B1

AN ONTOLOGICAL COMPARISON AND EVALUATION OF DATA MODELLING FRAMEWORKS

By

Simon Milton, BSc (Hons)

(Simon Kenneth)

Submitted in fulfilment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

School of Information Systems

UNIVERSITY OF TASMANIA

MAY, 2000

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

Simon Milton

HOBART, MAY 2000

Simon Milton

HOBART, MAY 2000

Abstract

DATA MODELLING FRAMEWORKS are used to construct models of reality for use in information systems, computer science and software engineering. Currently, there is a plethora of data modelling frameworks each possessing a view of the world, in that each data modelling framework has a distinct set of terms that are used to create models of reality. We are interested in finding a unifying framework in which to compare and study data modelling frameworks. We turn to the study of ontologies to find a theory with potential as a unifying framework. Ontologies are studied in philosophy and are concerned with 'what there is'. An ontology defines the categorial structure of reality and the terms that are fundamental for describing reality. Consequently, ontologies are ideally suited to comparing and evaluating data modelling frameworks, and an ontology has the potential to provide a unifying framework for data modelling frameworks. In this thesis we consider using ontology in this role and examine using a specific commonsense realistic ontology from philosophy in the role of a unifying framework.

We are seeking to establish the nature and degree of synonymity between the world view of a pragmatically selected ontology and of several representative data modelling frameworks. We propose and apply two qualitative methods to help us. By applying these methods, we can begin to understand the efficacy of our approach and the generality of using an ontology as a unifying framework.

We have found that there is reasonable commonality between the world view of the selected ontology and that of the data modelling frameworks we have studied. However, we have found areas in which specific data modelling frameworks do not match the world view of the ontology. We have also found areas in which aspects of specific data modelling frameworks are not supported by the selected ontology. Most notably, we have found grounds to question the rigid class hierarchies common in many object modelling frameworks. However, all data modelling frameworks appear capable of being extended to support the selected ontology. We have reason to believe that the selected ontology is thus an excellent candidate for further investigation as a unifying framework.

Acknowledgements

I WOULD LIKE to thank my supervisors, Chris Keen (School of Information Systems, The University of Tasmania) and Ed Kazmierczak (The University of Melbourne), for their professionalism and support during my study. Both have generously provided encouragement and engaged in stimulating interactions that have helped me through difficult points along the way. I am particularly grateful to Ed for his numerous suggestions and discussions during the development of this work and for reading everything so critically.

Thanks also to my colleagues and fellow students at The University of Tasmania: I especially wish to thank Cathy Urquhart, who helped me survive the rigours of study and showed me by example a new way through; Bernard Gunther, for his companionship in the department; and Linda Dawson for her level-headed and helpful suggestions in Hobart and Melbourne. I thank them all for their professional support.

I would also like to thank my very good and long-standing friend Glynne Ridgeway, for providing me with a refuge in which I could write without interruption either from students or from my adorable Abyssinian cats.

My parents have taught me consistently to follow my interests, to reach my goals and to succeed. They have shown consistent support and encouragement together with positive suggestions during periods of despair.

Lastly, but by no means least, heartfelt thanks to Philip for his love, patience and understanding through the years.

DECEMBER 1999

This thesis is dedicated to the memories of both my grandfather, Archibald Carl Milton ("The Mulga") and my uncle, Raymond Alan Milton, both of whom died during the writing and examination of the thesis. May they rest in peace.

MAY 2000

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

There are many data modelling frameworks in use today, each of which has a number of terms that are used to abstract from reality and build static and, in some cases, dynamic models of the world. The terms from a data modelling framework, together with their associated meanings, implies a view of reality. For example, the Entity-Relationship (ER) data modelling framework has a term called 'entity' that is used to represent significant independent things found in reality. Another term called 'relationship' relates two or more entities. Further, ER allows for the description of entity classes and the attributes that each member of a specific class may exhibit. Other data modelling frameworks differ from ER in the terms provided. The differences may be in the names given to terms that in other respects fulfil a very similar purpose to 'entity' from ER. Alternatively, there may be several terms in one data modelling framework to achieve the same modelling objective as one from ER. It may also be the case that the meaning behind terms (or groups of terms) may be different, in subtle or radical ways, from related terms in ER. The difference in terms is recognition that they all have a different view of reality, and by implication encourage modellers to view things differently when constructing abstractions using the terms. There are also claims of the 'superiority' of this data modelling framework over that, or that one data modelling framework is more 'natural' than another. In this thesis we begin to search for a unifying framework with which we can reconcile the different views of reality possessed by the many data modelling frameworks.

In philosophy, ontology is the study of world views or 'what there is' in the world, and it is a mature and advanced field of research. We turn to ontology in our search for a unifying framework for data modelling frameworks. The scope of this thesis is to work toward the goal of developing a unifying framework for the study of data modelling frameworks, and in it we seek to establish the nature and degree of synonymity between the world views of an ontology and of several data modelling frameworks. We propose two methods to help us determine the degree of synonymity. We then select an ontology based on a number of pragmatic criteria, and apply the methods to a number of representative modelling frameworks using the selected ontology. The methods use an ontology to compare and to evaluate data modelling frameworks. By applying these methods, we can begin to understand the efficacy of our approach, the role of ontology in developing theory, and to discuss further research that will help us reach our long-term goal of a unifying framework. The scope of this thesis is restricted to consideration of modelling languages that consider data or static structures and we use the term 'data modelling framework' in this context.

1·1 Data Modelling and Ontology

Many modelling frameworks have been proposed during the last three decades, beginning with semantic models (Hull and King, 1987, Peckham and Maryanski, 1988), and progressing through to object models (Banerjee et al., 1987, Cattell, 1991, Dobbie, 1990, Lécluse et al., 1988, Rumbaugh et al., 1991, Boman et al., 1997). Modern data modelling methods involve establishing an abstraction of data and information from the world which is acceptable and relevant for an organisation and its purposes. Semantic data modelling frameworks (Hull and King, 1987, Peckham and Maryanski, 1988) were motivated by the need to increase the semantic representation that was shown to be lacking in the relational model (Codd, 1970). Object models, in contrast, have been established partly to enable reuse of software components, and partly to improve the semantic content of modelling frameworks. Object models have been described as being more natural than other data modelling frameworks. There has been further influence on modelling frameworks by semantic nets and other aspects of artificial intelligence (Brodie, 1982, Hammer and McLeod, 1990, Hammer and McLeod, 1981).

The large number of data modelling frameworks available makes the task of selecting a modelling framework very difficult, and there is no universally accepted principle that determines the suitability of one data modelling framework over another. Additionally, claims of universality for, or quality of, specific data modelling frameworks are largely unsubstantiated. Nevertheless, data modelling frameworks are used for modelling reality, and for this reason we are interested in finding a unifying framework that can be used in conjunction with data modelling frameworks for the purpose of studying them. Ontology is useful for this purpose because of its role in Philosophy—that of considering the "furniture of the world" (Bunge, 1977). Philosophers who construct ontologies seek an understanding that "... deals with the most general properties of beings in all their different varieties" (Kim and Sosa, 1995). It is in this way that ontologies perform an analogous role to that played by data modelling frameworks.

Ontological studies is an ancient and mature field (Flew, 1989), dating back to Aristotle and Plato (Kim and Sosa, 1995). The study of ontologies provides us with a source of independent views on 'what there is' in reality, and it may, in due course, be able to provide us with a unifying framework for data modelling frameworks.

As a philosopher understands it, the study of ontologies deals with the 'categorial structure of reality' (Honderich, 1995). An ontology also provides a description of fundamental terms, which one uses to describe reality, and the ways in which these terms relate to the categories. In these ontologies basic questions are asked concerning the constitution of reality and what fundamental categories of things exist in reality, together with the terms that one needs to construct a description of a 'state of affairs'.

Other groups of researchers view ontologies as being restricted to a particular domain such as those used in artificial intelligence research (Vickery, 1997). The world view of a civil engineer is an example of a domain-specific ontology. In these types of ontology, categories are constructed that represent the vocabulary of the universe of discourse in which the group is likely to be interested.

We are interested in ontologies that are of the more general type and not confined to one problem domain, or that describe categories of a restricted sort. This is because data modelling frameworks are not restricted to one class of modelling problem and are needed to build representations of reality which are widely applicable in a similar way to ontologies constructed by philosophers.

An ontology tells us what terms are needed to describe 'what there is' (physical and non-physical) in reality and, more deeply, the ontology also provides the concepts that fully define each term. They also establish ontologically fundamental categories into which terms are placed, and thereby establishing what categories are needed. For example, an ontology may outline the fundamental categories as being 'individual', 'attribute', and 'event'. The terms may be 'individual', 'attribute', 'relation', 'class', 'set', and 'event'. Each term is then explained and related to the fundamental categories. Terms may define or elaborate upon specific categories, such as 'event', 'attribute', or 'individual', or their relationship to one of the categories is defined and explained. The terms are what is needed to adequately describe 'what there is'; The categories are the minimum number of ontological categories to which terms relate.

We are interested in exploring the role of ontologies that talk about the general nature of objects in reality as a unifying framework for data modelling frameworks. Data modelling frameworks also provide us with the terms with which we can build models of reality. For example, 'entity', 'entity class', 'relationship', 'attribute' are terms that can be used to describe 'what there is'. Each term has a meaning that governs the sort of things to which it refers when building models of reality.

Terms from both the ontology and the modelling frameworks have meaning through concepts that describe their nature. The totality of terms and concepts embodies the world view of a data modelling framework and an ontology. Our aim is to compare and contrast the views of reality as embodied in terms from various data modelling frameworks with the view of reality as found in terms from an ontology. For example, how similar are the terms 'individual' and 'entity'? In doing so, we can begin to examine the role that ontologies can play as a unifying framework for understanding data modelling frameworks.

We believe that a unifying framework for data modelling frameworks has three features or capacities. Firstly, it has a coherent set of concepts. Secondly, each data modelling framework can be expressed within the unifying framework. Thirdly, that the unifying framework can be used to analyse any data modelling framework. In this context 'analysis' means the ability to compare, evaluate, extend, and understand the view that each data modelling framework has with respect to the world.

In philosophy, there are many ontologies from which to choose. Further, there is little by way of objective measures of fitness. In fact, it could be said that without criteria to help with fitness for purpose, ontologies are all equal in quality. It is therefore important to select an ontology that is useful for our aims and objectives. Essentially, we use pragmatics to select an ontology.

1·2 Motivation

There is previous research using ontologies in information systems and data modelling. The work falls into three categories, some of which provides further motivation for our work.

Firstly, the tools and techniques used in systems analysis and design have been examined using ontology in order to discuss the quality with which they represent reality (Weber, 1997, Wand, 1996, Wand et al., 1995, Wand and Weber, 1995, Wand and Weber, 1993). Quality is defined in terms of the clarity and accuracy with which tools are capable of representing real world phenomena. In their work, clarity and accuracy of representation is categorised by examining, in specific ways, the mathematical mapping between grammars that represent firstly, the modelling tool under investigation (for example, ER, or data-flow) and, secondly, a selected ontology as the theoretical measure of representational fidelity. In comparing grammars in this way, there is a concentration on the construction of lawful models in a tool according to the syntax specified in the grammar. Thus, only the structure of the relevant tool is examined against the ontology. The work published to date, uses a scientific and systems oriented ontology by Mario Bunge (Bunge, 1977, Bunge, 1979). Our work differs in that we are seeking a unifying framework for qualitative studies of data models. We also consider a different ontology from Mario Bunge's. This thesis represents the first step towards achieving the goal of a unifying framework in that we compare the terms and associated concepts that constitutes the world view of an ontology with terms and concepts from several representative data modelling frameworks. Through the methods we propose and use in this thesis, we study the ontology of each data modelling framework and see how they compare with the world view embodied in our selected ontology. From this experience, we comment on the success or otherwise of our quest for a unifying framework based on the selected ontology.

Secondly, researchers in artificial intelligence have used domain-specific ontologies to guide the construction of models for knowledge representation (Vickery, 1997). Our work is different from this in that we seek to compare the world view implicit in specific data modelling frameworks using the world view found in a metaphysical ontology. We seek to compare and contrast the world view implicit in a modelling framework itself and not in a specific *model* constructed using a data modelling framework or to consider what a specific group considers to be the categories of 'what there is'. Further, we use ontology in the sense of it being the study of 'what there is' in a mature and long recognised area of philosophy, and not to create our own theory.

Thirdly, there has been work to categorise data modelling frameworks according to the ontology (or world view) embodied by an application domain (Mylopoulos, 1998), that is, to recommend certain data modelling frameworks according to the world view implicit in an application. This work is a pragmatic categorisation of data modelling frameworks and other modelling tools into a "categorization of information modelling techniques which classifies them according to their *ontologies*, i.e., the type of application for which they are intended to support, the set of abstraction mechanisms (or, structuring principles) they support, as well as the tools they provide..." (op cit.). It differs from our own work in that we seek a unifying framework for data modelling frameworks by considering a suitable ontology from philosophy and that we are working from the premise that all data modelling framework should be capable of modelling similar things. If we find that the modelling frameworks do not reflect a selected ontology, then further explanation and investigation is required.

We select an ontology using pragmatic criteria for the purpose of comparing and evaluating data modelling frameworks. By using the selected ontology in this way, we expect to be able to assess its usefulness in providing a theoretical basis for a unifying framework.

1·3 Focus of Research

We are interested in conducting ontological studies of a number of data modelling frameworks. Implicit in each data modelling framework is a view of the world that is encapsulated in the nature of the terms used to construct models. To illustrate this, we briefly consider three examples. Firstly, the entity relationship modelling framework discusses 'entity classes' and 'relationships' between members of entity classes. A member of an entity class (entity) must not be dependent, for its existence, upon any other element in the model. Each entity can be described using attributes, and each entity class has the same set of attributes for all of the entities within it. The set of attributes is the maximum number of attributes that can be exhibited by any entity in the class. Secondly, OMT consists of 'objects' that are grouped into 'classes'. Each instance in a specific class exhibits attributes. Attributes are specified at the class level. A rich and rigid class hierarchy is specified that allows inheritance of structure and behaviour. Associations are allowed between instances of certain classes and are specified through links. Thirdly, the functional data model discusses 'entities' (representing both 'objects' and 'attributes') and 'functions' between entities. Further, the type structures can represent generalisations of entities.

These three examples of data modelling frameworks all purport to be useful in constructing data models. Some people might claim that one of them is superior to the others. We may call this an example of a 'religious war' (my modelling framework is better than yours). Maybe the users of a specific data modelling framework claim it to be more natural than other frameworks. We are interested in discussing ways by which these claims can be rationalised, and we seek a unifying framework which can be used for debating such issues as those raised here.

An ontology of the type in which we are interested tells us what terms are required for modelling reality and, more deeply, the meaning behind each term by relating these terms to a specific theory of categories of 'what there is'. Through its terms, an ontology can be used to create an abstraction from reality. Data modelling frameworks provide us with terms we can use to build models or descriptions of reality. Our aim is to compare and contrast the views of reality as embodied by terms in various data modelling frameworks with the view of reality as found in the terms from an ontology and, in doing so, we can begin to evaluate the role ontologies may play as a unifying framework for data modelling frameworks.

Terms found in an ontology or a data modelling framework are given meaning through an associated concept. For example, 'individual' may be a term in an ontology. The term called 'individual' is given meaning through its associated concept that is in turn conveyed by the author of the ontology through his or her writing. In a data modelling framework there may be a term called 'entity'. It too has an associated concept that is described in the literature and in accepted usage. Each term from an ontology and a data modelling framework can be described in this way in that each term has an associated concept giving it meaning. It can be said that the world view of an ontology and of a data modelling framework is contained in the totality of terms and concepts for each. We are interested in examining and relating the world view of an ontology with that of data modelling frameworks.

The terms and concepts describing those terms are used for a qualitative comparison and evaluation of the world view of the ontology and the world view of the data model.

Once we have pragmatically selected an ontology as a potential unifying framework then we can use the ontology in several tangible ways.

We can use the terms and associated concepts from an ontology as the basis for an ontological comparison of a number of data modelling frameworks. The ontology, through its terms, guides the comparison. We begin by extracting the concepts from the ontology that are relevant for a comparison. We then compare the world view embodied in the ontology with the world views encapsulated by the respective data modelling frameworks through the concepts associated with terms in each data modelling framework.

Secondly, we can perform an ontological evaluation of a specific data modelling framework, and examine the similarities and differences between the world view as embodied in a data modelling framework and the world view as described by an ontology. In the evaluation we examine questions relating to the presence of concepts in data modelling frameworks that are absent from the ontology, and examine questions relating to the presence of concepts in the ontology that are absent in the data modelling framework. In this approach we extract a complete set of terms and concepts from the ontology and data modelling framework, and then examine similarities and differences between them on the basis of concepts. Of interest are the areas of agreement and difference between the concepts of the two. Differences, in particular, are analysed in order to more deeply understand them and to suggest parts of a data modelling framework requiring extension in order that it more closely relates to concepts from the ontology.

The world view embodied by the data modelling framework and that embodied by the ontology can be brought closer together, by modifying or extending the data modelling framework where a difference is found. There may be a number of extensions that result in a state of synonymity of world view between a data modelling framework and an ontology. The precise nature of the relationship between concepts from the ontology and the data modelling framework will determine the ease with which any extensions are likely to be made. Additionally, issues raised from examining the ontology with respect to widely used data modelling frameworks can be verified against a user community drawn from an industrial setting.

We examine the methods in greater detail in Chapter 3 where we develop the ideas outlined here and describe the methods, the results for which are then presented in chapters 6, 7, and 8. It is critical when presenting methods to also discuss the analysis of results from applying the methods. As we introduce the methods we also introduce the mechanisms that we use to analyse the results we obtain.

The selection of an ontology for the methods we describe in this thesis is critical for several reasons. Firstly, there is a multitude of ontologies, and each ontology has its own philosophical background that reflects the outlook of the author of the ontology and in turn will indicate its usefulness for certain tasks that we envisage. Secondly, not all ontologies will be suited to work with data modelling frameworks through the sorts of terms they use and categories that they determine. Thirdly, few ontologies deal with both dynamic and static terms describing reality. We consider past use of ontologies in information systems in Chapter 2 and discuss the criteria used in this thesis for a pragmatic selection of an ontology for the thesis in Chapter 4 where we also describe the selected ontology.

There are two research questions in which we are interested for this study. Firstly, we seek to establish the suitability of the selected ontology as a unifying framework. This can only be achieved by using the ontology in the ways that we've outlined and then analysing our experience. Secondly, we would like to know what ontology is implicit in the data modelling frameworks that we study.

Consequently, we examine the following research questions:

- How suitable is a specific ontology as a unifying framework for data modelling frameworks?
- Is this specific ontology tacit in data modelling frameworks?

1·4 Research Method

In any research there needs to be a clear research philosophy that logically leads to the methods applied and to the nature of knowledge that is gained from the experience of applying the methods. Also, an approach must be clearly shown so that others can follow a similar path should they so wish, and so that others can understand the approach based upon the premises.

There are three aspects to our study:

- The methods themselves that we use as a first step towards finding a unifying framework.
- The selection of the ontology for the study.
- The epistemology adopted that is appropriate for the methods we use and to assist us in making sense of the knowledge acquired, or the discoveries made through our investigation.

An ontology is used in this research to study a number of representative data modelling frameworks. We propose two methods that are qualitative and that we use to analyse and compare a number of data modelling frameworks and an ontology through the meaning of terms from the ontology and the data modelling frameworks. The two analytical comparative methods are described in detail in Chapter 3 and use 'concepts' (Kim and Sosa, 1995) that embody terms as a basis for comparison of an ontology and data modelling frameworks. Concepts are a way of thinking about terms from the ontology and from the data modelling frameworks. Some terms are complex, and for a complex term, its associated concept may be made up of a number of parts, and it is through the various parts that terms are given their fullest meaning. We utilise Umberto Eco's semiotic theory (Eco, 1976) to analyse and discuss our results.

The analytical methods we use explore the similarities and differences between concepts in the ontology and those in the data modelling frameworks. We recognise, and will cover in our critical analysis of the methods in Chapter 9, that replication of methods based on these ideas are likely to produce (we would hope small) differences in results. This is due to the philosophical and qualitative nature of this study, where we analyse and evaluate data modelling frameworks using concepts from an ontology with the aim of comparing world views.

Secondly, we propose a set of pragmatic criteria that we then use to select an ontology. That is, we select an ontology based upon the purpose to which it will be put, and consequently, the selected ontology must be able to fulfil the pragmatic criteria. We discuss this in detail in the first section of Chapter 4. An alternative approach is to nominate the ontology to be used.

Finally, we need to describe the way in which we expect to acquire knowledge or make discoveries through application of methods in this thesis. We expect to gain insight into five groups of subject. Firstly, we expect to gain knowledge about the use of ontology as a unifying framework and its use in the methods we apply. Secondly, we can conclude insights, based upon the selected ontology, into the specific data modelling frameworks we examine. Thirdly, we can form propositions, that require further investigation, concerning the way that modelling could be handled in practice by analysts in information systems modelling and representation. Fourthly, we can draw limited generalisations about central concepts that a modelling framework must have that are expressed in various terms found in the modelling framework. Lastly, we may be able to draw conclusions about the ontology that are implied in the data modelling frameworks studied.

In all of the categories we take the position that, given the subjects of our methods, and the fact that little is known about our approach, all findings are made in an *a posteriori* fashion and according to the epistemology of empiricism (Dancy and Sosa, 1992), in that it is through the experience of analysing data modelling frameworks using a pragmatically selected ontology and as specified in the methods, that we build knowledge about the specific ontology and data modelling frameworks and about the area of data modelling and ontology. The thesis is also deductive by nature, and we expect the knowledge to take the form of propositions (Dancy and Sosa, 1992). Extra research may be required to determine the implications of this knowledge.

This thesis is original in that we propose and develop the idea of ontology as a meta-theory or unifying framework for arguing about data modelling frameworks. In particular the thesis is original in the following ways. Firstly, the methods comparing an ontology with a number of data modelling frameworks are new because, to our knowledge, for the first time we propose and apply methods that recognise the qualitative nature of terms and concepts in an ontology and in data modelling frameworks, and that uses this trait in their procedure. Secondly, the ontology used in this thesis has not been used in published ontological studies in information systems more generally, or in published ontological studies of data modelling frameworks particularly. Thirdly, we claim through experience and argument that the ontology used in this thesis is suitable to be developed as a unifying framework for data modelling frameworks. Finally, there are a number of specific findings that emerge from our specific results that are outlined in the results chapters and in our conclusions.

1·5 Thesis Outline

In Chapter 2 we place our thesis in the context of ontology and of data modelling frameworks, and we define terms such as 'ontology' and 'data modelling framework'. We examine research that is related to our research goals, and discuss where our research fits with respect to these.

We proceed in Chapter 3 to describe in detail the methods that we use to investigate our research questions. We more formally define 'concept', the relationship between 'term' and 'concept', and how these relate to the world view embodied by a data modelling framework or an ontology.

In Chapter 4, we outline the pragmatics of selecting an ontology. That is, we select an ontology based on some criteria that in turn are based on how we intend to use the ontology, and we examine its philosophical heritage to ensure that it is appropriate for our work before describing the selected ontology in detail.

Chapter 5 contains a description of the data modelling frameworks that we use in this thesis. We are using the entity-relationship modelling framework, the functional data model, the semantic data model, NIAM, and OMT's object model. For each data modelling framework and for the selected ontology we go further and describe significant strengths, and weaknesses and list key terms with associated concepts and meaning that result from our analysis of them. As discussed these concepts are used in the methods we propose and so are useful to describe.

In Chapter 6 we report findings for specific applications of the method comparing the data modelling frameworks using the selected ontology. Chapters 7 and 8 contain findings resulting from evaluating the functional data model and OMT's modelling system respectively using our selected ontology. OMT includes concepts used to describe the dynamics of modelling, and uses its functional and dynamic models to do this. Dynamics is an aspect of modelling that is outside the primary focus of this thesis. Nevertheless, it is a criterion of selection that the ontology considers dynamics and it is for completeness that we include OMT's dynamics in the application of the evaluative method the results of which are described Chapter 8. We also describe the dynamics of OMT in that chapter, thus re-emphasising the fact that dynamics is not of primary focus in the thesis.

We then conclude with Chapter 9 where we critically reflect on the thesis. The critical analysis may conclude that our approach requires improvement, or that there is merit in examining a different ontology from the one selected that matches our criteria, or that our approach has yielded an ontology that seems useful to pursue in future research. We outline our findings and discuss knowledge that has arisen from the research and our progress towards answering our research questions. After describing the specific limitations of our research we outline further research that logically arises from our study.

CHAPTER 2

Background Literature

IN THIS CHAPTER we discuss the literature that is relevant for our study and its relationship with our goals. We begin by defining 'ontology' and by discussing work using ontology in information systems and related areas, and how this work relates to our research goals. In describing uses of ontology we outline a specific class of ontologies in which we are interested, and we differentiate between ontologies of the nature we desire and those that are for specific problem domains and various business applications. The latter group is most popular in artificial intelligence and may be called 'domain specific' ontologies.

In Section 2·2 we examine previous research which surveys, compares, categorises, and rationalises about data modelling frameworks. We also define the term 'data modelling framework' and specify the classes of data modelling frameworks under study in this thesis and list the specific data modelling frameworks we examine. The selected data modelling frameworks are described in detail in Chapter 5. The particular methods that are used to study the selected ontology with the chosen data modelling frameworks are described in Chapter 3.

2·1 Ontology

Ontology is the study of world views. It is the study of 'what there is.' Within the field of Ontology, an ontology written by a philosopher can be described...

> "... as a particular system of categories accounting for a certain vision of the world. As such, this system does not depend upon a particular *language*: Aristotle's ontology is always the same, independently of the language used to describe it." (Guarino, 1998)

An ontology is a world view and describes what is fundamental in 'what there is'. An ontology defines the terms used to construct a description of reality and how the terms are related. In this thesis we use this interpretation of the term 'ontology'. Ontology is sometimes used as a synonym for metaphysics. The word 'ontology' can be interpreted in a number of ways.

> "Metaphysics, most generally, the philosophical investigation of the nature, constitution, and structure of reality. It is broader in scope than science. e.g., physics and even cosmology (the science of the nature, structure, and origin of the universe as a whole), since one of its traditional concerns is the existence of non-physical entities, eg., God." (Audi, 1995)

Ontology is an important part of metaphysics. We can expand upon this by examining another aspect of ontology that is more specific to our particular needs. In these extracts, entity refers to anything that is, and has a meaning different from that in information systems.

> "[Ontology is t]he study of being in so far as this is shared in common by all entities, both material and immaterial. It deals with the most general properties of beings in all their different varieties" (Kim and Sosa, 1995)

> "Metaphysics can also be understood in a more definite sense, suggested by Aristotle's notion (in his *Metaphysics*, the title of which was given by an early editor of his works, not by Aristotle himself) of "first philosophy," namely, the study of being *qua* being, i.e. of the most general and necessary characteristics anything must have in order to count as being an entity (*ens*). Sometimes 'ontology' is used in this sense, but this is by no means common practice, 'ontology' being often used as a synonym for metaphysics" (Audi, 1995)

For our purposes, we find the following definition most helpful:

Definition 2·1—Ontology

> "Ontology, understood as a branch of metaphysics, is the science of being in general, embracing such issues as the nature of existence and the categorial structure of reality. ... Different systems of ontology propose alternative categorial schemes. A categorial scheme typically exhibits a hierarchical structure, with 'being' or 'entity' as the topmost category, embracing everything that exists." (Honderich, 1995)

Essentially, this is in alignment with our research goals, because ontology of this nature represents an independent theory through which the building blocks of reality are described. From here on, we use the word 'ontology' in this sense. Later in the thesis, we select an ontology of this type and nature, based upon pragmatic criteria.

In contrast, artificial intelligence (AI) researchers (Vet and Mars, 1998, Vickery, 1997) use ontology in a different but related way. Guarino goes on to explain that:

> "[o]n the other hand, in its most prevalent use in AI, an ontology refers to an *engineering artifact*, 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." (Guarino, 1998)

We call this type of ontology a domain-specific ontology. For example, *economics*, or *plant taxonomy* each has its own categories of terms and intended meaning for terms used in these fields. This approach and interpretation has been found to be particularly fruitful when considering domain-specific knowledge or world-views. For example, in expert systems, domain specific ontologies have been used most profitably when modelling expert knowledge.

To help clarify the distinction between the two interpretations of ontology that he defines, Guarino continues by saying:

> "[t]he two readings of 'ontology'... are indeed related [to] each other, but in order to solve the terminological impasse we need to choose one of them, inventing another name for the other: we shall adopt the AI reading, using the word *conceptualization* to refer to the philosophical reading. Specifically, two ontologies can be different in the vocabulary used (using English or Italian words, for instance) while sharing the same conceptualization." (sic) (Guarino, 1998)

We, in contrast, take an alternative approach and maintain the term 'ontology' for the former use. However, we do take note of the distinct uses of the word. The term 'conceptualization' is not recognised in this sense in philosophy. Conceptualism is a specific philosophical stance with associated issues and definitions. Conceptualisation, in contrast, and as mentioned in the quotation, is a term coined by the AI community in the way described above. Nevertheless, we note the problems with using the term, and continue.

The problem thus turns to defining what is meant by a conceptualisation. Informally, a conceptualisation is an intentional definition of what there is, as opposed to describing a specific state of affairs. A state of affairs may be described using terms that are consistent with a particular conceptualisation. For example, a conceptualisation that is a crude object model is that the world consists of 'objects' that are described using 'properties'. One state of affairs that adheres to this conceptualisation is a model of a university constructed using this modelling framework. A second example, that adheres to this conceptualisation, is a model of the European Community governance structure and interactions (Communities, 1999), again constructed using the modelling framework outlined. We can place Guarino's discussion concerning conceptualisation in the context of data modelling frameworks by likening the AI use of ontology to a specific model constructed in a particular data modelling framework (the two examples just cited). The philosophical use of ontology is akin to the world view embodied in the data modelling framework itself (such as the very crude and partial object model also described). That is, the world view implicit in the constructs and the meaning behind the constructs of a data modelling framework.

The term ontology is used by many in information systems and computer science. Firstly, we have the AI application of ontology as discussed above (the domain specific ontology). Secondly, the philosophical and more general interpretation of ontology is also used (the one called conceptualisation by Guarino; what we simply call 'ontology'). Mylopoulos (Mylopoulos, 1998) is one that uses this latter interpretation. He recognises that modelling techniques and more specifically the data modelling frameworks themselves have implied ontologies. For example, ER assumes that in the problem domain, entities exist with relationships between entities, and that entities are grouped into classes with entities in those classes possessing like attributes. Other tools and modelling methodologies can be similarly categorised according to 'ontology'. Mylopolous goes on and categorises modelling frameworks according to the type of ontology implied by an application domain, and identifies four ontologies—static ontology, dynamic ontology, intentional ontology, and social ontology. Instead of examining the frameworks in the context of an ontology written by a philosopher, he categorises modelling frameworks according to the ontology of application in which the framework is often found to be of use.

The categorisation used by Mylopoulos is general, but is constructed in a manner that does not take into account any categorisation from the philosophical study of ontology, and more generally it neither provides nor suggests a unifying framework for data modelling frameworks. Nevertheless, it is a recognition that there is an ontology implied by every data modelling framework. It is for this reason that it is important for us in that it generally supports our approach to data modelling frameworks in that we say that each framework possess an ontology. Mylopoulos' categorisation is a more useful categorisation of data modelling frameworks than has been constructed to date (Peckham and Maryanski, 1988), because of the fact that it recognises that each data modelling framework possesses an ontology. However, it begs the question of the origin of each ontology, and consequently it suggests that searching philosophy for relevant ontologies might be fruitful. Our research differs in that we seek a unifying framework for data modelling frameworks based on an ontology selected from philosophy, a mature intellectual endeavour that can yield an independent theory which could prove useful as a unifying framework.

Others in information systems have applied ontology using material explicitly from philosophy. Work commenced by Yair Wand and Ron Weber (Wand and Weber, 1989) is the most relevant example for the purposes of our research goals. In order to remind us of our objectives before considering their research in more detail, we restate our research goals and how they relate to ontology.

We intend to use an ontology as a unifying framework that suggests unifying principles for the design of data modelling frameworks. This means we are not interested in ontologies for individual problem domains like those in artificial intelligence, because data modelling frameworks are used to create models of reality for a great many problem domains or business applications. Instead, we are interested in ontologies as understood by philosophers. Ontologies establish the terms from which a description of reality can be built, and this characteristic is ideal for our purposes because we seek the unification data modelling frameworks that are used in a similar way. We intend to do this by comparing the world view found in specific data modelling frameworks with the world view embodied by a selected ontology. We expect to achieve this by comparing the terms used in data modelling frameworks with the terms used in an ontology.

We need to take into account two issues when considering the work by Yair Wand and Ron Weber and associated research in the context of our goals. We must discuss their research goals and methods used in their research, and how they relate to our objectives. We must also consider their choice of ontology and the ways in which the ontology has been used in their research.

Wand and Weber use a system-based ontology adapted from the ontology by Mario Bunge (Bunge, 1977, Bunge, 1979) in studying concepts fundamental to information systems, their structure, and their analysis. As summarised in (Wand et al., 1995, Wand, 1996), they have proposed a theoretical paradigm based on ontology to analyse modelling methods (Wand, 1996, Rohde, 1995, Wand et al., 1995, Wand and Weber, 1993, Wand and Weber, 1990, Wand and Weber, 1989, Green, 1996), discussed application of ontology to conceptual modelling (Wand and Weber, 1993, Wand et al., 1995, Wand, 1996), and interpreted specific semantic data models (data modelling frameworks, as we would call them) in the light of their selected (and modified) ontology (Wand, 1989, Weber and Zhang, 1996, Wand, 1996, Wand et al., 1995, Wand and Woo, 1993, Wand et al., 1993), including object models (Wand, 1989, Wand and Woo, 1993, Parsons and Wand, 1993). We discuss each of these areas of research in the following subsections. We conclude the discussion by examining, in more detail the ontology used by Wand and Weber and associated researchers. In the examination, we contrast that ontology with our objectives and outline our selection method.

Their research is motivated by a desire to better understand the process of analysing, designing, and implementing information systems, and so develop a theory that can be used as a basis for much of information systems analysis and design. It is for this reason they select an ontology that is described as a systems ontology, and are keen to integrate the various aspects of information systems analysis, design and implementation.

Ontological Studies of Systems Analysis and Design Methodologies

The main focus of the work by Wand and Weber is on systems analysis and design methodologies (Wand, 1989, Wand et al., 1995). They note, that despite the plethora of design methodologies, there does not exist a general theoretical foundation for tools and techniques used in systems analysis and design, and they state that a comparison of the many techniques is not possible without a unifying theory. Without such a theory, they say, it is also impossible to establish rules by which design of systems can proceed. It is for these reasons that Wand and Weber began to seek a "common paradigm" (Wand, 1989) for systems analysis and design.

Wand and Weber use a systems ontology to assess the modelling power of various methodologies in terms of a set of constructs based on an ontology, and the completeness of a constructed model can be assessed. We expand on these concepts below.

Wand and Weber argue "that systems analysis and design methods can be characterized in terms of:

1) the origin of the modelling constructs used in the methodology; and
2) how the methodology deals with structure and behaviour." (Wand, 1989)

We discuss each point. Firstly, they assert that a modelling method can be either information systems driven or reality driven, with the constructs being borrowed from the respective domains. Wand and Weber see this as a dilemma, in the sense that the modelling needs of one or other domain will be compromised by whichever approach is chosen, due to the fact that a modelling technique uses constructs that are from either the information systems domain or reality. Wand and Weber propose that each modelling method can be categorised as being either information systems driven, or reality driven. Secondly, it is recognised that models of both statics and dynamics are required.

Upon identifying a need for a general theory for analysis and design, Wand and Weber propose that it is useful to use the same constructs for both domains, and that ontology is a useful place from which to get such a set of constructs. In their work, they analyse methods for strengths, weaknesses, and completeness of system description "in terms of an ontologically-based common set of constructs" (Wand, 1989). We pause to point out that we too desire to regularise the things used to construct models so that both domains can be satisfied, and we are encouraged by their use of ontology for the purpose.

In the view of Wand and Weber (Wand, 1989), an analysis and design process creates an information system design that can be implemented. An information system is a human-centred creation of a real-world system as perceived by somebody, and built to deal with information processing functions in organisations. The process of constructing an information system is a transformation from human perceptions to an artifact *representing* these perceptions. This process involves three transformations. Firstly, analysis that creates, from human perceptions, a model of (perceived) reality. Secondly, design that transforms a model of reality to a model of the information system. Thirdly, implementation that transforms a model of an information system to an implemented artifact. It is in the first transformation that work by Wand and Weber has implications for our work, that is, the part dealing with the creation of a model of (perceived) reality. Much of their work deals with this first part of the process of analysis and design.

Their early work into the first aspect of analysis and design (Wand, 1989) resulted in an informal comparison between tools and the ontology, together with a categorisation of the domain in which the studied tools reside. The domain is categorised as being either information system, or real world. This informal comparison later progressed to be a more formalised (Wand and Weber, 1993) understanding of the representational clarity with which a model of reality is created from human perceptions using a specific tool. In this, they examine the grammars that information systems analysis and design methods provide to describe aspects of the real world. A grammar in this context "generates a language, which is a set of strings over some alphabet. ... In these grammars, sentences provide a graphical representation of some real-world phenomena" (Wand and Weber, 1993). Further, two mappings between the ontological constructs and the tool can be described. Firstly, a representation mapping, from a grammar representing the ontology (or real-world view) to the grammar representing the tool from analysis and design method (or model of reality) can be described. Secondly, an interpretation mapping from the tool's grammar to the ontology (or real-world view). One can then analyse each of the mappings to describe the quality of the representation and interpretation quality of the models that the analysis and design tool is likely to produce. Figure 1 shows the two mappings.

These mappings can be used to define terms to specify how well constructs from each map into the other. There is the idea of *ontological expressiveness* (Wand, 1996) that combines the ideas of ontological completeness and ontological clarity.

Figure 1—Ontology and Systems Analysis and Design Grammars: Representation and Interpretation Mapping
(Wand and Weber, 1993, Wand, 1996)

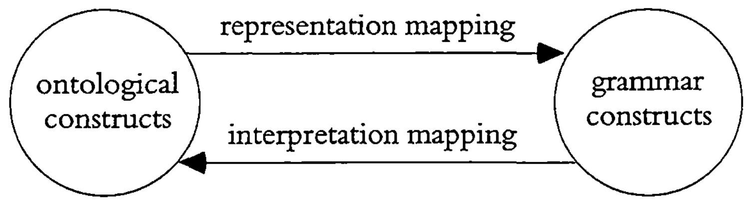

Firstly, *ontological completeness* is where a grammar from the analysis tool perfectly represents concepts from the grammar of the ontology. This is shown in Figure 2a, and is where all constructs of the ontology grammar are mapped to constructs of the analysis and design grammar. In contrast, *ontological incompleteness*, also known as construct deficit is where a construct from the ontology grammar has no construct into which it maps. This is depicted in Figure 2b.

Figure 2—Ontological Completeness
(Wand and Weber, 1993)

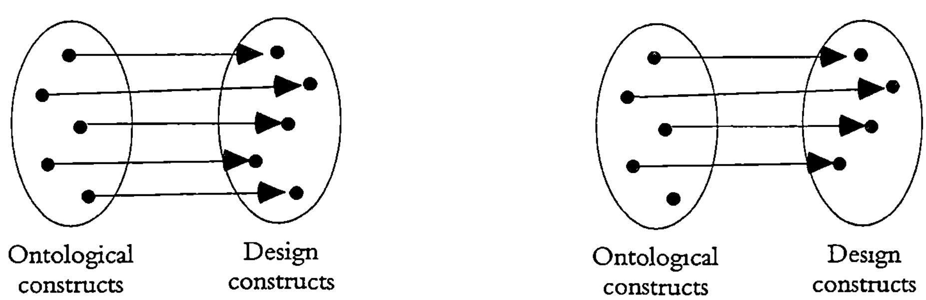

a—Ontological completeness **b—Ontological incompleteness or construct deficit**

Secondly, *ontological clarity* concerns the quality of the interpretation mapping. These are *construct overload* (Figure 3 a), *construct redundancy* (Figure 3 b), and *construct excess* (Figure 3 c). Construct overload is where several ontological constructs are mapped onto the single construct in the grammar. Construct redundancy is where several constructs from the tool's grammar map onto a single construct from the ontology's grammar. Construct excess occurs when the tool's grammar has a surplus of constructs with respect to the ontology's grammar. In other words, this is where, for a construct from the tool's grammar, there cannot be found an equivalent in constructs of the ontology's grammar. As referred to above, Figure 3a–c shows the various concepts associated with ontological clarity.

