep=",",header=F)

# iterate through each row in the file

for (li in 1:dim(alldata)[1]) {

  this_name=as.character(alldata[li,1])

  cat(sprintf("%s\n",this_name))

  # define matrix

#A=c(1,5,3,7,6,6,1/3,1/4,1/5,1,1/3,5,3,3,1/5,1/7,1/3,3,1,6,3,4,6,1/5,1/7,1/5,1/6,1,1/3,1/4,1/7,1/8,1/6,1/3,1/3,3
,1,1/2,1/5,1/6,1/6,1/3,1/4,4,2,1,1/5,1/6,3,5,1/6,7,5,5,1,1/2,4,7,5,8,6,6,2,1)

  A=alldata[li,]

  A[1]=NA #discard name

  A=A[which(!is.na(A))] #keep only the non-missing elements

  A=matrix(as.numeric(A),nrow=sqrt(length(A)),byrow=T) #convert to square matrix

  n=dim(A)[1]

  v=eigen(A)

  # max eigenvalue

  mv=v$values[1]

  # keep real part only

  mv=Re(mv)

  # calculate CI

  CI=(mv-n)/(n-1)

  CR=CI/RI[n]
```

```
cat(sprintf(' current CR of %.4f\n',CR))

# iterate through the responses and see which one gives the best improvement in CI

rind=which(A>1,arr.ind=T)

rind=cbind(rind,-1)

bestCI=CI

bestA=A

for (q in 1:dim(rind)[1]) {

	this_value=A[rind[q,1],rind[q,2]]

	thisA=A

	thisA[rind[q,1],rind[q,2]]=0

	thisA[rind[q,2],rind[q,1]]=0

	thisA[rind[q,1],rind[q,1]]=2

	thisA[rind[q,2],rind[q,2]]=2

	thisv=eigen(thisA)

	thisCI=(Re(thisv$values[1])-n)/(n-1)

	rind[q,3]=Re(thisv$values[1]) #keep track of eigenvalues

	if (thisCI<bestCI) {

	  bestCI=thisCI

	  bestA=thisA

	}

}

# sort rind[,3] and go through until 3 valid options found

sortidx=sort(rind[,3],index.return=T)

