THE NATURE AND ACQUISITION OF EXPERT KNOWLEDGE TO BE USED IN SPATIAL EXPERT SYSTEMS FOR CLASSIFYING REMOTELY SENSED IMAGES

UNIVERSITY OF TASMANIA

Paul Crowther

MSc, BSc(hons), Grad Dip Comp Stud,
Grad Dip Admin, Dip Teach (TAFE)

School of Computing
University of Tasmania
PO Box 1214, Launceston,
Tasmania, Australia, 7250

A thesis submitted in fulfilment of the requirements for the degree of
Doctor of Philosophy
University of Tasmania
12th June 1999

DECLARATION

I hereby declare that this submission is my own work and that, to the best of my knowledge and belief, it contains no material previously published or written by another person nor material which to a substantial extent has been accepted for the award of any other degree or diploma of the University or any other institute of higher learning, except where due acknowledgment is made in the text.

Paul Crowther
12th June 1999

AUTHORITY OF ACCESS

This thesis may be made available for loan and limited copying in accordance with the *Copyright Act 1968*.

Paul Crowther
12th June 1999

ABSTRACT

Knowledge engineering is the process of acquiring expert knowledge from human domain experts. In this thesis the emphasis is on the acquisition of geographic or spatial knowledge from experts involved in interpreting multi-spectral satellite images.

This thesis argues that spatial knowledge is primarily visual, hence tools to acquire it also need to be visual. Currently there is no methodology, other than ad hoc interview and protocol analysis, for acquiring expert knowledge of interpretation of satellite images. As a result, there cannot be an integrated knowledge acquisition toolkit, since this must be based on a formal methodology. This thesis offers a methodology to overcome this shortcoming and presents a series of tools to implement the methodology.

In the first part of the thesis the nature of geographic knowledge is investigated. A geographic knowledge classification scheme is presented as the basis of the work in the rest of the thesis. It is shown that geographic knowledge can be divided into a six level hierarchy:

- Primitive knowledge about point, line and areal objects,
- Relationship knowledge about the relationships between primitive objects,
- Assembly knowledge about related collections of primitive objects,
- Non-Visual knowledge of expert heuristics (knowledge of short cuts acquired by experience),
- Consolidation knowledge of how to resolve and evaluate conflicting information and
- Interpretation knowledge of how to combine the other knowledge types to produce a classified image.

This six level hierarchical classification of geographic knowledge forms the basis of the KAGES (Knowledge Acquisition for Geographic Expert Systems) methodology.

Traditional knowledge acquisition procedures are studied and their relevance to a geographic domain discussed. This includes both human interaction techniques such

as interviewing and automated knowledge acquisition methods such as neural networks and machine learning. It will be shown that although automated pattern recognition techniques are important, there is still a need to include knowledge acquired by human image interpreters in an automated image interpretation system.

There is a theoretical discussion of new techniques to acquire visual knowledge of the types identified in the KAGES methodology. It is shown that these methods can be combined into an integrated knowledge engineering toolkit to acquire geographic knowledge from satellite image interpreters. Not all geographic knowledge is visual however. Three types of non-visual knowledge, algorithmic, heuristic and temporal, are identified and investigated. The first two are implemented in the knowledge engineering toolkit described in this thesis.

It is shown that if there are multiple domain experts and multiple knowledge acquisition sessions multiple knowledge-bases will be produced. Techniques for the consolidation of these knowledge-bases is presented.

The final section of the thesis involves evaluation of KAGES. This is done in two ways: user evaluation and application of the methodology in two domains. The user evaluation of the KAGES methodology and toolkit involved a number of image interpretation experts from a variety of domains and currently using a variety of tools. They were questioned about the usefulness and useability of the KAGES toolkit.

The results of using the tools in the toolkit are evaluated by generating rules for two scenarios, one for sea ice identification and the other for crop recognition. The rules produced using the toolkit are compared with rules produced using other techniques. The effect of applying rules generated by the toolkit to classify images is compared with the results from other image classification methods.

ACKNOWLEDGEMENTS

The work reported in this thesis has been supported by an Antarctic Science Advisory Committee Research Grant and a Horticultural Research and Development Grant.

I would like to thank Steve Pendlebury (Bureau of Meteorology, Hobart), Dr. Robert Massom and Dr. Kelvin Micheal (Antarctic CRC, Hobart) for their assistance in providing glaciological knowledge. Also Dr. Rowland Lawrence, Chris Abbott (Tasmanian Institute for Agricultural Research) and Mark Brown (Department of Primary Industry and Fisheries) for their help with agricultural knowledge and the MIRC project. Ross Lincoln of the University of Tasmania deserves a special thanks as he was responsible for image acquisition (for which he was paid) and provision of an alternative classification method as well as evaluating KAGES (for which he was not paid).

As well as the above I would like to thank the following for evaluation and feedback on the KAGES toolkit: Assoc. Prof. Manuel Nunez, Prof. Jamie Kirkpatrick, Dr. Eleanor Bruce (Department of Geography and Environmental Studies, University of Tasmania), Mr Mark Chladil (Tasmania Fire Service), Mark Garvey (Country Fire Authority, Victoria), Guy Tuddenham and Dr. John Le Marshall (Bureau of Meteorology, Melbourne)

Staff and students of the School of Computing have been supportive throughout this project. I thank Head of School, Prof. Young Choi, for his continued support and advice. Jacky Hartnett, co-author of many of the papers I have written, has offered much valuable advice on which direction KAGES should be headed and without her half the papers produced would not have been written. Dr. Mike Cameron-Jones, as my associate supervisor, kept me on track while my supervisor was away and kept me focused. Finally many thanks go to my supervisor, Dr. Ray Williams, who has read so many drafts of this thesis he probably knows it off by heart. I appreciate his advice, guidance and influence over the last four years.

CONTENTS

	page
ABSTRACT	iii
ACKNOWLEDGEMENTS	v
LIST OF FIGURES	xi
LIST OF TABLES	xv
1. INTRODUCTION	1
1.1 Knowledge Engineering In Spatial Expert Systems	1
1.2 Statement of Thesis	6
1.3 Chapter Plan	7
1.4 Publications	8
2 THE NATURE OF SPATIAL KNOWLEDGE	9
2.1 Knowledge	9
2.2 Expertise and the Domain Expert	9
2.3 Expert Systems	10
2.4 Visual Knowledge	12
2.4.1 Geographic or Spatial Knowledge	12
2.4.2 Classification of Geographic Knowledge	13
2.4.3 Problems of Spatial Knowledge Classification Schemes	15
2.4.4 Unique Characteristics of Spatial Knowledge	16
2.5 Verification of Expert Interpretations	17
2.6 The Needs of a Visual Geographic Knowledge Acquisition System	19
2.7 A Proposed Classification Of Geographic Knowledge	19
2.7.1 Primitive Knowledge	20
2.7.2 Relationship Knowledge	21
2.7.3 Assembly Knowledge	21
2.7.4 Non Visual Knowledge	22
2.7.5 Temporal Knowledge	22
2.7.6 Algorithmic Knowledge	23
2.7.7 Heuristic Knowledge	23
2.7.8 Consolidation Knowledge	23
2.7.9 Interpretation Knowledge	24

2.8 Visual Knowledge and KADS 24
2.9 Conclusion 26
3 KNOWLEDGE-BASED SPATIAL INFORMATION SYSTEMS 27
3.1 Geographic Information Systems 27
3.1.1 Data Input and Verification 28
3.1.2 Data Output and Presentation 29
3.1.3 Data Transformation and Analysis 29
3.1.4 Data Storage and Database Management 30
3.1.5 Interaction With the User 33
3.2 Remote Sensing 36
3.2.1 Active and Passive Remotely Sensed Data 36
3.2.2 The Electromagnetic Spectrum 37
3.2.3 Problems With Remotely Sensed Data 38
3.3 The Need for Data Structures 39
3.3.1 Evaluation of Vector versus Raster Data Structures for this Study 39
3.3.2 Knowledge-base Structure for a Geographic System 40
3.4 Survey of Knowledge-based GIS and RSS 43
3.5 Conclusion 46
4 TRADITIONAL KNOWLEDGE ACQUISITION TECHNIQUES 47
4.1 Introduction 47
4.2 Manual Techniques 48
4.2.1 Interviews 48
4.2.2 Observation / Protocol Based Analysis 50
4.2.3 Repertory Grid 53
4.3 Automated Techniques 55
4.3.1 Machine Learning 56
4.3.2 Neural Networks 58
4.3.3 Supervised Versus Unsupervised Methods 59
4.3.4 The Importance of Training Data 60
4.4 Combination Of Knowledge Gained From Multiple Techniques and Multiple Experts 61
4.5 Knowledge Acquisition And GIS 63
4.6 Conclusion 65

5 THE SPATIAL KNOWLEDGE ACQUISITION PROCESS 66

5.1 Introduction 66

5.2 Scenarios For Assessing Spatial Knowledge Acquisition 66

5.2.1 Antarctic Sea Ice Classification 67

5.2.2 Crop Recognition 69

5.3 Primitive Knowledge Acquisition 71

5.3.1 Knowledge of Areal Features 71

5.3.2 Knowledge of Point Features 73

5.3.3 Knowledge of Line Features 73

5.4 Relationship Knowledge Acquisition 75

5.4.1 The Problem of Three Objects 79

5.5 Assembly Knowledge Acquisition 79

5.6 Non-Visual Knowledge Acquisition 80

5.6.1 Heuristic Knowledge Acquisition 80

5.6.2 Algorithmic Knowledge Acquisition 81

5.6.3 Temporal Knowledge Acquisition 82

5.7 Consolidation Knowledge Acquisition 83

5.7.1 Primitive Knowledge 83

5.7.2 Relationship Knowledge 87

5.7.3 Assembly Knowledge 88

5.7.4 Heuristic Knowledge 90

5.8 Interpretation Knowledge Acquisition 91

5.9 Conclusions 91

6 THE KAGES SPATIAL KNOWLEDGE ACQUISITION TOOLKIT 93

6.1 Overview 93

6.2 Primitive Knowledge: The Per Pixel Tool 95

6.2.1 Primitive Knowledge 95

6.2.2 Identifying Areal Features 96

6.2.3 Identifying Point Features 97

6.2.4 Identifying Line Features 98

6.2.5 Primitive Knowledge Data Structures 99

6.2.6 Primitive Knowledge Rules 100

6.3 Relationship Knowledge Acquisition: The Spatial Relationship Tool 101

6.3.1 Spatial Relationships Involving Lines 102
6.3.2 Spatial Relationships Involving Points 103
6.3.3 Spatial Relationships Involving Areas 103
6.3.4 Relationship Knowledge rules 104
6.4 Assembly Knowledge Acquisition: The Region Of Interest Tool 104
6.4.1 Assembly Knowledge Rules 106
6.5 Non Visual Heuristic Knowledge Acquisition: The Repertory Grid Tool 107
6.5.1 Analysis 110
6.6 Non Visual Algorithmic Knowledge Acquisition: The Band Calculator Tool 111
6.7 A Tool To Assist Machine Learning And Verification: The Point Data Tool 113
6.8 Consolidation Knowledge Acquisition: The Consolidation Tool 114
6.9 Knowledge Verification Tool 115
6.10 The User Interface 116
6.11 Conclusions 118
7 EVALUATION OF THE KAGES SYSTEM 120
7.1 Introduction 120
7.2 Methodology For Comparing The Knowledge Acquired By The System 120
7.2.1 General Approach 120
7.2.1 The Sea Ice Mapping Domain 121
7.2.1 The Crop Recognition Domain 122
7.3 Methodology For Evaluating The Useability Of The System 122
7.3.1 User Selection 122
7.3.2 User Testing Procedure 123
7.4 Results Of The Evaluation 124
7.4.2 Verification Strategies in Remote Sensing 124
7.4.2 Antarctic Sea Ice Knowledge 125
7.4.2 Crop Recognition Knowledge 135
7.4.3 Expert User Acceptance 138
7.5 Conclusions 141

8 CONCLUSIONS 143

8.1 Introduction 143

8.2 An Integrated Geographic Knowledge Acquisition Strategy 143

8.3 An Integrated Toolkit to Implement the Knowledge Acquisition Strategy 145

8.4 User-Centred Geographic Knowledge Acquisition 146

8.5 Future Work 147

8.5.1 Temporal Knowledge 147

8.5.2 Interpretation Knowledge 147

8.5.3 Machine Learning Integration 148

8.5.4 Data Mining 149

9 REFERENCES 150

GLOSSARY 164

APPENDIX A: OPERATION OF KAGES 165

A.1 Main Menu 165

A.2 Load SPOT Image 166

A.3 Display Bands 166

A.4 Band Calculator 167

A.5 Per Pixel Tool 168

A.5.1 Verification of Per Pixel Tool 170

A.6 Relationship Tool 172

A.6.1 Verification of Relationships 174

A.7 Region of Interest Tool 176

A.8 Repertory Grid Tool 178

A.9 Point Data Tool 183

APPENDIX B: THE MIRC DATA BASE 186

APPENDIX C: USER INTERVIEWS 190

LIST OF FIGURES

page

Figure 1.1 GIS, Remote Sensing and Expert Systems. 5

Figure 2.1 Essential components of a knowledge-based system (after Mockler, 1989). 11

Figure 2.2 The four-layer model of expertise (Wielinga et al, 1992). 24

Figure 3.1 Components of a GIS. 28

Figure 3.2 Types of vector representations. 32

Figure 3.3 Semantic Network. 41

Figure 3.4 Part of a frame for storing meta-data (from Miller, 1994, p. 148). 42

Figure 3.5 An integrated RSS, expert system and GIS (after Wilkinson and Burrill, 1991). 43

Figure 4.1 Relationship of traditional knowledge acquisition techniques. 47

Figure 4.2 Multilayer feedforward artificial neural network for land type classification (Leung, 1997, p205). 58

Figure 5.1 Context of Antarctic study area. 67

Figure 5.2 Antarctic study area: An area in Vincennes Bay near Casey Base. Note South is up the page. Taken from a NOAA AVHRR band 1 image viewed using KAGES. 68

Figure 5.3 Context of study area for MIRC. 69

Figure 5.4: The MIRC study area as seen on band 1 of a SPOT image taken on 2nd November 1996 using the KAGES tool. 70

Figure 5.5 Minimum bounding rectangles. The example on the right has a true MBR superimposed. 71

Figure 5.6 A 2D-projection picture (after Nabil et al, 1995). 72

Figure 5.7 The eight relations between two regions (Egenhofer and Sharma, 1993). 76

Figure 5.8 Some spatial relationships between line, point and areal objects. 78

Figure 5.9 Three Object Definition: Fence (line object) defined by interaction of two fields (area objects) and a gate (point object) defined by the interaction of a field (area object) and access track (line object). 79

Figure 6.1 Overview of the KAGES system showing the primary knowledge acquisition tools. 93

Figure 6.2 The band viewing tool showing the image histogram and band 3 of a SPOT image of Table Cape area in Tasmania taken on 24th December 1996. 96

Figure 6.3 Per Pixel Tool. The window on the left shows the histogram of pixel values for the entire image. The centre window is of the image itself while the right window contains a zoomed section of the image to allow for accurate selection. From band 1 of a SPOT image of Table Cape, Tasmania 2nd November 1996. 97

Figure 6.4 Relationship Tool. The first object (X) is open water defined in band 1. The second (Y) is sea ice identified from band 3 of the NOAA image. The results are shown superimposed on band 1. 101

Figure 6.5 The Region Of Interest Tool. The defined area in the top left has an MBR drawn around it. The pop-up window shows the object Types found within the area (Landsat band 1, 27th November 1996) 105

Figure 6.6 Repertory Grid Tool with menu options and the window to enter elements displayed. (NOAA VHRR image of Casey taken on 26th February 1998) . 108

Figure 6.7 Triad window for determining discriminators. 109

Figure 6.8 Ranking an element (pack_ice) against discriminators. 109

Figure 6.9 Repertory Grid of Antarctic elements (features) with 5 indicating complete agreement with discriminator (concept) and 1 indicating complete disagreement with the discriminator. 110

Figure 6.10 Band Calculator for use on a three band SPOT image set. 112

Figure 6.11 Consolidation of knowledge-bases. 115

Figure 6.12 KAGES main user interface with a SPOT image loaded. The, image is of Table Cape, North West Tasmania and was taken on 2nd November 1996. 117

Figure 7.1 Domain expert classification of sample points on the training image. 126

Figure 7.2 KAGES classification of sample points on the training image. Cells containing letters indicate misclassification. 127

Figure 7.3 Thermal band 3 of the test image derived from a NOAA AVHRR Image of Mawson/Davis taken on 4th March 1998 viewed using the KAGES Band Display Tool. 129

Figure 7.4 The classified test image produced using the KAGES generated rules. 130

Figure 7.5 Domain expert classification of sample points on the test image. 131

Figure 7.6 KAGES classification of sample points on the test image. Cells containing letters indicate misclassification. 132

Figure 7.7 Growth and cropping times for crops under study showing times of useable satellite passes. 135

Figure A.1 KAGES main menu with SPOT HRV image (24th December 1996) already loaded. 165

Figure A.2 SPOT image (2nd november 1996) with near infra red (band3) displayed using the Display Image Bands option. 166

Figure A.3 Band Calculator shown in conjunction with a SPOT image. (2nd November 1996) 167

Figure A.4 Per-Pixel Tool being used to define an area. The zoomed window shows a segment in the bottom right hand corner of the main image segment. 168

Figure A.5 Threshold Adjuster window showing the selected area has pixel thresholds between 67 and 81. 168

Figure A.6 Region definition (naming) dialog window displaying an area objects characteristics and allowing a user to enter a name. 169

Figure A.7 An image produced by applying all rules. 171

Figure A.8 Initial window from the Spatial Tool which is superimposed over the main screen. A user is required to pick 2 objects. 172

Figure A.9 Spatial Tool with the selected objects displayed. 173

Figure A.10 Spatial Rule Editor which is superimposed on top of the previous image showing the relationships between sea and trees. 174

Figure A.11 Effect of selecting Apply Rules / Apply Spatial Rule with the previous rule selected. 175

Figure A.12 Region of Interest Tool defining a rural area (SPOT band 1, 2nd November 1996). 176

Figure A.13 Region of Interest Editor showing object types found. 177

Figure A.14 Naming window for Region of Interest Tool. 177

Figure A.15 Repertory Grid Tool after New Grid option chosen. 178

Figure A.16 Repertory Grid Load Window allowing entry of image features. 179

Figure A.17 Triad (Get Discriminators) Window asking users to differentiate between features. 180

Figure A.18 Ranking Window with slider bars to input ranking. 181

Figure A.19 Repertory Grid results. 182

Figure A.20 Point Data Tool showing sampling points designated by X on the image and the Naming Window. (NOAA VHRR image of Casey taken on 26th February 1998) 184

Figure B.1 MIRC database. 186

LIST OF TABLES

page

Table 3.1 The electromagnetic spectrum covered by Landsat. 37

Table 3.2 SPOT and NOAA AVHRR spectral characteristics. 37

Table 6.1 Data structures for scene primitives. 99

Table 6.2 Partial transcript of a file generated using the Point Data Tool on a NOAA AVHRR image. 114

Table 7.1 Antarctic classification results using the KAGES knowledge-base on the training image. 128

Table 7.2 Comparison of Icemapper, KAGES and S performance on the training image. 128

Table 7.3 Antarctic classification results using the KAGES knowledge-base on the test image. 129

Table 7.4 Comparison of Icemapper, KAGES and S performance on the test image. 133

Table 7.5 Crop classification results using the KAGES knowledge-base. 137

Table 7.6 Summary of user responses to KAGES. 139

Chapter 1. INTRODUCTION

In this chapter the origins and confluence of remote sensing, expert systems and geography are discussed. The importance of knowledge acquisition for geographic expert systems is argued. A statement of thesis and a chapter plan is presented.

1.1 KNOWLEDGE ENGINEERING IN SPATIAL EXPERT SYSTEMS

"Geography is concerned with the description and explanation of the areal differentiation of the earth's surface." (Harvey, 1993, p3) This can be in terms of natural features (physical geography) or features which relate to people (human geography). The field of geography is wide as there are numerous subfields, for example environmental geography, biogeography, historical geography, climatology and geomorphology (many of which can themselves be subdivided). The one thing all of these specialist fields have in common is their use of spatial analysis (that is they study the way phenomena are distributed in space). The spatial distribution of phenomena is normally shown in the form of maps, which are two dimensional representations of parts of the earth's surface.

Maps have been used for centuries and although some map categories such as topographic maps are multi-purpose, they have always tended to relate to a particular theme. Initially each map was used for a specific purpose, but maps with different themes can be combined in a process known as overlay. Unwin (1996) cites early uses of map overlays in the mid-nineteenth century to link cholera outbreaks to water pumps (Snow, 1936). One map depicted the water reticulation system, the second map a plot of cholera outbreaks. Map overlays are now an essential analytical tool in Geographic Information Systems (GIS).

Data used in maps originally came from ground based cartographers who would survey a particular area. A major advance came in the form of airborne observations, especially from aircraft during the First World War when aerial photography was introduced. The data gathered using this method was based on reflected visible light.

To gain more information infra red spectral data was subsequently used and the visible spectrum was split into several distinct components.

The term remote sensing was introduced to distinguish multi-spectral imaging from aerial photography. Remote sensing is defined as:

"...the measurement of objects' properties on the earth's surface using data acquired from aircraft and satellites."
(Schowengerdt, R. A., 1997, p2)

where the term "properties" relates to electromagnetic properties. Use of sensors other than those sensitive to visual light began in the 1960's and space borne sensors appeared in the 1970's. With the launch of the first Landsat system in 1972, data from space borne multispectral sensors providing large scale ground cover became readily available.

Geographic Information Systems (GIS) are inherently linked to remote sensing systems (Star et al, 1997). GIS's are systems which allow a user to manipulate and analyse geographic data. They feature graphical user interfaces and integrated databases. Their development since the 1980's has been rapid. The number and range of applications that use GIS has increased markedly over the last few years (Star et al, 1997) and form the basis of the Tasmanian Governments' LIST (Land Information System Tasmania) component of the Future Directions Statement.

One problem with remote sensing systems is the resolution of the data they capture. Each image is made up of picture elements or pixels. In digital processing these are numbers, within a specified range, for example 0 to 256. Shades of gray or colours are assigned depending on the number. For example a value of 0 will be black and 256 will be white on a gray scale image. Currently, the best resolution is for each pixel on an image to represent an area of about 20 meters by 20 meters on the earths surface. Sensors with resolution of about one meter will soon be available. By contrast, in GIS the resolution is user controlled. Therefore transferring data from a remote system to a GIS causes problems of accurately positioning objects. Despite this, one of the great advantages of remotely sensed data is that it can be gathered in a systematic, timely fashion (Muller, 1988).

A second problem with remote sensing systems is that although large amounts of geographic data can be acquired, the expertise to interpret it is scarce. Therefore some means of automating tasks such as image interpretation and classification is required to make the best use of the data acquired (Openshaw and Scholten, 1994). Processing techniques using artificial intelligence are a possible solution. There are three areas of artificial intelligence which are used for image interpretation; expert systems, machine learning and neural networks.

The field of expert systems developed rapidly in the 1970's. Expert systems are programs which emulate an expert's problem solving methods (Pigford and Baur, 1990). An expert is often called the *domain expert* to indicate that a persons field of expertise is limited to a specific subject or domain. The expert's expertise is acquired in a variety of ways, including formal training. However the distinguishing factor of the way an expert gains expertise is through on-the-job experience. With experience a practitioner in a particular field develops rules of thumb and short cuts which enables them to solve problems more effectively than others in the field. An expert is only an expert because other practitioners in the field recognise them as such. An expert's knowledge may be supplemented with other information gained from records and the literature.

One of the early expert systems was the medical diagnostic system MYCIN. This was typical of the first generation expert systems which were predominantly text based diagnostic programs (Bonnet et al 1988). Another early system was Prospector which was used to predict the location of mineral deposits. Prospector was not a spatial expert system, but did analysis and prediction based on the results of analysis of geological samples and presented its results in map form. Later expert systems were developed for more complex tasks including planning and classification. Coupled with a graphical user interface, they are now being used for image analysis.

One of the major problems in developing expert systems is acquiring knowledge to code into them. This can be extremely time consuming, requiring numerous interviews with domain experts. The term *knowledge bottleneck* was coined to describe the problem (Weibel et al, 1995). A variety of techniques have been developed to overcome the knowledge bottleneck including fully automated techniques using machine learning and neural networks.

Knowledge used by expert interpreters of remotely sensed data is predominantly visual. Traditionally the knowledge in expert systems has had a semantic association. With visual media an expert is more likely to use the 'I know it when I see it' approach (Gupta et al, 1997). Experts see objects in images and can point to them and name them. To do this they may use specific bands of a set of satellite images or a combination of those bands. They may also work on histogram representations of an image. Hence an image interpretation expert knows what range the spectral signature (or reflectance value) of an object will fall into. They can recognise spatial relationships such as streams flowing in valleys and down slopes, and collections of objects, such as roads and buildings, which make up a settlement.

Because of its visual nature the acquisition of this knowledge from the expert is difficult using most traditional techniques. To overcome this difficulty a visual method of knowledge acquisition is needed. That is a method which directly captures the steps used by an expert to classify an image rather than requiring the expert to describe those steps.

Currently most automated systems classify images on a per pixel basis (Wilkinson, 1996). Each pixel or individual picture element that makes up an image is classified as being part of a particular object type and assigned a label. This technology is quite well developed and available in GIS and image processing packages. It is also the basis for pattern recognition systems using neural network and machine learning techniques.

A more difficult processing task is one which deals with spatial relationships between objects (Openshaw and Clarke, 1996). This has three components: the nearness of two objects, the orientation of two objects with respect to each other and the degree of overlap of two objects (Egenhofer and Sharma, 1993). A further complication occurs when the interaction of two objects defines a third object. For example where an object 'sea' meets an object 'land' a third object, a line, is defined which is 'shoreline'.

A second complex relationship occurs when a group of objects are combined into a single object. This process is called generalisation (McMaster and Mark, 1991) and is a technique that is used extensively in cartography. If it is impractical to show all the details of an object, or group of objects, so they are simplified. Hence a two track railway is often represented as a single track. A town is not shown as individual houses but as an object representing urban development.

In both these cases techniques which can handle knowledge beyond the per pixel level are needed. These require a system which can identify and manipulate spatial objects and groups of objects. To be able to do this an expert knowledge acquisition system needs to allow a domain expert to point and draw on an image. By doing this a user can identify the objects to be manipulated and define areas of the image to be generalised.

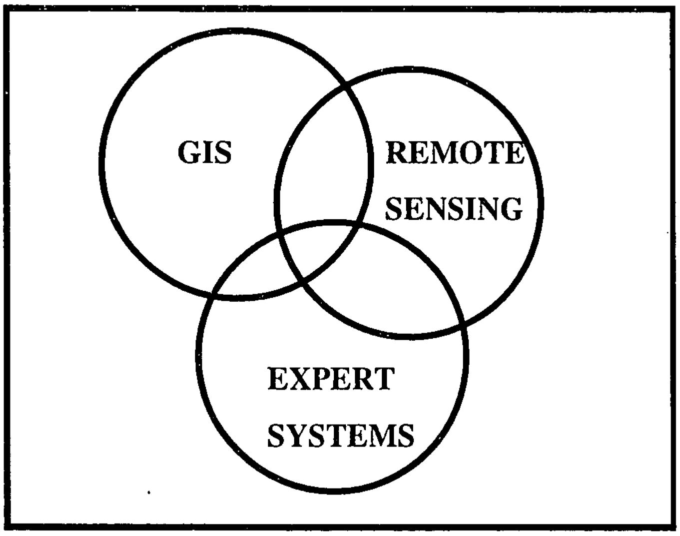

Figure 1.1 GIS, Remote Sensing and Expert Systems

The three subdisciplines which relate to this study are shown in Figure 1.1. Remote sensing is often closely tied to GIS. Images can be loaded into and analysed by GIS which have a raster (or pixel) processing ability (Star et al, 1997). There are expert systems, generally associated with image processing, which are used to directly classify remotely sensed data (Williams et al, 1994). Likewise there are expert systems coupled to GIS (Hartnett et al, 1994). This study is related to the overlap of the three areas. It is a study of the types of geographic knowledge which can be acquired from remotely sensed images and can be used in conjunction with a GIS.

1.2 STATEMENT OF THESIS

The aim of this thesis is to investigate knowledge acquisition techniques for expert systems used in conjunction with remotely sensed images and geographic expert systems. This includes an investigation of the nature of expert knowledge in a geographic context. A new classification of visual geographic knowledge consisting of *Primitive, Relational, Assembly, Non-visual, Consolidation* and *Interpretation knowledge*, is presented.

To validate this classification, tools and techniques for geographic knowledge acquisition from satellite images were developed for each of the knowledge acquisition categories. Traditional techniques of knowledge acquisition are investigated to determine their suitability in a geographic domain and to see how well they compare with techniques using primarily visual knowledge acquisition tools. A tool-kit testbench was developed to implement the tools within specific knowledge categories. The component tools are evaluated on a series of systems and the results analyzed both in terms of the accuracy of the rules produced and user acceptance.

It will be postulated that geographic knowledge has different characteristics from non-visual knowledge and therefore unique tools, reflecting the way the user analyses images, are required for knowledge acquisition. It will also be postulated that visual knowledge elicited from different sources and gathered by different elicitation techniques should be stored as separate knowledge-bases; then combined into a single consolidated knowledge-base prior to use. Despite the visual nature of geographic knowledge, it will also be shown that there is a need for non-visual knowledge to be integrated into a geographic knowledge-base.

1.3 CHAPTER PLAN

The second chapter investigates the nature of knowledge and expertise, in particular the unique properties of spatial and geographic knowledge and expertise. The role of the domain expert in terms of expert knowledge in relation to spatial systems and interpretation of remotely sensed data is studied and a new classification of geographic knowledge is presented.

The third chapter is a review of the theory and requirements of Geographic Information Systems and discusses the choice of data structures for use with expert systems which interact with them. There is an investigation of how the data structures can be utilized in a knowledge-base.

The fourth chapter explores the various traditional knowledge acquisition methods, the tools available and their applicability to geographic expert systems.

The fifth chapter introduces the scenarios to be used in testing the spatial knowledge acquisition methodology. There is a conceptual description of the tools and the features that are needed to acquire knowledge based on the classification system presented in Chapter 2. It will also describe the methods of verifying and combining spatial knowledge-bases.

The sixth chapter describes the implementation of a graphical geographic knowledge acquisition tool-kit, KAGES (Knowledge Acquisition for Geographic Expert Systems), based on the theoretical considerations presented in Chapter 5.

The seventh chapter presents the results of using the tool-kit in the domains described in Chapter 5. Knowledge acquisition using the methodology is compared, in terms of knowledge acquired, with other acquisition strategies. There is an evaluation of expert user acceptance of the various techniques and conclusions are drawn about knowledge acquisition for spatial systems.

The final chapter presents conclusions about geographic knowledge acquisition and suggests directions for future work.

1.4 PUBLICATIONS

Portions of this work have been published previously (Crowther and Hartnett (1996a), Crowther and Hartnett (1996b), Crowther and Hartnett (1997a), Crowther et al (1997), Crowther and Hartnett (1997b), Crowther (1998)). Note that the terminology used here supersedes the terminology defined previously.

Chapter 2. THE NATURE OF SPATIAL KNOWLEDGE

This chapter discusses knowledge and expertise in relation to visual systems and geography. There is an investigation of the role of expert systems in geography and the unique properties of spatial knowledge. A six level classification of visual geographic knowledge is presented which is used as the theoretical basis for developing knowledge acquisition tools.

2.1 KNOWLEDGE

Knowledge is a prerequisite for expertise in that first one must have knowledge, then after experience applying that knowledge one becomes recognised as an expert. Knowledge therefore is an awareness, familiarity and understanding acquired through education or experience. As such it is information which has been learned, perceived, discovered, inferred or understood (Nagao, 1988).

Visual and geographic knowledge introduces a spatial aspect to the definition. This knowledge, or spatial cognition, involves recognition of objects, patterns and the relationships between objects in space(Lloyd, 1997).

2.2 EXPERTISE AND THE DOMAIN EXPERT

An expert is a person who is regarded as a pre-eminent practitioner in their field of expertise (Agnew et al, 1994). Often called the domain expert (the term used in this thesis), they have in-depth knowledge acquired through training and experience. Meyer and Booker (1991) define an expert as

"... a person who has a background in the subject area and is recognised by his or her peers ... as qualified to answer questions" (p3)

Many of the methods experts use to solve problems are *heuristics* – rules of thumb, empirical knowledge or shortcuts developed by an expert through experience which may aid in a problems solution, but are not guaranteed to work (Giarratano and Riley, 1994). An important point of Agnew's paper and Meyer and Booker's Introduction is that expertise is socially selected. You are not an expert unless others regard you as an expert. Expertise in a specific domain is therefore rare.

Experts are generally only acknowledged experts in a particular field or domain. It is also possible that there is no one domain expert covering a particular domain. Instead there may be several experts in the same domain who as a group cover the domain of expertise. Another possibility is a group who are experts in overlapping subdomains (Barrett and Edwards, 1995). This is common when working in geographic domains where expertise in a variety of fields is combined into a geographic information system.

There are advantages and disadvantages to working with a group of experts. This is particularly so when their knowledge has to be combined. With one expert one is reliant on their expertise and their acknowledged status as an expert. With two experts one has the problem of conflict. With three or more experts one has the problem of minority interpretations, one of which could be the correct interpretation (Medsker et al, 1995)

Finally, the nature of expertise is temporal (Fuller, 1994). In technology based activities an expert must continually build their expertise, and the nature of that expertise will change. The proliferation of expert systems has had the effect of making expertise more widely available to users without skilling the users. The expert's standing may be eroded because of this. This is particularly true of emerging technologies where the early practitioners are regarded as experts, but their status as an expert is eroded as their skills become more commonplace.

2.3 EXPERT SYSTEMS

Expert systems contain representations of human expert knowledge which can be applied to a problem within a given problem domain. They are systems which emulate the result of human problem solving. The components of a typical expert system are shown in Figure 2.1. Knowledge is stored in a knowledge-base which is a collection of facts and heuristics. In expert systems this knowledge can be expressed as facts and rules, frames or semantic networks (see 3.3.2).

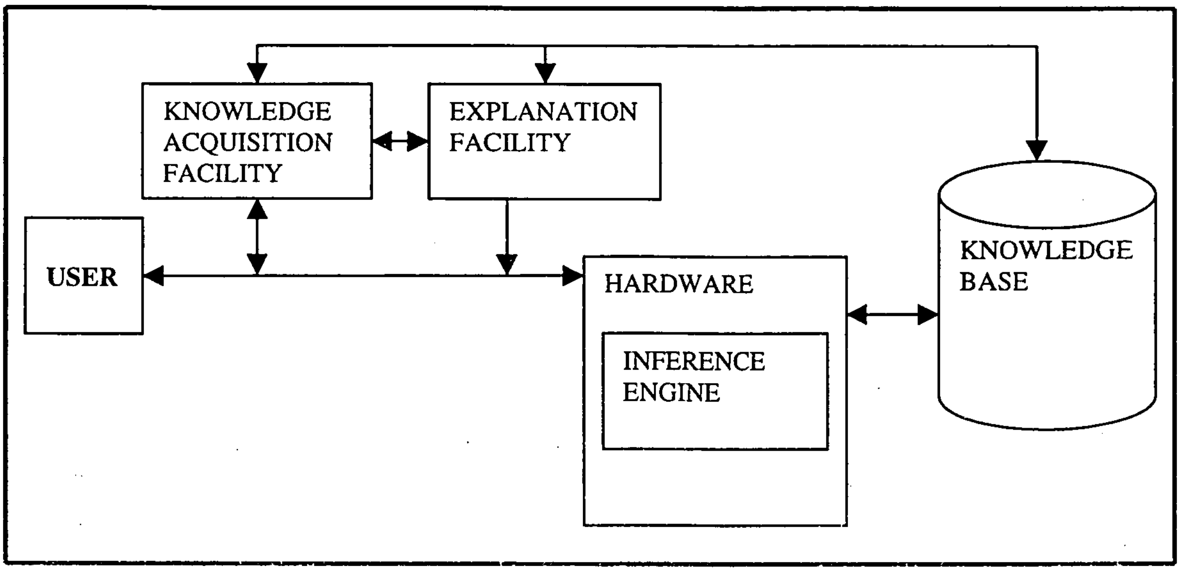

Figure 2.1 Essential components of a knowledge-based system (after Mockler, 1989)

The second major component of an expert system is the inference engine. The inference engine is a computer program which guides the manipulation of knowledge contained in a knowledge-base (Mockler, 1989). It is often supplied as part of a development package (Mockler and Dologite, 1992).

Hence the primary task of developing an expert system is the capture and coding of expert knowledge of a particular problem domain into a knowledge-base. This knowledge can come from a variety of sources; domain expert, journals, technical manuals, historic data, electronic databases and in the case of spatial systems, maps and GIS. In this thesis, the emphasis is on knowledge elicitation from the domain expert.

2.4 VISUAL KNOWLEDGE

Visual (including spatial) knowledge presents special problems for knowledge acquisition. Recognising visual features is easy for a human although the cognitive processing is complex (Lloyd, 1997). Describing those features without the use of diagrams is difficult. It is easy for a human expert to show what something looks like, but far more difficult to describe it in words, and more difficult again to describe it in terms of rules. (Kweon and Kanade, 1994). There has been work done in multimedia where sound and animation have been added to GIS, (Galetto and Spalla, 1996) but this still does not overcome the problem of different interpretations of the same feature.

Many experts draw diagrams to explain their reasoning (Cheng, 1996), or use diagrams to describe a process (Crowther, 1992). An ideal expert knowledge acquisition tool would be one which could capture this type of knowledge directly.

2.4.1 Geographic or Spatial Knowledge

It should be noted that in this thesis the term *Geographic Knowledge* is used extensively, while in scene analysis *Spatial Knowledge* is the term typically used and much work has been done on spatial knowledge acquisition. In this thesis *Geographic Knowledge* and *Spatial Knowledge* will be regarded as synonymous.

One of the primary ways of representing knowledge in geography is through the use of maps. Maps present knowledge naturally occurring in three dimensions, in a two dimensional graphic form. However maps are produced using information from images produced by sensors on aircraft or satellites, from photographic images and from ground (and sometimes underground) information. Each of these could be regarded as another dimension. The problem then becomes one of representing n-dimensional knowledge in a two-dimensional form (Crowther and Hartnett, 1996b). The information is then used to produce a map showing some specific characteristics of an area (land use, soil type, geology, vegetation cover for example). To produce maps experts use some or all of the information sources (dimensions) listed above.

Experts operating in different domains will use different sources and may produce different maps of the same area showing different features. By using expert system approaches it may be possible to make more use of all the dimensions of knowledge available and its interpretation by multiple satellite image domain experts (Tranowski, 1990).

Egenhofer and Mark (1995) use the term "Naive Geography" to describe the field of study that is concerned with formal models of the common sense geographic-world. They identify a number of elements of Naive Geography which have implications for developing geographic expert systems. These include:

- Naive Geographic space is two dimensional,
- The earth is flat as most representations (maps and images) are flat,
- Geographic space and time are tightly coupled,
- Geographic information is frequently incomplete,
- People use multiple conceptualizations of geographic space. That is, geographic space is regarded differently depending on the application,
- Geographic space has multiple levels of detail. For example, depending on scale, certain features may be grouped or generalised,
- Boundaries are sometimes entities and have a specific tag, but sometimes not,
- People have biases towards north-south and east-west directions and
- Distances do not add up easily because of order of magnitude reasoning

There is also a related problem - that of assigning definitions to features (Kweon and Kanade 1994). In geography most terms are described in natural language, but the definitions are often incomplete or open to interpretation. A method of knowledge acquisition that captures how users think about geographic space is needed.

2.4.2 Classification of Geographic Knowledge

Geographic knowledge differs from knowledge used in non-spatial expert systems in that geographic domain experts primarily use knowledge that is visually oriented.

Often multiple experts with different expertise are required to interpret images to make up a final composite image (Tranowski, 1990). There have been several papers which suggest how to identify and classify knowledge in geographic and spatial systems. McKeown et al (1989) identifies five types of spatial knowledge.

The five types are:

- **Type 1 Knowledge:** identifies scene primitives where a primitive is a readily identifiable object such as a road, a building or an iceberg.
- **Type 2 Knowledge:** is the knowledge of the spatial relationships between the scene primitives. For example, buildings are next to roads or icebergs are surrounded by water.
- **Type 3 Knowledge:** defines collections of objects which form spatial decompositions within the task domain.
- **Type 4 Knowledge:** consists of how to combine information from type 3 knowledge.
- **Type 5 Knowledge:** is used to resolve and evaluate conflicting information.

Tranowski also suggests a three level classification which is basically McKeowns Type 1, 2 and 3 knowledge. The classification is based on:

- The appearance of objects.
- Simple relationships between objects.
- More complex relationships between objects that describe a spatial pattern.

Since this classification is not an advancement on McKeown's, it will not be considered further.

One field of spatial systems receiving considerable attention using knowledge-based approaches is that of *generalisation* (Buttenfield and McMaster, 1991). Generalisation is the replacement of a group of features with a single generalised object.

Armstrong (1991) defines three types of knowledge which includes a non-visual class necessary for generalisation. They are:

- **Geometrical knowledge** which describes features in terms of location and density in a given area. For example, a feature may exist at a specific location and there may be a number of such features in the vicinity giving an indication of congestion.
- **Structural knowledge** which involves the intrinsic expertise of the domain expert to distinguish between features. This is also influenced by the purpose of the required map.
- **Procedural knowledge** which allows the control of the individual generalization operators and algorithms.

The classifications of McKeown and Armstrong are not mutually exclusive. They are in fact complimentary. The McKeown classification appears to expand on Armstrong's geometrical and structural knowledge classifications. Procedural knowledge is an essential component when image processing systems are considered and is not covered in the McKeown et al scheme.

2.4.3 Problems of Spatial Knowledge Classification Schemes

McKeown's classification is the most useful of the three so far discussed when considering the development of a geographic knowledge acquisition tool. However it does have one drawback in assuming all knowledge of a scene is visual. Non-visual knowledge also needs to be incorporated. One example is that of temporal knowledge where past classifications may affect those of the future. For example in agricultural systems, crop rotation and crop seasonality may be useful in producing a classification. Another example is the combination of satellite bands to highlight features; in other words Armstrong's geometric knowledge. Neither of these can be gained directly from the images, but are rather a non-visual aspect which needs to be added to the knowledge acquisition process.

The classification also lacks detail at certain levels. For example, what Type 1 Knowledge is essential and what is desirable when describing scene primitives? A new classification scheme is presented at the end of this chapter.

2.4.4 Unique Characteristics of Spatial Knowledge

Spatial knowledge has all the characteristics of knowledge in other domains but there are other additional characteristics which need to be considered including:

- Position,
- Shape,
- Size,
- Orientation,
- Connectivity,
- Containment and
- Proximity.

The relationship of one object to one or more other objects may determine the identity of a particular feature. It is also possible that features may form a classification hierarchy. A house may be part of a block of houses which may be a residential division of a town for example. This forms the basis of map generalisation.

When working with images the domain expert must rely primarily on visual interpretation of a particular image. Before satellite remote sensing images were available, aircraft images were used as the primary source of data. A major factor in the successful interpretation of these images was the experts' experience. The interpretation was then confirmed by looking at other images containing similar features in a different setting and by ground truthing. Remote sensing information allowed domain experts to investigate other characteristics by investigating individual pixel values and the combination of pixel values at the same point on different wavelength bands. It also allowed them to view distributions of pixel values over an image as a histogram. It is common for domain experts to stretch these histograms to highlight the areas of pixel values of most interest (Richards, 1993).

Many automated image classifiers use a parametric classification of pixels (Wilkinson, 1996). These classifiers rely on a pixel value alone rather than on identifying geometric patterns. This is partly due to the dichotomy between raster and vector representation of features. Raster representations are often classified pixel by pixel, while vector representations are more useful when defining spatial relationships. To put it another way some scene primitives may be identified by per-pixel classification, but beyond that spatial analysis of vectors is required.

In interpreting data, humans can rarely operate with more than two images simultaneously and as a result much information held in other dimensions is overlooked. GIS classification systems overcome that to a certain extent, but there are often large amounts of information which goes unchecked (Openshaw, 1993).

2.5 VERIFICATION OF EXPERT INTERPRETATIONS

Verification has been defined as building a product correctly (without errors) while validation is building the right product; one which meets the user's requirements (McGraw and Harbison-Briggs, 1989).

Verification and validation of knowledge-bases encompass four main activities (Meseguer and Preece, 1995):

- Inspection
- Static verification
- Empirical testing
- Empirical evaluation

Inspection aims to detect semantically incorrect knowledge in a knowledge-base and is a manual operation. This is normally done by either the original expert or, ideally, another expert in the knowledge domain. To aid this process the knowledge should be in a form easily understood by the domain expert.

Static verification is a check for anomalies in the knowledge-base. An anomaly is a static pattern in the knowledge-base structure which suggests an error in the encoded knowledge.

Empirical testing involves applying the knowledge-base to sample data sets. This would include applying the knowledge-base to a range of images after training on one or more image sets.

Empirical evaluation is the validation stage and determines the effectiveness of the operational expert system as far as the final user is concerned. It includes technical performance and applicability.

Verification of classified images presents unique problems. Ground truthing, which is the traditional way to verify geographic classification, is subject to error (Congalton, 1991). Firstly if the feature being studied varies rapidly over time, records of the feature may not exist unless it was being closely monitored. This is particularly true for atmospheric phenomena such as clouds and in agricultural systems around harvest time. Secondly it may be impossible or impractical to ground truth, in the case of mid ocean and polar features for example. Lastly it may be difficult to precisely locate a feature on the ground. In all cases there is some level of sampling and extrapolation. Sea ice ship-borne verification is inevitably a linear transect. Forestry sampling is done with random quadrants.

Verification must be systematic. Vicat et al (1995) suggest a verification model based on knowledge modelling and integrated into the knowledge-base construction life cycle. Rouge et al (1995) support this approach where a formalism based on the structure of knowledge is central to verification.

2.6 THE NEEDS OF A VISUAL GEOGRAPHIC KNOWLEDGE ACQUISITION SYSTEM

Given that geographic knowledge is visual and that domain experts work with images, a graphical system is required. Secondly, the level of computer familiarity may be low. Any system would therefore have to be easy to use and ideally operate in a way the domain expert understands (Kuhn, 1993). Since a domain expert can rarely work with more than two images at a time the system should also collect information from other bands and allow the user to rapidly change between bands. The knowledge generated by the system should be easily verifiable (although the problems of ground truthing are not going to be overcome).

2.7 A PROPOSED CLASSIFICATION OF GEOGRAPHIC KNOWLEDGE

The following classification is derived from and expands on those of McKeown et al (1989) and Armstrong (1991). It is more rigorous and incorporates non-visual knowledge. It consists of six levels of knowledge which are:

Primitive Knowledge about the identification of scene primitives. A primitive is a readily identifiable point, line or areal object which cannot be subdivided into smaller named entities. This includes knowledge about an object's size and shape if relevant.

Relationship Knowledge of the spatial relationships between scene primitives in terms of their proximity, orientation and degree of overlap.

Assembly Knowledge, used to define collections of objects which form identifiable spatial decompositions. This includes knowledge of the spatial density of primitives. This knowledge can be regarded as knowledge needed for generalisation.