Wand and Weber use concepts of ontological clarity and ontological completeness to establish a measure of the ontological expressiveness of an analysis and design tool as represented by its grammar as it compares with the grammar representing the ontology. Two analysis and design grammars are deemed to be equally expressive if they achieve the same measure of ontological completeness and clarity. That is, with respect to an ontology, each has the same measures of completeness and clarity. Wand and Weber make it clear (and this seems reasonable) that if the grammar's ontological clarity is deficient in some way, then the grammar must be modified to create, effectively, a new grammar. Further, they make the point that "a grammar that appears to be ontologically incomplete, may achieve completeness through being extensible. If a grammar is extensible, constructs in the grammar can be combined via production rules to form other constructs." (Wand and Weber, 1993)

Figure 3—Ontological Clarity

(Wand and Weber, 1993)

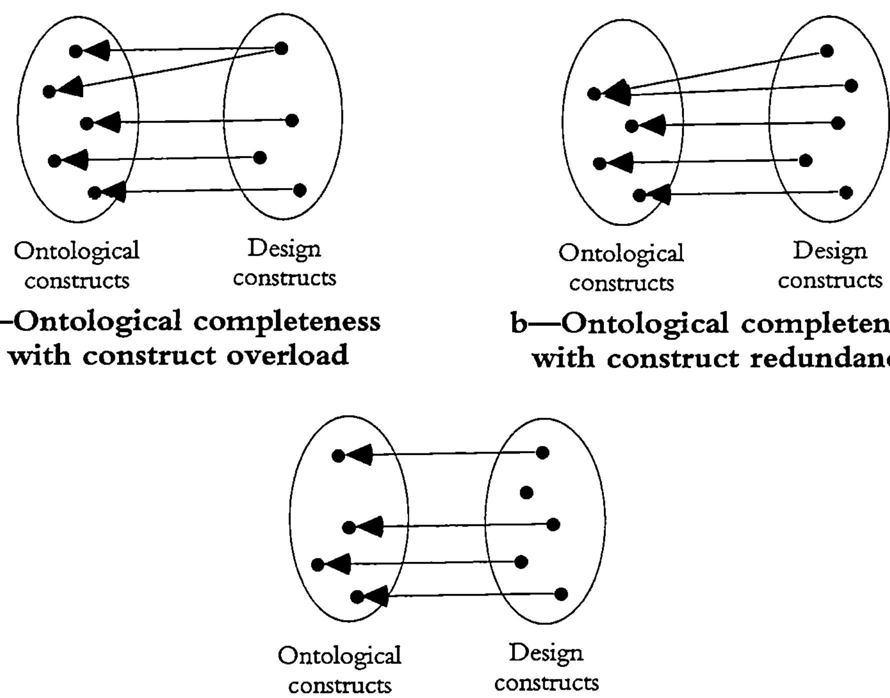

a—Ontological completeness with construct overload

b—Ontological completeness with construct redundancy

c—Ontological completeness with construct excess

Despite claiming to discuss the capability of tools to capture the same sets of semantics, the comparison is not on the basis of semantics of terms. Instead it is on the basis of structure of terms, or grammars that can be used to combine terms. It is the limitation presented by the grammar that limits the comparison. Further, these notions do not explore the more fundamental question of the differences or similarities in world view between the ontology and the various tools under examination. That is, these notions do not examine the differences in semantics between the various terms found in the ontology and in the tools. Instead a comparison of the mathematical constructs is performed. The research compares the differences between the ontology and modelling framework on the basis of constructs matching or there being a difference in the number and style of constructs in the modelling framework with respect to the ontology. Their approach also requires that the ontology selected be precise and defined mathematically. It is debatable whether formalisation really gains us anything in the context of a semantic comparison that we desire. It also requires that the modelling tool under investigation similarly be precisely defined. Formalising the tool in a grammar may compromise the meaning attached to parts of the tool and the meaning of constructs and concepts may be restricted by the formalising process. Our research differs from this in that we use the world view of the tools (in our case the data modelling frameworks) and compare this with the world view as encapsulated in an ontology. Given that the data modelling tools are imprecise and at best are semi-formal, a more pragmatic and qualitative comparison of world views is required than has been provided by this previous work.

Conceptual Modelling and Ontology

A more recent effort in an area related to modelling tools, is the work by Yair Wand and others concerning potential theoretical foundations for conceptual modelling in information systems development (Wand et al., 1995, Wand, 1996). One of the three foundations, as outlined in this work, is ontology. They argue that ontology, concept theory, and speech act theory all are potentially theoretical bases for conceptual modelling. Each is considered in turn in their work. Our interest is in their commentary on the use of ontology in this context because conceptual modelling is directly relevant with respect to the use of data modelling frameworks (or data models) and so is directly relevant to our work.

Conceptual modelling, according to work cited by Yair Wand and associates, is variously described as "a formal description of some aspects of the physical or social reality for the purpose of understanding and communicating" (Wand et al., 1995), or "an abstract model of the enterprise" (Wand et al., 1995). "A conceptual model should reflect knowledge about the application domain rather than about the implementation of the information system" (Wand et al., 1995). Clearly, from this, it is the construction of a model of an application domain that is critical. The implementation from the constructed models is not of interest to us here.

An important aspect of their 1995 paper (Wand et al., 1995) is that it considers a critical question with respect to conceptual modelling languages:

> "...how does one go about defining an 'appropriate' conceptual modelling language? The power of a modelling language lies in the semantics of its constructs. The semantics defines what individual constructs and their combinations 'stand for' in the modelled domain." (Wand et al., 1995)

Their paper considers each theory in turn and discusses in which ways they can be combined for the process of conceptual modelling. We seek to establish a unifying framework that embodies the semantics of what is required in order to understand 'what there is'. However, our intention is to consider only ontology, one class of modelling language that is mentioned in their work. We are interested in data modelling frameworks from both the semantic data modelling class and from the object modelling class of data modelling frameworks. It is for this reason that we note the statement "[t]o create a conceptual model we need a set of constructs to model real-world systems. For this we turn to ontology which can be defined as the branch of philosophy which deals with the order and structure of reality" (Wand et al., 1995). Once again, there is agreement with the approach of using ontology to examine what is needed to model reality. We need to consider their experiences, and address any issues they felt to be important. Our research must also complement their experiences, or it should extend their research for specific purposes.

The researchers mention three specific limitations of using ontology in the way envisaged (Wand et al., 1995). We paraphrase them here. Firstly, there is no ontology that is generally accepted. Secondly, that ontological models seem to assume an objective reality, while the world is only known through human perceptions. Thirdly, the specific model selected by Wand and his colleagues does not deal with the organisational and behavioural aspects of information systems.

In analysing these criticisms, several things can be said. Firstly, there is no measure of absolute 'goodness' of ontologies. It is therefore important to fully justify the appropriateness of using a particular ontology for a specific purpose. Further, each ontology has a heritage that shows its attitude to how reality is perceived. The ontology selected by Wand and Weber is a scientific ontology (Bunge, 1977). In this thesis, we choose an ontology different from Bunge's and justify that choice pragmatically. Conceptual modelling concerns modelling reality, and as such a unification of the world view of a data modelling framework with that from an ontology seems a useful goal. This differs from the approaches taken by Wand and Weber in that they have not studied ontology from the aspect of data modelling frameworks generally. Finally, if a consideration is taken to the philosophical heritage of an ontology, the second criticism from the paragraph above may also be ameliorated. This is because depending upon the philosophical heritage that its author holds to be true, an ontology may allow for a less 'objective' view of 'what there is'.

We believe that ontology is useful for our goals because it is independent of technology, the language contemplated, and is a mature field of study. The work by Yair Wand and associates supports our view and they have found ontology useful for related work.

Ontological Studies of Semantic Data and Object Models

Wand and Weber have studied data modelling frameworks in more depth than has been discussed above. They have examined relationships as used ER modelling (Chen, 1976, Wand et al., 1993). They have also examined the modelling methodology called NIAM (Nijssen and Halpin, 1989, Weber and Zhang, 1996), and explored a theoretical basis for object modelling frameworks (Wand and Woo, 1993, Wand, 1989). Further, they have considered in more depth the concept of class in object modelling frameworks (Parsons and Wand, 1993). The consideration of object modelling techniques and frameworks comes from criticism that object modelling lacks a theoretical basis (Li, 1991).

Bunge's ontology has been used to more deeply understand the appropriate uses of relationship when applying the ER modelling tool (Wand et al., 1993). In this work practitioners' use of relationship (particularly) from ER is analysed in the context of Bunge's ontology. Their analysis supports a precise definition of the relationship construct, and assistance is formulated, in the form of guidelines for practitioners wishing to better use relationships in projects, indicating when and where to use relationship as opposed to other elements of the ER modelling tool. The ontology does not have an element that corresponds to relationship. Instead, it is argued that a relationship is really one of two things. Firstly, it may represent a shared property between two entities. Secondly, it may represent a part-whole relationship between entities (otherwise known as aggregation or composition). This differs from our own research in that it studies a very specific element of a particular data modelling framework. Nevertheless, it provides more support for recognising distinct world views that are contained in the terms defined in ontologies and data modelling frameworks, and it also recognises that there is a level deeper than structure at which comparison may take place. To contrast with our work, we are attempting to use ontology as a unifying framework that can be applied in comparing and evaluating data modelling frameworks. Their work supports the combination of ontology and data modelling frameworks, indicated by comments such as "we believe the work demonstrates potential benefits of applying concepts from ontology to database design" (Wand et al., 1993). This would obviously flow through to considering ontology as a possible unifying framework between the reasonably large number of data modelling frameworks being used today.

More recently, Bunge's ontology has been used to evaluate NIAM (Weber and Zhang, 1996). In this research the comparison is in the form of grammatical comparison in a way similar to work reported in the preceding subsections (Wand and Weber, 1989, Wand and Weber, 1993, Wand, 1996), that is, in comparing structural elements from the ontology and from the selected modelling tool used in either conceptual modelling or systems analysis and design. Further, the deeper analysis of the semantic rather than structural differences and similarities is absent. There is no recognition of the implied world view of the data modelling framework and how it differs from or is similar to the world view as represented in the ontology. This is alluded to in the concluding remarks:

> "[i]mplicit in our prediction is a 'naïve' theory of semantics that attributes clearer meaning to one-to-one mappings between ontological constructs and grammatical constructs." (Weber and Zhang, 1996)

In this context, a construct is an element of the grammar representing the modelling tool (in this instance NIAM). This shortcoming is addressed in our approach, and we endeavour to compare the similarities and differences in world view as represented by the ontology and a modelling framework using concepts. The world view of a modelling framework or an ontology is in the language that is used to define or describe constructs (or what we call terms) of each. It is therefore through the concepts giving meaning to terms that the world view of modelling frameworks and ontologies can be compared or related at the semantic level.

The final aspect of the work coming from the use of Bunge's ontology is that concerning object modelling (Parsons and Wand, 1993, Wand, 1989, Wand and Woo, 1993). There has been a degree of apprehension that object modelling is not based on any strong theoretical bedrock (Li, 1991, Atkinson et al., 1989). Some object models represent little more than a baroque (or greatly decorated) entity-attribute-relationship (Griethuysen, 1982) style of modelling framework, or utilise tools that are not 'object oriented'. Additionally, there are divergent approaches to object modelling in the database systems community, the programming community, and the object systems community. Part of the early work by Wand and others has been in establishing a "formal model for objects" through ontology. This is closest to our objectives, in the sense that it is a unifying theory for object models (Wand, 1989, Takagaki and Wand, 1991). Further, related work concerns the formulation of rules or guidelines to aid the construction of object models (Wand and Woo, 1993). Three things can be said about this work in relation to our objectives. Firstly, since the formalisms and the analysis come from a specific ontology, namely Bunge's adapted ontology, we do not discuss the formalisms *per se* since a comparison of specific ontologies is beyond the scope of the thesis. However, there is a detailed discussion for our reflections later in the thesis where we relate our experiences with object modelling and the ontology selected for our work. Our objectives also lie with data modelling frameworks beyond object data models (semantic data modelling frameworks). However, we gain support, through their work, for using ontology in the work we are undertaking. Thirdly, in the next sub-section, we discuss the ontology used by Wand and Weber and subsequently by related researchers, and briefly discuss our method for selecting an ontology.

The Ontology Used in Past Studies (Bunge's Ontology)

Wand and Weber, in their work (and in work by others), have adapted an ontology by Mario Bunge (Bunge, 1977). Conceived and written during the late 1960s and 1970s, Bunge's ontology is exact, scientific, and well developed. In this work, Bunge is motivated to make more exact the ontological and metaphysical aspects of science. Part of his research during that decade or so is his ontology, which is a complete theory to encompass physical (or concrete) reality.

Reported in two volumes (Bunge, 1977, Bunge, 1979), the ontology is mathematically rigorous and highly formalised. In motivating his work, Bunge cites as a critical factor his dislike for most of the metaphysical work to that date, and he describes certain schools of philosophy in derisory terms. He keenly seeks an ontology that, once and for all, sets aside 'woolly polemics' and utilises formal tools and scientific principles in setting out 'what there is' in reality. By way of introduction, Bunge presents an essay (pp. 1–25 in Volume I) on what he sees as the key misconceptions that abound in metaphysics and where 'scientific ontology' can correct them. In the essay he betrays a keen objection to anything that is not 'physical' or part of the 'concrete world'. Only 'objects [that] can be studied scientifically' should be included in any theory about metaphysics.

Bunge states in his introductory essay that there are many different styles of ontology from which to choose, and he presents nine different philosophical approaches to the study of ontology. Realistically, for our work and that of Wand and Weber, very few of the styles of ontology are appropriate. However, it is necessary to make a selection of an ontology from the wide range available. There exist many debates concerning the approach that is embodied in any philosophical work, and ontology is no exception. We need to consider the selection made by Wand and Weber, and place our selection in that context. We discuss our pragmatic criteria for selection later in the thesis.

No ontology can claim to be the 'best' or even a 'good' one. It is therefore critical to select an ontology that suits one's needs. Wand and Weber have done that, citing several reasons for selecting Bunge's ontology. Bunge's ontology is comprehensive, well-formalised, and refers extensively to previous work in ontology (and philosophy of science). Further, Bunge's ontology is based on systems. It is also described by Bunge as being a scientific ontology that does not indulge in endless debate over esoteric issues and limits itself to concrete or physical modelling, and it also holds firmly that reality is determined by tight and formalised theory.

The requirement that the ontology be system based is not really of importance, because all of reality can be characterised as systems. For this reason, one could describe the systems to which Bunge's ontology is applied using many ontologies. Again, the fact that the ontology is formalised is not of particular importance. Formalisation helps us in situations where all the subjects of research are formalised using the same axioms or bases. In this research some of the subjects are semi-formal or appear formalised, but are in fact so in a superficial or contrived way, and the formalisation present does not progress beyond the structural level. Finally, the fact that in describing an ontology a philosopher engages in debate with others in his or her field is irrelevant to our considerations. Central to our deliberations is the suitability of the ontology to the goals of our research. Whether or not an ontology, as it is presented, discusses work by other philosophers is simply a side-issue and should not become a catalyst for distraction. In fact, selecting an ontology that is dogmatic in its assertion that all things behave according to preexisting theory may be dangerous.

We are a little more relaxed with our requirements than Wand and Weber, and we feel that selection should be based upon fitness for purpose. We are seeking a unifying framework for data modelling frameworks. It therefore makes sense that the selected ontology discuss items that are similar to those present in most data modelling frameworks. Further, we do not feel that it is essential to select a systems based ontology. We are keen to seek unification of data modelling frameworks, and these being semi-formal tools, it is therefore reasonable to widen the selection of ontology from those that are formalised to those that are less formal. We also seek to investigate ontologies that concern more than simply physical reality. This contrasts with Bunge:

> "... we maintain that the ontologist should stake out the main traits of the real world as known through science, and that he should proceed in a clear and systematic way. He should recognize, analyse and interrelate those concepts enabling him to produce a unified picture of reality."

Essentially, Bunge desires a unifying theory that brings together theories of the physical world as known through contemporary science. Moreover, his understanding of metaphysics is completely divorced from the more arbitrary nature of human existence.

It is also worth noting that Wand and Weber encourage use of alternative ontologies. This is partly motivated by the fact that no ontology can claim to be the best. They state that "... it would be of interest to examine other possible ontological approaches and compare their outcomes with those of the work described here."

(Wand, 1996) This invitation is further motivated through criticism of Bunge's ontology: "[o]ne possible criticism of the work is that Bunge's ontology is oriented towards the physical world and therefore does not provide for human perceptions and social context" (Wand, 1996). Some ontologies address the criticism relating to human perceptions through specific mechanisms, and groups of people can agree on an interpretation which can then be described. Clearly then, there is merit in considering an alternative ontology from the adapted Bunge ontology. It may be that the selected ontology is from a different school of philosophy from Mario Bunge's.

We intend to consider an ontology different from Bunge's, that we select based upon a number of pragmatic criteria (Dancy and Sosa, 1992). The selected ontology must be one that helps explain the statics and dynamics of reality and that is useful for our purposes. Chapter 4 contains a full discussion of the selection criteria for the ontology. After selecting an ontology we discuss the philosophical heritage of the selected ontology and whether it is compatible with our use of the selected ontology.

2·2 Data Modelling Frameworks

In this section we consider data modelling frameworks in the context of our goals. However, before doing so we need to define the term 'data modelling framework' so that we can use it with clarity in this thesis. Following this, we place our research in the context of the literature concerning the search for principles unifying data modelling frameworks, before selecting representative data modelling frameworks from the semantic and object data modelling classifications.

The term 'data model' means different things in different contexts. We can conceive of at least three meanings. Firstly, a database management system has a data model associated with its architecture that is adhered to by schemas of databases stored. For example, a relational database management system only allows schemas that adhere to the relational data model. Secondly, a specific conceptual model that is constructed and gives rise to a given schema can also be called a data model. For example, a university database would have a data model that describes significant entities in the university. For example, an entity called "The Faculty of Science", or "The Department of Philosophy." This data model could then be transformed and be implemented in a specific database management system. In fact, there could be many examples of schemas to describe a university (or even the same university) constructed by different practitioners. Thirdly, data model is referred to as the modelling framework that is used to construct conceptual models for a range of database systems or conceptual models of data often as part of systems analysis and design methodologies. The entity-relationship model, functional data model, and other similar models are examples of these. It is this third meaning that we intend in this thesis. In other words we are interested in data models that are used to construct specific conceptual models. That is, at the meta-level with respect to specific conceptual models. Further, to avoid confusion, the phrase 'data modelling framework' is used in this thesis whenever we refer to this idea. When quoting past work data model (unless otherwise stated) also refers to data modelling framework in the way that we have defined it here.

Definition 2·2—Data Modelling Framework

> A data modelling framework is a modelling tool that is used to construct conceptual models for a range of database systems or used to construct conceptual models of data as part of systems analysis and design methodologies.

There are many data modelling frameworks (Cárdenas and McLeod, 1990, Hull and King, 1987, Hull and King, 1990, Peckham and Maryanski, 1988), and despite this, there are recurring themes between modelling frameworks. Many discuss entities (Kent, 1978) or objects (Wand, 1989) and group them into entity/object classes, or entity/object types (King and McLeod, 1985). Modelling frameworks also allow the definition of attributes or properties common to entity classes, together with links or relationships between members of entity classes (Cárdenas and McLeod, 1990, Hull and King, 1987, Hull and King, 1990, Peckham and Maryanski, 1988, Kent, 1978). Since the advent of the first semantic data modelling frameworks in the 1970s (Stonebraker, 1994), there has been the occasional attempt to survey and to compare data modelling frameworks, but there has been little by way of a search for a unifying framework for data modelling frameworks, that is, a unifying framework based on sources independent from the data modelling frameworks themselves.

One of the first data modelling frameworks to be used for conceptual modelling is the relational model (Codd, 1970, Date, 1995). It frees designers from considering implementation details, such as file structures and access paths, when designing conceptual models. Nevertheless, it provides tools based on sound theoretical techniques that can be used in an implementation environment, and the mapping from a conceptual model to an implementation model is reasonably straightforward. However, there is agreement that problems exist with the relational model when used for conceptual modelling (Kent, 1979, Hull and King, 1987, Cattell, 1991, Peckham and Maryanski, 1988), in that the concept *relation* in the relational model is overloaded and is being used to represent both relationships and entities (King and McLeod, 1985). This makes it difficult for clients to make sense of the conceptual models constructed using the relational model. Further, the semantic overload increases the chance of mistakes being made or information failing to be modelled. Real world concepts are often difficult to represent using the relational model due to its inflexibility. These problems led to the development of the entity relationship (ER) modelling framework (Chen, 1976).

The ER modelling framework is one of the most popular conceptual modelling frameworks and is widely used in industry. The ER modelling framework was proposed by Chen (Chen, 1976) to abstract from the relational model, providing a richer set of terms for modelling, and, by providing separate terms for each of *relationship* and *entity*, alleviating the problems outlined above.

The ER modelling framework is part of a larger group of frameworks known as semantic data modelling frameworks (Stonebraker, 1994), that differ from the relational modelling framework in that an increased range of terms are available from which specific models can be constructed. The extra terms are all aimed at providing a richer set of primitives from which models of reality can be constructed and in order to increase modelling power. The number of modelling frameworks conceived under the category of semantic data modelling frameworks is large (Peckham and Maryanski, 1988, Mylopoulos, 1998). Some of the modelling techniques applied in these frameworks were adapted from programming languages and from artificial intelligence. The purpose of these modelling frameworks remains that of constructing representations or models of things from the real world that can be transformed into implementation models for specific database management systems or, more generally, information systems implementation environments.

The object oriented group of data modelling frameworks is sometimes described as a separate category of data modelling framework. These had their genesis partly in object oriented programming languages such as Smalltalk (Goldberg and Robson, 1983) and partly in research into abstract data types and artificial intelligence (Cattell, 1991, Hull and King, 1987, Hull and King, 1990, Peckham and Maryanski, 1988). In part, this class of data modelling framework is motivated by the requirement for extra modelling power found in modern applications such as *part design* (engineering), *document handling*, and *software engineering* (Cattell, 1991). Others see them as a variant on the semantic data modelling framework theme that in turn inherited various concepts from other sources (Brodie, 1982). Regardless of their categorisation, they can be seen as an important class of data modelling frameworks that must be considered.

Object modelling frameworks bring with them a range of concepts that have a degree of commonality with semantic data modelling frameworks. They discuss objects (similar to entities), attributes or properties, relationships or associations, classes and hierarchies. Further, according to their genesis (data modelling research *vs* object oriented programming) other aspects may be present, such as the need for explicit message passing between objects, or for operations and other dynamic features that are rooted in program execution. Some claim, without substantiation, that the concepts naturally reflect that which happens in reality, or that they are a natural way to model a given situation.

Increasingly, data modelling frameworks are being created with extra features being added or facilities being removed almost at whim, thereby emphasising the fragmented nature of the genesis of modelling frameworks. Some claim that a specific data modelling framework is natural or in some way better than others. Despite these problems, there has been little work undertaken to establish a unifying framework through which data modelling frameworks can be better understood or indeed more rigorously studied. Further, there has been little attempt to utilise reference disciplines that are mature and that deal with related areas, and those that have tried to establish a unifying framework usually do so by gathering the most popular features as found in data modelling frameworks (King and McLeod, 1985). There have been at least two examples of survey and classification work undertaken. In the late 1980s two surveys discuss and compare semantic data modelling frameworks (Hull and King, 1987, Peckham and Maryanski, 1988), and we discuss each of these in turn. We expand upon the second survey (Peckham and Maryanski, 1988) before discussing other related research that is more recent.

In the first survey (Hull and King, 1987, Hull and King, 1990) semantic data modelling frameworks are considered from the perspective of what they call the philosophy common to many frameworks in creating data models of enterprises. In this work (reprinted in an updated form in 1990) several specific modelling frameworks are explored and compared. Hull and King find that semantic data models have changed in that their focus had shifted from being used "to facilitate the design of database schemas" (Hull and King, 1990) to instead "[allowing] database designers to think of data in ways which correlate more directly to how data arise in the world" (Hull and King, 1990). Further, they generalise the typical semantic data modelling frameworks using a constructed modelling framework that they call GSM (Generic Semantic Model). It is constructed in the sense that GSM is loosely based on several modelling frameworks to illustrate key concepts. Additional to their consideration is a discussion of the parallels between semantic data modelling frameworks and each of object oriented programming languages and artificial intelligence modelling tools.

For our aims, it is important to note that their summary of key parts of semantic data modelling frameworks shows the commonality between many data modelling frameworks with respect to terms and concepts, in that objects or entities, attributes, relationships, and classes are often elements of these data modelling frameworks. Our research takes this a step further in seeking a candidate unifying framework from a recognised theoretical reference-point. We seek to use an ontology as a reference-point to establish a unifying framework that may help to bind the various semantic and object data modelling frameworks.

The second significant survey (Peckham and Maryanski, 1988) is different from the first and is more important for our research goals.

> "Hull and King (1987) take a slightly different approach: A model constructed with fundamental semantic data modeling components is used for comparison with each model. The Hull and King paper has a significant tutorial flavor, addressing the issue of semantic database modeling using a pedagogical data model. That work emphasises implementation aspects of database systems developed around semantic models, whereas this paper focuses more on conceptual modeling issues." (Peckham and Maryanski, 1988)

The more recent survey clearly recognises the conceptual modelling role that is increasingly being played by semantic data modelling frameworks, in that they are modelling reality. "The piece of the 'real world' that is represented by the database is commonly called an *enterprise*" (Peckham and Maryanski, 1988). Later they assert that

> "[t]his survey considers conceptual semantic models that make use of entities, relationships, and constraints to describe static, dynamic, and temporal qualities of an enterprise. The desired result is a representation of the enterprise that closely parallels the user's perception, without concern for the physical model." (Peckham and Maryanski, 1988)

They also note important contributions from related areas such as artificial intelligence and psychology in promoting key features in semantic database modelling frameworks, before proceeding to outline and describe certain central modelling concepts that are found in many semantic data modelling frameworks. These are entity/object, attribute, relationship, aggregation, generalisation, and similar concepts. In contrast with Hull and King, Peckham and Maryanski proceed to compare data modelling frameworks on the basis of the core concepts that they outlined as being common to most semantic data modelling frameworks. Reflecting on the results of their comparison they summarise

> "[m]ost agree that the following contribute to this objective [to facilitate the modeling of and the use of databases]:
>
> (1) A semantic model should provide relationships between data objects that support the manner in which the user perceives the real-world enterprise.
>
> (2) For these relationships a semantic model should contain semantics that specify the acceptable states, transitions, and responses of the database system." (Peckham and Maryanski, 1988)

They also found that there are significant levels of dispute concerning the independence of specific modelling frameworks from their implementation environments, and concerning the specific nature of relationships and levels of abstraction that are appropriate (such as generalisation, aggregation and so on). Further, few of the models provide for dynamic or event-related elements and they identify this as a key area for future research.

In the context of our research objectives, this work corroborates our approach in that the authors advocate the important and increasing role that these data modelling frameworks play with respect to modelling reality in terms of both dynamic and statics. Further, that implementation of a database (or information system) follows an understanding of data as it emerges from the real world, thus emphasising the importance of understanding part of the real world before recommending key implementation strategies, structures, and relevant efficiencies (as provided for in the implementation environment). In addition, their research has found unifying threads from the modelling frameworks themselves. In contrast, we seek a rationale for the unifying threads or at least a framework in which to study and explain them. Further, our comparison and evaluation is through the world view from the selected ontology and the world view as found in specific data modelling frameworks.

Clearly, the fact that semantic data modelling frameworks and indeed, object data modelling frameworks contain many common features is not in dispute. What we seek is a unifying framework that can be used to validate the approaches taken given the shift in focus from purely efficient database implementations that are (hopefully) accurate with respect to reality to a more balanced view that incorporates the importance of modelling reality before implementation. It is for this reason that we seek a unifying framework from ontology, or study of world views, a mature reference-point from philosophy.

Although not directly motivated from the viewpoint of data modelling, more recently, Mylopoulos examines a different aspect of the same family of data modelling frameworks. Mylopoulos (Mylopoulos, 1998) recognises that modelling techniques and more specifically the modelling frameworks themselves have implied ontologies. Some of the frameworks studied are data modelling frameworks. For example, ER assumes that in the problem domain (or what he calls the application domain), entities exist with relationships between entities. Further, that entities are grouped into classes with entities in those classes possessing like attributes. Other modelling frameworks would have different (some only slightly different) views of the application domain.

Mylopolous goes on and describes the categories into which data modelling frameworks would fit according to the ontology implied by an application domain. He identifies four ontologies—static ontology, dynamic ontology, intentional ontology, social ontology. Instead of using an ontology as the starting point for an objective theory, he categorises modelling frameworks according to the ontology of application in which the framework is often used. That is, they are categorised according to the ontology possessed by the modelling framework which in turn is categorised by the style of application to which it is put. This approach does not provide or suggest a unifying theory for data modelling frameworks generally. However, it is a recognition that there is an ontology implied by every data modelling framework, and in this sense ties in with our objectives, and it is a more useful categorisation of data modelling frameworks than has been constructed to date (Peckham and Maryanski, 1988) in that it recognises ontology (world view) as being important in categorising data modelling frameworks rather than simply calling it a style or philosophy of modelling (Hull and King, 1987). Nevertheless, this then begs the question of the origin of the ontology for each modelling framework. Moreover, it raises further questions regarding the justification or conscious reason for establishing these ontologies, and it also suggests that a search for a unifying framework based on an ontology from philosophy would be worthwhile pursuing.

For completeness, a number of data modelling frameworks from several categories are studied in this thesis. These span the breadth of data modelling frameworks, and are taken from the entity-attribute-association category (Griethuysen, 1982), the semantic data modelling category, and the object-oriented category. These categories are recognised as being significant in the literature (Hull and King, 1987, Cattell, 1991, Peckham and Maryanski, 1988). From these categories we select specific data modelling frameworks that are both popular in the literature, and are considered to be representative of the styles concerned.

We include five data modelling frameworks. Firstly, the entity-relationship model (Chen, 1976), is one of the first to begin to improve upon the relational model and is still extremely popular in industry. Secondly, the functional data model (Shipman, 1981) is one of the cleanest example of the functional-style of semantic data model, and it fits well with functional programming languages, a powerful computational model. Thirdly, we include the semantic data model (Hammer and McLeod, 1981) that is recognised to be an important modelling framework (Hull and King, 1987, Peckham and Maryanski, 1988). Fourthly, NIAM (Nijssen and Halpin, 1989). Finally, the object data model specified in OMT (Blaha and Premerlani, 1998) is included as a significant object model that is used in contemporary information systems development yet, in part, originating from the data modelling community.

In the following chapter we introduce the methods that are designed to address our research questions. In Chapter 4 we present the pragmatic selection of the ontology used in this thesis. In Chapter 5, we describe the specific data modelling frameworks used in this thesis.

Chapter 3

Using Terms and Concepts in Ontological Comparisons and Evaluations of Data Modelling Frameworks

In this chapter we consider the roles that an ontology can play in the analysis and design of data modelling frameworks. Recall from Chapter 2 that we have found motivation in the literature, through promising results in related areas, for using ontology in this way.

We envisage using an ontology in its role as a unifying framework in two ways. Firstly, we believe that an ontology can be used to compare modelling frameworks. Secondly, an ontology can be used to evaluate and suggest areas in which to strengthen modelling frameworks. It is through applications of these two methods that we propose to consider our research questions. We describe the methods in detail in this chapter, and report results from applications of these methods in chapters 6, 7, and 8.

Analysis of results is critical for any research method. It is for this reason that we also cover the analysis of results for the methods we describe. Not only does this indicate how analysis is to proceed, but also shows what sorts of conclusions can and cannot be drawn from applying these methods. The methods rely on terms and concepts from the ontology and data modelling frameworks respectively. It is for this reason that we define what we mean by 'term' and 'concept' and explore their role in the context of this study. It is the meaning behind terms revealed through concepts that constitute the various world views that can be found in an ontology or in various data modelling frameworks. By comparing the meaning of terms, rather than just structure or formalism, we open up consideration to include ontologies that have been rejected in information systems research previously due to their descriptive and often informal or semi-formal nature.

It is by using an ontology to compare data modelling frameworks and to evaluate specific data modelling frameworks that we can gain insight into the potential of using ontology as a unifying framework for data modelling frameworks. Further, we expect that through applying these methods we will build upon previous work in information systems that uses ontology.

We begin this chapter by considering the use of terms and concepts to relate an ontology with a number of data modelling frameworks. In doing this we utilise a recognised semiotic theory by Umberto Eco (Eco, 1976) to explain the relationships that exists between terms and concepts. We then describe our comparative and evaluative methods.

3·1 Relating Ontologies and Data Modelling Frameworks Through Terms and Concepts

We propose to use an ontology to compare and evaluate data modelling frameworks. We need to find a basis upon which to relate the world views embodied in the ontology and in the various data modelling frameworks.

An ontology provides us with the 'furniture of the world' and shows us the fundamental categories of things that exist in reality. It also describes basic terms that are needed to represent reality and relates these to the fundamental categories. Data modelling frameworks also provide us with the terms with which we build models of reality. Our aim is to relate the views of reality given by the ontology and data modelling frameworks in order to compare and improve data modelling frameworks.

We propose to conduct a comparison and an evaluation based upon terms as found in ontologies and data modelling frameworks and the concepts they embody. We wish to do this because data modelling frameworks and ontologies are often informal or semi-formal. Further, there is semantic depth contained in the descriptions of ontologies and data modelling frameworks. It is therefore insufficient or, at times, inappropriate to rely upon a mathematically formalised representation of an ontology or data modelling framework if that means that semantics are lost or in some way abbreviated. Further, formalisation in the form of grammars concentrates on structural issues while not acknowledging the role of semantics. Clearly, the challenge is to design methods that are rigorous and systematic so that they can be repeated by others, but also be flexible enough to capture the qualitative nature of the analysis that we wish to make.

There are three parts to describing the methods that we use. Firstly, we describe the extraction of terms and concepts from the ontology and data modelling frameworks. Each term is given meaning through the concept to which it refers and which is given in the descriptions of the modelling frameworks or the ontology. Concept is a recognised philosophical term that we introduce below. Secondly, we define the nature of the relationship between the concept that gives meaning to a term and the term itself. Further, we define the way we will refer to concepts and terms later in the thesis. Essentially, these are terminological issues. We use semiotics (Eco, 1976) to regularise terminology and intent. Thirdly, we describe the two methods themselves. We further utilise semiotics when analysing the results of the comparative and evaluative methods.

We describe the methods later in this chapter. We discuss the extraction of terms and concepts and related issues in the remainder of this section.

Extracting Terms and Concepts

The ontology and data modelling frameworks that we consider in this thesis consist of terms. It is through terms that the ontology and the data modelling frameworks can be used to describe, discuss, and analyse reality. For example, ER utilises terms such as 'entity', 'entity class', 'relationship', and 'attribute', in order to describe, discuss, and analyse reality. The term 'entity' is used to discuss things from the world that are distinct and don't depend upon other things for their existence. My cat is an example of an 'entity'. I, the author of this thesis, am an example of an entity. Things categorised as 'entities' can have 'relationships' with other entities. My cat has a relationship with me—she is owned by me. The point is that data modelling frameworks and appropriately selected ontologies have similar terms that we can use to describe, analyse, and explain 'reality' or 'what there is'.

Data modelling frameworks and most ontologies are semi-formal or descriptive. It is for this reason that terms used to describe reality are often not introduced formally but instead are explained using concepts that give terms meaning. 'Entity', a term used in ER, is described in such a way in original sources and texts. In order to utilise terms and concepts in our methods we need to explore the nature of terms and concepts more deeply. Nevertheless, a term is only a label that can be used in grammars in juxtaposition with other terms. The nature of terms is given by the concepts that each embodies. For example, 'entity' embodies a concept that tells me, broadly speaking, what it is that makes my cat an entity rather than a relationship.

Definition 3·1—Concept

> "[A] concept is a way of thinking about something—a particular object, or property, or relation, or some other entity." (Dancy and Sosa, 1992)

Essentially, we have a concept associated with each term that gives meaning to the term. For example, there is a concept related to the term 'entity' in ER that gives it meaning. It is through concepts that a rational analysis of data modelling frameworks using an ontology can be made. For example the term 'entity' in ER relates to terms such as 'individual' from an ontology, or 'object' from an object model. It is only through examining concepts associated with these terms that a sensible analysis can be made of the similarities and differences between the data models. However, we need a more precise understanding of the nature of concept so that we can express them and use them in our methods.

According to standard philosophy, a concept is expressed using what is called a 'possession condition' (Dancy and Sosa, 1992). A possession condition essentially embodies what is required so that a sufficient understanding of the concept can be conveyed to another person.

Definition 3·2—Possession condition

> "[a] statement which individuates a concept by saying what is required for a thinker to possess it". (Dancy and Sosa, 1992)

We use terms and concepts to analyse, compare, and evaluate the data modelling frameworks. Before any of these activities can take place we need to select the terms and their associated concepts, on which the analysis is to be made. Note however, that this selection has to be made for both the ontology and the data modelling frameworks. It is typically more straightforward to extract terms from an ontology than for data modelling frameworks, because for data modelling frameworks we need to extract terms and concepts from a number of available sources. We give an example of the extraction of terms from a data modelling framework later in this section.

There may be several possession conditions for the same concept. For example, a triangle can be thought of as a trilateral and consisting of three sides, or as an object consisting of three angles. Both of these describe the same concept. For our purposes, however, the concepts in which we are interested are described in recognised and identifiable sources, that is, other researchers can inspect those sources and may form concepts that are similar in perspective to ours. There is, however, no guarantee as to the degree of similarity between ours and theirs, and, we hope, that the concepts that give meaning to the terms will be only marginally different from one researcher to another for the very reason that there are recognised sources from which concept formation can flow. As a consequence the extraction of concepts is one of the areas where differences are sure to emerge between different researchers. However, we hold to the fact that there is no easy or automatic way by which this analysis can be achieved. The very nature of the subjects in this work mean that differences in interpretation are likely.

Before progressing to the methods we turn to semiotic theory to clarify terminology so that there is no confusion between our discussion of terms and of their associated concepts. We also use the theory to expand our understanding of terms and concepts so that we can comment on the relationship between terms in the ontology and the data modelling frameworks so that we can present and interpret our results in ways that are relevant to our study.

Understanding Concepts Through Semiotics

We use parts of Umberto Eco's semiotic theory of codes (Eco, 1976) to understand the relationship between terms (names of concepts) and their associated content or meaning (the concepts themselves). Later, we use the theory to help present, analyse, and discuss our results. This theory is also useful in that it helps to clarify terminology, and to set a typographical standard.

In a semiotic sense, terms signify (Eco, 1976) concepts. They are labels that attach to concepts. We utilise a typographical convention identical to that used by Eco to distinguish between the discussion of a term and the discussion of the the meaning of a term (concept) (Eco, 1976). For example, in cases where we refer to the name of the term 'entity', we use /entity/. Alternatively, when we refer to the meaning behind the term, we use «entity». In our discussion, the content of a term is its concept, that is, the way of thinking about the term. The name denotes the term's meaning. For example, the sign /entity/ denotes the meaning «entity» about our term 'entity'.

In semiotic theory signifiers may be attached to several concepts. For example, /orange/ simultaneously signifies a fruit as well as a particular hue. Further, signifiers may change their name or meaning as language changes, or may signify additional concepts. In our study, and as we discussed in the previous sub-section, the meaning of terms is restricted by the documents and context in which the terms are presented.

More formally, we adopt the reading that the terms and concepts that we have described fit Eco's theory of codes (Eco, 1976), in that they are examples of sign-functions. That is the term /entity/ has content or meaning «entity». Further, that there is a sign-function that describes the relationship between the two.

Definition 3·3—Sign function

> "A sign-function is realized when two functives (expression and content) enter into a mutual correlation" (Eco, 1976).

This then means that we are interested in relating terms from an ontology or a data modelling framework with their associated concepts. We do so by firstly, recognising the relationship between term and concept as a sign-function, and then realising that a term is the expression in a sign-function, while the content of a sign-function is its associated concept that embodies the meaning behind a term. For example, /entity/ is the expression of the sign-function while «entity» is the content or meaning of the sign-function.

Any semiotic system must define the role of an interpretant in that system. An interpretant defines the context in which the sign is placed, and consequently it clarifies the sign itself. The interpretant in turn can be seen as a sign in the semiotic system which requires an interpretant. The interpretant is essential for the consistency of a semiotic system in that it is the mechanism by which the system can check itself entirely by its own means. The consequence of this approach to interpretant is that there is infinite semiosis—we cannot ever be *certain* about meaning. In our study, interpretant is acknowledged in the analysis of our results by acknowledging the subjective nature of that analysis.

Some terms in our ontology and data modelling frameworks are described using compound concepts that can be subdivided into several parts. Each part defines further aspects to the meaning of the whole term. It is the concept that is compound. Using an example removed from information systems, the term /knife/ has an associated concept. That concept is compound. «knife» incorporates two parts, «blade» that also has the label /blade/ and «handle» that has the label /handle/. The core concept for /knife/ is «blade» for without such a part, the meaning of concept to which /knife/ refers is meaningless. A second example comes from information systems. /entity/ may consist of a number of possession conditions or descriptions that together describe «entity», but may individually have significance. For our purposes we could say that

/entity/ has two parts, a core concept (and associated meaning), and that concerning the identity of entities. That is, /entity core/ signifies «entity core», and is that concept that is at the heart of the /entity/ that if absent would render /entity/ meaningless. /Entity identity/ signifies «entity identity» and is the concept describing the need for each entity to be identified by a key. The whole is still referred to as /entity/ and means «entity». Analysis of a term can be undertaken at the part of the whole level of a term.

Definition 3·4—Core concept

> A core concept is one such that its absence will render the definition of the entire concept meaningless.