nvalid=0

ii=1

while (nvalid<3 & ii<length(sortidx$ix)) {

	 # find value that gives best CI for this entry

	 bestCR=CR

	 bestval=-1
```

```
            for (nv in 1:9) {
              thisA=A
              thisA[rind[sortidx$ix[ii],1],rind[sortidx$ix[ii],2]]=nv
              thisA[rind[sortidx$ix[ii],2],rind[sortidx$ix[ii],1]]=1/nv
              thisv=eigen(thisA)
              thisCR=(Re(thisv$values[1])-n)/(n-1)/RI[n]
              if (thisCR<bestCR) {
                  bestCR=thisCR
                  bestval=nv
              }
            }
            if (bestval != -1) {
              nvalid=nvalid+1
              cat(sprintf(' CR of %.4f by setting element [%d,%d] to value %d (current value
%d)\n',bestCR,rind[sortidx$ix[ii],1],rind[sortidx$ix[ii],2],bestval,A[rind[sortidx$ix[ii],1],rind[sortidx$ix[ii],2]]))
            }
            ii=ii+1
   }
}
```

Appendix 26 – Before and After Plots For Edited Expert Weight Data

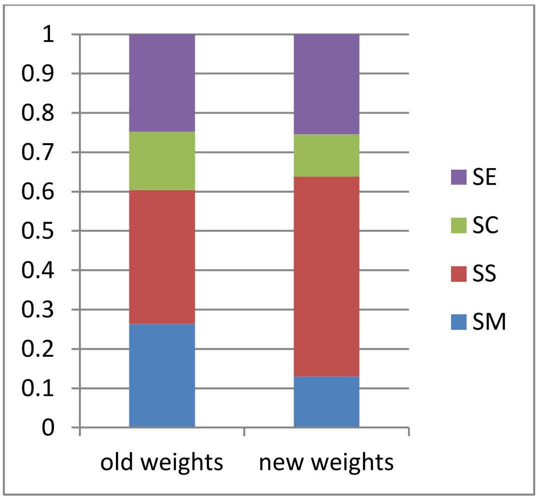

Expert KS (Structural Dimension)

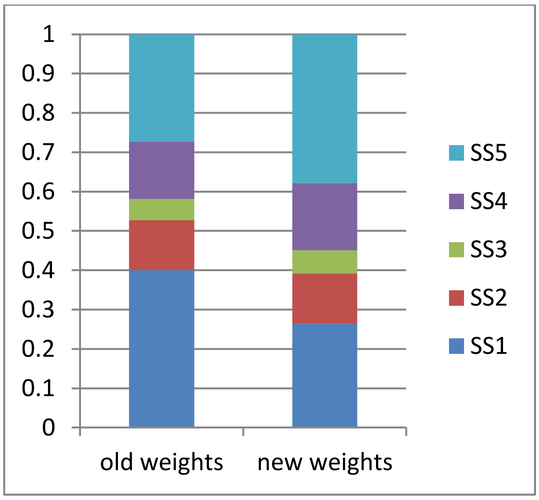

Expert WD (Structural Transparency)

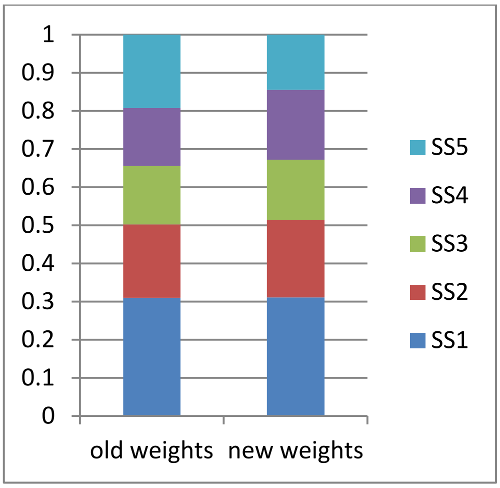

CA (Structural Transparency)

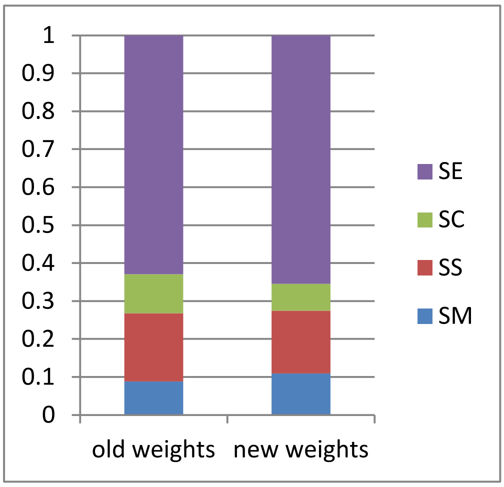

Expert CA (Structural Dimension)

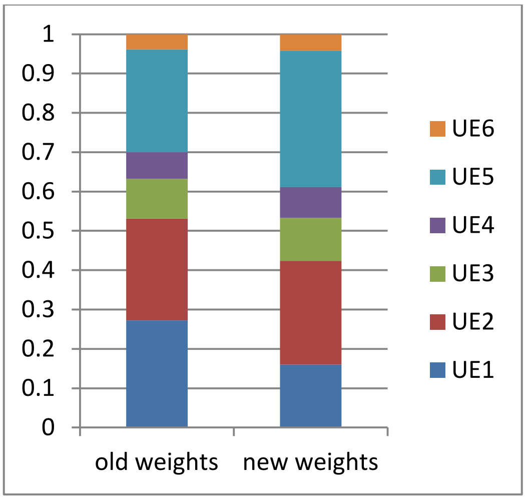

Expert CA (Ease Of Application)

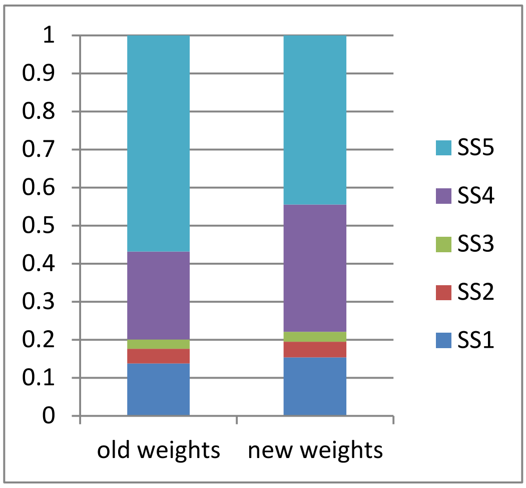

Expert DH (Structural Transparency)

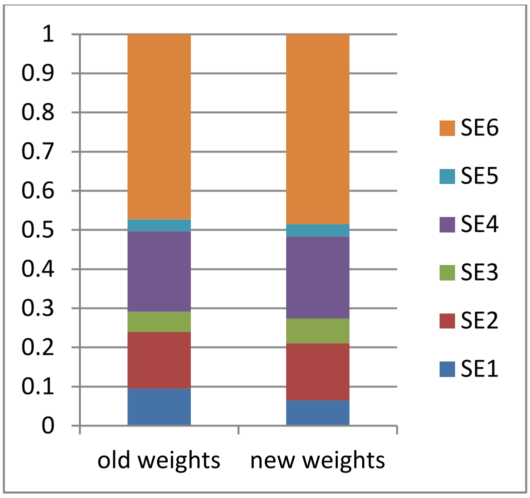

Expert DH (Structural Engineering)