Non-Visual Knowledge which helps refine classifications developed using visual knowledge including labelling of scene primitives and spatial relationships. It consists of:

- temporal knowledge of how a scene changes over time.
- algorithmic knowledge, including how to combine bands.
- heuristic knowledge of a non-visual nature.

Consolidation Knowledge used to resolve and evaluate conflicting information.

Interpretation Knowledge of how to combine the other five types of knowledge to produce a classified image.

2.7.1 Primitive Knowledge

A primary function of any automated image interpretation system is to identify the objects that make up a scene. An unclassified image consists of pixels each of which has a value assigned to this. From this image objects or scene primitives need to be extracted. Scene primitives fall into three categories; points, lines and areas. They can each be given a name and they cannot be subdivided. They are the basic building blocks of any GIS. The attributes of objects depend on their category and are summarised in Table 2.1

OBJECT CATEGORY	ATTRIBUTES			
Point	location			
Line	location	Length	shape	
Area	location (of centroid)	size	shape	spectral signature

Table 2.1 Primitive objects and their attributes

One of the primary aims of many systems is to identify and classify these scene primitives. Generally the first method used is a per-pixel classifier. This is followed by a segmentation program which assigns unique identities to each primitive of a class (Gerbrands, 1993).

2.7.2 Relationship Knowledge

Knowledge about the spatial relationship between scene primitives falls into three categories:

- Proximity of the objects,
- Direction of one object in relation to the other, and
- Degree of overlap of objects.

Proximity can be measured in a number of ways. For example, the proximity of the two areal objects' centroids or, the proximity of the closest edge pixels (assuming the two objects do not overlap). Generally the measurement is a fuzzy concept with terms like *near* and *very near* being used.

Direction is generally best measured in relation to two points; for example two centroids. It is normally not worth classifying to less than 45^{o} as users tend only to think in terms of north, south, east and west with north-west south-west south-east and north-east only rarely being used. (Egenhofer, 1995)

Degree of overlap is generally defined using the terms disjoint, touches, overlap, covers, and encloses. Eight basic forms have been identified (Egenhofer, 1991), but many more variations, especially with lines have been described.

2.7.3 Assembly Knowledge

Assembly Knowledge is knowledge about combinations of Primitive and Relationship Knowledge. At its most basic, it describes how related groups of scene primitives are grouped and classified as a larger entity. This may also include other 'Assembled' objects resulting in a hierarchy. These assembled objects or components can then be referred to as a single named entity.

Assembly Knowledge also includes:

- The density of components both by number and by area,
- The ratio of components (number of one object type in relation to others) and
- The required existence of certain key components and the possibility of optional components or atypical components.

Assembly Knowledge is also knowledge required for generalisation. That is simplification by replacing groups of features with a single feature. For example a group of objects representing paddocks could be replaced by a single object representing a farm.

2.7.4 Non-Visual Knowledge

Despite geographic and spatial knowledge being primarily visual, there are a number of non-visual aspects which need to be considered. These may be important in further refining the classification of an image. Armstrong's (1991) *Procedural Knowledge* has already been identified as a necessary inclusion in a classification. Other non-visual knowledge includes knowledge of change over time (temporal) and heuristic knowledge.

2.7.5 Temporal Knowledge

Temporal knowledge may have a visual aspect to it. For example a feature that is in the image in time 1 may not be present or may have moved in time 2. This is particularly true for meteorological applications. Other information however may only become apparent as a pattern when historical data is examined. Such information may include crop rotations, herbicide withholding periods, changes in demand for a particular crop. This knowledge may be held mentally by a domain expert or may come from analysis of data in a database system associated with a GIS. The term for acquiring knowledge from databases is *data mining* (Djoko et al, 1997)

2.7.6 Algorithmic Knowledge

Remote sensing in particular uses algorithms to analyse images. They include image pre-processing algorithms to carry out geographic registration and radiometric calibration, algorithms to combine satellite image bands into composite images and various statistical classifiers such as maximum likelihood and minimum distance classifiers (Lillesand and Kiefer, 1994).

2.7.7 Heuristic Knowledge

Heuristics are shortcuts which domain experts use to complete their tasks. The ability to use such heuristics is often what distinguishes an expert from other practitioners in a field. They may not be based on visible clues, but rather something the expert has found from experience. For example, potato crops are less likely to be planted to the east (with a prevailing westerly wind) of poppy crops because spray drift from poppies is harmful to potatoes. Heuristics are informal methods based primarily on human intuition (Bonnet et al, 1988). They may not work in all cases, but a greater problem is that they are often difficult for a domain expert to verbalise.

2.7.8 Consolidation Knowledge

Consolidation knowledge is the knowledge required to integrate the knowledge-bases which have been generated by the rest of the system. This includes the knowledge gained from a number of training images and from a number of domain experts. If more than one expert is available they may use different rules or visual clues to identify features. The same expert may also use different techniques at times. The knowledge acquired therefore needs to be checked for consistency. Possible outcomes include the most general or the most restrictive rule set. Knowledge at this level will identify and try to resolve conflicts within the component knowledge-bases and produce a single aggregated knowledge-base.

2.7.9 Interpretation Knowledge

Interpretation knowledge is knowledge of how to apply specific knowledge-bases to an unclassified image to produce a classified image for a specific domain. For knowledge-bases that span several domains this will include the type of classification that is being undertaken. For example a knowledge-base containing rules about soil type may be used along with other knowledge-bases in both vegetation and geomorphological classification systems. There needs to be meta knowledge on how this is to be done.

2.8 VISUAL KNOWLEDGE AND KADS

KADS (Knowledge Acquisition and Development System) has been presented as a development methodology for expert systems and classifies knowledge into a number of types (Breuker and Wielinga, 1987, Schreiber et al, 1993). Since this classification scheme is accepted as a standard in many areas it is useful to compare it with the classification of geographic knowledge presented above. The KADS model proposes a four layer model of expertise which is shown in Figure 2.2.

Knowledge category	**organisation**	**knowledge types**
strategic	strategies	plans and meta rules
↓ *controls*		
task	tasks	goals, control terms, task structure
↓ *applies*		
inference knowledge	inference structure	inference actions, roles and domain view
↓ *uses*		
domain	domain theory	concept, property and relation

Figure 2.2 The four-layer model of expertise (Wielinga et al, 1992)

The KADS *Domain Layer* is static knowledge describing a declarative theory of the domain. Knowledge at this level should be represented in a way that is independent of the way in which it is to be used. It should define the conceptualisation and declarative theory of the problem domain and provide the knowledge to carry out given tasks. The other layers contain knowledge to control the use of knowledge from the domain layer (Fensel and Van Harmelen, 1994).

The *Inference Layer* specifies how to use the knowledge from the domain layer. It restricts the use of the domain layer and abstracts from it. The *Task Layer* represents a fixed strategy for achieving problem solving goals. The final level of control knowledge is the *Strategy Layer* which involves knowledge of how to choose between various tasks that when completed successfully achieve the same goal. A formal specification language has been developed to record knowledge in each of the layers (Schreiber et al, 1994).

In terms of the suggested geographic knowledge classification, Primitive, Relationship and Assembly Knowledge are forms of knowledge at the Domain Level under the KADS methodology as is Heuristic knowledge in the Non-Visual category. This knowledge could be used in a variety of different ways to produce products showing different aspects of an area covered by an image

Consolidation Knowledge on the other hand requires knowledge of how the rules are to be applied and is knowledge at the Inference level. Algorithmic Knowledge which contains knowledge of image band combinations and when they should be applied is also at the Inference level. Knowledge at this level will be applied according to the use to be made of the final classified image.

The Task Level in the KADS system is represented by Interpretation Knowledge and shows how to apply the problem solving strategy to the whole image set. Depending on the objective of the system different strategies can be used, including the masking out of areas not relevant in a particular domain.

In the suggested geographic knowledge classification, there is no equivalent of the Strategy Layer as no attempt has been made to create a knowledge category which contains alternate ways of classifying images. This could change in the future when choices need to be made between using rules developed from machine learning, neural nets or rules elicited from domain experts to classify images.

2.9 CONCLUSION

Geographic knowledge used in interpreting remotely sensed images is essentially visual.

When developing a theory of geographic knowledge acquisition, it is desirable to develop a general model of the knowledge types which are to be acquired. Once this has been established, tools and techniques for eliciting this knowledge can be developed. The proposed six level geographic knowledge classification, which has equivalences in the KADS methodology, is such a model and will be used through the rest of the thesis.

Chapter 3. KNOWLEDGE-BASED SPATIAL INFORMATION SYSTEMS

This chapter investigates Geographic Information Systems (GIS) and their relationship with remote sensing and knowledge-based systems. The data structures and processing requirements necessary in a spatial system are considered.

To implement rules on spatial reasoning there must be an understanding of the way GIS information is stored and processed. This is the basis of selecting appropriate knowledge-base structures for use in a geographic environment.

The final section of this chapter is a survey of expert systems which have been developed and used in geographic and remote sensing environments.

3.1 GEOGRAPHIC INFORMATION SYSTEMS

It is not an aim of this thesis to develop a geographic information system (GIS). However, because a study of knowledge acquisition for expert systems used with GIS is the primary aim of the research, there must be a discussion of GIS systems. In such a discussion it is useful to begin with a definition:

"A geographic information system (GIS) is a computer-based information system that enables the capture, modeling, manipulation, retrieval, analysis and presentation of geographically referenced data" (Worboys, 1995, p1)

The fundamental part of a GIS is a database, and fundamental to that is the development of a data model. Another major component of the system is a user interface with which a user can input, manipulate and display information from the database in the form of a map or image. In current technology this takes the form of a graphical user interface. Therefore one needs to be able to input information into a system, store it and then manipulate it. One of the functions of an expert system linked to a GIS is to pre-process image information and make it available to the GIS.

Burroughs (1990) defines the following five components of a GIS:

- Data input and verification,
- Data storage and database management,
- Data output and presentation,
- Data transformation and
- Interaction with the user.

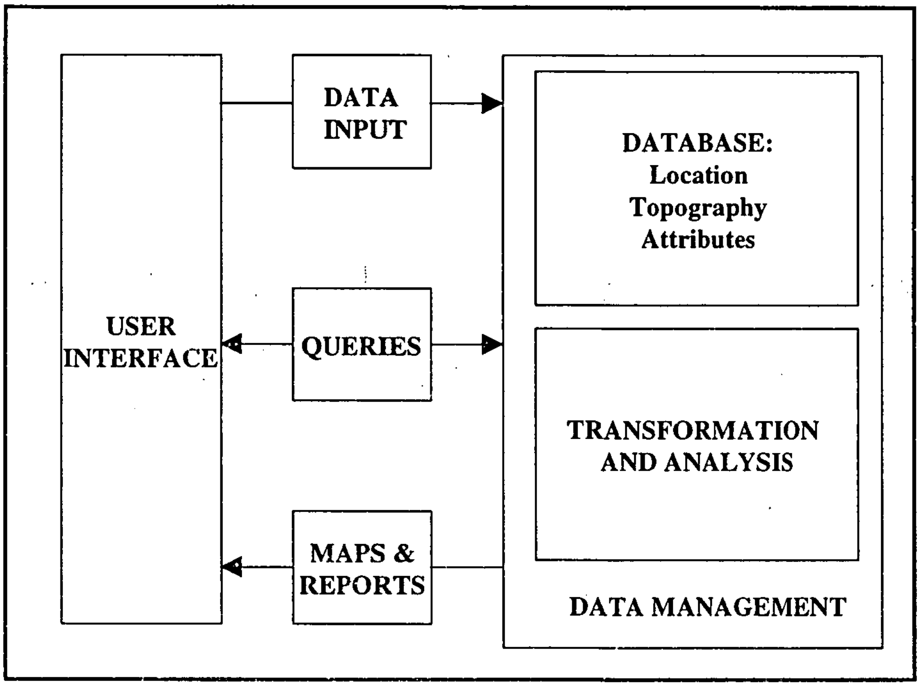

Figure 3.1 Components of a GIS

These are shown in Figure 3.1 where output and presentation is subdivided into *queries* and *maps and reports*. In this model *data storage* (or *database*) is grouped with *transformation and analysis* under *data management*.

3.1.1 Data Input and Verification

There are many sources of data that can be used as input to a geographic information system. These include traditional text files, field observation data, drawings, maps, tables and airborne and satellite sensors. Because of this, the input devices used include scanners and digitisers as well as traditional keyboards and storage systems. Input may also come from remote computer sites over communication lines.

One problem, which is not confined to GIS, is the quality of input data (Star and Estes, 1990) which in turn affects the quality of the output. Since GIS allow a user to combine data from a number of sources, many have built in verification routines which alert users to potential errors. Despite this, it is up to the user to record which are the more unreliable data sources. By using an expert system to preprocess some of this data, the quality of the data can be improved.

3.1.2 Data Output and Presentation

The data output component of a GIS incorporates many of the features typical of computer aided design packages (CAD). The output is displayed on a VDU screen or can be plotted for hardcopy. This output may be in the form of statistical reports, maps including three-dimensional topographic representations and various types of graphs.

3.1.3 Data Transformation and Analysis

This component of the GIS allows data to be transformed into information. Individual packages have different features. Burroughs (1990) classified data transformation functions into:

- Rotation,
- Stretching to a new scale,
- Transformation of scale and projections,
- Zooming,
- Joining,
- Polygon overlay,
- Smoothing and
- Data reduction or generalisation.

Since these functions are built into GIS it makes sense to use them before passing image data onto a coupled expert system (or knowledge acquisition system for the expert system). Among the techniques which can be used are smoothing to remove uncharacteristic values in the data, and scaling, which can be used when images from different sensors with different resolution are being used.

3.1.4 Data Storage and Database Management

There are two very different implementations of GIS on computers. The first uses a two dimensional grid with every location in the grid being georeferenced to the earth's surface. The grid location stores a number which represents some attribute of the earth's surface. This is referred to as a raster based system. Raster data structures are the simplest with data organised in a cellular data structure (Star and Estes, 1990). This cellular organisation is represented as an array over space with each cell containing a value for the parameter of interest. Several of these structures can be used to superimpose data. For example one array may contain elevation information while a second may contain vegetation types. From this it may be possible to relate vegetation changes to slope and aspect.

Simple raster arrays often have the rows of the array oriented east-west and the vertical dimension north-south. There are however two main limitations to these simple arrays. Firstly there is a finite ability to specify location and secondly the cells represented in the array may not be evenly spaced. This and the size of an area representing one cell affect accuracy. A further conceptual problem is that being pixel based, a raster system cannot represent objects beyond the pixel level. Rather pixels with particular values can be given a label assigning them to a classification. However raster systems have the advantage that they can build up map layers for each attribute under consideration. These layers can then be combined. Their second advantage is they can deal directly with satellite images which are pixel based. The reasons they are primarily used in this study are the ease of using them with data from image processing systems and the ease of using them to carry out spatial analysis on satellite images.

An alternative implementation is to have a grid which stores locational data according to its data type as objects. Hence point data is stored as an X Y coordinate. Lines are stored as a pair of X Y locations and area objects or polygons as a series of X Y coordinates. This is known as a vector based system. Vector data structures are based on elemental points or nodes whose locations are known to arbitrary precision. Hence a circle could be stored by specifying its centre and radius while a raster definition of the same object would require coding of all cells that correspond to the circumference. As well as this the nodes are used to carry metric information about the object. There are several forms of vector data structures available including :

- A whole polygon structure,
- A Dual Independent Map Encoding structure,
- A arc-node structure,
- An relational structure and
- A digital line graph.

(Star and Estes 1990)

These are shown in Figure 3.2. In a whole polygon structure, an image or layer is divided into a set of polygons (Figure 3.2 (a)) which are encoded as a series of locations defining a boundary. The Dual Independent Map Encoding (DIME) structure (Figure 3.2 (b)) is a structure where each polygon is uniquely identified. Nodes marking the point of joining of three or more polygons are encoded along with the location of the centre of the object (Bartelme, 1996). By coding two nodes and two adjacent polygon centroids the face or join between the polygons is defined. This was originally used for digitizing streets in urbanised areas with the nodes being street intersections and blocks being represented by polygon centroids (Star and Estes, 1990).

Arc–Node structures places objects in a hierarchy. Nodes are defined first, followed by lines which in turn build up polygons. Figure 3.2 (c) shows a polygon made up of a series of nodes connected by arcs. Another variation on the arc-node structure is the relational structure shown in Figure 3.2 (d) where each element in the structure is assigned a unique identification. In the example only nodes are shown. The location of the nodes are shown in one structure and their characteristics in another linked (related) by their identification number. Arcs are also given unique identification numbers but are defined in terms of their start and end node's identification numbers. Likewise polygons are defined in terms of their component arcs.

The final structure is the Digital Line Graph Structure. Like other structures, the location nodes are encoded. In addition there is an extensive coding scheme for attribute information of the elements (Figure 3.2 (e)). An example of this scheme is: the code 050 0004 indicates a hydrography feature (050) which is the point (node) of a stream entering a water body (0004).

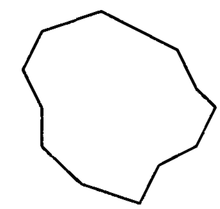

a) Whole Polygon Structure

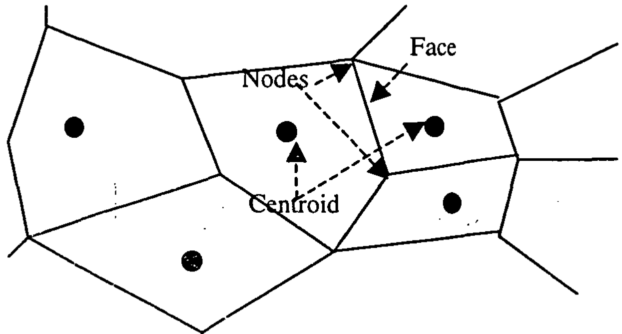

b) Dual Independent Map Encoding

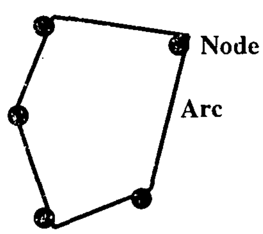

c) Arc – Node Structure

NODES

Number	Easting	Northing
1	126.5	578.2
2	218.6	581.2
3	224.2	470.4

TRAFFIC

Number	Control	Crosswalk
1	Light	Yes
2	Sign	Yes
3	None	No

d) Relational Structure

Major Code	:	050	Category:	Hydrography
Minor Code	:	0001 – 0099	Description:	nodes
Nodes	:	050 0001	Upper end of stream	
		050 0004	Stream entering water body	
		050 0005	Stream leaving water body	

e) Digital Line Graph

Figure 3.2 Types of vector representations

Satellite images are in raster form, therefore it is necessary to process them pixel by pixel which makes a raster based GIS attractive. Raster based GISs have the problem of only being able to represent features to an accuracy dependent on the area represented by a single pixel. Objects can only be indirectly defined as groups of pixels after an image has been segmented. These pixel groups can then have information about them stored, but not directly as part of the group.

This problem is solved by vector based systems which allow users to specify locational data as accurately as they wish (Breunig, 1996) and attach information about the object being defined directly to the objects nodes. Hence vector based GIS can produced finely detailed maps to high cartographic standards with information about the objects represented linked directly to the GIS database. Queries to a vector GIS database can be made across all the data sets within the databases using information attached to the vectors. Since satellite images are raster based, they must be first processed at the pixel level to segment them before the resulting pixel clusters are converted into vector objects.

Despite this vector – raster dichotomy, many of today's GIS have incorporated both representations. Transfer between the two representations is still fraught with difficulty however, especially when two vectors are closer to each other than the pixel resolution of the raster system (Ibbs and Stevens, 1988).

There are many commercially available GIS packages available. The one chosen to interface with the KAGES tool-kit, to be described in chapter 6, is IDRISI. This package has the advantage of being raster based but with vector capabilities. It is also low cost, it incorporates spatial analysis capabilities and it has the capability of interfacing with the Microsoft Access database system which is to be used in conjunction with a specific application of the tool-kit.

3.1.5 Interaction With the User

The user interface is probably the weakest area of most GIS in use today (Hazelhoff and Gunnik, 1994). One problem is that the interface may provide both the overall user interface and the means of creating queries and models. More traditional database systems such as Microsoft Access have a user interface, then separate interfaces for developing queries and models.

The display used in a typical GIS presents a visual display of geographic data in the form of charts, tables, maps, text or pseudo three dimensional images. The process of receiving data in these forms requires a human user to perform the tasks of detecting features, discriminating content and identification. Therefore it is up to system designers to present data in a form which will not be misinterpreted (Hearnshaw, 1994).

It should be noted that a visual interface is not the only means of interacting with a GIS. For example sound and movement are two other sensory modes which can be used to both input and represent data. In this thesis the discussion will be restricted to vision.

GIS provide a user with the means to visualize large volumes of spatial data. This has the advantage of allowing a user to rapidly comprehend the complex spatial relationships and hence improve performance on tasks involving that form of data. This leads to the importance of a user being able to interactively modify spatial data and observe feedback in real time (Medyckj-Scott, 1994).

An aim of designing a human computer interface is to make interaction between a user and the computer easy. That is the user should be concerned with the working on the data and not how the tasks are performed. To meet this aim three factors of, user, the task and the tool need to be considered.

The first factor, the 'user' is the most variable, because of variation in individual users skills, perceptions and needs. Because of this interface design is moving to accommodate users throughout their interaction with the system (Medyckj-Scott, 1990).

Tasks are activities which a user needs to fulfill to achieve an objective, Therefore there is a need for task analysis when building a system. This includes identifying:

- the types of tasks a user performs.
- the extent to which tasks vary from one occasion to another.
- how often tasks are performed.
- how tasks vary between different user groups.

The result should capture the similarities and differences among tasks and their interaction requirements.

A GIS is a tool designed to display and operate on spatial information. To be useable it must be able to transform data from an internal representation to one a user is familiar with and can interact with (Medyckj-Scott, 1994). This has two aspects, control representations and display representations. Control representations affect a users perception on how easy the tool is to use. They include functionality and utility. A tool should allow a user to focus on what their objective is rather than how the tool works. Display representations determine how the data is presented to the user. This should be in a form familiar to the user.

These considerations related to user interface design in typical GIS also have to be considered when developing a knowledge acquisition system where the domain experts may be GIS users. These users look for a utility which is familiar to them and provided by a GIS. For example, for many users the actual image from a satellite band is often of less interest than the histogram of the data of that image. Many users work directly with the histograms when classifying.

3.2 REMOTE SENSING

A simple definition of remote sensing is the:

"...collection of information about the properties of an object without physical contact being made with the object" (Mather, 1991, p 140).

In the present study a more useful definition is one which restricts it to:

"...measurement of electromagnetic properties of a surface or object without being in contact with it" (Davis and Simonett, 1991)

In recent years large volumes of data have become available from satellite based sensors. Information can be displayed as images, but those images can record reflectance and emittance over a number of spectral bands. In this study image sets from three different satellites have been used. These are:

- NOAAH AVHRR (National Oceanographic and Atmospheric Administration Advanced Very High Resolution Radiometer) ,
- SPOT HRV (Système Probatoire d'Observation de la Terre, Haute Résolution Visible) and
- Landsat .

3.2.1 Active and Passive Remotely Sensed Data

The data collected by passive Remote Sensing Systems (RSS) is from radiation recorded by sensors at various wavelengths in the visible through to the thermal infrared and microwave part of the electromagnetic spectrum (Tables 3.1 and 3.2). Visible light sensors collect information on solar radiation which is reflected back into space from earth. Infrared and microwave sensors collect data on energy emissions from the earth's surface. These sensors that record radiation emitted from the sun or the earth and are called passive sensors.

3.2.2 The Electromagnetic Spectrum

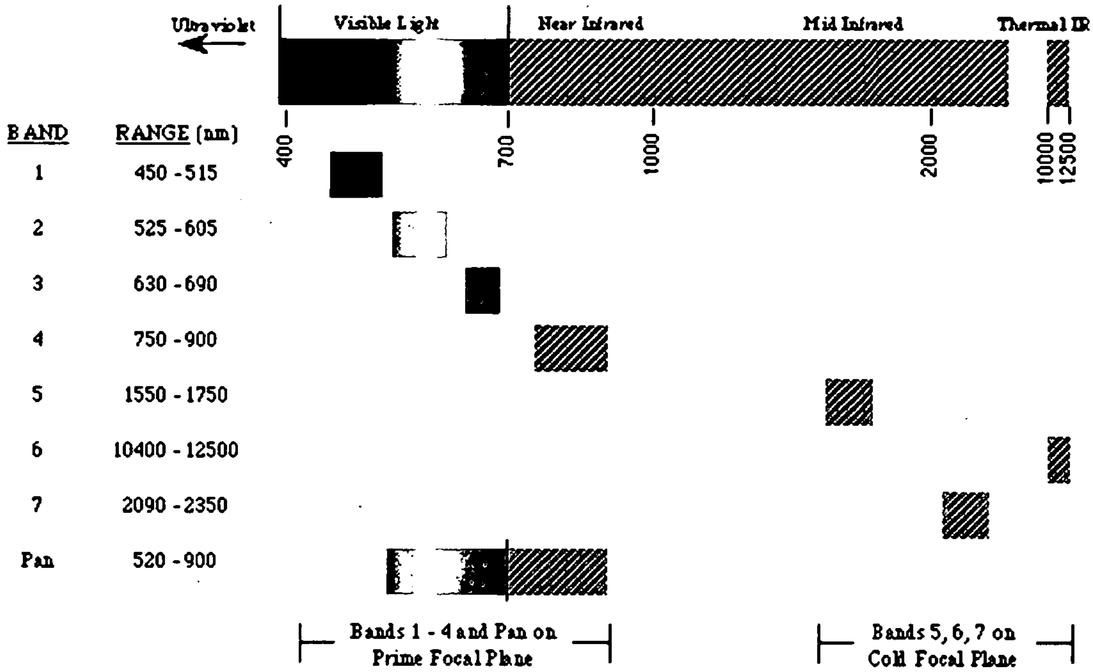

Table 3.1 The electromagnetic spectrum showing spectral bands covered by Landsat (from http//ls7pm3.gsfc.nasa.gov/mainpage.html)

SPOT 20 meter resolution		
Band 1	0.50 - 0.59 μm	Visible
Band 2	0.61 - 0.68 μm	Visible
Band 3	0.79 - 0.89 μm	Infrared
NOAA AVHRR 4Km resolution		
Band 1	0.58 - 0.68 μm	Visible
Band 2	0.72 - 1.10 μm	Near Infrared
Band 3	3.55 - 3.93 μm	Thermal Infrared
Band 4	10.50 - 11.30 μm	Thermal Infrared
Band 5	11.50 - 12.50 μm	Thermal Infrared

Table 3.2 SPOT and NOAA AVHRR spectral characteristics.

Active systems on the other hand generate their own energy, bounce it off the earth's surface and read the amount of energy which returns. Some microwave systems, most notably radar, are the most common active systems. These have the advantage that they are not affected by atmospheric conditions but they are not as readily available as passive systems.

Electromagnetic radiation can be thought of as consisting of waves of energy or as packets of energy called photons. This energy can only be detected as it interacts with matter (Sabins, 1997). Therefore the particle model is useful for understanding the way detectors measure and count photons of a particular frequency and assign a digital value to the count. For example, a light meter measures the intensity of light by the interaction of photons with the light sensitive photodetector in the meter. This produces an electrical signal proportional to the number of photons.

Table 3.1 shows the electromagnetic spectral ranges covered by Landsat with radiated energy separated into bands of specified frequencies. SPOT and NOAA AVHRR spectral ranges are shown in Table 3.2.

3.2.3 Problems With Remotely Sensed Data

One of the major problems with remotely sensed images is the amount of variability from one image to the next and the difficulty of correcting distortions inherent in the images. Some of the problems encountered include:

- The variability of atmospheric conditions,
- Varying illumination conditions depending on the season and time of day,
- Image registration difficulties and
- The effects of variation in land elevation.

In this thesis we are concerned with passive remote sensing and the energy reflected and or emitted from the earth's surface is far from uniform. Atmospheric conditions, time of day and season compound the problem. Hence there is need for calibration to eliminate or compensate for these effects. (Lillesand and Kiefer, 1994).

Georeferencing of satellite data is involves the assignment of coordinates to individual pixels relating each to a specific ground location. Although remotely sensed data from satellites is georeferenced, that referencing is subject to error. There are various algorithms available to correct for this. The simplest method is to extrapolate from points on the ground with known locations (by using a global positioning system for example).

3.3 THE NEED FOR DATA STRUCTURES

Data in any system needs to be organised so systematic processing of that data can take place. As the discussion above has indicated, Geographic Information Systems are no exception. There are several different ways to organise data and these will be discussed below. Basically there is need for a structure for the database and there is need for a structure for the knowledge-base.

3.3.1 Evaluation of Vector versus Raster Data Structures for this Study

The choice of a data structure is based on the processing requirements of a knowledge-based system. Given the different type of knowledge that can be stored (see section 2.6), there is a need for both pixel reasoning and reasoning about spatial objects. Although it is possible to have both vector and raster structures it is easier to standardise on one. In this case raster representation will be used as the primary method since most of the processing will be done via image processing software. However experts naturally work with objects (Couclelis, 1992). To accommodate this, areal objects will have a minimum bounding rectangle associated with them which will be stored along with the object's centroid. In the case of line and point objects, a vector type representation will be used. More complex objects will be stored in frame-like structures held in a separate database. For example a particular Assembly Knowledge object may be defined by having certain key Primitive objects as mandatory constituents (To be a settlement at least one church, hotel, store or service station is required). Reasoning would then be based on rules stored in a knowledge-base.

This combination of raster and simple vector structures has advantages for developing a generic spatial expert system as rules about specific features as well as information about objects could all be stored and applied to different images. Since the primary input data from the satellite image sets is raster based, raster processing rather than vector processing has been chosen.

3.3.2 Knowledge-base Structure for a Geographic System

Once a geographic data structure has been decided upon, the next step is to determine the structures necessary for data or knowledge to be stored in a knowledge-base. The knowledge can be represented using one of several knowledge representation schemes (Leung, 1997) including:

- Predicate logic,
- Rules,
- Semantic networks,
- Frames and
- Object oriented representations.

Predicate logic is the basis of logic programming languages such as PROLOG. Knowledge is represented in the form:

$$(\forall\ x)\ (\text{triangle}(x) \rightarrow \text{polygon}(x))$$

which means all triangles are polygons. This notation has the advantage of being concise and directly translatable into a programming language, but has the disadvantage of being difficult for users to understand and verify.

A common way of representing knowledge in an expert system is in the form of IF THEN rules which have the form:

IF condition THEN conclusion,

The predicate logic statement above now becomes:

IF X = triangle
THEN
X = polygon

One of the problems with spatial knowledge is finding a way to represent it so that it can be reasoned with and can be visually verified by a user who is unfamiliar with knowledge representation schemes (Shea, 1991). Rule based systems, where the rule can be verified by the user who can also see the rules effect, solve this problem and have been extensively used in spatial expert systems.

Semantic Networks are sometimes called propositional nets since the statements they represent are either true or false. Figure 3.3 is a simple semantic network

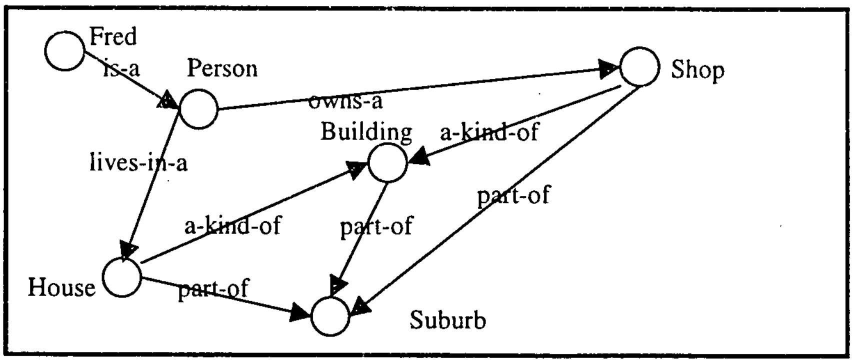

Figure 3.3 Semantic Network

It consists of nodes and arcs, with the nodes often referred to as objects and the arcs as links. Link statements fall into types such as IS-A and A-KIND-OF. Schematically this is a useful representation but is difficult to generate automatically.

Frames are a means of representing knowledge using a table-like structure (see Figure 3.4). They consist of slots which contain information about the concept described by the frame. Frames are useful for describing a subject which has much default knowledge. Rules and pointers to procedures can be placed in slots if needed (Giarratano and Riley, 1993). Frames can be linked together in a hierarchy which allows lower level frames to inherit characteristics defined in higher level ones.

```
Frame data history
        Slots
                Data title          =road map
                Scale capt          =50000
                Scale maxuse        =20000
                Data type           =vector
                Data source         =digitized
                Data dimension      = [500, 601]
                Data reliab type    = perkal
                Date acquired       = [01 01 1993]
                Data currency       = [5]
                Locn od             = [240.0  740.0]
                Locn acc            = [relative / absolute]
                Routine hisatory    = [[compiled [digitized aerial photography]] [CEC gazeteer]
Endframe;
```

Figure 3.4 Part of a frame for storing meta-data (from Miller, 1994, p. 148)

The choice of a knowledge-base structure is dependent on the type of system being developed. Representation in a knowledge-base has traditionally been done with frames and rules. This creates problems when dealing with visual knowledge where an object is best described in terms of a diagram.

Despite the problems raised above, the proposed representation to be used in this thesis is the IF – THEN rule. Such rules can be applied at both the object and pixel level and is verifiable by the domain expert. The advantage of rules over frames is that they are a more primitive structure, hence if knowledge is acquired and stored as rules, they can be incorporated in a frame-based system by placing them in an appropriate slot. To operate in the reverse way would require analysis of each slot in a frame-based system.

It is the role of a knowledge acquisition system to assist in acquiring knowledge to be used by an expert system. To be generic the representation needs to be as simple as possible. In a geographic context it is not the function of a knowledge acquisition system to identify individual objects and relationships. Instead the objective is to provide a representation of the knowledge required to identify individual features and relationships.

3.4 SURVEY OF KNOWLEDGE-BASED GIS AND RSS

There are three components to this discussion, remote sensing, expert (or knowledge-based) systems and GIS. In this study, expert system technology will be used to aid interpretation of remotely sensed data which will then be used in a GIS. Specifically, this study looks at building the knowledge-base.

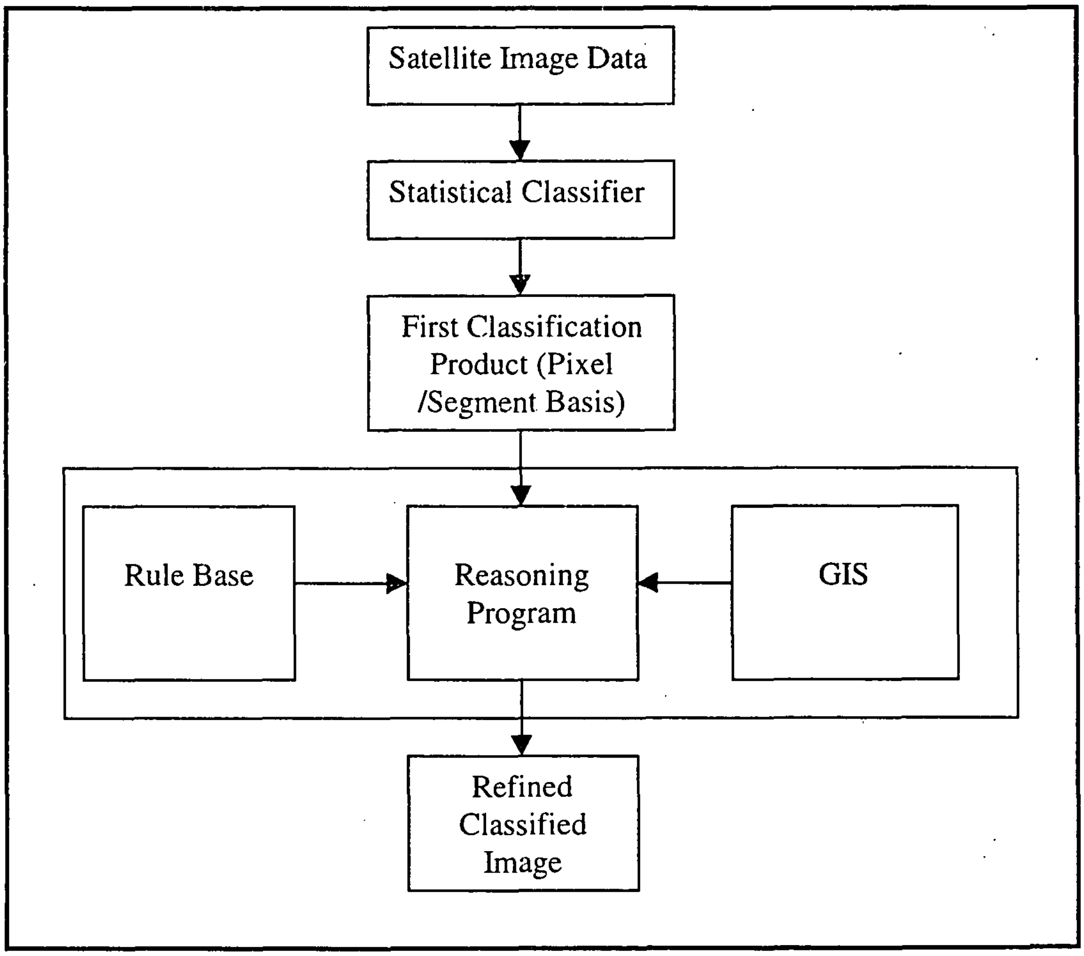

Figure 3.5 An integrated RSS, expert system and GIS (after Wilkinson and Burrill, 1991)

One problem with most GIS is that they contain insufficient image processing capabilities to handle remotely sensed data (Hinton, 1996). As a result there is normally some sort of coupling with an image processing system. This can range from loosely coupled systems where the GIS and image processing system are totally separate through to systems with a common interface and finally fully integrated systems (Figure 3.5). Most GIS and image processing systems are still at the first stage.

For systems at the first, loosely coupled stage of integration, an image must first be preprocessed using the GIS and the resultant raw bands exported to an expert system (which must have image processing capabilities). The bands are then processed by the expert system and the result, usually in the form of an image, is returned as a new layer to the GIS. A knowledge acquisition system operates in a similar way. However the results in the form of rules are stored in a knowledge-base rather than being passed back to the GIS.

A customised solution to classifying remotely sensed images is to avoid using a GIS altogether. Processing remotely sensed images can be done using an expert system acting as an image processing systems with no interaction with a GIS. The expert system accepts image input and produces a classified output. Icemapper (Williams et al 1994, Williams et al 1997) is typical of such a system. It accepts NOAA AVHRR data and performs a rule based classification building up evidence for the classifications. This is regarded as a first best guess with a high degree of accuracy. The system then allows a user to adjust any thresholds to tune the final image. This system does not currently have an independent rule base and is not linked to a GIS. Rather it is a stand alone expert image interpretation system.

Another approach to classifying remotely sensed images is not to use an expert system component and to rely on statistical classifiers. A typical approach to classifying remotely sensed images has been to use a raster based GIS and some of the standard band combination algorithms to do a supervised classification of images on a per pixel basis.

For example Schotten et al (1995) used ERS-1 (synthetic aperture radar) data and SPOT XS to discriminate between agricultural crops in the Netherlands.

The general approach was to:

- Derive the geometry of agricultural fields from topographic maps and a SPOT XS image and store it in a GIS,
- Process the ERS –1 data,
- Classify the images using a supervised maximum likelihood classifier and
- Store the resulting validated crop type in the GIS

An overall accuracy and reliability of 80% was obtained. This particular system had the added advantage of working with radar data which solved the problem of passive sensors being affected by cloud.

The most flexible systems are those which combine the features of expert systems, remote sensing systems and GIS. Hence a system would initially store image data and other map data as layers in a GIS. This information would then be processed using image processing techniques and the integrated expert system would work on the various layers in the GIS. The final results would then be stored in the GIS database.

An example of such a system is described by Kontoes et al (1993) where SPOT images and digitised soil and road network maps were placed into an ARC/INFO GIS. All the information was registered and held as data layers. The knowledge-base consisted of image context rules and geographic context rules. The result of the system was a thematic image depicting crop acreages.

The processing was done on a per pixel basis using class likelihood from a statistical image classifier to produce a first guess, the rules were then applied from the rule base to refine the classification. The result using the statistical image classifier alone was 64.5% overall, increasing to 77.3% over all classes when the rule base was applied. A similar system is described by Hartnett et al (1994).

Many of the current expert systems in this group have operated on a per pixel basis like the two mentioned above. The next generation of expert systems will incorporate spatial analysis functions (Openshaw and Clarke, 1996).

3.5 CONCLUSION

Most current GIS are not particularly effective in handling remotely sensed data. This is partly because most GIS are vector based, while the satellite images are raster based. More GIS are now being developed to handle both data forms and convert between them. Despite this, GIS tend to lack spatial analysis capability (Openshaw and Clarke, 1996). One solution to this is to couple an expert system with the GIS. The GIS could perform pre-classification, identifying scene primitives using statistical classifiers. The results could then be passed to an expert system to refine the classification and perform spatial analysis. The final classified data could then be passed back to the GIS.

To develop an expert system, knowledge acquisition is necessary. Currently this has been done either by fully automated (machine learning) techniques or completely manual techniques with rules being hand crafted into a customised system. Future GIS will have integrated knowledge-bases and improved spatial analysis capabilities.

Chapter 4. TRADITIONAL KNOWLEDGE ACQUISITION TECHNIQUES

The aim of this chapter is to examine established knowledge acquisition techniques which have been applied to other domains and investigate their applicability to spatial systems. Knowledge acquisition techniques fall into two main categories, manual and automated. The manual techniques themselves are divided into direct techniques (where the expert is asked directly) and indirect techniques (where the expert is asked to carry out a task which can be used to infer the knowledge indirectly). The automated techniques are divided into inductive, deductive and neural network methodologies.

4.1 INTRODUCTION

Traditional knowledge acquisition can broadly be classified into two groups of techniques (Figure 4.1). The first group requires interaction with a domain expert to extract expertise manually.

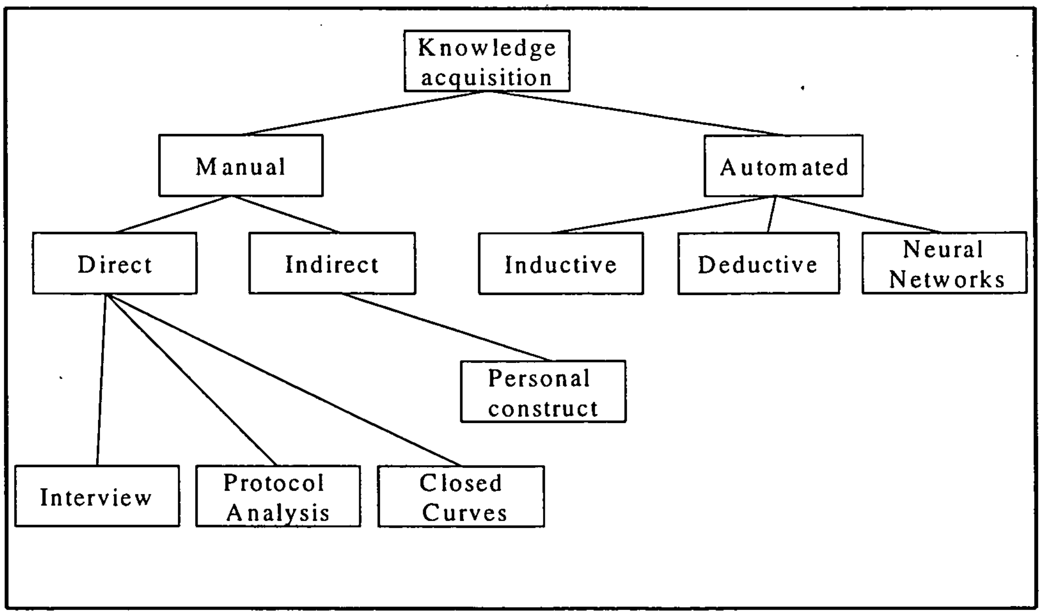

Figure 4.1 Relationship of traditional knowledge acquisition techniques

The second group of techniques is associated with machine learning. This chapter will describe, compare and contrast the various methods and evaluate their applicability in the development of a computer assisted spatial knowledge engineering system.

Traditional manual techniques are used extensively when working with a domain expert in a spatial context. This has not been due to their appropriateness, but rather the lack of a visual knowledge acquisition tool. These methods will be discussed and it will noted that even with computerised tools, many of the manual techniques are still required at various stages during the knowledge acquisition process.

It should be noted that these traditional techniques are used effectively in non-spatial domains and their rejection in a spatial domain in no way suggests they would not be useful outside the spatial context. The following discussion will assume a spatial framework and hence a domain expert working with primarily visual stimuli.

4.2 MANUAL TECHNIQUES

The term manual techniques is a bit of a misnomer as many of the techniques about to be described can have some degree of automation (Price 1990). They all have significant domain expert involvement, with the domain expert providing the expertise rather than learning from data. They can be classified as direct or indirect methods (Olsen and Rueter, 1987). The direct methods normally result in the generation of rules from information directly stated by a domain expert. Indirect methods result in a clustering of objects in the domain as well as the production of rules without directly asking a domain expert to describe or demonstrate problem solving methods.

4.2.1 Interviews

The classic means of acquiring expert knowledge has been through interviews with the domain expert. Much has been written about interviewing techniques and interview types (Welbank, 1990). Early checklists for the feasibility of the development of expert systems included statements such as 'An expert must exist and be available' and 'The expert should be a true expert' (Waterman, 1986).

Generally management of the organisation involved or the user representative would suggest who the expert in the field was and the knowledge engineer would make an appointment to interview them. In practice there is often more than one expert, and it is also common that no single expert can cover the entire domain (Medsker et al, 1995, Barrett & Edwards, 1995).

Whatever type of expert system is being developed and even if the final technique to be used is a machine learning technique, there is a need to interview the domain expert(s) at many of stages through the development of the system. Interviews with the domain experts are therefore an essential part of any expert system development.

The initial interviews with a domain expert would generally be of the unstructured type (Welbank, 1990) where the knowledge engineer would try to get an idea of the scale and scope of the problem. The primary aim of interviewing at this stage is to develop a feasibility report. The project would then either continue or be terminated. If the project continued the next series of interviews would become more and more focused with the knowledge engineer eliciting a series of heuristics which the domain expert used in the course of their problem solving. These would be added to a knowledge-base. Many early expert systems were developed using knowledge elicited primarily by interview and then hand crafted into a customised expert system.

The interviewing process has several problems associated with it which led to the term 'knowledge bottleneck' being coined for problems in the development of expert systems (Wooten and Rowley, 1995). They include:

- Experts who have trouble verbalising what they do,
- Experts who in practice do something different from what they say,
- Misinterpretation by the knowledge engineer and
- Reliance on shallow knowledge.