We need to use Eco's theory of codes further in that the idea of a semantic field helps us to explain how concepts are related to each other. We also need it to explain and analyse our results. In Eco's semiotic theory, each word in a language, through its meaning, spans a semantic field (Eco, 1976). Semantic field is not a tightly defined concept. However, when comparing languages one can compare the 'semantic field' spanned by related words from each language. In Latin «mus» spans the same semantic field described by «mouse» and «rat» combined. In a similar way to these words, each term from an ontology or data modelling framework spans part of a semantic field (Eco, 1976) through its associated concept. Therefore, through the idea of a semantic field we can express the similarities and differences between concepts in the ontology and those in the data modelling frameworks and thereby discuss and explain our findings. We discuss the issues concerning semantic field in the appropriate part of the comparative method that we introduce in the next section.

Since the terms and concepts that we are interested in are stable and are described using recognised sources that are identifiable and can be revisited for clarification, we need not explore Eco's semiotic theory beyond the introductory parts of his theory of codes. The remainder of Eco's semiotic theory concerns the role semiotics plays in social situations where words and meaning, as well as language more generally, are in a constant state of flux. This contrasts with our study where the data modelling frameworks and the ontology, have concepts and terms that are described in a limited number of sources.

In conclusion, there are three important points that come out of this discussion. Firstly, that concepts can be named and are the terms that we discussed earlier, concepts can be compound. Secondly, that we can find a sign-function that describes the relationship between the name of the concept (the term) and its meaning (or concept). Finally, that semantic field is a useful term to use in discussing and comparing concepts from an ontology and data modelling frameworks.

An Example Extracting Terms and Concepts

To help clarify the process involved in extracting terms and concepts from the ontology and the data modelling frameworks, we present an example extracting a term from ER, one of the most widely used data modelling frameworks.

The process commences by inspecting the original sources that are widely available. In the case of ER, the original paper by Peter Chen (Chen, 1976) is the original source. The original source is then used to extract the terms that are central to modelling. For some articles these terms are very easy to extract because there are explicit references to the importance of the key terms. For other sources, terms are mentioned in the description of the modelling framework and must be teased out in a process of review, reflection, and inspection of other sources other than just the earliest seminal articles. Whether the key terms are patently obvious or are extracted more reflectively, quotations supporting descriptions of terms are extracted from important sources where the quotations are available.

Once the terms have been extracted, possession conditions (or what we refer to as descriptions) for the concepts to which the terms refer are constructed. As tersely as possible, these convey to the reader the information necessary to understand the fundamentals of the concept. Where concepts are compound, involving several ideas, a core is identified.

Extracting part of our description that appears later in the thesis—ER models are centred around three fundamental concepts: entity-type, relationship-type, and attribute-type (Peckham and Maryanski, 1988). Members of entity-type are identifiable entities (either physical or conceptual) from the world. "An *entity* is a 'thing' which can be distinctly identified." (Chen, 1976) These are then classified into entity sets (Chen, 1976).

The term 'entity' refers to a concept that has two parts. Firstly, a core that describes the essential concept to which 'entity' refers. Secondly, a part that refers to the identity of entities in each class. Further, the description is a terse possession condition that enables others to understand the concept involved. A summary of the term and its concepts is shown in Table 1. Readers may notice that we have dispensed with the word 'Term' in the header of the table, and substituted 'Concept' to avoid verbosity in referring to terms and concepts. The consequence of this is that we can now refer to /entity/ where the name of the concept is needed (i.e. the 'term') and «entity» where the meaning of the term is needed.

Table 1—Extract from the Table of Concepts for the Entity-Relationship Modelling Framework

Concept	Part	Description
Entity	Core	ER allows for significant entities, or objects (either physical or conceptual) to be modelled. These must be grouped into entity classes. Each entity cannot depend upon other entities to be classed as an entity.
	Identity	Each member of an entity class must have a unique identity called a key.
...	...	...

This extract is for the Entity-Relationship modelling framework but illustrates the style that applies equally to the ontology or to one of the other data modelling frameworks. Interested readers will find these in this thesis where the candidate ontology or the data modelling frameworks are respectively introduced and described in Chapter 4 and Chapter 5.

3·2 Comparative Method

We envisage two ways in which concepts from the ontology can be used with data modelling frameworks. Firstly, the concepts in an ontology can be used in pairwise comparisons with data modelling frameworks. Secondly, concepts from the ontology can be used to evaluate specific data modelling frameworks. We explore the comparative method in this section before considering the evaluative method in the following section. In these methods we go beyond structural comparison to the more fundamental level where the semantics of terms and concepts from the ontology and data modelling frameworks are compared. In comparing the semantics or concepts we expect to compare the world views of each.

The methods are qualitative, that is, the information being compared is rich and descriptive. In methods of this nature, a high degree of trust is given to the researcher. This is because it relies on the experience of the researcher, and on his or her ability to examine deeply the issues that arise. Consequently, replication of the results will be more difficult than when using less qualitative approaches. We will revisit these issues upon conclusion of the thesis.

In the comparative method we use the ontology as a benchmark against which each framework can be measured, and as a reflection of reality. The chosen ontology presents us with a view of reality. We begin by selecting concepts from the ontology which are relevant for the comparison, and they form the basis for comparison. We then perform a pairwise comparison of each data modelling framework with the ontology based upon concepts from the ontology. For example, if we are examining four data modelling frameworks then we would conduct four pairwise comparisons with the ontology. In conducting the pairwise comparisons we are testing each framework against the selected and independent view of reality.

Each pairwise comparison of an ontology with a data modelling framework results in a series of graded indicators of agreement of the data modelling framework with the ontology. We are seeking a qualitative indication of agreement/difference between the ontology and each data modelling framework for each concept from the ontology. This is because a framework may support a concept but not in the full generality provided by the ontology. We use Eco's theory of codes and the idea of coverage of semantic field to more fully explain the nature of the graded indicator. We will discuss the issue of coverage of semantic fields a little later in this section. We introduce this aspect of analysis here rather than in the previous section, because it allows us to introduce concepts relevant to analysis close to the place where we first discuss results.

The pairwise comparison allows us to find a qualitative and relative 'goodness of fit' of the framework with the ontology, and consequently the 'goodness of fit' with the benchmark.

The comparison is motivated by a desire to understand the degree to which each data modelling framework reflects the concepts embodied by the selected ontology. Recall that this will be indicative of the capacity of an ontology to be a unifying framework for data modelling frameworks. Alternatively, one could say that we are examining the ontology of each modelling framework. Depending upon our results we may be able to judge if the selected ontology is likely to be a unifying framework. Once all pairwise comparisons are complete, they can be combined to form a grand comparison that is indicative of the likelihood of the ontology forming a unifying framework for data modelling frameworks. That is, if for many of the concepts from the ontology, there is a high level of agreement with all of the data modelling frameworks, then the ontology may be a suitable choice as to forming the basis for a unifying framework. Alternatively, if substantial and central concepts from the ontology are not supported by a number of data modelling frameworks, then the ontology may not be a suitable choice as a unifying framework.

Interested readers may see a preliminary presentation of results in (Milton et al., 1998). Previous comparisons of data models are not ontological by nature, for example, the surveys in (Hull and King, 1987, Peckham and Maryanski, 1988).

In each pairwise comparison, the researcher conducts a 'mind experiment' in which a comparison of each data modelling framework with an ontology occurs. We need to be able to convey succinctly what we have found, and consequently, we seek an indicator that shows the degree of the overlap between the semantic fields covered by concepts in the ontology and the semantic fields covered by concepts in the data modelling frameworks. It is critical to be aware that the indicator shows the nature of the overlap as the researcher sees it, and that it must be accompanied by an explanation of the results so that the nature of the coverage is explored and explained. This is why, when we present our results we follow it by a section in which the results are discussed. In the discussion some supporting arguments will be given.

Suppose we have a concept /c/ (from the ontology) and a specific data modelling framework. There may be three broad categories of results.

Firstly, the data modelling framework may have total overlap with respect to /c/. Total overlap may be provided by one concept (for example, /d/) or perhaps by several concepts (for example, /d/ and /e/). That is, there may be one concept or several concepts that together provide total overlap, in terms of semantic field, with the concept from the ontology.

The second possibility is where the overlap is partial. Further, it may be that the concept from the modelling framework can be extended to support the full generality of the concept from the ontology. It may also be the case that there is little prospect for extension due to certain concepts in the data modelling frameworks being contradictory with respect to the ontology in ways that make extension difficult.

Finally, it may be that there is no overlap at all between the data modelling framework and /c/ from the ontology.

Figure 4 shows the three categories pictorially, in the style of Eco's theory of codes. Firstly, Figure 4a shows full overlap where /c/ from the ontology is covered by one (/d/) or several (/d/ and /e/) concepts from the data modelling framework. This could be extrapolated to discuss more than two concepts from the modelling framework. Secondly, Figure 4b indicates partial overlap. The areas in grey show segments of the semantic field for the concept in the ontology that are not covered by concepts from the data modelling framework. Thirdly, Figure 4c shows no overlap whatsoever. In this last case, there is no concept in the data modelling framework that concerns the semantic field described in the concept from the ontology.

Because the comparison is also descriptive and is based on concepts (or meaning of terms) and not structure or form, one can discuss the ease with which a total overlap in semantic fields can be achieved, that is, achieved from the starting point of either no overlap or only partial overlap in coverage of semantic field (as shown in Figure 4b and Figure 4c respectively). While the coverage of a specific concept is depicted in this figure as a sharp rectangle, the edges are quite blurry. This betrays the very nature of semantic fields.

Figure 4—Degree of Overlap in Coverage of Semantic Field

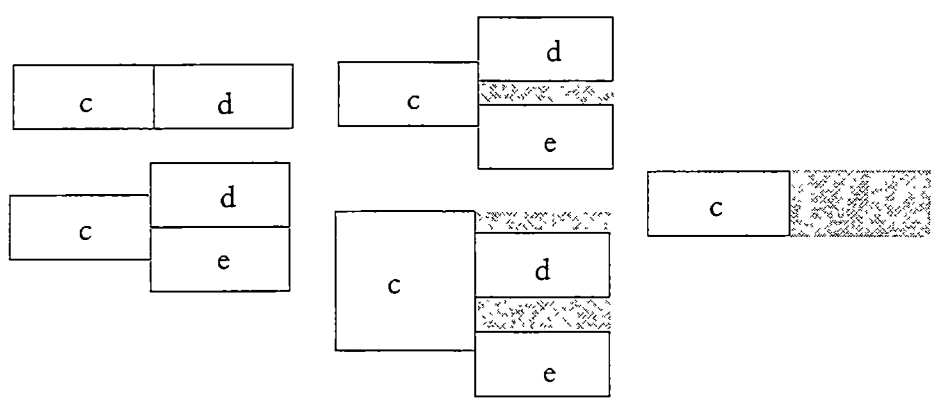

a—full overlap and coverage **b—partial overlap and coverage** **c—no overlap and coverage**

In the following sub-section we show how to analyse the results in order to carry out the comparison of the selected data modelling frameworks using the selected ontology.

Presentation, Analysis, and Discussion of Results

Results for the method are presented in a tabular form. Table 2 shows the basic idea. Concepts from the ontology are listed down the left of the table with the relevant data modelling frameworks constituting the columns with entries for each data modelling framework *vs* each concept from the ontology.

Table 2—Sample Results for a Comparison of Data Modelling Frameworks with an Ontology

Ontological Concept	Framework 1	Framework 2	Framework 3	Framework 4
A	√	√	√	√
B	$√_p$	X	√	$√_p$
C	$√_p$	X	X	$√_p$

Immediate feedback is provided in this table showing which concepts from the ontology are fully covered, partly covered or not covered at all by the respective data modelling frameworks. Each pairwise comparison is a column in the table. Each concept from the ontology is a row in the table. A 3-level indicator shows the level of coverage in the respective data modelling frameworks for the relevant concept from the ontology. These correspond directly to the degrees of overlap and coverage of concepts from the ontology as shown in Figure 4. Clearly, in the cases of full coverage or no coverage, the semantic field covered by the concept in the ontology is covered or not covered respectively with various concepts in the relevant data modelling framework. Full and no coverage is indicated using √ and X respectively, and are depicted in Figure 4a and Figure 4c respectively. Where partial coverage is evident, it means that concept(s) from the data modelling frameworks partly cover the semantic field described by the specific concept from the ontology and is indicated using $√_p$. This is depicted in Figure 4b. Discussion reveals the extent and nature of partial coverage for each concept.

Some of the concepts shown in Table 2 may be compound and there is a relationship between coverage in the parts of each concept and its whole. For example, suppose ontological concept A consists of two parts A_1 and A_2, then the coverage for the concepts as a whole are related in the way shown in Table 3. In cases of compound concepts discussion also reveals the extent and nature of the coverage.

Table 3—Consolidation of Coverage for Compound Concepts

A	√	$√_p$	$√_p$	$√_p$	$√_p$	X
A_1	√	√	√	$√_p$	$√_p$	X
A_2	√	$√_p$	X	$√_p$	X	X

Gaps in the coverage of the semantic space described by a concept from the ontology occur where either no coverage is evident (Figure 4 c) or where partial coverage of the semantic space described by the concept is evident (Figure 4b). The gap in these cases can mean one of two things.

Firstly, it could mean that there is no concept (or part of a concept) from the data modelling framework that has coverage in any way over the gap in the semantic field concerned. In this case, extension of the data modelling framework is likely to be relatively straightforward, although deeper analysis is required to determine the ease with which extension can occur.

Alternatively, it could be that there is a concept from the modelling framework that doesn't span the 'gap' in the semantic field, but instead spans a different semantic field that is contradictory with respect to the original semantic field from the ontology.

Definition 3·5—Contradictory Concepts

> Concepts are contradictory if they form 'pairs of oppositions constituting a semantic axis' (Eco, 1976)

We examine these relationships in more detail when we examine the evaluative method in the next section. Briefly, where a 'gap' in semantic coverage is found there may be a contradictory relationship between the concept from the ontology and some concepts from the modelling framework. That is, the meaning of concepts from the data modelling framework runs counter (in some way) to the concept from the ontology. This is an example of antimony.

Antimony is divided into three categories that fill out the 'semantic axis' (Eco, 1976); That which is 'complementary', 'contradictory', or 'contrary' (Eco, 1976). If the nature of the antimony is such that there are only terms that are contrary or contradictory to that from the ontology then extension of the modelling framework to fill the 'gap' in semantic coverage will be difficult because there are already terms that clash with that from the ontology. If, on the other hand, the terms are complementary then extension will take the form of simply adding terms to the data modelling framework that are synonymous with the 'gap' in the semantic field of the ontology.

Using these ideas we can discuss each set of results to analyse them more closely. In this thesis, where we report our findings of applying this method with five data modelling frameworks, we commence by presenting the indicative results, which are then discussed and analysed in the section immediately following. In the discussion of the indicative results, the meaning for each concept in the ontology can be used as a basis for explaining the relationship between concepts from the ontology and those from the data modelling frameworks. The discussion also enables us to reason about gaps in the overlap of semantic fields between the ontology and the data modelling framework. This in turn would allow us to comment on the likely ease with which one can extend a specific data modelling framework in order to improve the coverage of a specific semantic field.

This method compares data modelling frameworks against the theoretical benchmark of the ontology. This differs from existing comparisons of data modelling frameworks where categories are established for the frameworks or where application genres are created.

The ontology plays the role of an objective theory or of a benchmark in our comparison. Our experiences in conducting the comparison will enable us to comment on the prospect of using the ontology in the role of a unifying framework.

We have applied this method using the entity-relationship model, the semantic data model, the functional data model, NIAM, and OMT with our selected ontology, the results of which are presented and analysed in Chapter 6.

3·3 Evaluative Method

In addition to measuring the degree to which each data modelling framework fits the ontology, we may also wish to capture the world view of the ontology in our modelling framework. In order to achieve this we need to evaluate a data modelling framework. Essentially, we can perform an ontological evaluation of a specific data modelling framework in order to determine in what specific ways a modelling framework fails to fully reflect the world view of an ontology. Additionally, using the results and semiotic theory that we have introduced, we can comment on the likely ease with which a modelling framework may be extended to more fully reflect ontological concepts.

For us to extend a specific data modelling framework using an ontology, we need first to understand the differences and similarities between concepts in the data modelling framework and those in the ontology in a much deeper way than is required for the previous method. In the previous method we are interested in comparing data modelling frameworks through a series of individual comparisons. In this method, we need to understand much more deeply, the differences in concepts supported between an ontology and one data modelling framework. It is for this reason that we conduct a bidirectional comparison between an ontology and data modelling framework, as is illustrated in Figure 5.

Figure 5—Evaluating Data Modelling Frameworks

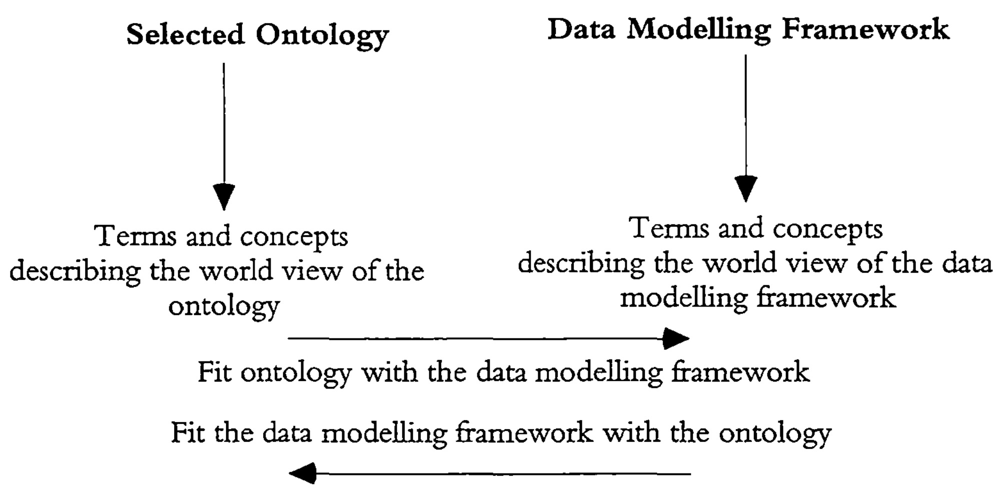

The bidirectional nature of the comparison between the concepts found in the ontology with those of the data modelling framework is shown by the two arrows in Figure 5. Essentially, we not only compare the data modelling framework with concepts from the ontology in similar ways to one of the 'individual comparisons' in the last method, but also compare the ontology with concepts from the data modelling framework. Changes are then suggested from the analysis of results of the comparison that intend to bring the ontology and data modelling framework closer together in terms of coverage in concepts. The motivation for a bidirectional comparison comes from our desire to bring a data modelling framework closer to the ontology in order to enable the modelling frameworks to be used for new application areas or for new purposes that are covered by the ontology. The evaluative method incorporating the bidirectional comparison is therefore useful in that it provides an assessment as to the likely ease with which any modifications to the data modelling framework can be made.

The method has two parts. Firstly, the (bidirectional) comparison. Secondly, the analysis and discussion of results. We deal with each of these separately below. We discuss the information that is yielded by the comparison near the conclusion of this section.

Bidirectional Comparison

In the previous method we compared a number of data modelling frameworks by conducting a series of individual comparisons in which each data modelling framework is compared with concepts from the ontology. This method has similarities with that method in that we use terms and concepts in comparisons, however, it differs from our previous requirements in that we do not wish to compare several data modelling frameworks with an ontology. Instead, we seek to determine in what specific ways a particular data modelling framework reflects the concepts from an ontology thus the method is evaluative by nature. Because the method is evaluative, we also need to be precise about the nature of the similarities and differences so that these can be corrected. This means that our approach to the method for this purpose is different from the previous application in three ways.

Firstly, we must perform two individual comparisons. The first compares concepts from the ontology with the data modelling framework similar to the previous method. We call this the 'forward' comparison. The second application compares concepts from the data modelling framework with the ontology and we call this the 'reverse' comparison. This bidirectional comparison enables us to determine in what specific ways the ontology and data modelling framework differ. Once we can describe and understand the differences, we can suggest where to extend the data modelling framework in order to bring its world view closer to that of the ontology. The comparison needs to be bidirectional because some of the concepts in the data modelling framework may cover concepts in the ontology where the ontology is not prescriptive or equivocates in some way. A deeper insight may also be gained into the modelling framework by viewing the ontology from the viewpoint of the data modelling framework. It also allows for judgement to be exercised as to whether a data modelling framework reflects a concept from the ontology.

Secondly, we must extract concepts from the ontology that relate to dynamic and static modelling. Some modern data modelling frameworks incorporate dynamic and static concepts, whereas the more traditional frameworks do not. Concepts from the ontology that are omitted from consideration are those that deal with purely philosophical questions. For example, those dealing with the existence of God, or the discussion of precisely where physical objects begin and end, need not be considered in this method.

Thirdly, the method involves a bidirectional comparison, and consequently results from two comparisons require analysis. It is for this reason that analysis involves an extra step whereby the individual comparisons are consolidated into a summary. We also seek to explain differences between an ontology and a data modelling framework in a deeper way than before. Consequently, in our discussion, which is very important in this method, we utilise our semiotic theory in order to explain specific differences between concepts in the ontology and the data modelling framework.

In the remainder of this section we consider the analysis required for this method.

Presentation and Discussion of Results from Each Comparison

Results for each part of the bidirectional comparison of a specific data modelling framework with an ontology are initially presented in a tabular form in a similar way to the previous method. With the comparison being bidirectional, two tables are presented, summarising a comparison of the data modelling framework with the concepts from the ontology, and a comparison of the ontology with the concepts from a data modelling framework. Figure 6 shows the basic idea and contains two tables representing the two comparisons. The left-hand and first comparison is of a data modelling framework with concepts from an ontology. The right-hand table shows the second comparison of an ontology with concepts from a data modelling framework. The concepts in each case are listed down the left of the respective tables with the cells of the respective tables showing the degree of coverage of the semantic field for the relevant concept. The left-hand and right-hand tables indicates the extent to which the semantic field of the concept is supported by concepts in the framework and the ontology respectively.

Figure 6—Sample Results for a Bidirectional Comparison of an Ontology and a Data Modelling Framework

Ontological Concept	Selected Framework
A	√
B	$\surd_P$
C	$\surd_P$

Framework Concept	Selected Ontology
I	X
II	$\surd_P$
III	$\surd_P$

Immediate feedback is provided in these tables showing which concepts are fully covered, partly covered or not covered at all by either the ontology or the data modelling framework (according to which is the table being inspected). Exactly as in the earlier method a 3-level indicator is used. The indicator corresponds directly to the categories shown in Figure 4 on page 39. We consider the nature of the two tables in turn.

The left-hand table in Figure 6, uses the 3-level indicator showing the level of coverage in the data modelling framework for the relevant concept from the ontology, that is, coverage, found in the data modelling framework, of the semantic field described by the respective concepts from the ontology. The 'reverse' comparison shown in the right-hand table in Figure 6 is where an ontology is compared with concepts from a data modelling framework yielding results using the same set of indicators as the 'forward' comparison with the necessary corrections made reflecting the fact that the roles are reversed with respect to the ways we have explained the nature of the results. Discussion of the indicative results follows the indicative results in the usual way. For a full description of the indicators, please refer to the comparative method described in the previous section.

In this method the discussions for each comparison, revealing and justifying the nature of similarities and differences, is more critical than in the comparative method, because we are evaluating a specific data modelling framework using an ontology to determine in what specific ways the data modelling framework is lacking in coverage of the semantic fields described by concepts from the ontology. We may also want to use these similarities and differences as a basis for suggesting ways to extend or in other ways amend a data modelling framework to cover areas where coverage of concepts from the ontology is partial or nonexistent. For this reason, we need to not only describe how the two differ, but also be specific in terms of the nature of any differences and elaborate upon the differences and to highlight areas where the coverage is lacking.

To facilitate the analysis of the two comparisons we use the discussions from each comparison in the consolidation and analysis of results.

Consolidation and Analysis of Results

Following the presentation and discussion of each part of the results for the bidirectional comparison of an ontology and a data modelling framework, a consolidation of the indicators of the coverage for each concept in the ontology can be made, which shows the degree to which concepts from the ontology are covered by the data modelling framework.

In the consolidated view, as illustrated by Table 4, the level of coverage is categorised to be at one of three levels. Firstly, 'full', where the full meaning of the concept is preserved in the data modelling framework. Secondly, 'partial', where the meaning of the concept from the ontology is only partly preserved by a concept (or concepts) in the data modelling framework. Further, extension or amendment of the data modelling framework may be required to 'fill the gap'. Thirdly, 'none', where there is no concept or concepts that provides coverage of the semantic field as specified for the specific concept from the ontology. Similarly, extension or amendment may be required to deliver the coverage that is similar to the data modelling framework.

Table 4—Consolidated Results Showing the Level of Coverage in a Data Modelling Framework of the Semantic Field for Concepts from an Ontology

Ontological Concept	Level of Coverage
A	full
B	none
C	partial

The 'forward' comparison reveals the degree and quality with which the data modelling framework covers concepts from the ontology. The 'reverse' comparison completes the picture and considers concepts in the data modelling framework and whether they cover parts of concepts from the ontology through unusual mechanisms or in subtle ways. A complete view as shown in the consolidated table is only possible after the 'reverse' comparison because in the 'reverse' comparison we take into account support that is not evident in the 'forward' comparison. We also examine the nature of the overlap in more depth.

Extension or amendment needs to be undertaken in the data modelling framework, where the consolidated results show that there is no coverage or only partial coverage evident in the data modelling framework for concepts certain from the ontology. The ease with which changes can be made depends upon the nature of any 'gap' in coverage. There are two possible interpretations.

Firstly, it could mean that there is no concept (or part of a concept) from the data modelling framework that has coverage over the 'gap' concerned in any way, and in this case, extension of the data modelling framework is likely to be relatively straightforward since the data modelling framework is likely to be orthogonal with respect to the ontology. A more thorough examination of the overlap is required to determine the ease with which extension can occur.

Secondly, it could be that there is a concept from the modelling framework that doesn't span the 'gap' in the semantic field, but instead spans a different semantic field that is contradictory with respect to the concept from the ontology. The relationship between the two is an example of one of the 'pairs of oppositions constituting a semantic axis' (Eco, 1976), that is, its meaning runs counter (in some way) to the concept from the ontology. This is called antimony. Using this idea of antimony we can comment upon the quality of the partial coverage as we discuss our results. It has clear implications for an extension of a specific data modelling framework in order that it covers the gap in the semantic field, and this is discussed below.

Antimony is divided into that which is complementary, contradictory, or contrary (Eco, 1976), and if the nature of the antimony is such that there are only terms that are contrary or contradictory to that from the ontology then extension will be difficult. If, on the other hand, the terms are complementary then extension will take the form of simply adding terms to the data modelling framework that are synonymous with the 'gap' in the semantic field of the ontology. In the context of this discussion, complementary antimony can also be described as orthogonality. These discussions will be important when analysing the consolidated results with a mind to strengthening the data modelling framework.

The method concludes with a discussion that interprets the consolidated results and comments on the nature and degree of support in the data modelling framework of concepts from the ontology and can be used to suggest amendments to a data modelling framework. In reacting to the results for an evaluation, there may be several ways in which a data modelling framework could be amended to better reflect the concepts from an ontology. When we apply the evaluative method we will discuss possible avenues for extension and amendment of data modelling framework under study. In line with our argument, however, the object of any extension or amendment is to support concepts from the ontology in a data modelling framework and thereby improve the fidelity with which the data modelling framework can represent 'what there is'.

The purpose of this evaluative method is not to perform an extension. Instead, it is there to highlights areas where the world view of the data modelling framework is different from the world view of the ontology. It also predicts the degree of difficulty likely to be encountered when making any necessary extension, and in supporting the concepts of interest within a given data modelling framework based upon an analysis of the nature of antimonies that may exist.

The results of applying this method with a selected ontology and some data modelling frameworks are presented in Chapter 7. In that chapter, we present our findings and recommendations after analysing the functional data model using this method. In Chapter 8 we apply this method using OMT's object model and the dynamic models for OMT. In each of these cases we use the method to suggest differences in the world view of data modelling frameworks with respect to the selected ontology.

CHAPTER 4

Selecting an Ontology

AN ONTOLOGY THEORISES about the categories of things that make up reality. There are many ontologies from which to choose, and no ontology is necessarily better than another. It is for this reason that we select an ontology based upon a number of criteria that fit within our research philosophy. In this chapter we discuss our criteria for pragmatically selecting an ontology, and we select one from the literature based upon these pragmatic criteria, before describing it.

Recall from earlier chapters that we need to select an ontology that has the potential to yield a set of unifying principles from which a unifying framework may eventually emerge. The potential for an ontology relies on it fulfilling certain pragmatic criteria. Our discussion is divided into six sections. Firstly, we examine the criteria by which we select an ontology. We then explore the philosophical heritage of the selected ontology, to make sure that the philosophical school from which the selected ontology is taken is not at odds with our purposes and is not going to present any difficult ontological or epistemological problems. We conclude the chapter by describing the selected ontology in three sections in which we introduce the major ontological categories, and we describe the static and dynamic parts of the ontology.

4·1 Selecting an Ontology by Pragmatism

Of the many ontologies that have been proposed, we are only interested in those that can be used for our purposes. You will recall from Chapter 1, that we ultimately intend to use an ontology as a unifying framework for data modelling frameworks. A practical consequence of our objectives is that we expect to be able to use ontology to compare data modelling frameworks and to evaluate data modelling frameworks.

We are considering ontology a unifying framework for data modelling frameworks. We believe that in order for a theory to have potential as a unifying framework for data modelling frameworks it must have three features or capacities. Firstly, the theory must have a coherent set of concepts. Secondly, each data modelling framework must be able to be expressed within the unifying framework. Thirdly, the unifying framework must be able to be used to analyse any data modelling framework. In this context 'analysis' means ability to compare, evaluate, extend, and understand the view that each data modelling framework has with respect to the world. The second and third traits give rise directly to a number of pragmatic criteria. In this context, pragmatic refers to our desire to use the ontology as a theory for understanding, evaluating, and extending data modelling frameworks.

Firstly, if the ontology is to eventually become the basis for a unifying framework, then it must discuss elements of a similar nature to data modelling frameworks. This is important because data modelling frameworks contain an array of elements that are used to construct models of reality.

From our knowledge of data modelling frameworks, we need concepts in the ontology that help to determine and describe significant objects or things in the world. Further, a description of an object or thing requires the determination of properties or attributes of those things. In addition, classification is used extensively in data modelling frameworks. It thus seems reasonable to expect classification structures of some sort which enables assemblies of things or objects to be described that represent useful or conceptually logical groups. Data modelling frameworks often have mechanisms by which things or objects are related. We therefore require the selected ontology to be able represent how things or objects are linked or related.

Additionally, there are some data modelling frameworks in which we are interested that go beyond statics to include dynamic aspects. Consequently, we also require an ontology that deals with statics and dynamics. Temporally, states of affairs change, and so a description of reality changes. A description of reality in which we are interested is one that adheres to an ontology. This in turn requires that the ontology according to which the description is fashioned, must deal with change, in that the ontology should handle both static and dynamic aspects of reality. A large number of ontologies deal with either dynamics or with statics. Fewer ontologies deal with both.

It is important to note that in order to fully satisfy this criterion, we must examine the ontology's philosophical heritage to ensure that it is compatible with this criterion. It could be argued that some style of realism is implied by this criterion. Nevertheless, we seek an ontology that assumes a philosophical heritage that is compatible with this criterion. It is therefore important for us to examine the philosophical heritage of our selected ontology before finally committing ourselves to its use. We consult sources in philosophy to guide our discussions (Honderich, 1995, Kim and Sosa, 1995, Audi, 1995, Dancy and Sosa, 1992). The discussion of this is in the next section of this chapter.

Secondly, the ontology should be set at a level which generalises from the linguistic expression of 'what there is'. Many ontologies are conceived by people expressing, through language, their understanding of 'what there is'. We desire an ontology that has abstracted away from these expressions and generalises in appropriate ways. Appropriateness in this context means that the generalities used in the ontology help directly in evaluating and discussing data modelling frameworks. In some ways this is a restatement of the first criteria but from the viewpoint of ontology. Before commencing the study, this criterion can only be used in a somewhat superficial manner, since it is only through using the ontology in the way suggested can we be sure that this criterion is fully satisfied. Nevertheless, given the range of ontologies from which to choose, this is one aspect to selection we feel is important to mention explicitly.

We have examined many ontologies in our search for an ontology to use in this thesis (Quine, 1960, Quine, 1969, Russell, 1908, Apel, 1980, Bunge, 1977, Bunge, 1979, Bühler, 1934, Carnap, 1963, Chisholm, 1996, Habermas, 1992, Kaminsky, 1982, Kattsoff, 1956, Mulligan, 1992, Milton and Keen, 1996). Most were discarded due to them failing to fulfil some or all of our pragmatic requirements.

We have found that the ontology recently published by Roderick Chisholm (Chisholm, 1996, Chisholm, 1992) qualifies according to our pragmatic criteria. We describe Chisholm's ontology in the last three sections of this chapter, but we briefly indicate how Chisholm's ontology meets our criteria below.

Firstly, we can say that Chisholm's ontology provides for individuals, attributes/properties, relations between individuals, classes and sets of individuals, and events/states. We have argued that this is important because data modelling frameworks have terms that relate to many, or all, of these. Consequently, it is likely that the modelling frameworks can be related to the ontology.

Secondly, Chisholm's ontology is set at a level of abstraction from language that is appropriate for our purposes. It is the case that many ontologies do not progress beyond the discussion of language. Chisholm's ontology is a generalised theory that is constructed to withstand social and cultural changes.

Thirdly, Chisholm's ontology not only models statics but also models dynamics. It supports changes in the exemplification of attributes together with a rich and powerful state and event structure.

Before commiting to the ontology we must check that the ontology is compatible with our research philosophy. This can only be verified by examining the philosophical heritage of the ontology.

4·2 The Philosophical Heritage of Chisholm's Ontology

We have selected an ontology written by Roderick Chisholm (Chisholm, 1996). Chisholm is well known for his contribution to philosophy in the late twentieth century.

In this section we discuss several philosophical aspects of our selection. We introduce the author of the ontology in order to ascertain their standing in the community of philosophers. We then utilise recognised sources to discuss the philosophical heritage that Chisholm assumes in constructing his ontology. In the discussion we will observe the likely impact on our aims and objectives should we select Chisholm's ontology.

> "**Chisholm, Roderick M**: American epistemologist and metaphysician who has been a seminal figure in contemporary philosophy, Chisholm has helped renew interest in metaphysics during the last third of the twentieth

century. Raising the art of philosophical analysis to new heights, his Socratic searches for analyses are legendary" (Kim and Sosa, 1995)

Clearly, Roderick Chisholm is an important contemporary philosopher. He remains active in his retirement and is Emeritus Professor at Brown University in The United States of America. We need, however, to examine his ontology with respect to its heritage before we can commit to selecting his ontology for our work.

We begin our discussion with a cautionary note. Western philosophy is complex and replete with discussions ranging from the existence of God to the workings of the universe. There are often discussions that border on the esoteric or that engage in what may be construed to be hair-splitting. We use several sources to help clarify terms, because terms that have popular meanings and that are seemingly innocent such as 'objective' and 'subjective' have multiple specific philosophical meanings and are so often misused in information systems and computer science.

We utilise a well-respected text (Flew, 1989) that introduces, both for philosophers and for the general reader, a history of the major arguments affecting western philosophy from the ancient Greeks to this day. We refer to more recent writing closer to the fields of information systems and computer science to help us to gain a different perspective on certain philosophical stances (Artz, 1997, Smith, 1995). Additionally, we use relevant reference material in which specific terms are defined and in which key people are described (Kim and Sosa, 1995, Audi, 1995, Honderich, 1995, Dancy and Sosa, 1992), and from which we get definitions for key epistemological and metaphysical concepts.

Each ontology that one considers assumes definitions, and uses terms, that are steeped in the western philosophical tradition. Moreover, the attitude taken to certain key questions by the author of an ontology reveals his or her philosophical outlook. This outlook is expressed in terms with deep philosophical meaning. For example, a philosopher may state in his or her ontology that they discount this or that aspect of *Plato's* writing. They may claim that their work is idealistic or, alternatively, realistic. While to novices this may be all very well, to the more aware observer such a claim betrays certain attitudes held by the author that affects diverse issues in specific ways.

Particularly, these attitudes affect the ontology in the ways in which sense-data about reality is interpreted or perceived by the observer. This last point is quite critical. It assumes a large amount of reasoning about a reality as it is perceived by an observer. Further, questions are automatically raised that discuss whether a reality exists separately from the observer, or whether it is in fact only in the mind of the observer. Suffice it to say that our interest is obviously not in these deep questions, but instead to determine whether the path of western philosophy found in our selected ontology is appropriate for our purposes. We need to do this without losing sight of the pragmatic requirements already outlined and without compromising our research goals. We also need to examine the philosophical stance taken in the selected ontology and be satisfied that it does not run counter to our research objectives.

In order to establish the philosophical stance that is taken by Chisholm, we examine the introductory essay in his monograph describing the ontology (Chisholm, 1996). There are two aspects to the philosophy in Chisholm's ontology. There is the brand of realism that he adopts, and there is the epistemological stance that determines, literally, the nature of the knowledge that he envisages will coexist with his categories.

In the opening comments in his monograph describing the ontology (Chisholm, 1996), in which he remarks upon the nature of the theory, Chisholm rejects the debate over much of Aristotle's metaphysics, but accepts some of the Platonistic theories:

> "And the theory that I defend rejects many of the metaphysical views that are generally associated with Aristotle: for example, the doctrine of form and matter, the distinction between substance and accident, and the 'moderate realism' according to which the only attributes that exist are those that are exemplified. Nevertheless, I will follow Aristotle and will make use of his insights throughout this book. The present theory is 'Platonistic': it is a form of *extreme realism*" (Chisholm, 1996)

This is a commentary on the nature of attributes and on substances. From another source referring to Chisholm, "... he combines two traditional forms of realism: Aristoteleanism about substances and Platonism about attributes" (Kim and Sosa, 1995). We can comment on its attitude to substance and attributes. The theory is Platonistic in the sense that attributes can be exemplified, may not be exemplified, or in some cases cannot be exemplified, but all are nonetheless attributes; This is counter to Aristotelean thinking. For example, being an elephant is exemplified. A pink elephant is not exemplified. A "round square" cannot be exemplified. In the ontology, attributes or properties themselves are abstract and endure. This is also contrary to Aristotlean belief. The theory is Aristotlean in a number of ways. The most critical for our concerns is that substances can be concrete (i.e. they come into being and pass away) or that are abstract (do not pass away and endure). A stone is an example of a concrete substance that can pass away from existence. An abstract substance could be God (subject to one's beliefs). We examine the implications of this philosophical attitude in our detailed description of the ontology. There is no evidence, based upon his stance on realism, that the ontology is not appropriate for our purposes.

Another aspect to his ontology is the epistemological aspect which outlines Chisholm's attitude to knowledge and knowledge acquisition. In a way, it also affects the style of realism that he adopts.

> "Our approach to philosophy is what Charles Sanders Peirce has called 'critical commonsensism.' This approach is based on faith in one's own rationality. Reason, as Peirce put it, not only corrects its premises, 'it also corrects its own conclusions' " (Chisholm, 1996)

Examining more deeply CS Peirce's critical commonsensism one can say that, in referring to Peirce:

> "In later writings his anti-Cartesianism took the form of 'critical commonsensism': our inquiries are guided by a slowly evolving body of vague commonsense certainties which are, in principle, fallible; rational self-control requires that we try to doubt these in order to establish that they genuinely form part of commonsense." (Dancy and Sosa, 1992)

A definition of commonsensism can also be found:

> "Commonsensism is the view that we know, most, if not all, of those things which ordinary people think they know and that any satisfactory epistemological theory must be adequate to the fact that we do know such things." (Dancy and Sosa, 1992)

Critical commonsensism differs from commonsensism (in the way in which commonsensism is described in the Companion to Epistemology) in that it demands a more rigorous standard of support for knowledge to be acquired. Hence the term 'critical'. Chisholm's ontology is an example of critical commonsensism. Nevertheless his ontology is categorised as being one of "extreme realism" (Chisholm, 1996). Does this then make the ontòlogy one of commonsense realism? We do find discussion of what is called commonsense realism. Further, we need to investigate this style of commonsensism in order to consider it with respect to critical commonsensism. Taking a step back, we can find a definition of realism:

> "Realism in any area of thought is the doctrine that certain entities allegedly associated with that area are indeed real. Common sense realism — sometimes called 'realism', without qualification — says that ordinary things like chairs and trees and people are real. Scientific realism says that theoretical points like electrons and fields of force and quarks are equally real. And psychological realism says mental states like pains and beliefs are real." (sic) (Dancy and Sosa, 1992)

However, Chisholm's critical commonsensism combined with his form of extreme realism means that his ontology adheres to a philosophy that may be a little different from commonsense realism. We need to examine the differences (if any) between critical commonsensism and commonsense realism. Barry Smith (Smith, 1995), also a prominent realist and who refers to Chisholm in his discussion, defines the school more clearly than the Companion to Epistemology, and in his article he outlines his support for the thesis that commonsense realism in its various guises seems to be useful in cognitive science. Indeed he argues for two radical theses that state:

> "1 *Uniqueness*: that what we shall call the common-sense world is, modulo certain trivial differences of emphasis and calibration, culturally invariant; and
>
> 2 *Autonomy*: that this common-sense world exists independently of human cognitive activities" (Smith, 1995)

This certainly is a radical pair of theses. This is an argument in favour of the uniqueness of the commonsense world (not the interpretations of that world). It also states that we can separate the cognitive processes of humans from the description of the world. It does not mean that the commonsense world as he sees it does not include human elements. It does, however, support the notion that the sorts of things that are used to describe elements in the commonsense world do not have to change significantly between cultures (except by way of minor calibration). This is a promising start because modelling frameworks such as those we study should also be culturally invariant, at least with respect to the things which we are trying to model.