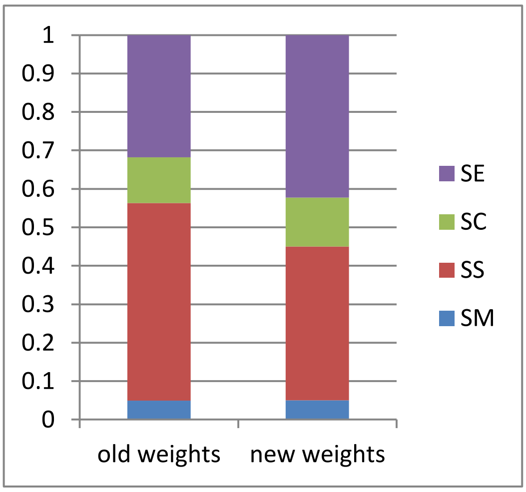

Expert DH (Structural Dimension)

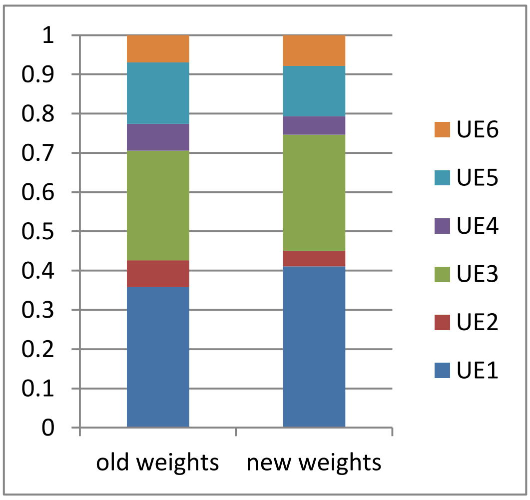

Expert DH (Ease of Application)

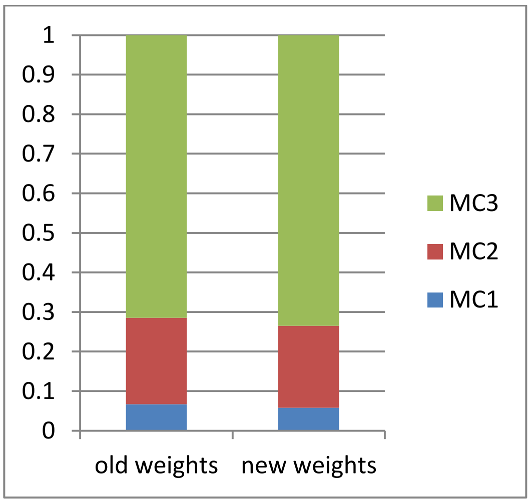

Expert DH (Curation)

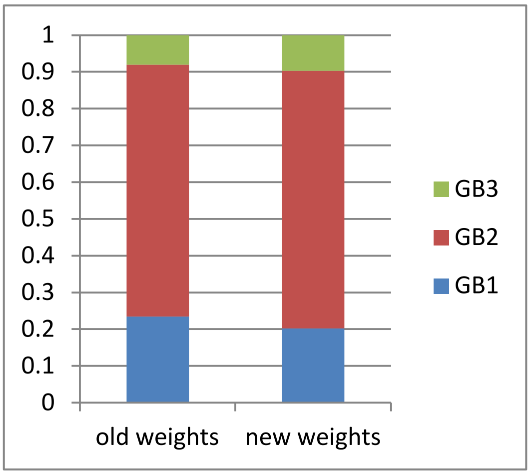

Expert DH (Behaviours)

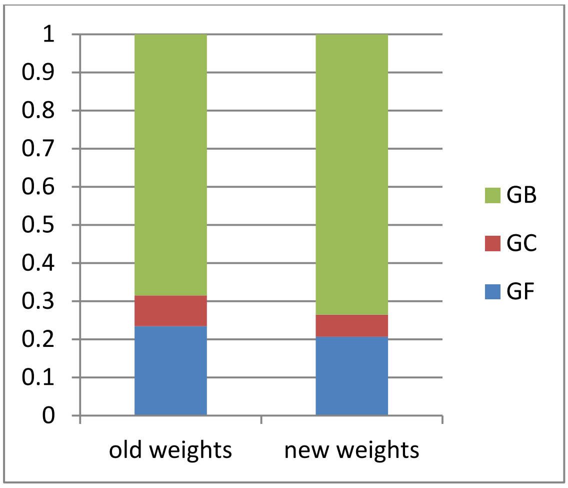

Expert DH (Governance)

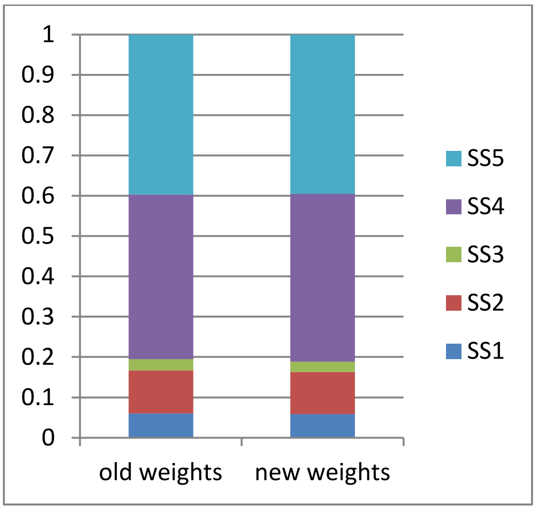

MH (Structural Transparency)

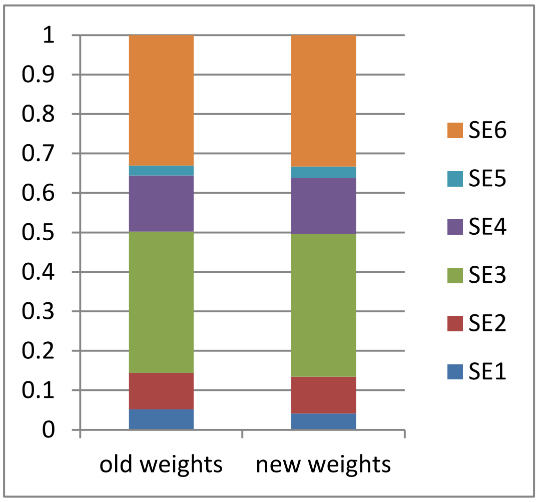

Expert MH (Structural Engineering)

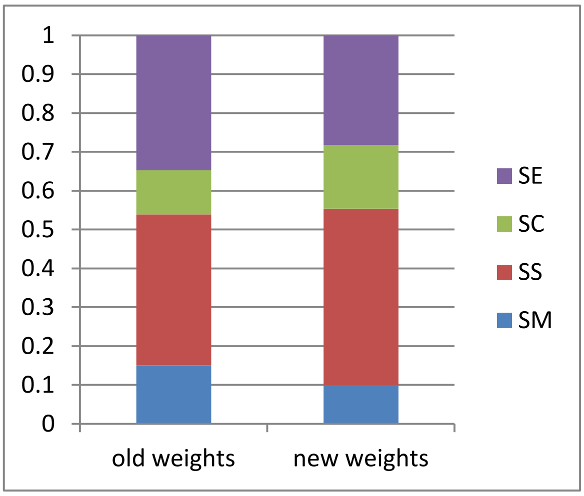

Expert MH (Structural Dimension)

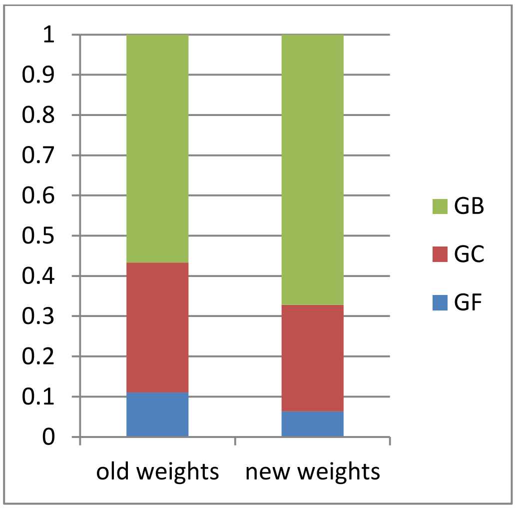

Expert MH (Governance)

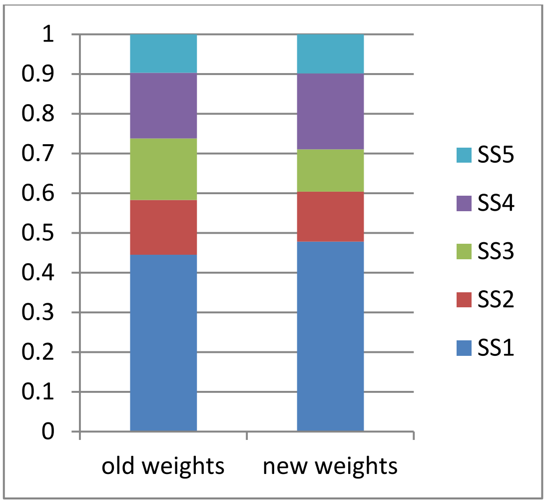

JH (Structural Transparency)