As a result of these problems there has been a move to search for other methods of knowledge acquisition. Having stated that, it should be noted that a large percentage of knowledge acquisition is still done by the traditional interview technique (Cullen and Bryman 1988). This is despite the availability of systems which try to use domain experts for only occasional guidance (Sleeman and Fraser, 1996)

Interviews are an important knowledge acquisition tool for spatial expert systems. Initial interviews give an indication of the nature of the problem, the time taken by a domain expert to produce a solution and an overview of how that solution is arrived at. Problems mentioned above with the technique are compounded in the spatial domain since domain experts who manually interpret images normally use the images to describe how they work. The problem of verbalising a visual problem and its solution is difficult, even for a domain expert who can communicate effectively since the method used for gaining and transferring knowledge is not primarily verbal Interviews cannot be totally eliminated from knowledge acquisition, even with a computerised tool-kit. A four phase approach to interviews has been suggested by Wooten and Rowley (1995) The phases are:

- Descriptive Elicitation where the knowledge engineer is learning the domain, language and important cues and labels used by the domain expert,
- Structured Expansion where relationships between domain concepts are expanded,
- Scripting where the procedural knowledge of solving the problem is studied and
- Validation where the knowledge acquired is assessed for accuracy.

This kind of interview strategy is necessary in the development of spatial expert systems. However, a combination of this strategy with some of the other methods discussed below is needed to be fully effective in a visual domain.

4.2.2 Observation And Protocol Based Analysis

Observation is the technique of watching the expert in action as he or she solves a problem (Neale, 1989). On its own this may not be particularly useful as the knowledge engineer may not be able to interpret what the domain expert is doing.

Protocol analysis involves one extra step where the domain expert is encouraged to verbalise decisions taken on the way to a solution (Burton et al, 1987).

Observation and protocol analysis overcome one of the problems mentioned in interviewing, that of the expert actually using different heuristics when solving a problem as compared with those used when asked to state how a problem was solved. Senjen and Mee (1993) used the technique of participant observation where a knowledge engineer became part of the domain experts team for an extended period of time. The advantage was that they became familiar with the way the expert operated and the cultural setting they operate in.

There are still disadvantages however. The knowledge engineer being on site when an 'interesting' case occurs (which may be quite rare) is not always possible unless the knowledge engineer is also on call when the expert is on call. One way around this is to get the knowledge engineer to redo some atypical or interesting cases with the domain expert. In some case this may still have the problem of removing some of the stimuli which may have been present during a 'live' problem solving session. The expert may have revised their method because of hindsight, specific atypical cases being the ones that are most likely to be remembered.

For spatial systems, observation of the experts problem solving technique is a useful method because historic images are usually available and their interpretation would not normally be different from that of a 'fresh' image. The domain experts can also describe what they consider to be both typical and interesting cases. Building a computer assisted observation/protocol function is difficult since observation usually means looking at and recording. However it may be possible to present a digitized image and by a combination of other techniques be able to 'record' the actions undertaken by the expert, in effect 'observing' the techniques being used. The philosophy behind this knowledge acquisition technique could be applied to a graphical spatial acquisition tool and will be expanded on in chapter 5.

Interrupt analysis is a variation on observation where a domain expert continues with a task until they get to a point that the knowledge engineer does not understand, at which point they are interrupted. The methodology then moves towards the interview techniques. The obvious disadvantages of this method is that it can only be done on past or hypothetical cases. Once an interruption has occurred it may be impossible for the expert to resume the consultation. In the case of spatial systems where there is the ability to use images in a less pressured environment this technique can be very useful.

Once again building a computer assisted tool to perform interrupt analysis is difficult as the expert is expected to operate in a verbal way interpreting what they see on an image or group of images. Hence a mechanism which allows the recording of the experts manipulation of an image would be most appropriate if this type of tool was to be included in a tool-kit.

A variation on protocol analysis is the technique of asking the domain expert to draw *closed curves*. This techniques has a lot of promise for any spatial system as it requires the domain expert to delineate objects or areas in two dimensional space. By drawing a curve around related objects experts can indicate examples of spatial relationships. Considering this kind of operation is the basis of any spatial system, it makes sense to use it or some derivation of it to develop a geographic knowledge acquisition system.

In a computer-aided system an expert could do this in two ways. Firstly the expert could use a mouse-driven program to outline a specific spatial object. A problem could occur if the object boundaries are not clear cut. This method only defines the boundaries of objects and says little about the characteristic of the object itself. A better method may be to get an expert to 'point' to a typical example of a spatial feature and used the information from that point (subject to appropriate thresholds) to find similar areas. That is the expert determines what to look for and the system then finds similar areas, effectively drawing closed curves. These can then be reviewed by the expert and thresholds adjusted where necessary. Rules and objects could then be generated based on the threshold information. Finally the most important thresholds could be identified by the domain expert.

4.2.3 Repertory Grid Analysis

One of the most popular indirect knowledge acquisition techniques is *repertory grid analysis*, which is a technique derived from Kelly's (1955) personal construct theory. It is the basis of a number of computer assisted knowledge acquisition tools such as AQUINAS (Boose & Bradshaw 1987).

The theory of personal constructs was proposed in the context of psychotherapy. In terms of clinical psychology a patients personal-social behavior is influenced by their internal representation of their feelings towards other individuals who play an important role in their life. These feelings are developed based on past interactions, experiences and perceptions. These internal representations were elicited by Kelly using repertory grid techniques, with treatment being based on the results. (Garg-Janardan and Savendy, 1990).

Perceptions are represented by what Kelly called constructs. Constructs are bi-polar concepts which can be used to discriminate between events. That is, similarity or lack of similarity can be represented. These events are called elements and can be objects, situations or even individuals. With a group of elements, inter-element similarities and differences are perceived. Based on past experiences new or current elements are rated according to the constructs. Constructs themselves are interrelated and may be represented by hierarchies or networks.

This theory has been validated by several researchers, including Mair (1966), and it has been concluded that:

- Individuals do represent their environment using constructs,
- Constructs are organized in interrelated structures which change from time to time,
- The repertory grid technique elicits these constructs accurately and reflects the changes in an individuals construct system over time and
- The grid technique elicits the true structure and organization of the individual's construct system.

For use in expert system construction, experts are required to identify discrete region classifications (*elements*) which become the column headings of the grid. Groups of three of these are then taken and the expert is asked to identify what differentiates one from the other two elements. These differentiates are known as *constructs*. All elements are then rated against the construct as either totally belonging or not belonging to groups on a scale of 1 to 5. This construct then becomes the label for a row of the grid. Cluster analysis (such as Johnson hierarchical clustering, Olson and Rueter, 1987) is used on the two dimensional grid which results. Patterns and associations of the elements and constructs are identified and rules are generated.

Rules can be generated by concentrating on extremes of rating in grids. For example with a grid based on a bi-polar rating of 1 to 5, rules would concentrate on concepts which were at the extremes. These can be refined by finding the concepts which are best at differentiating between elements. To this end the between-concepts and between-elements matrices are of use. Elements which are very dissimilar are easy to distinguish between. Elements which are very similar on the other hand are more difficult. Hence although a computerized system can automatically generate initial rules, there is still the need for a human expert to refine these rules.

The analyzed grids can be inspected and new concepts generated for grouped concepts which are similar. For example with two (or more) concepts at a very high level of similarity as distinguished by cluster analysis, the expert could be asked to name a new concept which incorporates those being grouped. Once this has been done more rules can be generated to better reflect the experts reasoning.

It should be realized that the grids generated are dynamic and it should be possible for a domain expert to add both new concepts and elements. It should also be possible to combine the knowledge of several domain experts held in several grids into a single grid (Boose and Bradshaw, 1987). This can be done by identifying similarities and differences between experts grids. One would expect some of the elements identified by multiple experts to be in common as would some of the concepts (however labels could be different). On the other hand, the more experts involved, the more chance there will be dissenting opinions (Gaines and Shaw, 1993).

Within a GIS this technique has significant promise as an alternative or complement to classifying with graphical tools and as a means of identifying deeper knowledge. Different spatial areas can become the basis of the elements and serve as the prompts to get the domain expert to distinguish between them and generate constructs. Spatial relationships may also be elicited with size, orientation and proximity being concepts.

Knowledge and rules about the expert's methods can therefore be generated, even though the expert may have trouble verbalizing knowledge in a traditional interview session.

One problem with using a single grid to represent a problem is that it can become both unwieldy and tedious for an expert to use when eliciting constructs and also when rating those constructs against elements. Therefore a hierarchy of grids can be developed. This is particularly useful in classification of complex images where one region may be subdivided into subregions.

The advantage of a repertory grid is that it can be regarded as a computerised structured interview manager. When coupled with a graphical user interface where images can be directly used for prompting and manipulation, it can be a useful addition to a tool-kit and should be incorporated. A full discussion of the implementation of this tool will appear in a later chapter.

4.3 AUTOMATED TECHNIQUES

Machine learning has been seen as the solution to the problem of the knowledge acquisition bottleneck (Cullen and Bryman 1988). Two general approaches are used in machine learning. The first is the inductive approach where rules are extracted from representative example data and secondly there is the deductive approach where knowledge is deduced from theories or rules. Some systems combine the two approaches. Artificial neural networks are automated knowledge acquisition techniques which are based on processes thought to underlie the operation of neurons in a biological brain.

4.3.1 Machine Learning

Rule induction is a machine learning technique where a program is supplied with a table of examples which it then uses to create a decision tree. That is a process of general inference from particular instances (Jeng et al, 1996). The data used to create the table is known as training data. The decision tree is then applied to other data and the results analysed for accuracy.

Kodratoff et al (1994) identifies three goals for machine learning applications which are intended to:

- Detect similarities in a data set including clustering and pattern recognition,
- Acquire knowledge to revise and complete knowledge-based system or create a model and
- Classify data into appropriate categories.

All of these have applications to the visual geographic knowledge domain, especially in identifying scene primitives. The first goal includes pattern detection where pixels are clustered. The last goal, classification, on the other hand starts with a given set of classes and assigns pixels to the most appropriate class.

Quinlan's ID3 (Quinlan 1986) algorithm is the most often quoted example of these techniques where a set of example cases are taken and converted into a decision tree to differentiate between cases. The objective is to determine a decision procedure which will allow entities to be assigned to categories or classes on the basis of attribute classes. This technique has been very successful where there is a large training set and significant criteria have been hard to find. The main problem occurs when some of the important examples are not included in the training set. Another major problem of particular significance to a remote sensing domain is noise (errors) in the training set (Vrtacnik and Dolnicar, 1995).

Rule induction has been used in situations where there has been a spatial context. Kumar et al (1994) used Quinlans C4.5, the descendant of ID3, to induce rules about weather forecasting. The results were competitive and matched human forecasters in performance. The results were a forecast of occurrence and depth of rainfall rather than a map, so the applicability to the current study is limited. Another limiting factor is the amount of ground truthed training data available. In Kumar's study a 10 year data set of 2,663 items were available, as was a 30 year data set. In many case data sets over this sort of time period are difficult (and expensive) to acquire.

A second use of rule induction has been to take a training set of data to improve an expert system which had already been built by domain experts to classify salinity (Eklund and Salim, 1993). As well as reorganising and improving the performance and efficiency of the expert system the technique was able to discover new domain knowledge in the form of classifiers. This system also used C4.5 and achieved an order of magnitude improvement over the user developed system.

Another approach to machine learning has been rule deduction which is based on applying general rules to specific situations. For example the general premise:

Anyone who can program is intelligent

And the specific premise:

Paul can program

Can be combined using deduction to conclude that:

Paul is intelligent (from Giarratano and Riley (1994), p120)

This can be very useful in a when a large body of knowledge is available and also forms the basis of explanation based learning (Arciszewski and Ziarko 1992). Systems of this type developed in the spatial domain include Palermo, (Matwin et al, 1995) a system which uses transformational analogy, derivational analogy and goal regression to develop forestry management plans. Transformational analogy, also known as case based reasoning adapts solutions from similar problems. Derivational analogy uses cases to choose among competing directions along search paths. Goal regression is a search based problem solving approach.

4.3.2 Neural Networks

Neural networks are based on the supposed neuron structure of the brain (Openshaw, 1993). The brain consists of an extremely large network of highly interconnected neurons with each neuron being a simple summative device. The main feature of artificial neural networks are their ability to learn from training examples enabling them to be used as a form of automated knowledge acquisition technique.

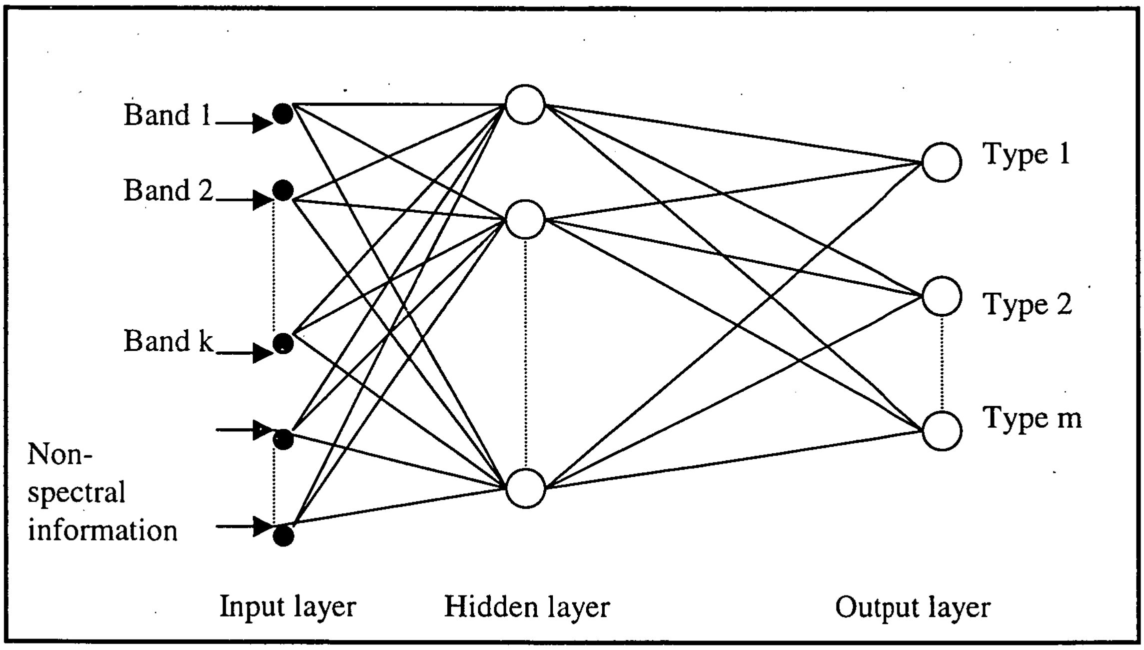

Figure 4.2 Multilayer feedforward artificial neural network for land type classification (Leung, 1997, p205)

There are several different implementations of neural networks which fall into two main categories, feedforward nets and recurrent nets (Leung, 1997). A simple feedforward net is shown if Figure 4.2. The input layer consists of values presented to the network. In Figure 4.2 this is a pixel by pixel input from the available image bands and other layers from a GIS. The output layer contains the classification types for the image. Hidden layers encode intermediate objects or variables.

Recurrent schemes have feedback as well as a feedforward connections. In other words there are connections back from intermediate nodes. Two variations are bidirectional associative memories (Leung, 1997) and Hopfield nets (Raghu and Yegnanarayana, 1997).

In theory a neural network can be trained to represent and model almost any complex system, no matter how difficult the task may appear to more conventional approaches. They can be taught to recognise patterns and analyse data by letting the data define the structure within it (Openshaw et al, 1991).

Neural network techniques have been used for classification of remotely sensed images with varying levels of success. However Skidmore (1995) states the often quoted advantages of neural networks "such as parrallelism, speed and trainability, were more than negated by the variable and unpredictable results we generated".

Better results were obtained by Carpenter et al (1997) who used a neural network based on adaptive resonance theory. A hybrid system incorporating maximum likelihood classification gave results of 60% accuracy. This does not however compare well with expert system approaches, for example Skidmores (1989) Eucalypt Forest classifier which achieved a 75% accuracy.

Despite this apparent poor performance, neural network technology does hold promise (Civco, 1993). A possible use of this kind of processing is to use the results as a first guess from which an expert system can be developed (Khosla and Dillon, 1993). It would then be up to the domain expert to make adjustments where they saw fit.

4.3.3 Supervised versus Unsupervised Methods

The first step in supervised training is to generate examples from a manually outlined group of pixels on a multispectral image. These groups are representative of the various land cover types present in the image. This is the training stage. Next, in the classification stage, these serve as input to the machine learning algorithms or are used in other classification methods (Sabins, 1997). There are many classification techniques which can be used and they include:

- The Minimum Distance to Means Classifier,
- The Parallelepiped Classifier and
- The Gaussian Maximum Likelihood Classifier.

(Lillesand and Kiefer, 1994)

Unsupervised classifiers do not use training data as the basis for classification. Unknown pixels in an image are aggregated into classes based on natural clustering. They are classified on purely spectral information. The image analyst is then required to compare the results with some reference data to assign labels to the classes. There are numerous clustering algorithms which can be used for unsupervised classification (Lillesand and Kiefer, 1994).

4.3.4 The Importance of Training Data

One of the strengths of a remotely sensed data set is that it represents a complete spatial population. However the data set used for training a machine learning system is usually a spatial sample. That sample must be chosen with care (Curran and Williamson, 1986). Machine learning techniques are dependent for their accuracy on the quality of the training data rather than on the algorithm used for classification. The training data set needs to be representative of the whole area to be classified. The populations of pixels used for training must be statistically significant and follow a Gaussian distribution. (Buttner et al, 1989). This means that there is a need to know the minimum number of observations required to characterise a particular site to an acceptable level of error.

4.4 COMBINATION OF KNOWLEDGE GAINED FROM MULTIPLE TECHNIQUES AND MULTIPLE EXPERTS

One problem with using a variety of techniques for knowledge acquisition (or using multiple domain experts) is that a series of knowledge-bases will be developed. At some stage inconsistencies between these knowledge-bases needs to be resolved. In the SPARTEX system multiple knowledge-base structures are used to refine classifications (Crowther et al, 1994), but the knowledge in those knowledge-bases is consistent. It is very probable that the knowledge generated by a tool-kit using different tools at different times with different experts interpreting different images will be in conflict.

Although many early practitioners recommend a single domain expert for knowledge acquisition (Boy, 1996), using multiple experts has many advantages (Crowther 1992) including:

- No one expert knows the total domain,
- Experts have complimentary knowledge,
- Experts can prompt each other and
- Experts can validate each others knowledge.

Chao and Salvendy (1994) did a statistical study and concluded that an optimum number of experts from whom up to 90% of knowledge could be extracted was 3 to 4 depending on which techniques were being used. Where experts have complimentary rather than overlapping knowledge, combination of knowledge-bases is relatively straight forward. Overlapping knowledge however requires aggregation and consolidation (Mak et al 1996)

Problems can arise where knowledge from multiple experts is to be combined and there are dissenting opinions. It should be noted that the dissenting opinions are not necessarily wrong but may be an alternate way of regarding the domain under investigation. Techniques including repertory grid analysis can take this into account.

It should also be noted that the more experts there are involved, the more chance there is of dissenting opinions.

Repertory grids have some well defined methods to combine grids developed from several experts. The result is a single knowledge-base that is a combination of knowledge.

Other traditional approaches to knowledge consolidation include Dempster-Shafer theory, certainty factors and fuzzy logic. In Dempster-Shafer theory evidence is reinforced for a particular solution. Certainty factors require a domain expert to assign a level of certainty to outcomes, or this can be done computationally. Fuzzy logic allows the development of fuzzy sets to represent boundaries where there are no crisp delineations.

The Group Elicitation Method or GEM (Boy, 1996) is a six phase process for use with multiple experts consisting of the following steps:

- Develop a statement of the problem and choose the group of problem solving experts,
- Generate of viewpoints from the participants,
- Reformulation the viewpoints into more elaborate concepts,
- Generate relationships between concepts,
- Derive a consensus and
- Critically analyse the results

The Delphi Technique mentioned by DeMers(1986) is used for eliciting and processing the opinions of a group of experts. Attributes identified by experts are compared pairwise and ranked. These rankings are then weighted and a group ranking calculated. The results are then presented to the experts for reevaluation. The process is continued until all domain experts agree on the final attributes.

Both of these techniques require the multiple domain experts to be present at the same time. They require verbal interaction between participants. This is one of the problems discussed by Medsker et al (1995) who found that knowledge acquisition and consolidation was more complex with multiple experts who were dispersed (and could not meet face to face).

Mak et al (1996) suggests using more machine centred methods, ranging from machine learning techniques and neural networks discussed above through to *discriminant analysis*. Discriminant analysis is used in two group situations. A linear combination of the set of group characteristics is developed which will provide the maximum differentiation among the groups.

4.5 KNOWLEDGE ACQUISITION AND GIS

GIS have several unique problems when compared with other systems. Many of these were discussed by Srinivasan and Richards (1990a). The main problems identified by them relate to multi-source data. First there is the data picked up from sensors with differing characteristics. For example Landsat and SPOT satellite data have different resolutions and use different wavelength bands. The number and variety of sensors is likely to increase further in the future, adding to the complexity. The two satellite systems rarely produce images on the same day so even when they are geo referenced there may be anomalies due to changes on the ground (for example harvesting). Secondly there is the desire to add what the authors describe as 'ancillary' sources which include topographic information and existing maps, but could also include domain expert generated heuristics.

Another consideration with knowledge acquisition is that no one approach will work all the time. Instead, a tool-kit needs to be assembled which will allow a knowledge engineer to acquire knowledge-based on the situation, the expert, available time and possibly the domain (Trimble and Cooper 1987). Further it may be necessary to plan the knowledge acquisition based on the above factors. In traditional knowledge acquisition this would start with interviews (which almost always will be necessary, even in a spatial domain) and then move on to literature searches, repertory grids or more interviews. The initial set of interviews, which would normally be carried out during a feasibility study, would then point the way to other more appropriate knowledge acquisition strategies.

In most geographic applications an expert will start analysis by viewing some form of two dimensional image representing a ground area. This image may be an actual image created from one or more sensor bands on a satellite, or a histogram representation of the actual image. Various areas will then be defined as belonging to various classifications. Expertise can be captured directly from closed curves drawn on the image or picking spectral threshold values from the histogram. However this technique may only capture a proportion of the experts knowledge.

Other knowledge acquired by the expert through experience may be non-visual. Such unstructured heuristics are important and need to be captured. There may also be calibration algorithms and other theoretical knowledge that needs to be added to the knowledge-base.

It may be possible to develop some form of focused session based on what was produced by the closed curve exercise. This would involve asking the expert why a particular region was classified as such, perhaps presenting readings from other sensors and asking for range limits. It may also produce the basis for some form of grid (Boose and Bradshaw 1987). It should however be up to the expert (or knowledge engineer) to decide which method of acquisition should be used and what order the other tools should be used in. As part of the process there should also be some consistency check so conflicting knowledge can be detected and resolved as early as possible.

The approach taken in the SPARTEX system (Williams et al 1994) has been to develop a series of rule groups which refine initial image data to a final solution. Based on this approach it may be possible to use different tools to develop different knowledge-bases which are then linked. It should also be possible to capture knowledge in a form which the system can use almost directly (for example CLIPS rules). It will however require a specialised visual knowledge acquisition tool.

4.6 Conclusion

Traditional knowledge acquisition tools have been used successfully to develop expert systems for use with GIS and RSS data. Of these, the fully automated techniques of rule induction and neural networks have been used most consistently. The manual techniques, particularly interviewing, protocol analysis and closed curves, have been used to develop hand crafted expert systems which have been quite successful.

Repertory grids seem to be a particularly useful tool to be included in a geographic knowledge acquisition tool-kit (Waters, 1989). This technique provides a non-visual means of requiring a domain expert to distinguish between geographic features. Such a system can be used in a structured interview by asking a domain expert to identify features, then to identify the means of distinguishing between them.

User interaction is also necessary for automated techniques. Although machine learning techniques have been integrated with non-automated techniques (Nedellec, 1995) it is not the aim of this thesis to investigate this aspect. The emphasis here is on interaction with the domain expert. In spite of this there needs to be a mechanism to couple with machine learning systems. To achieve this, a mechanism to allow an expert user to develop and acquire training data should be included, particularly for supervised systems where training pixels need to be isolated and labeled. Hence a tool which can either randomly or systematically select and label training data should be included.

Chapter 5. THE SPATIAL KNOWLEDGE ACQUISITION PROCESS

This chapter investigates the process necessary to acquire spatial knowledge and the techniques needed. It is shown that different tools and techniques are required to elicit knowledge in the six classifications described in section 2.7. These tools are integrated into a knowledge acquisition tool-kit, KAGES (Knowledge Acquisition for Geographic Expert Systems), which will be described in Chapter 6. The scenarios used to evaluate the knowledge classification scheme and the KAGES tools based on it are introduced.

5.1 INTRODUCTION

Much of spatial or geographic knowledge is visual. Hence it is natural for a user to show or draw things on a map or image. A spatial knowledge acquisition tool therefore needs to be able to directly capture the expert image interpreter's actions from their manipulation of the image.

Not all spatial knowledge is visual however. An image interpretation expert may use heuristics to identify a feature or confirm its existence. Likewise knowledge of algorithms is not visual. These types of knowledge have traditionally been acquired by interviews. Interviews however have severe limitations (Section 4.7.1) and a more systematic method of knowledge acquisition is required.

5.2 SCENARIOS FOR ASSESSING SPATIAL KNOWLEDGE ACQUISITION

Two scenarios were chosen for detailed testing of the knowledge acquisition methodology, the type of knowledge it generated and user reaction to an interactive knowledge acquisition tool.

The two scenarios involved the development of :-

- An Antarctic sea ice classification system and
- A system for differentiating agricultural crops and land use in Tasmania.

Results obtained in these two scenarios were used to demonstrate the capabilities of the methodology and the potential for using the knowledge acquisition tool-kit in other remote sensing domains.

5.2.1 Antarctic Sea Ice Classification

The knowledge acquired using the knowledge acquisition methodology in this domain was compared with that acquired for a system called Icemapper which was developed using traditional knowledge acquisition techniques (Williams et al, 1997).

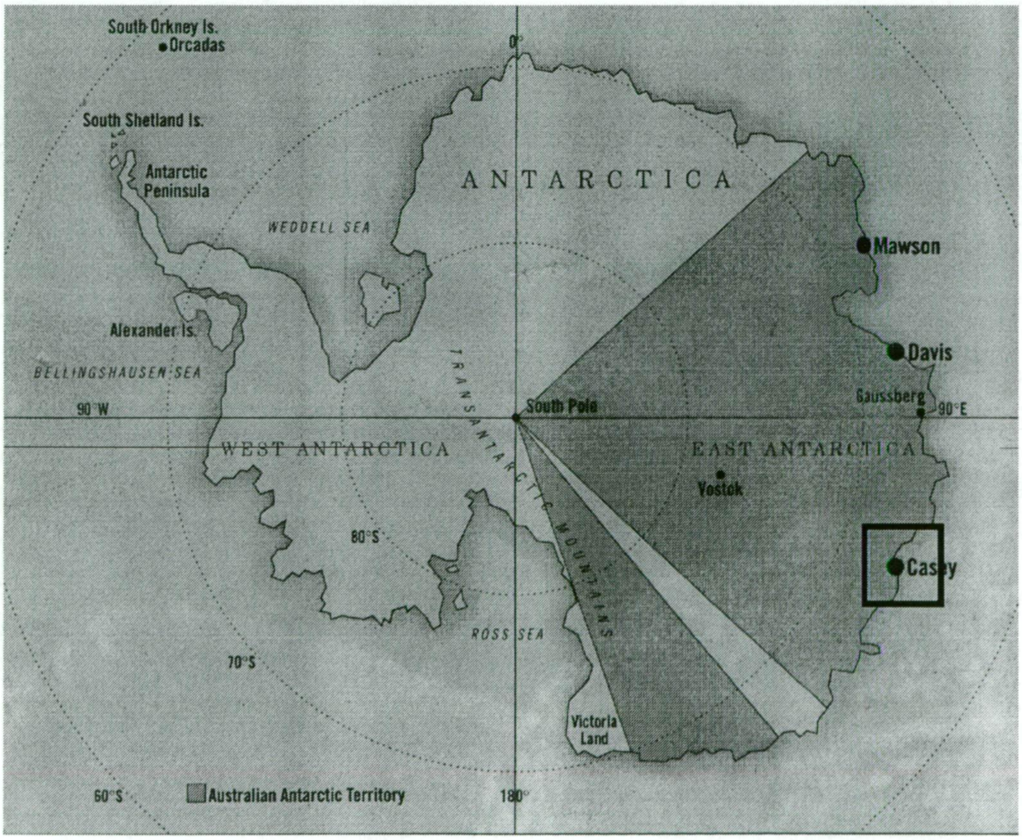

Figure 5.1 Context of Antarctic study area (from 'The Frozen Continent', *The Sunday Age*, 6th April 1998)

Icemapper is an operational expert system used for identifying sea ice features in Antarctica. The system has a number of objectives. Of immediate concern is the location of open water leads and thin ice which can be used by the Australian National Antarctic Research Expedition (ANARE) supply vessel. Other uses for Icemapper include global warming studies based on the seasonal expansion and contraction of sea ice.

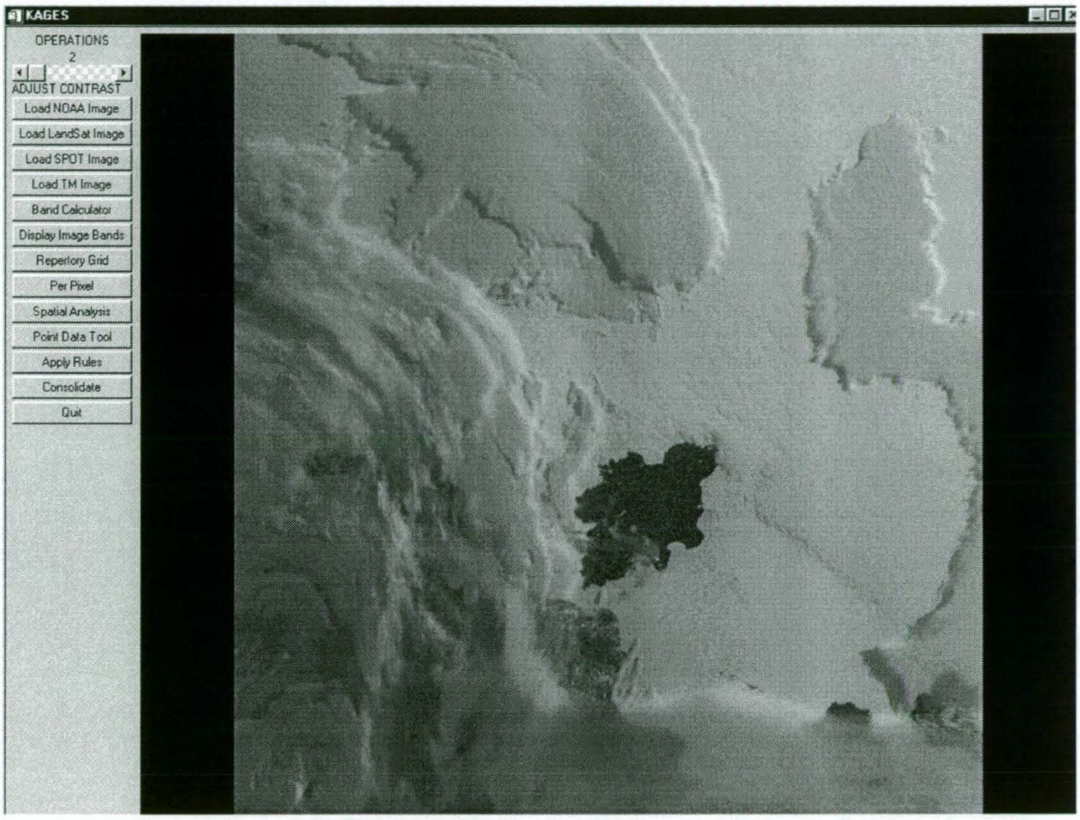

Figure 5.2 Antarctic study area: An area in Vincennes Bay near Casey Station. Note South is up the page. Taken from a NOAA AVHRR band 1 image viewed using KAGES

It is a stand alone system developed in IDL (Interactive Data Language) and is not tied to a GIS. However it serves as a useful standard because the rules encoded in the system were acquired by interview and protocol analysis and by searching the literature for appropriate algorithms. In terms of the proposed knowledge classification scheme, the knowledge in Icemapper is of the Primitive and Non-Visual Algorithmic type.

The data used in this field scenario is from an area around Australia's Casey Antarctic Station in Vincennes Bay (Figure 5.1). The area is depicted using NOAA AVHRR images (Figure 5.2). An image interpretation expert (the same one used for the original Icemapper rules) was available to both use and comment on the KAGES approach, as well as to advise on the rules which were produced.

5.2.2 Crop Recognition

Knowledge for crop recognition is to be used in a local crop mapping project called the Multi-Temporal Imaging for Remote Sensing of Crops Project, (MIRC). The MIRC project is primarily GIS based. The study is based in an area of Northwest Tasmania around Table Cape (Figure 5.3). In 1996 a series of Landsat TM and SPOT satellite images (Figure 5.4) were obtained of the region which consists of sea, forest, agricultural land and urban development (the township of Wynyard).

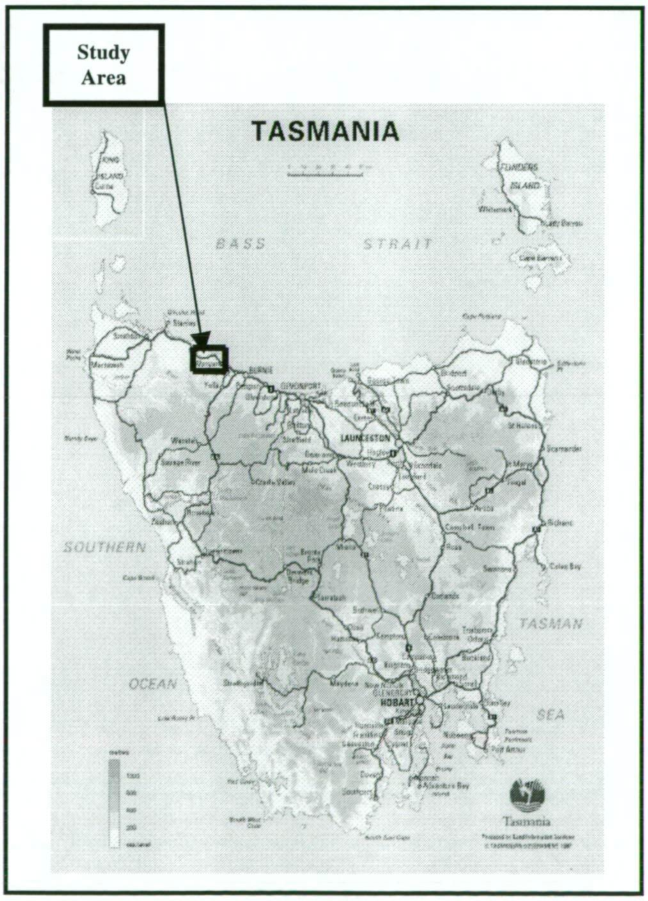

Figure 5.3 Context of study area for MIRC

The objective of the MIRC project was to create a system which could differentiate between the various crops being grown in the region. This would be used to give an indication of the relative areas under different crop types. Crops grown in the region include poppies, peas, potatoes, carrots, onions, brassicas, pumpkins, cereals, beans, pyrethrum, tricalate and corn, much it for the processing and export industry. A future development of MIRC will be to extend it to include the recording of crop health.

This project was significantly different in method from the Antarctic project because there were no expert satellite image interpreters available. As a result KAGES was initially used as an image analysis program. The rules for crop discrimination which were generated were then checked against ground truth and against an independent classification study using a statistical classifier (details in chapter 7).

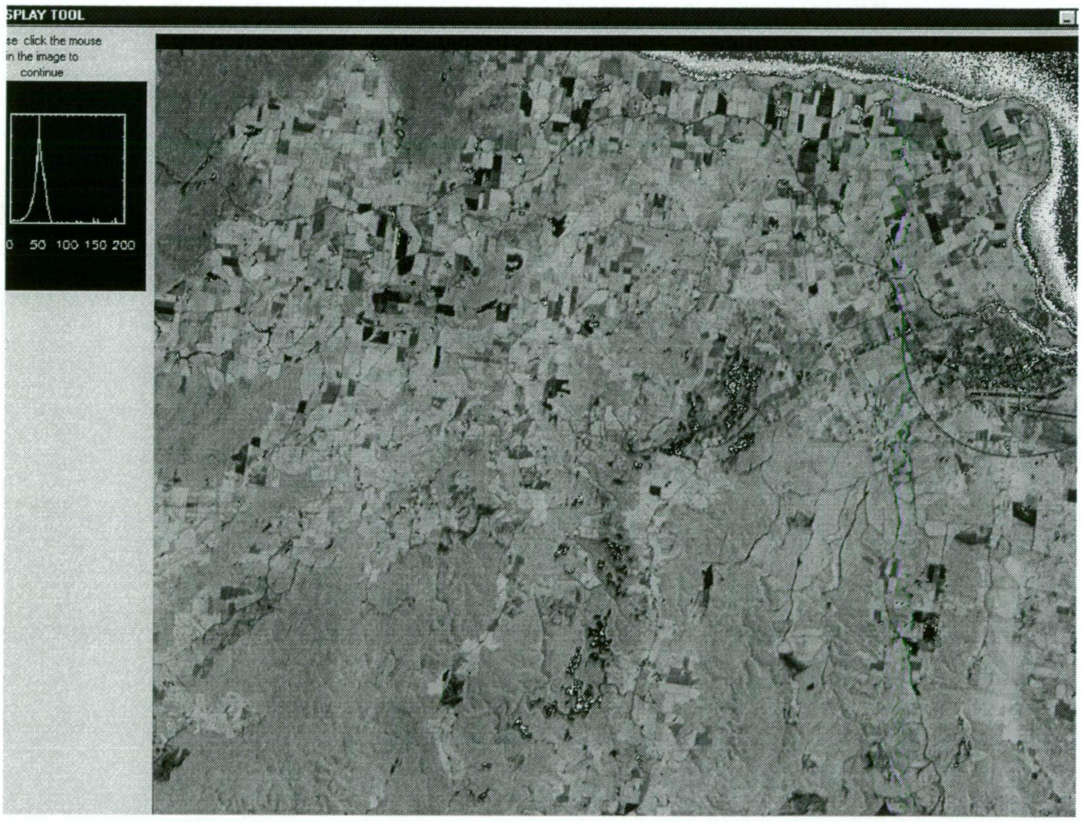

Figure 5.4 The MIRC study area as seen on band 1 of a SPOT image taken on 2nd November 1994 using the KAGES tool

5.3 PRIMITIVE KNOWLEDGE ACQUISITION

The basic components of any spatial system are its scene primitives; basic units which cannot be further subdivided into smaller units and which together constitute a scene. These are the basic building blocks that are used for reasoning about relationships in scenes. Scene primitives are point, line or areal objects which cannot be further subdivided into other objects. They are initially identified by their spectral signature. While this identification method using spectral signatures is quite well developed it is not totally reliable (Wilkinson,1996).

5.3.1 Knowledge of Areal Features

Areal feature primitives are two dimensional objects identified within a band or band combination of a remotely sensed image. They have characteristics which make them homogenous in terms of their pixel value. The usual way for a classifying expert to identify them is to point at them then study their spectral characteristics. This is generally done with reference to a histogram of pixel values for the image as a whole and of the object in particular.

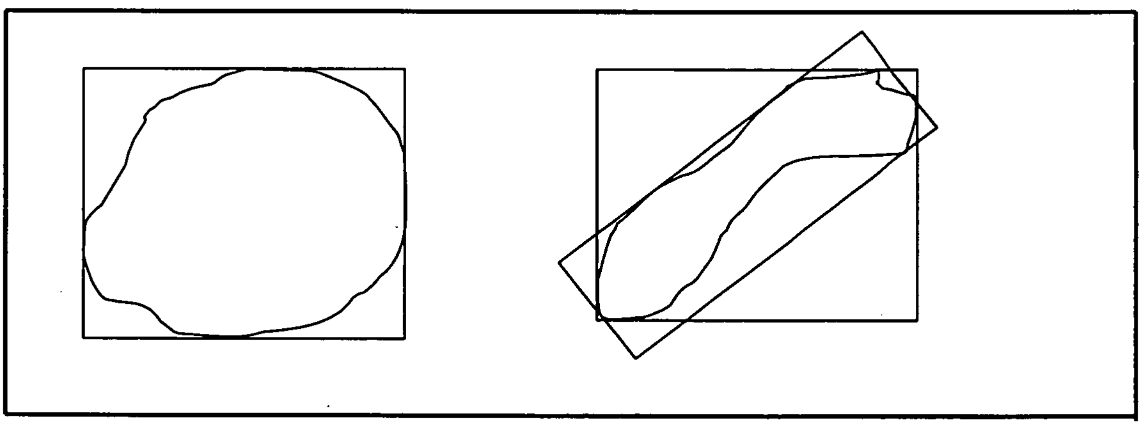

Figure 5.5 Minimum Bounding Rectangles. The example on the right has a true MBR superimposed.

Because of the often irregular nature of the boundaries of areal objects, there is a need to provide some generalised description of their extent. At the simplest level this is done with a *minimum bounded rectangle* (MBR) (Chang and Jungert, 1996). An MBR

(Figure 5.5) has its boundaries parallel to the vertical and horizontal axes of the image, thus ignoring the orientation of the object. Hence a long thin object at 45 degrees to the axes has a very large MBR compared to its actual area. Despite this they can be used to produce topological relationships for example the 2D Projection Interval Relationship (2D-PIR) described by Nabil et al (1995).

2D-PIR would represent each relationship between spatial objects A and B as an ordered pair (χ, ψ), where χ is the relationship between the objects along the X axis and ψ is the relationship along the Y axis. To write a string each of the various constructs of Allen Intervals (discussed fully in section 5.3) is given a symbol to describe the χ and ψ relationships. Hence in Figure 5.6 a relationship between A and B would be represented as (<, **O**) with < meaning A is before B on the X axis, and **O** meaning A overlaps B on the Y axis.

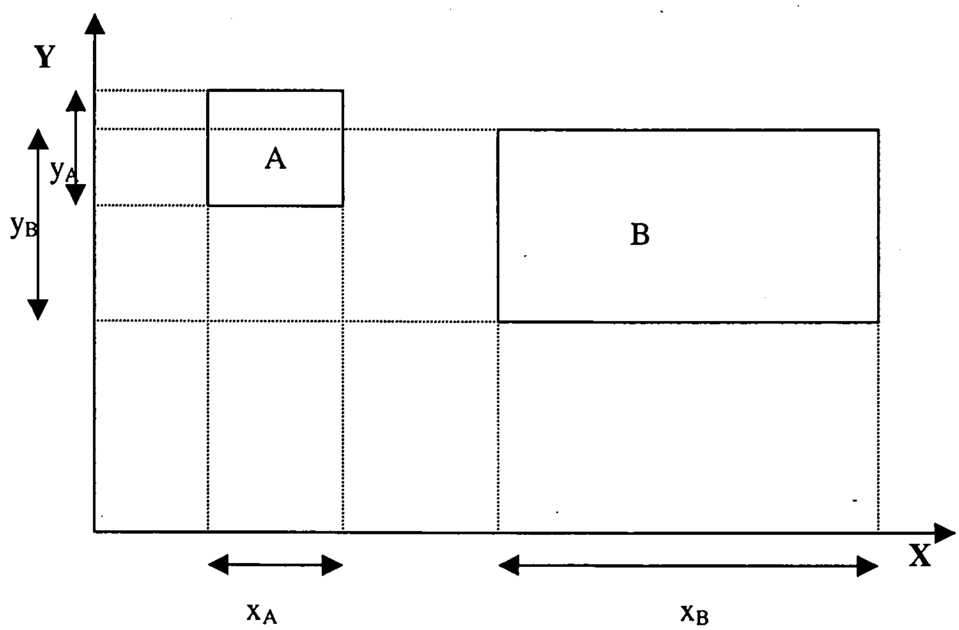

Figure 5.6 A 2D-projection picture (after Nabil et al, 1995)

A further refinement is to use *true MBRs* whose boundaries are governed by an orientation angle (Nabil et al, 1995). To provide an even more accurate description a vector polygonal structure can be used (Jungert, 1993).

In this study simple MBRs will be used for the following reasons

- Simplification of calculation,
- Simplification of display and
- Orientation is not definitive of all objects and can be calculated when required.

5.3.2 Knowledge of Point Features

Point features are generally subpixel and are defined in terms of an absolute location. They may also be defined in terms of interaction between two other features (Palmer, 1984). For example a level crossing (point) can be defined where a road (line) crosses a railway (line) or an entrance gate (point) where a road (line) enters a national park (area). Much work has been done in producing semantic descriptions of this type of point feature (Gould et al, 1996). A visual rather than semantic representation may be more universally applicable. That is the appearance of a spatial relationship may identify an object rather than its name.

From the above discussion it can be seen that point features can be defined in two ways. Firstly an expert classifier can point to where the point is, or in a geo-referenced system where a pixel's actual location is known, give its actual location. One problem with this second method is it is subject to error of up to 100 meters (Harris, 1997), even if a state-of-the-art *Global Positioning System* (GPS) is used. This error can be reduced to a meter if the differences from a known ground station are used.

5.3.3 Knowledge of Line Features

Lines are one dimensional objects when considered in terms of image processing. Generally they are an object which is very narrow, often only a pixel wide (Laurini and Thompson, 1992). Occasionally, as in the case of a boundary they may be zero pixels wide with definitions being in terms of the two objects which meet. A second more important attribute is that they are many orders of magnitude longer then they are wide. It should be noted that there are occasions were the difference between line and areal objects is difficult to state. For example where a river widens near its mouth, or where a road becomes significantly wide when a zoom function is used. For the purposes of this thesis the original definition of a feature a few pixels wide or less will be used.

Line objects can be of several types. They can be a simple line which is at least a pixel wide with pixel characteristics such as reflection. For this type of line there are a number of line following classifiers available (e.g. Gruen and Li, 1995). This type of classifier uses actual pixel values to follow a line from some user defined staring point. The result is a set of pixel locations (which can be quite large).

A second type of line object is one which is a boundary and hence is subpixel. Again there are a number of algorithms which identify these type of lines by edge enhancement (Drewniok, 1994). These edges can be stored as a series of raster points in a data set.

The third type of line is human defined and generally is related to some socio-political boundary. Examples include international and state political boundaries and cadastral data including such things as the location of local government services and private land boundaries.

For non-straight lines which are identified using a line following algorithm a pixel by pixel trace is generated producing a raster data set of points identifying the object. Storage can be saved in the case of fixed lines which are wide enough to have a spectral signature by storing thresholds, an end point and a location so the line can be regenerated. Temporally varying lines such as cloud edge shadow also need information identifying which image they came from. For these objects, data needs to be stored so the whole line can be regenerated.

Straight lines are easier to deal with and are better suited to vector representation with end and intersection points being stored. A very sinuous line or one which requires the actual feature to be represented are better defined as a raster representation. This of course can cause problems when a vector based line and a raster based line are compared. A solution is to convert the vector representation of the feature to a raster representation. The system which will be described in chapter 6 uses both raster and vector representations and does conversions where necessary.

5.4 RELATIONSHIP KNOWLEDGE ACQUISITION

Relationship knowledge acquisition is a process which allows users to select pairs of scene primitives and generate their spatial relationships. These relationships include the degree to which one object overlaps another and its proximity and orientation with respect to the other object. By-products from the process may be more scene primitives (for example point objects where lines intersect other objects, or line objects where two areal objects have a boundary). (Palmer, 1984). The spatial relationship between two objects may also add evidence to the classification that one or both of the objects has been assigned.