Before one can be satisfied that a commonsense realistic ontology is suitable for our purposes we need to be satisfied that by adopting a commonsense realistic stance we are still able to have a socially agreed understanding of what there is, and that we do not require an understanding of the cognitive processes of human beings in order to understand what there is in the commonsense world.

In order to satisfy ourselves that the commonsense realistic outlook satisfies these two goals we can examine a definition of commonsense realism taken from Barry Smith (Smith, 1995):

> "The thesis that there is only one world towards which natural cognition relates is a central plank of what philosophers in the course of history have identified as the doctrine of common-sense realism. It is a doctrine according to which:
>
> (a) we enjoy in our everyday cognitive activities a direct and wide-ranging relational contact with a certain stable region of reality called the common-sense world;
>
> (b) our everyday cognitive activities rest upon a certain core of interconnected beliefs called "common sense" which is in large part *true* to the common-sense world as it actually is, not least in virtue of the fact that such beliefs and our associated cognitive capacities have arisen through interaction with this world;
>
> (c) this common-sense world exists autonomously, which is to say independently of our cognitive relations to it. Indeed from the perspective of common-sense realism the common-sense world exists entirely independently of human beings. Partial evidence for this thesis is provided by the fact that palæontology and related disciplines describe this world as it was before human beings existed. Of course this world would lack theoretical interest in a universe populated exclusively by creatures with cognitive capacities radically different from those of human beings. But what these disciplines describe is, nonetheless, such as to exist independently." (Smith, 1995)

A careful reading of the extract reveals that it leaves room to subsume the scientific reality or outlook while still allowing for a commonly held view or socially agreed reality and it reassures us that we don't require human cognition for this world to exist. This frees us from difficult questions that have been posited by idealists of various hues where the nature, and even the existence, of reality is only in the mind. It doesn't mean that the human beings in which we have an interest are not part of the world, because human beings are clearly part of this commonsense world. It is also worth pointing out that we are discussing the epistemological stance that Chisholm takes to his ontology. The ontology itself describes what is fundamental in constructing reality, not an epistemological system. Chisholm describes his epistemology extensively elsewhere in the literature.

Additionally, this school of philosophy allows for a difference between the reality and the appearance of reality. Often this is called the error that is involved in making sense of reality.

> "Thus common sense is not, in spite of its reputation, naïve; it draws a systematic distinction between *reality* and *appearance*, or in other words between

> the way the world is and the way the world *seems* or *appears* via one or other of the sensory modalities and from the perspective of one or other perceiving subject in one or other context. The thesis that there is only one world towards which natural cognition relates must thus be understood as being compatible with the thesis that there are many different ways in which the world can *appear* to human subjects in different sorts of circumstances." (Smith, 1995)

It is important, however, not to ignore the success in describing reality through science, and so we need to describe the relationship between the world as seen in common sense and that described by physics, the most closely related science. Later in this paper, Smith relates commonsense realism with physics:

> "The common-sense realist must confront the question of the relation between the common-sense world and the world that is described in the textbooks of standard physics. Here again a number of different philosophical alternatives have been mapped out in the course of philosophical time, including that view that it is the common-sense world that is truly autonomous while the world of physics is to be awarded the status of a cultural artifact. Here in contrast, we assume a thesis to the effect that the commonsense world overlaps substantially with physical reality in the more standard sense." (Smith, 1995)

This is important because physics, and indeed science generally, cannot be discounted in the quest for determining what there is. Science has been very successful in determining what there is in the physical world and is in part responsible for constructing models of physical reality. As paradigm shifts in science come and go, so will the models we use to describe physical reality. Paradigm shifts clarify our understanding of the physical world often importantly at the margins. We are experiencing the ramifications of just such a paradigm shift that began a little under a century ago. Issues such as string theory, quarks, and quantum mechanics generally have arisen from this paradigm shift, just as notions of mass, force, momentum, gravity, and planetary motion accompanied the Newtonian paradigm shift of several centuries ago. One therefore needs to be careful not to overstate the importance of science when examining ontology. Nevertheless, commonsense realism has a place for such development and clearly refuses to discard scientific realities in order to allow for social structures and understanding.

To date, the only ontology that has been adapted and used in related work is a scientific ontology by Mario Bunge (Bunge, 1977, Bunge, 1979). Bunge's scientific ontology does not well reflect the interpretive aspects in the critical commonsense realistic heritage and so an ontology of the commonsense realistic school may be a good ontology to integrate both the scientific and interpretive views required for information systems and computer science (Milton and Kazmierczak, 2000).

Summarising, we argue that the philosophical school to which Chisholm's ontology adheres seems compatible for our purposes. Chisholm's ontology is one of critical commonsensism and it appears to be related to commonsense realism. We do not claim that critical commonsensism is the only possible philosophical stance that is suitable. However, we have argued that it is worth pursuing the selected ontology since its philosophical stance is not at odds with our objectives. In the remaining sections of this chapter we introduce and describe Chisholm's ontology.

4·3 Chisholm's Ontology

In the following sections we describe Chisholm's ontology (Chisholm, 1992, Chisholm, 1996). It is the ontology selected for this study on the basis of pragmatics. The task of describing an ontology is a daunting one. Due to space constraints, we must necessarily be brief when compared with the original work. Interested readers are encouraged to read the monograph by Chisholm (Chisholm, 1996) for a more comprehensive treatment of the ontology. We illustrate Chisholm's ontology using examples where it is helpful. Through this description, we reveal the terms and associated concepts embodied in Chisholm's ontology. Roderick Chisholm has written extensively in the areas of ontology, metaphysics, and epistemology (Chisholm, 1989a, Chisholm, 1957, Chisholm, 1976, Chisholm, 1979, Chisholm, 1982, Chisholm, 1989b, Chisholm, 1992, Chisholm, 1996) and these works provide a backdrop to the 1992 and 1996 treatments.

After covering the categories established by Chisholm, we proceed in two parts. Firstly, we examine the concepts and terms concerning statics and consider individuals/things, attributes/properties, classes and sets, and relations. Secondly, we examine dynamic concepts and terms from the ontology and describe states, and events. States and events encompass the most obvious dimension of model change. However, there are other aspects to model flexibility and we are interested in the ability of the static model to change structure when reality, or rather our interpretation of reality, changes significantly.

Chisholm's ontology, like many mainstream metaphysical ontologies, concerns the "categories" of 'what there is'. That is, it presents categories into which entities from the world can be placed. We begin by introducing those categories. This gives us structure into which the concepts and terms can be seen to fit. The methods developed in Chapter 3 use both terms and concepts, and in that chapter we discuss how terms and concepts from the ontology are to be used. In that discussion we deliberately avoided talk of any one ontology since we envisage that our methods can be used with any ontology meeting our pragmatic requirements.

Throughout our description there will arise occasions where expressions will need to be defined. Usually these expressions have a specific meaning in philosophy that may be in conflict with common usage either in the general community or in the information systems or computer science communities. An example is 'entity'. Philosophers use entity differently from its usage in information systems. In Chisholm's ontology, entities describe all things that make up the furniture of the world, with the term 'individual' representing discernable objects within that world.

Chisholm's categories are organised into a taxonomy which is shown in Figure 7. The taxonomy shows what is fundamental to the ontology, but does not show the terms used to describe a state of affairs. The taxonomy doesn't show the relationships between the categories. For example, it does not show the relationship between individual and attribute. Relationships such as these are defined in the detail of the ontology. Chisholm adheres to the theory that establishes the dichotomy dividing the world into entities that are 'contingent' and don't have to exist, and 'necessary' entities that must exist in order for his theory to be consistent. The categories that play important roles in our methods are shown in bold type-face in this figure. Our overall aims lead us to these categories and their sub-categories. They will play the most critical role in both the comparison and the extension because they are similar to those used in data modelling frameworks.

Figure 7—The Categories for Chisholm's Ontology

(Chisholm, 1996)

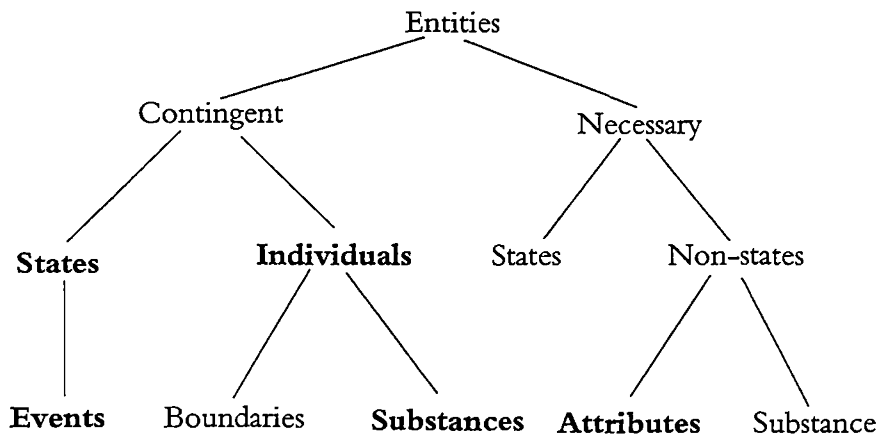

In many ontologies "[e]ntities are held to exist necessarily if natural processes will not lead to their cessation, contingently if such processes will lead to their cessation" (Honderich, 1995). For Chisholm, attributes are necessary because they do not need to be exemplified in order to exist as we have discussed in an earlier section of this chapter. This is influenced by the Platonistic view of attributes that Chisholm holds. Individuals on the other hand are contingent as they may be brought into being or cease to exist. There is a substance such as God (depending upon personal belief) that is non-contingent. States concern the exemplification of attributes by substances and will be contingent or non-contingent depending upon the nature of the substance concerned.

In defining the categories and concepts therein more formally, Chisholm takes a number of terms from philosophy and either uses them to establish the background for the ontology, or declares them "undefined" while at the same time explaining his attitude to them. Chisholm's ontology has eight such "undefined locutions." In this context, 'locution' means phraseology. Chisholm uses them to explain his ontology. We examine each after listing them.

Our eight undefined categorial locutions, then, are the following:

(1) *x* is necessarily such that it is *F*;

(2) *x* is a state of *y*;

(3) *x* is a constitutent of *y*;

(4) *x* believes that something is *F*;

(5) *x* senses *y*;

(6) *x* wholly precedes *y*;

(7) x spatially overlaps with y;

(8) it is a law of nature that p. (Chisholm, 1996)

At this stage, it is unclear as to what each is referring. For example, (6) refers to a concept called 'wholly preceding'. This could refer to either time or space. In fact it refers to time or more particularly to the ordering of events. Other concepts are unclear due to overlaying of philosophical expressions or concepts. For example, (1) talks about "necessarily such that." This is a well recognised logical expression in philosophy (Dancy and Sosa, 1992).

We describe the eight locutions below. These have been expanded with the help of Chisholm's monograph and of Flew (Flew, 1989) together with our other references (Dancy and Sosa, 1992, Honderich, 1995, Audi, 1995, Kim and Sosa, 1995). The importance of these locutions is not apparent at this stage. Some of these locutions will be actively used in establishing an understanding of the ontology for our purposes. Others will only be of passing interest.

(1) x is necessarily such that it is F

'Necessarily' is a standard term in philosophy and it refers to the logical (or metaphysical) state of a term. "A necessary statement (or proposition) is one which must be true—where this 'must' may be understood as being expressive of a logical necessity..." (Honderich, 1995). It is a term of logic not of causality. It applies "to a whole proposition (*dictum*) when it is not possible that the proposition be false (the proposition being *de dicto* necessary)" (Audi, 1995)

Chisholm also says that:

> The usual rules relating "necessarily" and "possibly" should apply; thus, "necessarily" implies "possibly", "not possibly not" implies "necessarily," and "not necessarily not" implies "possibly." (Chisholm, 1996)

Theoretically, 'x is necessarily such that it is F' may only be replaced by a predicate in which no variables occur freely (uninstantiated). Further, possibly is used through the ontology and also has a specific meaning. However, this meaning is in concert with common understanding, and so it is not in need of further explanation by Chisholm. Chisholm uses this definition to explain such things as the necessary nature of attributes. Its importance is as a logical and metaphysical expression through which definitions in the ontology can be understood and explained.

(2) x is a state of y (being a state)

State is not a widely recognised term in metaphysics or philosophy (as opposed to 'state of affairs'). However we can find references to state in concert with event. "...[M]any theories of events include *states* that consist in things' having (or retaining) properties..." (sic) (Honderich, 1995). Chisholm is one such theorist, and therefore is required to present and define the term. It refers to states of individuals, and states of other states. The example cited by Chisholm is "If you are thinking, then there is a state that is you thinking" (Chisholm, 1996). It can be applied equally to more mundane things. That is, states involving attributes more familiar to people generally. The critical element is that state involves an individual and an attribute (or property).

In Chisholm's ontology, states are required to order things temporally (the 'before' and 'after'). It is through the nature of the attributes concerned that governs the nature of the state. Further, and as a consequence causal relations are defined with the help of the concept of state. That is, one state contributing causally to another state.

(3) *x* is a constitutent of *y*

Individuals (defined later), "may be divided into (a) those that are necessarily such that they are constituents and (b) those that need not be constituents." (Chisholm, 1996). This locution allows us to talk about individuals made up of other individuals, thereby saying 'that *x* is a constituent of *y*'.

(4) *x* believes that something is *F*

It is through thinking or believing that an observer can establish that a thing exhibits an attribute. Indeed, Chisholm uses this as a basis for defining the concept called 'attribute'. It is required in order that someone can *believe that something is F,* it is essential that they have already *thought about something being F.* In many ways, this is not of critical importance to our use of the ontology. It does, however, highlight a concept of philosophical importance.

(5) *x* senses *y*

Described by Chisholm as being *psychological* in nature, this locution stresses the need to have a being "capable of thinking" in order to be able to sense anything.

(6) *x* wholly precedes *y*

This is a temporal concept that allows us to reason about the ordering of states. If *x* and *y* are states, then this comments on the temporal ordering of *x* and *y*. These are critical to our understanding of dynamics in Chisholm's ontology. Once wholly precedes is defined the remaining relationships follow (Russell, 1946).

(7) *x* spatially overlaps with *y*

This penultimate locution is "designed to throw light on the spatial nature of material things and on the distinction between such things and the boundaries they contain (surfaces, lines, and points)" (Chisholm, 1996). It concerns what it means to have a boundary between material things. That is, in determining where one material thing ends and another begins.

(8) it is a law of nature that *p*

Described as a "nomological concept". That is, 'lawlike generalization' (Audi, 1995). This deals with the naming of laws of nature. Chisholm sees this dealing with statements such as "It is a law of nature that *p*." where "'*p*' can be replaced by any well formed declarative English sentence". Presumably, a sentence in another language would also suffice for speakers of that language. Chisholm goes further, and declares that one cannot have these as constants. That is, our naming of generalisable 'laws of nature' may change as our models of science also change, to either a large or small extent, through experiment and observation. Essentially, we mean that these are nomic or "... meaning *scientifically* lawlike..." (Honderich, 1995). They are subject to paradigm shifts and similar scientific influences.

Having introduced some undefined concepts, we now describe the static and dynamic concepts in Chisholm's ontology that are of interest in our work. As the terms are introduced, we describe the concepts giving meaning to the terms. We use terms and concepts in the manner in which they were introduced in Chapter 3. At the end of each of the following sections we present, in tables, terms and concepts for the statics and dynamics of Chisholm's ontology. In the description of the key concepts we will, where appropriate, refer to the locutions we have introduced in the last few pages. We attempt to illustrate the concepts by way of examples. Further, we include concepts that play a role in our unifying framework. There are sections of the ontology that are epistemological and therefore not of interest to us for our goals and objectives.

4·4 Static Concepts in Chisholm's Ontology

Chisholm's ontology centres on *individuals* and the *attributes* they exhibit. Further, Chisholm holds that attributes are fundamental to his ontology. As we will see later, he reduces other terms that we need so that they are defined only using attributes. The terms and associated concepts of 'individual' and 'attribute' have descriptions that show not only their individual disposition, but also their roles in *sets*, *classes*, and *relations*. We also describe these below. These terms are not *fundamental* to the ontology but nevertheless are important terms that are discussed and are appropriate for our goals. We now examine each of the key terms concepts in turn.

Individual

An *individual* is a discernable and transient object. It need not be material (or physical) in nature. A person is an individual. Examples of individuals are an accountant named *Freda*, the annual financial statements for *Ericsson*, and *Orly International Airport*.

As we shall see later, individuals are identified using attributes that only they exemplify. Further, individuals may have *constituents* thereby giving them structure. Constituents may be other individuals (called *parts*) or may be boundaries (the other constituents). For example, consider *Orly Airport*. It has several rent-a-car franchises, bars, restaurants, departure gates, each of these are parts of *Orly Airport* and are also individuals. In this example, most of these can be further sub-divided.

On the other hand spatial substances have boundaries. A boundary is either a surface, line, or point. For example, Orly Airport may have as its constituents surfaces that help to identify it as a spatial object. That surface is a boundary and is in turn made up of a number of surfaces, lines, and points.

Attribute

An individual may exemplify attributes. *Orly Airport* is very busy; *Nokia's* balance sheet is good; *Freda*, our accountant, is of age 43. Some attributes *may never* be exemplified and others *cannot* be exemplified. For example, Orly Airport *may never* be green. We can be sure that Orly Airport *can never* be a liquid.

Two attributes may be equivalent in the sense that if one is exemplified by an individual the other is also. This relationship is called *conceptual entailment.* This can be illustrated by considering Orly Airport. The attribute very busy may involve a conceptual entailment with the attribute of having over a certain number of aircraft movements an hour.

Chisholm allows for compound attributes. Further, compound attributes may be consist of other compound attributes or simple attributes. He suggests that an attribute may be the conjunction or disjunction of several attributes. For example, the attribute of 'being good' with respect to Nokia's financial statements may be the conjunction of being in surplus (profit) and being of good credit rating. Chisholm also indicates that there exist alternative mechanisms for providing compound attributes.

Philosophically and logically, it makes little sense to talk about *when* an attribute came into being. We will discuss passing away and coming into being in the next section, when we discuss dynamic elements in Chisholm's ontology. However, Chisholm asserts that attributes are enduring, thus avoiding the problem of declaring when an attribute comes into being. For example, when did the attribute 'being green' first come into being? Since we cannot know and since raising its genesis brings about certain problems it is better to adopt the position that attributes are non contingent, they exist perpetually.

Classification

In Chisholm's ontology, attributes are used to restrict membership of sets and classes. Further, Chisholm reduces discussion of classes to discussion of attributes by adopting Russell's reduction of classes to attributes (Russell, 1908). This has the effect of building classes and sets from individuals through the exemplification of attributes and not by constructing elaborate class structures. In this way flexibility is maximised. For example, suppose we are maintaining a taxonomy of plants. Periodically, the taxonomy may change quite drastically without a change in the majority of attributes exhibited by the plants involved. Using Chisholm's ontology classes can change radically through a change in membership criteria based on attribute exemplification.

Classes and sets can be selected based upon attributes that are conjunctions and disjunctions of other attributes, and in this sense complex class relationships can be realised, that are essentially class structures. The central point remains, that individuals come together to form classes and are fundamental to the ontology. Classes are reflections of attributes exemplified by individuals due to the fact that they exemplify the attributes that are used to select the class.

Relation

Individuals may be related. Chisholm, on the subject of relations says: "To know what relations are, we must understand the concept of the direction of relations." Chisholm means that relations may not be reciprocated. For example, I may be interested in a job with Nokkia. Nokkia may not be interested in employing me. It is for this reason that relations are *unidirectional.* Further, in the ontology, relations are attributes in the form of ordered pairs. This results from the fact that ordered pairs are related to sets of a specific form; Therefore ordered pairs can be reduced to being represented by attributes through mechanisms already introduced.

Table 5—Concepts for Statics in Chisholm's Ontology

Concept	Part	Description
Individual	Core	Chisholm allows for discernible and transient objects. These are called individuals. Individuals come into being (are created) and pass away (destroyed). In this sense they are transient.
	Identity	Each individual possesses an attribute (or several attributes) that uniquely identifies it.
	Structure	Individuals may have constituents. These are either other individuals (known as parts) or boundaries (the other constituents). Individuals that make up parts of others are still thought of as being individuals.
Attribute	Core	Attributes are exhibited by individuals. They are central to Chisholm's ontology, after individuals. Further, attributes are enduring, in the sense that they don't come into being and don't pass away. Further, attributes must be loosely coupled with individuals.
	Equivalence	Attributes can be equivalent in the sense that if something exhibits one attribute then it exhibits the other. This is called *conceptual entailment.*
	Complexity	Attributes may be simple or complex. Complex attributes are combinations of either simple or other complex attributes. The mechanism suggested by Chisholm is one involving conjunction and disjunction of attributes. He feels there may be *other* ways of providing for this complexity.
Classification (classes and sets)	Core	Classes and sets are provided using attributes, in the ontology. Specifically, it is through the attributes that membership of classes is determined.
Relation	Core	Individuals may be related. Specifically, relations are attributes (an ordered pair). The ontology requires that attributes that identify the participating individuals are required. Further, that the relations are unidirectional (not bidirectional).

For an ordered pair to represent unidirectional relations, attributes that uniquely describe each individual need to be found. For example, suppose that Freda (our accountant) is recruited to audit Nokia's books then there is a relation between Freda and Nokia. Consequently, an attribute consisting of an ordered pair of attributes identifying each of Freda and Nokia would have to be exemplified by Freda. A reciprocal attribute representing the reverse relation would need to be exhibited by Nokia. In the simplest case an individual may be related to another (binary). More complex relations between three individuals (ternary) or more (n-ary) are allowed. Mathematically it is proven that these all can be reduced to a series of binary relations (Quine, 1960).

Summary of Statics

Summarising, «attributes» as revealed in Chisholm's ontology is used to define a number of key terms and concepts. The nature of attributes is such that they may conceptually entail each other and they must be considered to endure (they are necessary). Further, they are used as the basis for sets, and classes and for the provision of relations between individuals. Propositions and facts can also be represented by attributes.

4·5 Dynamic Concepts in Chisholm's Ontology

Chisholm's ontology represents a significant shift in realistic thinking. "In metaphysics, Chisholm is open and precise about his ontological commitments, which, according to his most recent writings, include only attributes and individual things" (Honderich, 1995), and as we have seen so far, this is reflected in his ontology. Essentially, this means that to describe 'what there is' and the nature of change, one only needs attributes and individuals. By implication, and as we shall see here, this means that there is nothing ontologically fundamental about time, and that time as a measure or marker is not one of the fundamental terms or categories of 'what there is'. This in turn means that we must allow for temporally oriented attributes instead of time so that one can talk about temporal relations (ordering) of the exemplification of attributes by individuals (state), and of changes in state.

Chisholm envisages events without time. "... we reject the view of Jaegwon Kim, according to which an event is 'a structure consisting of a substance, a property and a time' " (Chisholm, 1996). Based on Chisholm's understanding, and using a recognised source, an event is "... defined as a change (for example, the loss or acquisition of a property by something) or composite of changes. However, many theories of events include states that consist in things' having (or retaining) properties" (Honderich, 1995). Commenting on his ontological choices and its implications for events he continues by saying "... [s]uch an approach to events has the advantage of not requiring that only substances be the substrates of events. It also allows us to say that those states that consist of one state contributing causally to another state are also events" (Chisholm, 1996).

For Chisholm, it means that event and state can be unified because time is no longer required as part of an event, and that states allow for modelling dynamics, because events are a specialisation of state (essentially a higher-order type of state). Consequently, it is the concept of temporal relations instead of time that is critical to being able to order states, and by implication events.

We begin by discussing the reasons why time is rejected by Chisholm, and the role that temporal relations fills in a theory without time. We then introduce /state/, the most difficult concept in Chisholm's ontology.

Temporal Relations

Time is often thought of as being required in order to individuate and identify events. For Chisholm, time is not an ontologically fundamental concept, and instead, he uses the concept of temporal relations between states, to order states. Chisholm sees no reason for specific times and dates such as "18:00 on January 1, 1982" to exist as separate *fundamental* entities in reality, because there is little useful one can say about that time other than things that are culture specific, and there is nothing ontologically important about a particular date and time. For example, the Gregorian calendar may mean very little to a person in a Muslim culture or to someone born into a culture where the Julian calendar is used. We could construct a different, yet equally valid subdivision of time, and therefore, any subdivision and marking of this nature is arbitrary and not fundamental to 'what there is.'

Instead of time there are "temporally oriented attributes", as Chisholm puts it, that things exemplify. It is still critical to order states and thereby changes in state to enable a meaningful discussion of change, and Chisholm removes time from ontological consideration by insisting that all attributes have a temporal orientation when exemplified by an individual. In modelling, we are often most interested in attributes that are oriented to the present, and individuals that exemplify an attribute at the present time, are said to exemplify the attribute now. Other attributes may be oriented to the past or to the future. Together, these enable us to discuss past, present and future state.

As we shall see, there are attributes that form part of state and of events. Chisholm uses his formulation of event and state that are unified, together with Russell's results (Russell, 1908) concerning the temporal ordering of events, based upon an undefined but explained concept of wholly preceding that we introduced at the beginning of this chapter, and temporally oriented attributes in order to remove the need for time from events. However, as Chisholm indicates, he is deviating from related philosophical views, so we take the stance that it is the temporal ordering of states and events that is required and covered by this concept. Since many philosophers have not yet rejected time as being fundamental, we view his section on temporal relations as a reiteration of the need to temporally order events and state.

State

In removing time from ontological consideration Chisholm "replaces the undefined concept of *a time* with the undefined concept of *a state*" (Chisholm, 1996), and so state is one of the "undefined concepts" mentioned in our introduction to this chapter, and is responsible for a number of disparate aspects of exemplification and change in exemplification of attributes. There are five parts to /state/ that reflect these disparate aspects. Firstly, /state core/, that outlines the core of the concept. Secondly, /state structure and parts/, deals with the complexity of the structure of state. Thirdly, /state event/, concerns states that are events. Fourthly, /state causal relationships/, a concept that enables description of 'generalizable laws' that we described earlier, and finally, /state internal and external attributes/ for individuals that have ceased to exist but still exemplify attributes in the eyes of those who still discuss them.

/state/ is, by far, the most powerful of the concepts covered thus far in this thesis. Being undefined, /state/ is difficult to explain, but nevertheless, it mirrors the simplicity of the remainder of the ontology by using individual and attribute as the foundation for state, and consequently only one concept is needed for individual dynamics and state. Most states involve the exemplification of attributes by individuals and, further, that event and state are in the same category in the ontology. Despite /state/ being simple, it is difficult to grasp the full ramifications of the way in which Chisholm structures state and it is also difficult to explain it. We have taken care in our attempt to expose the subtleties of the concept. However, we are comparing the meaning of concepts as they are explained by Chisholm, and a more detailed investigation of /state/ is flagged for future research. We now discuss the five parts to the concept.

States enable the exemplification of attributes by individuals, and must also enable events. States begin and end. They can endure or they may be instantaneous. They may be compound, consisting of parts and fitting together in a way so as to form what is called a 'temporal whole' of states, and compound states are sometimes described as a 'history' of an individual and of its parts. For example, the history of a country could be a complex consisting of the states of all individuals (people, cities, animals, etc.) within its jurisdiction. Some types of states are events, and this unification of state and event is made possible by the assertion that there are events without times.

In its simplest form, state involves a contingent substance (individual) and a property or attribute. For example, the individual known as *Nokia* is presently financially viable. One may say that there is a state involving Nokia and the attribute or property of financial viability. Indeed in this example, Chisholm would call Nokia the *substrate* of the state and the attribute being financially viable its *content*. In this context substrate means the basis or bedrock or the subject of the state. The substrate may be an individual. In contrast, content means the adjective or description of the state. That is, the content may be an attribute, the real crux of the state of the thing involved. Generalising, we can say that states in the ontology consist of a substrate and a content, and it is the complexity and nature of the substrate and content that determines the complexity and nature of the state itself.

States *begin* and *end*. Clearly the ending of a state and the beginning of another involves change. We can describe, with respect to states, *beginnings*, and *endings*. At this point we pause to also consider 'coming into being'. This is used to define not only the beginnings and endings of states, but also the coming into being of individuals. A state or individual, comes into being if there are (present tense) no attributes that it had (past tense). Further, the 'passing away' (or 'ceasing to be') is described similarly to be that there are no attributes that it will exemplify. For example, Nokia is said to come into being if there are no attributes that it had. If it had already come into being, then it would be the case that there are attributes that it formerly had. Further, the orientation of a state or exemplification of attributes to the present, past, or future, will depend upon the orientation of its attributes, an issue we encountered at the end of the last subsection, and which clearly has an affect on this part of /state core/. 'Ceasing to be' is similarly defined in terms of its attributes. For example, should it happen that Nokia be closed as a business then it will be said to have ceased to be if there are no attributes that it will have. This is, in turn, affected by internal versus external attributes, because we may still discuss Nokia after it has closed. We introduce this as the last concept in /state/ (below).

States may be instantaneous (where the state is simultaneously an ending and beginning) or states may be enduring. An enduring state of a thing is what is called by Chisholm a 'temporal whole'. This means the state may have as parts that, if in existence simultaneously, would have contents that are incompatible. In other words, the parts of a state change. The parts may be incompatible. For example, your life is an example of a temporal whole, and is an enduring state. At some point you may be financially wealthy, and at others you may be financially poor. These two states (part of the whole) are incompatible, but are temporally related by the temporal relations we have already described that enable us to discuss changes in state. An example of an enduring state is the financial history of Nokia. It could be encapsulated by one enduring state consisting of many states. Some states are past oriented, some are present oriented. As we progress the substance (and associated substances) undergo change whereby we pass from one present state to another.

/state core/

«state core» is a very broad part of /state/ in that it describes what is needed in terms of support so that if such a part of the concept were not present, the concept from the ontology would not be realised and be rendered meaningless.

The core of /state/ is that firstly, the exemplification of attributes by individuals must be supported. Secondly, that instantaneous and enduring events as represented in a subcategory of state in the ontology must be supported. An event is "... defined as a change (for example, the loss or acquisition of a property by something) or composite of changes. However, many theories of events include states that consist in things' having (or retaining) properties" (Honderich, 1995), and changes in exemplification of attributes and a composite of changes that are enduring or instantaneous must be supported. An example of an event is President Clinton speaking. We discuss the elegance of Chisholm's unification of event and state in a later part of the concept. Finally, states, events, and individuals, all being contingent things must be able to be come into being, and cease to exist.

/state structure and parts/

/state structure and parts/ concerns the relationship between an enduring state and, firstly, the temporal wholes (states) that make up its content, and, secondly, the states of the parts of individuals (or more generally of the substrate) involved in the state. This encapsulates simultaneously, state complexity as reflective of individual (or substrate) structure together with state complexity that reflects temporal wholes. This concept of structure and parthood can be described as:

State x is part of state y = Df Either
(1) the substrate of x is a proper part of the substrate of y, or
(2) the content of x is part of the content of y

Firstly, we deal with substrate parthood. In simple cases, the substrate of a state is a substance; You will recall that a substance is a contingent individual that is ontologically independent. This means that for *Nokia*, part of its history (an enduring state) is the history of its *Accounts* department, because that department is part of Nokia. Recall the discussion of constituents and parts in the statics of the ontology.

On the other hand, the parthood for content requires further description because it concerns attributes and simultaneously past, future, and present state. Part of the content of a state can be defined as:

The content of state x is part of the content of state y = Df
(1) the content of y implies the content of x; and
(2) the content of x does not imply the content of y

This describes the relationship between the attributes exemplified by a substrate over time. However, the change results in some of the content remaining. For example, suppose Nokia is in the enduring state of being financially viable (call the content of this *f*). Suppose in the past they've not only been financially viable, but that they also have the most cash on hand of any Norwegian company (call this *c*). A change occurred. Nokia is still financially viable, but they are no longer the company with most cash on hand, however they are now the best company with respect to safety record for their employees (call this *s*). Clearly, there is a part relationship here between the content of the whole (*f*) and its parts (*c* and *s*). It is a part relationship. That is, *c* and *s* are both parts of *f*. The content *c* is the two attributes, being financially viable, and being most cash on hand in Norway; The content *s* is being financially viable and having best safety record in Norway. The enduring state applying to Nokia is the state of being financially viable. The relationship between past and present states is achieved by using for the respective contents (c and s) past oriented and present oriented attributes respectively.

/state event/

Not only are contingent individuals such as Nokia the substrate of a state, but a state can be the substrate of another state. Therefore the key point to note is that the parthood of substrate includes the parthood of state (a reflexive element) since a state can be the substrate of another. This sort of state is called a 'higher order state'. Chisholm sees a particular need for first-order and second-order states. A first-order state is one where the substrate of the state is a substance or individual and the content is an attribute. A second-order state is the state of a first-order state. An example of this second type is that of a causal relationship where one state contributes causally to bring into being a second state. An event may be either a first-order state or a second-order state. The content of a second-order state is a state. Further, Chisholm hold that common usage restricts events to these two orders of state.

Most first-order states are involved in attribute exemplification. In order for a first-order state to be classed as an event, one must be careful not to include what Chisholm calls non-events, and to help distinguish between the two we present contrasting examples. Firstly, Nokia being financially viable is a first-order state, and is clearly not an event. However, Nokia being in the state of shifting headquarters is also a first-order state, and is clearly an enduring event.

A second-order state is an event, in that it is a state of a state. For example, suppose two people, Mathilde and Petra interact, and we may have the second-order state "[Mathilde being angry] contributing causally to [Petra being upset.]" [Mathilde being angry] is the substrate; [Petra being upset] is the content (a state). This can be described as an example of a epitome of an event.

/state/ can nest infinitely, and orders higher than two are possible, but Chisholm sees no practical application of orders higher than two.

/state causal relationships/

Often events are described to have recurred. Chisholm rejects that notion in favour of the idea that the exemplification of attributes recur. For example, I may repeatedly serve a ball in tennis game. Some may say that the event of me serving recurred. Chisholm would say that the attribute of serving being attributed to me recurred, not the event. Furthermore, Chisholm holds that *causation* is in fact what people really imply when they discuss recurring events, and that causation is not a conjunction of event templates. To further describe causation we utilise the last undefined philosophical locution (*#8*) on page 59 concerning the *laws of nature* (lawlike generalisation). In rejecting the repetition of events, and further, that events cause others, Chisholm needs to name laws of nature so that an event may lead causally to another event or state according to those laws. These laws of nature cannot exist *a priori*. Evidence or support must be found in reality to support them. This is the crucial distinction between laws of nature, and laws of logic. Logic relies on *a priori* assumptions not empirical evidence.

Causal conditions can be established based on the locution concerning laws of nature. Further, a minimal sufficient causal condition stipulates the events that must occur before the causal relationship can be acted upon. Further, no subset of that minimal set can cause the event. Chisholm formalises it in the following way:

S is a minimal sufficient causal condition of *E* = Df
(1) *S* is a sufficient causal condition of *E*; and
(2) no subset of *S* is a sufficient causal condition of *E*

Furthermore:

S is a sufficient causal condition of *E* = Df *S* is a set of events that is such that it is causally, but not logically, necessary that *E* occurs either when all the members of *S* occur or later

Again, one can also describe a partial cause for an event:

C is a partial cause of *E* = Df *C* is a member of a minimal sufficient causal condition of *E*

It is entirely possible that an event may have a number of minimal causal conditions. Also, an event can be called causally determined if there exists at least one sufficient causal condition. Further, some events may be causally determined where others may not be so determined. We mean that a determined event is one where there is at least one sufficient causal condition. In contrast, an undetermined event is one for which no sufficient causal condition is found.

/state internal *vs* external attributes/

A natural question in response to the 'ceasing to be' of individuals concerns individuals that have ceased to be, but about which we still speak, and further, that we believe exemplify attributes. This then deals with the rationale by which this is possible if the individual does not exist. For example, the mythical city Atlantis has long since ceased to exist (if indeed it ever did). We, nonetheless, attribute properties to that city. That is, we can describe its state as people attribute properties to it. Chisholm deals with this by establishing a distinction between internal and external attributes. An external attribute requires something attributing a property to that thing. For an internal attribute to be exemplified, it requires that the thing that exemplifies the attribute to be a substance and exist irrespective of anything attributing the property.

Table 6—Concepts for Dynamics in Chisholm's Ontology

Concept	Part	Description
Time and temporal relations	Core	Temporal relations order states. Chisholm rejects time as being ontologically fundamental and uses mechanisms different from time to order states.
State	Core	Individuals exemplify attributes and that exemplification is a state. Events represent changes in state. Individuals, and states are contingent and therefore cease to exist, and come into being.
	Structure and parts	A state consists of a substrate (*usually* a substance) and a content (*usually* an attribute). Both of these can be divided further into parts according to need. For example a substance has parts. It's state will be divided in similar ways to reflect this. Similarly, a state's content reflects the attributes the substance exhibits. A state's content may change during a period, but some attributes may continue to be exhibited. In this way, parthood for content is defined. See description of Chisholm's statics for a fuller explanation.
	Event	Events are states. Specifically, they are appropriate first-order and all second-order states. Second-order states have structure much the same as first-order states (substrate and content are more broadly defined). Further, they don't recur.
	Causal relationships	These describe 'laws of nature'. That is minimal sufficient conditions that will lead to an event occuring. This does not imply that there are templates of events that permit them to recur. Only attribute exemplification recurs.
	Internal vs external attributes	Individuals, like states, come into being and pass away. However, some individuals may still exhibit attributes after they have passed away. These are called external attributes, since they need something else to have the experience of still attributing a property after the individual has ceased to exhibit internal attributes.

Summary of Dynamics

Summarising, Chisholm's ontology provides for /state/ that unifies what others call event and state. Further states can be described to have order. First-order states that are appropriate and all second-order states are events. To paraphrase Chisholm, one could say that these are events "*par excellence*". There is structure in state. Further, a temporal whole can be defined that is an enduring state. One example of such a temporal whole is the history of an individual. Causal relationships describe laws of nature. These laws are established using empirical evidence. Further, events are not thought of as recurring, it is instead the exemplification of attributes that recurs. Finally, individuals can come into being and pass away. It is therefore required, for individuals that have passed away and for which one still feels exhibits attributes, the vocabulary concerning internal attributes must be described.

CHAPTER 5

The Candidate Data Modelling Frameworks

WE HAVE SELECTED five data modelling frameworks from the literature to use in this thesis. Firstly, the Entity-Relationship (ER) model (Chen, 1976). It is one of the most important entity-relationship-attribute style of modelling framework and is still extremely popular in industry. Secondly, the Functional Data Model (FDM) (Shipman, 1981) is the cleanest example of the functional-style of semantic data model. Thirdly, we include the Semantic Data Model (SDM) (Hammer and McLeod, 1990, Hammer and McLeod, 1981) recognised to be an important model (Hull and King, 1990, Hull and King, 1987, Peckham and Maryanski, 1988) and may be considered to be a significant model with respect to object data model development. Fourthly, NIAM (Nijssen and Halpin, 1989). Finally, the object model specified in OMT (Blaha and Premerlani, 1998) is included as a significant object model that is used in contemporary information systems development yet, in part, originating from the data modelling community.

This chapter divides naturally into five sections reflecting the number of specific data modelling frameworks selected. Each section proceeds with a general discussion of the modelling framework under examination and concludes with a discussion of strengths and weaknesses. Further, a list of terms and concepts is given in tabular form as specified in Chapter 3. The terms described and associated concepts emerge from original sources. Where appropriate supporting evidence is quoted from sources.

All data modelling frameworks are subject to the comparative method. The results for the comparison are presented in Chapter 6. Only the functional data model and the OMT family of models are subject to the evaluative method. Results for these applications are found in chapters 8 and 9 respectively.

5·1 Entity Relationship Model

ER modelling frameworks (sometimes called Entity-Attribute-Relationship models) were first proposed in the late 1960s, and the most popular of which is the model proposed by Chen (Chen, 1976). Chen's modelling framework has since been augmented with extra constructs (Batini et al., 1986, Teorey et al., 1986) to form the Extended Entity-Relationship Model. For the sake of completeness we will include elements of the extended model.

ER models are centred around three fundamentals concepts: entity-type, relationship-type, and attribute-type (Peckham and Maryanski, 1988). Members of entity-type are identifiable entities (either physical or conceptual) from the world. 'An *entity* is a "thing" which can be distinctly identified.' (Chen, 1976) These are then classified into entity sets (Chen, 1976). Some ER approaches allow disjoint entity sets—entities may be members of more than one entity type, although the semantics of such entities are strained. Attributes are properties held by all members of a given entity-type or relationship-type. Relationship-types are defined over one or more entity-types. 'A *relationship* is an association among entities.' (Chen, 1976) In other words they *relate* entities. This is a concept which is fundamental to all ER approaches. Some attribute or attributes forms a 'primary key' that is 'needed to identify the entities in the entity set' (Chen, 1976).

Relationship-types are of several dimensions, characterised by the number of entity-types allowed to participate in the relationship-type. Additionally functionality (1-1, 1-m, m-n nature of the relationship-type) together with partial or total participation in a relationship is described as cardinality. Cardinality specifies the minimum and maximum levels of participation in relationships for members of entity-types).