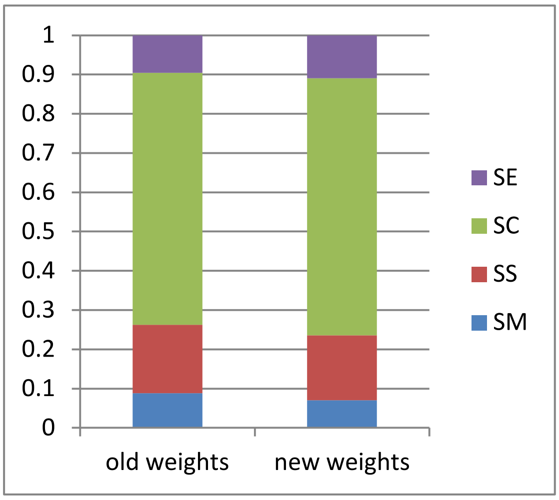

Expert JH (Structural Dimension)

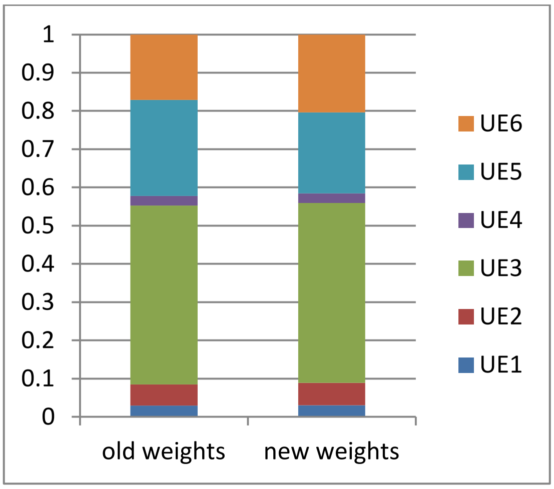

Expert JH (Ease of Application)

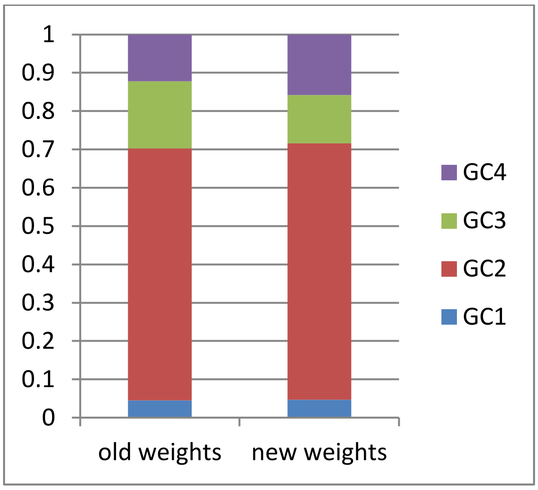

Expert JH (Community)

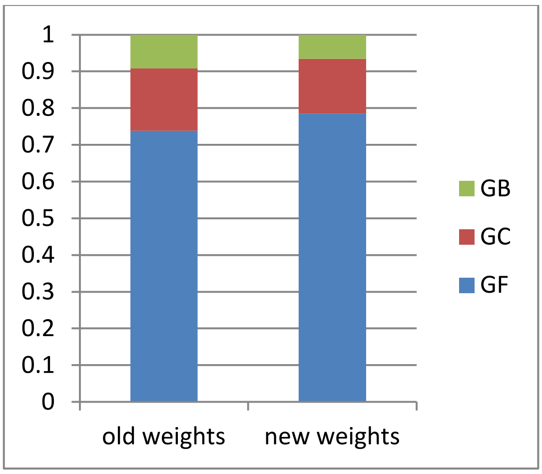

Expert JH (Governance)

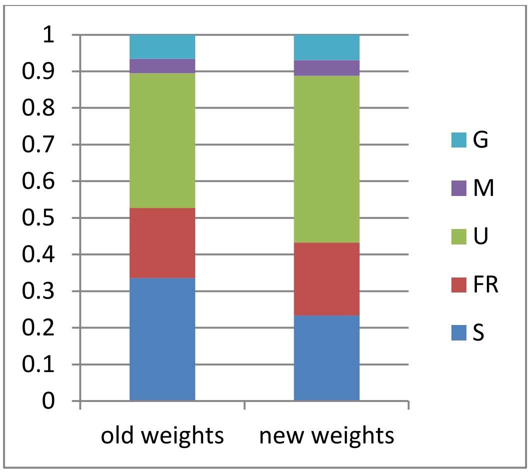

Expert JH (All Dimensions)

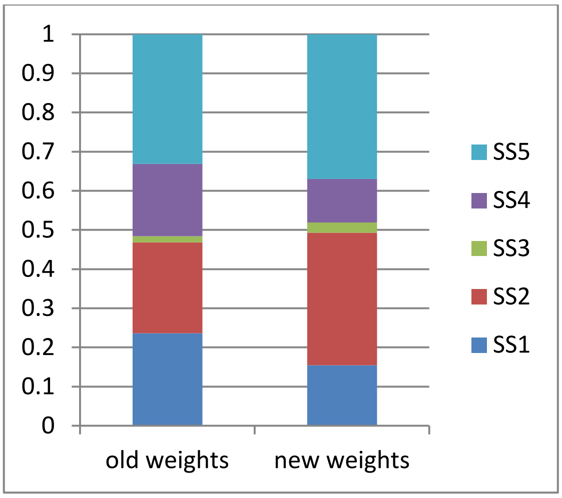

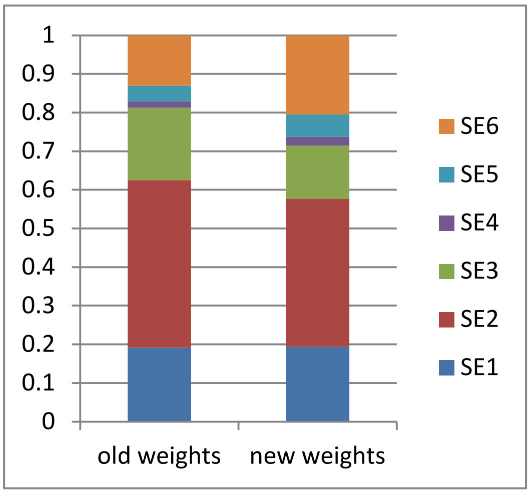

RA (Structural Transparency)

Expert RA (Structural Engineering)

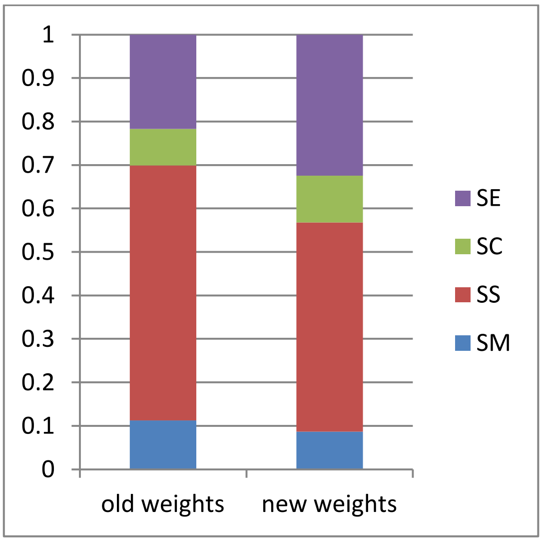

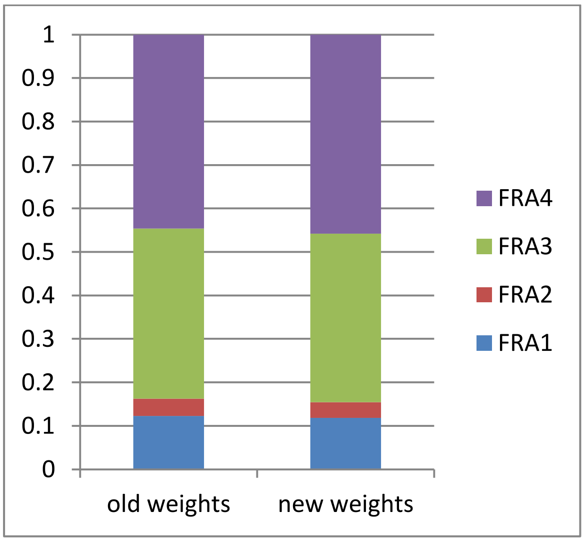

Expert RA (Structural Dimension)

Expert RA (Application Relevance)

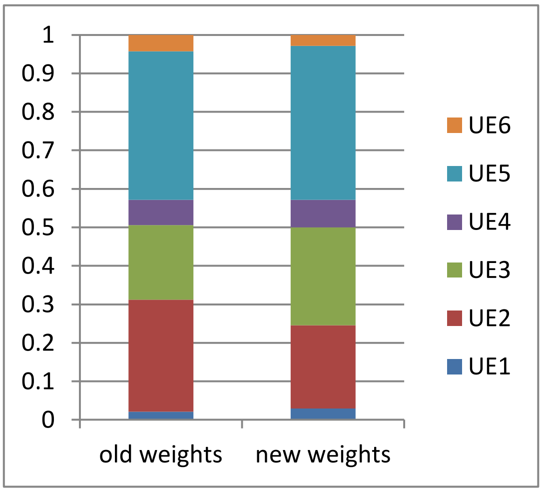

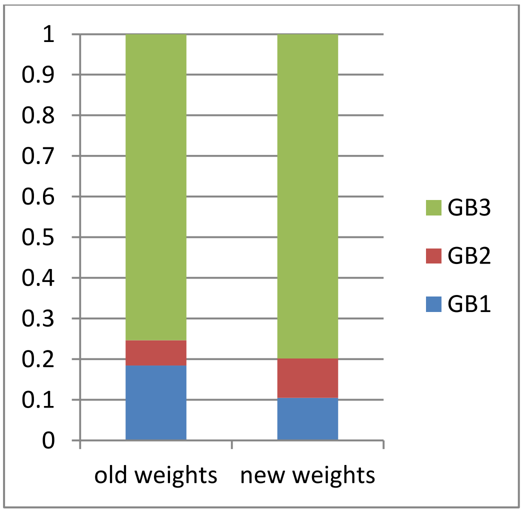