A common method of comparing the spatial relationship of two objects uses Allen intervals. Allen intervals have both a spatial and a temporal dimension to them.
The temporal relationships are:

Before	Meets
Overlaps	Finished By
Contains	Equals
Started By	Starts
During	Finishes
Overlapped By	Met By
After	

For example, in agriculture, a field that has been cropped can have the temporal relationships MEETS to show its transition to fallow. Fallow is also AFTER cropping. Hence a rule developed from successive images during the cropping period would determine that fallow would be a valid classification if the previous classification for the same paddock was a specific kind of harvestable crop.

Although temporal relationships will not be used directly in the spatial tool, but these types of intervals will be used in an associated database which will later be used to extract historical data (Appendix B). This data will then be used to develop a temporal knowledge-base.

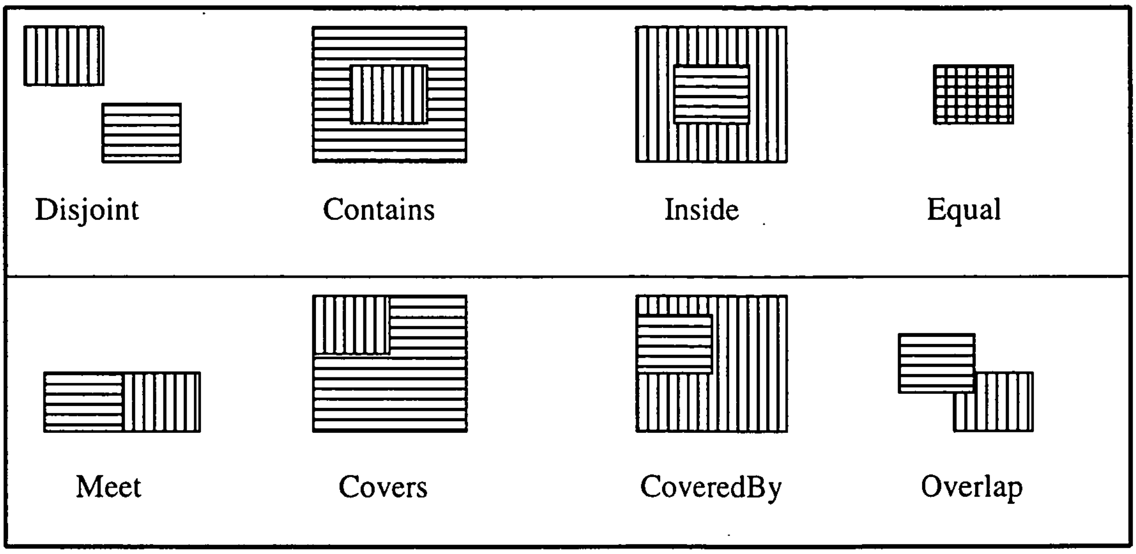

Figure 5.7 The eight relations between two regions (Egenhofer and Sharma, 1993)

The spatial relationships of Allen intervals are shown in Figure 5.7. For example in an agricultural system most paddocks MEET another paddock, but rarely OVERLAP. However a paddock may be subdivided so the super-paddock may CONTAIN two or more sub-paddocks. Lastly a paddock may be COVERED BY cloud.

There are numerous spatial logic methods based on these relations, for example Cui et al (1993), Clementini et al (1993), Chang and Jungert (1996), Smith and Park (1992) and Grigni et al (1995). All of these present a script for describing spatial relationships. The one problem with all of these is they provide a script which is ideal for computerised reasoning but has drawbacks for knowledge acquisition where a domain expert needs to verify, preferably a plain English statement, what the system is interpreting as their actions.

Orientation of one object in relation to another can be considered in several ways. One way is from a view at a point on an image. For example one object can be behind another object or left of an object (Hernandez, 1993). Another method is to use compass directions when describing the orientation of one object in respect to another as suggested by Antony and Emmerman (1986), Freksa (1992) and Sharma and

Flewelling (1995). Various parts of objects have been used to calculate orientation. These include using object centroids, the corners of their MBRs and intersecting sectors (Sharma and Flewelling, 1995). In all the papers reviewed, orientation was specified to within 45°. The directions defined were north, south, east, west, north-east, south-east, south-west and north-west.

Most reasoning in images is in terms of areal two-dimensional objects. Inclusion of line and point objects complicates reasoning considerably. In the simplest case a line or point can be completely disjoint from an areal object or it can be contained within an areal object but not touching a boundary. In the first case reasoning about the proximity of the nearest point of a line and the orientation of the line relative to the areal object is fairly straight forward, as is reasoning about a point. Special cases include a line that is disjoint from an area but encloses it (for example an airport perimeter road).

Lines are particularly complex. Consider, for example, a road following a field boundary. In this case there are no pixels in common between the two objects (hence they are disjoint); however there are pixels adjacent to each other. Since they are not all adjacent, a system would conclude that the road partially follows the field boundary.

In the case of a line or point enclosed in an areal object (Figure 5.8 (b), (h)) it is possible to reason about the location of the line or point within the object relative to its centroid. Hence the line or point object may be in the south east quadrant (Figure 5.8 (h)).

The most complex relationships occur where a line interacts with the boundary of an areal object. A line may bisect an area (Figure 5.8 (e)), start/terminate in an area (Figure 5.8 (d)), start/terminate on a boundary (Figure 5.8 (a)) or a combination of the previous two (Figure 5.8 (f)). With these there is a simple interaction at the boundary with one or two contact points. More complex interactions occur when a line touches but does not terminate at or enter an area (Figure 5.8 (c)), a line follows a boundary for all or some of its length and a line continuously passing in and out of an area. In this regard, if the pixels on the boundary of the area are considered to be a line feature, a reasoning method for two lines could be used.

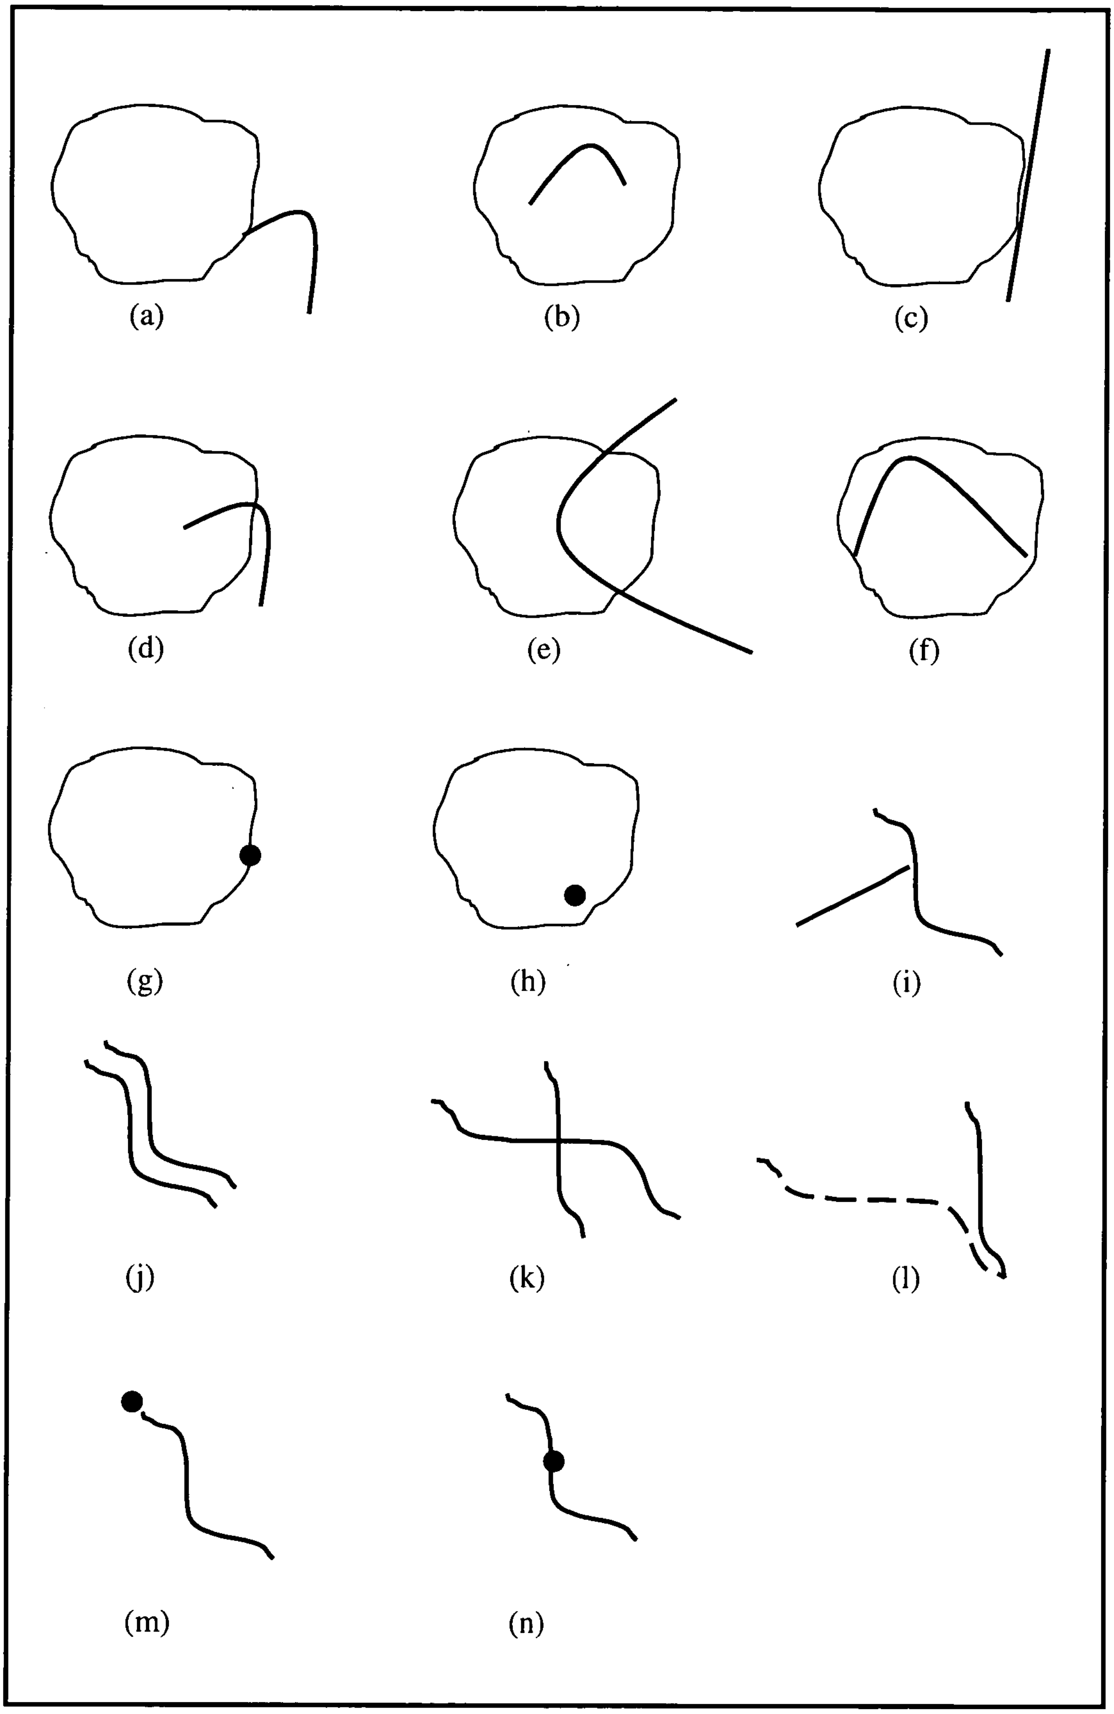

Figure 5.8: Some spatial relationships between line, point and areal objects

5.4.1 The Problem Of Three Objects

In many cases point objects are defined in terms of the interaction of two of the other classes of object (Lee and Chin, 1995). For example if two lines intersect and the two lines are roads the intersection is a cross roads, or if one line is a road and the other is a river, the intersection is a bridge.

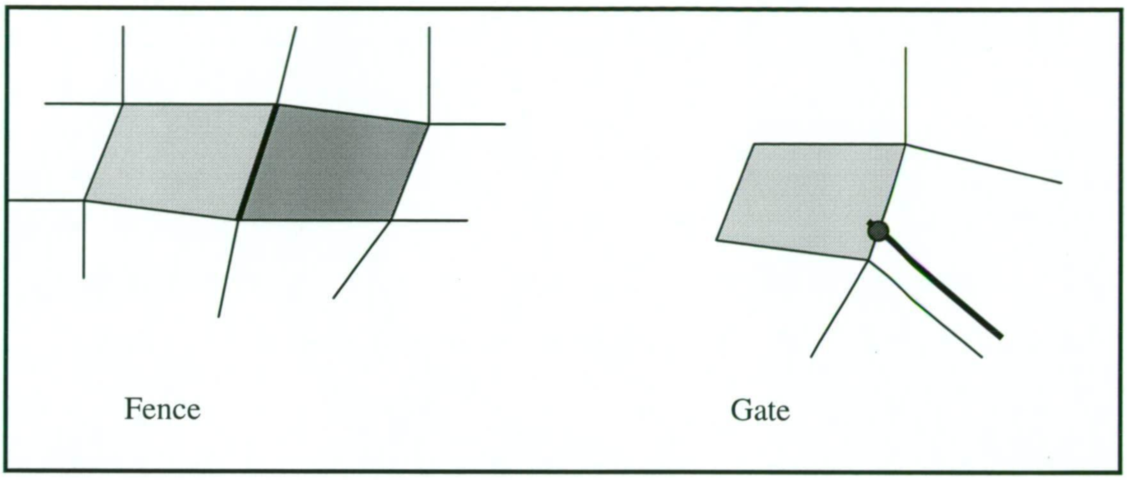

Figure 5.9 Three Object Definition Fence (line object) defined by interaction of two fields (area objects) and a Gate (point object) defined by the interaction of a field (area object) and access track (line object)

The same is also true for certain line objects. For example if areaA is a field and areaB is a field and areaA touches areaB, then the region of touch is a line object fence (Figure 5.9).

These type of rules are visual (in that a domain expert can see them and point them out). However they are also of a common sense heuristic nature and are common to most classification tasks. It is therefore suggested that they be built into common knowledge-bases available to all systems.

5.5 ASSEMBLY KNOWLEDGE ACQUISITION

Assembly knowledge is knowledge which allows an image classifier to group scene primitives into a larger object. This process is known as generalisation (Beard, 1991). Most generalisation models are based on scaling and feature removal in an automated way (McMaster, 1991). Generalisation includes the processes of *simplification, classification, symbolization* and *induction.* Simplification involves determining the important characteristics of the data (and possible exaggeration of these important characteristics). Classification is the ordering and grouping of data. Symbolization is the process of graphically encoding the grouped characteristics into a single grouped object. Induction is replacement of the original features by the new generalised object in the image.

5.6 NON-VISUAL KNOWLEDGE ACQUISITION

5.6.1 Heuristic Knowledge Acquisition

Heuristics are recognized as the basis for an expert's expertise. They comprise expert opinion, subjective judgments, expert forecasts, best estimates and educated guesses (Meyer and Booker, 1991). Heuristics are gathered through experience and can be very difficult to capture using traditional knowledge acquisition techniques. In allowing the expert to classify the image manually and by recording the steps in that classification process the heuristics used can be automatically captured. The tools discussed in sections 5.2 to 5.4 do this by allowing an image classifier to directly manipulate an image. However there may be other techniques a human classifier uses which are not visual. Therefore a non-visual heuristic acquisition method is needed.

Repertory grids based on Kellys personal construct theory (described in section 4.2.3) are the basis of several interview managers (Boose and Bradshaw, 1987). The technique assists in identifying both objects in a domain and distinguishing between those objects. They can be either drawn by hand or elicited using a computer program.

During expert system construction, experts are required to identify discrete classifications (*elements*) which become the column headings of the grid. In geographic systems these elements are usually scene primitives (but could be groups of primitives as defined by the Assembly Knowledge). Groups of three of these are then taken and the expert is asked to identify what characteristics differentiate one element from the other two. This is called the *triad* method of comparison. These characteristics are known as *concepts* and become row labels. All elements are then rated against the construct as either totally belonging or not belonging to groups on a scale of 1 to 5. Cluster analysis, such as Johnson Hierarchical Clustering (Olsen and Rueter, 1987), is used on the two-dimensional grid which results. Patterns and associations of the elements and constructs are identified and rules are generated.

Knowledge elicited by repertory grid analysis may be of the primitive, relationship and assembly types as well as the non-visual algorithmic and heuristic types. It is possible to develop grids which are used specifically for one of those levels. Repertory Grids can also be used to elicit hierarchical information (Boose, 1990). Items identified as Primitive Knowledge are usually part of a larger group of items and some domain experts prefer to begin by identifying broader groups of items which make up a scene (Assembly Knowledge). The component elements then need to be identified. In other cases discrete items are identified (Primitive Knowledge) and then grouped.

It is unnecessary to force a domain expert into exclusively identifying scene primitives or Assembly Knowledge groupings. Typically an initial consultation will result in a mixture of spatial knowledge types being identified. Further consultation will result in knowledge being refined.

Repertory grids are a useful tool for eliciting spatial knowledge from domain experts who interpret satellite imagery. They help in the construction of rules which both classify and discriminate between geographic features. They can also be used to combine the knowledge of several domain experts from different disciplines constructing rules which can be used to produce composite maps.

5.6.2 Algorithmic Knowledge Acquisition

To interpret multi-band satellite images, interpreters often combine bands in various ways to highlight particular features. Sometimes bands from several image sets are also combined. Common algorithms used are those for calculating Normalized Differential Vegetation Index (NDVI) and Sea Surface Temperature (SST), but the particular algorithm(s) used depend on the subject of the classification.

For example Chuinsiri et al (1997) in a study of delimiting crops uses

$$((TM4 - TM\ 3)/(TM4 + TM3) + (\sqrt{}\ (TM3^2 + TM4^2) + TM2$$

which is NDVI plus a Brightness Index plus the Landsat TM green channel (band 2). TM2, TM3, and TM4 are Landsat bands 2,3 and 4 which are green, red and near-infrared respectively.

The number of available algorithms is such that hard coding each of them into an automated system restricts the range available and increases complexity. Each result has to be stored as an extra dimension in an image processing package or layer in a GIS. If several of these band combination algorithms are required the data structure involved in storing all the layers becomes quite large and unwieldy.

A solution to this problem is to develop a tool which allows an image interpreter to combine raw bands as they work and interactively define the algorithm which they want to apply. When a suitable band combination is found, the characteristics of the feature identified along with the algorithm used to combine bands can be recorded. Only the raw bands of the image need to be permanently stored at any time, along with a result layer representing classified features.

5.6.3 Temporal Knowledge Acquisition

Temporal knowledge acquisition requires a system to store information over time and then use that information in a temporal analysis. Many spatial data sets have a temporal component (Unwin, 1996). For example in agricultural systems there are variations in how much of a particular crop is planted, how well it grows and when it is harvested. In a sea ice identification system the extent of sea ice varies seasonally and there may also be longer term trends in the extent of the ice sheet.

Design of a spatio-temporal database requires the ability to include discrete simple events at a specific time and more complex cases where a series of events occur within an interval (Story and Worboys, 1995).

The applications of temporal data have been to detect change (for example Carlotto, 1997) where differences are detected in subsequent images, and to use overlays to determine correlations in spatio-temporal patterns (Unwin, 1996).

Once a significant amount of spatio-temporal data has been collected the database can be analysed for patterns. This process is known as data mining. It is:

"...the extraction of implicit knowledge, spatial relations, or other patterns not stored explicitly in spatial databases" (Koperski and Han, 1995)

For this type of knowledge acquisition to be successful, a long period of data acquisition is necessary. In most remote sensing applications, this data is not yet continuously available in a systematic form.

5.7 CONSOLIDATION KNOWLEDGE ACQUISITION

During knowledge acquisition for Geographic Expert Systems, several knowledge-bases can be created. These can contain different rules which define the same object due to knowledge acquisition from more than one expert, knowledge acquisition at more than one session, knowledge acquisition using different images or a combination of all three. Knowledge of how to consolidate these knowledge-bases, resolving conflicting and incomplete rules, is required.

5.7.1 Primitive Knowledge

Primitive knowledge typically consists of rules relating to the image band or band combination used to identify the scene primitive, the spectral signature of the primitive, the type of object (area, line or point) and other information such as object size or orientation

Primitive rules are of the from :

IF Band = 1
AND max_pix_val <= 22
AND min_pix_val >= 18
AND month = 12
AND type = areal
THEN feature = open_water

There are two broad ways knowledge-bases consisting of such rules can be combined. Firstly the most restrictive set of rules can be generated or alternatively the most general set of rules can be produced. For example combining the two rules:

IF Band = 1
AND max_pix_val = 20
AND min_pix_val = 10
AND month = 12
AND type = areal
THEN feature = sea

IF Band = 1
AND max_pix_val = 18
AND min_pix_val = 8
AND month = 12
AND type = areal
THEN feature = sea

would result in

IF Band = 1
AND max_pix_val = 18
AND min_pix_val = 10
AND month =12
AND type = areal
THEN feature = sea

in the first case, and

IF Band = 1
AND max_pix_val = 20
AND min_pix_val = 8
AND month =12
AND type = areal
THEN feature = sea

in the second case. A third approach is to apply some form of certainty factor to the rules obtained. The two rules then become:

IF Band = 1	**IF Band = 1**
AND max_pix_val = 18	**AND max_pix_val = 20**
AND min_pix_val = 10	**AND min_pix_val = 8**
AND month = 12	**AND month = 12**
AND type = areal	**AND type = areal**
THEN feature = sea CNF 100	**THEN feature = sea CNF 75**

and both are held in the knowledge-base. One advantage of this method is that it avoids having an expert assign certainty values to each solution, as the system calculates them during knowledge consolidation. The most restrictive rule is given a certainty of 100 and more general rules a value between 50 and 100 depending on the number of unconsolidated rule cases. For example in the above case with only 2 rules to consolidate a certainty of 75 has been assigned. This will avoid the problems noted by Mak et al (1996) related to an expert thinking qualitatively rather than quantitatively, or having to resort to a knowledge 'czar' or preeminent expert as suggested by Barrett and Edwards (1995).

A more complex problem occurs when there is a mix of bands involved in identifying an object. For example a crop of peas may be identified on band 7, of a Landsat TM image, but to distinguish them from beans which have a similar response on that band, a temporary NDVI layer is needed to complete the discrimination. In this case a restrictive combination of ANDing the two rules is taken.

If the rules are

IF Band = 7
AND max_pix_val = 45
AND min_pix_val = 38
AND month = 11
AND type = areal
THEN feature = peas

IF Band = abs(b3-b4)/(b3+b4)
AND max_pix_val = 12
AND min_pix_val = 0
AND month = 11
AND type = areal
THEN feature = peas

They are combined into the single rule:

IF (Band = 7
AND max_pix_val = 45
AND min_pix_val = 38)
OR (Band = abs(b3-b4)/(b3+b4)
AND max_pix_val = 12
AND min_pix_val = 0))
AND month = 11
AND type = areal
THEN feature = peas

Hence only pixels which meet both criteria are selected as being of object type peas. Rules are combined by searching for the result. Hence in the above case the knowledge-base is searched for peas. Rules are only combined if they are of the same type and are in the same time period for example rules referring to areal features cannot be combined with those referring to line or point features even if the result has the same name.

5.7.2 Relationship Knowledge

The Relationship Knowledge rules in the knowledge-base fall into two categories, firstly those that are of the canonical method for heterogeneous reasoning type (Sharma and Flewelling, 1995). Secondly Relationship Knowledge rules which are user generated, that is a user chooses two features and the system generates relationships.

Examples of the canonical method for heterogeneous reasoning type are:

IF feature_X OVERLAPS feature_Y
THEN feature_Y OVERLAPPED BY feature_X

and

IF feature_X SOUTH OF feature_Y
  AND feature_Z SOUTH OF feature_Y
  AND feature_Z NORTH OF feature_X
THEN feature_Z BETWEEN feature_X and feature_Y

Since these rules are not user generated they can be assumed to be logically consistent and in the least redundant form.

Consolidation of user generated Relationship Knowledge usually require checking for logical consistency. For example consider the rule:

IF feature_X = open_water
  AND feature_Y SOUTH OF feature_X
  AND feature_Y OVERLAPS feature_X
THEN feature_Y = cloud

and a second rule

IF feature_X = open_water

AND feature_Y NORTH OF feature_X

AND feature_Y OVERLAPS feature_X

THEN feature_Y = cloud

These are inconsistent in the directional clause since a single rule cannot contain feature_Y being both NORTH OF and SOUTH OF feature_X. If there is more than one example of one of these rules then there is an increased possibility that it is the correct interpretation. In spite of this, user involvement to mediate is probably necessary. A system may help by generating messages providing an indication of which of the rules found is most likely to be correct.

5.7.3 Assembly Knowledge

Assembly knowledge can be used for map generalisation. Map generalisation requires the grouping of scene primitives so they can be represented as a single named object (McMaster, 1991). Since Assembly Knowledge is primarily a combination of scene primitives, two options are available. Either the most restrictive rule can be produced or the most general rule can be produced to describe the assembly object. The most restrictive rule requires an exclusive definition of the components of the object. General rules are generated by using clauses common to all rules being consolidated. For example consider the two rules:

IF AREA > 1000

AND AREA <5000

AND Feature_1 = pasture

AND Area_Feature_1 > 250

AND Feature_2 = peas

AND Area_Feature_2 < 250

AND Feature_3 = potatoes

AND Area_Feature_3 < 500

THEN Assembly_Feature = Cropping_Farm

and a second rule

IF AREA > 1000
AND AREA <5000
AND Feature_1 = peas
AND Area_Feature_1 < 250
AND Feature_2 = cereal
AND Area_Feature_2 > 500
AND Feature_3 = potatoes
AND Area_Feature_3 < 500
AND Feature_4 = poppies
AND Area_Feature_4 <100
THEN Assembly_Feature = Cropping_Farm

which would result in

IF AREA > 1000
AND AREA <5000
AND Feature_1 = peas
AND Area_Feature_1 < 250
AND Feature_2 = cereal
AND Area_Feature_2 > 500
AND Feature_3 = potatoes
AND Area_Feature_3 < 500
AND Feature_4 = poppies
AND Area_Feature_4 <100
AND Feature_5 = pasture
AND Area_Feature_5 > 250
THEN Assembly_Feature = Cropping_Farm

In the most restrictive case and:

IF AREA > 1000
AND AREA <5000
AND Feature_1 = peas
AND Area_Feature_1 < 250
AND Feature_2 = potatoes
AND Area_Feature_2 < 500
THEN Assembly_Feature = Cropping_Farm

in the most general case when consolidated.

The general rule case provides the domain expert with the most flexible solution. The domain expert can then be provided with the capability of modifying the consolidated rule by adding, deleting or modifying clauses as required. This process allows for generalisation of the form used by Rigaux and Scholl (1994).

5.7.4 Heuristic Knowledge

Unlike knowledge captured from other tools, heuristic knowledge derived from the repertory grid tool uses the domain expert's terminology to differentiate between features on an image. Any reference to pixel values will tend to be in terms of fuzzy concepts such as high or low value. Hence rules are more likely to be of the form:

IF albedo = low
AND area = small
AND feature-origin = man-made
THEN feature = reservoir

The features which domain experts are most likely to identify with repertory grids are equivalent to those identified using Primitive and Assembly Knowledge and could be directly consolidated with them. However this would require classifying the rules generated by the repertory grid. It is therefore better to deal with them separately and use a consolidation technique similar to that used for Primitive Knowledge.

To deal with fuzzy concepts procedure to define the limits of the fuzzy sets is required. This will have to be set up on a case by case basis because small in the above example would have a different definition in a different context.

5.8 INTERPRETATION KNOWLEDGE ACQUISITION

Interpretation knowledge is meta-knowledge in that it is knowledge about where to apply knowledge. In the KADS methodology, the equivalent knowledge is defined as being a Task Layer, and is concerned with how to apply the available knowledge-bases to a specific problem (Fensel and Van Harmelen, 1995). In the case of spatial systems where there may be several digitised data sets and images as well as several knowledge-bases, it means determining which knowledge-bases are relevant to the particular activities being carried out. In the case of agricultural systems, it may be necessary to determine region similarity before applying a knowledge-base developed in one area to another because of differences in soil, climate and horticultural practice.

Acquiring Interpretation Knowledge has not been a significant problem to date because most spatial systems have been developed with a specific application in mind. For example crop recognition systems are developed within a specific region and are only intended for application within that region. Similar development techniques may be applied elsewhere but the actual knowledge acquired will generally be significantly different for different regions and needs to be acquired on a region-by region basis.

5.9 CONCLUSIONS

Spatial knowledge as defined in Chapter 2 can be acquired using both visual and non-visual tools. Acquisition of Primitive, Relationship and Assembly Knowledge is primarily visual and knowledge acquisition tools for these types of knowledge should reflect that. Consolidation Knowledge is more task oriented (Fensel and Van Harmelen, 1994) in that it represents a fixed strategy for combining knowledge from multiple knowledge-bases and multiple domain experts. Interpretation Knowledge is strategic knowledge, and although the systems reviewed are very specific, its importance in choosing which knowledge-bases to apply in particular cases will grow as more general spatial knowledge-based systems are developed.

The Non-Visual Knowledge tools are necessary because of the fact that some geographic knowledge cannot be drawn or seen on an image. Such tools are capable of capturing algorithms and heuristics.

Temporal knowledge requires the implementation of a database to trace spatial change over time. Once the database has been established and sufficient data has been collected, data mining techniques can be used. One problem with this technique is, that depending on the spatial features being studied, it may take several years of data acquisition before useful results can be produced.

Chapter 6. THE KAGES SPATIAL KNOWLEDGE ACQUISITION TOOL-KIT

This chapter will investigate the tools required to acquire five of the six types of geographic knowledge introduced in Chapter 2. The development of a tool-kit for testing the tools is described. The final form of the tool-kit, KAGES (Knowledge Acquisition for Geographic Expert Systems), and the implementation of its components is presented. A detailed discussion of how to use the various tools in the tool-kit is given in Appendix A.

6.1 Overview

The first five of the six levels of geographic knowledge are acquired by various tools which have been combined into an integrated tool-kit. The practical implementation has been designed in such a way that new tools can be added as they are developed. The reverse is also possible. A tool can be extracted and customised for a specific task on a particular system.

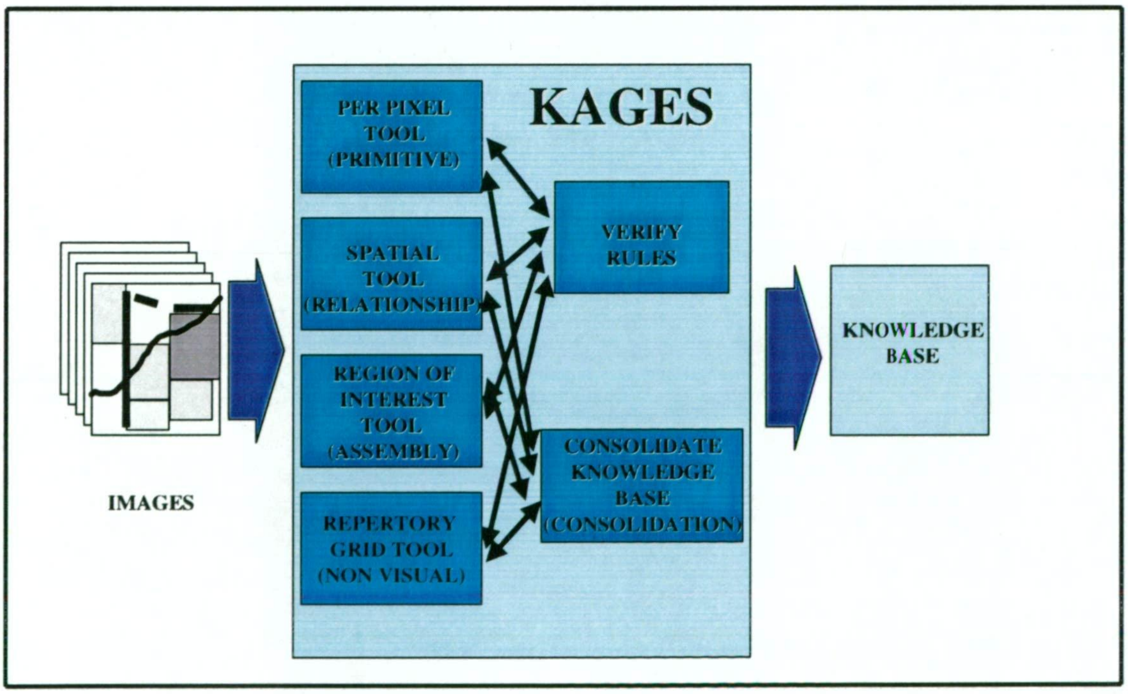

Figure 6.1 Overview of the KAGES system showing the primary knowledge acquisition tools

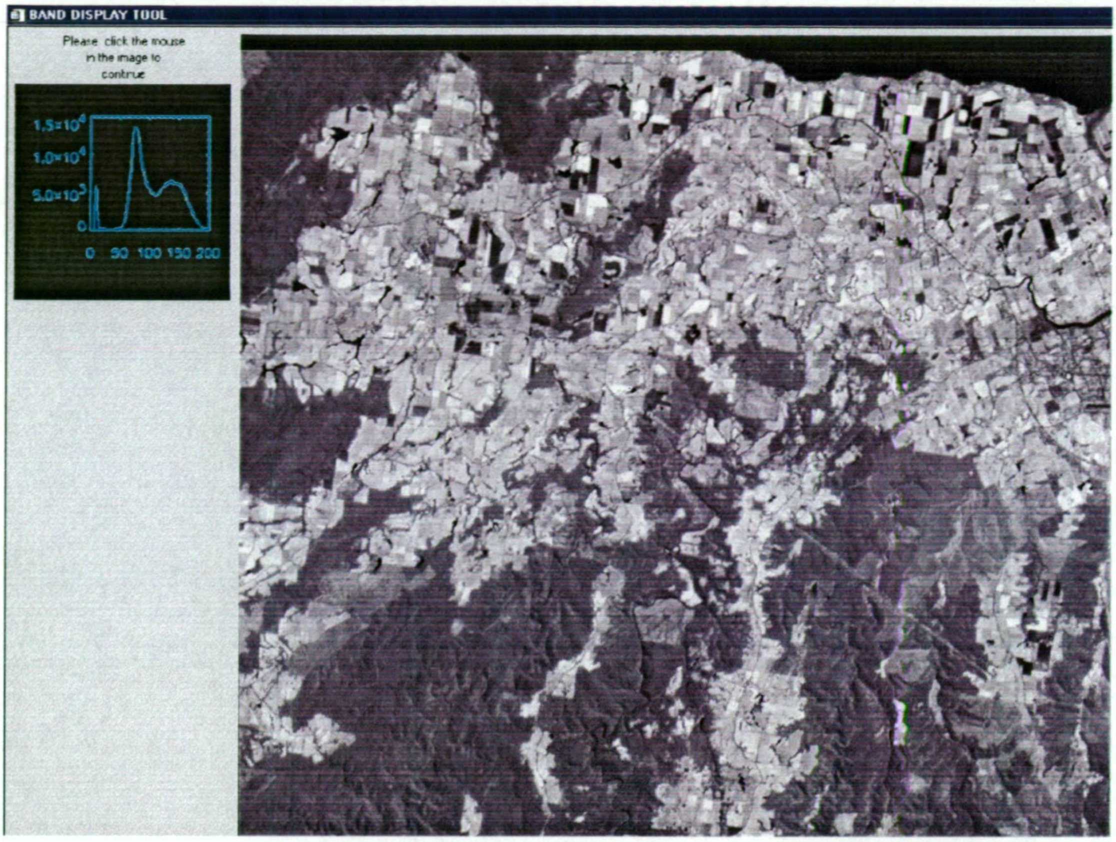

Figure 6.2 The band viewing tool showing the image histogram and band 3 of a SPOT image of Table Cape area in Tasmania, taken on 24th December 1996

6.2.2 Identifying Areal Features

Areal features are two dimensional features which can be defined on an image band or band combination. On that band or band combination they have pixel values which fall between two thresholds. Although a domain expert may define a feature directly in terms of its spectral signature, the usual way for a user to describe a feature is to point to it.

Figure 6.3 shows the general layout of the per pixel tool which is used by the domain expert to identify domain primitives. The system automatically determines a spectral range of pixel values of an object, once the expert has pointed to it. If the expert wants to increase or decrease the thresholds, they can adjust them using slider bars. All pixels with a value within the threshold range and which are contiguous with the initial pixel are then grouped and highlighted.

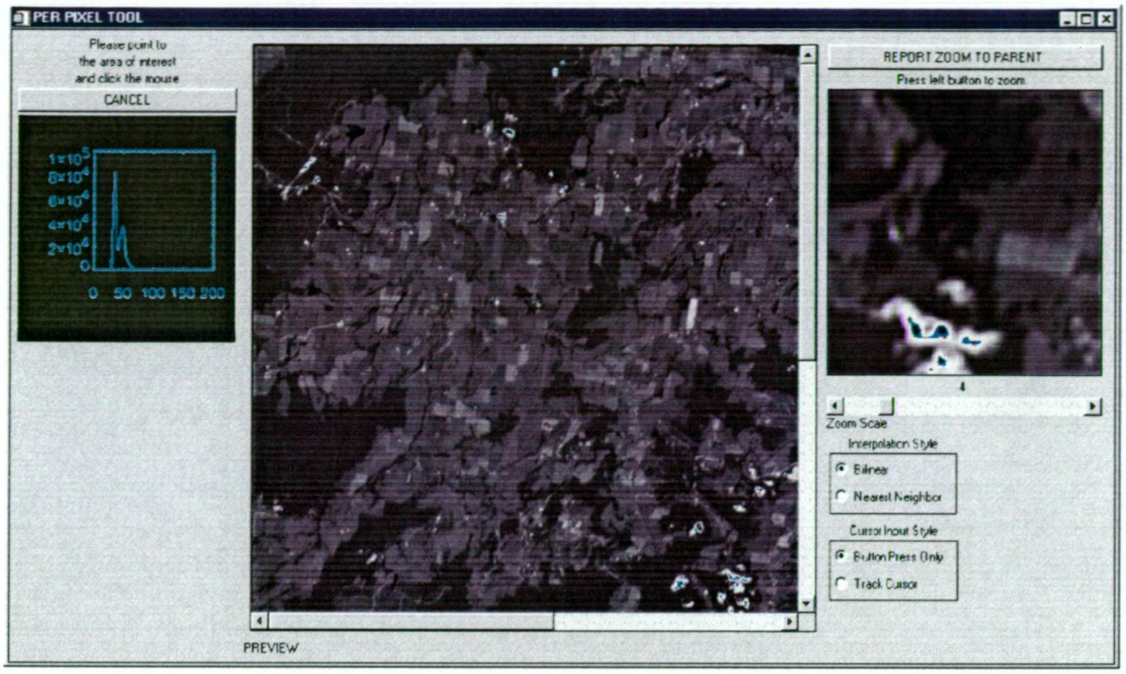

Figure 6.3 Per Pixel Tool. The window on the left shows the histogram of pixel values for the entire image. The centre window is of the image itself while the right window contains a zoomed section of the image to allow for accurate selection. From band 1 of SPOT image of Table Cape, Tasmania taken on 2nd November 1996

A Minimum Bounding Rectangle (MBR) is drawn around the feature and a histogram of pixel values of the feature is also displayed in a format similar to that shown in Figure 6.3. If the threshold values are too high or too low the expert can reset thresholds; otherwise the feature is named and the information about extent and thresholds is recorded when the accept and return button is pressed. A more detailed description of the operation of the tool is provided in Appendix A.

6.2.3 Identifying Point Features

Point features are the simplest features in images since they are subpixel and less than a pixel in size. Because of this they are also the most difficult to locate. One method of acquiring knowledge about a point feature is to have the domain expert use a pointing device such as a mouse to indicate its location. The point is displayed on the image and the user can then give the point a name and type.

The second method of identifying point features results from the use of the spatial analysis tool where intersections between other features can be named and stored. This will be discussed further (Section 6.4.4).

6.2.4 Identifying Line Features

Lines are probably the most complex objects to work with. KAGES attempts to gain as much information as possible about lines and then allows the expert user to edit out irrelevant or spurious information (Crowther and Hartnett, 1996). It achieves this objective in two steps. Individual lines are identified with information about their length and orientation. Secondly a rule defining the line is displayed to the user who can then interactively edit out irrelevant clauses.

Individual line objects are identified using two different methods. A line following algorithm is available which automatically follows lines from a given point and requests information about names. This produces a raster data set of pixels along the line together with an average trend of the line (which may or may not be relevant).

The line following algorithm relies on a user picking two points on a line they can identify in an image. The average of the two pixel values is used to do a pixel by pixel trace by finding adjacent pixels with similar characteristics. An end point is reached when there are no further pixels of similar characteristics. To speed up the tracing, pixels lying in the same direction as that section of the line previously identified are tried first. The algorithm also prevents backtracking down the line. The resultant pixel set is stored as an array, each element of which contains the coordinates of each pixel in the set. This array can be recalled at any time, along with information about average trend, start point and length of the line being represented.

In the case where the resultant line is not what the user expected, the definition can be abandoned. This is usually a result of setting threshold values for the line pixels too high or too low. Once the user is satisfied with the result they are asked to give a name for the line and specify the line type.

A second method is to manually trace a line using a pointing device. This is a useful alternative where the line involved either cannot easily be picked up by the line follower or the line is not pixel based (for example cadastral data such as a municipal boundary). In this case the resultant line is a vector although the number of points involved is directly related to speed and operation of the mouse. Gaps in the line are then filled in for display purposes.

6.2.5 Primitive Knowledge Data Structures

Laurini and Thompson (1992) define five categories of information for spatial entities. These are:

- An identifier,
- A locator,
- The character of the entries,
- The role behaviour or function of the entry and
- Spatial properties of the entry.

Of these, only the fourth need not be stored as part of the feature identification. It would be a necessary part of a GIS database however (Appendix B).

Points:	point name, point type, location, band used, image
Lines:	line name, line type, spectral value (0 if irrelevant), band used, actual line, points, image
Areas:	area name, area type, spectral signature, band used, MBR coordinates, centroid coordinate, area, image.

Table 6.1 Data structures for scene primitives

Despite requiring these common categories of information, data structures in the system differ in detail with the type of feature. Hence there are different data structures for points, lines and areas (Table 6.1). All objects have a specific name (identifier) and a type name (character). For example a road may have a specific name of Midlands Highway and a type name of A-route. Other fields common to all types of objects are the image id and the band or band combination (from the Band

Calculator, section 6.7) used in identification of the object as well as a method of locating the object (position).

Depending on the object, spatial properties such as the surface area and orientation of an areal object or the length and orientation of a line object may also be stored. The form of the data structure is transparent to the user who can retrieve and manipulate these objects using the name they assigned to them.

6.2.6 Primitive Knowledge Rules

The rules generated from the primitive tool depend on the type of object being investigated. The form of the rule is based on the data structure used to store information about scene primitives. Unlike the data structure, the rule is available to the user to view and modify.

Rules have the form:

```
IF      band = 1
        AND pixel >= 18
        AND pixel <= 22
        AND type = area
        AND area > 1000
        AND month = 1
        AND image = NOAA
THEN
        feature = open_water
```

This method of knowledge representation has been chosen because it:

- Is a natural way of expressing rules of thumb,
- enables modular organisation of knowledge and
- has a restricted syntax.

Leung (1997)

6.3 RELATIONSHIP KNOWLEDGE ACQUISITION: THE SPATIAL RELATIONSHIP TOOL

One of the features missing from many geographic information systems is a method of defining spatial relationships (Openshaw, 1991). The lack of this tool has limited the use of many of these systems. It is something which will be included in new generation GIS.

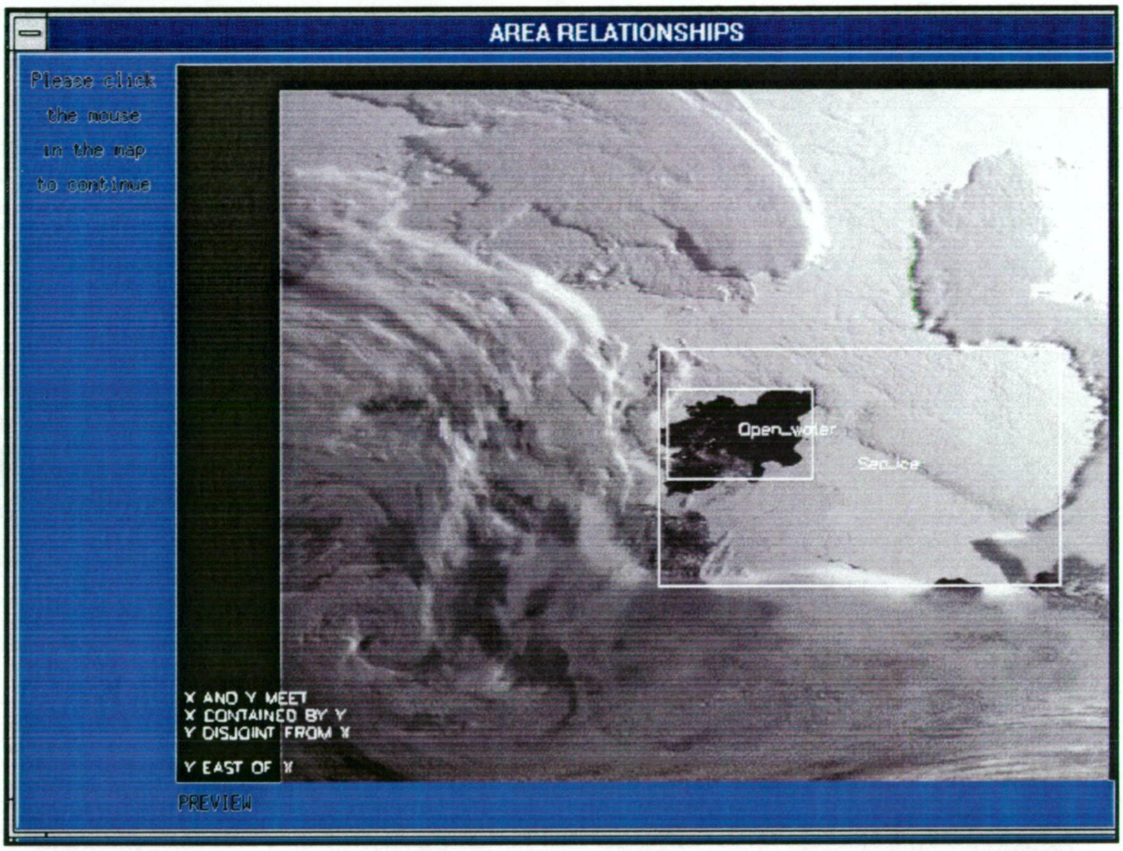

Figure 6.4 Relationship Tool. The first object (X) is open water defined in band 1. The second (Y) is sea ice identified from band 3 of the NOAA image. The results are shown superimposed on band 1.

The method of acquiring Relationship Knowledge depends on the type of object involved. In the case of area objects (Figure 6.4) this is done initially by comparing the Minimum Bounding Rectangles (MBR's) (Chang and Jungert, 1996) of the objects involved. If they are close so that the MBR's overlap, the objects themselves are compared to see if they overlap or touch. In the case of point and line objects the actual features are compared.