Table 7—Concepts for the Entity-Relationship Model

Concept	Part	Description
Entity	Core	ER allows for significant entities, or objects (either physical or conceptual) to be modelled. These must be grouped into entity classes. Each entity cannot depend upon other entities to be classed as an entity.
	Identity	Each member of an entity class must have a unique identity called a key.
Relationship	Core	Relationships are defined over entity classes. Members of specified classes may *participate* in the relationship.
	Cardinality	Each relationship has cardinality. This specifies the minimum and maximum numbers of entities from the participating entity classes that can engage in relationships.
Generalisation	Core	A special category of relationship that allows for *is-a* types of relationships. It allows for classification of entity classes, although it is forced in this technique.
Attribute	Core	Attributes are specified over entity and relationship classes. They are tightly bound to these classes.

Recently, it has been recognised that relationship types are overloaded when specifying what are essentially, 'is-a' relationships. In response to this concern a special concept known as *is-a* has been included in many ER approaches. Nevertheless, '... the ER model was the first semantic model centred around relationships, not attributes.

It views the world as consisting of entities and relationships among entities. Both entities and relationships may have single-valued printable attributes.' (Hull and King, 1987)

ER is popular because of its easy-to-understand concepts but suffers from the drawback of rigidity. The rigidity concerns the tight binding between entities and the attributes they present. Further, it is difficult and forced to represent entities that are members of several classes simultaneously. The advantage of this modelling framework is that it is simple containing a small number of easy to understand concepts. It is debatable whether people find them easy to use or to construct models using the modelling framework. Significant work is being undertaken to address the quality of models constructed using ER. Other limitations concern the rigidity that is present due to the class focus of the modelling techniques, and the tight binding between entity classes and the attributes exhibited by entities in the class.

5·2 Functional Data Model

The functional data modelling framework (Kerschberg and Pacheco, 1976, Shipman, 1981) is one of the simplest in the family of semantic data models, and is also "recognized as the first semantic model centred around functional relationships, that is, attributes" (Hull and King, 1987). 'The basic constructs of [FDM] are thc *entity* and the *function.*' (Shipman, 1981)

Although not initially obvious, entities are of two types. Firstly, they can be identifiable objects or things in a way similar to other modelling frameworks. We call these object-type entities. Secondly, they can be attributes or properties. These, we call attribute-type entities. An example of an object-type entity is the pilot of an aeroplane. The number of years a pilot has been flying is an example of an attribute-type entity. Further, these are associated with an object-type entity.

'... function, in general, maps from a given entity into a *set* of target entities.' (Shipman, 1981) These are used to express relationships between object-type entities (object to object mapping) or to associate attributes with object-type entities (object to attribute mapping.) These relationships are unidirectional. Bidirectional relationships between object-type entities are modelled using inverse functions. Further, the meaning of the function is determined by the participants in the function. For example, exemplification of attributes is shown using a function mapping from the object-type entity to the attribute-type entity concerned. In the functional data model attribute-type entities are separate from the object-type entities they describe. The two are associated using functions. Essentially, 'FDM connects objects directly with attributes without the use of intermediate constructs such as aggregation or grouping.' (Hull and King, 1987)

Functions can also be used to mimic a number of other modelling terms. Firstly, they can implement a form of aggregation of entities. 'Multiple argument functions... can be used to represent aggregate relationships between multiple entities' (Peckham and Maryanski, 1988) This is achieved by essentially declaring mapping from a combination of two or more object-types to an attribute type thereby declaring attributes of the more complex object. Secondly, classification (including generalisation) can be mimicked using functions. Essentially, '... we have implied a number of subtype and supertype relationships.' (Shipman, 1981) More flexible classification based on criteria (through attributes) is also possible using functions. Elsewhere, '[t]he model does not provide explicit means for generalisation and classification, although the user may define functions representing these abstractions.' (Peckham and Maryanski, 1988)

One key strength of FDM is its simplicity, in that '[t]he designers of this system [FDM] found, however, that limiting the constructs to entity and function provide a direct and simple language for data definition and manipulation.' (Peckham and Maryanski, 1988) Essentially, only two constructs are used to construct models. Further, all attributes and relationships are described using functions. There is a separation between the concepts for entity and relationship. Overloading of entity to include both attribute-type and object-type entities is unusual, but in practice does not seem to present problems, and it is in this way that attributes are loosely coupled with entities.

Table 8—Concepts for the Functional Data Model

Concept	Part	Description
Entity	Core	An entity is either a significant thing in reality, or it is an attribute. Firstly, identifiable objects are entities (object type). Secondly, the values of attributes are entities (attribute type).
	Object type	This models significant and distinguishable objects from reality.
	Attribute type	This type of entity records attributes of object-type entities. These are loosely coupled with the object-type entities they describe.
Function	Core	Functions map from entity to entity. It supports several types of concepts. Firstly, relationships; Secondly, attributes; Thirdly, classification; Fourthly, crude aggregation.
	Relationship mapping	Functions may map from object-type entities to other object-type entities. These are relationships. They are unidirectional. Bidirectional relationships are modelled using inverse (or reverse) functions.
	Attribute mapping	Mappings from object-type entities to attribute-type entities represents the exemplification of attributes by object-type entities.
	Aggregation mapping	Mappings from two or more object-type entities to one attribute-type entity. Each mapping representing an attribute of the aggregated entity.
	Classification mapping	Mapping from object-type entities to a list of object-type entities. Can be used to group like entities or to select entities based on some qualification. This can mimick generalisation/specialisation.

5·3 Semantic Data Model

The semantic data model (SDM) (Hammer and McLeod, 1981) was designed to improve the modelling of semantics in data modelling frameworks. It is one of the first data modelling framework to incorporate the concept called *class*. '... SDM was among the first published models to emphasize the use of the grouping constructor ...' (Hull and King, 1987) Further, it incorporated an increased number of features compared with other semantic data modelling frameworks where the focus tended to be on a reduction in the number of constructs.

'The following principles of database organization underlie the design of SDM.

(1) A database is to be viewed as a collection of *entities* that correspond to the actual objects in the application environment.

(2) The entities in a database are organized into *classes* that are meaningful collections of entities.

(3) The classes of a database are not in general independent, but rather are logically related by means of *interclass connections*.

(4) Database entities and classes have *attributes* that describe their characteristics and relate them to other database entities. An attribute value may be derived from other values in the database.

(5) There are several primitive ways of defining interclass connections and derived attributes, corresponding to the most common types of information redundancy appearing in database applications. These facilities integrate multiple ways of viewing the same basic information, and provide building blocks for describing complex attributes and interclass relationships.' (Shipman, 1981)

Essentially, SDM is centred on abstract *entities* that can be collected into classes of meaningful collections of entities. Classes are named and have entities as members. Essentially, '[c]lassification and association have greater emphasis in SDM than aggregation and generalisation. An SDM database is a collection of entities (instances) organised into classes or types.' (Peckham and Maryanski, 1988) Entities can be concrete objects, events, categorisations (or higher order entities), or names. Attributes are defined in terms of classes. Firstly, member attributes are used to describe members of a class. Further, class attributes describe properties that the class as a whole possess. 'The designer defines classes and within this framework specifies member and class attributes, interclass connections, and derivations.' (Peckham and Maryanski, 1988)

There are base classes and nonbase classes. Base classes are modelling primitives and entities only really exist in these base classes. 'Base classes are mutually disjoint in that every entity is a member of exactly one base class.' (Shipman, 1981) Base classes have associated groups of attributes called *identifiers* that can be used to uniquely identify member entities. For example, there may be two alternative identifiers for entities in a base class, and in this case two groups of unique attributes are defined. Further, base classes can be defined to either allow or disallow duplicate members. Nonbase classes do not have a separate existence from the base classes in that non-base classes are related back to base classes. "A nonbase class is one that does not have independent existence; rather, it is defined in terms of one or more other classes." (Shipman, 1981) Nonbase classes are linked to base classes using interclass connections.

Interclass connections fall into two categories. 'There are two main types of interclass connections in SDM: the first allows subclasses to be defined and the second supports grouping classes' (Shipman, 1981). In the first case it is a connection where the entities of the subclass are a subset of the entities in the superclass. A predicate is used to restrict membership when compared with the superclass. The predicate may take one of four forms. Firstly, a predicate on the member attributes of the superclass. Secondly, 'where specified' is used when the subclass is an *ad hoc* collection of entities from the super class. Thirdly, as an intersection, union, or difference of two other classes. Fourthly, members of a class that are subjects of an attribution in a third class. For example, if we are constructing a subclass 'poor drivers' we can take them from the class 'drivers' but only selecting those that are mentioned in an attribute 'driver name' in a third class 'accident'.

The second type of interclass connection is a grouping connection where a grouping class can be constructed. In a grouping class the members are in fact classes themselves. There are three types of grouping class. Firstly, an attribute can be used to select the classes that are members of the grouping class. For example, vehicle types may be class that holds all the vehicle classes possible from an underlying class. Moreover, there is an attribute that is used to separate members of those classes. There can be overlapping classes if the attribute used is multi-valued. Secondly, the grouping classes members can be enumerated where they are well known ahead of time but there is no easy selection attribute available. Thirdly, the user can completely control the members of the grouping class. Although permitted, it is advisable that only one interclass connection be used for any nonbase class.

Attributes in SDM are associated with classes: '... each class has an associated collection of attributes.' (Shipman, 1981) Each attribute has a unique name. Further, they can have a value which is either an entity in the model or a collection of entities. The value of an attribute is selected from its underlying value class. This contains the permissible values for the attribute concerned. Attributes can be either multi-valued or single-valued. Single-valued attributes acquire it's value from one member of its value class. A multi-valued attribute may take its value from members of it's value class. Further, attributes can be specified to be *mandatory* (or not), or to be *not changeable.* An attribute can either apply to members of the class (a member attribute) or to the class as a whole (a class attribute). Moreover, a member attribute can be required to be exhaustive of its class. In other words, that every member of the value class of attribute must be the value of some entity. A multi-valued attribute may be specified as nonoverlapping. Meaning that the value of the attribute for two different entities cannot have any entities in common. An attribute can be defined in terms of other information in the model through a number of simple and complex method. The specific mechanisms are member attribute derivation that follows a recognised method for evaluation and application; Class attribute interrelationships; Attribute predicates for subclass definitions; Attribute inheritance.

SDM has the advantage of allowing a high degree of freedom with construction of classes together with the flexibility of constructing a wide variety of classes from base classes. 'SDM is unique in that it provides a rich set of primitives for specifying derived attributes and subtypes.' (Hull and King, 1987) It suffers from being somewhat difficult to manage with particularly complex attribute derivation rules. Further, it is possible that a modelling example would not fit the class rigidity that is presented, since it is the formation of the base classes that determines many of the features of derived classes. Flexibility, that is required as models that are constructed change, is not high with SDM. Further, its complexity makes for difficulty in model construction due to the wide choice of constructs, and it has been remarked that '... in a model with many constructs such as SDM, the designer is continually forced to choose from among a variety of ways of representing the same data.' (Hull and King, 1987)

Table 9—Concepts for the Semantic Data Model

Concept	Part	Description
Entity	Core	SDM has entities. Entities can be objects, events, categorisations (known as higher-order entities), or names.
	Class	Entities are collected into classes of "meaningful collections of entities." These classes are named. Two types of classes are established. Base classes are fully independent of any other class. Non-base classes are derived from other classes. The data is only extant in base classes.
	Identity	Base classes use identifiers that are groups of attributes (see below) that uniquely identify member entities. There may be several different unique identifiers specified.
	Interclass Connection	These connect non-base classes to base classes. There are two types of interclass connections. Firstly, a subclass connection is where entities of the subclass are a subset of the entities in the superclass. A predicate is used to restrict membership of the subclass. Secondly, a grouping connection establishes a class whose members are classes themselves.
	Subclass	These are selected in three ways. Firstly, using a predicate on the member attributes of the superclass. Secondly, 'where specified' to create an *ad hoc* collection. Thirdly, from a third class where membership is specified using an attribute's value for members of the third class.
	Grouping Class	This establishes a class whose members are classes themselves. Three ways of selecting members is possible. Firstly, by attribute where an attribute establishes the member classes. Secondly, enumeration. Thirdly, user control.
Attribute	Core	Attributes are specified over classes of entities. That is, they are tightly bound to these classes. Further, they may apply to members of the class (member attribute) or be for the class as a whole (class attribute)
	Values	Values of entities are taken from a value class according to rules. Attributes can be single valued (one value from the value class) or multi-valued (several values from the value class)
	Restrictions	Attributes may be restricted according to recognised rules. Specifically, these concern what values may be selected for an attribute, how the values are calculated, etc.

In closing, it is worth noting the similarities in approach between ER, FDM, and SDM. 'SDM, like FDM and unlike the ER Model, is centred around attributes, but it is richer (and thus more complex) than either FDM or the ER model. (Hull and King, 1987)

5·4 NIAM

NIAM (Nijssen and Halpin, 1989) is a fact oriented design method. A subset of the method concerns data modelling. We concentrate on the elements which could be described as being responsible for data modelling.

A NIAM data model consists of facts, entities, constraints, and derivation rules. Each of these have several facets. For example, a fact encompasses several aspects of the model. It could be that a particular object is involved in a relationship. Further, it could be a property that describes an object. There are also a myriad of constraints that can be specified for a schema.

Facts describe elements of a data model. 'Our first step is to *begin with familiar examples* of relevant information and *express these in terms of elementary facts*' (Nijssen and Halpin, 1989). Firstly, they can record properties that an object (entity) possesses. Secondly, they could register the participation in a relationship by an entity. Essentially facts are assertions made about elements of the model. 'Basically, *an elementary fact asserts that a particular object has a property, or that one or more particular object participate together in a relationship.*' (Nijssen and Halpin, 1989) Facts can be derived by being based on other facts in the model. For example, in an economic model, gross domestic product (GDP) for a nation would be derived in a complex way from a number of different facts in the model.

'An entity may be a tangible object (e.g. the person Felix) or an abstract object (e.g. the subject CS112).' (Nijssen and Halpin, 1989) Entity types define groups of like entities. 'Each entity is an instance of a particular **entity type**' (Nijssen and Halpin, 1989) One begins determining the types of entities by describing examples using facts and other information. Entities can be thought of as members of entity class determined by the entity type. Each entity is identified using a label. A label is like a unique combination of properties that name it. A label is something that is attached to uniquely identify something. '... a **label** is used to denote a particular object' (Nijssen and Halpin, 1989) However, the distinction between a label and another form of identity is somewhat forced. Suffice it to say, that entities can be uniquely identified in NIAM. Sometimes labels are important enough to be modelled as separate entities of sorts.

In most cases in NIAM, properties appear to be able to be loosely coupled with entities. That is they can be considered to be conceptually separate from the entities they describe. Presumably, this means that properties have an existence separate from entities. Practically, and in implementation, it is probably unlikely to be the case. It is, however, somewhat unclear as to whether properties in NIAM are truly loosely coupled with respect to entities. We are particularly concerned whether the same property definition could be used by several entities simultaneously. We assume that this is possible in modelling and therefore that attributes are truly loosely coupled with respect to entities.

Table 10—Concepts for NIAM

Concept	Part	Description
Fact	Core	Not essentially a part of the modelling framework, but underpins the philosophy of the framework in key ways. Most notable is the basis for facts in predicates. From these information about other concepts are gathered for specific modelling cases.
Entity	Core	NIAM allows for significant entities, or objects (either physical or conceptual) to be modelled. These are grouped into entity classes. Each entity cannot depend upon other entities to be classed as an entity.
	Identity	Each member of an entity class must have a unique identity called a label. It may be several properties (see below).
	Classification	Essentially a fact gathering technique for determining the sorts of allowable properties, nevertheless identification of classes based on the attributes exhibited or by selection is important in the *process* of modelling in NIAM
	Constraints	Constraints on entities and their classes can be expressed in NIAM. Essentially, where alternative entity types are permitted, and exclusive selection is made based on a group of choices. Further, constraints on allowable entities in a class can be defined. Additionally, limits can be placed on the numbers of entities allowed in a specific class. Constraints on the behaviour or selection of properties for an entity within a class can be made.
Relationship	Core	Relationships are defined over entity classes. Members of specified classes may *participate* in the relationship. Each relationship is bidirectional.
	Constraints	Each relationship may have a number of constraints placed upon it. Specifically these concern the minimum and maximum numbers of entites allowed to participate in the relationship. These constraints may extend to exact numbers of each allowed to participate.
Property	Core	Properties help describe entities. In NIAM, they are loosely bound to entity classes.
	Complexity	Properties may be calculated based on other properties. Further, it is not unforseen that properties could be complex in nature. NIAM does not explore such possibilities since it is as much a method for business data modelling as it is a modelling framework.

Entities can be divided into subclasses (or subtypes). Interestingly, '[i]n NIAM, subtypes are introduced to express constraints on what is recorded in the database, rather than to provide a complete picture of natural classification in the real world.' (Nijssen and Halpin, 1989). Nevertheless, the subclass structures it has is based on set theoretic basics. Essentially, it differentiates between those properties that are only for certain subclasses of entity and the rest. Further, it constrains descriptions for entities where appropriate. Essentially, the class structure exists as subtypes after description. In other words, it is an optimisation technique rather than a modelling technique. It appears that the intent of NIAM through this discussion is to allow classes to emerge from entities.

Relationships are between entities. The entities involved are said to participate in the relationship. Each entity involved in such a way is said to play a role in the relationship. 'Basically, a role is a part played by an entity in some relationship' (Nijssen and Halpin, 1989). Further, relationships in NIAM are bidirectional by default.

Additional to the static elements described here, there are a number of constraints that are designed to control and maintain the veracity of information modelled and to convey the meaning of specific models created using the data modelling framework. There are uniqueness constraints for participation in relationships; Entity-type constraints limiting the values allowable for particular entities; Mandatory/optional roles for properties; Occurrence frequencies where strict limits on the numbers of entities that are allowed; Equality constraints where for one property of an entity is dependent upon the existence of another property for a specific entity (for example, if a width is specified for a building, then one would expect height also); Exclusion constraints show where an entity can play only one role of several indicated; Subset constraints in which a property may be optional for a certain type of entity (also for sets/classes); Homogeneous binaries constraints where entities of the same type are participating in a relationship.

5·5 OMT's Object Model

OMT has an object model as part of its methodology (Blaha and Premerlani, 1998, Rumbaugh et al., 1991), that is essentially the same as UML. We refer to the object model as being part of OMT. Further, OMT's dynamic and functional models are used to model dynamics in addition to the object model. We use these models in the thesis. Further, we describe OMT's object model here, but we delay introduction of OMT's dynamic model and functional model until we use them in Chapter 8, since we only require the dynamics of OMT for the evaluative method.

OMT's object model is centred on the idea of *classes* of *objects*. A *class* is a description that defines all objects in the class. All objects in a class have similar attributes (properties), behaviour, relationships, and semantics. Further, it can be said that objects are *instances* of the classes to which they are member. Each object in a class has an identity, and it is 'that property of an object which distinguishes each object from all others.' (Blaha and Premerlani, 1998) For example, the class *Aeroplane* has attributes such as *wing-span, number-of-engines*, etc. An instance of this class may be *Boeing 737-400* with a specific wing-span and with two engines with a particular range of power. The object Boeing 737-400 has its aircraft type through which it is identified. Further, an attribute 'is a named property of a class that describes a value held by each object of the class.' (Blaha and Premerlani, 1998) In OMT the object model is dominated by structural constructs. Attributes assume lesser importance with respect to classes and relationships. An instance of an attribute is called a value.

"An operation is a function or procedure that is applied by or to members of a specific class." (Blaha and Premerlani, 1998) They can also be defined to apply to an entire class of objects (this is a much rarer occurrence). They are often used to calculate values based on one or more objects. Functions can be viewed in a similar way to attributes. Operations can also have side-effects, by changing the state of an object. Further, operations play a role in the dynamic model (see Chapter 8, where the dynamics of OMT are described and used.) A method is a specific implementation of an operation for a class. In modelling, methods assume least importance since it is the conceptualisation of the operation that is critical, not its implementation. Operations apply for all objects in the specific class. Operations can be inherited from classes higher in the hierarchy. Further, inherited operations can be overridden, or more precisely redefined, to provide different localised behaviour for objects lower down the class hierarchy.

Objects can be related, and in OMT, these are through *links* and *associations*. A relationship is a physical or conceptual connection between objects, and must relate at least two objects. An association is a *description* of a group of links with common structure and semantics. A link is an *instance* of an association. For example, aircraft types may be *owned by* specific airlines. This is an example of an association. A specific link may be between our Boeing 737-400 aircraft type which may be owned by a specific airline such as *Air France*. Further, there is a constraint called *multiplicity* that specifies the maximum number of instances of a class that can participate in an association with a single instance from another class. Multiplicity can be indicated using a zero, exactly one, many (zero or more). Alternatively specific numbers can be indicated (such as 2 to 4.) If more than one association is between the same classes, then roles are specified to avoid confusion. Association classes are associations that can participate in subsequent relationships, and are like associations that are of a higher order than binary associations.

Attributes may be specified for associations. For example, the association between aircraft type and airline may have an attribute that indicates the number of that type operated by the company. Specifically, the link between *Air France* and the *Boeing 737-400* may have a value of 30, thus showing the number of that type of aircraft owned by the company.

The ideas *Generalisation* and *Specialisation* allows for subclasses of classes to be created. Subclasses more specifically define the objects that are members. For example, the class of aircraft type may have as a subclass the categories or specialisations of jet-aircraft, and propeller-aircraft. The subclass represents a specialisation of the superclass. Alternatively, the class is the *generalisation* of its subclasses. Critical to this sort of class hierarchy is simple or single inheritance. Specifically, it is the relationship between classes of objects where attributes, operations, and associations are inherited by subclasses directly from a superclass. Inherited properties may be overridden by subclasses. More complex is multiple inheritance where properties from several superclasses are inherited by a subclass through multiple inheritance.

A strong form of association is that of aggregation. Specifically, aggregation is where an object is made up of a number of parts. Together it may be said that the parts aggregate to form the whole. A popular example of this is a car. It consists of a body, steering wheel, transmission, windows, seats, wheels, etc. Together these aggregate to form the car. Aggregation differs from generalisation (see above) in that we are not simply talking about the specialisation of a more general class. Instead we are discussing the structure of an object. That is, its specific parts.

Table 11—Concepts for the Statics of OMT (OMT's Object Model)

Concept	Part	Description
Object	Core	UMLs OM supports objects. These are grouped into classes for which common attributes, behaviour, relationships, and semantics can be defined.
	Attribute	Attributes are specified over object classes. Specifically, they are tightly bound to the classes. They are conceptualised for a class and then operationalised in an instance of the class (object).
	Identity	Each member of a class must has a unique identity within that class. The identity is a subset of the attributes defined for an object.
	Operation	An operation may be applied to or by objects in a class. They are functions or procedures. These apply for all objects in a class. They can play the role of a derived value for an object. Alternatively, they can specify changes in state for an object (see the dynamic model later in the thesis.) Note: a *method* is an implementation of an operation for a specific class.
Link-Association	Core	Links between objects are supported in UMLs OM. The concept embodied in links are instantiated through associations. Attributes can be specified for links.
	Multiplicity	In a similar way to other modelling frameworks, UMLs OM require the specification of multiplicities for links. These indicate the minimum and maximum number of objects that can be linked in the way indicated.
	Aggregation	A specialised form of link or association. It is from instance of object to instance of object showing part-whole relationships. That is, where an object is made up of a number of different objects. Each part-whole relationship is modelled using a different link and association.
Class Hierarchy	Core	An elaborate class hierarchy is established allowing for specialisation /generalisation. Properties, operations, and associationss may be inherited from superclasses. Further, they may be overridden according to need.

The key advantage with the object model for OMT is the breadth of concepts available for modelling. Further, it has more tightly defined semantics when compared with other modelling frameworks. Despite this, object models, suffer from having a rigid class hierarchy with attributes of objects within these classes being spread through the hierarchy. OMT's object model is no exception. Flexibility with this is very limited. For example, should the class hierarchy for plant taxonomy popular presently be replaced by a more newly (and perhaps more complex) accepted hierarchy, OMT's object model and other object models would require significant reorganisation of its schema. This rigidity is reminiscent of that found in ER modelling frameworks. In addition to this rigidity, there is often the claim that object orientation is somehow 'natural' for modelling reality. This claim is difficult to refute. Anecdotally, there is evidence to suggest that modellers do not find them particularly natural.

Chapter 6

An Ontological Comparison of Data Modelling Frameworks with Chisholm's Ontology

We have reason to believe that data modelling frameworks can be compared with an ontology. We have described a method that uses concepts and terms as a basis for comparison. We have selected, based on pragmatics, an ontology by Roderick Chisholm, and we have applied the method using that ontology with five representative data modelling frameworks—the entity-relationship model, the semantic data model, the functional data model, NIAM, and OMT's object model that is compatible with UML. In this chapter, we present and discuss the results of an ontological comparison of five data modelling frameworks with Chisholm's ontology according to the method described in Chapter 3. In Section 6·1 we present the results for the comparison in the manner shown in Chapter 3. We then discuss and analyse our results in Section 6·2 before reflecting on our experiences in applying the method.

We have found that the world view of the ontology resonates to a reasonable degree with the world views of the selected data modelling frameworks. By conducting the comparison we have gained a deeper understanding of ontologies like Chisholm's and of the method. We did, however, encounter the following issues that may affect future applications of the method. Firstly, the process was intensive and required a great deal of familiarity with the ontology and the data modelling frameworks. Secondly, there is scope for misjudgement due to the qualitative nature of the methods. However, we have taken care to ensure that it is likely that only a minor variation on what is presented will be found.

6·1 Results

According to the method we need to extract concepts from the ontology in order to conduct a comparison of a number of data modelling frameworks. This was done in Chapter 4 at the time of describing the ontology.

Since most data modelling frameworks do not talk about dynamics, in order to compare them we only need to use the static concepts from Chisholm's ontology. Chisholm's ontology establishes a number of static concepts that are summarised in Table 5 on page 62.

In the method, each data modelling framework is compared with concepts from the ontology. The individual comparisons are then summarised. We have conducted a comparison using five representative data modelling frameworks and according to the method that we described in Chapter 3. The summary of the results of this comparison are shown in Table 12. We call this table the indicative results of the comparison. For a full description of the symbols used in the table and for an explanation of their relationship to semiotic theory, please refer to the comparative method in Chapter 3.

Table 12—Results of the Comparison of Selected Data Modelling Frameworks Using Chisholm's Ontology

Ontological Concept	ER	FDM	SDM	NIAM	OMT/ UML
Individual	$\surd_p$	$\surd_p$	$\surd$	$\surd_p$	$\surd$
Attribute	$\surd_p$	$\surd_p$	$\surd_p$	$\surd_p$	$\surd_p$
Classification	$\surd_p$	$\surd$	$\surd_p$	$\surd_p$	$\surd_p$
Relation	$\surd_p$	$\surd$	$\surd$	$\surd_p$	$\surd_p$

In the following section we discuss these results in detail and describe the relationship between the world views of the data modelling frameworks and the ontology.

6·2 Discussion of Results

The indicative results in Table 12 show a good degree of coverage for a number of concepts from Chisholm's ontology by all the data modelling frameworks. However, each one of the concepts from the ontology fail to be fully covered by at least one of the data modelling frameworks. In this section we examine each concept and discuss the specific differences between Chisholm's ontology and each of the data modelling frameworks.

In the discussion we examine the indicative results to determine more clearly and in what ways specific concepts are covered in the data modelling frameworks and why the coverage is to the extent that it is indicated. It also allows us to more deeply discuss the relationship between concepts from the ontology and the data modelling frameworks. We proceed by discussing each concept in turn in sub-sections. Where concepts are compound we elaborate upon the indicative results, shown in Table 12, and expand the results according to part. This affects /individual/, and /attribute/. We conclude the section with a summary.

Individual

The semantic field defined by «individual» is generally well covered by all of the data modelling frameworks. That is the data modelling frameworks, through its own distinct set of concepts, covers a semantic field similar to that covered by /individual/ from the ontology. However, /individual/ is not supported in fullest generality by most of the frameworks, with only OMT and SDM showing full coverage.

/Individual/ is a compound concept consisting of three parts. Firstly, /individual core/. Secondly, /individual identity/. Thirdly, /individual structure/. Table 13 shows these parts and their degrees of coverage found in each data modelling framework.

Table 13—Expanded Results for /Individual/

Part for /Individual/	ER	FDM	SDM	NIAM	OMT/ UML
Core	√	√	√	√	√
Identity	√	√	√	√	√
Structure	X	$\surd_p$	√	X	√

Inspecting the table, we see that all data modelling frameworks examined have concepts that totally overlap with the semantic field for /individual core/. «Individual core» is that an individual that can be identified and is not ontologically dependent upon other individuals can be represented. In this case «individual core» is fully supported by all of the modelling frameworks examined, with each of them supporting the concept in similar ways. In ER, the meaning is preserved in /entity core/. FDM provides an equivalence through its concepts /entity core/ and /entity object-type/. /entity core/ in SDM, supports /individual core/ and maintains the meaning of the concept. /Entity core/ in NIAM as well as /object core/ in OMT also provide an equivalent coverage of the concept.

In a similar way, all frameworks support the semantic field defined by «individual identity» and can uniquely identify individuals using attributes, but vary in the mechanisms by which identity is achieved. ER, SDM, and NIAM all support this concept through their own concepts called /entity identity/. OMT has a similar concept called /object identity/. Further, each of these data modelling frameworks use attributes to provide for the meaning of «individual identity». FDM is structurally dissimilar compared with the other modelling frameworks but nevertheless allows for the unique identification of object-type entities using attributes.

In contrast, only two of the data modelling frameworks have concepts covering the full meaning «individual structure». This means, for those data modelling frameworks that do not have coverage of «individual structure», individuals, described by the relevant term, can not easily be subdivided into parts as stipulated by «individual structure» in Chisholm's ontology.

The only two data modelling frameworks that give full coverage to this concept are OMT and SDM through their respective concepts /link-association aggregation/ and /attribute/. In OMT, an object can linked with others each with their own structure and state. Similarly SDM can use attributes to link with entities in other classes thereby giving part-whole relationships of the type desired by the ontology. FDM has a form of aggregation, /function aggregation mapping/ that goes part of the way to supporting the concept. Essentially, it allows several object-type entities to aggregate and share an attribute. However, the intention behind /individual structure/ is for a full part-whole relationship with a further intention to utilise this in dynamic concepts that is lacking in FDM. It is for this reason that we render the support partial. The likelihood with which FDM is to be extended is still reasonably good since the concept used to provide this inferior type of aggregation does not run counter to the part-whole form of aggregation supported in the ontology.

The remaining data modelling frameworks do not cover the semantic area designated for /individual structure/ because of their reliance on a form of tight classification within which each individual equivalent resides. Encouragingly, neither do they have concepts that run counter to the concept, and consequently, coverage of this semantic field by amending or extending the data modelling frameworks should be relatively straightforward. Presently, ER and NIAM can be used in a crude way to represent part-whole relationships, but these were not intended by the original modelling framework and so are not considered to adequately support the semantics.

/individual/ is to a large extent supported by all of the modelling frameworks studied. We can therefore conclude that insofar as this concept is concerned, they show a high degree of overlap with the ontology we have used in this study.

Attribute

A quick inspection of the results table for the second concept, /attribute/, shows quite a different story from that for /individual/ with all modelling frameworks failing to provide full coverage for the concept. The major issue here is that Chisholm's ontology contains an important emphasis with respect to /attribute/. Like many realistic ontologies, attributes or properties are universals and so are enduring. By implication, conceptually speaking, attributes must be loosely coupled with individuals, and attributes must not be tightly bound to entities. Only two data modelling frameworks have concepts that cover this aspect of Chisholm's ontology in its full generality. Further, all of the other data modelling frameworks would have difficulty in reflecting Chisholm's ontology because they have concepts that are antonymous in a contradictory way with this aspect of the ontology in that loose coupling is not observed.

Recall that /attribute/ is compound and has three parts. Firstly, /attribute core/. Secondly, /attribute equivalence/. Thirdly, /attribute complexity/. We examine each part separately. Table 14 shows these parts of the concept and their degrees of coverage found in each data modelling framework.

Table 14—Expanded Results for the Concept /Attribute/

Part for /Attribute/	ER	FDM	SDM	NIAM	OMT/ UML
Core	$\surd_p$	$\surd$	$\surd_p$	$\surd$	$\surd_p$
Equivalence	X	X	X	X	X
Complexity	$\surd_p$	$\surd$	$\surd$	$\surd_p$	$\surd$

Firstly, /attribute core/ concerns the necessity that individuals exemplify attributes. All data modelling frameworks clearly support that aspect of «attribute core». However, the nature of /attribute core/ concerning the enduring nature of attributes renders the coverage of the semantic space for /attribute core/ not to the full generality for three of the data modelling frameworks. All of them support the intention behind the concept that allows for an individual to exemplify attributes. In ER, and SDM, /attribute core/ provides the exemplification of attributes by entities. FDM utilises two concepts, /entity attribute-type/ and /function attribute mapping/ to provide exemplification. SDM has /property core/ and OMT has /object attribute/ to support the concept /attribute core/. NIAM utilises /property core/.

Additional to exemplification, another aspect to the core concept is that attributes are enduring. In philosophical terms this means that attributes are universals (Honderich, 1995), and consequently, it makes no sense to discuss the beginning or ending of attributes. If we take this interpretation of enduring attributes then the Functional Data Model supports attribute endurance. NIAM allows this through its emphasis on facts as a guide for describing entities and for determining properties. None of the other data modelling frameworks adhere to this requirement. A further dimension to this aspect of «attribute core», and as a direct consequence of attribute endurance, is loose coupling of attributes to individuals. FDM and NIAM exhibit this trait. However, all of the other data modelling frameworks fail with respect to this. It is worth noting that some object models also adhere to this specification; OMT is not one of these object models. FDM with attribute-type entities clearly separate from object-type entities embodies the requirements. NIAM similarly emphasises the separate nature of entities and attributes or properties. Most of the other data modelling frameworks emphasise the link between object or entity and the attributes that are used to describe the object, and the modelling frameworks emphasise that attributes are common to all members of a particular class of objects or entities.

As a philosophical stance taken in Chisholm's ontology, it seems reasonable that the attributes themselves are enduring. This is an entirely different issue than that of modellers being able to construct a data model with or without such a requirement of endurance. Further, if we do take the view that we need to think of attributes as enduring in data modelling, we need not extend this idea to implementation.

Secondly, we have /attribute equivalence/, where attributes with different names may mean the same thing. Chisholm calls this 'conceptual entailment'. None of the data modelling frameworks have concepts that overlap with «attribute equivalence». We flag this for future discussion in the following section, since it may be that such a requirement is only of philosophical value. That is, it is used to avoid difficult questions regarding the equivalence of attributes. However, pragmatically it seems such a requirement may be of use, particularly when querying a database, where knowledge of the conceptual equivalence of attributes may help the formulation or execution of queries.

The third part of the concept is /attribute complexity/. That is, compound attributes of some nature must be provided by the frameworks, that may be constructed from either simple attributes or from other compound attributes. Chisholm is not prescriptive in the nature of the mechanism for supporting this concept, but notes that conjunction and disjunction of attributes would be a way through which compound attributes could be supported. Further, that conjunction and disjunction of this type would be of either simple attributes or compound attributes. FDM, SDM, NIAM, and OMT's object model covers this semantic space regarding such complexity to varying degrees of generality, and, although not explicit in the original work, ER allows a type of complexity in that compound attributes are permitted, but these are simple part relationships and cannot be nested indefinitely, or independently of the entity class in which it is defined, and thereby do not go very far to provide the power encompassed by «attribute complexity» and the support while partial is very weak indeed. NIAM provides for fact types that are compound in ways very similar to ER.

FDM and SDM have quite complex attribute construction permitted as part of /entity attribute-type/ and /attribute/ concepts respectively. In FDM, the type structure that is open for attributes allows for complexity in that one can construct complex attributes that have parts as intricate as Chisholm's model. In SDM it is the rich combination of attributes that provides the coverage of the concept. /Object operation/ in OMTprovides for a complexity, since that concept allows for calculations based on many and varied sources. It is therefore the combination of the two concepts /object attribute/ and /object operation/ that support is provided in OMT.

Summarising, we see that the core of this concept is supported to a large degree by all modelling frameworks. There are some issues requiring further discussion concerning conceptual entailment and the enduring nature of attributes. Conceptual entailment and universality of attributes is something that has not been pursued in modelling. However, it is a central plank in Chisholm's ontology and is also common in the realistic philosophy that we have examined. It is also clear that attribute structure needs to be examined in more depth for the data modelling frameworks studied. Despite this, there is considerable support for /attribute/.

Classification

/Classification/ in the ontology is determined through attributes. Chisholm argues that in order that classes and sets can be represented there is no requirement for extra things in reality, other than individuals and attributes, and further, that classes themselves are not ontologically fundamental. He also demonstrates that membership of classes and sets by individuals must arise from the attributes that they exhibit. In other words, classes are found in reality but they are defined in terms of the attributes that are exhibited for individuals to be members of the each class.

A logical extension of the requirement that sets and classes be established from attributes is that this provides us with a high degree of flexibility. In other words, it frees us from the rigidity of establishing classes from the outset when constructing a model using a specific data modelling framework. Further, individuals can be members of several classes simultaneously, and can move freely between classes.

FDM has a flexible approach to the concept of classification because /function classification mapping/ uses attributes to select classes in a way that resonates well with /classification/ in the ontology. This in turn results from FDM having attributes loosely coupled with entities.

In contrast, the rigidity in ER through /entity core/, and OMT's object model through /object core/, where each concept embodies what is essentially a premeditated attitude to classes, renders them largely incompatible with the concept but to varying degrees. For OMT's object model, we mean 'premeditated' from the viewpoint that classes are defined early with rich class hierarchies and that objects can only exist at one place in the hierarchy. In the case of ER, entities are hide-bound by their original entity classes. Flexibility in this context is very problematic, requiring a wholesale redrafting of classifications that are long established. The nature of the two concepts is such that they cover a semantic area that is contradictory or contrary to /classification/ from the ontology and so extension of both of these modelling frameworks would be difficult.

Interestingly, SDM, and NIAM come reasonably close to covering the notion embodied by «classification» in Chisholm's ontology. The process by which NIAM determines classes (as described by «entity classification») reflects to a reasonable degree, the intention contained in the ontology. However, resulting from the process are entities of the nature described through «entity core», and consequently, NIAM does not support the full generality of Chisholm's ontology because its support for this concept is in the way it allows facts about entities to be gathered and only in part carries this through to models constructed. The result is a tight type structure. Amid the complexity with which classes are provided for in SDM we can find the type of class/set environment that Chisholm advocates, in that SDM partly uses attributes for establishing some classes, and that base classes contain the data in SDM. However, the nature of /entity subclass/ in SDM is such that we cannot say that SDM fully supports /classification/ from the ontology. This is because in /entity subclass/ membership can be *ad hoc* when using the 'where specified' option. This approach is not supported by /classification/ from the ontology.

The modelling frameworks that have concepts dealing with classification and that have concepts that are contradictory with respect to the ontology are likely to be very difficult to extend in order to cover the semantic space defined by /classification/. We can say that classification is important in that Chisholm requires the term in order to describe 'what there is' and that classification is part of its realistic heritage. All modelling frameworks support classification. For this reason we can say that they all have some overlap with the ontology. However, there is no support for individuals being forced to exist in classes. Instead, attributes exhibited by individuals are used to select members for classes. Individuals may be a member of several classes simultaneously. The issue of classification is an area for further research in that there are many interesting issues that need investigation with specific data modelling frameworks.

Relation

The concept /relation/ in Chisholm's ontology views relations as attributes and that they must be unidirectional, consequently, in order for a modelling framework to cover the semantic space defined by /relation/, concepts must be provided so that relations can be directed from one individual to another. FDM and SDM satisfy this requirement through /function relationship mapping/, and /attribute/ respectively. OMT, NIAM, and ER all have similar concepts, /link-association core/ in the case of OMT, and /relationship core/ in the cases of NIAM and ER, but they cover different semantic fields in at least one critical aspect, in that they fail to provide for the unidirectional nature of relationships. Nevertheless all modelling frameworks recognise the importance of modelling relationships between entities.

The matter, however, is complicated in the case of OMT by the fact that, when describing the dynamics of OMT, the state of an object includes the links it possesses, which in turn implies that the instantiation of relationships is, in fact, unidirectional. We rate the coverage of the semantic field defined by «relation» from the ontology, in OMT to be similar to, but not to the full generality of, Chisholm's ontology because we cannot be certain that the intention is for unidirectional relations. In some implementations of a model designed using OMT a unidirectional link may eventuate whereas in other implementations of the same model may result in a bidirectional link. OMT could be amended reasonably easily to provide such a functionality because the concept in OMT is not antonymous in a way that is contrary or contradictory with respect to the ontology.

In the NIAM and ER modelling frameworks, some existing concepts dealing with this are clearly antonymous in a way that is contradictory with respect to /relation/ and therefore it is difficult to see how relations of a unidirectional nature can be provided in these modelling frameworks. The relationship between concepts in NIAM and ER are contradictory with respect to the unidirectional nature of the ontology (with respect to /relation/); Bidirectional relations are clearly implied.

SDM provides for relationships between entities in /attribute/ which are unidirectional, and all attributes in SDM have the capability of relating two or more entities. It is for this reason that we judge that SDM covers the same semantic area as the ontology.

Chisholm's approach to relations raises interesting questions regarding the criticality of modelling relations as attributes. It is likely that a significant part of this issue is philosophical and to a lesser extent concerns quality in modelling. Alternatively, it may simply be symptomatic of Chisholm's desire to include, in his ontology, only categories that are necessary, and then to relate terms such as relation to the categories whenever possible. This aspect of the ontology is an area that we intend to investigate in future work.