Expert RA (Ease of Application)

Expert RA (Behaviours)

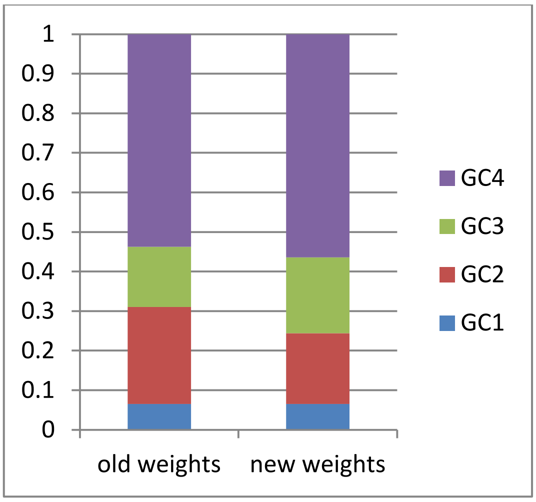

Expert RA (Community)

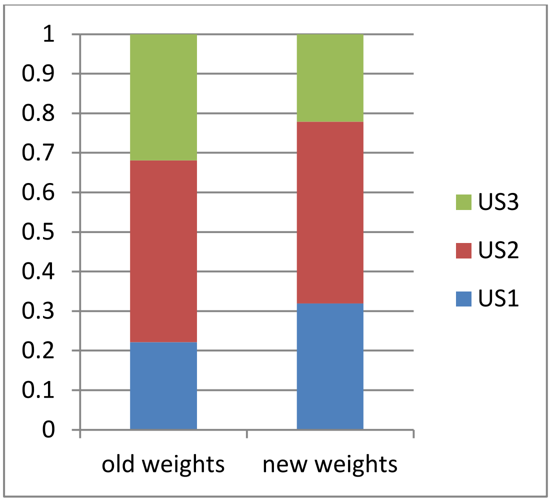

Expert RA (Sustainability)

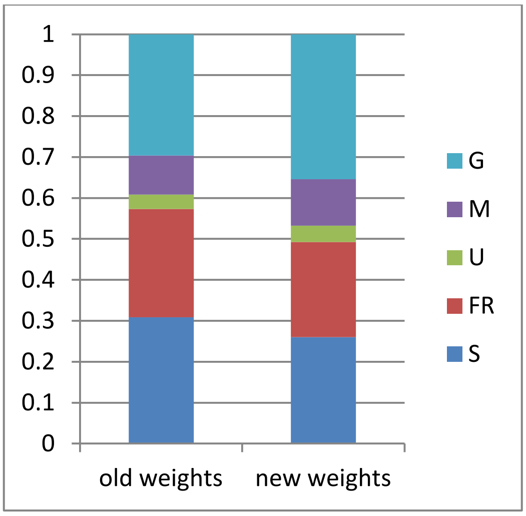

Expert RA (All Dimensions)

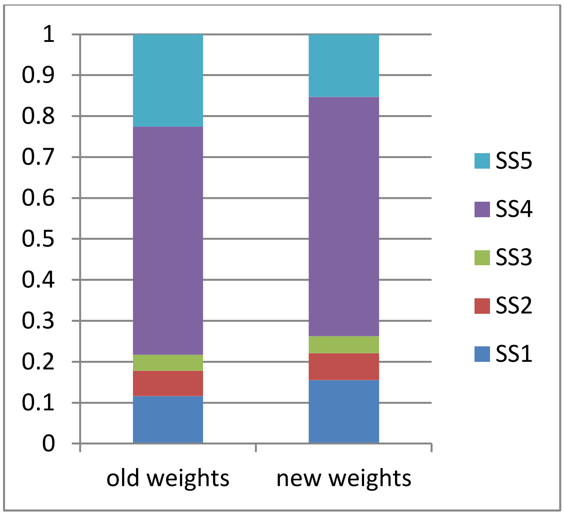

Expert LL (Structural Transparency)

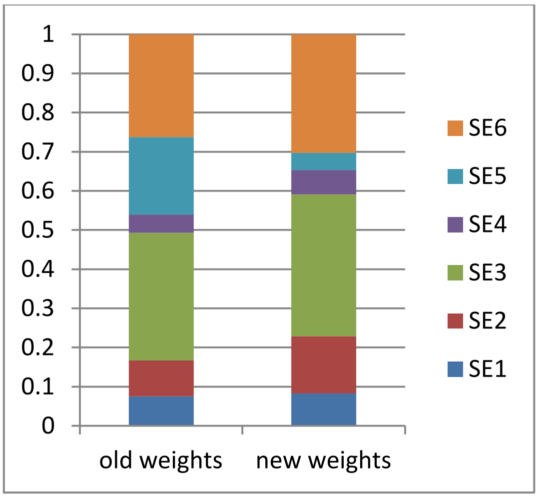

Expert LL (Structural Engineering)

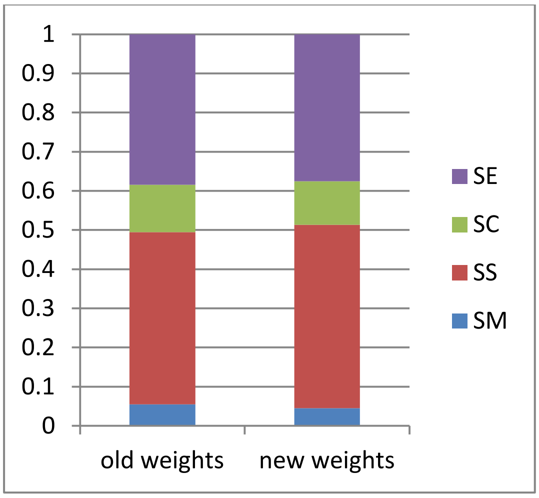

Expert LL (Structural Dimension)

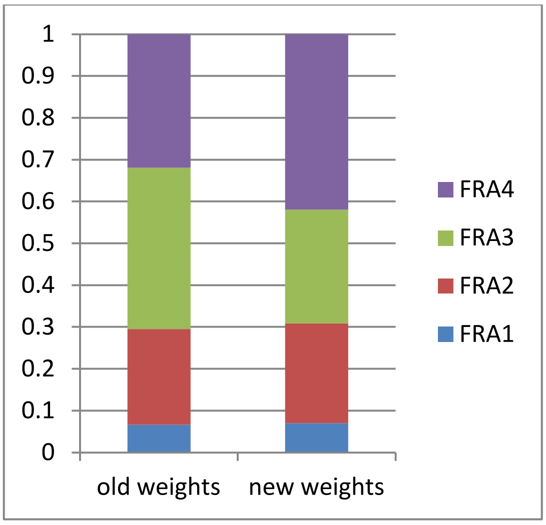

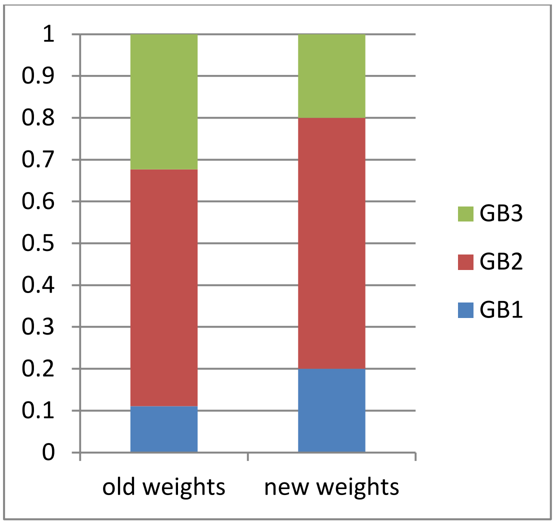

Expert LL (Application Relevance) **Expert LL (Behaviours)**

Appendix 27 – Plotted Measures of Central Tendency

(For Each Dimension & Sub-category)