Implementation of the tool to compare two objects uses Egenhofer's (1991) eight point classification of overlap consisting of the following overlap clauses: *disjoint, touches, overlaps, contains, covers, equals, is contained by* and *is covered by.* Orientation is based on the direction of the second object's centroid in relation to the first to within 45°. The distance between the centroids of the objects is used for nearness. This is a fuzzy concept and is classified in terms of being the same point, very near or near.

A rule editor is provided to allow a user to check the spatial relationships KAGES has found. These relationships are in an IF THEN rule format which is easy for the user to verify. The expert can then either accept the relationships or remove those which are considered irrelevant or are due to chance.

6.3.1 Relationships Involving Lines

Reasoning about two discrete lines is difficult. Lines unlike other objects can have a beginning and an end, they can have a direction or trend, and they can, due to their interaction with other objects define point and even areal features (Fleck, 1996). Lines do not always have a beginning or an end (a circuit being an obvious example) or a direction.

The approach taken in the KAGES system to incorporate spatial reasoning with lines was to deal with simple cases first. The most simple case involves two lines which do not intersect or enclose and where proximity and orientation of one relative to the other is determined. A slightly more complex case involves determining the level of parallelism between two lines (for example where power lines follow a road).

The next level of complexity involves a simple intersection of two lines or one line terminating the other line (for example a road T-junction). More complex interactions occur where lines follow each other for all or part of their length, (for example when a road partially follows a river for part of its length. In this case there is more than one pixel of the first line adjacent to pixels of the second line. In the most complex cases the lines are intertwined.

6.3.2 Relationships Involving Points

Two points have the simplest number of relationships. The points are either at the same location or they are not. Proximity can likewise be easily determined as the distance between the two points. Lastly the direction of one point relative to the other is a simple calculation.

Points in relation to an area can be described in terms of their location relative to the boundary. A point can be outside an area and not touching it, outside an area and touching the boundary, on the boundary, inside the area and touching the boundary, or inside the area and disjoint from the boundary. Distance can be calculated in terms of distance from the areas boundary, or distance from the areas centroid. Orientation is in relation to the area object's centroid.

There are similar relationships between points and lines. A point can be on a line, next to it or located at an end.

6.3.3 Relationships Involving Areas

The majority of work in spatial relationships has been in terms of area-area relationships. Most of this in turn has been based on Allen intervals as discussed above. When dealing with relationships between primitives in a single image frame, the temporal aspects of the Allen classification are not needed. However when combined with a database the resultant relationships can be used. The temporal aspects of knowledge of areal features are being considered as part of a process of developing a non-visual tool which uses databases.

The current system however uses only the spatial Allen interval relationships in this particular tool. As well as the degree of overlap, the system also determines the proximity and direction of one centroid from another.

Relationships with line and point objects have already been discussed. In both cases these types of objects are dealt with in terms of their relationship to the areal objects boundary.

6.3.4 Relationship Knowledge Rules

Relationship knowledge is stored directly as rules in the following form:

IF feature_X = primitive_1
[AND feature_Y *DIRECTIONAL_CLAUSE* feature_X]
[AND feature_Y *PROXIMITY_CLAUSE* feature_X]
[AND feature_Y *OVERLAP_CLAUSE_1* feature_X]
[AND feature_Y *OVERLAP_CLAUSE_2* feature_X]
THEN
Feature_Y = primitive_2

For example:

IF feature_X = open_water
AND feature_Y NORTH OF feature_X
AND feature_Y OVERLAPS feature_X
THEN
feature_Y = cloud

This knowledge is in a form which is directly verifiable by the domain expert. Modules also exist to parse these rules for automated verification across an image.

6.4 ASSEMBLY KNOWLEDGE ACQUISITION: THE REGION OF INTEREST TOOL

A third tool available to the expert allows the interactive definition of a region of interest on an image (Figure 6.5). This tool is required because users group primitives into identifiable units which can then be referred to as a whole; a process known as generalisation (Beard, 1991). The component primitives may or may not be related in terms of spectral pixel values or other characteristics.

Figure 6.5 The Region Of Interest Tool. The defined area in the top left has an MBR drawn around it. The pop up window shows the object types found within the area. (Landsat band 1, 27th November 1996)

The system allows a user to draw a closed curve (Olson and Rueter, 1987) to define a region of interest; then determines all object classes which fall within the region and displays them. It also provides the user with information on the composition of the region of interest by providing the percentage area each of the component classes cover. The user can give the region a name and interactively edit membership. This allows the expert to generate assembly relationships between point, line and area features identified on different bands or band combinations in an image.

This tool is visual and allows a user to isolate groups of scene primitives into assemblies of related objects with a specific identification. The tool is manual and allows a user to visually generalise. It allows a user to nominate what are essential components, what are optional (but support the grouping) and what are irrelevant (does not matter if they are there or not). As a result a rule would define the maximum and minimum size of an assembly object, the essential components, the optional components and the density thresholds of each of the components.

The actual processing of this kind of knowledge is similar to processing a image in that after identifying the assembly object by tracing it, the system identifies scene primitives within its boundary then segments the area. Once this has been done the population and area density of various object can be calculated.

The Region Of Interest tool can be used to group features with widely differing (or in the case of point objects zero) pixel values into an areal classification. It can also be used to determine what objects are within or near a particular feature.

6.4.1 Assembly Knowledge Rules

Rules of this type are made up of a list of the components of the assembly object, their characteristics and the characteristics of the object as a whole

IF [Area_Component_1 = Area_Object 1]
[AND Area_Component_1_Area *conditional_clause* Area_object_1]...
[AND Point_Component_1 = Point_Object 1]...
[AND Line_Component_1 = Line_Object 1]...
[AND Object _Area *conditional_clause* real number]
THEN
Assembly_Object = Assembly_Object_Name

One or more of these clauses need to be key clauses identifying an object or objects. These key clauses must be both essential and sufficiently characteristic to make searching an image for them, and hence the assembly object, as simple as possible.

An example of an Assembly Knowledge rule is:

IF AREA > 1000
AND AREA <5000
AND Feature_1 = pasture
AND Area_Feature_1 > 250
AND Feature_2 = cereal
AND Area_Feature_2 > 500
THEN Assembly_Feature = Cropping_Grazing_Farm CNF 75

Which contains two areal object types pasture and cereal

6.5 NON-VISUAL HEURISTIC KNOWLEGE ACQUISITION: THE REPERTORY GRID TOOL

Repertory grids and Personal Construct Theory were discussed in detail in Section 4.2.3. The system allows for a domain expert to develop a new grid (and hence knowledge-base) or to review an existing one. The grids are stored as data sets consisting of three arrays containing *elements* which are the objects identified on an image, *concepts* which are used to differentiate between the elements, and *ratings* which measure how closely an element displays a concept. For ease of reading, the term concept will be replaced with *discriminator* as this gives a better indication of its use. All can be modified or expanded.

The expert is presented with a graphical user interface which displays the image to be analyzed. An initial menu gives the user a choice of developing a new grid or loading one from file.

If the user opts to develop a new grid they are presented with a data entry window (Figure 6.6). The domain expert is prompted to name the regions or objects on the image. These become the elements used in repertory grids.

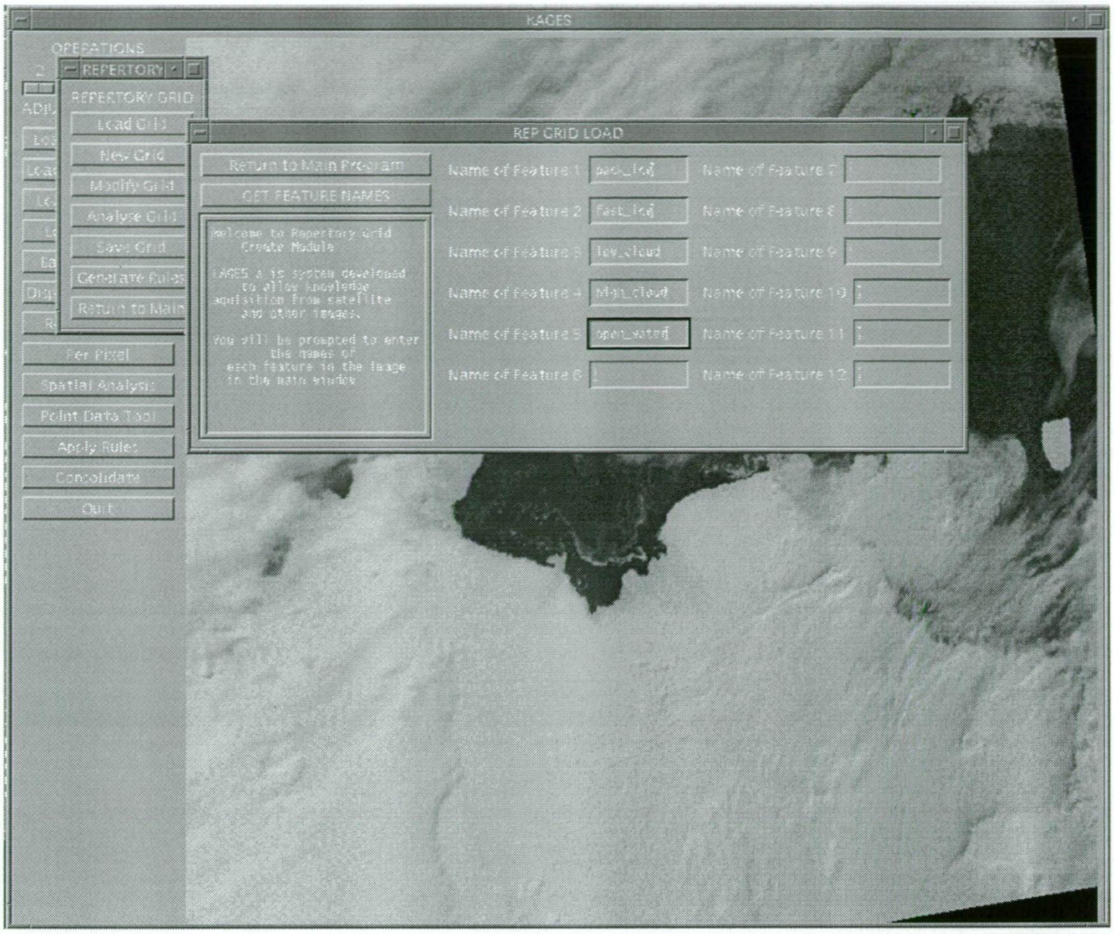

Figure 6.6 Repertory Grid Tool with menu options and the window to enter elements displayed (NOAA VHRR image of Casey taken on 26th February 1998)

A triad method is used to determine discriminators (Figure 6.7). A domain expert is presented with three of the objects they have identified. They are then asked to define a discriminator which will differentiate one of the objects from the other two. This forces the expert to state why there is a difference and not just name features. This process is repeated with various combinations of the elements the user has defined.

Spatial relationships as well as threshold information can be used by the expert in this delimitation. Since domain experts use visual analysis, they are more likely to give a spatial explanation than a threshold value. Threshold values will be expressed in fuzzy terms such as high or low rather than absolute terms.

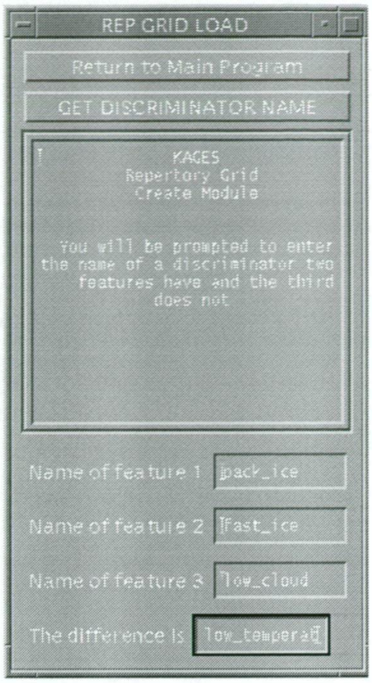

Figure 6.7 Triad window for determining discriminators

Following the identification of the discriminators the expert then ranks each of the regions or objects against each of the discriminators on a scale of 1 (does not exhibit any of the discriminator) to 5 (exhibits discriminator completely) using slider bars (Figure 6.8).

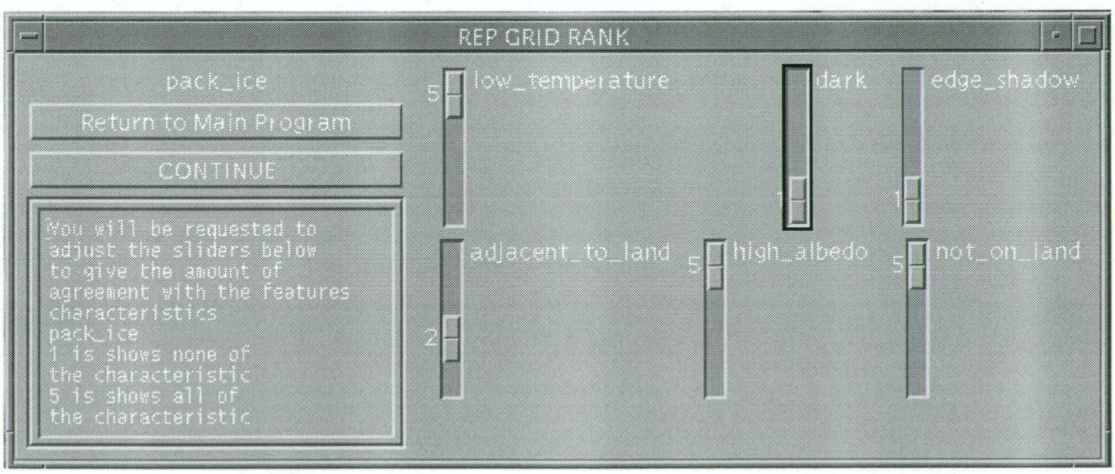

Figure 6.8 Ranking an element (pack_ice) against discriminators

The repertory grid produced at the end of this process (Figure 6.9) can be immediately analyzed using clustering techniques, saved to disk or converted to rules. Any combination of the three processes may be chosen.

ELEMENTS

pack_ice	fast_ice	low_cloud	high_cloud	open_water	land	**DISCRIMINATORS (CONCEPTS)**
5	5	4	5	1	2	low_temperature
1	5	2	1	5	4	adjacent_to_land
1	1	2	1	5	4	dark
5	5	4	5	1	3	high_albedo
1	1	5	5	1	4	edge_shadow
5	5	3	3	5	1	not_on_land

Figure 6.9: repertory grid of Antarctic elements (features) with 5 indicating complete agreement with discriminator (concepts) and 1 indicating complete disagreement with the discriminator.

The maximum number of regions which can be included in one grid is 12. If there are more than 12 regions, some grouping of elements needs to be considered and a hierarchy of grids developed. This in itself is a spatial relationship as it is likely to identify sub regions. That is a relationship between Primitive and Relationships or Relationships and Assembly Knowledge.

6.5.1 Analysis

To aid with the knowledge acquisition, the data in the grids is analyzed using Johnson Hierarchical Clustering Techniques (Olsen and Rueter, 1987). The result of this is displayed as a clustering graph which domain experts may use as the basis for describing a classification scheme (Tranowski, 1990). If such a classification can be identified it may become the basis of further discussions with domain experts over the rating of certain regions or objects and their relationship to other objects, as well as the usefulness of certain discriminators.

The system can generate rules at any time once the grid has been developed. Currently these concentrate at the extremity of the ratings. Hence a rating of 1 means the concept is in favour of the item while a rating of 5 gives a negated antecedent.
For example the Figure 6.6 grid would generate a rule for pack_ice as follows:

IF low_temperature
AND NOT adjacent_to_land
AND NOT dark
AND high_albedo
AND NOT edge_shadow
AND NOT_on_land
THEN feature = pack_ice

Once rules have been generated these can be discussed with the domain expert. In some cases further concepts may be needed to delineate some of the image elements. For example, in Figure 6.6, pack_ice and fast_ice are very similar when threshold values are considered. The only way to tell them apart is by use of spatial relationships. Pack ice is ice that floats on water as rafts of ice. Fast ice is attached to the land (fastened). Per pixel classification cannot delineate ice with these classifications and groups them together

6.6 NON-VISUAL ALGORITHMIC KNOWLEGE ACQUISITION, THE BAND CALCULATOR TOOL

The KAGES Band Calculator (Figure 6.10) allows a user to combine image bands from either SPOT, Landsat or NOAA AVHRR images and immediately see the result of applying a band combination algorithm which has been interactively entered as an equation. The algorithm can then be named and stored for reuse or can be discarded without having to hard code it into the system. Once the algorithm has been stored, it can be retrieved and used in other classification systems and the results of applying it can be viewed at any time.

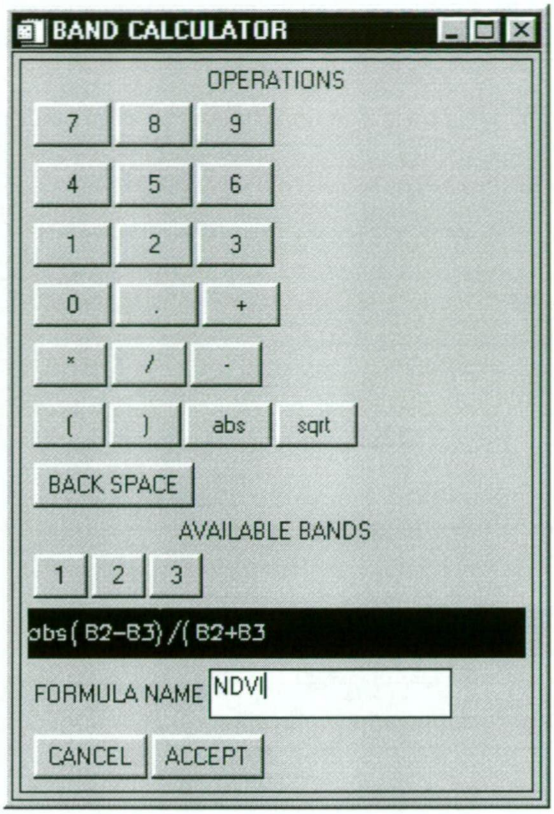

Figure 6.10: Band Calculator to generate a Normalised Difference Vegetation Index for a three band SPOT image set.

As the name suggest, the tool looks like a calculator to the user. It has extra buttons to represent the raw bands of the satellite image being analysed as well as extra buttons representing a selection of useful mathematical functions. The calculator is written in Research System Inc's IDL language and its results are also executable by IDL. Therefore the algorithms can be placed as executable variables within code which classifies images or code which is used for knowledge elicitation in the rest of KAGES.

The band combinations are used for identifying scene primitives using other tools in the KAGES tool-kit. If an areal object is identified using the Normalised Difference Vegetation Index (NDVI) which has been entered via the Band Calculator, that object's name, pixel thresholds (using the NDVI calculated pixel values), size and the band combination algorithm are stored as a rule. That rule can then be applied either individually to similar objects for verification, or as part of the rule set used for the classification of an entire image. The results can also be used to generate rules representing spatial relationships between two objects, where one object is identified using one combination of bands and the other object is identified using a different band combination.

6.7 A TOOL TO ASSIST MACHINE LEARNING AND VERIFICATION: THE POINT DATA TOOL

This tool was developed in direct response to a user request for a tool to add data to a decision tree package, S (Chambers and Hastie, 1992). It was then found that the same tool could be used to create training data for use with machine learning systems. This means that KAGES can be loosely coupled to a third party machine learning package.

The tool's operation gives the user the choice of automatically sampling an image or allowing a user to sample the image in a directed way. In the former case the user can specify the coarseness of the sampling grid and then is asked to supply names at each sampling point. The user can switch between bands to decide on their naming. It is also possible to switch between automatic sampling and user-directed sampling. This feature was added because if a coarse grid was chosen for automatic sampling, some important features could be missed.

In the user directed version the user still has the ability to switch between bands, but the sample points are located using the mouse. It is possible to combine the resultant files to produce a composite sample.

At each point sampled the tool writes data to an output file in the form of:

- A pixel location (X,Y coordinates),
- Pixel values (on all bands in the image) and
- A user classification of the pixel.

Band 1	2	3	4	5	classification	pixel location	
58,	58,	139,	152,	157,	cloud,	178,	716
46,	41,	141,	152,	155,	cloud,	356,	716
58,	55,	124,	171,	177,	cloud,	534,	716
66,	62,	112,	166,	171,	cloud,	712,	716
48,	44,	133,	155,	158,	cloud,	178,	895
56,	54,	121,	170,	173,	cloud,	356,	895
69,	65,	163,	169,	172,	continental_ice,	534,	895
78,	73,	162,	175,	177,	continental_ice,	712,	895
19,	15,	120,	118,	121,	water,	499,	461

Table 6.2 Partial transcript of a file generated using the Point Data Tool on a NOAA AVHRR image. The first 8 were generated using a regular grid while the last is the result of user directed sampling

This tool allows the production of a quality spatial sample of data (Table 6.2) to be used for training. The sample size used will vary depending on the characteristics of the feature being studied and the resolution of the sensor (Curran and Williamson, 1986). The quality of the spatial sample can have a major effect on the result achieved using machine learning techniques. The two methods available in this tool, sampling by grid and user directed sampling, allow a training data set with a sample over the whole image to be produced. This can be followed up with more intensive sampling of features of interest if necessary.

6.8 CONSOLIDATION KNOWLEDGE ACQUISITION: THE CONSOLIDATION TOOL

The purpose of the Consolidation Tool is to combine knowledge-bases which have been generated at different sessions. Multiple knowledge-bases for a single domain generated by an individual tool may be the result of there being more than one domain expert or more than one set of training images (Barrett and Edwards, 1995). The results from these various sessions are held in different knowledge-bases which need to be combined into a single knowledge-base. The process of consolidation of each of the knowledge-base types is described in Section 5.7.

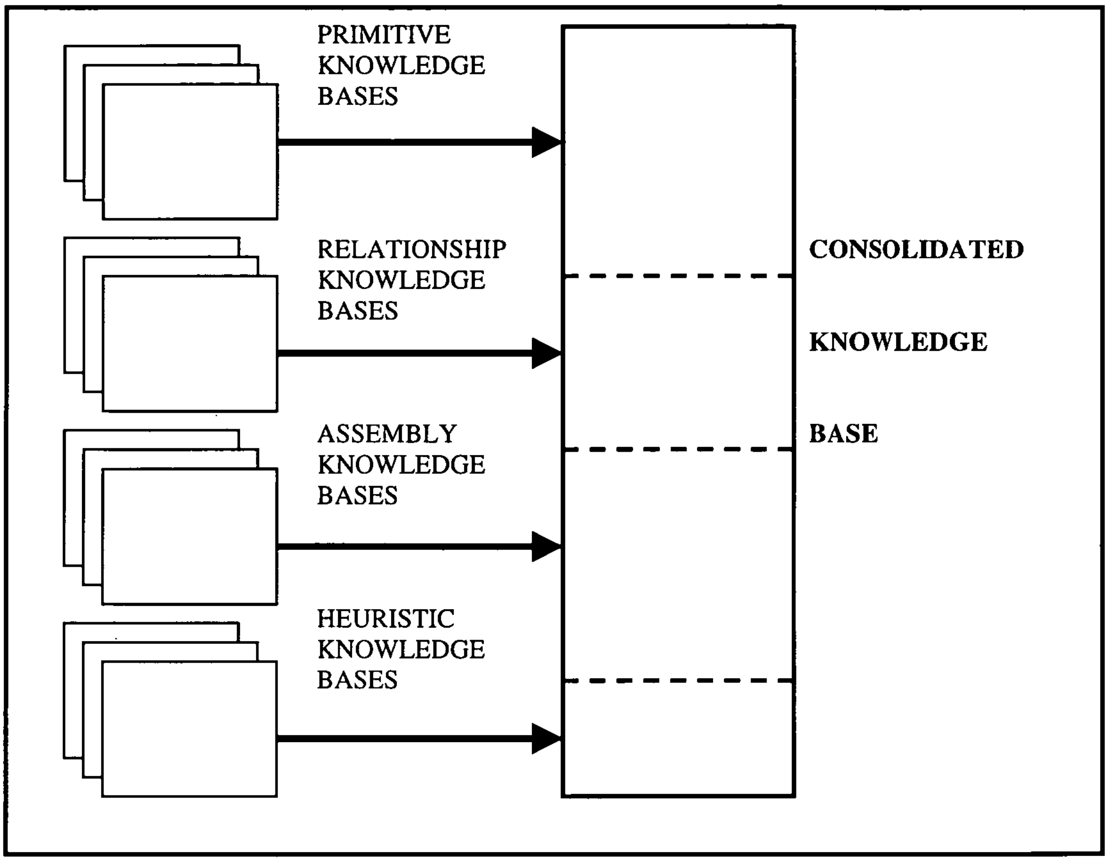

Figure 6.11 Consolidation of Knowledge-bases

Once each of the individual geographic knowledge-bases containing knowledge acquired by using the individual tools have been resolved, the resultant knowledge-bases are combined into a single segmented production knowledge-base of Primitive, Relationship, Assembly and Heuristic Knowledge (Figure 6.11).

6.9 KNOWLEDGE VERIFICATION TOOL

Once rules have been generated by the system they need to be verified. Verification of geographic knowledge requires a combination of knowledge-based system verification and spatial verification. The formal methods of knowledge-base system verification (Meseguer and Preece, 1995) need to be combined with methods for assessing the accuracy of classifications of remotely sensed data (Congalton, 1991).

The method adopted here involved the following verification steps:

- Visual inspection of rules by the domain expert.
- Applying a single rule to the entire training image and comparing with ground truth.
- Applying all rules from an unresolved knowledge-base to the training image and comparing with ground truth.
- Applying all rules from a consolidated knowledge-base to the training image and comparing with ground truth.
- Applying rules from the consolidated knowledge-base to other images and comparing with ground truth.

During the knowledge engineering sessions, the first three steps are used. Once sessions are complete and the knowledge-bases have been consolidated the fourth step is applied. During most knowledge engineering sessions, only the first two steps are used because all subsequent steps involve complete image classifications which are time consuming and therefore are best done between sessions with the resultant classifications being checked by the domain expert at the following session.

It should be noted that the tool is not designed to be a full image classifier. However elements of a classifier must be built in for verification. As a result, Primitive Knowledge is always applied first followed by Relationship Knowledge. For ease of verification each of the Relationship rules are applied separately and information on the number of correct classifications produced compared with of the number of times the rule fired, the number of true clauses in the rule and the number of positive outcomes from the rule, is presented.

6.10 THE USER INTERFACE

The user interface is of primary importance with any tool (Openshaw and Clarke, 1996). To produce a general purpose Spatial Knowledge Acquisition system some level of interactive graphical interface needs to be incorporated because at least one image (that under consideration) will need to be displayed. Several approaches could have been taken to interface design. One approach uses a single primary window with operators along the top and tools to use with those operators down the left hand side (McMaster and Mark, 1991). Tranowski (1990) uses a simple approach with operators which expand into tools for specific tasks when selected. This method was also used by Avouris and Finotti (1993).

The approach taken with the KAGES toolkit was similar. The IDL Graphical User Interface (GUI) was used to develop an interface with the same general layout as that of the previously developed Icemapper system (Williams et al, 1997). This has a primary set of operators down the left of the screen. These call sub windows and menus as required (Figure 6.12). The development environment is extremely modular with the overall system broken into many modules, some of which are common to Icemapper. The interface relies on the use of features such as drawing tools, slider bars, dialog boxes and buttons to interact with the knowledge engineer or domain expert. Typing is kept to a minimum and is generally restricted to naming. This met the criteria listed by Medjckj-Scott (1994) in terms of functionality and utility and the ability to display data in a form familiar to GIS and remote sensing users.

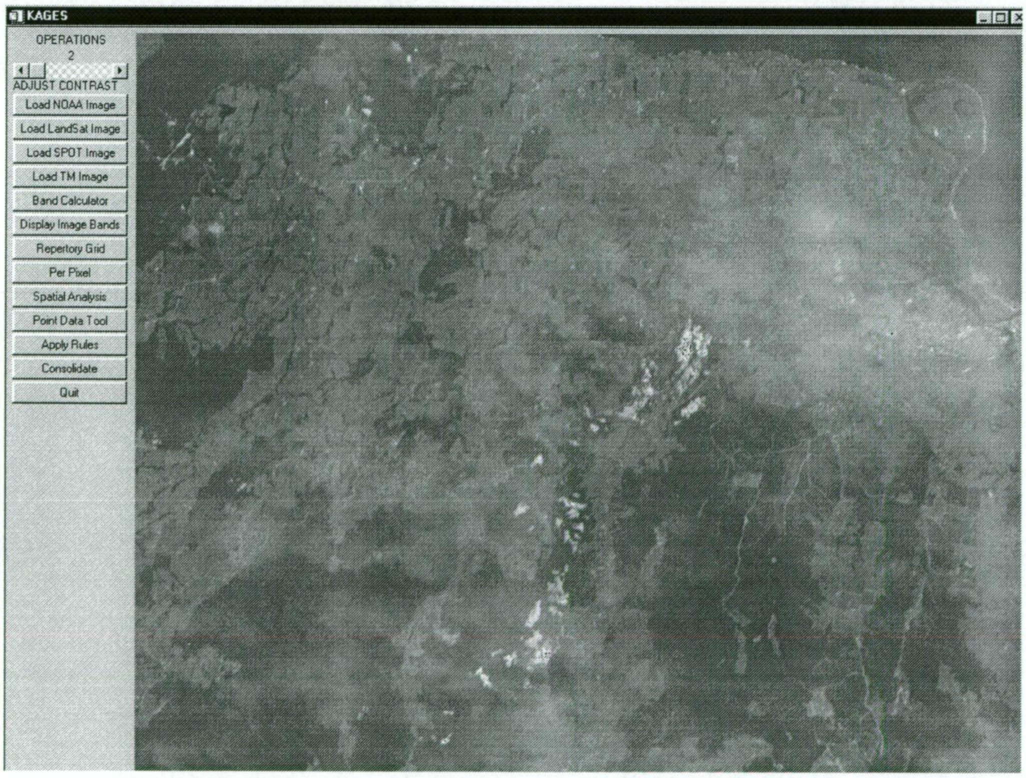

Figure 6.12: KAGES main user interface with a SPOT image loaded. The image is of Table Cape, North West Tasmania and was taken on 2nd November 1996.

The initial action of the system is to load an image to be used as a prompt or work surface by the domain expert with band 1 of the image being displayed in the display window. Other bands can be displayed when required as the image set is held as a three dimension array with the first dimension representing the image band. The other two dimensions in the array are representations of the images. This structure can be viewed in the traditional GIS manner as a series of layers representing each band of the image. Extra temporary layers are added to store results and band combinations. These temporary layers are only held in memory while they are required for an operation, then discarded. It improves the overall efficiency of the system if the image structure is kept to a minimum size.

The repertory grid tool is the only tool from the traditional suite to be implemented in KAGES and is the only tool not requiring the user to directly manipulate an image. In this case the image merely serves as a prompt. The KAGES user interface therefore reflects the function of the tools in the tool-kit and varies according to the tool being used, adopting Inria's (1991) approach of not using an identical interface tool for all applications. The KAGES tool interfaces have many common features, but with other features unique to particular tools within the tool-kit.

6.11 CONCLUSIONS

The knowledge needed to analyse geographic images is primarily visual, which means a view of both an image and a histogram of the pixel distribution of the image as a whole is required. This view is needed for each of the raw bands in a satellite image set, and for combinations of those bands.

The tools required to capture geographic knowledge need to allow an expert image interpreter to manipulate images via a graphical user interface, and to capture those manipulations at the Primitive, Relationship and Assembly Knowledge levels.

The tool-kit containing these tools needs to be structured, allowing a user to identify scene primitives, then identify higher level relationships between scene primitives of spatial relationships and generalisation.

Not all geographic knowledge is visual however. Algorithms for combining image bands and processing images are not visual. Some heuristics for discriminating objects are not visual. Temporal relationships identified for data in an associated database are not visual. However in the case of eliciting heuristics about geographic phenomena, a visual prompt to users can be useful, if not essential.

The knowledge acquired by using the tool-kit may come from many different sessions with multiple domain experts. This knowledge is stored in knowledge-bases which needs to be resolved. The method chosen for combining knowledge-bases has been to use certainty factors for Primitive and Assembly Knowledge. User intervention is available when there is a conflict the system cannot resolve or the domain expert disagrees with a consolidated rule. In the case of Relationship Knowledge, the system detects conflicts, but then requests user intervention to resolve them.

Verification of the knowledge acquired requires the techniques of both knowledge-based system verification and classification accuracy of remotely sensed data. Hence a combination of inspection, testing and ground truthing was adopted.

The user interface of a geographic system is of primary importance to its usability. An interface for acquiring visual knowledge is however different from one used for acquiring non-visual knowledge. The approach taken was for each tool in the tool-kit to have its own user interface with a similar look and feel linked back to the primary menu screen.

Chapter 7. EVALUATION OF THE KAGES SYSTEM

This chapter describes the process of validating the knowledge acquisition methodology itself and presents the results obtained using that methodology. KAGES is used as the testbench for this exercise. In developing field test strategies three main considerations are taken into account; users, image availability and the availability of ground truthed field data. The results of field testing in the sea ice and crop classification domains is presented.

7.1 INTRODUCTION

KAGES is a system designed to elicit knowledge in accordance with the six level classification system of geographic knowledge proposed in Chapter 2. To achieve this the toolkit has been evaluated in two different domains and demonstrated to a variety of users many of whom work in other domains using either remote sensing technology, GIS or a combination of both.

7.2 METHODOLOGY FOR COMPARING THE KNOWLEDGE ACQUIRED BY THE SYSTEM

7.2.1 General Approach

The KAGES methodology was evaluated using the following five criteria:

1. How well did the expert like the technique?
2. How much knowledge was elicited?
3. How 'good' was the knowledge?
4. Which tool was used?
5. How well did the system work in the different domains?

The first criterion is user centred. Users were interviewed and asked a set of questions about their views of the toolkit and its applicability to their domain. The second and third criteria are system oriented. The results of using the KAGES methodology were compared with the results of using other methodologies. Criterion four is a measure of the effectiveness of the knowledge classification scheme. The fifth criterion was evaluated by applying the methodology to two distinct domains.

7.2.2 The Sea Ice Mapping Domain

Rules generated by the KAGES methodology were compared with rules in the Icemapper system described in Chapter 5. Those rules were obtained by interviewing experts and by searching the literature for appropriate rules and algorithms. There is therefore a direct comparison at the Primitive Knowledge level as both systems have rules designed to do per-pixel level classification. Acquisition of Algorithmic Knowledge can also be directly compared. Beyond that, comparison with Icemapper is difficult since it does not use higher levels of knowledge.

Rules at the Primitive Knowledge level were also compared with the rules obtained using the S statistical package (Chambers and Hastie, 1992). Data was sampled across an image and an expert asked to manually classify the sample points. The results were then used to create a decision tree. The KAGES toolkit contains a tool to sample images, hence, although the S tool generates rules, KAGES can assist it by preparing data sets for it to analyse.

The algorithm used by S attempts to partition the space of predictor variables (in this case pixel values from NOAA AVHRR bands and band combinations) into homogeneous regions to which a classification label can be attached. It begins with the full set of pixel values and partitions them into two sets, such that the members of those two sets are most different. These two sets become new nodes for further partitioning. This continues until the number of members of a set reaches some predetermined minimum size or the members of the set have a low variance.

7.2.3 The Crop Recognition Domain

In this domain, rules generated by KAGES were compared with results obtained using a GIS based statistical classifier. This classifier incorporated a clustering algorithm based on nearest neighbour techniques. It combined the information from five satellite images, three SPOT and two Landsat. It also had available to it an image from February which was not available to the KAGES system.

Before the statistical classifier was applied, the user screened out areas of the image which were not of interest for the crop classification study. These included areas affected by cloud, urban areas, sea and forest. The remaining pixels in the image were then clustered and, using the training data, had labels of *peas, poppies, potatoes, pyrethrum, cereal, pasture, onions* and *beans* applied.

The results of this classification system were compared with the results of applying rules generated by KAGES.

7.3 METHODOLOGY FOR EVALUATING THE USEABILITY OF THE SYSTEM

7.3.1 User Selection

One of the primary requirements for a knowledge acquisition tool which accepts direct user input is for the user interface to be effective and comfortable for the user. The users of the KAGES system are expert image interpreters. They fall into a number of distinct groups:

- GIS users who manipulate satellite images,
- Image interpreters who use data visualisation or other tools and
- Image interpreters who use predominantly manual techniques.

Users of these different types came from a number of different organizations and required different features in a system such as KAGES. The organisations included:

- The Bureau of Meteorology,
- The Tasmanian Institute for Agricultural Research,
- The Tasmanian Fire Service,
- The Tasmanian Department of Primary Industry and Fisheries,
- The University of Tasmania School of Geography and Environmental Studies,
- The University of Tasmania Central Science Laboratory and
- The University of Tasmania School of Computing.

Development was therefore incremental with numerous visits to the various experts to get their feedback on the system. Initially only a subset of the users was used to develop the system. The system was then demonstrated to a wider group of users and their feedback obtained. Finally, the system was taken to a series of overseas venues for comment by a wider group of users.

7.3.2 User Testing procedure

A group of expert users from the areas identified above were asked to comment on the KAGES system, the knowledge classification scheme and the useability of KAGES in relation to systems they were already using. This questioning took the form of a structured interview rather than a questionnaire because of the number of subjects involved and their diverse backgrounds.

The initial questions asked were:

- Is this a useful tool ?
- Which aspects of the tool are most useful ?
- What does it provide which your current system does not provide ?
- What would you like to see it provide ?
- Which aspects of the system duplicate what you already have ?
- What aspects of the system don't you like ?

These were asked as open ended questions. All the users were asked these primary questions, although their answers often led to other discussions. The questions were asked in conjunction with a demonstration of KAGES as it was at that time. As a result, as field testing proceeded, more refined versions were demonstrated.

Each demonstration took about two hours with a follow up interview after the inclusion of new features requested by the user. In some cases this cycle continued for up to five refinement stages.

7.4 RESULTS OF THE EVALUATION

7.4.1 Verification Strategies in Remote Sensing

All systems involving classification of remotely sensed images have difficulties when verification is required. A satellite passes over a region at a particular time and produces a set of images. To verify the results of using automated classifiers on those images, information about the actual feature on the ground needs to be recorded. In many domains this information may also vary over time. As a result there is a need to have a strategy of ground truthing in place.

Both scenarios used for testing the KAGES methodology relied on ground truthing for verification of the results, however in the Antarctic domain there are problems with its collection as discussed below. The scenarios were both dynamically changing systems with the possibility of rapid change over a short time interval. Both were also very different in the way ground truthing had to be conducted.

Ground truthing in Antarctica has many difficulties compared with other domains. The continent is extremely sparsely populated and there are few permanent ground stations. In the case of sea ice studies there are even fewer. As a result most ground truth comes from either airborne surveys, or from shipboard observations.

Shipborne observations are necessarily biased. The ship can only provide a line transect through the ice at a specific time which may or may not coincide with a cloud free satellite image. Secondly the ship sails through thin ice or open water, so observations are mainly going to be of that type. Thicker or older ice can only be ground truthed by surface expeditions or aerial survey.

Observations have been made from the *Aurora Australis*. However these have yet to be correlated with the satellite images obtained during the period of the voyage. This is currently the basis of a separate study at IASOS (Institute of Antarctic and Southern Ocean Studies). For the study described in the thesis, all verification relied on the manual interpretation of the test images by the domain expert.

The crop recognition domain does not present the same verification problems as the Antarctic domain does. However there are still many difficulties. These include:

- the identification of individual paddocks,
- the temporal nature of crops,
- the planting periods of crops and
- the small size of paddocks.

To solve some of these problems a sample of 187 paddocks was taken during the 1996/1997 growing period. 75 of these were used as training data on the Landsat and SPOT images.

Individual paddocks on the ground were identified, digitised using a GIS and assigned identification numbers. The type of crop growing in the paddock was also identified on the ground and the planting date of the crop where available. Although the contents of the paddocks were identified, this was done in a fairly ad hoc way over a period of two months. There was also some concern that there may have been some incorrectly identified paddocks. No notes were taken on weed infestations or crop health.

7.4.2 Antarctic Sea Ice Knowledge

Within the Antarctic sea ice identification domain two images from different time periods (26th February and 4th March 1998) were used as samples for comparison.

Both images were acquired during the Antarctic summer when visible NOAA VHRR images are useful. Although cloud was present in most of the images, an ability to distinguish between cloud and ice was important.

Figure 7.1 is the result of using a sample grid over an image which was manually classified by an image interpretation expert. 769 points were sampled. This image was then used by KAGES to generate rules based on the users interpretation methods. This included a variety of image bands and band combinations. These rules were then applied to the whole of the image. The classified image was then sampled using the same grid as in Figure 7.1 and the results are shown in Figure 7.2.

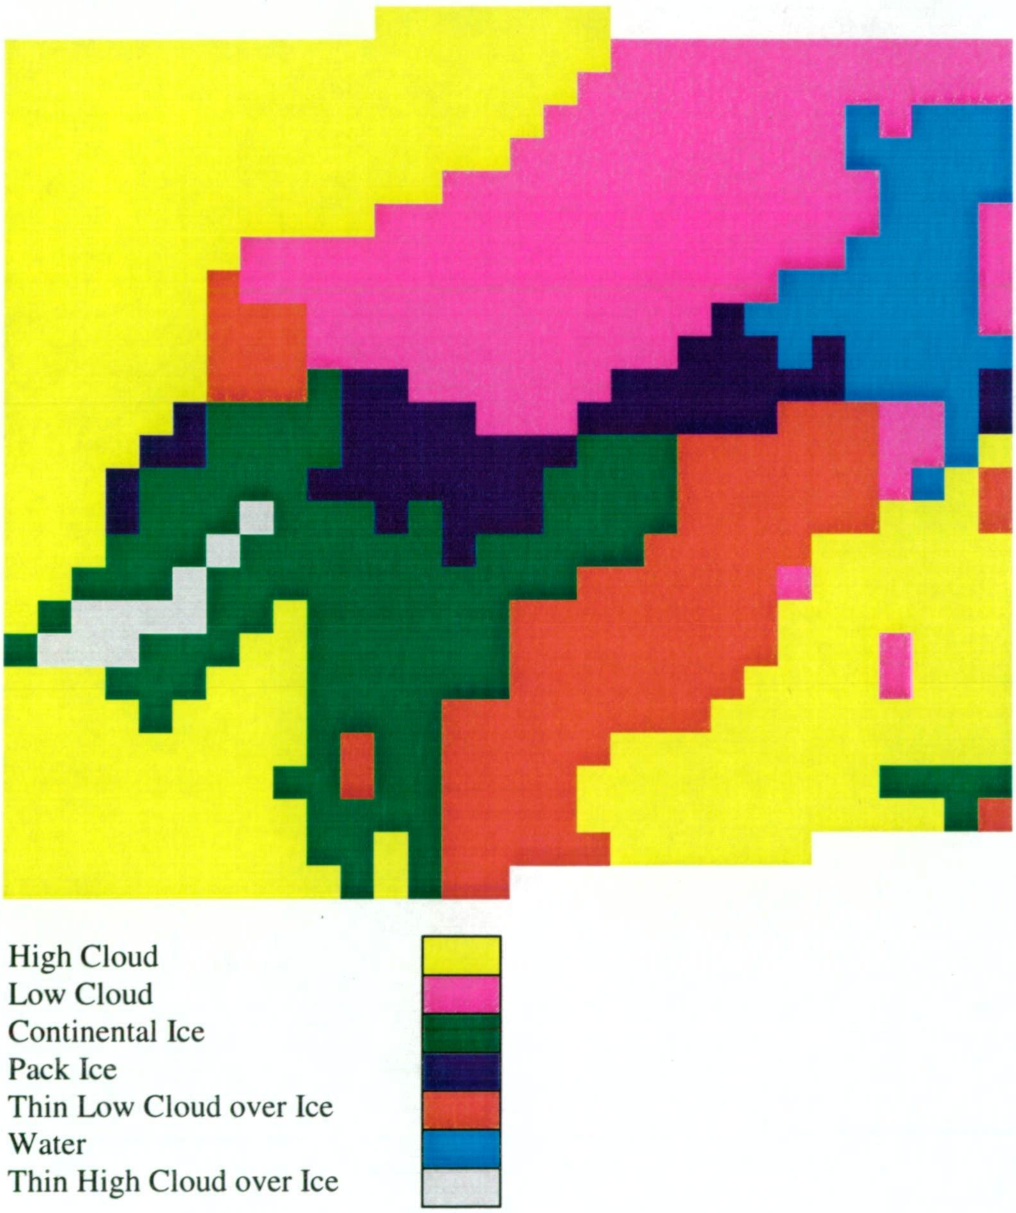

Figure 7.1 Domain expert classification of sample points on the training image

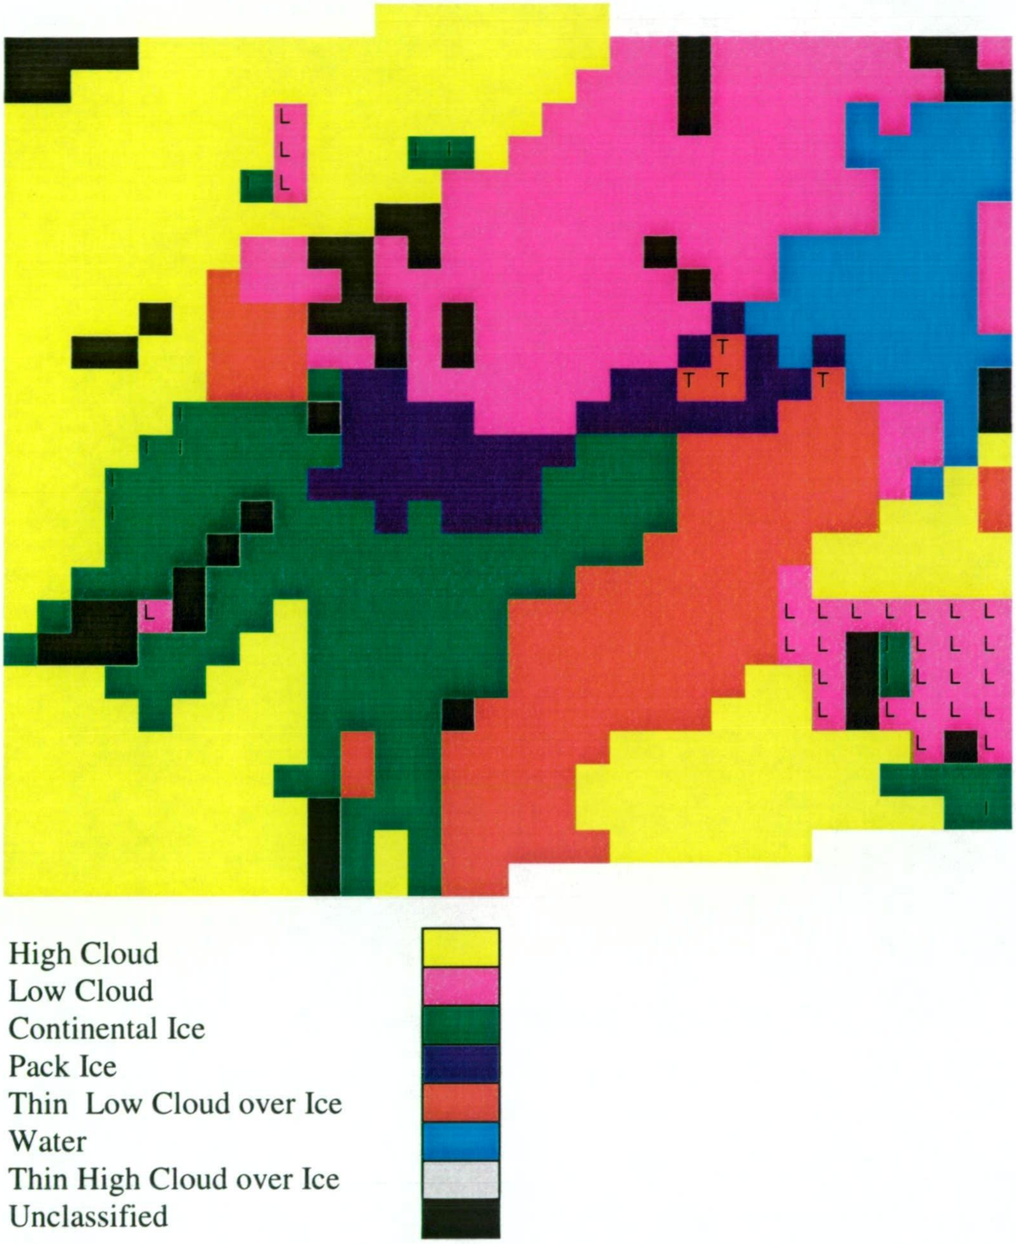

Figure 7.2 KAGES classification of sample points on the training image. Cells containing letters indicate misclassification, L = low cloud, I = continental ice, W = water, T = Thin High Cloud over Ice.