Relations are common in realistic ontologies. Unidirectional relations between individuals are critical for Chisholm's ontology. All modelling frameworks provide for relations. Few of them allow for unidirectional relations.

6·3 Reflections

We begin by summarising Chisholm's ontology as it relates to modelling and comment on the levels of support we have found for that world view in the modelling frameworks. Following this we discuss some specific issues that arose through the comparison. Finally, we draw some specific conclusions about the modelling frameworks and Chisholm's ontology.

Chisholm's ontology views the world as a collection of individuals and relations between them, and the ontology uses attributes to describe both individuals and relations. Attributes are universals and endure, and, consequently they are loosely coupled with individuals. Attributes are also used to determine class and set membership. Our comparison suggests that this is to a large extent a similar world-view as those imparted by the data models and there is a significant level of agreement with the ontology and the modelling frameworks that we've studied, but the data models lack the full generality of Chisholm's ontology. The major departures are in the more subtle parts of the nature of relations and the implications of a lacking of loose coupling between individuals and attributes, particularly implications concerning classification.

Classification in the ontology is evident through the attributes exemplified by members of classes. In the ontology, classes are related to each other by the intersections and unions of the attributes used to select them and thereby can simulate class hierarchies. This approach is entirely different from the most common classification approaches used by most data modelling frameworks where instead, rich and rigid class hierarchies are prevalent.

The consequence of these departures from the ontology is that it is likely one can model a narrower range of situations using the studied data modelling frameworks than Chisholm's ontology, although this requires further investigation. Further, Chisholm's ontology has the potential to change our view of data modelling by its increased flexibility achieved through bidirectional relations and through its loose-coupling of attributes with respect to individuals. In turn, this has positive implications for the flexibility of models which are subject to radical or ongoing change. It is the formation of classes through attributes as a direct consequence of loose coupling that is of most beneficial for flexibility.

We have found that ER, OMT's object model, and NIAM do not support such class flexibility. This is principally because of tight coupling between individuals and attributes found in ER, OMT's object model, and by practice in NIAM.

We found that FDM captures the fundamental nature of Chisholm's ontology more closely than the other modelling frameworks and, due to its evident simplicity, has more potential to be able to support other elements presently not supported that are directly related to loose coupling of attributes and individuals and to classification.

Its simplicity means that there are few, if any, concepts in FDM that are antonymous with respect to concepts from the ontology in either a contradictory or contrary manner.

SDM's complexity with respect to its class system makes it a difficult modelling framework to use to fully express Chisholm's ontology. Nevertheless, it would be interesting to investigate SDM further.

We have reason to highlight two related aspects of Chisholm's ontology for future study directly as a result of conducting this comparison, the conceptual entailment and universality of attributes. The support for conceptual entailment was entirely missing from all modelling frameworks. Conceptual entailment is a concept that may have much to offer the pragmatics of modelling and perhaps of implementation. Conceptual entailment results directly from treating attributes as universals because in making them universals they assume a much more important role in modelling 'what there is'. Universality of attributes is widely accepted in realistic philosophy. We can see benefits in modelling flowing from the ability to make explicit conceptual relationships between various attributes that are, in turn, evident in reality. We can also see implementation benefits by being able to query based upon conceptual entailment relationships between attributes. However, this in itself is a complex area for future research.

Concluding, we can see from the results that the modelling frameworks share, to a large degree, the world view of the ontology. The areas of departure tend to be of the nature of a difference in emphasis rather than complete absence of support. Also, all concepts had a high degree of emphasis with respect to their core. There are, however, some issues that need investigation. The importance of conceptual entailment and the implications of the universality of attributes need to be investigated in practical situations, and there are possible efficiency gains by examining this. Apart from the universality of attributes, the area of most concern is that of classification. Clearly, the rigidity of class construction and the presence of rigid class hierarchies is not supported in the ontology. As implementation efficiencies these may be acceptable. As modelling features there appears to be little support in traditional realistic philosophy for such an approach.

CHAPTER 7

Results of Evaluating the Functional Data Model using Chisholm's Ontology

IN THIS CHAPTER, we present the results of evaluating the functional data model (FDM) using Chisholm's ontology. The evaluative method is in two parts. Firstly, there is a bidirectional comparison between the ontology and FDM. Following the application of the bidirectional comparison, we consolidate the results, and then use the consolidated results to suggest which parts of FDM need to be extended to bring FDM closer to the ontology in terms of the coverage of concepts from the selected ontology by FDM, thereby bringing the world view of the data modelling framework closer to that of the ontology. We also comment on the likely ease with which extension is possible.

We have found that regarding the static terms of Chisholm's ontology, there is good agreement with FDM. However, and not unexpectedly, FDM does not support the dynamics of Chisholm's ontology, and we found that extension or augmentation is required in certain specific areas. We expect that the ease with which extensions can be provided will depend upon the relationships between concepts in FDM and the ontology, and we comment on this aspect in this chapter. Future work will involve extending FDM in the ways outlined in this chapter.

We proceed by comparing Chisholm's ontology (Chisholm, 1996) and the functional data model (FDM) (Shipman, 1981) according to the method described in Chapter 3.

The method requires us to conduct the comparison in two parts, and reflecting this, we present our results of applying the method with FDM in two sections. Firstly, we compare FDM with concepts from the ontology. Following this we conduct the 'reverse' comparison and compare Chisholm's ontology with concepts from FDM. In each part of the bidirectional comparison, we pause to discuss, analyse and reflect on the results. We use the discussion to explain our finding and to help in the consolidation of the results later in the chapter. It also enables us to gain a deeper understanding of the similarities and differences in concepts between Chisholm's ontology and FDM. We present the consolidated results in a later section of this chapter.

Readers should be familiar with FDM and Chisholm's ontology. FDM is introduced in Chapter 5, whereas Chisholm's ontology is described in Chapter 4.

7·1 Results Comparing FDM with Chisholm's Ontology

We commence by comparing FDM with Chisholm's ontology. Recall that each part of the bidirectional comparison is similar to one of the pairwise comparisons from the comparative method.

In this section we present the first part of the comparison, and show the results of comparing the concepts found in FDM with concepts from Chisholm's ontology. We have already extracted the static and dynamic concepts for Chisholm's ontology, and these are shown in Table 5 on page 62, and Table 6 on page 69 respectively. An understanding of FDM is required for the comparison. FDM is summarised in Chapter 5.

We use the concepts from Chisholm's ontology one at a time and relate them with FDM in the ways described in Chapter 3. These indicative results are then summarised in a table. In addition to an indicative findings of similarity and difference, we discuss the substance of the differences and similarities in a number of subsections following. Each subsection deals with a specific ontological concept.

A summary of the results of this comparison is shown in Table 15. As we have done previously, this table shows indicative results which we then explain and discuss. The 3-level indicator that we use is described fully in Chapter 3.

Table 15—Indicative Results of Comparing the FDM with Chisholm's Ontology

Ontological Concept	FDM
Individual	$\surd_p$
Attribute	$\surd_p$
Classification	√
Relation	√
Temporal Relations	X
State	X

The indicative results in Table 15 show a degree of support for a number of concepts from Chisholm's ontology by FDM. We proceed to discuss, the concepts from the ontology and how they relate to FDM. In our discussion we rely heavily on the semantic field delineated by each concept in the ontology and how they relate to

FDM. The discussion also enables us to understand more deeply the reasons why a concept is supported to the extent that it is by FDM.

We expand upon the indicative results, shown in Table 15, where complexity in the concept necessitates us to examine the parts of the relevant concept. In Chisholm's ontology (including dynamics) there are three compound concepts. Namely, /individual/, /attribute/, and /state/. We conclude the section with a summary.

Individual

FDM has good semantic coverage for /individual/. However, /individual/ is not supported in its fullest generality. /Individual/ in Chisholm's ontology has three parts. Firstly, /individual core/. Secondly, /individual identity/. Thirdly, /individual structure/. Table 16 shows these parts and their levels of coverage found in FDM. The results are shown using the same key as used previously.

Table 16—Expanded Results for /Individual/

Part for /Individual/	FDM
Core	√
Identity	√
Structure	$\surd_p$

Inspecting the table, we see that FDM covers «individual core». In this context, «individual core» means that significant individuals that are not ontologically dependent on other individuals can be described. In FDM the concepts /entity core/ and /entity object-type/ provide for coverage of an identical semantic area.

Similarly, FDM covers «individual identity» by enabling the selection attributes for individuals that are possessed by only one individual.

However, FDM does not completely cover the semantic area defined by the third concept, «individual structure». It is true that FDM allows a form of aggregation, through /function aggregation mapping/, but it does this in a way that is separate from the term that concerns object-type entities. In so doing it fails to allow an individual that is the aggregation of several others to exist independently as an individual, but instead simply allows attributes to be declared for an aggregated individual. It is for this reason that we find that FDM does not provide for the full generality of /individual structure/ as stipulated in Chisholm's ontology.

Summarising, we can see that FDM supports /individual/ to a large degree. The implications and ramifications of partial support for /individual structure/ are discussed later in the chapter.

Attribute

Being a compound concept, /attribute/ in Chisholm's ontology has three parts. Firstly, /attribute core/. Secondly, /attribute equivalence/. Thirdly, /attribute complexity/. We examine each one in turn. Table 17 shows these concepts and their levels of coverage in FDM, and we can see that FDM supports two of the three parts of the compound concept.

Table 17—Expanded Results for /Attribute/

Part of /Attribute/	FDM
Core	√
Equivalence	X
Complexity	√

Firstly, /attribute core/ partly concerns the fact that individuals exemplify attributes, and FDM covers this aspect through its own concepts /entity attribute-type/ and /function attribute mapping/. Exemplification of attributes by individuals is achieved by mapping object-type entities to attribute-type entities using an attribute mapping function. However, in Chisholm's ontology, /attribute core/ extends beyond this to encompass two extra elements. Firstly, Chisholm argues that a core of the concept is that attributes are enduring. Secondly, as a direct consequence, that attributes and individuals are loosely coupled.

Attribute endurance concerns the fact that it makes no sense to discuss the beginning or ending of a specific attribute. For example, was it last Tuesday when the attribute of 'being green' was first described? Alternatively, was it three centuries ago? Clearly, it then follows that the attributes themselves are enduring even if no observer perceives them for a long period. FDM covers the semantic area concerning enduring attributes. It does this by predetermining 'types' for data, and allowing for these to be universally available. The concept of enduring attributes is essentially philosophical. Consequently, if we do take the view that we need to think of attributes as enduring, we need not necessarily carry this through to an implementation environment (eg. a database management system). However, there may be some positive outcomes. It is interesting to note that some object modelling frameworks and their implementation environments (Milton, 1993) go some way to supporting this aspect of the concept.

A related aspect to enduring attributes is the loose coupling of attributes to individuals. This is a direct corollary of attribute endurance because if attributes endure for periods infinitely longer than any individual, then the two cannot be tightly coupled, they must exist separately. Consequently, Chisholm's ontology requires that attributes be loosely coupled with entities (individuals). FDM covers an equivalent semantic area by supporting two types of entity, /entity object-type/ and /entity attribute-type/ and allowing mapping between the two. Therefore in FDM, attributes are separate from object-type entities.

In the ontology, /attribute equivalence/ is also referred to as conceptual entailment.. This is where different attributes exhibited by individuals may refer to the same idea. They may differ not in the meaning of the attribute, but instead in the name given to the attribute. FDM does not have concept(s) that cover the semantic area defined by this concept. It may be that such a requirement is only of philosophical value, in that it is used to avoid difficult questions regarding the equivalence of attributes that goes beyond the equivalence of name or label. However, pragmatically it seems such a requirement may indeed be of particular, and maybe practical, use. Particularly, querying of an implemented database adhering to a model may be enhanced by recording this type of equivalence. We reflect on such things at the conclusion of this thesis.

Thirdly, we have /attribute complexity/. In the ontology, compound attributes are constructed from either simple attributes or from other compound attributes. Chisholm is not prescriptive as to the nature of the mechanism for supporting this concept. However, in his description, he notes that conjunction and disjunction of attributes would be one such mechanism by which compound attributes may be supported. Further, that conjunctions and disjunctions of this type are of either simple attributes or compound attributes. FDM supports such complexity, by providing a complex type structure for attributes in the modelling framework, that utilise simpler types or can be based, in turn, on other complex type. Essentially, /entity attribute-type/ uses type structures to form compound attributes that can then be linked with entities through /function attribute mapping/. It is for this reason that we rate the coverage of /attribute complexity/ by FDM to be full.

Classification

In Chisholm's ontology, this term is related to the category (and term) /attribute/ in that Chisholm argues, apart from attributes, no other fundamental categories need to be defined to represent terms concerning classes or sets, and that attributes are sufficient to represent classes and sets. Further, he demonstrates that membership of classes and sets by individuals must arise from the attributes that are exhibited by the individuals concerned. That is, classes are formed from individuals that exhibit specific attributes. It follows that any rich class structure is extremely flexible, since the class structure is supported through the structure of attributes that are used to select it.

Increased flexibility is a logical extension of the requirement that sets and classes be assembled from attributes, in that it frees us from the rigidity of establishing classes from the outset as is popular in some specific data modelling frameworks. Further, individuals can be members of several classes simultaneously, and can move freely between classes. It is important to view this as a modelling requirement. Implementation issues need to be considered carefully in this light.

FDM supports this concept, by having attributes loosely coupled with entities. Consequently, in FDM, classification can be based on attributes that object-type entities exhibit by defining functions that collect object-type entities based upon the attributes they exhibit. Classification, based on aggregated individuals may be more difficult due to the lacking in full generality for /individual structure/, however, this does not limit support for the concept /classification/ in general.

Relation

Chisholm's ontology sees relations as special attributes exhibited by an individual, and that they are directed from an individual to another, that is, they are unidirectional. Relations other than diadic (involving two individuals) can also be represented, by a series of diadic relations.

FDM supports «relation», as expressed by Chisholm's ontology, in its full generality, and it does this through a combination of its concepts /function/ and /entity-type/. Functions can relate two or more entities, and /entity/ includes /entity object-type/ for the FDM equivalent to /individual/, and /entity attribute-type/ for the FDM equivalent to /attribute/. Functions are used to map between entities. Where a function maps between two objects of entity-type, a relation is specified, and in its simple form, it relates two entities or individuals. Where it maps between an entity-type and an attribute-type, it represents an attribute being exemplified by the object. Consequently, it can be said that relations in FDM can be seen as a special type of attribute. Essentially, FDM supports the concept /relation/ from the ontology through its own concepts of /entity/ and /function/ and in its fullest generality.

Temporal Relations

/Temporal relations/ concern relating events over time, and the concept is required because Chisholm's ontology defines events and state that do not use the concept of absolute time. Examples of absolute time are the date/time descriptions with which we are all familiar such as, for example, *January 1, 1983*. It makes no sense to ask about the fundamental nature or to describe anything fundamental about the time *January 1, 1983*, and it is simply a chronological marker. Other markers and measures can be (and have been) contrived to achieve the same purpose.

Instead of using an arbitrary marker of time to sequence events, the ontology uses the concept of temporal relations to order the occurrence of events. Attributes are temporally oriented to the past, to the future, or to the present, and can then be used in conjunction with individuals and state to order events.

FDM does not cover the semantic space defined by «temporal relations», because FDM does not deal with dynamics. In this context we mean dynamics in terms of changes in state and the description of changes in state. FDM is, instead, a traditional and static data modelling framework, and is not integrated with a dynamic modelling framework.

State

Recall that /state/ consists of five parts. Firstly, there is /state core/. Secondly, /state structure and parts/. Thirdly, /state event/, in that events are a specialisation, or subcategory, of state. Fourthly, /state causal relationships/. Finally, /state internal *versus* external attributes/.

Table 18—Expanded Results for /State/

Parts for /State/	FDM
Core	X
Structure & parts	X
Event	X
Causal relationships	X
Internal vs. external attributes	X

FDM does not support any of the parts of /state/ because FDM does not deal with dynamics and therefore has no need to discuss state or changes in state. Not unexpectedly, most of the work required in strengthening FDM will be in this area. However, it is likely that, because there is no dynamic modelling in FDM, that such a strengthening may not be as difficult as it may first seem, because the absence of concepts in FDM dealing with dynamics means that there are no concepts that are antonymous with respect to the dynamic concepts from the ontology.

We refer readers to the description Chisholm's ontology contained in Chapter 4, particularly the description of event and state that is contained in the discussion of dynamics. We will discuss the central issues concerning extensions to FDM in the concluding section of this chapter.

7·2 Results Comparing Chisholm's Ontology with FDM

In accordance with the method in Chapter 3, in this section, we present results of the 'reverse' comparison where we compare the ontology with concepts from FDM.

We proceed in the same way as in the previous section. However, in this section, we use concepts from FDM, and we present and analyse the results of comparing Chisholm's ontology with those concepts. You will recall, that we have already extracted the concepts for FDM in Chapter 5, and they are presented in Table 8 on page 74. You may recall from Chapter 5 that the number of concepts in FDM is small when compared with other data modelling frameworks. Consequently, FDM is less complex conceptually when compared with most other semantic data modelling frameworks of its time.

Summarising the method as described in Chapter 3, we use the concepts from FDM one at a time to compare Chisholm's ontology with FDM. In a way similar to the last section, results are presented in a table, and in addition to the indicative findings of similarity and difference contained in that table, there is an exploration of the substance of the differences and similarities in the following subsections, with each subsection covering one of the concepts from FDM.

The indicative results of this comparison are shown in Table 19. These are presented in accordance with the method.

Table 19—Indicative Results of Comparing Chisholm's Ontology with Concepts from FDM

Concept from FDM	Chisholm's Ontology
Entity	√
Function	√

The indicative results show that there is full support in Chisholm's ontology for the concepts from FDM. In this 'reverse' comparison, we are seeking to ensure that we take into account any subtle ways in which FDM supports concepts from Chisholm's ontology. We are far less interested in the indicative results themselves, and consequently, it is the discussion that emerges from examining these results that is much the more important. Upon first examination, the simplicity of FDM makes it less surprising that it is so fully supported by the ontology.

In the usual way, we proceed by examining each concept and by discussing the specific similarities and differences that we've found between these and Chisholm's ontology. We expand upon the indicative results, for compound concepts to show the degree of coverage found in the ontology by the parts of the relevant concepts. This further expansion is only necessary for complex concepts of which FDM has two—/entity/ and /function/. We conclude the section with a summary.

Entity

/entity/ consists of three parts. Firstly, /entity core/. Secondly, /entity object-type/. Thirdly, /entity attribute-type/. We present, in Table 20, indicative results for /entity/, based on these three parts.

Table 20—Expanded Results for /Entity/

Parts for /Entity/	Chisholm's Ontology
Core	√
Object-type	√
Attribute-type	√

We can see from this table that there is total overlap in semantic field evident in Chisholm's ontology for «entity core». However, the overlap is not contained within a single concept in the ontology. «Entity core» in FDM concerns the representation of objects (like individuals) and attributes. FDM's «entity core» is supported by /individual/ and /attribute/ from Chisholm's ontology.

Exemplification of attributes in FDM, however, requires more than simply /entity/. Nevertheless, «entity core» in FDM discusses type hierarchies and structures. This resonates with attribute structure in Chisholm's ontology.

Secondly, «entity object-type» is fully supported in Chisholm's ontology by the /individual/. We do not say that /individual/ is fully covered by the concept, but that /entity object-type/ is fully covered by parts of /individual/.

Finally, we found that «entity attribute-type» is fully supported in Chisholm's ontology, through the concept, /attribute/ that covers, in part, the semantic field defined by «entity attribute-type».

Function

/Function/ is used in FDM to represent a number of modelling fundamentals concerning attributes, relationships, aggregation, and classification. /Function/ consists of five parts. Firstly, /function core/. Secondly, /function relationship mapping/. Thirdly, /function attribute mapping/. Fourthly, /function aggregate mapping/. Finally, /function classification mapping/. In Table 21, expanded results for /function/ are provided. These are based on these three parts of the concept.

Table 21—Expanded Results for /Function/

Part of /Function/	Chisholm's Ontology
Core	√
Relationship mapping	√
Attribute mapping	√
Aggregation mapping	√
Classification mapping	√

/function core/ is the mapping from entity to entity. In mathematics, a function is equivalent to an ordered pair, in that it represents a connection between a domain value and a range value. In FDM, it plays the role of linking entities of various types in order to represent such things as relationships, exemplification, classification and types of aggregation. The ontology resonates with these more than it does with the mathematical concept of a functional mapping between entity types. However, in its more fundamental sense, a function in mathematics can be represented as a relation. In the ontology, a relation is an ordered pair, and an ordered pair can be used to represent a relation. A function can be written as an ordered pair. Then, as a consequence for the ontology's support for ordered pairs through /relation/, we can say that the ontology supports /function core/.

There are four parts to /function/ beyond the core, and these represent other dimensions to the relationship with the ontology beyond a purely mathematical one. FDM uses functions to represent relationships between two or more object-type entities (/relationship mapping/), and to represent mapping from object-type entities to attribute-type entities (/attribute mapping/). This covers the ontological concepts for «relation» and «attribute exemplification» respectively. Functions are also used to deliver two further concepts. Firstly, /function aggregation mapping/ allows for a crude form of aggregation to be specified that is covered by /individual structure/ in the ontology. Secondly, /function classification mapping/ is also fully covered by concepts from the ontology through tbe concept dealing with the membership of classes being based upon the attributes exemplified by individuals.

In summary, /function/ is required in FDM in order to cover part of the semantic field of both «relations» and «attribute exemplification» from the ontology. That is, the part of each that either joins two object-type entities in a relation, or indicates the exemplification of attributes. In FDM functions map from entity to entity; and entity encompasses both individual (object) and attribute. Thus, functions associate individuals with individuals, and individuals with the attributes they exemplify. Further, they deliver classification, and a crude form of aggregation of object-type entities.

7·3 Consolidation of Results and Implications

Having reported the results of a bidirectional comparison of FDM using concepts from Chisholm's ontology, we now consolidate these results so that specific areas where FDM fails to cover the semantic areas defined by concepts from the ontology can be identified and addressed. Clearly, the bulk of information for this consolidation is found in the 'forward' comparison of FDM with Chisholm's ontology (the first comparison reported in this chapter.) However, some insights may arise from the 'reverse' evaluation.

We consolidate the results in Table 22, and show, using a three-level indicator that is consistent with the method, the level to which the semantic fields as specified by the concepts in the ontology are covered by FDM. This consolidation shows the relationship between the world view as encapsulated by the ontology and that found in FDM in tabular form.

Table 22—Consolidated Results for Comparing FDM with Chisholm's Ontology

Ontological Concept	Level of Coverage
Individual	partial
Attribute	partial
Classification	full
Relation	full
Temporal relations	none
State	none

The gaps in coverage in FDM fall into two groups. Firstly, concerning the static concepts from Chisholm's ontology. Secondly, the dynamic concepts from Chisholm's ontology. The latter betraying a complete lack of support for dynamic terms in FDM, due to FDM being a traditional data modelling framework.

We examine the static concepts /individual/, /attribute/, /classification/, and /relationship/. Using Table 22 we can see that FDM provides full coverage for «classification», and «relation» (however, coverage for /classification/ is in actual fact affected by partial coverage of «attribute» and «individual»), and it has partial coverage of «attribute», and «individual».

The ease with which extensions can be made to cover the areas of partial coverage will depend upon the orthogonality between the concepts in the ontology and those already provided by FDM. That is, partial coverage implies that there are gaps in the coverage in FDM of semantic fields described by concepts from the ontology. If there are concepts in FDM that are antonymous in a contradictory way with respect to the semantic area of certain gaps, then extension of FDM will be difficult. If the antonymy is one of a complementary nature, then extension will be easier. If there is no relationship between the semantic area of the gap and other concepts in FDM, then extension is likely to be straightforward. However, this is also likely to be affected by the conceptual complexity of FDM.

Upon deeper investigation as outlined in a previous section, «individual structure», and «attribute equivalence» are the only areas not fully covered by FDM.

There is no relationship between /attribute equivalence/ with any concept in FDM. That is, there is, indeed, a true gap in the semantic coverage of FDM with respect to this /attribute equivalence/. It is likely, therefore, to be easier to augment FDM with new concepts that aim to provide coverage for «attribute equivalence».

With respect to /individual structure/ the story is different., because FDM has concepts in this semantic area. The support that is currently in FDM is a crude form of aggregation that is provided independently from that which defines the equivalent to individuals in FDM. In order to fully cover the semantic area of /individual/ from the ontology, all aspects of /individual/ will need to be treated in a similar way in FDM. That is, to regularise the treatment of object-type entities with respect to structure. Given that FDM already has concepts that are antonymous with respect to «individual structure» it is likely to be difficult to provide FDM with support for that concept in its fullest generality.

FDM does not support either of the dynamic concepts from Chisholm's ontology, and it lacks any coverage for «temporal relations» and «state». It is likely that adding this functionality will be more straightforward than for the static terms, due to the likely orthogonality between FDM and these concepts, in turn due to the fact that FDM does not support dynamic terms. That is, FDM does not have any concepts that are antonymous in any way with the semantic areas covered by the dynamics of the ontology.

The 'reverse' comparison failed to reveal very much of use showing instead full coverage of concepts from FDM in the ontology. This is because of the clean and simple nature of FDM and may in turn indicate that FDM is likely to be a good candidate for extending to provide a fuller degree of support for the ontology. We indicate our intention to pursue FDM as a candidate for extension in order to determine if this is indeed true.

7·4 The Extensibility of FDM

We have found that FDM does not reflect the world view of the ontology in certain areas, in that FDM lacks coverage of the semantic areas defined for the conceptual entailment of attributes («attribute equivalence»), and for compound individuals («individual structure») as defined in the ontology. We have also found that there is no coverage of the semantic area defined for the dynamic concepts /state/, and /temporal relations/. By reflection, support for /classification/ is restricted by partial coverage of the semantic are defined by «individual» and «attribute» along the lines outlined.

Three discrete areas from the ontology show partial or no support

- states and events (or the dynamics of reality)
- conceptual entailment (or attribute equivalence)
- individual structure (essentially aggregation)

Recall that in Chapter 3 when we described this method, that extending a specific data modelling framework proceeds in two stages.

Firstly, an assessment of the degree of difficulty with which we are likely to be faced when extending or strengthening a framework. This is accomplished using the analysis technique we outline in the method. We mean analysing the gaps in coverage of the semantic area defined by the partially covered concept from the ontology to determine if any relationship exists with parts of concepts from the data modelling framework. This is done to determine whether matters are complicated by an antonymous relationship between certain semantic areas covered by concepts in the ontology and concepts in the data modelling framework. In cases where antonymous relationships that are contradictory or contrary exist extension to cover the semantic areas that are missing will be difficult. Where contrary or no relationship exists, then extension may be more straightforward, depending upon the conceptual complexity of the data modelling framework.

Secondly, having assessed the likely difficulty of extension, several changes to the modelling framework can be outlined, or approaches recommended that may lead to a similar outcome.

The difficulty that one faces in strengthening FDM in each of the areas outlined above depends upon the nature of existing coverage for the semantic area each of the concepts.

The first area is that concerning state and events. FDM has no concepts in this area due to its nature. It is a traditional data modelling framework. Further, since FDM was conceived, there have been significant advances in the related area of functional languages. It seems reasonable that a functional language with state, such as

ML[1] would serve as a mechanism by which aspects of dynamic behaviour could be incorporated into FDM. An area requiring further research is the nesting of state and the unity between state and event in ways that mimic the ontology. The clean nature of functional languages (with state) should prove useful in casting the dynamics of FDM without destroying its own simplicity.

Secondly, /attribute equivalence/ is essentially related to meta data within FDM. Meta-data of this nature is not considered in our description of FDM. However, it is true to say that /attribute equivalence/ could be dealt with using some mechanism of defining synonyms within FDM. In any case, there is nothing that conflicts in FDM with this concept, in that there is no part of a concept that is antonymous with «attribute equivalence». It is for this reason that extension to FDM in this area is likely to be straightforward.

Thirdly, /individual structure/ is by far the most difficult area in which to suggest extensions in FDM. This is because, in part, the concept providing limited structure for individuals (namely, /function aggregation mapping/ is antonymous with respect to part of the semantic are defined by «individual structure» in the ontology. Moreover it is antonymous in a contradictory way. The mechanism by which structure is provided in FDM assumes that the aggregated object-type entity (in FDM parlance) is not a fully-fledged individual when compared with the individuals that establish the aggregation. That is, the aggregation has no separate existence. This therefore breaks the very underpinning of /individual/ that is described in the ontology. It is in this way that FDM contradicts part of the semantic area defined by /individual structure/. Further research is required to determine the specific ramifications of this departure. Nevertheless, because /entity object-type/ in FDM is not responsible for the structural issues concerning individuals, it becomes difficult to see in what way «individuals structure» can be covered in FDM. Further investigation into modern functional techniques may uncover possible avenues for extension.

Nevertheless, the simple nature of FDM is something that we use to predict that extension and augmentation of FDM is to be relatively straightforward. We also expect that, in using functional languages, the simplicity and cleanliness of FDM can be maintained while at the same time providing for dynamic concepts.

7·5 Reflections

We have found that the world views of FDM and Chisholm's ontology have a reasonably large degree of overlap for the statics of Chisholm's ontology, and that the concepts from the ontology have been found to cover similar semantic fields to those covered by concepts from FDM. However, we have found (not unsurprisingly) that the dynamics of Chisholm's ontology are very poorly covered by FDM, since FDM is a traditional semantic data modelling framework, and does not deal with dynamics.

In response to our findings of gaps in the coverage of semantic fields defined by concepts from the ontology, we have examined the nature of the gaps, and we have

[1] ML is a functional language that incorporates the manipulation of state. It was developed by Robin Milner at The University of Edinburgh.

argued that FDM can be strengthened with a relative degree of ease in order to bring its world view closer to that of Chisholm's ontology. Moreover, we have identified parts of the statics of FDM that are in need of extension in order for it to better cover concepts from Chisholm's ontology. We expect that only one part of FDM, concerning the structure of ontologically independent individuals, will be difficult to extend, and we further predict that strengthening FDM with respect to dynamics, and concerning attribute equivalence, where there is little overlap between Chisholm's ontology and FDM, will be relatively straightforward because there are no concepts in FDM that are antonymous with respect to concepts from the ontology.

We expect that, given the areas and level of agreement between the world view found in Chisholm's ontology and that in FDM, the data modelling framework will be well suited to development that will bring it closer to Chisholm's ontology (regarding world view). There is also great potential in utilising functional languages that include state, such as ML, to bring FDM much closer to the ontology.

We will reflect upon our experience in using the method in, when we critically analyse this and the comparative method. In that analysis, we suggest ways to make the process easier.

CHAPTER 8

Results of Evaluating OMT's Models using Chisholm's Ontology

IN THIS CHAPTER we present the results of an evaluation of OMT using Chisholm's ontology.

We described the statics of OMT that are compatible with UML in Chapter 5, but we have not yet described OMT's dynamics, and so we begin this chapter by doing so. Following this we conduct an evaluation of OMT using Chisholm's ontology. The evaluative method involves a bidirectional comparison between OMT and Chisholm's ontology. It is for this reason that we present the comparison in two sections. In the first of the two sections we present results for a comparison of the concepts of OMT with Chisholm's ontology. In the second section, we present the results for a comparison of the concepts of Chisholm's ontology with OMT. In each of the two sections we analyse the results for that section. Consolidation of our bidirectional comparison follows, before we summarise our findings indicating the areas requiring extension and the likely extensibility of the modelling framework. Finally, we reflect on the experience.

We have found a significant overlap in concepts between the ontology and OMT. The dynamics of OMT, while not being as elegant as the ontology in the way that concepts are represented still has a significant degree of overlap with the ontology. The consequence is that this may affect OMT's ability to express reality elegantly. We expect that modification of OMT to bring it closer to the ontology in terms of full coverage of concepts is likely to be more difficult than with FDM because in OMT, there are a larger number of concepts supporting dynamics when compared with the ontology.

After describing the dynamics of OMT, we proceed by comparing Chisholm's ontology and OMT according to the method described in Chapter 3.

We conduct the comparison in two parts, and reflecting this, we present the results in two sections. Firstly, we compare OMT with concepts from the ontology. Following this we conduct the 'reverse' comparison and compare Chisholm's ontology with concepts from OMT. After each part of the comparison, we pause to discuss the results, as specified in the method. We use this discussion to explain our findings and to help in the consolidation of the results which we present later in the chapter.

Readers should be familiar with the ontology and with OMT. We remind the reader that OMT is introduced in two places in the thesis. Firstly, in the following section of this chapter (OMT's dynamics), and in Chapter 5 (OMT's statics that are compatible with UML) where the statics of all the data modelling frameworks used in this thesis are described. Chisholm's ontology is described in Chapter 4.

8·1 OMT's Dynamic and Functional Models (the Dynamics of OMT)

Object models differ from other data modelling frameworks in that they typically incorporate more than simply static terms, and include terms that describe dynamic aspects of reality. Some incorporate these in object specification, by including concepts such as methods or operations that describe more complex attributes and behaviour. According to Blaha et al. (Blaha and Premerlani, 1998), the modelling concepts for dynamics is mostly contained in OMT's dynamic model (Rumbaugh et al., 1991) with little in the functional model. In describing the dynamics of OMT, we adhere to the format used in Chapter 5.

There are also several concepts in the object model for OMT that are related to dynamics, namely operations and their implementation counterparts called methods, that are then further described in OMT's functional model. We shall clarify the roles played by operations and methods during our description.

OMT's Dynamic Model

OMT's "dynamic model characterizes the temporal interaction of objects and their responses to events." (Blaha and Premerlani, 1998). It does this by considering states of objects in the system, and further, by considering transitions in state, and defines related dynamics of object models. We proceed by describing the concepts in OMT's dynamic model, namely /event/, /state/ and /concurrency/, all of which are compound concepts.

The first two concepts in OMT's dynamic model are /event/ and /state/. Firstly, events handle stimulus that is external to an object, and which usually comes from another object or the external environment. Secondly, a /state/ describes all of the values of attributes and instantiations of associations (links) of a specific object. The relationship between event and state can be best characterised by the separation in responsibility between the dynamic model, and the object model. In OMT's object model, one specifies the objects in a system together with attributes and the links objects may have. It represents the structure for a snap-shot of a specific model. Contrasting with this, the dynamic model shows what sorts of states are allowed and how a model's state can change. A state shows the actual values for attributes of a specific object, and the links or instantiations for links, and events enable changes in state.

Change often results from an event stimulating an object. An event may prompt a change in state, or it may result in another event being sent to the originator or to a third object. In OMT, events and states can be generalised, and categorised, that is, they can be described in conceptual terms, in that states can be summarised in what are called state diagrams, and events can be shown using state transition diagrams.

In OMT, "an event is something that happens at a point in time." (Rumbaugh et al., 1991), and conceptually, an event has no duration and is instantaneous, and it can also be thought of as a one-way transmission of information from one object to another. An event may convey values, or may simply indicate that something has occurred. Two events can be categorised as either being unrelated to each other or related to each other. Two unrelated events are called concurrent events, in recognition of their independence in a causal sense. In contrast, if an event logically follows another event, then the two events are said to be related, or more exactly, they can be considered to be causally related to each other.

In OMT, each event is considered to be unique, however, for ease of modelling, in OMT events are grouped into event classes. Event and event class are different ideas, one being the idealised version (event class) of an actualised reality (event), but each are often referred to using the same name, *event*. It is the context that helps to distinguish between the two uses. This is directly parallel to the idea of an object and an object class that we encountered in describing the statics of OMT in its object model, in that it is often the case to use the term object to refer to its class.

Each event has attributes, one of which is the time at which the event occurs. Other attributes convey information about the event. For example, if an event shows *Air France* flight number *AF0221* departing *Lyon* at *18:03* on *March 3, 2000*, then it would have certain attributes. Firstly, the time would be explicit, 18:03 on March 3, 2000, as well as implicit, in the sense that every event has an associated time. Details such as the city of origin and flight number would also be attributes of the event.

A state is described to be "an abstraction of the attribute values and links of an object" (Rumbaugh et al., 1991), and it is through the categorisation based on the current values for certain attributes of an object together with specific links it possesses that one determines the state of an object at any instant. For example, a person could be considered to be in the state of being financially wealthy, if he or she has savings to the value above a certain amount of money (perhaps based on average national or regional savings levels).

In OMT, state is used to help specify an object's response to events. It does this by allowing the specification of a subset of attributes for an object that "affects the gross behavior of the object." (Rumbaugh et al., 1991). The consequence is that one can categorise the state and the subset of attributes can be used when examining the effect of events on objects in a specific class. Further, a number of attributes can be analysed to determine the typical states an object will assume, and then, based on the values for attributes and permitted changes in state through the effects of events, the likely paths of transition to other states can be described. By predetermining the likely states for an object in this way, one can then specify ahead of time how an object will react to an event that will change the state of the object in known ways. A specific object's response to an event will depend upon the specific values held by the object, but qualitatively, the response will be similar to other objects in the same categorisation of state. In OMT, when describing a state in this way, one specifies how an object in that state will behave in response to an event based on values of important attributes.

There is an alternative view that helps distinguish between state and event, whereby the difference between event and state can be described in the following way: "A state corresponds to the interval between two events received by an object. Events represent *points* in time; states represent *intervals* of time" (Rumbaugh et al., 1991) (in this case, the emphasis is ours), and means that events enable changes in state, whereas states endure between events. Another way of looking at state is to say that an object is in a state when it satisfies some condition, however complex the condition.

The idea of an event being instantaneous is not entirely accurate, in that events in OMT may activate operations in objects. Recall from the object model, that operations can be used to specify what an object does in response to an event or what an object does upon finding itself in a specific state after a transition in state occurs. Consequently, it can be said that events can trigger operations, and so, operations are used to describe what an object does in response to a state or an event.

Operations are categorised as being either activities (longer than instantaneous) or actions (instantaneous). Activities are often triggered by an object reaching a particular state, and so may be described as being tied to state. Actions, on the other hand, are associated with events, in that an action is stimulated by an event and itself is instantaneous. However, the distinction between the two is a moot one, since an action may results in a change in state that in turn may induce an activity to take place. Activities can be continuous operations. For example, a person will breathe for their entire life (until they die). An activity may cease after a specified time, or an activity may persist until an event terminates it by changing the state of an object. An action, since it is associated with an event, is instantaneous. For example, a classic form of an action is that our *Air France* flight may arrive at *Paris Charles de Gaulle II*. In response to the event of docking at the air bridge (passenger connector with the terminal building,) an action may be to highlight on the monitors in the terminal building that the flight has docked and is allowing passengers to leave the aeroplane. Both types of operation can be incorporated into state diagrams. Actions also include more "house-keeping" matters such as setting attributes, or generating other events.

OMT uses state diagrams to relate states of objects to events involving those objects. State diagrams are seen not as tightly formalised structures but instead as guides to the behaviour of a system and show the state sequence that results from an event sequence. Each event in OMT is unique, however, in the context of state diagrams, event is being used to denote the generics of a specific event class. A state diagram shows "the state sequence caused by an event sequence. If an object is in a state and an event labelling to one of its transitions occurs, the object enters the state on the target end of the transition." (Rumbaugh et al., 1991). A transition is an arc leading from one state to another, and signifies change. Conditions can be specified on arcs that indicate when the transition can fire. For example, only when the aeroplane has all of the passengers that are expected can our Air France flight depart from *Lyon*.

States and events can be structured in ways analogous to objects, through the generalisation of events and states, and event and state concurrency results from the aggregation of objects. Recall from modelling fundamentals that generalisation is where a classification logically divides into a number of distinct sub-classes and one effectively selects the most appropriate subclass for an object. Aggregation is where the whole is made up of a number of parts. We consider each in turn.

OMT supports generalisation of events and states. Event generalisation allows the grouping of events into a hierarchy, where there is "inheritance of event attributes" (Rumbaugh et al., 1991) from classes higher in the hierarchy. Time is inherited by all events in the hierarchy. The root of the hierarchy is the class *event* for which *time* is the only attribute. The example cited by the authors in (Rumbaugh et al., 1991) is that of input device event specification. In this context, an input device is one connected to a personal computer (such as a mouse, or keyboard). These hierarchies enable handling of state transitions differently according to which level of abstraction is required and enables triggering of transitions at different levels of the hierarchy.

State generalisation is where the state of an object is described at any one time by one of its substates, and it is best described using a nested state diagram where the substates of an object are shown. For example, a traffic light moves between the states, red, green, and amber. Each substate at any point represents the state of the object and is adequate for representing the object's state, and the object is allowed to be in only one of these substates. This is also useful when the substates themselves have some degree of complexity yet the state of the whole is fully defined by one of them, and the state of the object moves between the various substates. Each substate can be further nested.

Concurrency concerns the simultaneous activity of things, and in OMT, concurrency comes in the form of either aggregation concurrency or concurrency within an object.

Aggregation concurrency emerges from two facts. Firstly, from the fact that a dynamic model consists of a description of a number of concurrent objects, and that "the objects in a system are inherently concurrent and can change state independently" (Rumbaugh et al., 1991) due to their independence from one another. Secondly, the aggregation of objects leads to further concurrency. Our aeroplane is a classic example of aggregation in that it is made up of a number of different parts (aggregation), each part has a number of individual states and events that enables transitions between states. The engines have their events and state, the ailerons on the wings and tail have their events and state. They may coordinate, but they can exist quite separately and they can change state quite separately. A related example is that of an air space controlled by air traffic control. Each part of the system is complex and many objects make up the whole. Each object has its own behaviour. One of the aircraft in this controlled air space may well be our Air France aeroplane.