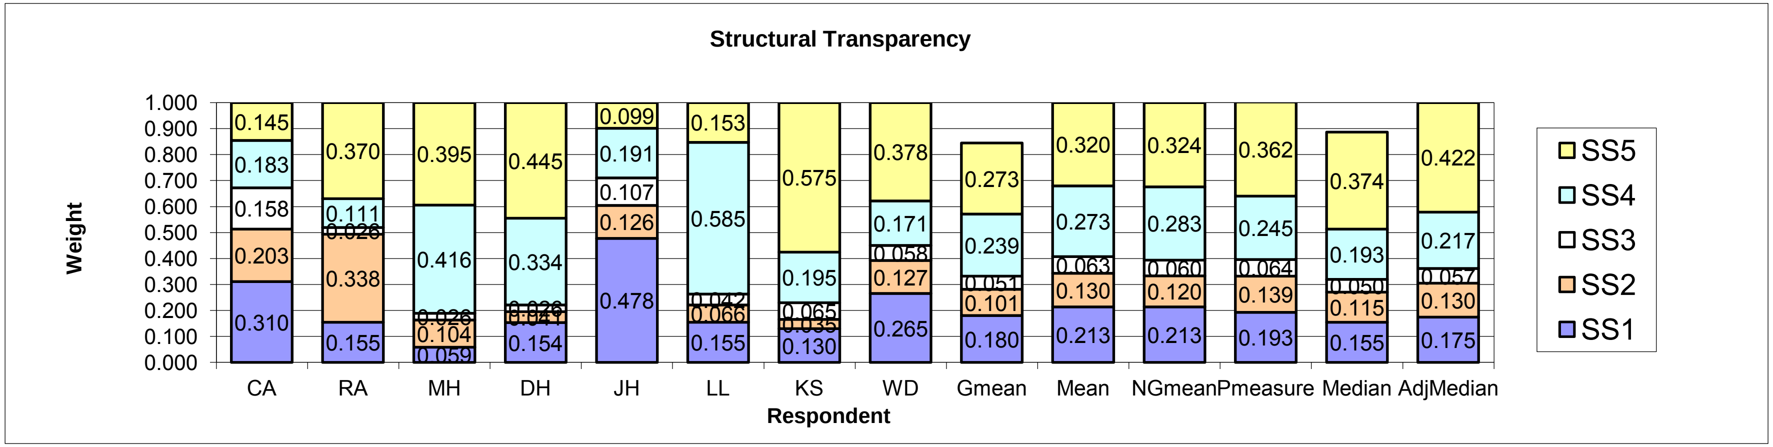

Structural Transparency Weight Data

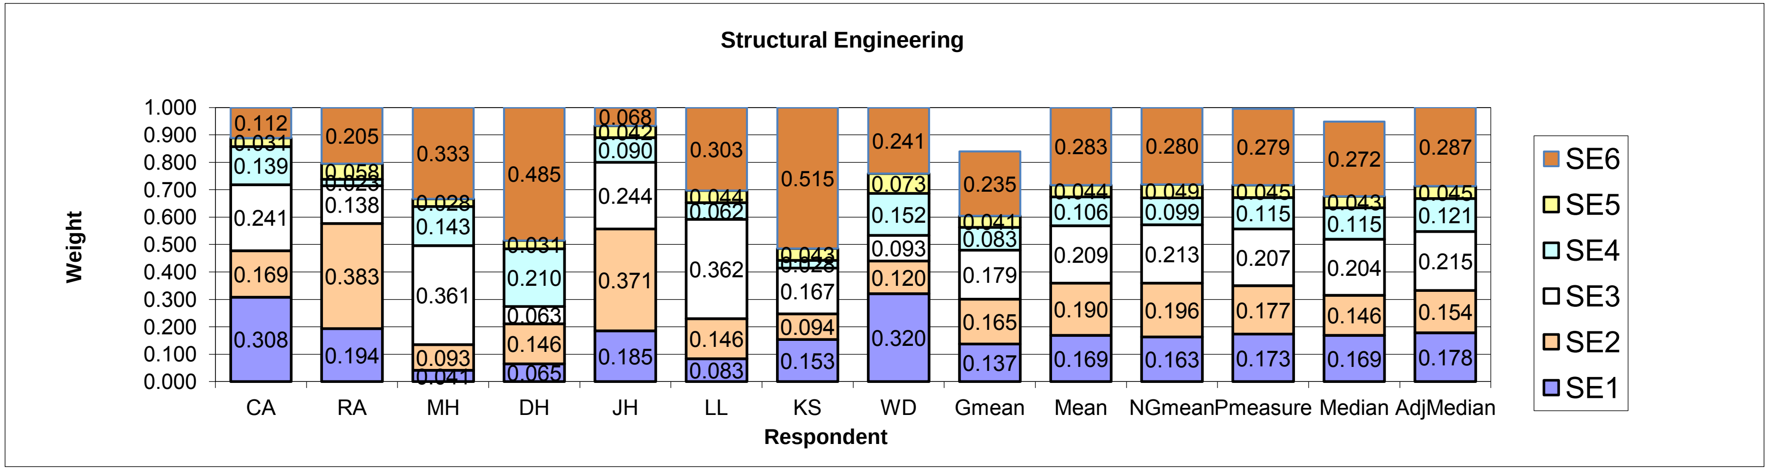

Structural Engineering Weight Data

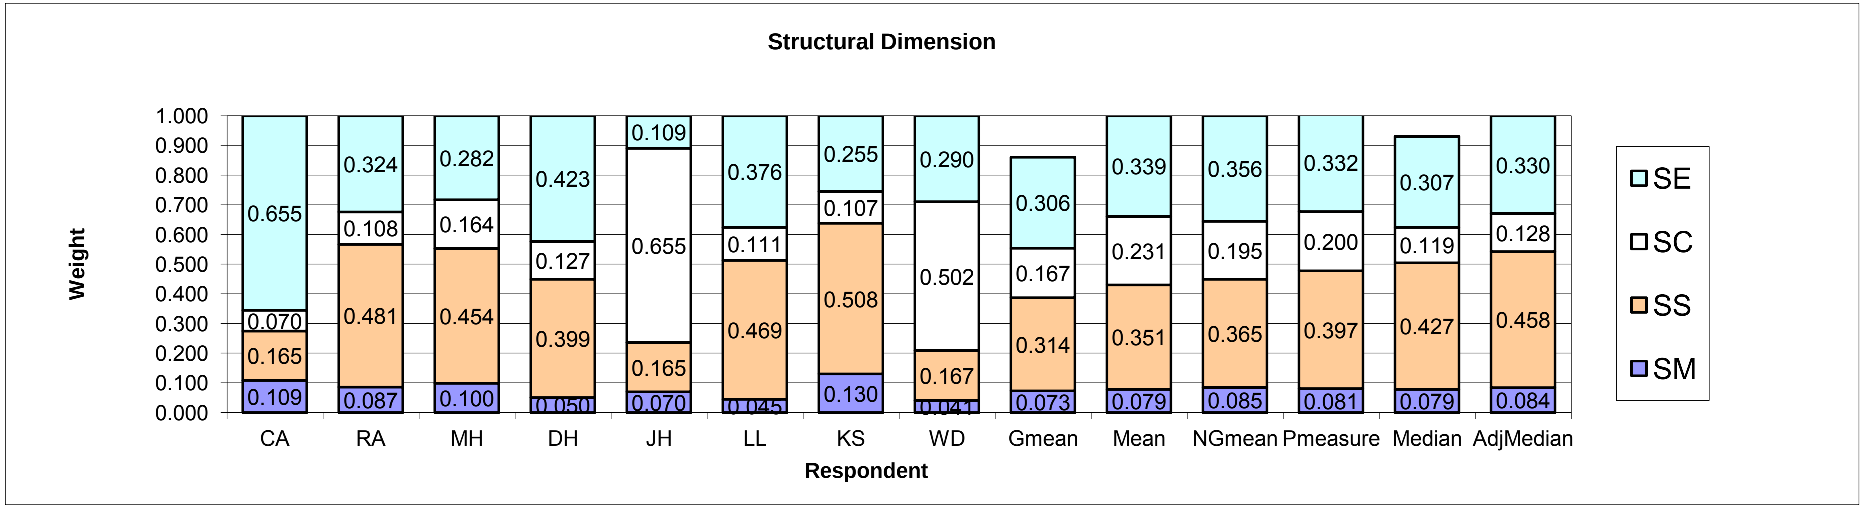

Structural Dimension Weight Data

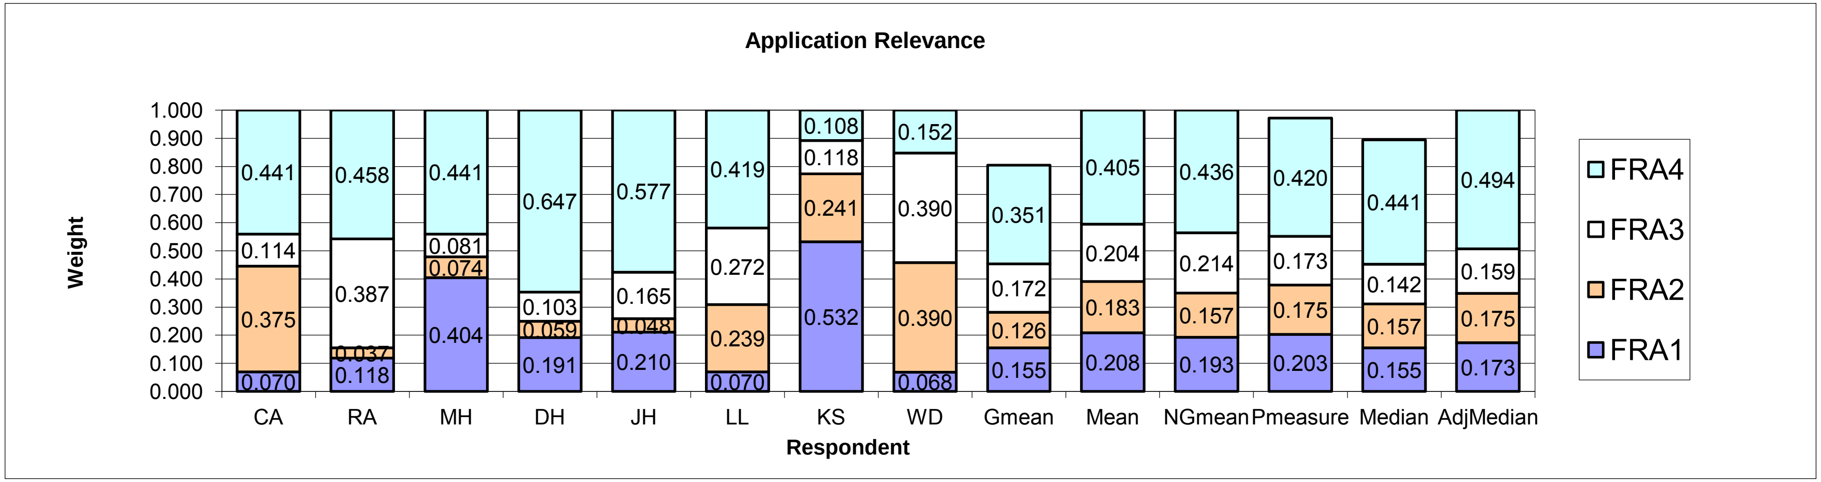

Application Relevance Weight Data

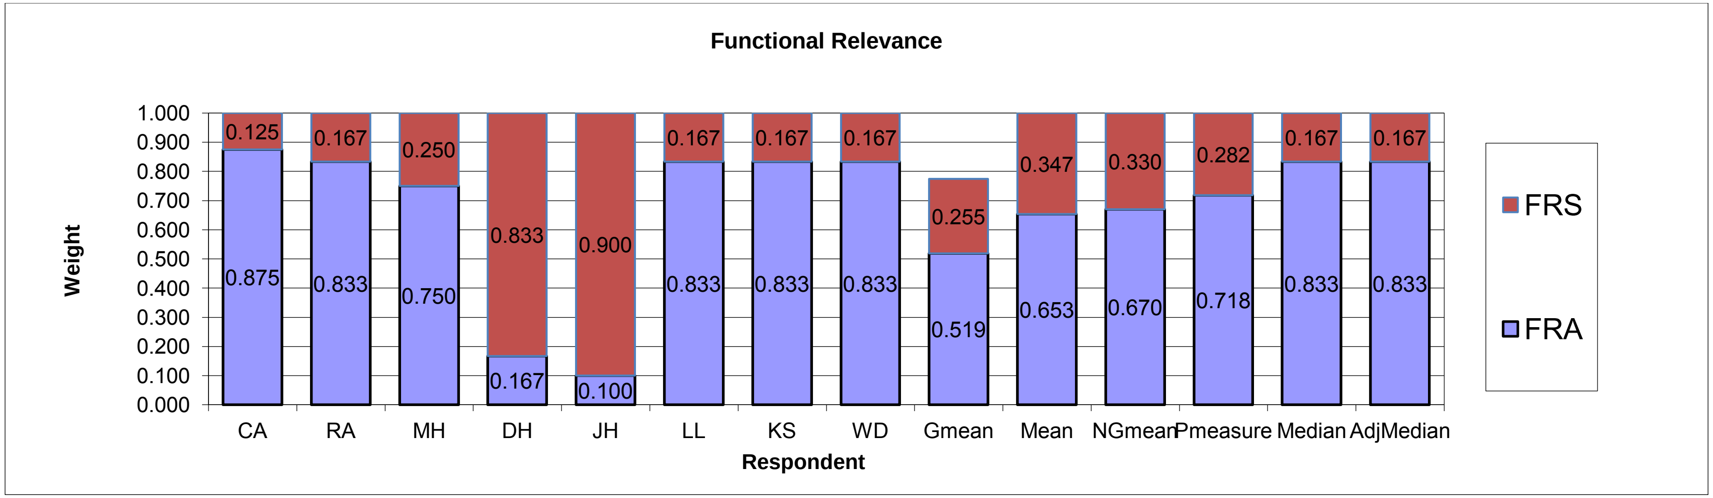

Functional Relevance Weight Data

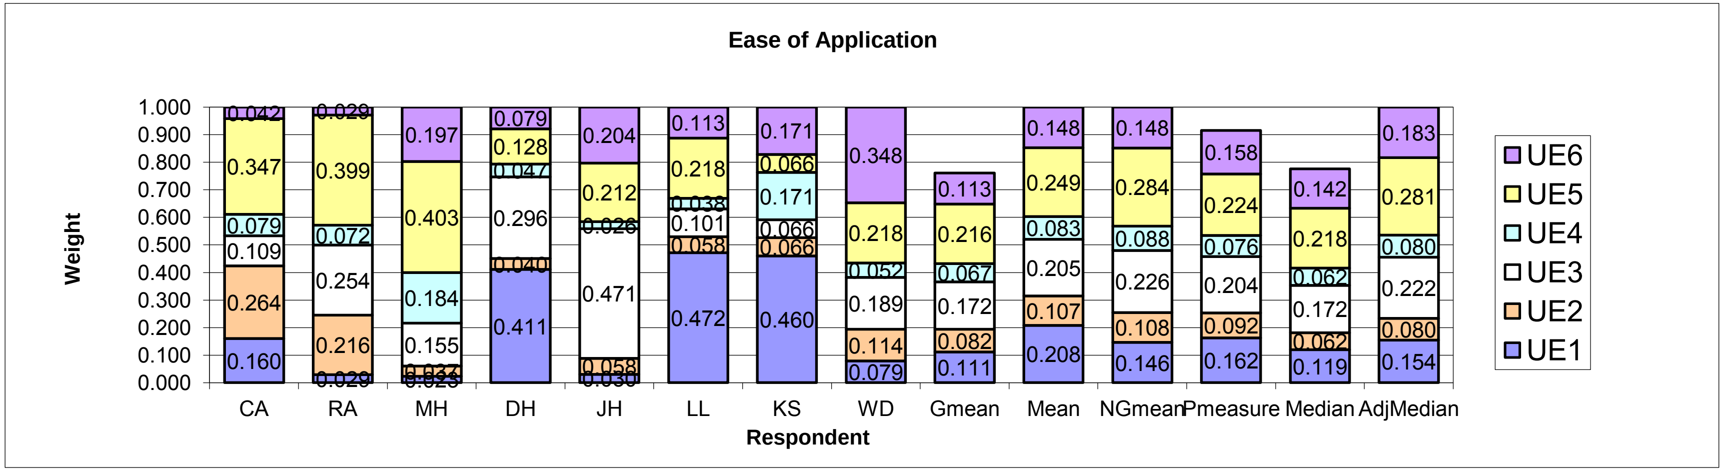

Ease of Application Weight Data

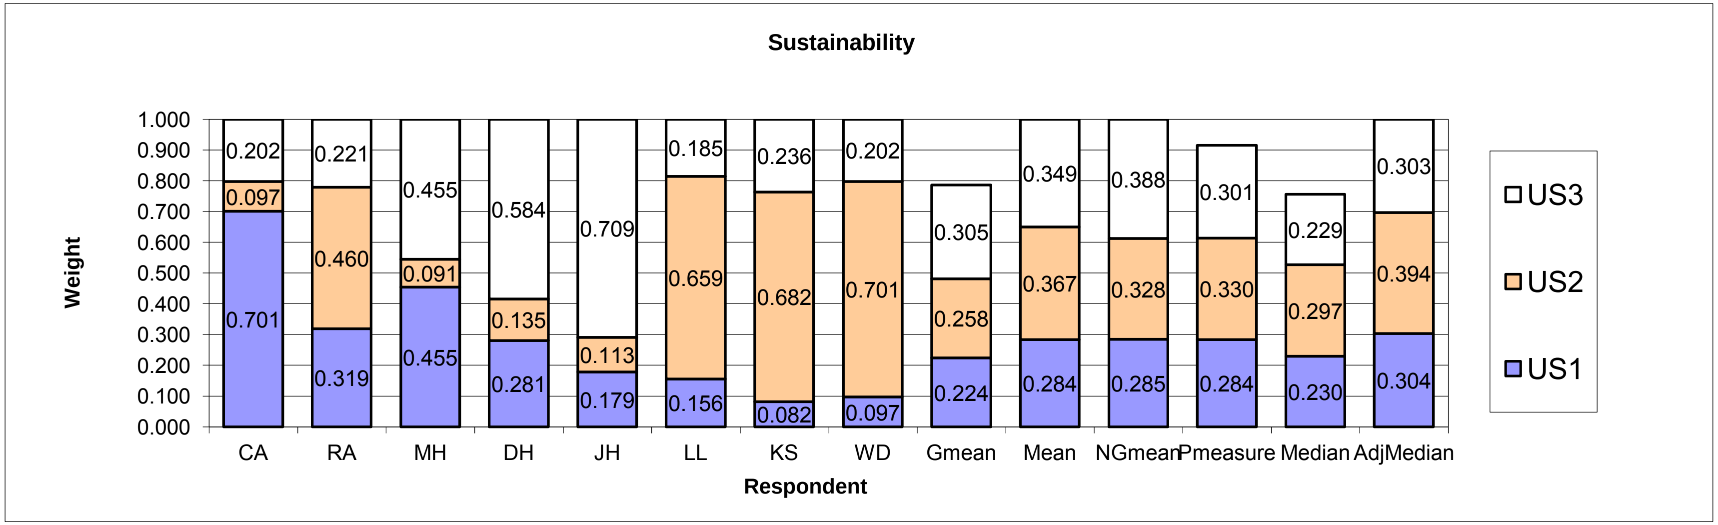

Sustainability Weight Data

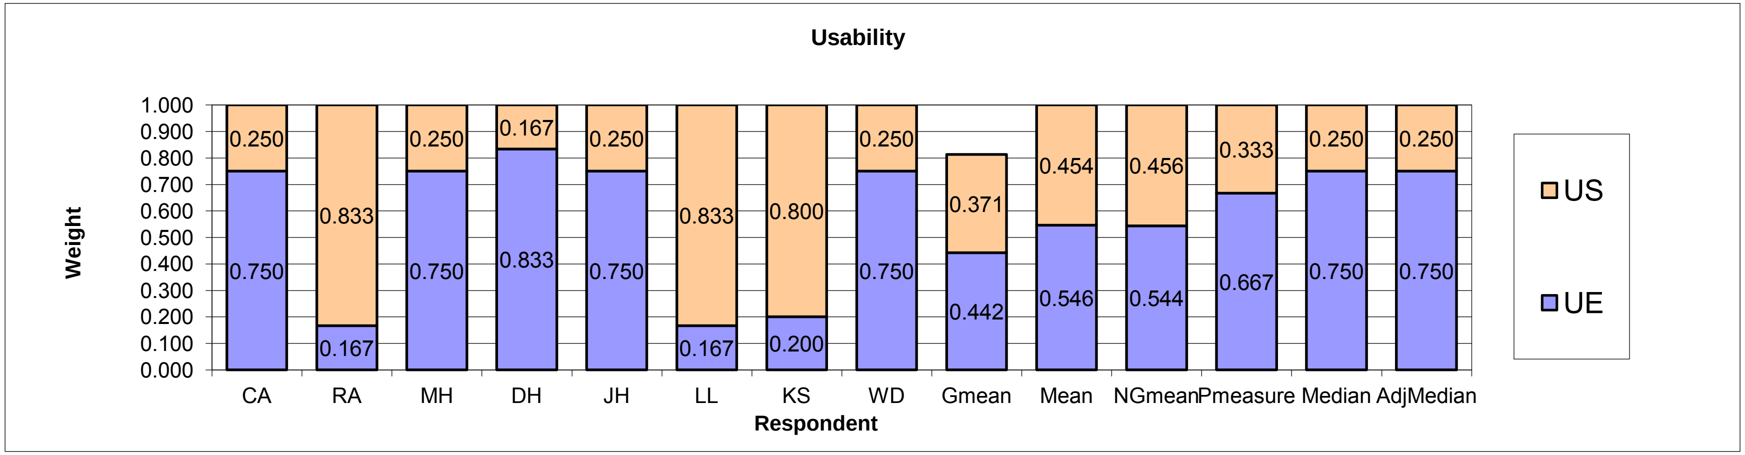

Usability Weight Data

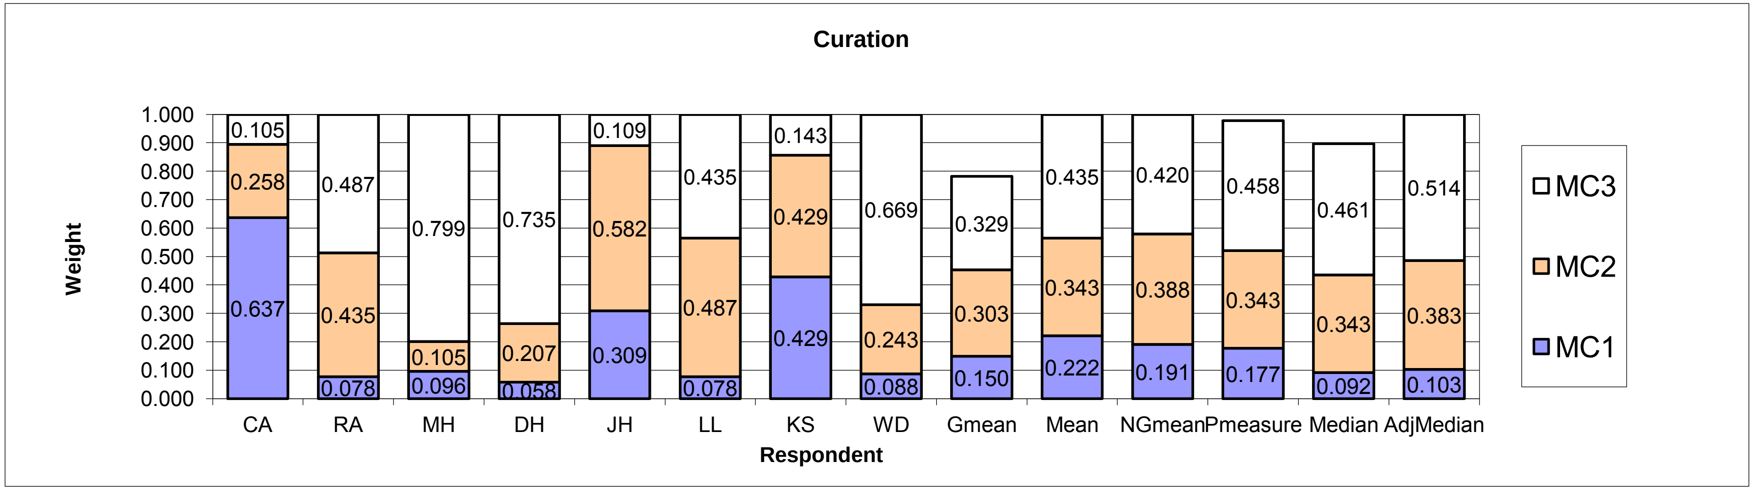

Curation Weight Data

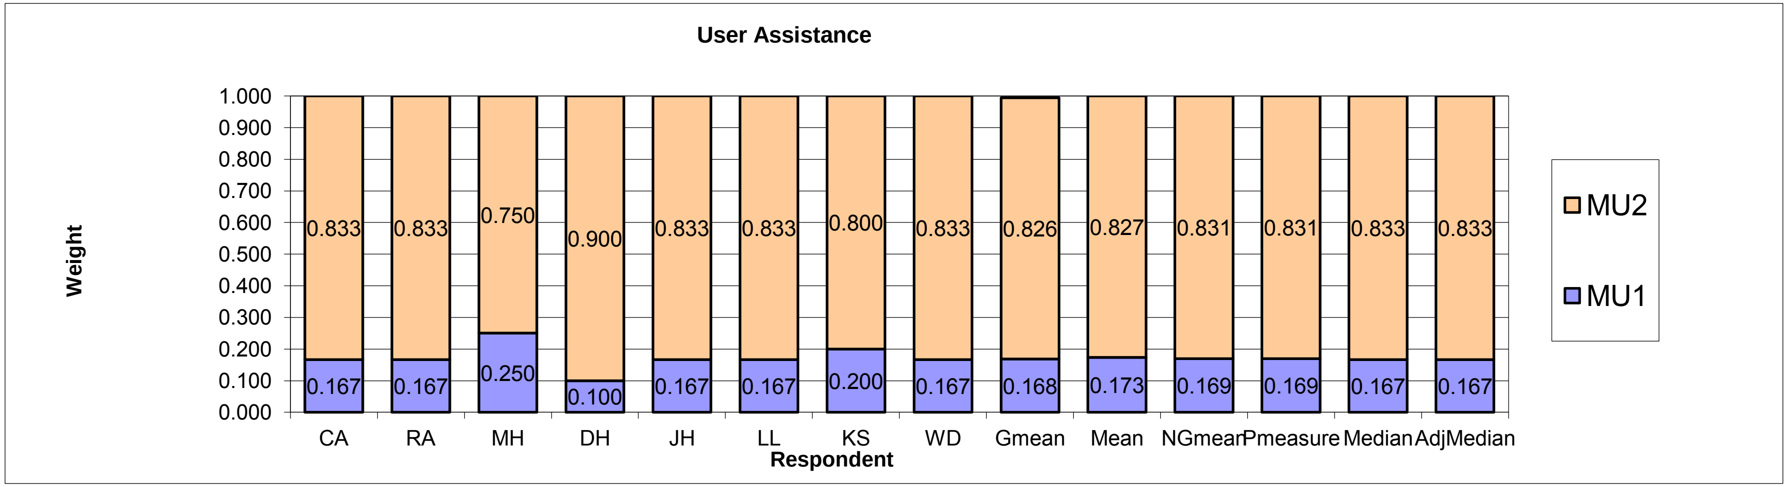

User Assistance Weight Data