Table 7.1 is a summary of the image classifications produced using KAGES on the NOAA image of 26th February 1998. The rules generated by S, and those in the original Icemapper system were then applied to the training image. The resulting classified images were then sampled using the sampling grid. A comparison of the results of the three classification techniques is summarised in Table 7.2.

					MISCLASSIFICATIONS							
CLASS	ACTUAL	FOUND	MISSED	CORRECT	HC	LC	CI	PI	TLIC	W	THIC	UC
HC	310	267	43	86.1%		26	3					14
LC	150	128	22	85.3%			2					20
CI	113	111	2	98.2%								2
PI	51	40	11	78.4%			5		4			2
TLIC	88	87	1	98.9%			1					
W	47	47	0	100.0%								
THIC	10	0	10	0.0%								10
	769	680	89	88.4%								

Table 7.1 Antarctic classification results using the KAGES knowledge-base on the training image. HC = high cloud, LC = low cloud, CI = continental ice, PI = pack ice, TLIC = thin low cloud over ice, THIC = thin high cloud over ice.

	ACTUAL	KAGES	S	ICEMAPPER
HIGH CLOUD	310	86.1%	52.3%	0.0%
LOW CLOUD	150	85.3%	80.7%	98.7%
CONTINENTAL ICE	113	98.2%	93.8%	53.1%
PACK ICE	51	78.4%	51.0%	51.0%
THIN CLOUD OVER ICE	88	98.9%	38.6%	97.7%
WATER	47	100.0%	100.0%	100.0%
THIN HIGH CLOUD OVER ICE	10	0.0%	80.0%	0.0%
OVERALL	769	88.4%	65.5%	47.7%

Table 7.2 Comparison of Icemapper, KAGES and S performance on the training image

Since KAGES was trained on this image one would expect its classification accuracy to be high and this was indeed the case. The S rules misclassified high cloud as thin high cloud over ice in 78% of the high cloud error cases and pack ice as water in 88% of the pack ice error cases, accounting for 52% of all error cases. KAGES better performance was due to it being used to specifically generate rules on these features.

The original Icemapper rules had similar problems to S, compounded by completely failing to classify high cloud, the majority of which came out as low cloud and thin low cloud over ice.

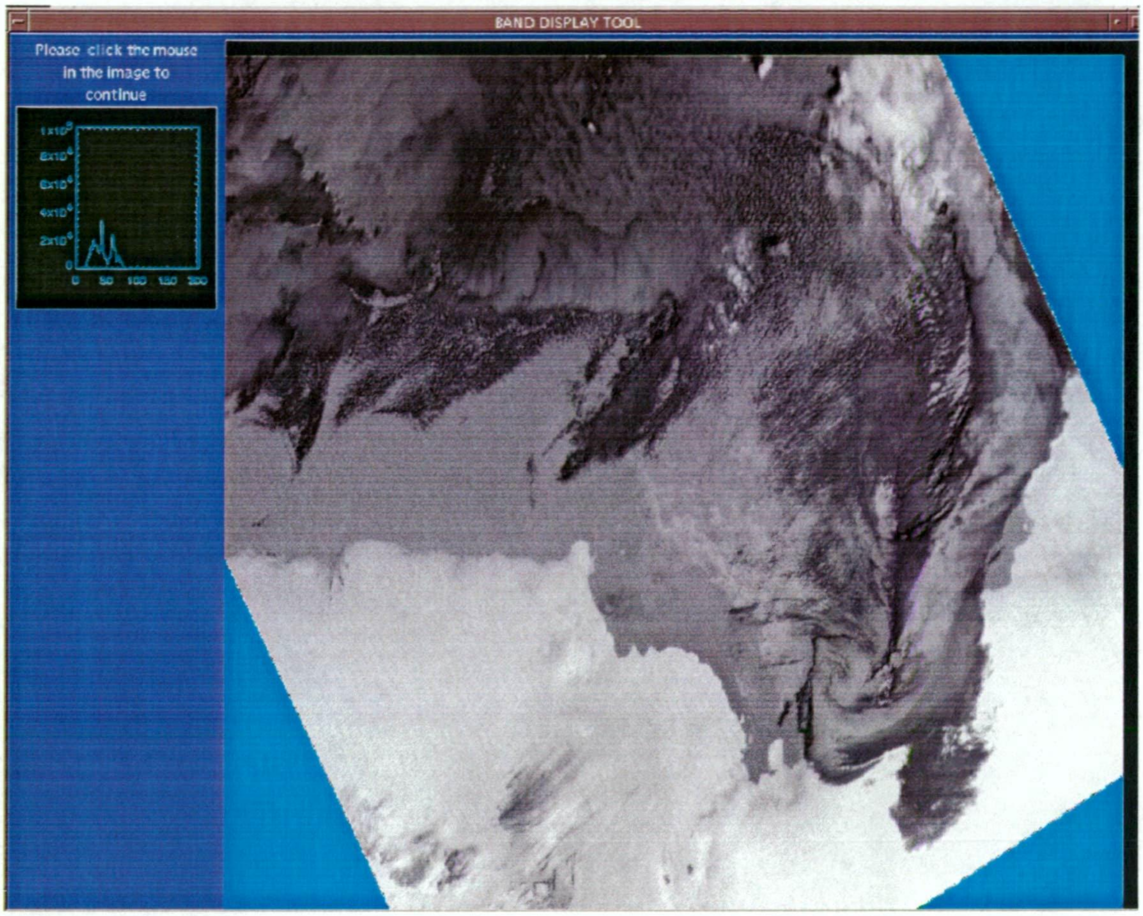

Figure 7.3 Thermal band 3 of the test image, derived from a NOAA AVHRR image of Mawson/Davis taken on 4th March 1998 viewed using the KAGES band display tool

The second Antarctic image (Figure 7.3) was used to test the KAGES generated rules and compare their performance with the other two rule sets. The resulting classified image is shown in Figure 7.4. The domain expert's classification of the image sampled using a regular grid of 634 points at the same spacing as the training image is shown in Figure 7.5. The KAGES classification results in terms of the sampling grid are shown in Figure 7.6. The major misclassification is high cloud and thin high cloud over ice being classified as low cloud.

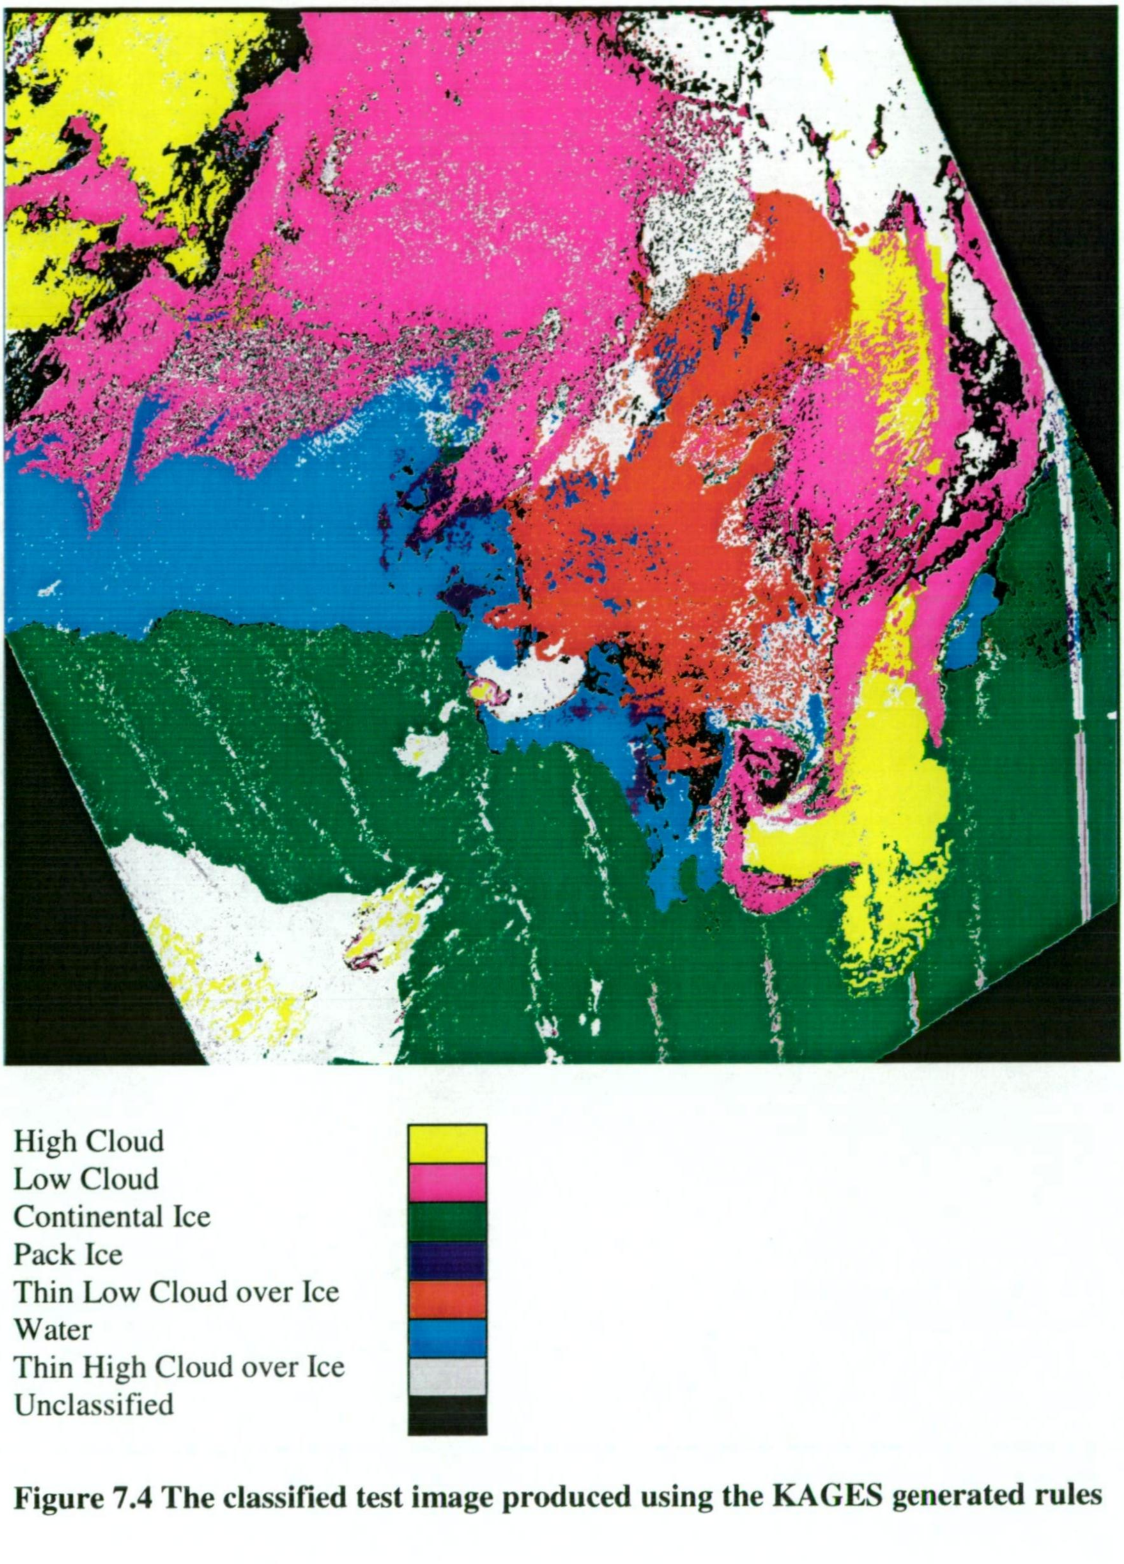

Figure 7.4 The classified test image produced using the KAGES generated rules

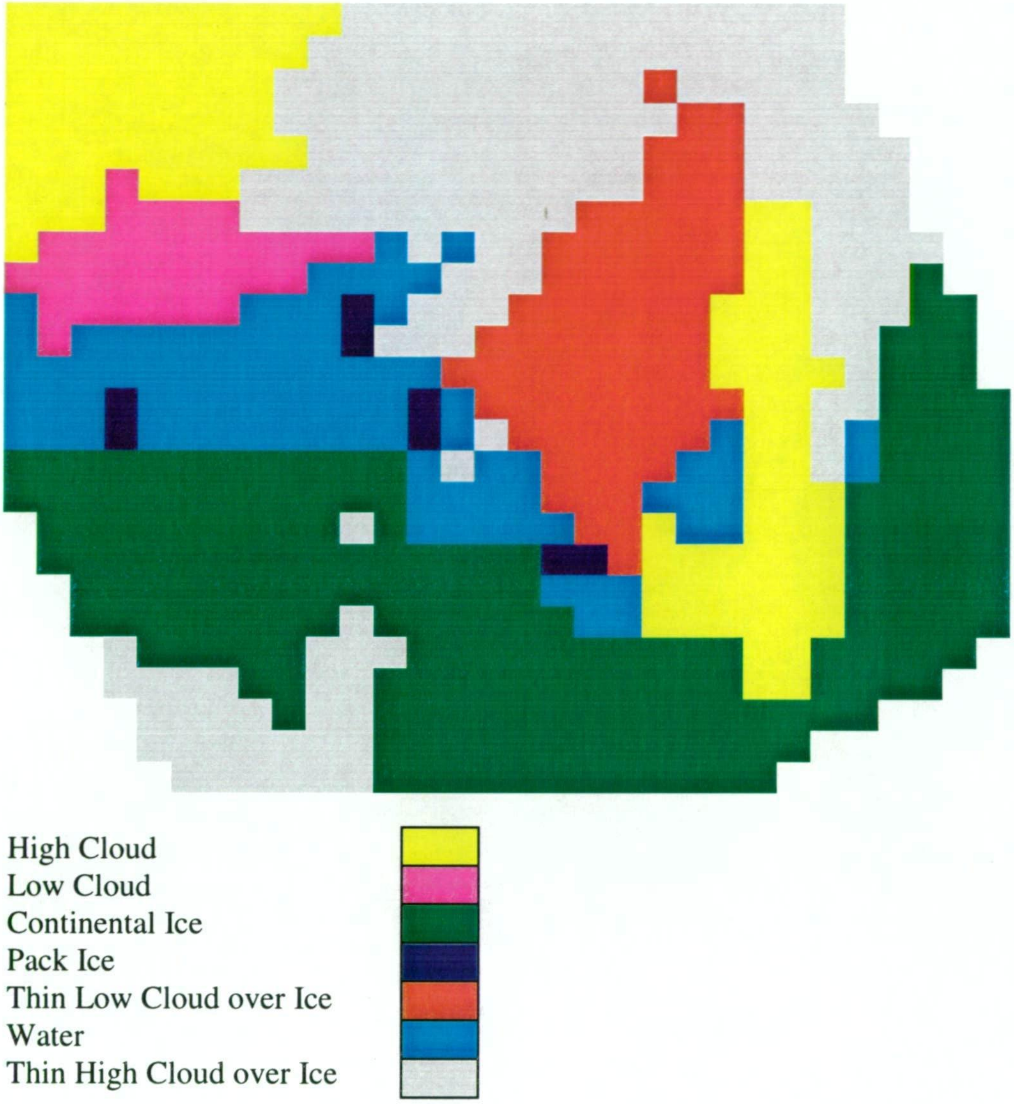

Figure 7.5 Domain expert classification of sample points on the test image

					MISCLASSIFICATIONS							
CLASS	ACTUAL	FOUND	MISSED	CORRECT	HC	LC	CI	PI	TLIC	W	THIC	UC
HC	105	74	31	70%		23	4					4
LC	31	30	1	97%								1
CI	167	167	0	100%								
PI	8	6	2	75%						2		
TLIC	72	72		100%								
W	84	83	1	99%			1					
THIC	167	79	88	47%		88						
TOTAL	634	512	122	81%								

Table 7.3 Antarctic classification results using the KAGES knowledge-base on the test image. HC = high cloud, LC = low cloud, CI = continental ice, PI = pack ice, TLIC = thin low cloud over ice, THIC = thin high cloud over ice.

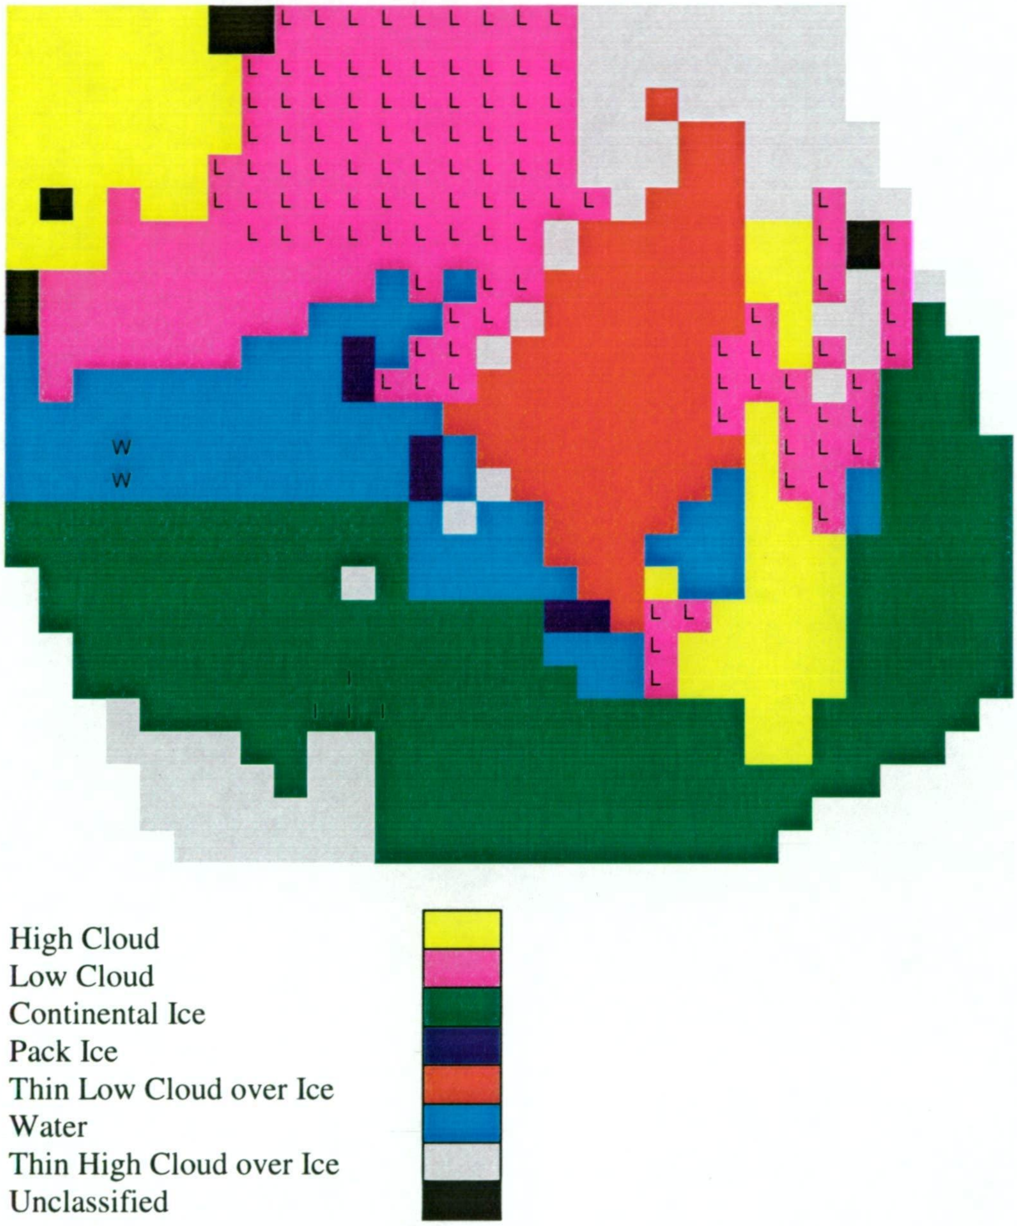

Figure 7.6 KAGES classification of sample points on the test image. Cells containing letters indicate misclassification, L = low cloud, I = continental ice, W = water, P = pack ice.

A summary of the performance of KAGES rules and an analysis of misclassifications is shown in Table 7.3. Classified images were also generated using the S and Icemapper rules. These classified images were sampled using the grid and the results are shown in Table 7.4

	SAMPLE	KAGES	S	ICEMAPPER
HIGH CLOUD	105	70%	47%	0%
LOW CLOUD	31	97%	90%	90%
CONTINENTAL ICE	167	100%	98%	94%
PACK ICE	8	75%	100%	88%
THIN CLOUD OVER ICE	72	100%	71%	39%
WATER	84	100%	95%	99%
THIN HIGH CLOUD OVER ICE	167	47%	65%	5%
OVERALL	634	81%	77%	49%

Table 7.4 Comparison of Icemapper, KAGES and S performance on the test image

The results of classifying the second image with the S rules and with KAGES generated rules are very similar. High cloud still presents problems with the S rules with 96% of high cloud misclassification being as low cloud. KAGES rules performed better in high cloud identification. The KAGES rules misclassify thin high cloud over ice as low cloud in all error cases of that type. S rules performed better in thin high cloud over ice identification.

The Icemapper rules also have problems with high level cloud. All of the high level cloud was missed with 95% of it being classified as low cloud. Thin high cloud over ice was misclassified as low cloud in 75% of this classes error cases. This may be due to Icemapper originally having rules which placed cloud of all types in a single class. Although later rules were added to correct this, low cloud was kept as the default cloud class.

The testing reported above was performed using per-pixel reasoning because both Icemapper and S rules are of the Primitive Knowledge type. Relationship Knowledge is not used in Icemapper or S. KAGES on the other hand was used to generate this type of knowledge about the relationships of shadow to cloud banks and the relationship of land to water to delineate continental ice and sea ice.

KAGES appears to have improved on the original pixel level knowledge acquired by interview and coded into Icemapper. It has the added capability of generating extra rules about spatial relationships which can be used to refine classification. The relationship with the S package is more complex since KAGES can be used to generate the training set needed by S including any Algorithmic Knowledge necessary for band combinations. The performance of S and KAGES generated rules appears to be very similar on the images used.

The main advantages of the KAGES methodology over the traditional knowledge acquisition techniques evaluated in this study were:

- Knowledge was acquired more quickly. Traditional methods require acquisition, encoding, testing and user verification of rules, usually over several interviews with the domain expert. KAGES combines these steps for rapid development and feedback.
- Knowledge requiring algorithmic manipulation of band combinations was easier to acquire, enter and verify. Icemapper uses band combinations, but these require hard coding and an image representation is not produced for the user.
- Knowledger could be validated by previewing individual rules. Individual rule clauses can be inspected and modified interactively. The results of modifications can also be seen immediately. Validation and modification of Icemapper rules, on the other hand require significant programmer involvement.
- The rules generated by KAGES when applied to an Antarctic image, give classifications which are comparable with those generated by S and better than those acquired by interview.

7.4.3 Crop Recognition Knowledge

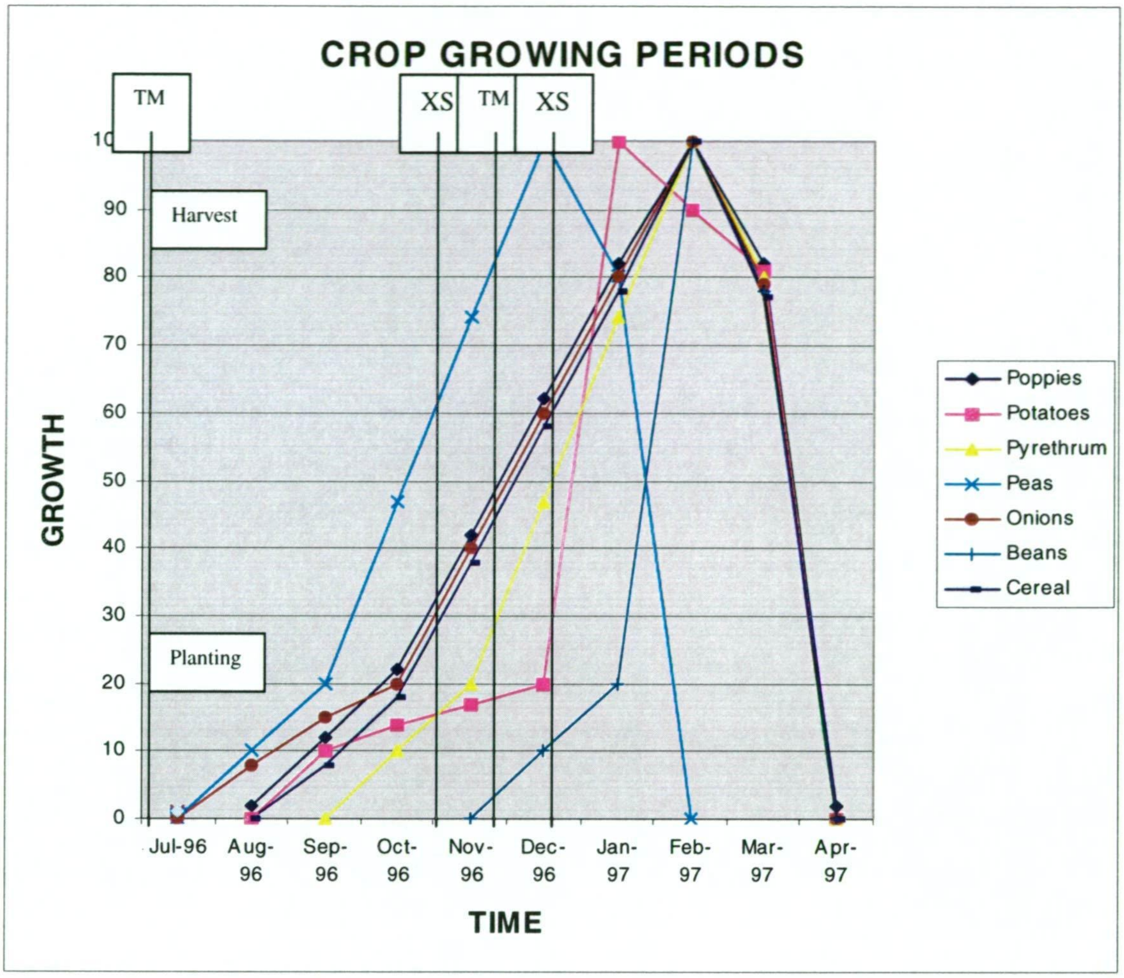

Figure 7.7 Growth and cropping times for crops under study showing times of useable satellite passes. TM is Landsat, XS is SPOT.

Significant problems were experienced with image acquisition in the crop recognition domain. The 1996/97 growing season was particularly cloudy and there were many very cloudy days on satellite passes. As a result, only four images, two Landsat (6^{th} July 1996 and 27^{th} November 1996) and two SPOT (2^{nd} November 1996 and 24^{th} December), were useable. The timing of these are shown superimposed on the growth / harvest chart of various crops in Figure 7.7. All of the four images had potential problems in correlation with ground truth.

The 6^{th} July 1996 Landsat image provided a winter view that included fallow areas and winter crops such as brassicas. Since ground truthing did not commence until November 1996, this image was of limited use. It was used to delimit forest areas which remained more or less constant during the 10 months of the study.

The 2nd November 1996 SPOT image was at the early stage of the growth cycle for most crops. Depending on the planting time, many crops had not developed a large ground cover. As a result some paddocks appear as bare soil or are influenced by weeds growth.

The 27th November 1996 Landsat image was more useful, as many crops were established by this stage although late potatoes were still being planted. Hence some of the problems associated with the November SPOT image were overcome.

The 24th December SPOT image was acquired during the middle of the growth cycle for most crop types, although some crops such as poppies still had low ground cover. Peas were being harvested at this stage and could have been removed before the image was produced, again causing discrepancy with ground truth.

Six crop types; poppies, pyrethrum, peas, potatoes, onions and beans were used for classification testing. Several other agricultural land cover features were also sampled (including cereals, brassicas and pastures) but the numbers involved were small. Also land cover such as forest was detected and included for screening purposes.

Only Primitive and Algorithmic Knowledge was used in this particular study because:

- There were no agricultural image interpretation experts available so KAGES was used as an image analysis system using the training sample to discover rules
- Spatial relationships were not obvious and without an expert were unobtainable
- Although Assembly Knowledge about urban areas could have been obtained, it was not required for the objectives of the crop identification project

Table 7.5 gives a summary of the classification performance of KAGES in comparison to ground truth. The numbers in the table refer to the number of paddocks classified or misclassified. The label at the beginning of each row represents the crop label assigned by the KAGES knowledge-base.

					MISCLASSIFICATIONS					
CLASS	**ACTUAL**	**FOUND**	**MISSED**	**CORRECT**	**B**	**ON**	**P**	**POT**	**POP**	**PY**
B	13	6	7	46%		3	1		3	
ON	19	12	7	68%	1		4			2
P	20	16	4	80%	1	2				1
POT	26	17	9	65%	1	1	1		5	1
POP	38	18	20	47%	2	1	9	8		
PY	13	6	7	46%	1		3	1	2	
TOTAL	129	75	54	58%						

Table 7.5 Crop Classification results using the KAGES knowledge-base. B= beans, ON = onions, P = peas, POT = potatoes, POP = poppies, PY = pyrethrum

The best performance was for peas (80%) which were classified using a Landsat band combination algorithm:

abs(Band_5 – Band_7)

From Figure 7.1 it can be seen that peas are the first crop to be harvested in the area and the Landsat image (27^{th} November 1996) used for identification was taken at the beginning of the harvest season when peas would have been at their maximum ground cover. This problem with satellite passes and identification may also account for the performance with beans, which had only just begun their growth cycle and may have been affected by weeds.

The worst performance was with poppies where there were problems distinguishing them from potatoes and peas. Errors in the classification of potatoes and poppies seemed systematic, with both showing similar levels of misclassification. One reason for this may be that both of these crops were both at early stages of their growth cycle with large areas of bare soil still visible in the paddock.

The reason for the high number of misclassifications of poppies as peas was the following heuristic:

IF classification of pixel = peas
AND classification of pixel = poppies
THEN classification = peas

Which was provided by KAGES based on the training data provided for it. This rule was subsequently proved incorrect. It may have been produced as a result of bare earth appearing in paddocks where pea crops, identified within the paddocks by the ground truthing process, had been harvested just prior to the satellite pass.

By comparison a statistical classifier using Principle Components Analysis achieved an overall accuracy of 59.4%. The statistical classifier had, in its training set, an extra image from late February which had not been made available to KAGES.

7.4.4 Expert User Acceptance

The KAGES toolkit was developed using prototypes which were gradually modified to meet user requirements. It should be noted that where users required a specific feature to be added, or where there was a feature they did not like, action was taken to rectify the problem. Table 7.6 is a summary of the answers provided by the various users to the open-ended questions described in section 7.3.2. Appendix C contains transcripts of some of the interview sessions.

The four users who were expert in the use of GIS all commented that the Per Pixel tool for acquisition of Primitive Knowledge was no real advance over existing GIS packages which had spectral signature manipulation implemented. However all agreed that unless objects on an image had already been assigned labels using a GIS, and the labeled segmented image passed to KAGES, such a tool was essential. In spite of this negative comment, when the Per-Pixel tool was viewed as independent of a GIS, its implementation with both a histogram and image display was well received. One user would have liked contrast stretching included.

Type Of User	Tool Assisted Image Interpretation	GIS User	GIS User	GIS & Data Visualisation Image Interpretation
Is this a useful tool?	Yes	Yes	Yes	Yes
Which aspects of the tool are most useful ?	Pointing, Point data gathering.	Spatial relationships	Spatial relationships, Heuristics.	Heuristics, Spatial relationships.
What does the tool provide that your current systems does not?	Outlining Spatial relationships	Spatial relationships, Grouping, Heuristics.	Spatial relationship, Heuristics Grouping	Heuristics, Spatial relationships.
What would you like to see the tool provide that it does not?	Easier band changes.	Histogram, Contrast stretching.	Histogram.	Histogram contrast stretching.
What aspects of the system duplicate what you already have?	X	Pixel level classification.	Pixel level classification.	Pixel level classification.
What aspects of the tool don't you like?	X	X	X	Inability to define primitives from histogram.

Type Of User	Image Interpreter (manual)	GIS User	Data Visualisation Image Interpretation
Is this a useful tool?	Yes	Yes	Yes
Which aspects of the tool are most useful ?	Pointing, Outlining Composites.	Spatial relationship, Heuristics, Grouping.	Pointing, Heuristics, Outlining.
What does the tool provide that your current systems does not?	N/A	Grouping, Heuristics.	Outlining, Spatial relationships, Heuristics.
What would you like to see the tool provide that it does not?	X	Spatial concentration (Assembly).	X
What aspects of the system duplicate what you already have?	N/A	Pixel level classification.	X
What aspects of the tool don't you like?	X	Rule editor a bit simple.	X

Table 7.6 Summary of user responses to KAGES

Users who were not familiar with GIS packages (including the manual interpreter) liked the operation of the Per Pixel tool. The ability to quickly select bands, point to an object and adjust thresholds until a feature was defined, then apply that to the whole image was mentioned as being a useful feature.

The most positive comments came about the Relationship, Assembly and Heuristic tools which, according to the GIS users interviewed, had no equivalent in current GIS packages. Different users suggested a range of enhancements to these tools, many of which were implemented.

The heuristic tool based on repertory grid analysis was seen as innovative, since much knowledge held by domain experts is not in fact visual and some means other than interview was needed to elicit that knowledge. Some users found the elicitation of classes and attributes a bit time consuming and said they would prefer a quick entry mechanism for adding extra classes and attributes to save going through the triad comparison each time. They liked the ranking system using slider bars and indicated that they found it quick and easy to use.

All saw significant potential in the Relationship Tool to determine spatial relationships, especially if coupled to a temporal database with which changes in these relationships could be investigated. This tool was seen as an advance over existing GIS technology and something lacking in existing geographic expert systems which still concentrate on pixel-by-pixel processing. A suggested enhancement was a facility for calculating the degree of overlap where two areal objects shared pixels.

The Region of Interest tool for acquisition of Assembly Knowledge was seen as a good way of extracting generalisation knowledge. One GIS user suggested that more knowledge could be collected by this tool if it could calculate areas and concentrations of objects in an area selected by this tool. This suggestion was implemented.

In terms of the overall operation of the tool, all users liked the layout of the user interface and the linking of the various user windows. An initial comment from GIS users was that a histogram of the image should be shown as well as the image, and a histogram of pixels in an object be shown when an object was delineated. Both of these were subsequently added to the user interface.

In summary, the positive aspects of the final user interface were that it was seen to be intuitive and easy to use, consistent in the rules it produced and minimised the amount of typing required in the knowledge acquisition process.

The main negative aspects were that it duplicated some of the features already available in GIS and that the rule editor attached to several of the tools lacked sophistication.

7.5 CONCLUSIONS

The geographic knowledge classification scheme which formed the basis of the philosophy behind each of the tools in the KAGES toolkit was validated by user testing. The result of using the toolkit on the two test scenarios was to produce rules comparable with alternative methods but in a more efficient manner. In the case of crop recognition, where rules were applied to produce a classified image, the results were comparable with those of a statistical classifier. In the Antarctic domain the results of applying KAGES generated rules were as good as results of applying two alternative knowledge bases.

According to GIS users, Primitive Knowledge acquisition using the KAGES methodology provided no major advance over the statistical spectral classifiers available in GIS. Users not familiar with GIS were pleased with the tool and found it easy to use. The inclusion of a *Primitive Knowledge Acquisition Tool* is necessary to provide the basic spatial objects needed for acquiring subsequent higher level spatial knowledge in situations where a GIS classification and segmentation process has not been performed, or where the knowledge is to be used without reference to a GIS. It is also a useful tool to allow a user to point out important primitive features.

The *Spatial Relationship Tool* for Relationship Knowledge acquisition and the *Region of Interest Tool* for Assembly Knowledge acquisition were particularly well received by users. According to those surveyed, there was no equivalent tool in the systems they were familiar with. Although enhancements were suggested, these tools were regarded as a real advance in knowledge acquisition from satellite images.

The non-visual tools (the Band Calculator and Repertory Grid) were also well received. The Band Calculator was seen as a quick and easy way to enter algorithms and get an instant result. The Repertory Grid Tool was regarded as good way to acquire knowledge which was non-visual. Even though users agreed that most geographic knowledge is visual, there is still a significant amount of non-visual knowledge needed to analyse satellite images.

Chapter 8. CONCLUSIONS

This chapter presents the conclusions about the KAGES (Knowledge Acquisition for Geographic Expert Systems) knowledge acquisition methodology and toolkit. It presents the argument that the knowledge classification for geographic systems presented in Chapter 2 is valid. It draws conclusions about the knowledge acquisition tools described in Chapter 6 for acquiring the six types of knowledge identified in this thesis and discusses the advantages of using an interactive system. The chapter concludes with a statement of projected future research.

8.1 INTRODUCTION

Conclusions about the KAGES methodology fall into a number of categories. These include the nature of geographic knowledge, the need for a geographic knowledge acquisition methodology analogous to the more general KADS methodology and the usefulness of a series of tools developed to implement the methodology.

8.2 AN INTEGRATED GEOGRAPHIC KNOWLEDGE ACQUISITION STRATEGY

Without an integrated knowledge acquisition strategy, the development of expert systems for use with GIS (Geographic Information Systems) and RSS (Remote Sensing Systems) is constrained. The knowledge acquisition bottleneck becomes tighter because there is no systematic way of collecting knowledge about particular aspects of an image. To acquire knowledge directly from an image interpretation expert requires tools which capture knowledge as the expert works. Without tools of this type a knowledge engineer is reduced to using interview and other off-line techniques which require interpretation of the users actions. Working with these techniques one tends to identify individual image objects, and seldom move beyond that.

One of the aims of Schreiber et al's (1993) KADS (Knowledge Acquisition and Design System) methodology was to present a generalised knowledge acquisition model that details various types and uses of knowledge. The KAGES methodology presents a complementary model for use with geographic or spatial systems. Without an integrated structuring of geographic knowledge, the production of spatial knowledge acquisition tools will be an ad hoc process.

The KAGES methodology assumes that geographic and spatial knowledge is primarily visual, a view supported by McKeown et al (1989). The implementation of tools for such a methodology must also be visual and capture users' knowledge as they analyse an image. Image interpreters initially identify areal, point and line objects either by their appearance or by their spectral signature, and then indicate them by pointing at them. This process includes knowledge of scene primitives or *Primitive Knowledge.*

Information on proximity, orientation and overlap is *Relationship Knowledge.* Users can name the relationship between two objects by pointing at the two objects. The relationship is calculated and although two users may provide a different label for the relationship, its definition rather than its name is definitive.

Grouping objects to replace them with a composite or generalised object is also a visual process with a user drawing around primitives which should be grouped. *Assembly Knowledge* therefore involves knowledge of groups of objects and how they make up a generalised object.

These three tasks of an image interpreter involve the first three knowledge types, Primitive, Relationship and Assembly. They are visual and should be acquired in a visual way. Although these types of knowledge can be used to classify an image, further refinement requires *Non-Visual Knowledge.*

Many experts rely on composite images to make features more visible. The Band Calculator Tool provides an easy method to obtain, name and store *Algorithmic Knowledge* for this purpose.

Scenes can change over time in a systematic way. Some of these changes may be described verbally by the image interpretation expert. Other ways of acquiring *Temporal Knowledge* include analysing the contents of a database to discover patterns.

Experts may also rely on non-visual heuristics. This *Heuristic Knowledge* can be acquired by interview, but a more systematic method is to use tools such as repertory grids where a user is asked to differentiate between objects.

The remaining two knowledge types are at a higher level. *Consolidation Knowledge* (knowledge of how knowledge-bases are to be combined) is at the *Inference Level* in the KADS methodology. *Interpretation Knowledge* (knowledge of how to apply a knowledge-base to a specific domain) is at the KADS *Task Level.*

8.3 AN INTEGRATED TOOLKIT TO IMPLEMENT THE KNOWLEDGE ACQUISITION STRATEGY

The KAGES methodology provides a framework for the classification of geographic knowledge which has been proven sound by the implementation of tools to capture knowledge of each class. This linking of a knowledge classification methodology and a toolkit containing tools which capture users actions, provides an integrated geographic knowledge acquisition system.

The toolkit is integrated to allow a user to choose the tool most suitable for solving a particular knowledge acquisition problem. In the case of a toolkit like KAGES some tools are required before other tools can be used. For example the visual tools for Relationship and Assembly Knowledge acquisition are dependent on the Per-Pixel Primitive

Knowledge acquisition tool to define scene primitives. The Per-Pixel tool, in turn, may require algorithms entered via the Band Calculator to produce composite images.

Only the Heuristic Tool can operate more or less independently of the others since this only uses an image as a prompt to the user. A normal sequence of operations would still see this tool being used after the visual tools to confirm or refine classification rules.

The final level of integration is with the Consolidation Tool, which combines any multiple knowledge-bases created by the other tools in the toolkit. The result of using this tool is the production of an integrated unambiguous knowledge-base.

8.4 USER-CENTRED GEOGRAPHIC KNOWLEDGE ACQUISITION

This thesis has postulated that geographic knowledge is primarily visual and therefore tools designed to acquire geographic knowledge must also be visual. As a result the user interface is of primary importance for the acceptance of such a tool. The tools available within KAGES have been shown to be intuitive providing very rapid feedback on user actions.

This thesis is concerned with the acquisition of geographic knowledge rather than the development of a full inference engine to apply that knowledge. However it is necessary for a user to see the effect of the rules that the system has generated.

The users' reaction to knowledge acquisition using KAGES has been positive. Most of the features in KAGES were not available in packages they were currently using. Although the Per-Pixel Tool did not provide any significant advances over features available within current GIS, users seemed impressed with the other tools. However, unless segmented images classified at the pixel level are imported, users agreed there is a need for the Per-Pixel Tool.

The overall user interface was seen as being easy to use. The operation was similar to many GIS and image processing packages. The manipulation of images was facilitated in an intuitive way with user suggestions being incorporated into the final design.

8.5 FUTURE WORK

The research work reported in this thesis could be extended in at least four significant directions. Three of them are related to non- visual knowledge and have the potential for a long-term study.

8.5.1 Temporal Knowledge

The basis of an investigation into temporal aspects of knowledge acquisition has already been established in the form of a database for use with the MIRC (Multi-temporal Imaging for Remote sensing of Crops) study. The initial study for crop recognition has now been expanded. A new study area to the east of the one reported in this thesis has been chosen along with a second area around the Cole River in Southern Tasmania.

This aspect would require an extension to the knowledge acquisition methodology which would allow the temporal aspects of Allen intervals to be incorporated and would include a study of state change of objects between images (Story and Worboys, 1995). For example some crops are usually replanted with a particular following crop. Some crops are perennials such as pyrethrum.

8.5.2 Interpretation Knowledge

Interpretation Knowledge is knowledge of how to apply a knowledge-base to produce a classified image in a particular situation. It is analogous to KADS task knowledge and as such it is meta-knowledge. Current geographic knowledge-bases have been built with a specific classification task in mind. Using the same knowledge-base for different tasks requires meta-knowledge about which databases are suitable for what application.

An extension to this work is the development of an equivalent to the KADS strategic knowledge. In this case it is knowledge of choosing alternative ways of attaining the same goal. Currently there are too few geographic knowledge-bases for this to be of great concern. However given the growth of use of GIS and RSS and the need for expert systems technology to aid in the interpretation of data this is likely to change. Knowledge of the existence and purpose of knowledge-bases and the means of combining knowledge-bases created for different scenarios to produce a new classifier will become necessary.

8.5.3 Machine Learning Integration

Machine learning and neural network classifiers suffer from the same limitations as statistical classifiers in that they produce classified images based on Primitive Knowledge. Fischer (1998) comments that neural networks have the potential to improve spatial data analysis tasks. However it is also noted that to date, the potential has not been realised.

Integration of machine-based methods with knowledge gained directly from the expert user may be a solution to the problem. This approach has been investigated in non-spatial domains (Faure et al, 1993) but not in the spatial domain. There is a need to investigate the use of objects labeled by machine learning and neural network methods and apply spatial knowledge of the Relationship and Assembly levels to refine the analysis of images. KAGES is already capable of supplying training data through its Point Data Tool. The next stage would be a more direct coupling with machine learning or neural network systems.

8.5.4 Data Mining

One aspect of spatial knowledge acquisition not fully explored in this thesis has been that of data mining. This area is distinct from temporal analysis in that non-temporal relationships can also be found. Data mining is essentially a database analysis technique. Spatial data mining is the extraction of implicit knowledge, spatial relations or other patterns not explicitly stored in a spatial database (Koperski and Han, 1995). Techniques which can be used to accomplish data mining include artificial neural nets, decision tree induction and nearest neighbour methods (Mejia-Lavelleand and Rodriguez-Ortiz, 1998).

During this study a database was established to record crop location and growing characteristics for the MIRC project. Insufficient data had been entered as yet to allow data mining to be investigated, but future expansion of this database will open up the possibility of data mining in the future.

Chapter 9. REFERENCES

Agnew, N. M., Ford, K. M. and Hayes, P. J., 1994, Expertise in Context: Personally Constructed, Socially Selected, and Reality-relevant?, *International Journal of Expert Systems*, Vol. 7, No.1, pp. 65 - 88.

Antony, R. and Emmerman, P.J., 1986, Spatial Reasoning and Knowledge Representation, in *Opitz (ed.), GIS in Government,Hampton, VA.* , pp. 795 - 813.

Arciszewski, T. and Rossman, L.A., 1992, *Knowledge Acquisition in Civil Engineering*, American Society of Engineers, New York.

Armstrong, M. P, 1991, Knowledge Classification and Organization, in Buttenfield B.P. and McMaster, R.B., Map Generalization: Making Rules for Knowledge Representation, Longman, Ch. 5, pp.86 - 102.

Avouris, N. M. and Finotti, S., 1993, User Interface Design to Expert Systems Based on Hierarchical Spatial Representations, *Expert Systems With Applications*, Vol. 6, pp.109-188.