There can also be concurrency within an object, where groups of attributes and links can be found that are semi-autonomous, and their states and events can therefore be defined separately. This is a concept for which it is difficult to find similar examples in reality but is a little like part-whole relationships but where the parts are not external to the object but instead are defined within. Examples that best illustrate this idea are ones where the object itself is essentially a process, such as a rubber in the game of Bridge (Rumbaugh et al., 1991).

Summarising, all of the specification mechanisms here are concerned with describing repeatable events and states, and are similar to causal relationships, where states are used to predict changes in state. Nevertheless, it is important to note that specific events as they occur in OMT are thought of as being distinct, and in this way are independent of an event class. Further, an event, in OMT, does not repeat.

The relationship between the OMT's object model and its dynamic model is one based on the class structures within which objects sit in that the dynamic model is where one specifies the allowable sequences of changes in state for objects defined according to the object model. A state of an object is the specific instantiations of attributes and associations for the object. Associations are known as links when instantiated. Events can be thought of as the activation mechanism for operations on the object model.

The class hierarchy from the object model also influences the dynamic model. For example, in cases where overriding is not utilised, an object of a subclass inherits the behaviour of its superclass, and in this way the dynamic model structure is restricted and constrained by the object model structure. Additionally, in OMT objects cannot be in different classes simultaneously. It is therefore necessary to model differences in state for a specific object by the allowable states for a specific class of objects.

OMT's Functional Model

OMT's functional model describes the operations of objects defined during the design of a specific object model using OMT's object model, and from OMT's dynamic model. It "... specifies operations from both the object and dynamic models. As such, it defines the computations that objects perform." (Rumbaugh et al., 1991) In their later book this is elaborated, in that "[t]he functional model defines the operations that objects perform and to which they are subjected. Operations arise from the object model from queries, updates, derived data, and constraints. Operations arise in the dynamic model from events, guard conditions, actions, and activities" (Blaha and Premerlani, 1998). It seems reasonable from this, that the functional model is not going to give us any additional information concerning the modelling of dynamics. Instead, it gives us details about specific algorithms for specific operations, and is something in which we are not interested.

Summary

Table 23 summarises the dynamic elements of OMT, taken from the dynamic model.

Table 23—Concepts for the Dynamics of OMT

Concept	Part	Description
Event	Core	Events in OMT are instantaneous. Further, they occur at a *time*. Each event is unique, although they may be generalised and represented in an event class. Events convey information to objects. It is a one-way transmission of information from one object to another. Some are signals, others carry attributes as information.
	Relating	Events may be related. Alternatively, they may be unrelated. Related events are similar to causal relationships.
	Operations	Operations are triggered in objects in response to events. They can be instantaneous or enduring. An activity takes time to complete (is not instantaneous) and is in response to being in a state. An action is instantaneous and is in response to an event.
	Generalisation	Event generalisation allows the inheritance of attributes from events higher in the hierarchy of events, in much the same way as objects themselves. This permits abstraction of events at different levels. Further, time is an attribute inherited by all events from the root event (called *event*).
State	Core	According to the values of certain attributes possessed by an object, an object can be said to be in a certain state. For each state only the attributes directly determining it are important.
	Transition	The movement of objects between states is triggered by events and external stimulus. This is called a state transition and can be summarised diagrammatically
	Generalisation	OMT uses nested state diagrams. This way state generalisation can be used to describe the states of an object. Further, the nested nature allows for complexity to be hidden until needed. Further the state of the object at any time is one of the substates defined in the generalisation.
Concurrency	Core	OMTs dynamic model supports concurrency. It comes in two forms. Aggregation concurrency and concurrency within an object. The more common is aggregation concurrency
	Aggregation	This is where an object consists of several other objects. Each object will have its own state and behaviour.
	Within object	Concurrency within an object concerns objects where there are several reasonably independent substates that are not objects in their own right, yet have reasonably independent behaviour.

8·2 Results Comparing OMT's Models with Chisholm's Ontology

In this section we present the first part of the bidrectional comparison, and show the results of comparing OMT with concepts from Chisholm's ontology. We have already extracted the static and dynamic concepts for Chisholm's ontology, and these are summarised in Table 5 on page 62, and Table 6 on page 69 respectively. The statics for OMT are described in Chapter 5, while the dynamics of OMT are described in the first section of this chapter.

We use the concepts from Chisholm's ontology one at a time, and relate them with OMT in the ways described in Chapter 3. A summary of results of this comparison is shown in Table 24, and this table shows the indicative results of the comparison which we explain and discuss. The 3-level indicator that we use is described fully in Chapter 3.

Table 24—Results of Comparing OMT's Object, Dynamic, and Functional Models with Chisholm's Ontology

Ontological Concept	OMT (Object, Dynamic, & Functional Models)
Individual	$\surd$
Attribute	$\surd_p$
Classification	$\surd_p$
Relation	$\surd_p$
Temporal Relations	$\surd_p$
State	$\surd_p$

The indicative results in Table 24 show a significant degree of support in OMT for a number of concepts from Chisholm's ontology. We proceed by discussing each concept from Chisholm's ontology and by examining specific differences between Chisholm's ontology and OMT.

In discussing each concept from the ontology, we evaluate and analyse OMT to determine the nature of any semantic coverage that is evident. We expand upon the indicative results for compound concepts /individual/, /attribute/, and /state/. We conclude the section with a summary.

Individual

/Individual/ is fully supported by /object/ in OMT. Recall that /individual/ in Chisholm's ontology is a compound concept and has three parts. Firstly, /individual core/. Secondly, /individual identity/. Thirdly, /individual structure/. Table 25 shows these parts and their levels of support found in OMT. The results are shown using the same key as used previously.

Table 25—Expanded Results for /Individual/

Part of /Individual/	OMT
Core	√
Identity	√
Structure	√

Firstly, by inspecting the table, we see that OMT covers the semantic field defined by «individual core» through /object core/, in that there is support for the representation of significant objects or entities that are ontologically independent in a way that is identical to the view outlined by the ontology.

Secondly, individuals can be uniquely identified in OMT. OMT achieves this through /object identity/ and by determining attributes that collectively identifies a specific individual. Since OMT is organised into classes, that the collection is defined at the class level, but applied to each object in a specific class. The fact that attributes are defined at the class level does not lessen the degree of support for «individual identity».

Finally, OMT covers the semantic field for «individual structure» through /link-association aggregation/, and means that in OMT, individuals can easily be subdivided into parts as stipulated in Chisholm's ontology. The modelling concept of aggregation is precisely the same as «individual structure». In this case we are seeking support for constructing individuals that are based in part upon other individuals, and each part is an object and so has an existence in its own right. In line with the ontology, OMT also treats the aggregation instantiation very much on an object by object basis.

Summarising, OMT adequately covers the semantic field defined by «individual».

Attribute

The compound concept /attribute/ has three parts. Firstly, /attribute core/. Secondly, /attribute equivalence/. Thirdly, /attribute complexity/. Table 26 shows these concepts and the level of coverage found in OMT. This table shows that not all concepts are supported by OMT. We examine each part in turn.

Firstly, /attribute core/ concerns the fact that individuals can exemplify attributes, and means that individuals can be described in detail using attributes. OMT clearly supports this concept through its own concept /object attribute/. However, in Chisholm's ontology, /attribute core/ extends beyond this to encompass two extra elements. Firstly, Chisholm argues that attributes are enduring. Secondly, that attributes and individuals are loosely coupled.

Table 26—Expanded Results for /Attribute/

Part for /Attribute/	OMT
Core	$\surd_p$
Equivalence	X
Complexity	√

Attribute endurance concerns the fact that it makes no sense to discuss the beginning or ending of a specific attribute. For example, was it last Tuesday when the attribute of 'being green' was first described? Alternatively, was it three centuries ago? Clearly, it then follows that the attributes themselves are enduring even if no observer perceives them for a long period. OMT doesn't support this. Some object models do support this concept to its fullest generality.

«Attribute core» covers another aspect that flows directly from attribute endurance, namely, loose coupling of attributes and individuals. In contrast to the ontology, OMT requires that attributes and objects be tightly coupled. OMT's object model emphasises the bond between object class and the attributes that are used to describe objects in that class, and emphasises the fact that attributes are common to all members of a particular class of objects.

The second part to /attribute/ is /attribute equivalence/. In the ontology this is 'conceptual entailment', and is a semantic equivalence between attributes. It is most useful when crossing cultural boundaries or where there are several terms for the same attribute, particularly if some are lesser known or archaic. OMT does not have concepts covering the semantic field for «attribute equivalence». We can also say that there is no concept covering a semantic area related in an antonymous way to «attribute equivalence». It may be that /attribute equivalence/ is only of philosophical value, in that it is required in order to avoid difficult questions regarding the equivalence of attributes. However, such a requirement may be of practical use when querying an implemented model in a specific database system, where the conceptual entailment of attributes may enable the formulation of more powerful queries.

Thirdly, Chisholm's ontology has /attribute complexity/. «Attribute complexity» means that compound attributes of some nature must be supported, and can be constructed from either simple attributes or from other compound attributes. The ontology is not prescriptive as to the nature of the mechanism for supporting this concept, but notes that conjunction and disjunction of attributes would be one way through which compound attributes could be supported.

OMT's object model covers the same semantic field «attribute complexity». It is the concepts /object attribute/ and /object operation/ that provide for complexity in attributes in such a way as to cover the semantic area defined by «attribute complexity». «Object attribute» allows for a structure to be declared based upon simpler attributes, and «object operation» allows for calculations of a complex nature to be performed.

Classification

OMT partly covers the semantic field defined by «classification». At length, Chisholm argues that in order to represent classes and sets, there is no requirement for anything more than /attribute/, and that /attribute/ is sufficient for representing classes and sets, and he demonstrates that classes and sets arise from the attributes that are exhibited by individuals. Consequently, it is attributes exemplified by the member individuals that defines membership of classes, and the individuals exist separate from the classes to which they are members. Individuals may also be members of several classes simultaneously.

A logical consequence of the requirement that sets and classes be established from attributes is that of flexibility. This requirement frees us from the rigidity of establishing the membership and attributes that apply to classes from the outset. Individuals can also be members of several classes simultaneously, and can move freely between classes, because the individuals do not exist thanks to the classes but are instead fundamental to 'what there is'.

It is clear that OMT partly covers the semantic field defined by «classification». The rigidity in the object model, with essentially a premeditated attitude to classesmeans that the gap in coverage of the semantic field defined by «classification» betrays a concept that is antonymous in a contradictory way to «classification». We mean 'premeditated' from the viewpoint that classes are defined early with rich (and somewhat rigid) class hierarchies, and the flexibility evident in the ontology in this context is very problematic, requiring a wholesale redrafting of classifications in OMT that are long established.

Chisholm's ontology allows the description of class hierarchies through judicious selection of attributes for classes that are related in this way. However, since the individual exists separate from the classes, and are members of classes thanks to specific exemplification of attributes, modelling flexibility is maximised in the ontology.

Relation

Chisholm's ontology views /relation/ as an attribute of an individual that is unidirectional. OMT's object model covers the semantic field covered by the concept through «link-association». In /link-association/ members of the classes appear linked through what appear to be bidirectional associations. The matter, however, is complicated by the fact that, when describing the dynamics of OMT, the state of an object includes the links it possesses, which in turn implies that the instantiation of relationships is, in fact, unidirectional. We rate the coverage of the semantic field defined by «relation» from the ontology, in OMT to be similar to, but not to the full generality of, Chisholm's ontology because we cannot be certain that the intention is for unidirectional relations. In some implementations of a model designed using OMT a unidirectional link may eventuate whereas in other implementations of the same model may result in a bidirectional link. OMT could be easily amended to provide such a functionality because the concept in OMT is not antonymous in a way that is contrary or contradictory with respect to the gap in coverage of the semantic field defined by «relation».

Temporal Relations

«Temporal relations» is a recognition that there is not only a chronological sequence and ordering between states and events, but also that there is nothing ontologically fundamental about time as we record or talk about it. We have found that there is significant commonality between OMT and the ontology with respect to /temporal relations/ in that OMT temporally orders events and states. However, since the mechanism by which the ordering is provided is time as an explicit marker, we rate the coverage of the semantic field defined by «temporal relations» as partial.

It is clear that a system of ordering states and changes in state, such as Chisholm uses the concept, is supported by OMT's models, however, it is also clear that in OMT, time is the mechanism used to model state and therefore to model temporal relations between states and events. Temporal relations are, however, a recognition that states and state changes do occur in a specific sequence that is obviously closely related to a chronological sequence, however, that does not indicate the existence, in fundamental terms, of a specific measurement of time. We draw the conclusion that the inclusion of time in OMT's models is really another implementation technique providing for what is a very closely related concept to temporal relations. An important aspect to /temporal relations/ is the chronological ordering states and events and not the mechanism by which the ordering is provided. However, the unification of event and state and associated benefits will not be realised until OMT adopts an alternative approach to chronological ordering that does not rely upon arbitrary time, such as the approach taken in Chisholm's ontology. The full implication of the unification of event and state is a topic of future research.

Summarising, OMT partly covers the semantic field defined by «temporal relations», however time as used in OMT is a related but relatively unsophisticated way of providing for temporal relations. The coverage, by OMT, of the semantic field defined by «temporal relation», is therefore partial. The full implications of OMT's style of state and event and how it interacts with temporal relations needs further investigation.

State

Recall that /state/ is a compound concept and has five parts. Firstly, /state core/. Secondly, /state structure and parts/. Thirdly, /state event/. Fourthly, /state causal relationships/. Finally, /state internal *vs* external attributes/. Table 27 shows these concepts and the degree of overlap in the semantic field for the concept from the ontology, in OMT.

It is pleasing to see from these results that there is a level of support for all but one of the parts of /state/, and as it shall be seen, the degree of support is encouragingly high. However, there are certain critical areas in which OMT fails to cover the semantic field defined by «state», and the nature of coverage is such, in some areas, as to indicate that extension may be difficult. We proceed by discussing each of the concepts in turn.

Table 27—Expanded Results for /State/

Part for /State/	OMT
Core	√
Structure & parts	$\surd_p$
Event	$\surd_p$
Causal relations	√
Internal vs. external attributes	X

/State/ is a good example of where the names of concepts are not relevant in comparing meaning. For example, in the ontology «state» covers the semantic field defined by «state» and «event» in OMT. In our discussion it will become clear that we are not interested in where coverage for the ontology is made manifest in OMT but instead we are interested in whether, and the clarity with which coverage of the semantic field defined by concepts from the ontology is given in OMT. In this case we are concerned with how faithfully the simplicity of the ontology with respect to state is reproduced by concepts in the ontology (whatever the names that are given to terms that define the semantic area equivalent to the concepts from the ontology). This is important since the clarity with which coverage is provided gives an indication as to the degree to which the power of the ontology is reflected in the modelling framework.

/state core/

In the ontology, /state core/ holds that the exemplification of attributes, changes in the exemplification of attributes and other events must be covered. Additionally, the contingent nature of individuals and states, in that they come into being and pass away, must also be covered.

We have found that OMT covers the semantic field defined by «state core» through parts of its own concepts /event/ and /state/, and we have found that coverage is to the fullest generality of Chisholm's ontology. In OMT, individuals can exemplify attributes, and also changes in the exemplification of attributes as well as events can be represented. Individuals, and states as they are represented in OMT can come into being and pass away. In OMT states may be enduring or instantaneous in a similar way to the ontology, and these are supported through the concepts /event core/, /event operations/ as well as /state core/.

/structure and parts/

Reflecting the nature and structure of /state/, in that it unifies state and event, our discussion of /structure and parts/ consists of two parts, in that we firstly discuss the issues concerning the structure of the states of individuals, before considering the structure of events.

Recall, that generally, state consists of a substrate, and a content. In the simplest case the substrate is an individual and the content is the temporal whole that represents attributes exemplified by the individual, and that each of the content and substrate reveals further structure. Firstly, we can say that the parts of the individual have their own state that is directly reflective of the structure of the individual. This style of compound state is largely supported by concepts from OMT, and further, that /link-association aggregation/ together with /object attribute/ and /concurrency/ delivers this aspect of the concept in OMT. Secondly, the temporal wholes that an individual assumes over time, using attributes that are oriented temporally, each of which describes state for chronological sequences of its history, is not supported in OMT, because OMT is rooted in the present with little by way of history being represented.

However, there are more complex states, and /state structure and parts/ reflects this, in that a state may be an event. In the ontology, events are either first order states of a specific sort, or second order states.

Events that are first order states can be thought of as special cases of that which we considered above. For example, a person breathing is an example of an event that is a first order enduring state involving an individual and an attribute. OMT supports this form of /state/ through the concept /event operations/ with activities and actions.

In contrast, events that are second order states are states of states, and this latter part of structure is supported in OMT's concepts /event relating/ and /state generalisation/. We are interested in the fact that events of this nature are possible, and are not interested in the particular mechanisms. The structure of each state involved in this complex is then considered in turn.

Summarising, the simplicity of /state structure and parts/ and associated power of the concept is not realised in OMT, and for this reason, we rate support for the concept not to the fullest generality to that which is evident in the ontology, in the sense that there is coverage of the meaning of events of this nature, but the unity of event and state in the ontology allows a powerful economy of terms when specifying states and events. This simultaneously provides for economy of expression as well as full transference of semantics concerning execution of change as it influences other individuals' states. The full impact of this part needs to be further researched.

/state event/

Events are happenings that affect individuals. In the ontology, events are a specialisation of state in that /state/ is used to cater for the exemplification of attributes that betray happenings as well as static exemplification. For example, Fred walking is an example of an event in the ontology. A non-happening is Fred being tall. These are both examples of what Chisholm calls 'first-order states'. They are the more common form of event. An event is a specialisation of state. That is, in the ontology, events are either first order states of a specific sort, or second order states.

An example of an event that is a first order enduring state involving an individual and an attribute is Fred breathing. OMT supports this form of /state/ through activities and actions., in that activities and actions represent simple events that individuals engage in from time to time. An activity describes special sorts of behaviour of an individual that essentially are a kind of enduring state. That is, where something happens for a period of time that does not affect any other states. OMT supports these first-order states.

In contrast, events that are second order states are states of states. This style of event is supported by a variety of means in OMT and by a variety of mechanisms where objects and events can cause various changes in states of other objects. These are described using state transitions that specify the effect state and changes in state will have. In the ontology, events of this type are one-off instantiations of nomological laws (that are known to those operating in the world). Therefore in OMT the facility of these higher order events are supported in OMT. That is, the meaning of an instantiation of a law by specific states is supported in OMT. We are not interested that it requires several mechanisms by which the support is made possible. For example, *(flight 451 crashing) contributing causally to (the bridge collapsing)* is an example of a second-order state. Here we see two event-type states linked. The parentheses have been added for illustrative effect. This is because this style of event is supported using state-transition diagrams together with actions and operations in OMT. We conclude that this style of event is supported in OMT.

Summmarising, the unification of event and state has powerful ramifications, in that it allows for complex discussion of an individual's history. Also, a nomological law in the ontology can be matched by a state that then gives rise to a second-order state. Within that state one can see the effect that law has on the individual concerned. In OMT, the contrast is that event is generalised and used as a template together with state transitions to show the same change, that may result in actions or activities being initiated, or values changed. By viewing state and event in the way Chisholm does, all of these complex effects can be powerfully represented using a unified term.

Consequently, we can say that OMT provides substantial overlap with the semantic field defined by «state event», and it appears that most of that which is intended by the ontology is supported in some way in OMT. However, a significant departure is the lacking in unity between event and state and we therefore rate the coverage of the semantic field defined by «state event» to be partial. The full implications of this are not apparent, but instead will be pursued in future research.

/causal relations/

The fourth part of /state/ concerns causal relations, and OMT supports the full generality of causal relationships, in that it recognises that certain events will lead causally to other events happening, and this is represented in OMT as state transitions of various forms. Individuals assume certain states, and from these states, and subject to certain stimuli, it will assume another state that can be predetermined from the original state given the stimuli.

/internal *vs* external attributes/

Finally, recall that Chisholm considers the difficulty in allowing individuals, that have ceased to exist, to continue exemplifying attributes, and he develops the concept of internal attributes vs external attributes. Internal attributes are exhibited by individuals that exist. External attributes are those that require another individual to ascribe the attribute to an individual that has ceased to exist. OMT does not cover the semantic field defined by «internal vs external attributes».

8·3 Results Comparing Chisholm's Ontology with OMT's Models

The second part of this method requires a comparison of the concepts found in OMT with Chisholm's ontology.

We proceed in the same way as in the previous section. However, in this section, we present and analyse the results of comparing the concepts found in OMT's models with Chisholm's ontology. The static and dynamic concepts for OMT can be found in Table 11 on page 82, and Table 23 on page 114, respectively. Chisholm's ontology is introduced and described in Chapter 4.

We have compared Chisholm's ontology with OMT's statics and dynamics, as contained in OMT's object, and dynamic models. A summary of the results of this comparison are shown in Table 28. These, in accordance with the method, use a 3-level indicator to show the level of coverage in Chisholm's ontology of the semantic field defined by the relevant concept from OMT. In addition to the indicative findings of similarity and difference contained in that table, there is an exploration of the substance of the differences and similarities, in the subsections following the table, with each subsection dealing with the concepts from OMT in turn.

Table 28—Results of Comparing Chisholm's Ontology with OMT's Models

OMT Concept	Chisholm's Ontology
Object	$\surd_P$
Link-Association	$\surd_P$
Generalisation	$\surd_P$
Event	$\surd_P$
State	$\surd_P$
Concurrency	$\surd_P$

The indicative results contained in Table 28 show a degree of coverage for a number of concepts from OMT in Chisholm's ontology. It is the discussion in this 'reverse' comparison that yields more information than this table, and we proceed to examine each concept and discuss the specific differences that we've found between concepts in OMT and Chisholm's ontology.

In our discussion we use the semantic field of each concept in OMT to more deeply discuss the extent to which a concept is covered by the ontology, and the discussion enables us to understand more deeply the reasons why a concept is supported to the extent that it is by the ontology. We expand upon the indicative results, shown in Table 28, for the compound concepts /object/, /link/association/, /event/, /state/ and /concurrency/.

Object

The concept/object/ is compound and has four parts. Firstly, /object core/. Secondly, /object attribute/. Thirdly, /object identity/. Finally, /object operation/. The indicative results for /object/ are contained in Table 29. We discuss each of these in turn below.

Table 29—Expanded Results for /Object/

Part of /Object/	Chisholm's Ontology
Core	$\surd_p$
Attribute	$\surd_p$
Identity	√
Operation	√

In OMT objects are assembled into classes. All member objects in a class have similar attributes, behaviour, links or relationships, and semantics. Member objects are often called instances of objects, or just objects.

The semantic field described by «object core» is partly covered by concepts from Chisholm's ontology. However, there is one area where Chisholm's ontology has concepts that cover a semantic area that is antonymous with respect to OMT. This makes Chisholm's support for «object core» lacking in full generality. Specifically, object classes in OMT are rigid, in that objects are determined to be members of a specific class, and membership of more than one class is not possible. An object must reside in exactly one class. Additional to this, object classes form complex class hierarchies. This is clearly antonymous in a contradictory way to the ontology, where «individual core» and «classification» are not linked, with individuals being quite separate from the classification structures that may exist.

Secondly, according to /object attribute/ objects can have attributes, which are defined at the object class level. Objects, when they exemplify attributes are said to have specific values for specific attributes. Attributes in OMT may be computed, or calculated based on values of other objects or using some algorithm. Links with other objects can be thought of as being similar to attributes and conceptually may be thought of as logical extensions to attributes in that they form part of the state for an object. /attribute/, /individual structure/, and /relation/ from the ontology are the similes for these concepts. However, in OMT, there is tight coupling between attributes and classes of objects, in that attributes are declared at the class level, thus restricting the attributes that can be exemplified by objects in the class to the 'lowest common denominator.' This means that the concept in OMT is in part antonymous with respect to the ontology, in that the ontology requires loose coupling of individuals with attributes. In summary, «object attribute» is partly covered by concepts from Chisholm's ontology, and that it would be difficult to extend coverage to be full.

The semantic field covered by «object identity» is fully covered by concepts in the ontology by «individual identity». In OMT a group of attributes exemplified by a specific object, are used to identify that object.

Finally, we have /object operation/. Operations can be used to calculate values. Alternatively, they can specify changes in state for an object through an event activating an operation, or, an operation can be activated through the state of an object. Operations can either be actions (instantaneous) or activities (enduring) respectively. Actions often involve changing the state of the object, or sending an event to the originator of the event that activated the action, or to a third object. Activities are activated when an object is in a specific state and can cease after a period of time, or endure until the state is left by the object. These do not have direct and single similes in Chisholm's ontology. However, operations of this type help cover part of «state causal relations» in that they determine the effect of an event. Secondly, they help to cover the necessary requirements to enable «state core» to the extent that enduring states can be defined. It is for these reasons, that Chisholm's ontology fully covers the semantic area of «object operation».

Summarising, we can say that there is a significant degree of coverage of the semantic area defined by «object». However, the rigid classification structures evident in OMT together with the tight binding of attributes and classes means that full coverage is going to be difficult to achieve due to antonymous relationships between these ideas and parts of concepts from the ontology.

Link-Association

/link-association/ in OMT allows the specification and representation of relationships between objects, and, as a specialisation of association and link, aggregation. It is compound, consisting of three parts, /link-association core/, /link-association multiplicity/, and /link-association aggregation/. The results of comparing the concept with Chisholm's ontology is shown in Table 30.

Table 30—Expanded Results for /Link-Association/

Part of /Link-Association/	Chisholm's Ontology
Core	$\surd_p$
Multiplicity	X
Aggregation	$\surd$

The core of the concept is not as clearly defined as one may think from inspecting the definition of the object model, in that associations are clearly seen as bidirectional. However, upon examination of the dynamic model, it becomes apparent that the state of an object is in part defined by the links an object possesses. This tends to suggest that links are similar in nature to attributes and are unidirectional. Further, aggregation, as we shall see shortly, is a specialised form of attribute. This also implies that links are unidirectional. If this latter interpretation is taken, then Chisholm's ontology supports this aspect of OMT through the concepts of «relation» and «individual structure», in Chisholm's ontology.

This doubt about the unidirectional, or bidirectional nature of associations and links, however, leads us to rate the covering of the equivalent semantic area in the ontology as lacking the full generality of the concept from OMT. It is predominately because the implementation of the models in OMT do not insist upon a unidirectional relation between objects, that we then conclude that there is a contradictory antimony between this and a semantic area from the ontology.

Secondly, /link-association multiplicity/ has no similar concept in Chisholm's ontology, because relations are defined individual by individual in the ontology. The idea of multiplicity is instead a concept derived from several others in the ontology. It is not to say that one cannot derive the idea for a specific state of affairs. This contrasts sharply with OMT, where objects are clustered into classes, and therefore it makes sense for multiplicity in participation to be specified, because of the need to describe how many objects from specific classes can have links with other classes. Consequently, there is no coverage for «link-association multiplicity» in the ontology and individuals can relate to others as needed.

Finally, in OMT objects can aggregate using a specialised form of association (in concept) and link (in instance.) This is not covered by «relation» in the ontology, but instead by «individual structure», in that individuals are made up of other individuals. The constituent individuals have existence beyond the whole. «link-association aggregation» is fully supported by Chisholm's ontology. The impact this coverage has, in turn, on events and states is described later in this section.

Summarising, we can see that two parts of this concept have either total or near-total coverage of their semantic area in Chisholm's ontology. However, full coverage of the concept is unlikely since there is an antonymous relationship between a concept in the ontology and the seemingly bidirectional nature of «link-association core».

Generalisation

/Classification/ in Chisholm's ontology does indeed cover a related semantic area as «generalisation» in OMT. Coverage in the ontology is not to the fullest generality of «generalisation» because there is no support in the ontology for a rigid class hierarchy into which individuals are slotted, and through which attributes and behaviour are inherited. Classes representing the «generalisation» relationship are supported to a limited degree in the ontology through careful choice of attributes that define the various classes that are related in this way. Inheritance, however, is not supported.

Additionally, in OMT an object cannot be a member of more than one object class in the hierarchy. For example, a specific match of football is an example of the class football game. This, in turn, is a specialisation of a game. However, while our specific game or match of football may inherit attributes from game, and may be considered a game, it is a member of the class football game only. In the ontology, an individual can be member of several classes simultaneously, since the classes are built from the individuals through the attributes each individual exemplifies.

Event

/event/ is a compound concept and has four parts. Firstly, the /event core/. Secondly, /event relating/. Thirdly, /event operations/. Finally, /event generalisation/. Table 31 shows the results from comparing Chisholm's ontology with OMT for /event/. We discuss each part in turn.

Table 31—Expanded Results for /Event/

Part of /Event/	Chisholm's Ontology
Core	$\surd_p$
Relating	$\surd$
Operations	$\surd_p$
Generalisation	$\surd_p$

Chisholm covers a semantic field that is similar to that for «event core». However, Chisholm does not support OMT's concept to its fullest generality. Recall, that in OMT, events are unique, instantaneous, and require time as an attribute. Chisholm's ontology covers events that are also unique, and may be instantaneous, but in contrast, the ontology does not require time to be part of an event. It is for this reason, that support is not to the fullest generality.

The semantic field defined by «event relating» concerns mechanisms by which one can order events. OMT relates events where one event can cause another event to occur. In this way the two events are causally related. Event in this context is a generalised event encapsulated in state/event diagrams. This «event relating» is covered in Chisholm's ontology through «state causal relationships». Further, these give rise to a special class of event in the ontology that instantiates the law, giving rise to a new state. It is for these reasons that we can say that the semantic field defined by «event relating» is covered in Chisholm's ontology to the fullest generality.

Chisholm's ontology does not directly cover the semantic field defined by «event operations». However, upon closer inspection, one sees that operations are an extension of event and change in state, in that operations are triggered either by events or by changes in state for an object. A change in state for an object may be an effect of an event influencing that object in any case. Operations triggered by events are actions and are instantaneous. These may send events or change state of the object concerned. Events sent in this manner may affect either the originating object (of the event in the first place) or a third object. Those operations triggered by changes in state are called activities and endure. Further, they endure until either a specified period has elapsed, or the object concerned has changed state.

In the sense described, /event operations/ is related to /event/ in Chisholm's ontology, and the role of /event operations/ in the dynamic model is similar to one event contributing causally to another in that an event may cause a change in state, or trigger an action in OMT that contributes causally to another event happening. By inference, an event can effectively endure, by causing a change in state that then triggers an activity for a period and this is covered by «event core» in the ontology. There is, however, no elegant unity between event and state in OMT. And so, the parallels cease upon deeper inspection. However, it can be said that Chisholm's ontology partly covers the semantic area defined by «event operation».

There is no coverage of «event generalisation» in Chisholm's ontology. Firstly, the structure of state (and therefore event) is quite different in the ontology when compared with event generalisation. In OMT the generalisation of this form once again forces a hierarchy of convenience that helps implementation but does not aid modelling of reality. In summary, there is no coverage of this concept in the ontology.

State

/state/ has three parts. The parts firstly, /state core/, secondly, /state transition/, and thirdly, /state generalisation/. Table 32 shows the results for this concept.

Table 32—Expanded Results for /State/

Part of /State/	Chisholm's Ontology
Core	√
Transition	X
Generalisation	X

The semantic field defined by «state core» is completely covered by various concepts in the ontology. /state core/ in OMT discusses objects being in particular states according to the values of indicative attributes. The semantic field for /state core/ in the ontology by /state core/. State describes the values of links and attributes of an object.

The relationship between event and state is covered by /state transition/ in OMT, where the effect of an event on an object in a specific state can be specified by showing transitions to other states. The ontology does not cover the same semantic field as «state transition» because it does not discuss the mechanisms by which changes in state occur, beyond those generated by causal laws.

Recall that state generalisation is where the state of an object is entirely one of its substates. A substate may inherit transitions from the state of the object. This is a kind of 'or' relationship between states of an object, in that the object will be in one of the states specified. The concept of generalisation of states from OMT does not have an equivalent in Chisholm's ontology. In the ontology, different attributes would be used in this situation. The concept of state generalisation as specified in OMT is not covered in the ontology.

Concurrency

/concurrency/ in OMT is another compound concept and has three parts. Firstly, /concurrency core/. Secondly, /concurrency aggregation/. Thirdly, /concurrency within object/. Table 33 shows the results for this concept. We discuss each of the elements of vocabulary below.

Table 33—Expanded Results for /Concurrency/

Part of /Concurrency/	Chisholm's Ontology
Core	√
Aggregation	√
Within object	X

Chisholm's ontology does not contain an 'execution' or processing system as such. Implied in its description, however, is the realisation that state can change for an individual in an independent way from another individual, in that the state of an individual can be independent from other individuals. The parts of an individual can similarly be independent. It is for this reason that Chisholm's ontology covers the semantic field described by «concurrency core» to the fullest extent.

Secondly, «concurrency aggregation» in OMT concerns the independent nature of the parts of an object as specified by an aggregation association. Chisholm explicitly supports such a notion, in that the state of each part of an individual can be considered to be an independent state, because each individual can be thought of as being independent, and each part can be affected independently since each is an individual.

In contrast, there is no coverage in Chisholm's ontology for «concurrency within object».

8·4 Consolidation of Results and Implications

We have conducted a bidirectional comparison of OMT with Chisholm's ontology in accordance with the method found in Chapter 3, and in this section we consolidate the results from the previous sections.

We show the consolidated results in Table 34. Using a three-level indicator that is consistent with the method it shows the level to which the semantic fields as specified by the concepts in the ontology are covered by OMT. This consolidation shows the relationship between the world view as encapsulated by the ontology and that found in OMT in tabular form.

The gaps in coverage in OMT fall into two groups. Firstly, concerning the static concepts from Chisholm's ontology. Secondly, the dynamic concepts from Chisholm's ontology. The latter is likely to be quite complex to analyse given the complexity with which these concepts are developed in OMT.

Table 34—Consolidated Results Showing Levels of Coverage in OMT for Concepts from Chisholm's Ontology

Ontological Concept	Level of Coverage
Individual	full
Attribute	partial
Classification	partial
Relation	full
Temporal relations	partial
State	partial

We firstly examine the static elements /individual/, /attribute/, /classification/, and /relationship/. From Table 34 we can see that OMT provides full coverage for «individual», and «relation». However, it has partial coverage of «attribute», and «classification».

The ease with which extensions can be made to cover these areas will depend upon the orthogonality between the required semantics and those already provided by OMT. No or partial coverage implies that there are gaps in the coverage of semantic fields described by concepts from the ontology. If there are concepts in OMT that are antonymous in a contradictory or contrary way with respect to the semantic area of certain gaps, then extension of OMT will be difficult. If the antimony is one of a complementary nature, then extension will be easier. If there is no relationship between the semantic area of the gap and concepts in OMT, then extension is likely to be straightforward. However, this is also likely to be affected by the conceptual complexity of OMT.

«attribute core» and «attribute equivalence» contain the 'gaps' in semantic coverage by OMT of /attribute/. Both of these have gaps that define a semantic area that contradicts a semantic area outlined in OMT. Firstly, attributes are tightly coupled with objects in OMT. This makes provision of loose coupling very difficult. Secondly, at this stage, there is no provision for «attribute equivalence» in OMT. Further, with the object-centric only attitude of OMT, it will be difficult to provide such semantic equivalence.

«classification» covers a semantic area that is contradicted by concepts in OMT. This will also be an area where extension will be quite difficult. This difference stems from OMT's rigid class structures. Some object models have a more flexible attitude toward class structures. We will discuss some options in the next section.

OMT's dynamics are expected to be difficult to augment in order to bring the concepts covered closer to that which is intended by the ontology. One obstacle is the requirement in OMT for an arbitrary measure of time to record the chronological ordering for events and state. Once the arbirary measure of time has been removed the descriptions of event and state the unification evident in the ontology can begin to be provided in OMT. The second difficulty in extending OMT emerges from the generally fragmented way in which state and event are supported. It is for this reason that the unification of state and event as described in Chisholm's ontology is difficult to provide in OMT and is the subject for future research.

8·5 The Extensibility of OMT

In a similar way to the previous chapter where we examined FDM, we can say from the outset that it is this 'forward' comparison of OMT with Chisholm's ontology that yields the majority of the information required for understanding and analysing the relationship between the ontology and OMT.

Our work has revealed two areas of major difference. Firstly, we have found that the use of rigid class hierarchies in OMT is not supported in Chisholm's ontology. That is, /classification/ does not require individuals be within one class that in turn forms part of a generalisation/specialisation hierarchy. This has ramifications for static modelling and for model change. In OMT's object model, class hierarchies of objects representing generalisation are specified early. Further, attributes are tightly bound to objects through the classes to which any specific object is a member. Moreover, objects cannot be members of more than one class, and cannot move between classes. When the class structure inevitably changes, there may be significant problems.

Secondly, OMT's dynamic model has a similar semantic coverage of «state». However, its support for the various parts of /state/ is such that it fails to achieve the advantages of unifying /state event/ with /state/. That is, /state event/ is a specialisation of /state/ and utilises similar structures. The fullest ramifications of unity will be subject to further research.

We have identified a number of specific areas in OMT's modelling tools that must be addressed in order to bring the world view of OMT closer to that of the ontology. For these areas, there are a number of specific extensions or augmentations of OMT that will help to bring its world view closer to that of the ontology.

There are two areas into which deficiencies fall. Each must be further analysed and discussed in order to determine the ease with which one can extend or otherwise strengthen OMT so that it is closer to Chisholm's ontology in terms of emphasis and features.

Firstly, the rigid class hierarchy in OMT is not compatible with Chisholm's ontology. The rigid class system, where objects cannot be member of several classes presents a significant difficulty with respect to Chisholm's ontology. We also found that, contrary to the ontology, in OMT attributes must be tightly bound to classes. Some object models support loose coupling; OMT's object model is not one of these object models.

Secondly, the nature of dynamic modelling in OMT presents further difficulties with respect to Chisholm's ontology. We have found that this may be due to the implementation focus of dynamic modelling in OMT. That is, there is no clear separation between modelling and implementation. This is a significant point that requires much deeper discussion.

8·6 Reflections

The application of this method with OMT and Chisholm's ontology, yields some critical insights into class hierarchies as they are traditionally envisaged in object modelling. Further, it may be that the influence of implementation evident in object modelling generally has affected this aspect of OMT. We believe that further research is crucial into the application of these ideas to practice. That is, discovering if Chisholm's ontology's view of classes is reflected in industrial situations.

We have also found that the application of the method to OMT is more difficult than with a traditional data modelling framework. This is because OMT has a complex and fragmented method of expressing dynamics.

CHAPTER 9

Conclusions

WE COMMENCED THIS printed work with a statement of thesis and an outline of our research philosophy. We argued that ontology, of the style we defined, is useful as the basis for a unifying framework for data modelling frameworks. This thesis represents the first step in finding a unifying framework for data modelling frameworks. In Chapter 2 we related our study to the literature in the sense of indicating how our research is placed with respect to past work. Chapter 3 saw the introduction of qualitative methods of comparison and evaluation that use terms and concepts as a basis. After describing an ontology by Roderick Chisholm, that we selected using pragmatic criteria, we described the data modelling frameworks used in the study. We then applied the methods the results of which we reported in Chapters 6, 7, and 8, and we reflected on the results pertaining to the individual methods as they applied to the subjects under study in those chapters. In this chapter we draw the thesis to a conclusion in that we critically analyse our study and form propositions in accordance with our research philosophy outlined in Chapter 1.

We begin this chapter by critically reviewing our achievements in this thesis. We then reflect upon our results and draw conclusions of a general nature about data modelling frameworks and we reflect upon our research questions. We expect to gain knowledge about the use of an ontology of the nature of the one selected with data modelling frameworks in the way we have argued, and we can conclude specific insights, based upon the selected ontology, about the data modelling frameworks we have studied. We also form propositions, some of which require further investigation, concerning the way that modelling could be handled in practice by analysts in industry. We draw limited generalisations about concepts that a modelling framework must support that are expressed in various terms found in the modelling framework.

Towards the end of the chapter we return to our research questions and comment on our progress towards answering them. We then highlight key constraints and limitations of the research before looking forward to opportunities for logical extensions to this work. We conclude the thesis with some closing remarks in which we discuss the potential wider implications of this ontology for the theoretical basis of information systems.

9·1 A Critical Analysis of Our Study

Our long-term aim is to find a unifying framework for data modelling frameworks. This thesis represents a first step towards that long-term goal. Our major contribution in this thesis is the proposal of one such framework based on Chisholm's ontology and a demonstration of its use in comparing and evaluating data modelling frameworks.

The comparative and evaluative methods introduced and used in this thesis are novel in that they recognise that data modelling frameworks have ontological terms of a qualitative nature and that these terms can be used as a basis for relating the modelling frameworks to an appropriately selected ontology. The methods also recognise that the tools under investigation are not always formalised but are often semi-formal by nature. In the next section of this concluding chapter we will discuss our findings and revisit our research questions. In this section we critically analyse the various parts of our study and comment on the capabilities of the ontology and of the methods and we suggest ways by which the process can be improved.

The two major aspects to the study that require critical analysis are the selection of the ontology and the methods themselves.