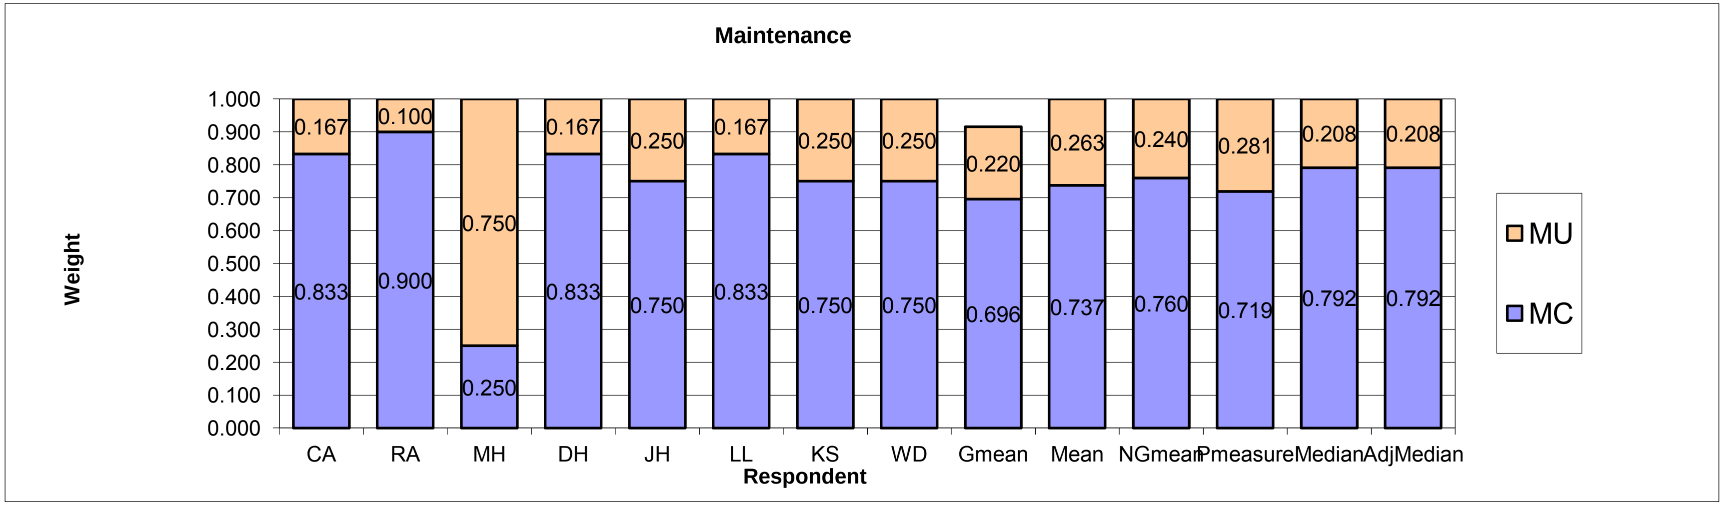

Maintenance Weight Data

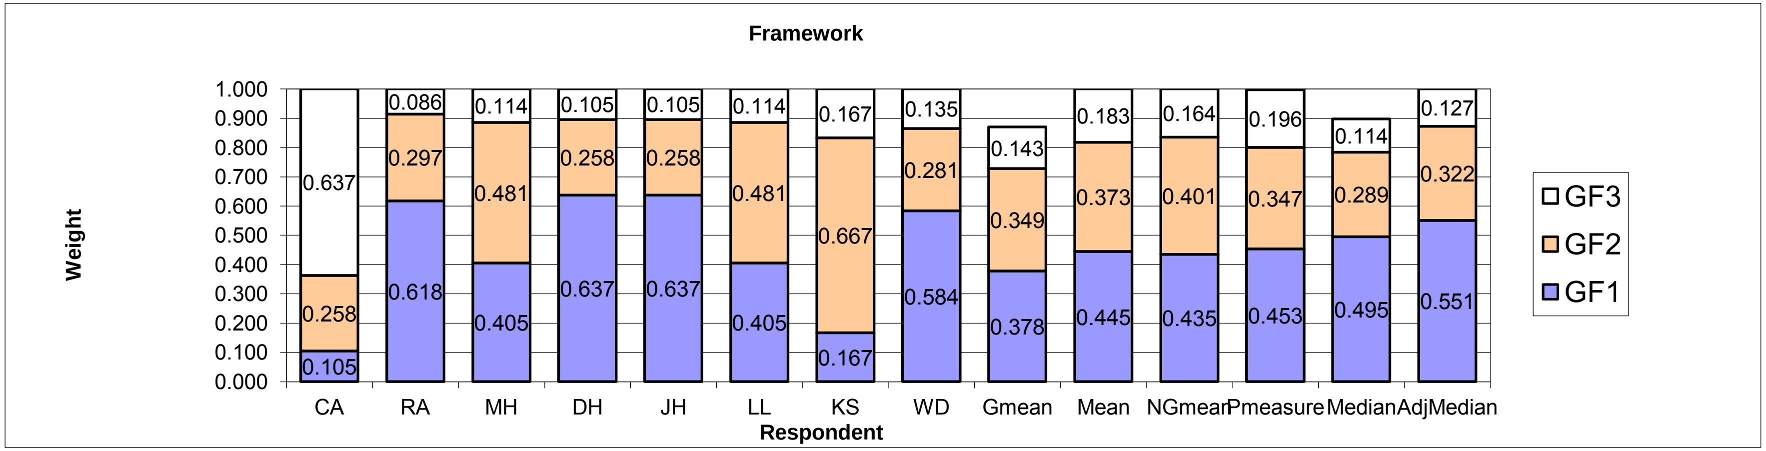

Framework Weight Data

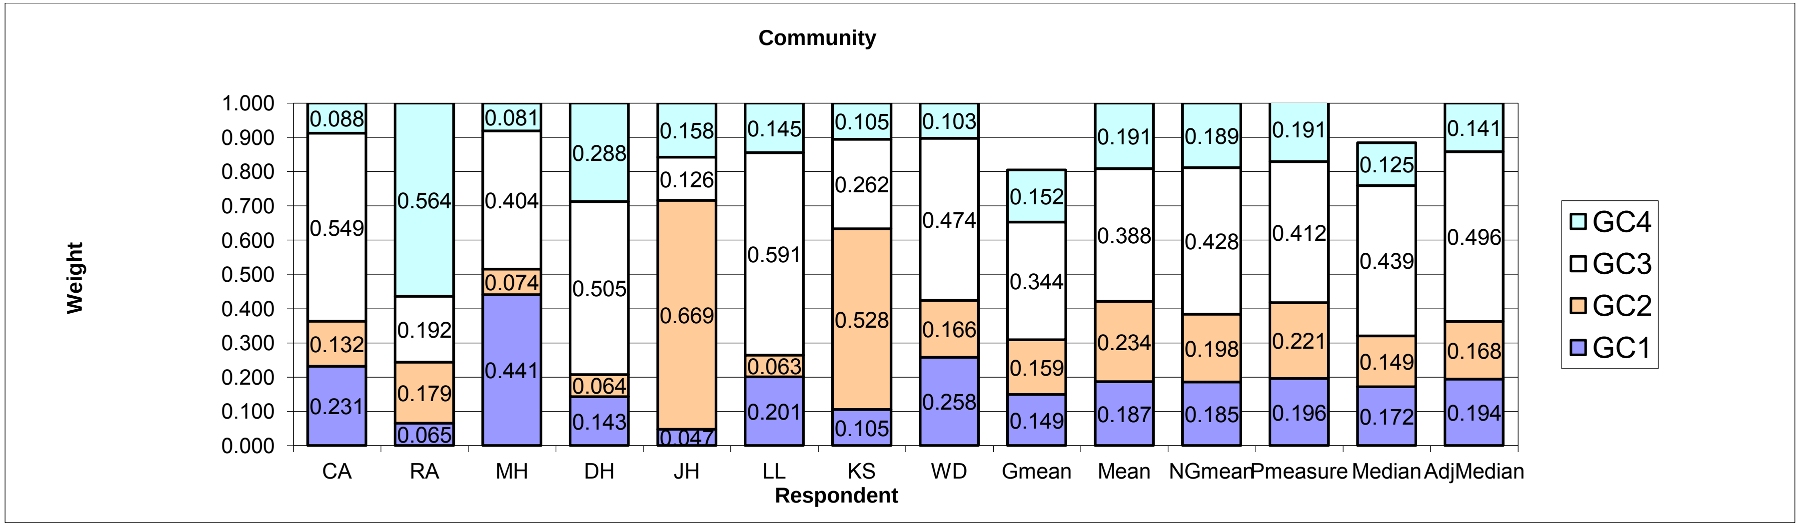

Community Weight Data

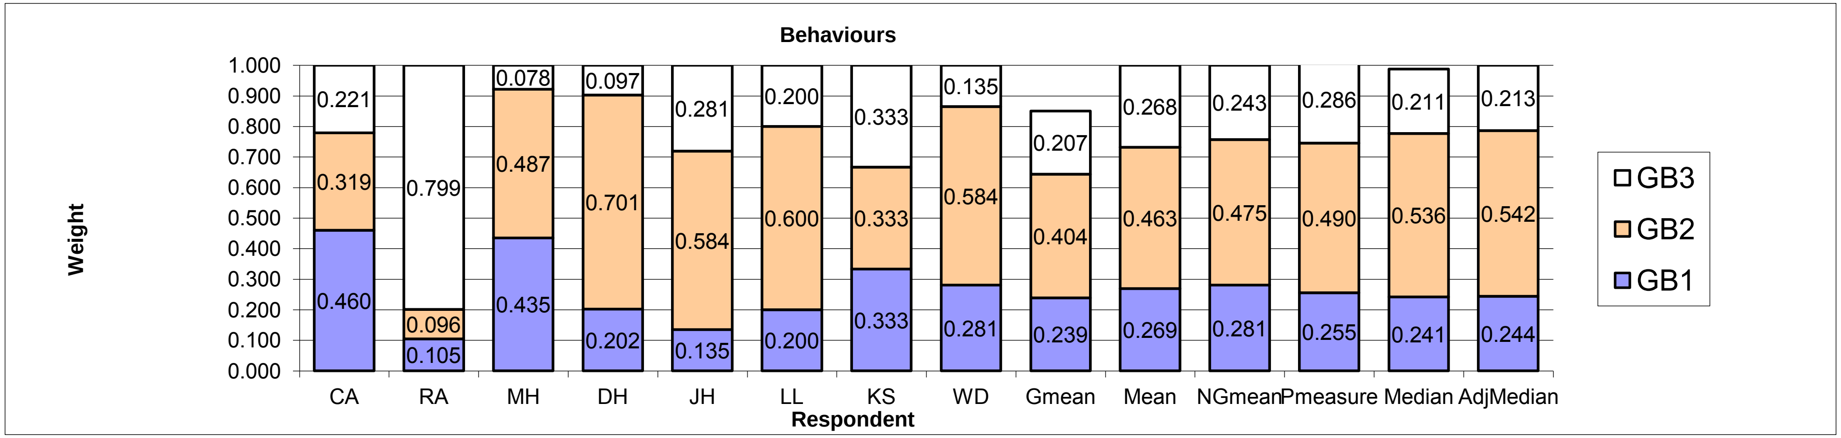

Behaviours Weight Data

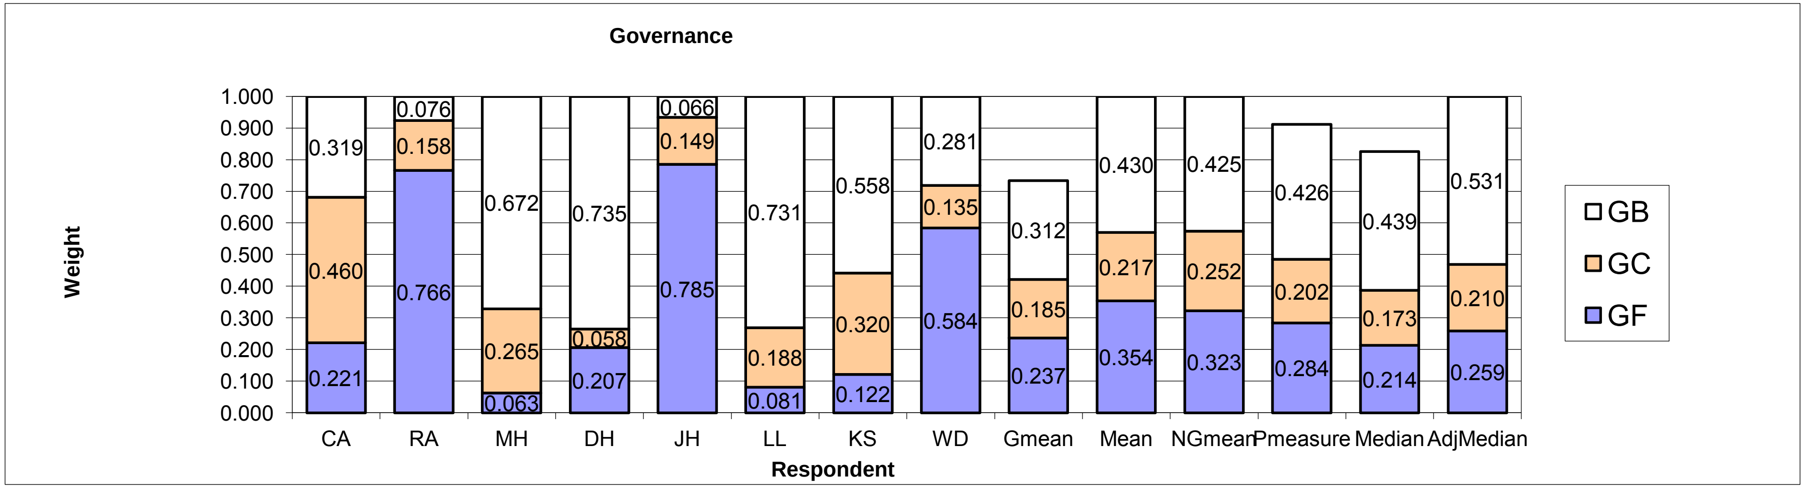

Governance Weight Data

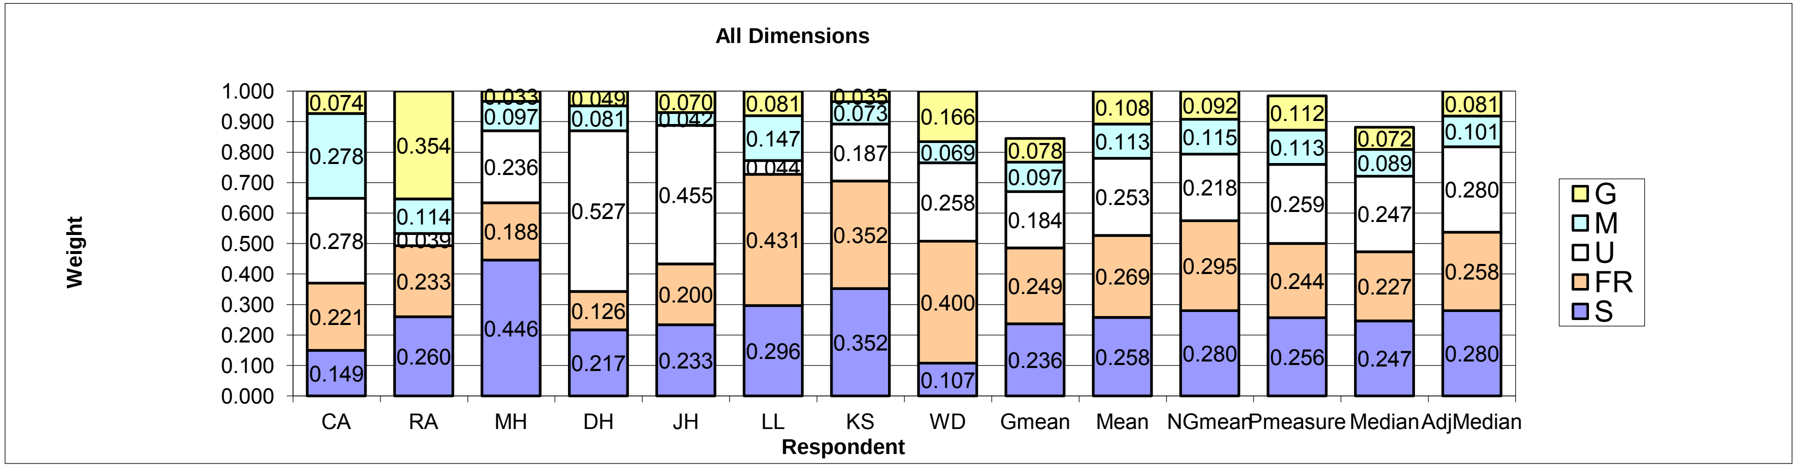

All Dimensions Weight Data

Appendix 28 – R Script For Performing Multi-Dimensional Scaling On Expert Weight Data

```
library(MASS)

x=read.table('c:/Users/kim_fin/Documents/DocumentspostNovember2009/MDS_matrix.csv',sep=',',header=F)

x=as.matrix(x)

fmds <- isoMDS(dist(x))

fmds$stress

plot(fmds$points, type = "n")

text(fmds$points, labels = as.character(1:nrow(swiss.x)))
```