Barrett, A. R. and Edwards, J. S., 1995, Knowledge Elicitation and Knowledge Representation In A Large Domain With Multiple Experts, *Expert Systems With Applications*, Vol. 8, No.1, pp. 169 - 176

Bartelme, N., 1996, Data Structures and Data Models in GIS, in *Mussio, L., Forlani, G. and Crosilla, F. (eds.), Data Acquistion and Analysis for Multimedia GIS, Springer Wien, New York,* pp.359 - 376.

Beard, M. K., 1991, Constraints on Rule Formation, in *Buttenfield B.P. and McMaster, R.B.(eds), Map Generalization: Making Rules for Knowledge Representation, Longman,* Ch. 7, pp.121 - 135.

Bonnet, A., Haton, J-P., and Truong-Ngoc, J-M., 1988, *Expert Systems : Principles and Practice*, Prentice Hall, New York.

Boose, J. H., 1990, Uses of Repertory Grid Centred Knowledge Acquisition Tools for Knowledge-Based Systems in *Boose, J. H. and Gaines, B.R.(eds.), Data Acquistion and Analysis for Multimedia GIS, Springer Wien, New York*, pp. 61 – 83.

Boose, J. H. and Bradshaw, J., M., 1987, AQUINAS: A Knowledge Acquisition Workbench for Building Knowledge-Based Systems, *Proceedings of the First European Knowledge Acquisition Workshop,* A6, Reading UK.

Boose, J. H. and Gaines, B.R. ,1990, *Data Acquistion and Analysis for Multimedia GIS*, Springer Wien, New York.

Boy, G., 1996, The Group Elicitation Methods: An Introduction, *Proceedings of the Tenth European Knowledge Acquisition Workshop, Lecture Notes in A.I. 1076, Advances in Knowledge Acquisition*, Nottingham, pp. 290 – 305.

Breuker, J. and Wielinga, B., 1987, Knowledge Acquisition as Modeling Expertise: the KADS Methodology, *Proceedings of the First European Knowledge Acquisition Workshop,* Reading, B1.

Breunig, M., 1996, *Integration of Spatial Information for Geo-Information Systems: Lecture Notes in Earth Sciences 61*, Springer, Berlin.

Burrough, P. A., 1990, *Principles of Geographic Information Systems for Land Resource Assessment,* Oxford University Press.

Buiten, H. J. and Clevers, J. G. P. W. (eds), 1993, *Land Observation by Remote Sensing: Theory and Applications,* Gordon and Breach.

Burton, A. M., Shadbolt, N. R. Hedgecock, A. P. and Rugg, G., 1987, A Formal Evaluation of Knowledge Elicitation Techniques for Expert Systems: Domain 1, *Proceedings of the First European Knowledge Acquisition Workshop,* Reading, D3.

Buttenfield B. P. and McMaster, R.B. (eds), 1991, *Map Generalization: Making Rules for Knowledge Representation,* Longman.

Buttner, G., Hajos, T. and Korandi, M., 1989, Improvements to the Effectiveness of Supervised Training Procedures, *International Journal of Remote Sensing*, Vol. 10, No. 6, pp. 1005 - 1013.

Carpenter, G. A., Gjaja, M. N. Gopal, S. and Woodcock, C.E., 1997, ART Neural Networks for Remote Sensing: Vegetation Classification from Landsat TM and Terrain Data, *IEEE Transactions on Geoscience and Remote Sensing,* Vol. 35, No. 2, pp. 308 - 325.

Chambers, J. M. and Hastie, T. J., 1992, *Statistical Models in S,* Wadsworth and Brooks, Ch. 9.

Chang, S. and Jungert, E., 1996, *Symbolic Projection for Image Information Retrieval and Spatial Reasoning,* Academic Press, London.

Chao C. and Salvendy G., 1994, Percent Of Procedural Knowledge Acquired As A Function Of Number Of Experts From Who Knowledge is Acquired for Diagnosis, Debugging and Interpretation Tasks, *International Journal of Human-Computer Interaction*, Vol. 6, No. 3, pp. 221 - 234

Cheng, P. C-H, 1996, Diagrammatic Knowledge Acquisition: Elicitation, Analysis and Issues, *Proceedings of the Tenth European Knowledge Acquisition Workshop, Lecture Notes in Artificial Intelligence 1076, Advances in Knowledge Acquisition,* Nottingham, pp.179 - 194.

Chuinsrs, S., Blasco, F., Bellan, M-F and Kergoat, L., 1997, A Poppy Survey Using High Resolution Remote Sensing Data, *International Journal of Remote Sensing*, Vol. 18, No. 2, pp. 393 - 40.

Civco, D. L., 1993, Artificial Neural Networks for Land-Cover Classification and Mapping, *International Journal of GIS*, Vol. 7, No. 2, pp.173 - 186.

Clementini, E., Di Felice, P. and Van Oosterom, P., 1993, A Small Set of Formal Topological Relationships Suitable for End-User Interaction, in *Symposium on Spatial Databases '93: Advances in Spatial Databases, Abel, D. and Ooi, B. C. (eds)*, Singapore, Springer-Verlag, pp. 277 - 295.

Congalton, R.G., 1991, A Review of Assessing the Accuracy of Classifications of Remotely Sensed Data, *Remote Sensing of the Environment*, Vol. 37, pp. 35 - 46.

Couclelis, H., 1992, People Manipulate Objects (but cultivate fields): Beyond the Raster-Vector Debate in GIS, *International Conference on GIS: Theories and Methods of Spatio-Temporal Reasoning in Geographic Space, Frank, A. V. Camperi, I and Formentini, U. (eds)*, Pica, Italy,*)*, Springer Verlag, pp. 65 - 77.

Crowther, P., Hartnett, J. and Williams, R. N, 1994, Documentation of Design and Structure in Expert Systems Rule Base, *Proceedings of the 7th Australian Joint Conference on Artificial Intelligence*, Armidale, NSW, pp. 568-574.

Crowther, P. and Hartnett, J., 1996a, Using Repertory Grids for Knowledge Acquisition for Spatial Expert Systems, *Proceedings of the 4th Australian and New Zealand Conference on Intelligent Information Systems*. Adelaide.

Crowther, P., Hartnett, J., 1996b, Knowledge Acquisition for Expert Systems used with a Geographic Expert Systems, *Proceedings of the 24th Annual International Conference and Technical Exposition of the Australasian Urban and Regional Information Systems Association*. Hobart.

Crowther, P. and Hartnett, J. 1997a, The Problem with Lines - Spatial Reasoning with One Dimensional Objects in a Two Dimensional Representation of Three Dimensional Space, *Proceedings of the 8th Joint European Conference and Exhibition on Geographic Information Systems*. Vienna.

Crowther, P. and Hartnett, J. 1997b, Eliciting Knowledge With Visualization - Instant Gratification For The Expert Image Classifier Who Wants To Show Rather Than Tell. *Proceedings of GeoComputation 1997*, Otago, New Zealand

Crowther, P., Hartnett. J. and Williams, R. N., 1997, A Visual Tool For Capturing The Knowledge Of Expert Image Interpreters: A Picture Is Worth More Than A Thousand Words, *Proceedings of the International Geoscience And Remote Sensing Symposium*, Singapore.

Crowther, P., 1998, Knowledge Acquisition from Satellite Images for Geographic Expert Systems – Show or Tell?, *Proceedings of the International Conference on Modeling Geographic and Environmental Systems with Geographic Information Systems*, Hong Kong, June.

Cui, Z., Cohn, A. G. and Randell, D. A., 1993, Qualitative and Topological Relationships in Spatial Databases, in *The Proceedings of the 3rd International Symposium on Large Spatial Databases Advances in Spatial Databases, Abel, D. and Ooi, B.C. (eds)*, Singapore, Springer-Verlag, pp. 296 - 315.

Cullen, J. and Bryman, A., 1988, The Knowledge Acquisition Bottleneck: Time for Reassessment, *Expert Systems*, Vol. 5, No. 3, pp. 216 - 225.

Curran, P. J. and Williamson, H. D., 1986, Sample Size for Ground and Remotely Sensed Data, *Remote Sensing of the Environment*, Vol. 20, pp. 31 - 41.

Davis, F. W. and Simonett, D. S.1991, GIS and Remote Sensing, in *Maguire,D.J., Goodchild, M.F. and Rhind, D.W. (eds.), Geographic Information Systems: Principles and Applications, Vol. 1*, Longman, London.

DeMers, M. N., 1986, A Knowledge Base Acquisition Strategy for Expert Geographic Information System Development, in *Optiz, B.K. (ed.), GIS in Government,* Hampton, VA, pp. 837 – 850.

Ding, Y. and Desham, P. J., 1996, Spatial Strategies for Parallel Spatial Modelling, *International Journal of GIS*, Vol. 10, No. 8, pp. 669 - 698.

Djoko, S., Cook, D. J. and Holder, L. B., 1997, An Empirical Study of Domain Knowledge and Its Benefits to Substructure Discovery, *IEEE Transactions on Knowledge and Data Engineering*, Vol. 9, No. 6, pp. 575 - 585.

Drewniok, C., 1994, Multi-Spectral Edge Detection. Some Experiments on Data from Landsat-TM, *International Journal of Remote Sensing*, Vol. 15, No. 18, pp. 3743 - 3765.

Egenhofer, M. J. 1991, Reasoning About Binary Topological Relations, in *Advances in Spatial Databases, The Proceedings of the 2nd International Symposium on Large Spatial Databases, Gunther, O and Schek, H.J. (eds)*, Springer Verlag, pp. 143 - 159.

Egenhofer, M. J. and Herring, J. R.(eds), 1995, *Advances in Spatial Databases, The Proceedings of the 4th International Symposium on Large Spatial Databases,* Springer, Portland, ME, USA.

Egenhofer, M. J. and Mark, D. M., 1995, Naïve Geography, *Proceedings of the Conference on Spatial Information Theory '95, Frank, A. U., and Kuhn, W. (eds)*, Semmering, Austria, pp. 1 – 16.

Egenhofer, M. J. and Sharma, J., 1993, Topological Relations Between Regions in IR2 and ZZ2, *in Advances in Spatial Databases, The Proceedings of the 3rd International Symposium on Large Spatial Databases Advances in Spatial Databases Abel, D. and Ooi, B.C. (eds)* Singapore, Springer-Verlag, pp. 316 - 336.

Eklund, P. W. and Salim, A., 1993, An Experiment with Automated Acquisition of Classification Heuristics from GIS, *Proceedings of LIM, GIS and Advanced Remote Sensing Conference*, Sydney, Australia, pp. 83 - 93.

Faure, C., Frediani, S. and Saitta, L., 1993, A Semi-automated Methodology for Knowledge Elicitation, *IEEE Transactions on Systems, Man and Cybernetics*, Vol. 23, No. 2, pp. 346 – 356.

Frank, A. V. Camperi, I and Formentini, U. (eds), 1992, *Proceedings of the International Conference on GIS: Theories and Methods of Spatio-Temporal Reasoning in Geographic Space, Pica, Italy*, Springer-Verlag.

Fensel, D. and Van Harmelen, F., 1994, A Comparison of Languages which Operationalize and Formalise KADS Models of Expertise, *Knowledge Engineering Review*, Vol. 9, No. 2, pp.105 - 146.

Fischer, M. M.,1998, GIS and Spatial Analysis: Towards a New Generation of Spatial Data Analysis Models, *Proceedings of International Conference on Modelling Geographical and Environmental Systems with Geographic Information Systems*, Hong Kong, Vol. 2, pp.474 – 580.

Fischer, M. M., Scholten, H.J. and Unwin, D. (eds), 1996, *Spatial Analytical Perspectives on GIS GISDATA IV*, Taylor and Francis.

Fleck, M. M., 1996, The Topology of Boundaries, *Artificial Intelligence*, Vol. 80, pp. 1 - 27.

Freksa, C., 1992, Using Orientation Information for Qualitative Spatial Reasoning, *Proceedings of the International Conference on GIS: Theories and Methods of Spatio-Temporal Reasoning in Geographic Space, Frank, A. V. Camperi, I and Formentini, U. (eds)*, Pica, Italy, Springer-Verlag, pp.162 - 178.

Fuller, S, 1994, The Constitutively Social Character of Expertise, *International Journal of Expert Systems*, Vol. 7, No. 1, pp. 51 - 64.

Gaines, B. R. and Shaw, M. L .D., 1993, Knowledge Acquisition Tools based on Personal Construct Psychology, *The Knowledge Engineering Review*, Vol. 8, No. 1, pp. 49 - 85.

Galetto, R. and Spalla, A., 1996, Multimedia Technology: A New Opportunity for GIS Improvement, in *Mussio, L., Forlani, G. and Crosilla, F. (eds)*, Springer Wien, NY, pp. 1 – 7.

Garg-Janardan, C. and Salvendy, G., 1990, A Structured Knowledge Elicitation Methodology for Building Expert Systems, in *Boose and Gaines, Data Acquistion and Analysis for Multimedia GIS*, Springer Wien, New York, pp. 85 – 113.

Gerbrands, J. J., 1993, Digital Image processing, *in Buiten and Clevers*, pp. 275-285.

Giarratano, J. and Riley, G., 1994, *Expert Systems: Principles and Programming (2^{nd} edition)*, PWS, Boston.

Gould, M., Nunes, J., Comas, D., Egenhofer, M .J., Freundschuh and Mark, D. M., 1996, Formalising Informal Geographic Information: Cross-Cultural Human Subjects Testing, *Proceedings of 2^{nd} Joint European Conference and Exhibition On GIS*, Barcelona, Spain.

Graham and Llewelyn-Jones, 1988, *Expert Systems: Knowledge and Decisions*, Chapman and Hall.

Grigni, M., Papadias, D. and Papadimitriou, C., 1995, Topological Inference, *Proceedings of the 13^{th} International Joint Conference on Artificial Intelligence*, Montreal, pp. 901 - 906.

Gruen, A. and Li, H, 1995, Road Extraction from Aerial and Satellite Images by Dynamic Programming, *ISPRS Journal of Photogrammetry and Remote Sensing*, Vol. 50, No. 4, pp.1 - 20.

Gupta, A., Santini, S. and Jain, R., 1997, In Search of Information in Visual Media, *Communications of the ACM*, Vol. 40, No. 2, pp. 35 – 42.

Hallett, S. H., Jones, R. J. A. and Keay, C. A., 1996, Environmental Information Systems Developments for Planning Sustainable Land Use, *International Journal of GIS*, Vol. 10, No. 1, pp. 47 - 64.

Harris, R, 1997, *What's New in Satellite Technology*, Technical Report, Department of Geography, University College, London.

Hartnett, J., Crowther, P. and Williams, R. N., 1997, Three Spatial Analysis Techniques which Incorporate Expert Knowledge Applied to the Selection Of the Best Site for a Forestry Plantation, *Proceedings of the 8^{th} Joint European Conference and Exhibition on Geographic Information Systems*, Vienna.

Hartnett, J., Williams, R. N. and Crowther, P., 1994, Per Pixel Reasoning Using a GIS Closely Coupled to an Expert System to Produce Surface Classifications Based on Remotely Sensed Data and Expert Knowledge, *Proceedings of the 5^{th} Joint European Conference and Exhibition on Geographic Information Systems*, Paris.

Hartnett, J., Williams, R. N., Crowther, P., and Hill, K.,1996, Providing a Graphical User Interface for Customised GIS Applications, *Proceedings of the 7^{th} Joint European Conference and Exhibition on Geographical Information*, Barcelona, Spain.

Harvey, D., 1973, *Explanation in Geography*, Edward Arnorld, London.

Hazelhoff, L. and Gunnik, J. L., 1994, Linking Tools for Exploratory Analysis of Spatial Data with GIS, *Proceedings of the 5th Joint European Conference and Exhibition on Geographic Information Systems*, Paris.

Hernandez, D, 1993, Maintaining Qualitative Spatial Knowledge, *Proceedings of the Conference on Spatial Information Theory*, Elba Island, pp. 36 – 53.

Hernshaw, H. M., Psychology and Displays in GIS, in Hernshaw, H. M. and Unwin, D.J., (eds.), *Visualisation in Geographic Information Systems*, Wiley, Chichester, pp. 193 - 199.

Hernshaw, H. M. and Unwin, D.J., (eds.), 1994, *Visualisation in Geographic Information Systems*, Wiley, Chichester.

Hinton, J. C., 1996, GIS and Remote Sensing Integration for Environmental Applications, *International Journal of GIS*, Vol. 10, No. 7, pp. 877 - 890.

Ibbs, T. J. and Stevens, A., 1988, Quadtree Storage of Vector Data, *International Journal of GIS*, Vol. 2, No. 1, pp. 43-56.

Jeng, B., Liang, T. P. and Hong, M., 1996, Interactive Induction of Expert Knowledge, *Expert Systems with Applications*, Vol. 10, No. 4, pp. 393 - 401.

Jungert, E., 1993, Symbolic Spatial Reasoning on Object Shapes for Qualitative Matching, *Proceedings of the Conference on Spatial Information Theory*, Elba Island, pp. 444 – 462.

Kelly, G. A., 1955, *The Psychology of Personal Constructs*, New York: Norton.

Kendall, E. A. and Senjen, R., 1993, A Hypertext-based Knowledge Organization Tool, *Proceedings of the 6th Australian Joint Conference on Artificial Intelligence*, pp. 139 - 144, Melbourne.

Khosla R. and Dillon, T., 1993, Perspectives for Integration of Artificial Neural Networks with Expert Systems, *Proceedings of the Australian and New Zealand Conference on Intelligent Information Systems*, Perth, Western Australia, pp. 624 - 628.

Kilpatrick, D. and Williams, R., 1995, Unsupervised Classification of Antarctic Satellite Imagery Using Kohonen's Self-Organising Feature Map, *Proceedings of IEEE International Conference on Neural Networks*, Perth, Western Australia, Vol.1, pp. 32 - 36.

Kodratoff, Y., Moustakis, V. and Graner, N, 1994, Can Machine Learning Solve My Problem?, *Applied Artificial Intelligence*, Vol. 8, No. 1, pp. 1 - 32.

Knotes, C. Wilkinson, G. G., Burrill, A., Goffredo, S. and Megier, J., 1993, An Experimental System for the Integration of GIS Data in Knowledge-Based Image Analysis for Remote Sensing of Agriculture, *International Journal of GIS*, Vol. 7, No. 3, pp. 247 - 262.

Koperski, K. and Han, J, 1995, Discovery of Spatial Association Rules in Geographic Information Databases, *Advances In Spatial Databases, Proceedings of 4th International Symposium on Large Spatial Databases,* Springer, Portland, ME, USA, pp. 47 - 66.

Kuhn, W., 1993, Metaphors Create Theories for Users, *Proceedings of the Conference on Spatial Information Theory,* Elba Island, Italy, pp. 366 – 376.

Kumar, V. R., Chung, C.Y.C, and Lindley, C.A., 1994, Towards Building an Expert System for Weather Forecasting Operations, *Expert Systems with Applications*, Vol. 7, No. 2, pp. 373 - 381.

Kweon, I. S. and Kanade, T., 1994, Extracting Topographic Terrain Features From Elevation Maps, *CVGIP: Image Understanding*, Vol. 59, No. 2, pp.171 - 182.

Mejia-Lavalle, M. and Rodriguez-Ortiz, G, 1998, Obtaining Expert System Rules Using Data Mining Tools From A Power Generation Database, *Expert Systems with Applications*, Vol. 13, No. 1/2, pp. 37 - 42.

Mussio, L., Forlani, G. and Crosilla, F.(eds), 1996, *Data Acquisition and Analysis for Multimedia GIS,* Springer Wien, N.Y.

Laurini, R. and Thompson. D., 1992, *Fundamentals of Spatial Information Systems*, Academic Press, pp. 39 - 42.

Lee, Y. C. and Chin, F. L., 1995, An Iconic Query Language for Topological Relationships in GIS, *International Journal of GIS*, Vol. 9, No. 1, pp. 25 - 46.

Leung, Y., 1997, *Intelligent Spatial Decision Support Systems*, Springer, Ch. 2, pp. 11 - 58.

Lillesand, M. T. and Kiefer, R.W., 1994, *Remote Sensing and Image Interpretation (3rd edition)*, John Wiley and Sons.

Lloyd, R., 1997, *Spatial Cognition: Geographic Environments*, Kluwer Academic Publishers, Dordrecht.

Maguire,D. J., Goodchild, M. F. and Rhind, D. W.(eds), *Geographic Information Systems: Principles and Applications, Vol. 1*, Longman, London.

Mair, J. M. M., 1966, Prediction of Grid Scores, *British Journal of Psychology*, No. 57, pp.187 - 192.

Mak, B., Bui, T. and Blanning, R., 1996, Aggregating and Updating Experts' Knowledge: An Experimental Evaluation of Five Classification Techniques, *Expert Systems With Applications*, Vol. 10, No. 2, pp. 233 - 241.

Mather, P. M., 1991, *Computer Applications in Geography*, John Wiley and Sons, Chichester.

McGraw, K. L. and Harbison-Briggs, K., 1989, *Knowledge Acquisition: Principles and Guidelines*, Prentice Hall, New Jersey.

McKeown, D. M., Harvey, W.A. and Wixson, L.E., 1989, Automating Knowledge Acquisition for Aerial Image Interpretation, *Computer Vision* and *Graphics*, Vol. 46, pp. 37 - 81.

McMaster, R. B. and Mark, D. M., 1991, The Design of Graphical User Interface for Knowledge Acquisition in Cartographic Generalization, *GIS/LIS'91 Proceedings*, Atlanta, pp. 311 - 320

McMaster, R. B., 1991, Conceptual Frameworks for Geographic Knowledge, in *Buttenfield, B. P. and McMaster, R. B. (eds.), Map Generalization: Making Rules for Knowledge Representation*, Longman, Ch. 2, pp.21 - 39.

Medsker, L., Tan, M. and Turban, E., 1995, Knowledge Acquisition from Multiple Experts: Problems and Issues, *Expert Systems With Applications*, Vol. 9, No. 1, pp. 35 - 40.

Medyckj-Scott, D. J., 1994, Visualization in Human-Computer Interaction in GIS, in Hernshaw, H. M. and Unwin, D.J., (eds.), *Visualisation in Geographic Information Systems*, Wiley, Chichester, pp. 200 – 211.

Medyckj-Scott, D. J. and Blades, M., 1991, Cognitive representations of space in the design and use of geographical information systems, in Raper, J. (ed.), *Three Dimensional Applications in GIS*, Taylor and Francis, London, pp. 80 – 98.

Meseguer, P. and Preece, A. D., 1995, Verification and Validation of Knowledge-Based Systems with Formal Specifications, *Knowledge Engineering Review*, Vol. 10, No. 4, pp. 331 - 343.

Meyer, M. and Booker, J., 1991, *Eliciting and Analysing Expert Judgement, A Practical Guide, Knowledge Based Systems Vol. 5*, Academic Press, London, Ch. 1, 11.

Miller, D. R..,1994, Knowledge-Based Systems Coupled with Geographic Information Systems, *Proceedings of the International Workshop on Advanced Research in GIS*, Nievergelt, J. (ed). Springer-Verlag, Monte Verita, Ascona, Switzerland, pp. 143 – 154.

Mockler, R. J., 1989, *Knowledge Based Systems for Management Decisions*, Prenyice-Hall, New Jersey.

Mockler, R. J. and Dologite, D. G., 1992, *Knowledge-Based Systems - An Introduction to Expert Systems*, McMillan.

Mussio, L., Forlani, G. and Crosilla, F. (eds), 1996, *Data Acquistion and Analysis for Multimedia GIS*, Springer Wien, New York.

Muller, J-P. A. L., 1988, Key Issues in Image Understanding in Remote Sensing, *Phil. Transansactions Royal Society London*, A 324, pp. 381 – 395.

Nabil, M., Shepherd, J. and Ngu, A. H. H., 1995, 2D Projection Interval Relationships: A Symbolic Representation of Spatial Relationships, *Advances In Spatial Databases, Proceedings of 4th International Symposium on Large Spatial Databases,* Springer, Portland, ME, USA, pp. 292 - 309.

Nagao, M., 1990, *Knowledge and Inference*, Academic Press, San Diego.

Neale, I. M., 1987, First Generation Expert Systems: A Review of Knowledge Acquisition Methodologies, *The Knowledge Engineering Review*, Vol. 3, No.2, pp.105 - 145.

Olson, R. J. and Rueter, H. H., 1987, Extracting Expertise from Experts: Methods for Knowledge Acquisition, *Expert Systems*, Vol. 4, No. 3, pp. 152-168.

Openshaw, S. and Clarke, G., 1996, Developing Spatial Analysis Functions Relevant to GIS Environments, in *Fischer, M. M., Scholten, H.J. and Unwin, D. (eds), Spatial Analytical Perspectives on GIS GISDATA IV,* Taylor and Francis, Ch. 2, pp. 21 - 37.

Openshaw, S., Wymer, C. and Cross, A., 1991, Using Neural Nets to Solve Some Hard Analysis Problems in GIS, *Proceedings of the 2nd European Conference on GIS*, Brussels, pp. 795 - 807.

Openshaw, S., 1993, Modelling Spatial Interaction Using a Neural Net, *GIS Spatial Modelling and Policy Evaluation*, Springer-Verlag, pp. 147 - 164.

Openshaw, S., 1991, Developing Appropriate Spatial Analysis Methods for GIS, in *Geographic Information Systems: Principles and Applications, Maguire, Goodchild and Rhind (eds.), Vol. 1,* pp. 389 - 402.

Openshaw, S., 1993, Some Suggestions Concerning the Development of Artificial Intelligence Tools for Spatial Modelling and Analysis in GIS, *GIS Spatial Modelling and Policy Evaluation*, Springer-Verlag, pp. 17 - 33.

Openshaw, S. and Scholten, H., 1994, Spatial Analysis and Geographic Information Systems: A Challenge for the 1990's, *Workshop on Spatial Analysis, Proceedings of the 5th Joint European Conference and Exhibition on Geographic Information Systems,* Paris.

Optir, B.K. (1986), *Geographic Information Systems in Government,* Hampton, VA.

Palmer, B., 1984, Symbolic Feature analysis and Expert Systems, *Proceedings of the International Conference on Spatial Data Handling*, pp. 465 - 479.

Pigford, D.V. and Baur, G., 1990, *Expert Systems for Business*, Boyd and Fraser, Boston, MA., Ch. 2.

Price, C. J., 1990, *Knowledge Engineering Toolkits*, Ellis Horwood, New York.

Quinlan, J. R., 1986, Induction of Decision Trees, *Machine Learning*, Vol. 1, pp. 81-106.

Raghu, P. P. and Yegnanarayana, 1997, Multi-spectral Image Classification Using Gabor Filters and Stochastic Relaxation Neural Network, *Neural Networks*, Vol. 10 No. 3, pp. 561-572.

Raper, J. (ed.), 1990, *Three Dimensional Applications in GIS*, Taylor and Francis, London.

Research Systems, Inc., 1994, *Interactive Data Language Version 3.6*

Richards, J. A., 1993, *Remote Sensing Digital Image Analysis*, Springer-Verlag, New York, Chs. 4, 6 and 8.

Rigaux, P. and Scholl, M., 1994, Multiple Representation Modelling and Querying, *Proceedings of the International Workshop on Advanced Research in GIS*, pp. 59–69.

Rouge, A., Lapicque, J. Y., Brossier, F. and Lozinguez, Y., Validation and Verification of KADS Data and Domain Knowledge, *Expert Systems with Applications*, Vol. 8, No. 3, pp. 333-341.

Sabins, F. F., 1997, *Remote Sensing: Principles and Interpretation (3rd edition)*, Freeman and Co New York.

Schotten, C. G. J., Van Rooy, W. W. and Janssen, L. L. F., 1995, Assessment of the Capabilities of Multi-Temporal ERS-1 SAR Data to Discriminate Between Agricultural Crops, *International Journal of Remote Sensing*, Vol. 16, No. 14, pp. 2619 - 2637.

Schowengerdt, R. A., 1997, *Remote Sensing, Models and Methods for Image Processing (2^{nd} edition)*, Academic Press.

Schreiber, G., Wielinga, B., Akkermans, H. Van de Velde, W. and Anjewierden, A., 1994, CML: The CommonKADS Conceptual Modelling Language, *Proceedings of the 8^{th} European Knowledge Acquisition Workshop*, Hoegaarden, Belgium,

Schreiber, G., Wielinga, B. and Breuker, J., 1993, *KADS - A Principled Approach to Knowledge-Based System Development*, Academic Press, San Diego

Senjen, R. and Mee, W., 1993, Epistemological Implications of the Use of Participant Observation in Knowledge Engineering, *Proceedings of the 7th Australian Joint Artificial Intelligence Conference*, Melbourne, pp. 160 - 165.

Sharma, J. and Flewelling, D.M., 1995, Inferences from Combined Knowledge about Topology and Directions, *Advances In Spatial Databases, Proceedings of 4th International Symposium on Large Spatial Databases,* Springer, Portland, ME, USA, pp. 279 - 291.

Shea, K. S.,1991, Design Considerations for an Artificially Intelligent System, in *Buttenfield, B.P. and McMaster, R.B. (eds), Map Generalization: Making Rules for Knowledge Representation,* Longman, Ch. 1, pp. 3 - 20.

Skidmore, A. K., 1995, Neural Networks and GIS, 1995, *The Australasian GIS Applications Journal*, No.11, May-July, pp. 53 - 54.

Skidmore, A. K., 1989, An Expert System Classifies Eucalypt Forest Types Using Thematic Mapper Data And A Digital Terrain Model, *Photogrammetric Engineering And Remote Sensing,* Vol. 55, No. 10, pp. 1449 - 1464.

Sleeman, D. and Mitchell, F, 1996, Towards Painless Knowledge Acquisition, *Proceedings of the Tenth European Knowledge Acquisition Workshop, Lecture Notes in A.I. 1076, Advances in Knowledge Acquisition*, pp. 262 - 277.

Smith, T. R. and Park K. K., 1992, Algebraic Approach to Spatial Reasoning, *International Journal of GIS*, Vol. 6, No. 3, pp. 177 - 192.

Srinivasan A and Richards, J., 1990, Knowledge-Based Analysis of Multi-Source Data, *Proceedings of 5th Australasian Conference on Remote Sensing*, pp. 978 - 986.

Srinivasan A and Richards J., 1990, Knowledge-Based Techniques For Multi-Source Classification, *International Journal of Remote Sensing*, Vol. 11, No. 3, pp. 505 - 525.

Star, J. L. and Estes, J. E., 1990, *Geographic Information Systems: An Introduction*, Prentice Hall, New Jersey.

Star, J. L., Estes, J. E. and McGuire, K. C., 1997, *Topics in Remote Sensing 5: Integration of Geographic Information Systems and Remote Sensing*, Cambridge University Press.

Story, P. A. and Worboys, M. F., 1995, A Design Support Environment for Spatio-Temporal Database Applications, *Proceedings of the Conference on Spatial Information Theory, Frank, A. U., and Kuhn, W. (eds),* Semmering, Austria, pp. 413 – 442.

Tranowski, D. 1990, A Knowledge Acquisition Environment for Scene Analysis, in *The Foundations of Knowledge Acquisition - Knowledge Based Systems Vol. 4, Boose, J. and Gaines, B. (eds.)* Academic Press, pp. 339 - 355.

Trimble, E. G. and Cooper C. N., 1987, Experience of Knowledge Acquisition for Expert systems in Construction, *Proceedings of the First European Knowledge Acquisition Workshop*, Reading, Great Britain, C5.

Unwin, A., 1996, Exploring Spatio-Temporal Data, in *Fischer, M. M., Scholten, H. J. and Unwin, D. (eds.), Spatial Analytical Perspectives on GIS, GISDATA IV,* Taylor and Francis, Ch. 7, pp.101 - 110.

Unwin, D., 1996, Integration Through Overlay Analysis, in *Fischer, M. M., Scholten, H. J. and Unwin, D. (eds.), Spatial Analytical Perspectives on GIS, GISDATA IV,* Taylor and Francis, Ch. 9, pp.129 - 138.

Vicat, C., Brezillon, P. and Nottola, C., 1995, Knowledge Validation in the Building of a Knowledge-Based System, *Expert Systems with Applications,* Vol. 8, No. 3 pp. 391 - 397.

Vrtacnik, M. and Dolnicar, D., 1995, Evaluation of Automatic Rule Induction Systems, *Expert Systems with Applications*, Vol. 8, No. 1, pp. 77 - 87.

Waterman, D. A., 1986, *A Guide to Expert Systems*, Addison-Wesley, Reading.

Waters, N. M., 1989, Expert Systems Within a GIS: Knowledge Acquisition for Spatial Decision Support Systems, *Proceedings of the ACTES National Conference,* pp. 740 - 759, Otowa.

Welbank, M., 1990, An Overview of Knowledge Acquisition Methods, *Interacting with Computers*, Vol. 2, No. 1, pp. 83 - 90.

Wiebel, R., Keller, S. and Reichenbacher, T, 1995, Overcoming the Knowledge Acquisition Bottleneck in Map Generalization: The Role of Interactive Systems and Computational Intelligence, *Proceedings of the Conference on Spatial Information Theory, Frank, A. U., and Kuhn, W. (eds)*, Semmering, Austria, pp. 139 - 156.

Wielinga, B. J., Schreiber, A. T. and Breuker, J. A., 1992, KADS: A Modelling Approach to Knowledge Engineering, *Knowledge Acquisition*, Vol. 4, No. 1.

Wilkinson, G. G. and Burrill, A., 1991, Integration of GIS-Derived Spatial Data with Satellite Imagery for Improved Land Cover Mapping Using an Expert System, *Proceedings of the 2nd European Conference on GIS*, Brussels, April 2 - 5, pp. 1240 - 1249.

Wilkinson, G. G., 1996, A Review of Current Issues in the Integration of GIS and Remote Sensing Data, *International Journal of GIS,* Vol.10, No. 1, pp. 85 - 101.

Williams, R. N., Crowther, P. and Pendlebury, S., 1997, Design and Development of an Operational Ice Mapping System for Meteorological Applications in the Antarctic, *Proceedings of the International Geoscience And Remote Sensing Symposium,* Singapore.

Williams, R., Crowther, P. and Micheal, K., 1994, Interpretation of Remotely Sensed Sea Ice Images for Climatological Studies using an Expert System, *Proceedings of the 7th Australasian Remote Sensing Conference*, Melbourne.

Wooten, C. T. and Rowley, T. H., 1995, Using Anthropological Interview Strategies to Enhance Knowledge Acquisition, *Expert Systems With Applications*, Vol. 9, No. 4, pp. 469 - 482.

Worboys, M. F., 1995, *GIS: A Computing Perspective*, Taylor and Francis.

GLOSSARY

2D-PIR	2-Dimensional Projection Interval Relationship
AVHRR	Advanced Very High Resolution Radiometer
GIS	Geographic Information System
GPS	Global Positioning System
IASOS	Institute of Antarctic and Southern Ocean Studies
KAGES	Knowledge Acquisition for Geographic Expert Systems
KADS	Knowledge Acquisition and Design System
MBR	Minimum Bounded Rectangle
MIRC	Multi-temporal Imaging for the Remote sensing of Crops
NDVI	Normalised Difference Vegetation Index
NOAA	National Oceanographic and Atmospheric Administration
RSS	Remote Sensing System
SPOT HRV	Système Probatoire d'Observation de la Terre, Haute Résolution Visible
TIAR	Tasmanian Institute of Agricultural Research

APPENDIX A. KAGES OPERATION

KAGES (Knowledge Acquisition for Geographic Expert Systems) is an interactive knowledge engineering tool which captures expert interpretation knowledge from classifiers of remotely sensed satellite images. Currently NOAA AVHRR, SPOT and Landsat images can be processed.

The following is an operational description of the use of the tool using SPOT and Landsat TM images of the North West Coast of Tasmania based near Table Cape.

A.1 MAIN MENU

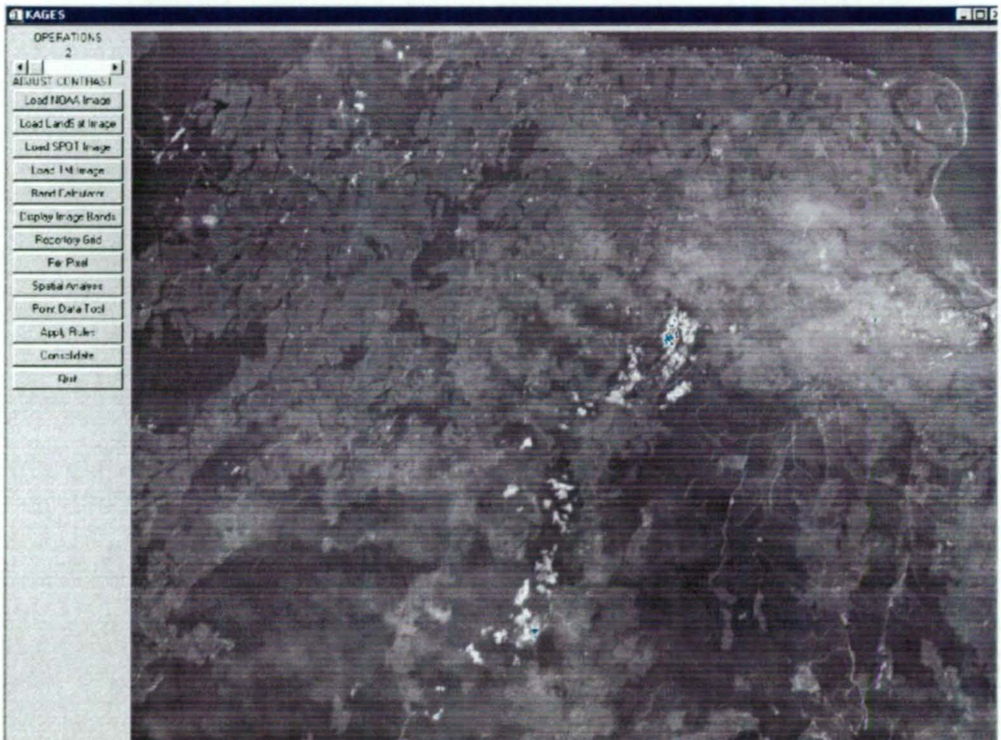

Figure A.1 KAGES main menu with SPOT HRV image (24th December 1996) already loaded

In the above screen shot, an image has already been loaded by selecting the **Load SPOT Image** button. The initial image band displayed is Band 1. Other image bands can then be previewed by selecting the **Display Image Bands** option. **Contrast Adjustment** is provided at this level. Once adjusted here, the contrast value is passed to the rest of the system. In the current system contrast cannot be adjusted in lower level modules.

A.2 LOAD SPOT IMAGE

Load SPOT Image is typical of the **Load** menu options. The common dialog box is a standard IDL load file procedure. It allows a user to select a file containing one band of a multi spectral image. Files need to be in the form X#, where X is any file name and # is the particular band number. In the case of SPOT this is in the range 1 to 3. The number represented by # is then stripped off and replaced by 1 and that file loaded. This is repeated with the # being replaced consecutively by the other band numbers for that image set. Band 1 is then displayed in the main menu window.

A.3 DISPLAY IMAGE BANDS

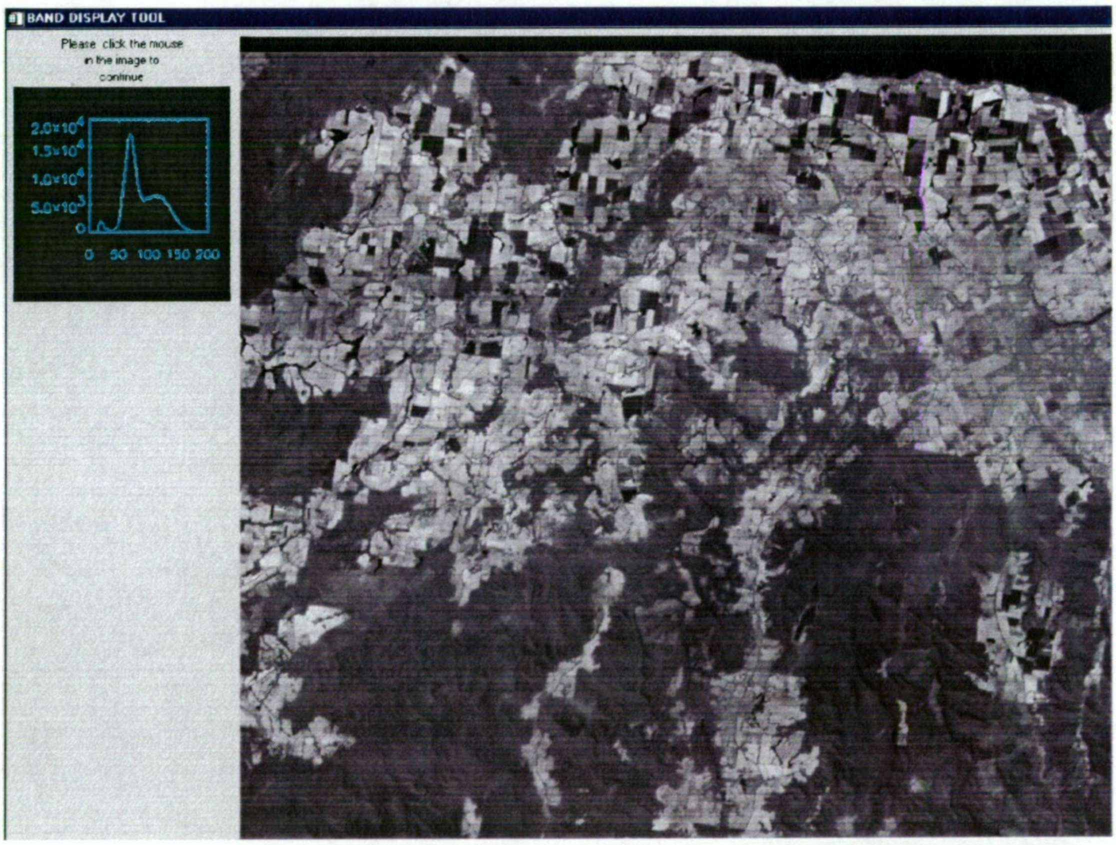

Figure A.2 SPOT image (2nd November 1996) with near infra red (band 3) displayed using the Display Image Bands option

This option allows the display of individual satellite image bands and produces both an image and a histogram of pixels over the whole of the image band. If the contrast is too dark or too light it can be adjusted at the main menu screen. This option can be used in conjunction with the **Band Calculator** (see below) to produce images and histograms based on band composites. Using this facility a user can preview images before using one of the other tools to construct rules.

A.4 BAND CALCULATOR

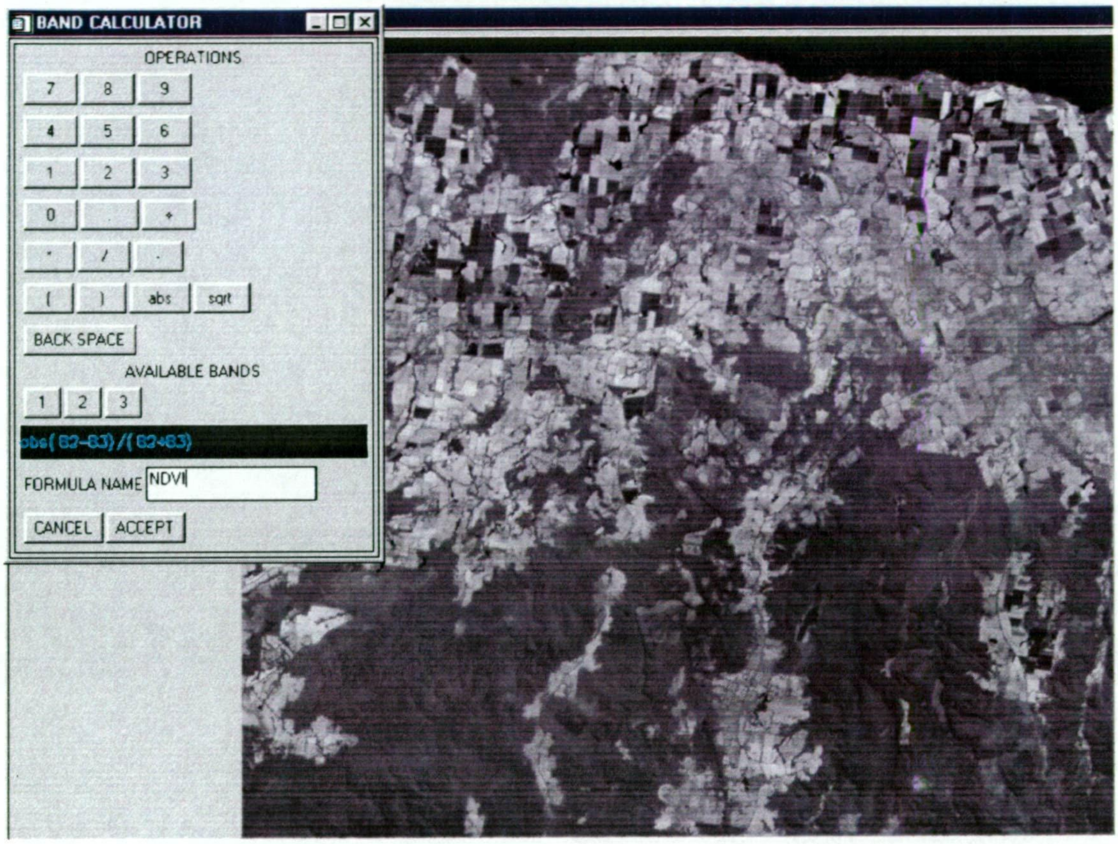

Figure A.3 Band calculator shown in conjunction with a SPOT image (2nd November 1996)

When the Band Calculator button is selected, the band calculator appears superimposed over the initial band 1 image. It is then a case of selecting the appropriate bands to include in an algorithm, then naming the algorithm.

In this example the band calculator has been selected from the main menu and used to enter the NDVI algorithm to aid with crop identification. Once an algorithm has been entered it is stored in a file which is available to all applications. This can be accessed when bands are selected for display in conjunction with other tools such as the **Per-pixel Tool.** When used with other tools in the tool kit, the algorithm is converted to executable IDL code.

A.5 PER-PIXEL TOOL

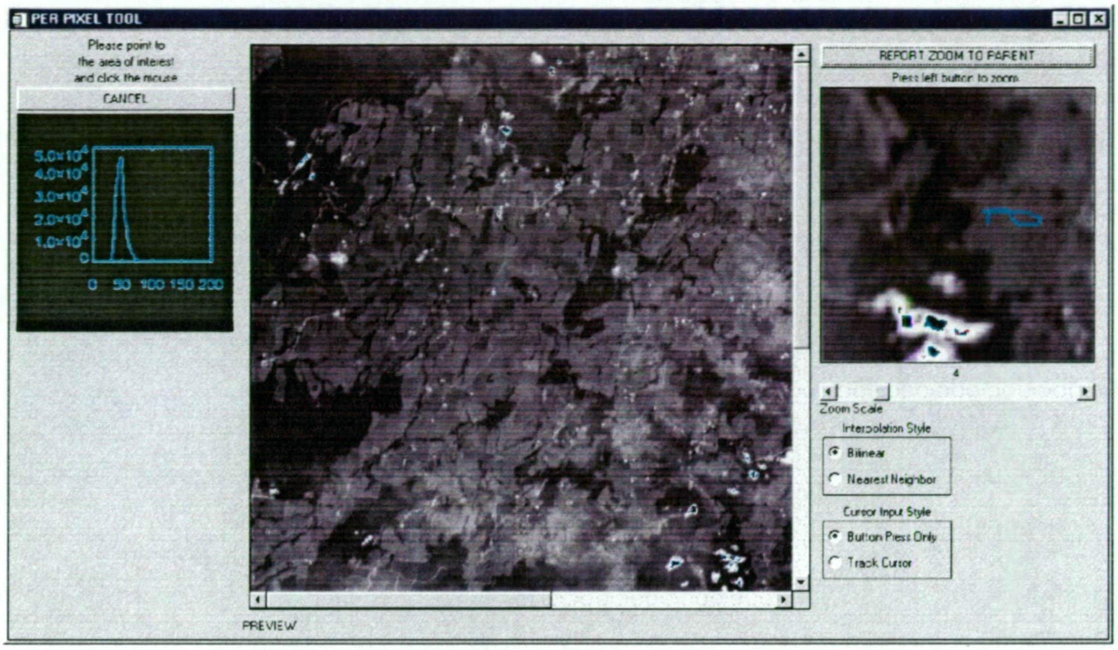

Figure A.4 Per-pixel tool being used to define an area. The zoomed window shows a segment from the bottom right hand corner of the main image window.