The Ontology

The methods we have proposed rely greatly on the terms and concepts of the selected ontology. The nature and depth of the findings are thus limited by the nature of the terms and concepts found in the ontology. The ontology that we have selected using pragmatic criteria, is by Roderick Chisholm. It is a realistic ontology that uses individual substance and attribute as its core. It also discusses the terms class and set, relation, and state. Its economy of expression means that event and state can be unified representing a shift with respect to other recent commonsense realistic positions. This unification of event and state, because it departs from related realisms limits the conclusions and it is an important avenue for further research. We have found that because of the quality of terms and concepts found in the ontology it was very useful in analysing the data modelling frameworks and that there is a good degree of overlap between the five data modelling frameworks that we have studied and Chisholm's ontology. The overlap is not total but does cover the core parts of terms and concepts from the ontology. We discuss generalisations from our results in the next section.

We have found that Chisholm's attitude to event and state to be difficult to use in practice but that this is not serious enough to warrant rejecting the ontology for future use. The dynamic terms of the ontology are difficult to use not due to their incompatability with the data modelling frameworks but due to their simplicity and power compared to the complex and fragmented nature of dynamics in many modelling frameworks. This offers scope for rationalising and drawing together the diversity of dynamic concepts in existing data modelling frameworks. While Chisholm's ontology has been useful in our study, we leave an open invitation for others to apply different ontologies in ways similar to those presented here.

We have found the ontology to be at an appropriate level of abstraction so that we can successfully analyse and discuss the data modelling frameworks in this study.

However, there is no fundamental provision within the ontology for constraints as seen in data modelling frameworks such as OMT and ER, such as limits on participation in relationships or associations and similar concepts. The attribute and individual centred ontology has less need for such restrictions since there is no class binding individuals in the ontology.

Concluding, we can say that there is evidence of significant overlap with data modelling frameworks in that the terms and concepts in Chisholm's ontology are present in the data modelling frameworks studied. Further, the terms and concepts in the data modelling frameworks can be mapped into those of Chisholm's ontology and so the ontology does present as a suitable unifying framework. We have found that further and deeper studies are required to fully examine Chisholm's dynamics. Chisholm's statics are a good foundation for a unifying framework and we had comparatively less difficulty in drawing conclusions concerning the statics of Chisholm's ontology.

The Methods

The methods developed as part of this study represent a significant advance over previous approaches in that they recognise the important role played by terms and concepts in qualitatively comparing data modelling frameworks and ontologies. We have argued that studies of data modelling frameworks can be undertaken using ontological terms and concepts. In the methods we analyse the relationship between terms in a data modelling framework and an ontology using a restricted part of a semiotic theory. By using semiotic theory in this way we highlight the importance of meaning of terms in the ontology and data modelling frameworks. This in itself is an important advance over past approaches.

Each of the methods in this thesis is based upon a comparative core in which terms and concepts are used to explain the relationship between subjects that have ontological elements. Comparisons of this nature are limited by the terms and concepts obtained from the ontology and to a lesser extent from the data modelling frameworks. For many concepts and terms the comparison is relatively straightforward. For example, /individual/ is a term from the ontology that compared easily with all of the data modelling frameworks that we studied. In contrast, our study of /event/ from the ontology was difficult, because the fullest ramifications of the simplicity and power of the concept from Chisholm's ontology could not easily be gauged against the fragmented nature of OMT, the framework that we studied with a dynamic part.

An important contribution of the methods proposed and used in this thesis, is the recognition of the roles that terms and their associated concepts can play in ontological studies. The ontological comparison of data modelling frameworks explicitly requires terms and associated concepts from the ontology in order for the comparison to proceed. The evaluative method requires that terms and concepts be extracted from both the subjects of the method, that is, for a selected data modelling framework and for a selected ontology. It is therefore important that we review the process involved in extracting terms and concepts, and, where applicable to suggest alternatives given our experience.

The data modelling frameworks and the ontology are analysed into constituent terms and concepts. Using important primary sources, the researcher establishes appropriate terms and concepts for the ontology and data modelling frameworks. The definitions we use for concept and term are not controversial and we utilised recognised philosophical sources in reaching them.

Concept formation is subjective in that, despite care being taken by researchers, two different people are unlikely to agree in every aspect of interpretation. However, such concept formation is fundamental to philosophical and qualitative work.

It would be interesting to investigate alternative ways of concept formation. For example, Roderick Chisholm, the author of the ontology, is still alive and productive at Brown University, albeit in his role as Professor Emeritus. Similarly, those who originally designed the modelling frameworks are also likely still to be alive. It would be interesting to interview these people so that a more direct avenue to the terms and concepts may be found. The major drawback to this approach is that neither group of people have an understanding of the aims and objectives of the research, but would, of course be informed of the goals and directions of the research prior to interview. It would be particularly interesting if one had the opportunity to interview the author of the ontology, Roderick Chisholm since, by far, his work is the most comprehensive and dense contrasting sharply with the sources for the data modelling frameworks.

In our discussion of our method, the word 'compare' means two different things. Firstly, pairwise comparison of terms and concepts from the ontology with each of the data modelling frameworks. The second meaning is in describing the method overall wherein we compare a number of data modelling frameworks by utilising the ontology as a benchmark, after the individual comparisons have been completed. This second meaning is reflective of the role the ontology may play as a unifying framework. Recall that the comparative method is in two parts. Firstly, one takes the terms and concepts from the ontology and uses these in pairwise comparisons with each data modelling framework to find the degree to which each of the selected data modelling frameworks have terms that cover a similar semantic area as those from the ontology. An indicator is used to describe the similarities and differences in the meanings of related terms, that is based on a well-reasoned and comprehensive theory developed by Umberto Eco. Eco developed his theory through an examination of semiotics across a variety of disciplines. We utilise part of Eco's theory of codes in that we consider the degree to which terms in the respective data modelling frameworks cover the same semantic area as is covered by terms from the ontology. A discussion of the precise similarities and differences is also recorded. The use of semiotics as a mechanism to explain and interpret our results is an important advance in that it takes us from a structural comparison to one where the semantics of terms can be explored.

Secondly, the comparison of data modelling frameworks, as opposed to the comparison of an individual modelling framework with concepts from an ontology, is performed in an indirect manner. That is, by inspecting a set of consolidated results one can see which candidates are the ones that best fit the ontology. Further, one can see on a term by term basis the degree to which each data modelling framework fits the ideal of the ontology. That is, we use the ontology as a benchmark.

There are two parts to this method that need critical analysis. Firstly, the pairwise comparison of the ontology's terms with each data modelling framework. Secondly, the role of the ontology as a 'benchmark' or unifying framework against which each modelling framework is examined and therefore the basis for the comparison.

The comparison of the ontology with each of the modelling frameworks requires a deep understanding of the ontology and of each data modelling framework. Further, it is on the basis of the terms described in the ontology that the comparison with each candidate modelling framework is undertaken. It takes time to understand an ontology as dense as Chisholm's and to relate each term and its associated concept to the modelling frameworks. However, we cannot conceive of an easier way in which to undertake such a work. Further, the results of the comparison may vary to a degree between researchers. This, however, is the nature of the type of research and is not of itself grounds enough to doubt its applicability. There is no easy way to perform this comparison, and there is no algorithmic comparison of concepts. However, given the types of research questions that we pursue in the research, our approach is quite appropriate. Further, the approach has the potential to greatly enhance our understanding of the nature of data modelling frameworks.

The role of the ontology as a benchmark is not really part of our critical analysis, in that one objective of this research is to discover whether the ontology is a suitable candidate as a unifying framework for data modelling frameworks. Further, we argued that if this were to be the case then the ontology should be able to be used to discuss terms and concepts from each data modelling framework. Comparison is one such purpose to which a unifying framework can be put. Thus, the role of the ontology as a benchmark is not part of the method requiring analysis, but instead is a manifestation of our original goals and objectives and needs to be further discussed when we reconsider our research questions in the light of our results.

The evaluative method is similar to the comparative method in that it explicitly involves extraction of terms and concepts from an ontology and, additionally, from a selected data modelling framework, followed by a bidirectional comparison between an ontology and a specific data modelling framework. It is bidirectional in the sense that we firstly compare a data modelling framework with an ontology and secondly, compare the ontology with a data modelling framework. The nature of each comparison is similar to each pairwise comparison in the comparative method. There are three major differences with respect to the comparative method. Firstly, all of the terms and concepts from the ontology are used (concepts concerning dynamics in addition to statics). Secondly, the comparison in this method is bidirectional. Thirdly, analysis is more extensive and of a different nature due to the amount of information resulting from the bidirectional comparison.

Firstly, we need to extract dynamic parts of the ontology. There is a specific issue in the candidate ontology, in that the selected ontology has a particularly terse nature to the description of state and event. There is an elegance in the unification of state and event into closely related ideas combined into the single concept, where event is a special case of state, and represents a significant change with respect to accepted realistic thinking, and we expect this aspect to have a significant impact on, and contribution to, information systems. The fullest implications of the concept are difficult to fathom, and we found that this section of the ontology was very difficult to conceptualise. We have taken our understanding to a depth sufficient for an evaluation of the modelling frameworks that we have selected, but more work is needed to gauge the full ramifications of the concept. Further, this part of the ontology was difficult to apply in the method. It was further complicated by the semi-formal nature of dynamic concepts in OMT. Later in the chapter, we highlight further research into the dynamic parts of Chisholm's ontology beyond that which is required in this thesis.

Apart from involving a more complete set of concepts from the ontology when compared with the comparative method, the evaluative method is a bidirectional comparison followed by an analysis of the areas in which 'gaps' are perceived in the coverage offered by the modelling framework in the semantic area covered by the ontology.

It is through using a method that a critical analysis is possible. We have applied the evaluative method with two quite different data modelling frameworks. Firstly, with FDM, a very simple and clean modelling framework; Secondly, with OMT, a popular and rich object modelling framework. Further, we stated in Chapter 3 that the bidirectional comparison in the evaluative method is useful because of the need to judge whether the semantic field covered by concepts from the ontology ran counter to concepts from the data modelling framework under study, that is, where the semantic area of a concept from the ontology is not covered by a term or terms from the modelling framework, we need to know if there is a concept from the data modelling framework that covers a related semantic area that is contrary or contradictory with respect to the concept from the ontology.

Using the bidirectional comparison with FDM was straightforward. Further, it yielded some insight into the nature of the coverage of semantic areas. However, and this is likely to be due to its simple construction, little of substance was learned by the 'reverse' comparison. However, this may be indicative of the ease with which extensions to FDM may be made to bring it closer to the ontology.

By contrast, the 'reverse' comparison was quite useful in the evaluation of OMT. In this application we found some extra qualitative information that added to evidence uncovered in the 'forward' comparison. Although it is pure conjecture, we believe that due to the complex nature of OMT, particularly with respect to dynamic issues, combined with the terse nature of Chisholm's attitude to state and event that resulted in the reverse comparison adding to the information that became evident from the 'forward' comparison. We found it was only through the 'reverse' comparison we were able to evaluate aspects of the dynamics in OMT that were difficult to judge in the 'forward' comparison, that is, the 'reverse' comparison enabled us to better analyse OMT to conclude that its coverage of concepts from the ontology was adequate as far as it went. Further research is needed to investigate whether the full nature of event and state in the ontology can be replicated with OMT.

Having applied this part of the method with different styles of modelling framework, we found the 'reverse' comparison was useful for a complex modelling framework such as OMT, where more subtle elements of the modelling framework can be examined with respect to the ontology. We also found that it was of limited value for a simpler modelling framework such as FDM, since the reverse comparison added little information to that found through the 'forward' comparison.

The consolidation of results was easily achieved. In both cases, the forward comparison yielded the bulk of the information, with the reverse comparison adding detail to the extensibility of the frameworks under study and the ease with which extensions are likely to be possible. Once again, the utilisation of semiotic theory in explaining the ease with which extension of a data modelling framework is likely is an important contribution.

Summary

We can conclude from applying these methods the approximate degree of overlap in concepts between the ontology and data modelling frameworks. Using the degree of overlap for core concepts indicates a more narrowly defined measure of overlap that may be a more indicative measure of ontological support in a modelling framework.

Summarising, we have introduced two methods that use terms and concepts as a basis for a qualitative comparison. By introducing these methods we are expanding the types of ontology that can be used in these types of studies and we are recognising that many tools used in information systems are semi-formal. The use of Chisholm's ontology has been successful with a degree of difficulty encountered with the dynamics of the ontology. We expect this to be an important part of future research.

We have also used an ontology that largely has been successfully applied using five data modelling frameworks. We have also found that the methods have helped to contribute insights into the data modelling frameworks we have studied.

9·2 Our Findings

In this section we reflect upon our results and draw conclusions in the form of propositions as outlined in our research philosophy that represent our findings.

We have compared ER, FDM, SDM, NIAM, and OMT with Chisholm's ontology. Two of these, FDM and OMT, have been evaluated in detail. We have found a large degree of similarity between these modelling frameworks and the ontology. We can see five areas where the ontology departs from the modelling frameworks.

Firstly, in the ontology, attributes are considered to be quite separate from the individuals that exemplify them, and attributes also endure. One can describe this as loose coupling between individuals and attributes. In contrast, some of the data modelling frameworks group individuals into homogeneous classes and by so doing restrict the range of attributes that each individual can exhibit.

Secondly, in data modelling frameworks, the equivalent of individuals are seldom allowed to be members of several classes simultaneously. The only exception to this is in cases where class hierarchies around a common family of classes is established or where individuals are distributed or dispersed across several classes. This is different from the ontology in which individuals can be member of more than one class simultaneously because classes are established based upon attributes that member individuals exhibit. In other words, the classes are not repositories for individuals, they filter individuals.

Thirdly, OMT, like many object modelling frameworks, establishes a rigid hierarchy of classes based upon inheritance of both behaviour and state. There is no support for rigidly defined class hierarchies of this nature. This does not mean that class hierarchies of this nature are not appropriate in implementation, or that such hierarchies are not necessarily found in reality. This aspect of object modelling may be important in implementation due to efficiency requirements. However, the ontological studies we have performed reject them as an important element in making sense of 'what there is'. Where these hierarchies are seen to be important, they instead can be synthesised based upon careful selection of attributes that determine relevant classifications, and by selecting classes using specific groups of attributes, a similar effect to that found in many object modelling frameworks is possible to achieve. However, in modelling reality, the class hierarchy is not critical. Instead, it is the identification of the various classes, some of which may form part of a class hierarchy, that is important.

Fourthly, in contrast to class hierarchies there is support for part-whole or aggregation structures in the ontology, known in philosophy as mereology. This also supports related research that uses a different ontology. Many of the modelling frameworks supported this to at least a crude degree, such as with the ER modelling framework. Others had more sophisticated methods of providing this, such as with most object modelling frameworks or with SDM.

Fifthly, OMT, the only model studied with a dynamic model, provides significant semantic coverage for the concepts of event and state. However, the ontological terseness with respect to event and state, where event is a specialisation of state is not replicated in OMT. Chisholm is noted for his departure from the realistic norm in this aspect of his ontology. The simplicity of the concept that results from the unity of event and state is something that may help to clarify the role of the various dynamic process-oriented techniques that are popular today. It may also aid the modelling process in hybrid systems where analog mechanical, and digital systems interact with the environment, in that, Chisholm may provide a high level at which to reason about state and event that in turn can be mapped to a much lower level and modelled using either another ontology of a much more fine-grained nature or by a relevant modelling scheme.

The conclusions of a pragmatic nature that we have raised here require investigation in practical situations in order to clarify their role and relative importance. Of specific interest to us is the nature and role of classification and class hierarchies in practical situations seen by clients rather than information systems professionals.

Our long-term goal is to establish a fully elaborated unifying framework for data modelling frameworks based on ontology, and this research is the first step towards it. We set out to investigate two research questions that consider the role of ontology as a unifying framework. 'How suitable is a specific ontology as a unifying framework for data modelling frameworks?', and 'Is this specific ontology tacit in data modelling frameworks?'

The suitability of the ontology as a unifying framework, assuming that the ontology accords with our pragmatic criteria, is judged by examining the results of our methods and comment on the results of the comparison and of the evaluations. We can also examine the desirable features that our unifying framework must possess.

Recall that in Chapter 1 we indicated that a unifying framework should fulfil three criteria. That the unifying framework have a consistent set of concepts; That each data modelling framework can be expressed within the framework; That the framework can be used to compare, evaluate, extend, and understand the world view of each data modelling framework.

We have used Chisholm's ontology in methods with what we claim to be representative data modelling frameworks. We can say that the ontology does consist of a set of consistent concepts. These will be used as the basis for a unifying framework. We have found that each data modelling framework can be related to terms and concepts from the ontology and so can be expressed in the language of the unifying framework. We also found that even where a modelling framework appeared to be structurally different from the ontology it could still be compared to the ontology and a significant degree of overlap with it was found. Finally, we have successfully applied methods that use the ontology as a benchmark by which to compare a number of data modelling frameworks and to evaluate specific data modelling frameworks. For these reasons, we have reason to believe that Chisholm's ontology is a suitable candidate for a unifying framework. Further research is required to examine its role more deeply and to deepen our understanding and application of it. We also need to examine specific differences that have been observed in all data modelling frameworks. For example the lacking in support of /attribute equivalence/ by all data modelling frameworks.

The second research question concerns the ontology implicit in the studied data modelling frameworks. We have found that there is a significant degree of overlap between the ontology and data modelling frameworks studied, but that the overlap is not total. In order to answer this second research question, we must examine the philosophical nature of the ontology that we have used, and the specific areas in which overlap between the ontology and modelling frameworks is not total. We also need to discuss the implications for the observed lacking in total coverage.

As we noted in Chapter 4, Chisholm's ontology is one of commonsense realism and is categorised by Chisholm as being a realistic ontology. We have found, in our reading (Audi, 1995, Flew, 1989, Honderich, 1995), that the terms used by Chisholm are reasonably widely supported, and the style of realism upon which his ontology is based has a high degree of consensus between philosophers with respect to what categories of 'entities' exist in the world (Milton and Kazmierczak, 2000). We are confident that this core of consensus, which is considerable, forms a good starting point from which to progress towards a detailed and consistent unifying framework. We have found the following form such a consensus.

That

- there are individuals that are ontologically independent
- individuals exhibit attributes
- individuals may consist of parts that are each individuals (known as mereology)
- individuals relate to each other
- individuals congregate into classes that are selected based upon the attributes exhibited by members of the class
- events represent changes in state

More centrally, there is a realistic core that forms a subset that can be summarised in the following proposition. At this stage we cannot be certain that event and state as Chisholm depicts it has wide support in the realistic philosophical literature, and it is for this reason that we exclude it from the proposition.

Proposition A

"That a unifying framework includes the following terms

1) Individuals that are ontologically independent
2) Attributes that are exemplified by individuals
3) Part-whole relationships between individuals
4) Relations between individuals
5) Classes of individuals selected on the attributes exemplified"

In addition to these terms, Chisholm's ontology departs from a realistic core in specific ways, the unification of state and event being the most radical. Essentially these departures stem from the ontological commitments that are made by Chisholm.

> "In metaphysics Chisholm is open and precise about his ontological commitments, which, according to his most recent writings, include only attributes and individual things" (Honderich, 1995)

By using attributes exhibited by individuals to provide for terms such as relations, classes, and state, Chisholm is departing from mainstream realism. It would be interesting to see if this view of 'what there is' is reflected by clients.

Proposition B (to be tested)

> "That attribute based, rather than class based, data modelling frameworks more naturally accommodate and reflect client views of 'what there is'"

The unification of state and event also flows from this terseness in categories by enabling temporally oriented attributes to replace the need for time as a measure in events and so to enable event to be classified as a special case of state. Temporal ordering of attributes allows for a chronology to be constructed.

Proposition C (to be tested)

> "That temporally oriented attributes together with the unification of state and event more naturally represents the view of dynamics held by clients than current approaches in information systems modelling"

Despite these areas of departure, we have found that Chisholm's ontology is a sound basis for a unifying framework. The methods we have applied in this research are based on semantic comparisons between terms prevalent in each of the data modelling frameworks and the ontology. Combining observation with method, it is therefore logical to draw the conclusion in answer to our research question that the five data modelling frameworks studied in this thesis have a tacit ontology that concords with Chisholm's ontology, and because Chisholm's ontology adheres to commonsense realism and our methods show a high degree of semantic overlap with concepts from the ontology, that the modelling frameworks we have studied have a tacit ontology of commonsense realism.

Proposition D

> "That the modelling frameworks we have studied show an ontology of commonsense realism"

Further we conjecture that other data modelling frameworks will also exhibit a tacit ontology of commonsense realism.

Proposition E (to be tested)

> "That data modelling frameworks beyond those we have studied also exhibit a tacit ontology of commonsense realism"

We can also posit that it is highly unlikely that data modelling frameworks have a tacit ontology of idealism. We can see where commonsense realism fits with other styles of realism and indeed with idealism:

> "Common-sense realism is a special brand of realism in general, or in other words of the view that the world exists independently of our cognitive relations to it. Realisms differ in their accounts of what the world is that thus exists independently. They are opposed to idealisms of the various sorts, for example to linguistic idealism, a view to the effect that the world exists (or has the structure which it has) in virtue of the language we use to speak about it. Subjective idealists hold similarly that the world exists in virtue of our mental activity. Note that all idealists—to the extent that they embrace a position that is capable of being coherently formulated—are in fact also realists of a certain stripe, since all idealists hold that there is something (be it mind, language, or Absolute Spirit) which enjoys autonomous existence." (Smith, 1995)

We can then, in turn, conclude the following:

Proposition F

> "That the modelling frameworks we have studied are unlikely to show an ontology of idealism"

9·3 Limitations and Constraints

There are four significant limitations and constraints that can be placed on the work we have done. Firstly, the methods are such as to make it difficult to wholly replicate the results in fine detail and in this sense are subjective. Secondly, they are time consuming to execute. Thirdly, our conclusions are restricted to one class of ontology and that the ontology may not be the only one that fits the modelling frameworks studied. Fourthly, we are limited to conclusions about the data modelling frameworks under study. We discuss each of these in turn.

It is difficult to replicate the results that we have established here. The methods involving the data modelling frameworks and the ontology are predicated upon a deep understanding of the concepts behind terms used in each, and reflect the descriptive nature of many data modelling frameworks and ontologies. Each candidate is described in recognised original sources, and in this sense, the sources can be revisited by others. However, despite care being taken in extracting concepts, there are likely to be differences between our interpretation of some subtleties when compared with others. This is only a serious problem in this style of research if we present conclusions that rely upon sections of the study in which interpretation is likely to be quite different. However, we have found that the nature of the style of ontology involved is such that there are some limited generalisations that have become quite obvious after applying the methods, and we have already covered them in this chapter. Further, there are clear places where the ontology departs from accepted realistic stances. We have also noted these in this chapter.

The second limitation concerns the fact that the methods require intense reflection upon the candidates, and it is for this reason that the methods are time consuming to apply. The ontology, as is characteristic for much of Roderick Chisholm's writing, is extremely terse and replete with meaning making it a difficult subject to manage. We are confident that our major conclusions faithfully reflect the core focus of the ontology. There are sure to be subtleties of the ontology that we have not fully appreciated. However, at some stage the process of reflection and investigation has to cease and findings made. It will be at the further stage of formulating a unifying framework for data modelling frameworks that the fullest subtleties of the ontology will need to be appreciated.

Chisholm's ontology is an example of a realistic ontology that comes from critical commonsensism or commonsense realism. Based upon this heritage, there are some things that he holds in common with others. There are key areas wherein he departs from orthodoxy. We have discussed these in our contribution towards a unifying framework (in the previous section). Therefore, our contribution is constrained by the class of metaphysics from which the ontology comes, and we cannot definitively state that this ontology is the only one supported by the modelling frameworks studied. This is true for much of the work involving ontologies of any nature. However, we will discuss how this research fits within the context of other ontological studies involving scientific realistic ontologies.

Finally, our findings with respect to the research questions under investigation can only go so far as to conclude about the data modelling frameworks that we investigate, and our findings are limited to the modelling frameworks studied. We have, however, claimed that the group of five are generally representative of data modelling frameworks.

In conclusion, our contribution needs to be read in the light of these constraints. Due to the nature of the candidates involved, the constraints we have covered cannot be easily avoided. Our conclusions and contribution have been considered at length and we have been at pains to ensure that we do not make sweeping statements where there is little by way of evidence to support them.

Like many qualitative studies, a high degree of trust must be given to the researchers. However, since this is not a study in a social setting, the qualitative methods can result in generalisations that are normally much more difficult to achieve when studying information systems social settings. This is because our subjects are not people or groups undergoing the constant influences of society that may be localised, but instead our subjects are contained in recognised sources and, in the case of the data modelling frameworks, are accepted widely in information systems.

9·4 Further Research

Our research provides many avenues for further research. Some future work consists of logical extensions of the theoretical work. Others are applications of theory to practical settings because it is only through observation in industry can we further develop the practically applicable aspects of the work.

Most critically, work needs to commence in using the ontology to create a detailed unifying framework. In so doing, it is important to explicitly note those parts of the ontology that depart from accepted forms of realism. The most noted departure is the unification of state with event. The full implications of this on modelling state and changes in state is in need of further research. Similarly, conceptual entailment of attributes also needs further consideration. It is possible that each of these will provide an interesting extension to many modelling frameworks.

We are also keen to investigate the reception aspects of classification scheme in Chisholm's ontology by practitioners, and equally importantly clients, in industry. This could involve a series of surveys together with more detailed focus group research. We also hope to apply the theory in practical settings perhaps using action research.

There is also scope for further analytical studies. Apart from studying more modelling frameworks, it would be interesting to widen investigation to include modelling systems that better integrate dynamics and statics than does OMT, and we have in mind dynamic modelling systems such as Petri nets. These systems fully integrate state and transitions in state. Recently, Petri Nets have been extended to include object oriented features, and they are being used in information systems modelling (Keen and Lakos, 1994, Keen and Lakos, 1993, Keen and Lakos, 1995). Coupled with an underlying mathematical rigour through transformation to coloured and then to black and white Petri nets, object Petri nets represent an elegant process modelling framework. There is also a high degree of unity between the static and dynamic features of the object Petri net modelling system. The ontological analysis of object Petri nets may be more appropriately performed using Bunge's ontology, another realistic ontology that has been applied in related research.

9·5 Closing Remarks and Summary

This thesis represents an enrichment of the ontological studies in information systems in two important ways. We have introduced and applied new qualitative methods that recognise the importance of concepts in ontological studies concerning modelling in information systems. As a direct consequence of our qualitative methods we have opened consideration of ontologies to include those that are more descriptive and less mathematically formalised. The methods explicitly recognise the semi-formal nature of many of the tools used in modelling in information systems.

We have applied the methods using a number of representative data modelling frameworks and have found that they exhibit a tacit ontology of realism. We have also concluded that Chisholm's ontology is a potential unifying framework for data modelling frameworks. The next step is to formalise the unifying framework based on Chisholm's ontology and to investigate the relationship between Chisholm's ontology and Bunge's ontology. Bunge's ontology is another realistic ontology used in the information systems literature but, in contrast with Chisholm's ontology, tends towards scientific realism.

We conjecture that Chisholm's ontology, due to its philosophical heritage, is likely to be more successful than Bunge's ontology in theorising about modelling in the application or real-world domain, where social issues dominate. In contrast, we expect Bunge's ontology due to its mechanistic and fine-grained nature, to be more successful in theorising about the implementation domain or the application domain, in cases where there is an absence of social or human issues.

Concluding, we note that this study further supports the use of ontology in information systems and that it should be seen as being complementary to studies using Bunge's ontology. The use of ontology to theorise about information systems modelling and representation is indeed of some considerable merit and, we hope, will lead to increasingly successful information systems.

Much more broadly, we also feel strongly that research is increasingly being conducted in isolation in that awareness of related fields and disciplines is limited. The unity of human endeavour towards the acquisition of knowledge has been significantly fragmented in the late twentieth century. Information systems and computer science have suffered because of it. It is our hope that this research will, in some small way, advance the cause of the 'university' as it once was to again recognise that we are all studying 'what there is'. For information systems, being an area that is so often reliant on research from related disciplines, this hope is most important for ensuring quality in its own research.

References

Apel, K.-O. (1980) *"Wittgenstein and the Problem of Hermeneutic Understanding" Towards a Transformation of Philosophy,* Routledge & Kegan Paul, Boston.

Artz, J. M. (1997) A Crash Course in Metaphysics for the Database Designer. *Journal of Database Management,* 8 (4), pp. 25–30.

Atkinson, M., Bancilhon, F., DeWitt, D. and Zdonik, S. (1989) The Object-oriented Database System Manifesto, In Proceedings of *International Conference on Deductive Object-Oriented Databases* Kyoto, pp. 40–57.

Audi, R. (Ed.) (1995) *The Cambridge Dictionary of Philosophy,* Cambridge University Press, Cambridge.

Banerjee, J., Chou, H.-t., Garza, J. F. and Kim, W. (1987) Data Model Issues for Object-Oriented Applications. *ACM Transactions on Office Information Systems,* 5 (1), pp. 3–26.

Batini, C., Lenzerini, M. and Navathe, S. B. (1986) A comparative analysis of methodologies for Database schema integration. *ACM Computing Surveys,* 18 (4), pp. 323–346.

Blaha, M. and Premerlani, W. (1998) *Object-Oriented Modeling and Design for Database Applications,* Prentice Hall, Upper Saddle River.

Boman, M., Bubenko, J. A., Johannesson, P. and Wangler, B. (1997) *Conceptual Modelling,* Prentice Hall, London.

Brodie, M. L. (1982) On the Development of Data Models In *On Conceptual Modelling* (Eds, Brodie, M. L., Mylopoulos, J. and Schmidt, J. W.) Springer-Verlag, New York, pp. 19–47.

Bühler, K. (1934) *Sprachtheorie,* Gustav Fischer, Jena.

Bunge, M. (1977) *Treatise on Basic Philosophy: Vol. 3: Ontology I: The Furniture of the World,* Reidel, Boston.

Bunge, M. (1979) *Treatise on Basic Philosophy: Vol. 4: Ontology II: A World of Systems,* Reidel, Boston.

Cárdenas, A. F. and McLeod, D. (1990) An Overview of Object-Oriented and Semantic Database Systems, In *Research Foundations in Object-Oriented and Semantic Database Systems* (Eds, Cárdenas, A. F. and McLeod, D.) Prentice Hall, Englewood Cliffs, 341 pages.

Carnap, R. (1963) *The Logical Construction of the World,* University of California Press (Berkeley).

Cattell, R. G. G. (1991) *Object Data Management,* Addison-Wesley, Menlo Park.

Chen, P. (1976) The Entity-Relationship Model—Toward a Unified View of Data. *ACM Transactions on Database Systems,* 1 (1), pp. 9–36.

Chisholm, R. (1957) *Perceiving: A Philosophical Study,* Cornell University Press, Ithaca.

Chisholm, R. (1976) *Person and Object: A Metaphysical Study,* Open Court, La Salle.

Chisholm, R. (1979) Objects and persons: revisions and replies. *Grazer Philosophische Studien,* 7/8 pp. 317–388.

Chisholm, R. (1982) *The Foundations of Knowing,* University of Minnesota Press, Minneapolis.

Chisholm, R. (1989a) *On Metaphysics,* University of Minnesota Press, Minneapolis.

Chisholm, R. (1989b) *Theory of Knowledge,* Prentice-Hall, Englewood Cliffs.

Chisholm, R. M. (1992) In *Language, Truth, and Ontology* (Ed, Mulligan, K.) Kluwer Academic Publishers, Dordrecht, 211 pages.

Chisholm, R. M. (1996) *A Realistic Theory of Categories—An Essay on Ontology,* Cambridge University Press.

Codd, E. F. (1970) A Relational Model for Large Shared Data Banks. *Communications of the ACM,* 13 (6).

Communities, E. (1999) Welcome to Europa, Vol. 1999 The European Union.

Dancy, J. and Sosa, E. (Eds.) (1992) *A Companion to Epistemology,* Blackwell Publishers, Oxford.

Date, C. (1995) *An Introduction to Database Systems,* Addison-Wesley, Menlo Park.

Dobbie, G. (1990) Object Oriented Database Systems: A Survey, Department of Computer Science, The University of Melbourne, Melbourne.

Eco, U. (1976) *A Theory of Semiotics,* Midland, Bloomington.

Flew, A. (1989) *An Introduction to Western Philosophy: Ideas and Arguments from Plato to Popper,* Thames and Hudson, London.

Goldberg, A. and Robson, D. (1983) *Smalltalk-80, The Language and its Implementation,* Addison-Wesley, Reading, MA.

Green, P. (1996) An Ontological Analysis of Information Systems Analysis and Design (ISAD) Grammars in Upper Case Tools, The University of Queensland.

Griethuysen, J. J. v. (1982) Concepts and Terminology for the Conceptual Schema and the Information Base: report by ISO-TC97-SC5-WG3, International Organization for Standardization, New York.

Guarino, N. (1998) Formal Ontology and Information Systems, In Proceedings of *Formal Ontology in Information Systems* (Ed, Guarino, N.) IOS Press, Amsterdam, Trento, Italy.

Habermas, J. (1992) *Postmetaphysical Thinking: Philosophical Essays,* Polity Press, Oxford.

Hammer, M. and McLeod, D. (1981) Database Description with SDM: A Semantic Database Model. *ACM Transactions on Database Systems,* 6 (3), pp. 351–386.

Hammer, M. and McLeod, D. (1990) Database Description with SDM: A Semantic Database Model, In *Research Foundations in Object-oriented and Semantic Database Systems* (Eds, Cárdenas, A. F. and McLeod, D.) Prentice Hall, Englewood Cliffs, 341 pages.

Honderich, T. (Ed.) (1995) *The Oxford Companion to Philosophy,* Oxford University Press, Oxford.

Hull, R. and King, R. (1987) Semantic Database Modelling: Survey, Applications, and Research Issues. *ACM Computing Surveys,* 19 (3), pp. 201–260.

Hull, R. and King, R. (1990) A Tutorial on Semantic Database Modeling, In *Research Foundations in Object-Oriented and Semantic Database Systems*(Eds, Cárdenas, A. F. and McLeod, D.) Prentice Hall, Englewood Cliffs, 341 pages.

Kaminsky, J. (1982) *Essays in Linguistic Ontology,* Southern Illinois University Press, Carbondale and Edwardsville.

Kattsoff, L. O. (1956) *Logic and the Nature of Reality,* Martinus Nijhoff, The Hague.

Keen, C. and Lakos, C. (1993) A Methodology for the Construction of Simulation Models using Object-oriented Petri Nets, In Proceedings of *European Simulation Multiconference on adapting object-oriented development methodologies to the design of discrete event simulation models.*

Keen, C. and Lakos, C. (1995) An Analysis of the Design Constructs Required in Process Modelling.

Keen, C. D. and Lakos, C. A. (1994) Information Systems Modelling using LOOPN++, an Object Petri Net Scheme, In Proceedings of *4th International Workshop on Dynamic Modelling* Delft.

Kent, W. (1978) *Data and Reality: basic assumptions in data processing reconsidered,* Elsevier Science Publishers B.V., Amsterdam.

Kent, W. (1979) Limitations of Record-Based Information Models. *ACM Transactions on Database Systems,* 4 (1), pp. 107–131.

Kerschberg, L. and Pacheco, J. E. S. (1976) A functional database model, Pontificia Univ. Catholica do Rio de Janeiro, Rio de Janeiro, Brazil.

Kim, J. and Sosa, E. (Eds.) (1995) *A Companion to Metaphysics,* Blackwell Publishers, Oxford.

King, R. and McLeod, D. (1985) Semantic Data Models, In *Principles of Database Design,* Vol. I (Ed, Yao, S. B.) Prentice-Hall, New Jersey, pp. 115-150.

Lécluse, C., Richard, P. and Velez, F. (1988) O2, and Object-Oriented Data Model, In Proceedings of *Advances in Database Technology* 88, pp. 556–562.

Li, Q. (1991) Object Data Model = Object-oriented + Semantic Models. *Computer Standards and Interfaces,* 13 pp. 99–103.

Milton, S. (1993) OBST, Department of Computer Science, The University of Tasmania, Hobart.

Milton, S. and Kazmierczak, E. (2000) Enriching the Ontological Foundations of Modelling in Information Systems, In Proceedings of *IS Foundations - Ontology, Semiotics and Practice, Proceedings* (Ed, Dampney, C. K.) Lighthouse Press, Macquarie University, Sydney, ISBN .

Milton, S., Kazmierczak, E. and Keen, C. (1998) Comparing Data Modelling Frameworks using Chisholm's Ontology, In Proceedings of *6th European Conference on Information Systems*, Vol. I Aix-en-Provence, pp. 260–272.

Milton, S. and Keen, C. (1996) Linguistic-based Information Systems Modelling, In Proceedings of *7th Australasian Conference on Information Systems*, Vol. 2 (Eds, Keen, C., Urquhart, C. and Lamp, J.) The University of Tasmania, Hobart, pp. 473–483.

Mulligan, E. K. (1992) *Language, Truth and Ontology,* Kluwer Academic Publishers, Dordrecht.

Mylopoulos, J. (1998) Information Modeling in the Time of the Revolution. *Information Systems,* 23 (3/4), pp. 127–155.

Nijssen, G. M. and Halpin, T. A. (1989) *Conceptual Schema and Relational Database Design: A Fact Oriented Approach,* Prentice-Hall, New York.

Parsons, J. and Wand, Y. (1993) Object-Oriented Systems Analysis: A Representation View, Faculty of Commerce and Business Administration, The University of British Columbia, Vancouver, 22 pages.

Peckham, J. and Maryanski, F. (1988) Semantic Data Models. *ACM Computing Surveys,* 20 (3), pp. 153–189.

Quine, W. V. O. (1960) *Word and Object,* MIT Press, Cambridge Massachusetts.

Quine, W. V. O. (1969) *Ontological Relativity and Other Essays,* Columbia University Press, New York.

Rohde, F. (1995) An Ontological Evaluation of Jackson's System Development Model. *Australian Journal of Information Systems,* 2 (2), pp. 77–87.

Rumbaugh, J., Blaha, M., Premerlani, W., Eddy, F. and Lorensen, W. (1991) *Object-Oriented Modeling and Design,* Prentice-Hall, Englewood Cliffs, New Jersey.

Russell, B. (1908) Mathematical Logic as Based on the Theory of Types. *American Journal of Mathematics,* XXX pp. 222-263.

Russell, B. (1946) In *Essays on Logic and Knowledge* Allen and Unwin, London, pp. 347–363.

Shipman, D. W. (1981) The Functional Data Model and the Data Language DAPLEX. *ACM Transactions on Database Systems,* 6 (1), pp. 140–173.

Smith, B. (1995) Formal Ontology, Commonsense and Cognitive Science. *International Journal of Human-Computer Studies,* 43 (12), pp. 641–667.

Stonebraker, M. (Ed.) (1994) *Readings in Database Systems,* Morgan Kaufman Publishers.

Takagaki, K. and Wand, Y. (1991) An Object-Oriented Information Systems Model Based on Ontology, In Proceedings of *Object Oriented Approach in Information Systems.*

Teorey, T. J., Yang, D. and Fry, J. P. (1986) A Logical Design Methodology for Relational Databases Using the Extended Entity-Relationship Model. *ACM Computing Surveys,* 18 (2), pp. 197–222.

Vet, P. E. v. d. and Mars, N. J. I. (1998) Bottom-Up Construction of Ontologies. *IEEE Transactions on Kowledge and Data Engineering,* 10 (4), pp. 513–526.

Vickery, B. C. (1997) Ontologies. *Journal of Information Science,* 23 (4), pp. 277–286.

Wand, Y. (1989) In *Object-Oriented Concepts, Applications, and Databases* (Ed, Won Kim, F. L.) Addison-Wesley, Reading, pp. 537-559.

Wand, Y. (1996) Ontology as a Foundation for Meta-modelling and Method Engineering. *Information and Technology Software,* 38 pp. 182–287.

Wand, Y., Monarchi, D. E., Parsons, J. and Woo, C. (1995) Theoretical Foundations for Conceptual Modelling in Information Systems Development. *Decision Support Systems,* 15 (1995) pp. 285—304

Wand, Y., Storey, V. C. and Weber, R. (1993) Analyzing the Meaning of a Relationship, Faculty of Commerce and Business Administration, The University of British Columbia, Vancouver, 31 pages.

Wand, Y. and Weber, R. (1989) In *Information Systems Concepts: An In-depth Analysis* (Eds, Falkenberg, E. D. and Lindgreen, P.) Elsevier Science Publishers B.V., Amsterdam, pp. 79–107.

Wand, Y. and Weber, R. (1990) An Ontological Model of an Information System. *IEEE Transactions on Software Engineering,* 16 (11), pp. 1282–1292.

Wand, Y. and Weber, R. (1993) On the Ontological Expressiveness of Information Systems Analysis and Design Grammars. *Journal of Information Systems,* 1993 (3), pp. 217–237.

Wand, Y. and Weber, R. (1995) On the Deep Structure of Information Systems. *Information Systems Journal,* 1995 (5), pp. 203–223.

Wand, Y. and Woo, C. C. (1993) Object-oriented Analysis—Is It Really That Simple, In Proceedings of *Third Annual Workshop on Information Technologies and Systems* (Eds, Hevner, A. and Kamel, N.) Orlando, Florida, reprint.

Weber, R. (1997) *Ontological Foundations of Information Systems,* Buscombe Vicprint, Blackburn, Victoria.

Weber, R. and Zhang, Y. (1996) An analytical evaluation of NIAM's grammar for conceptual schema diagrams. *Information Systems Journal,* 1996 (6), pp. 147–170.