There are three components in the Per-pixel tool. The first allows for manipulating areal objects, the other two manipulate point and line objects.

The screen shot above shows the areal knowledge acquisition tool. The initial Per-pixel tool screen allows a user to move the zoom window and zoom in on a feature to aid location. Once that has been done the user "scribbles" over the object of interest. The system then returns the maximum and minimum pixel values of the scribble and uses those to initialise the threshold slider bars

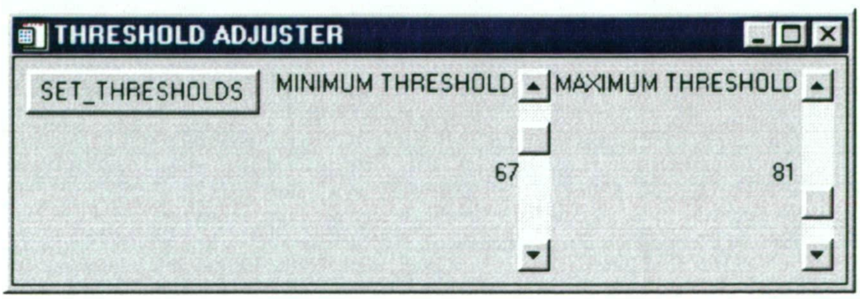

Figure A.5 Threshold adjuster window showing the selected area has pixel thresholds between 67 and 81

The **THRESHOLD ADJUSTER** window is then displayed and if a user knows the thresholds are set too high or too low they can be adjusted at this point. Once the user is satisfied the **SET_THRESHOLDS** button is pressed.

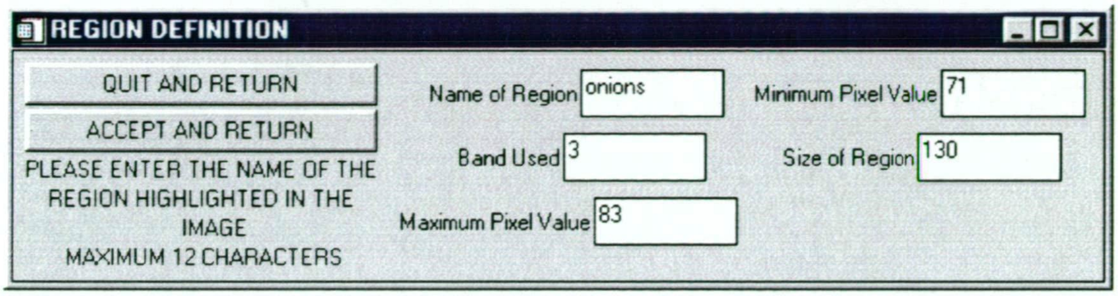

Figure A.6 Region definition (naming) dialog window displaying an area object's characteristics and allowing a user to enter a name

Once the threshold has been set the system locates all pixels within that range contiguous with pixels initially defined by the expert. The system draws a minimum bounded rectangle around the object in the main (central) window and displays a histogram of the pixels in the object in the histogram window.

The **REGION DEFINITION** dialog box is then displayed providing information on the pixel threshold values, the area of the region (in pixels) and the band used. The user can then either enter a name or label for the object, or abandon the operation.

The following rules are copied from the contents of the file containing the generated rules:

```
RULE_1
IF Band =1
AND pixel_value<= 27
AND pixel_value >= 24
AND type  = 3
AND area > 1000
AND month = 12
AND image = sp
THEN FEATURE_NAME = trees
```

RULE_2
IF Band =abs((B3-B2)*100)/(B2+B3)
AND pixel_value <= 16
AND pixel_value >= 6
AND type = 3
AND area > 100
AND month = 12
AND image = sp
THEN FEATURE_NAME = potatoes

RULE_3
IF Band =abs(B4-B3)/(B3+B4)
AND pixel_value<= 10
AND pixel_value >= 0
AND type = 3
AND area > 100
AND month = 12
AND image = tm
THEN FEATURE_NAME = poppies

The first rule uses a raw SPOT band (Band 1). Rules 2 and 3 use composite images defined by the Band Calculator. Rule 3 in this case is defined from a Landsat Image, Rule 2 from SPOT.

Line and Point feature rules are generated in a similar way and have a similar form.

A.5.1 Verification of Per-pixel Rules

While developing Per-pixel rules, the system only displays the instance of the object being selected as a typical example of the object class. Therefore to see the effect of applying the rule to the entire image the **Apply Rules** option in the **Main Menu** can be chosen. This option allows a user to apply both Per-pixel rules and Relationship rules (which will be described below) for verification.

The effect of applying Per-pixel rules can be seen by choosing the **Apply Rules** option from the **Main Menu**. The user then has the choice of applying either **a single rule**, or **all rules**. Choosing **single rule** will present the user with a menu of the primitive objects which have been defined. When an object type is selected, the system selects all pixels which satisfy the clauses and displays them.

Applying all rules produces a per-pixel classification of the image by applying the rule for each of the primitive objects identified in the image.

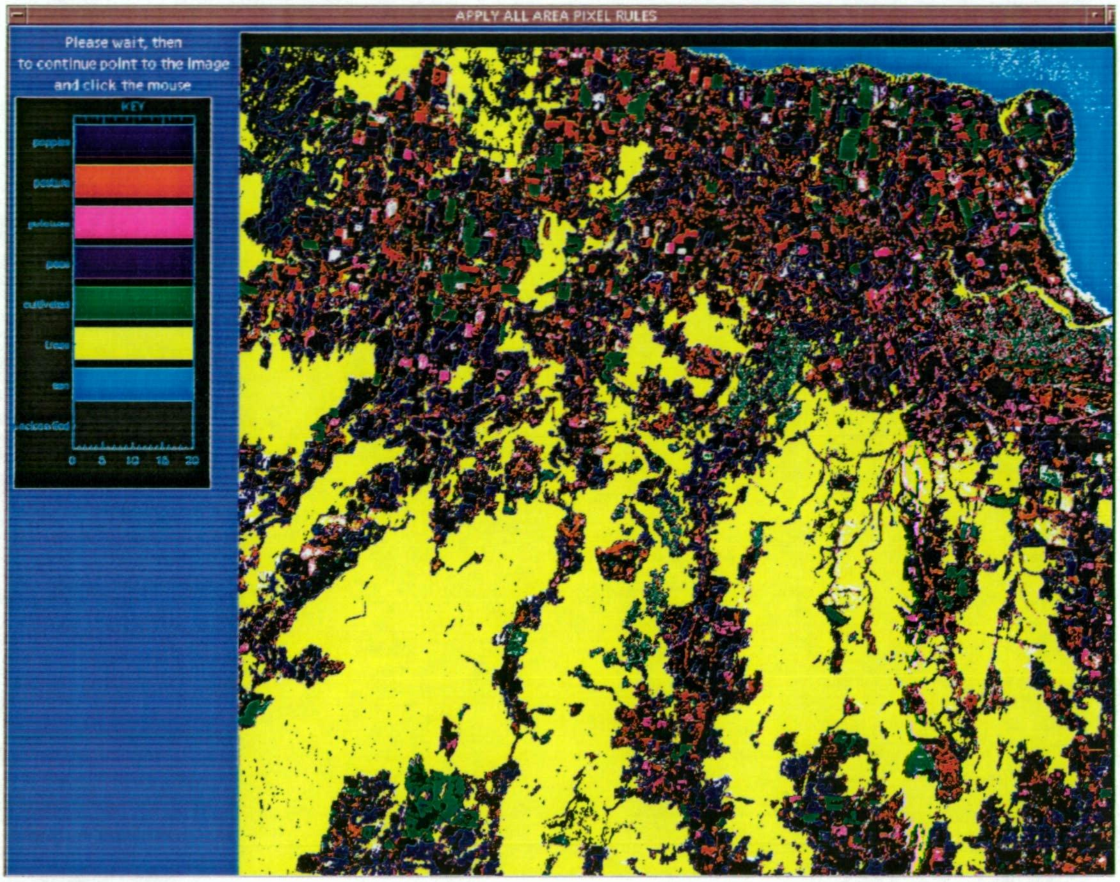

Figure A.7 An Image produced by applying all rules

Therefore, although KAGES is designed to be a knowledge acquisition system, it can do limited classification via its verification options. At present the system applies only Per-pixel rules and can produce images such as the one above. Applying spatial rules would allow areas which have been classified twice (coloured pale mauve in the example) to be further refined.

A.6 RELATIONSHIP TOOL

Once more than one object has been defined, the Relationship Tool can be used to determine spatial relationships. This is selected from the main menu by pressing the **SPATIAL ANALYSIS** button. The Spatial Tool window is displayed and the user selects two objects, then presses **OK**

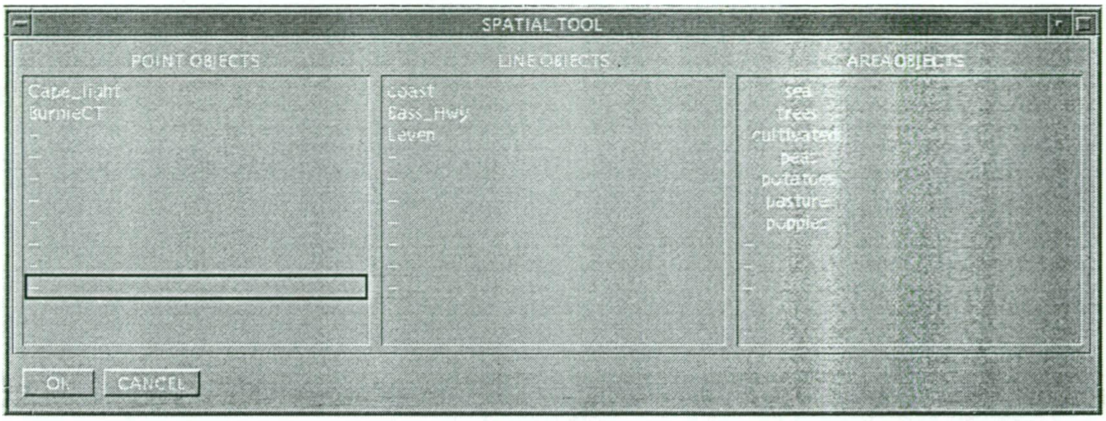

Figure A.8 Initial window from the spatial tool which is superimposed over the main screen. A user is required to pick 2 objects

The system first displays the two selected objects on the image by applying their associated per-pixel rule. Since only two instances of the primitive object are used, ones which are in a typical visual relationship with each other should be selected. The system then draws a minimum bounded rectangle around them before investigating directional, overlap and nearness characteristics.

Direction is calculated in terms of the orientation of the objects centroids. This is to within 45° hence typical values are SOUTH and NORTH_EAST. Overlap characteristics are calculated using the minimum bounded rectangles in the first case, then if they overlap, the objects' boundaries within the overlap of the MBR are checked. Most objects in the agricultural example do not overlap, but many are adjacent to other objects.

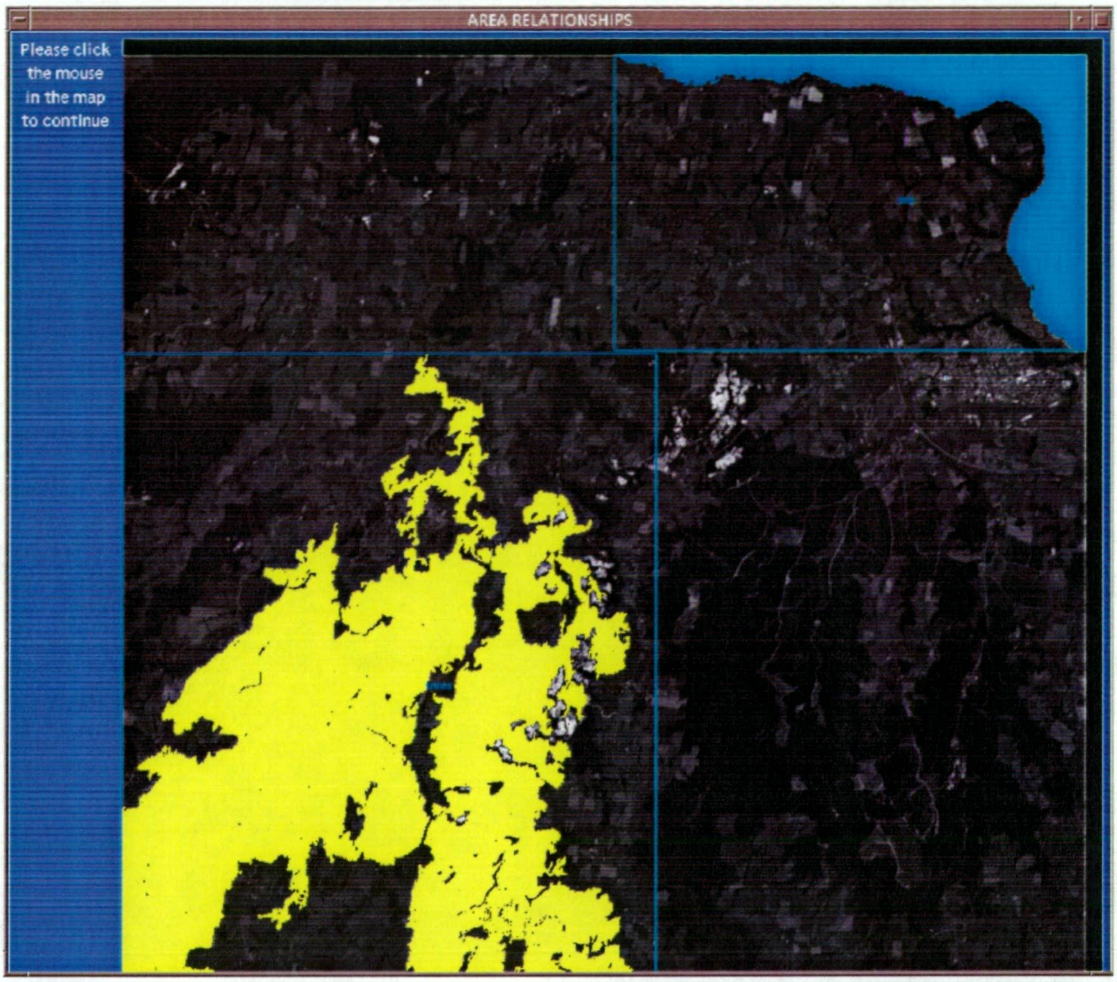

Figure A.9 Spatial tool with the selected objects displayed

In the above example the two objects being used for relationship calculations are trees and sea. One problem with trees is that when the primitive level rule is applied some pixels are classified as trees even though they appear offshore. This effect is probably due to the influence of submerged sea grass. To overcome this, a spatial rule is required.

The above image shows relationships between two areal objects, in this case trees (forest) and an area of sea. Both have a minimum bounded rectangle associated with them and are labelled.

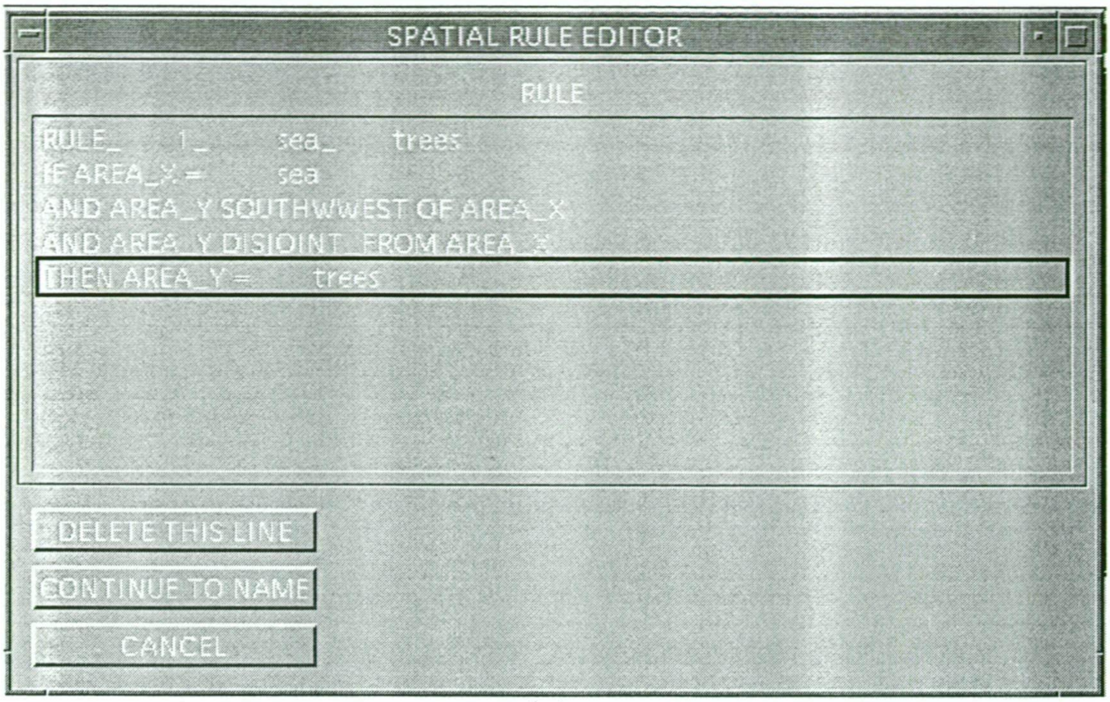

Figure A.10 Spatial Rule editor window showing a rule which expresses a possible relationship between sea and trees

Once the system has completed the task of working out the relationships, the **SPATIAL RULE EDITOR** window automatically appears. This displays all the relationships calculated by the system in the form of a rule. The user has the option of modifying this by deleting clauses that have arisen due to chance. If the rule is acceptable the user can select the **continue to name** option allowing the rule to be saved.

A.6.1 Verification of Spatial Relationship Rules

As well as visually checking the clauses of the rules produced, the effect of applying the rule to the entire image can also be seen. The **apply rules** option can be chosen from the **main menu**. From there **apply spatial rules** can be chosen and the rule to be applied is then selected from a dialogue window.

The system applies the pixel level rules to define the two objects in the spatial rule. The image is then segmented and labels attached to each of the object types. The rule is then applied to all pair combinations of the objects.

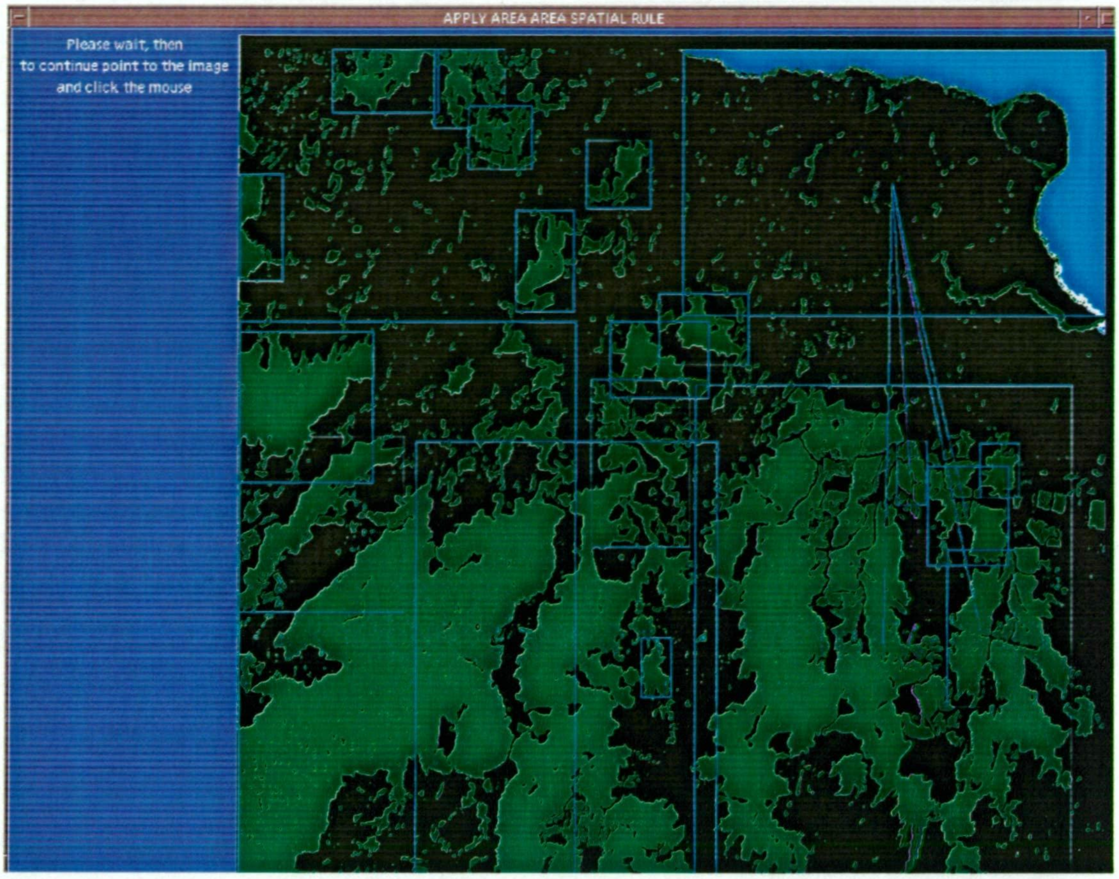

Figure A.11 Effect of selecting Apply Rules / Apply Spatial Rule with the previous rule selected

The effects of applying a spatial rule are shown above. In this case the relationship is between areas of forest and the sea. Only areas greater than 1000 pixels are considered in this case. These are shown with a minimum bounded rectangle drawn around them. Those areas that fulfil the conditions of the rule are shown with their centroids linked. The program also outputs a summary of what happened. In this case:

Total Directional Clause True	6
No Distance Clause	
Total Disjoint Clause True	20
Total Overlap Clause True	0
Total Rules Fired	20
Total rules True	6

The verification process suggests that the rule is too restrictive as there are areas of trees identified on the image which do not match the rule. Ground truth revealed that these areas really were trees. This problem could be corrected by defining a second spatial rule to allow for areas of trees which are also to the south west of the sea area. During consolidation these two rules would be combined with a directional clause:

...

AND (AREA_Y SOUTH OF AREA_X

OR AREA_Y SOUTH_WEST OF AREA_X)

...

A.7 THE REGION OF INTEREST TOOL

The Region Of Interest Tool allows a user to investigate the grouping of objects in an image by tracing an area which may contain one or more scene primitives. This is an option chosen from **Spatial Analysis** in the main menu. The Object Grouping Tool window is displayed and the user uses the mouse to trace an object. Once the trace is compete the area delimited is highlighted and a minimum bounded rectangle drawn around it

Figure A.12 Region of Interest tool defining a rural area (SPOT band 1 taken on 2nd November 1996)

Once the above operation is complete the system displays the **ROI CHECK** window containing a list of object labels representing object types located in the user defined region. The user then has the option of deleting labels which are not typical of the region. In the case of the MIRC project this tool was not used.

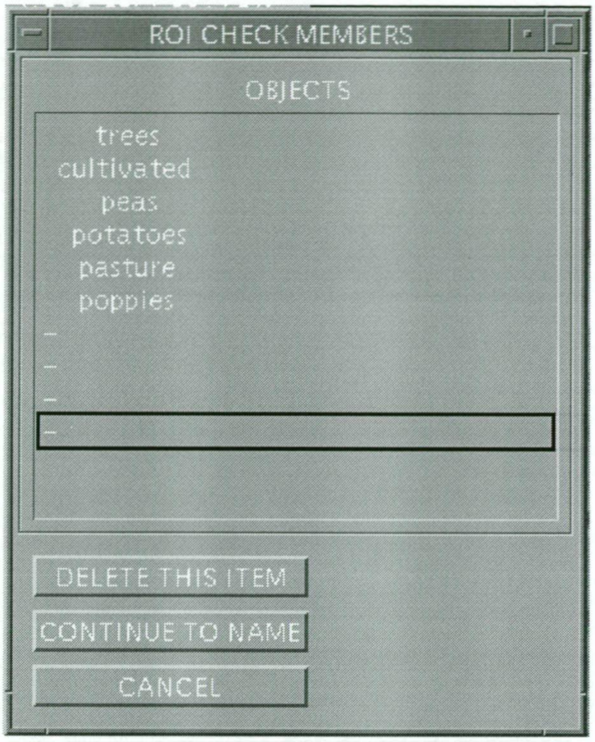

Figure A.13 Region of Interest editor showing object types found

Although the example shown here is only an illustration of how the tool works and does not define a ground truthed object, it illustrates how a particular farm type could be represented. In this case the farm would be an example of mixed cropping. Others with a large percentage of pasture would be more likely to be sheep or cattle production. Once the user is satisfied with the list they can proceed to name the region.

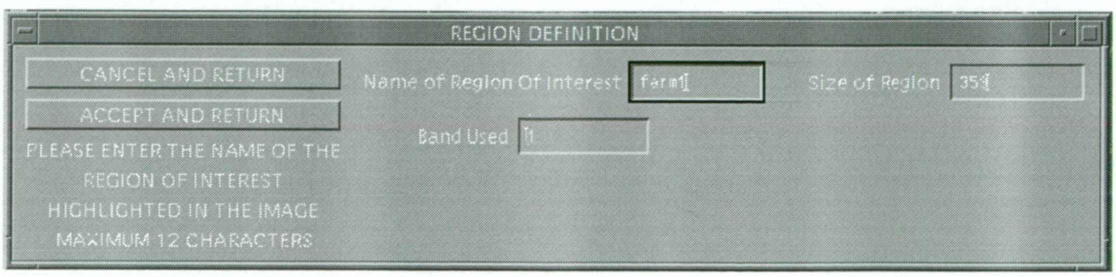

Figure A.14 Naming window for Region of Interest tool

A.8 REPERTORY GRID TOOL

Figure A.15 Repertory Grid Tool after the *New Grid* option has been chosen.

The Repertory Grid Tool can be selected from the main menu after an image is loaded. This image serves as a background and prompt for the repertory grid dialog. There is no manipulation of the image. The screen shots below are normally superimposed on the image as in the example below

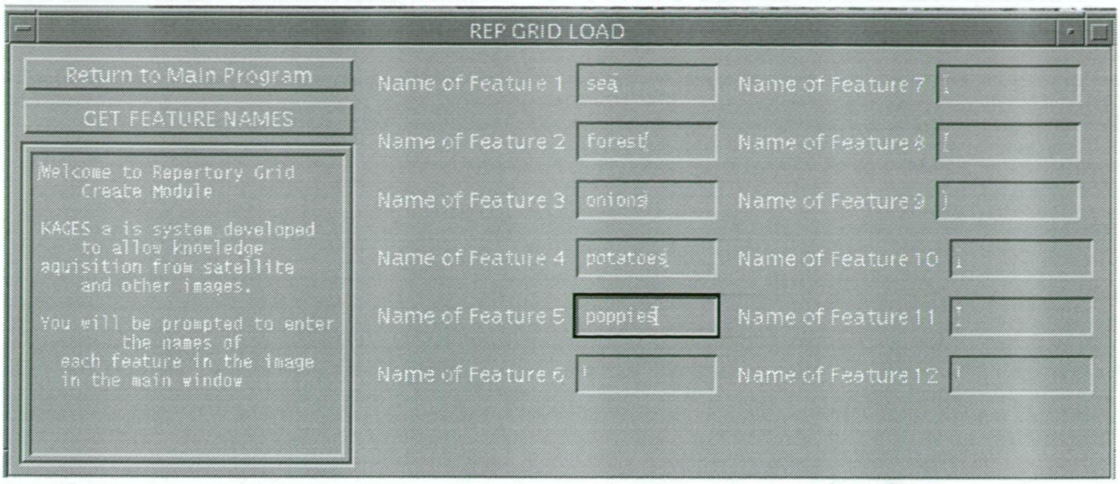

Figure A.16 Repertory Grid Load window allowing entry of image features

Repertory Grid options are:

- Load a Grid,
- New Grid,
- Modify Grid,
- Save Grid,
- Analyse Grid,
- Generate Rules and
- Return to Main KAGES menu.

When a the **New Grid** option is selected from the Repertory Grid main menu the above screen is displayed. This requires an image interpretation expert to name features on the image. To help the expert, band 1 of the currently loaded image is displayed as a backdrop. This is purely as a prompt and cannot be manipulated by the user.

The grid can be updated if the user later remembers a feature they forgot to include in the list, or includes a feature they decide they want to delete by using the **Modify Grid** option. The order and type of feature is unimportant in the repertory grid process.

Once all feature names have been entered by the domain expert, the next step is activated by pressing the **GET FEATURE NAMES** button. This establishes the names entered as column headings of the grid. To abort at this point a user presses the **Return to Main Program** button

Figure A.17 Triad (Get Discriminators) window asking a user to differentiate between features

When **GET FEATURE NAMES** is selected, groups of three objects are displayed to the user who is then asked to name an attribute two of the features have but the third lacks. This process, known as a triad comparison, continues with similar screens until a series of attributes or discriminators have been entered. These become the row labels of the grid. If the user wishes to add extra discriminators, the **MODIFY GRID / ADD DISCRIMINATOR** option at the main Repertory Grid Menu can be used.

In the above example the crop onions has a relatively small area compared with the areas of forest and sea in the image.

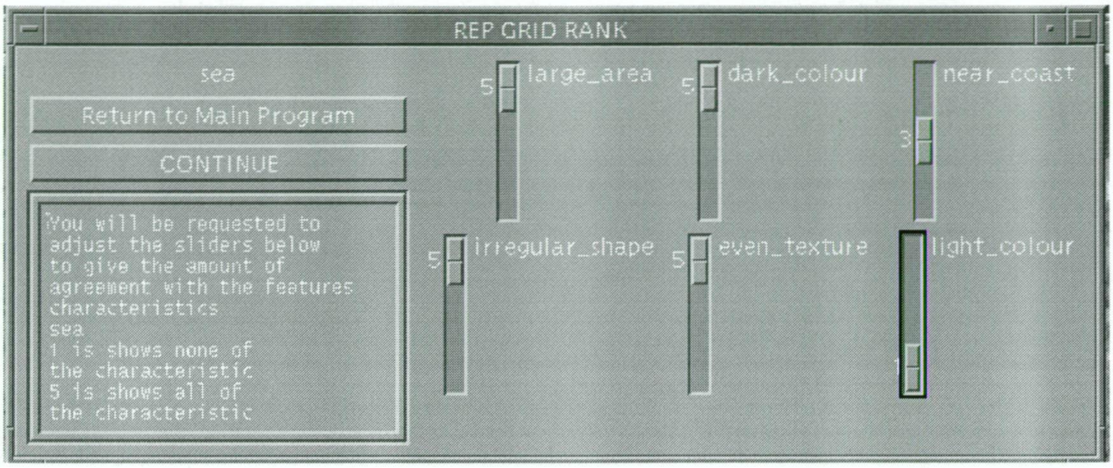

Figure A.18 Ranking window with slider bars to input ranking

Once all the discriminators have been added, each object or feature named by the user on the first screen is compared with each of the discriminators. This is done with the above screen. The order of adjusting the slider bars is unimportant.

In the above example the feature **sea** is analysed. Since the discriminator **early_cropping** is irrelevant in this particular case, the slider bar is left at 3. Any slider bar left at 3 is not regarded as significant when the **generate rules** option is selected.

```
forest
|  sea
|  |  onions                    FEATURES
|  |  |  peas
|  |  |  |  potatoes
|  |  |  |  |  poppies
5  5  1  1  1  1  large_area
1  1  4  5  4  4  high_albedo
1  1  5  5  5  5  small_area              DISCRIMINATORS
1  1  5  5  5  5  regular_area
1  5  5  5  2  4  even_texture
3  3  4  5  1  2  early_cropping
```

Figure A.19 Repertory Grid results

The result of the above steps is to produce the grid shown in figure A.18. The setting of each of the slider bars becomes a value in the appropriate cell of the grid. This grid can then be stored as a text file for future manipulations, analysed for clustering of both features and discriminators or used to generate rules.

Rules are displayed on at a time on screen for user verification. These can then be either accepted, modified or rejected. The final rule base is then saved as a text file. The following three example rules are transcribed from the contents of the file which contains the output of the above grid. In this case all the rules were accepted without modification.

RULE_1
IF large area
AND NOT high_albedo
AND NOT small_area
AND NOT regular_area
AND NOT even_texture
THEN OBJECT = forest

RULE_2
IF large area
AND NOT high_albedo
AND NOT small_area
AND NOT regular_area
AND even_texture
THEN OBJECT = sea

RULE_3
IF NOT large area
AND small_area
AND regular_area
AND even_texture
THEN OBJECT = onions

A.9 POINT DATA TOOL

The Point Data Tool allows an image interpreter to create training data at various locations across an image. This can be done in either automated sampling mode or manual mode. In automated mode the system asks the user to select the sampling density using slider bars. The system then samples the image set using the sampling grid and requesting the user to name the sample points.

In manual mode the system operates in a similar way but allows the user to select the points to be named. In practice the user started with the regular grid then filled in unsampled features in manual mode.

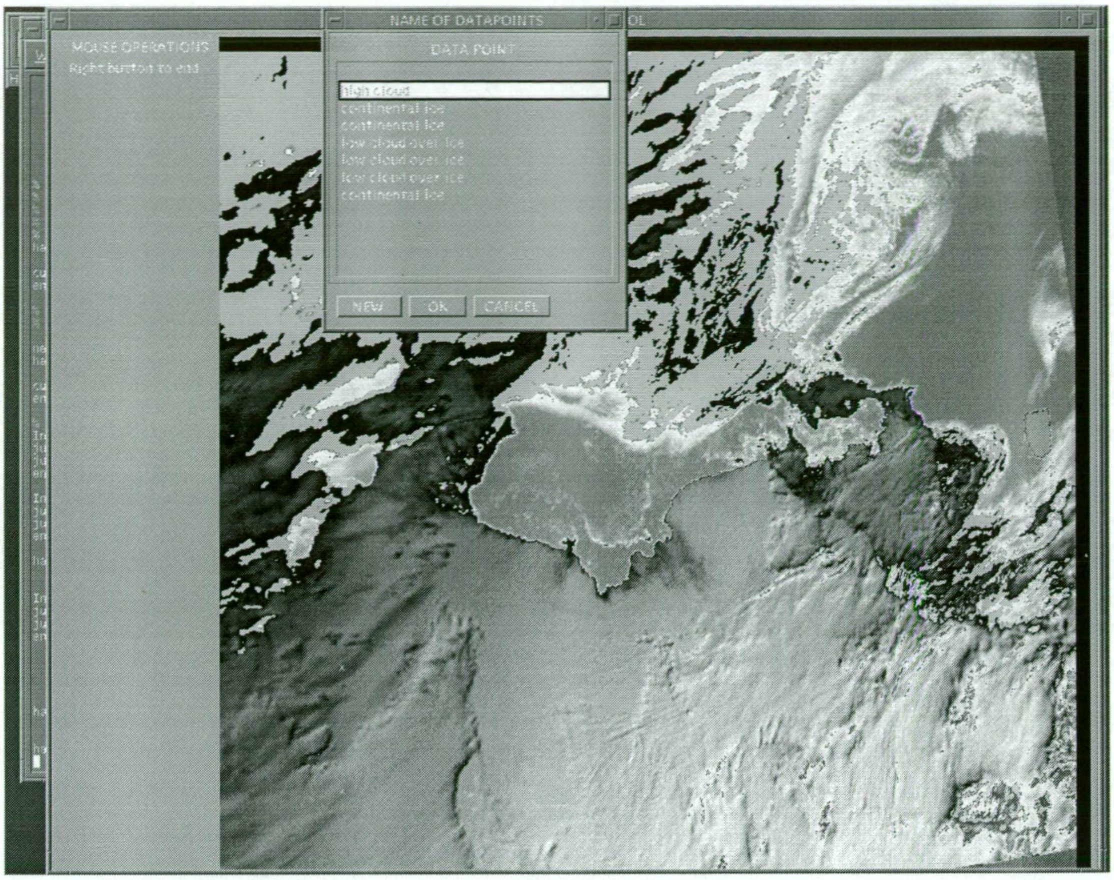

Figure A.20 Point data tool showing sampling points designated by X on the image and the naming window (NOAA VHRR image of Casey taken on 26th February 1998)

As each point is selected either by the system or the user, the user is requested to name it. To do this a user can select a previously entered label by highlighting it and selecting **OK** (as in Figure A.19) or select the **NEW** button to enter a new label for the sample point.

Choosing the **Cancel** button will allow the user to:

- change image bands,
- change to user selected sampling or
- end and save the results of sampling

A copy of the output from sampling Figure A.20 is:

BAND						Location	
1	**2**	**3**	**4**	**5**	**class**	**X**	**Y**
141,	127,	4,	85,	85,	high cloud,	127,	128
153,	137,	48	79,	80,	continental ice,	254,	128
166,	146,	27,	75,	75,	continental ice,	381,	128
159,	137,	13,	77,	77,	low cloud over ice,	508,	128
178,	155,	17,	80,	80,	low cloud over ice,	635,	128
143,	122,	12,	76,	79,	low cloud over ice,	762,	128
139,	115,	41,	67,	68,	continental ice,	889,	128

APPENDIX B. MIRC DATA BASE

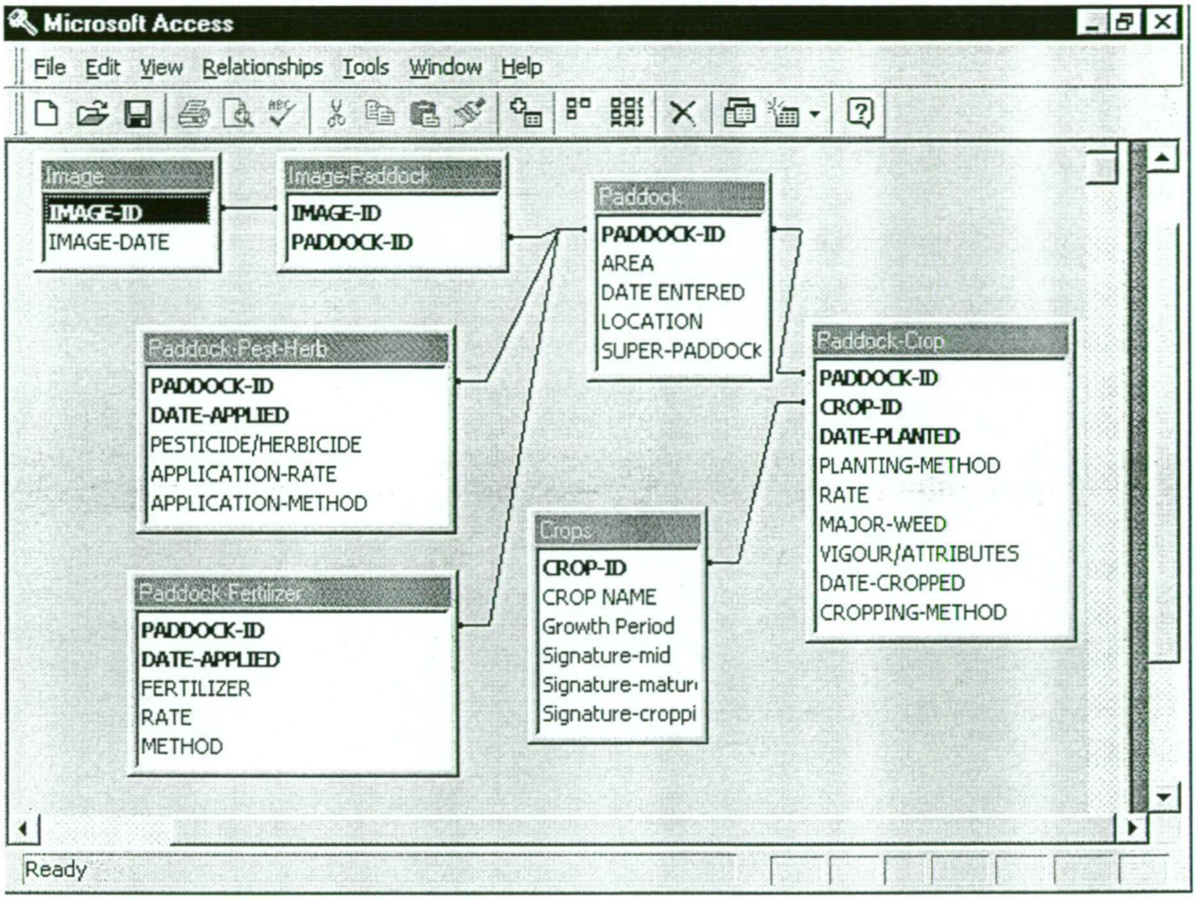

Figure B.1 MIRC database

The above screen shot is of the entity – relationship table of the relational database for the MIRC 2 project. MIRC 2 is a larger scale investigation funded by a HRDC (Horticultural Research and Development Corporation) grant with the same aims of the MIRC project described in the thesis. It contains data for paddocks in two areas; the first being on Tasmania's North West Coast centred on the Devonport area, the second being in the South East of the state centred on the Coal River Valley.

Data has been entered for the 1997 / 98 growing period, but it is unlikely that any analysis of the data will be carried out before three growing seasons of data have been entered.

Details of each of the tables and their attributes shown in the above table are as follows:

Paddock – Crop Table

This table provides the details of a crop planted in a particular paddock at a particular time along with an indication of crop health including weed infestation. The first three fields are key fields and unique temporal identification. Details are:

Paddock-ID – A field identifying individual paddocks and linking back to individual paddock details

Crop-ID – Identifies the type of crop in the paddock and links to individual crop details

Date-Planted –An estimation of the date when the crop was planted.

Planting-Method – Optional field detailing the way planting was done, for example type of seed drill and depth.

Rate – Optional field detailing rate of sowing.

Major Weed – Optional field containing major weed(s) including regrowth of previous crop.

Vigour/Attributes - Optional comment field detailing crop health.

Date-Cropped – Date the crop was harvested. This may only be an approximate date.

Cropping-Method – Optional field containing information on the type of harvester and any details on the methods used.

Image Table

This table contains a unique identification number for each image used in the study and the date it was obtained

Crops Table

This table relates a Crop-ID to an actual crop name. This table once created requires little maintenance, except to add new crops. Details are:

Crop–ID – Unique identifier of a crop.

Crop Name - Crop name – generic name rather variety (for example potatoes, not Russet Burbank).

Growth Period – Approximate period since planting in weeks.

Signature-mid – Spectral signature (including band) of a typical example of the crop midway through its growth cycle.

Signature-mature - Spectral signature (including band) of a typical example of the crop when mature (generally flowering).

Signature-cropping - Spectral signature (including band) of a typical example of the crop at harvest, but before being cleared.

Paddocks Table

This table contains features of individual paddocks. This table may require maintenance as new paddocks may be created by either land clearing or subdividing an existing paddock. Details are:

Paddock-ID – Unique identifier for each paddock in the study.

Area – Estimated area of paddocks in square meters.

Date Entered – Date the paddock was entered into the system. This is particularly important for subdivided paddocks which may only exist for a single growing season and are therefore temporal in nature.

Location – Geographic location of paddocks centre (if available).

Super-Paddock – If this is a subdivided paddock, this field contains the Paddock-ID of the undivided parent paddock.

Image – Paddocks Table

This table links paddocks to images. Images which are affected by cloud will not show all paddocks. The table is designed to be used to aid filtering when temporal data is required. Hence images which contain a paddock required as part of the study can be identified.

Paddock–PestHerb Table

This table is used to keep track of applications of pesticides and herbicides to individual paddocks.

Paddock-ID – Paddock to which pesticide or herbicide has been applied.
Date-Applied – Approximate date of the application.
Pesticide/Herbicide – Name of the pesticide or herbicide used.
Application Rate – Rate in grams per square meter at which the pesticide or herbicide was applied.
Application Method – Method by which the herbicide or pesticide was applied (Includes aerial and boom sprays).

Paddock – Fertilizer Table

Thiis table is used to keep track of applications of fertilizers to individual paddocks.

Paddock-ID - Paddock to which pesticide or herbicide has been applied.
Date-Applied - Approximate date of the application.
Fertilizer - Name of the fertilizer used.
Rate - In terms of grams per square meter at which the fertilizer was applied..
Method - Method by which the fertilizer was applied (Includes aerial and boom sprays).

APPENDIX C. SAMPLE INTERVIEW TRANSCRIPT

The following is an edited extract from an interview with an image interpretation expert. This particular expert had skills in both satellite image interpretation and GIS. The interview was conducted while the KAGES toolkit was still being developed, so some of the suggestions were incorporated in the system

Q You have now seen a demonstration of the tool and used some of its features, do you think it is a useful tool?

A The answer is yes. It contains a lot of features which are not part current GIS packages and require quite a bit of work if you are going to do it with current image processing tools.

Q Which aspects of the tool are most useful?

A The method you use for sorting out the spatial relationships between two objects is a good idea. Particularly the way it allows you to see all relationships.

The repertory grid is an interesting idea. I can see that it can give an alternative method of gaining knowledge now the way it comes up with rules has been explained.

The region of interest tool is also a quick way of grouping objects, although I think it may need more work to be completely useful. It doesn't actually replace everything in the group with a single object.

Q What does the tool provide your current system does not?

A My current tool basically works at what you call the scene primitive level. Most of the things you have shown me apart from that are not available in my current system.

Q What would you like to see the tool provide that it does not?

A The histogram display is good but it would be nice to be able to adjust thresholds using the histogram as well as slider bars, both in the initial definition and during modification.

Make sure the full range of image processing tools are available. This is one of the main shortcomings of the current range of GIS packages.

It would also be useful if the system could generate more statistics. For example the degree of overlap of the boundaries of two objects. That is if a polygon is of a particular type and it overlaps another polygon, how big is the polygon formed by the overlap, and what percentage of the original polygons does it represent.

Q I had a similar comment about the Region of Interest Tool. Do you think it is useful to record information on the make up of objects within a subregion?

A Definitely, also it would be good if the Region of Interest Tool could be expanded to allow automatic replacement of a group of objects by a single object.

Q Which aspects of the system duplicates what you already have?

A The main one is the Per Pixel tool. The spectral stuff has been pretty much fully developed in GIS systems, as a result your system does not really offer any advance. But if you don't have a starting point which defines scene primitives, I supposed you do need something of that type before you can do very much more. I would look at importing information from a GIS which already has labelled polygons identified from spectral signatures

Q Which aspects of the tool don't you like?

A Apart from the stuff I mentioned I would like to see the tool provide, a couple of features could be improved. The part where it shows rules then allows you to delete lines. That is very restrictive. A better editor is needed there. I would like to be able to do modifications more easily