



Eliciting Knowledge with Visualization - instant gratification for the expert image classifier who wants to show Rather than tell

Paul Crowther and Jacky Hartnett
 Artificial Intelligence and Spatial Systems Research Group
 Department of Computing
 University of Tasmania
 P.Crowther@utas.edu.au

Presented at the second annual conference of GeoComputation '97 & SIRC '97, University of Otago, New Zealand, 26-29 August 1997

ABSTRACT

GIS systems the world over are awash with data that experts can classify visually. This process is time consuming and costly. Expert Systems have been built which attempt to at least pre-classify images and hence speed up the process. To build these systems it is necessary to elicit information from the human expert classifiers in order to assist the classification of these many hundreds of images. Traditionally this knowledge has been captured through interview and protocol analysis. However, this required either the expert classifier to describe verbally what they were seeing or the expert systems developer (knowledge engineer) to interpret what they were being shown.

To overcome this problem, a visual knowledge acquisition tool, KAGES (Knowledge Acquisition for Geographic Expert Systems), was developed. Impetus to the development of this tool was given to our group by the need to classify many remotely sensed images of Antarctica in order to provide information on global climate change and Southern Ocean currents.

1 INTRODUCTION

This paper describes a tool for acquiring knowledge from expert image interpreters by allowing them to demonstrate their expertise. To do this the tool must work quickly in order to provide the user with rapid feedback on the knowledge acquired. KAGES was developed on a work-

station using Research Systems Inc's (1994) IDL image processing package to overcome this problem. The matrix handling features of IDL have allowed a fast and efficient system to be developed which allows a human classifier to identify features of interest by pointing or drawing on an image displayed on a computer screen. KAGES captures the knowledge underlying the identification of these features in the form of production rules. These can include rules that describe the spatial relationships between two of the identified feature types as well as rules that identify features in terms of their spatial relationships within a group of features.

KAGES is a toolkit which provides a series of knowledge acquisition techniques including an interview manager, several graphical acquisition tools and a rule editor. The captured classifier knowledge is held in a series of knowledge bases which are then consolidated and checked for consistency and redundancy. The result is a knowledge base which can be viewed and reviewed by the human classifier.

Without exploiting the speed of current workstations, this type of computing would not be feasible since interactive graphical knowledge acquisition of visual knowledge is computationally expensive. KAGES operates on all bands of a satellite image and overlays the results on a composite image. The data structure which is being manipulated in memory is therefore an array of 1000 X 800 X 6 in the case of NOAA images. This can be even larger when more





than one version of the composite image is required, adding extra dimensions to the array structure. Since the operation is pixel-based and the results are held as a raster format, there is a large memory requirement for individual objects as well. Users require instant feedback about the results of the operation of the various tools to allow for quick verification. As a consequence, zoom and scaling functions require fast processor speed as these are required to operate on all dimensions and objects. Processor speed is also at a premium when the spatial analysis tool is operating, as all spatial relationships, including overlap, proximity and orientation, are captured.

Currently KAGES operates as a serial processor, but many of its functions are well suited to parallel processing. This is particularly true of the spatial relationship and the knowledge base consolidation tools which have a number of distinct independent functions.

2. VISUAL KNOWLEDGE

Knowledge is understanding, awareness, or familiarity acquired through education, or experience, anything that has been learned, perceived, discovered, inferred or understood and the ability to use information (Nagao, 1988).

Maps present knowledge naturally occurring in three dimensions in a two-dimensional graphic form. However, maps are produced from information from ground (and sometimes underground) survey, from images produced by sensors on aircraft or satellites, and from photographic images. Each of these could be regarded as another dimension. The problem, then, becomes one of representing n dimensional knowledge in a two dimensional form. 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. By using expert system approaches it may be possible to make more use of all the dimensions of information available and the interpretation of that into the knowledge of multiple domain experts.

Acquiring knowledge from multiple expert classifiers also introduces another 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. This interpretation may also be culturally based (Clementini et al, 1993). However, the visual definition of the feature in a graphical form is more concrete and less subject to interpretation. What it looks like defines it, rather than what it is called.

Spatial knowledge here will be defined as knowledge of entities in two dimensions (as in a map) as distinct from three dimensions where research is more directed towards vision and recognition.

McKeown et al (1989) identifies five types of knowledge used to identify specific spatial relationships. The five types are:

- Type 1 Knowledge: identifies scene primitives where a primitive is a readily identifiable object such as a road or a building.
- Type 2 Knowledge: is the knowledge of the spatial relationships between the scene primitives, for example buildings are next to roads, 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.

This classification has become the basis of the tools and techniques described in this paper. Further sub classification has been necessary due to the characteristics of the object(s) being investigated. Hence, lines have different characteristics from areal scene primitives at the Type 1 level.

3. A VISUAL KNOWLEDGE ACQUISITION SYSTEM

KAGES consists of tools to capture knowledge of the first three types using visual tools. This overcomes some of the problems associated with verbal descriptions and



definitions. Type 5 knowledge is addressed by a module which combines knowledge gained from different image interpreters, different images and different sessions into a consolidated knowledge base. The system interacts with the user in order to resolve inconsistencies. This tool will not be discussed in detail in this paper.

An interview manager based on repertory grids and personal construct theory (Kelly, 1955) is also provided (Crowther and Hartnett, 1996). This tool was developed to test users' reaction to both a visual and a text based tool (although it uses visual cues) and provide an alternative knowledge acquisition technique to deal with non visual knowledge. This was the tool for experts who would rather *ëtelli* than *ëshowí*.

3.1 TYPE 1 TOOL

Domain primitives are features to which an image interpreter can point and give a name. They may be point, line or areal features. Each requires different processing as each has different properties. The most simple features are point

features which are subpixel in nature and which are not generally identifiable by pixel threshold signatures. These features generally are a fixed point such as a building. The expert interpreter identifies these objects by selecting an image band and pointing to their position. The name of the point feature, its location and the identification of the image on which it was defined are stored.

Areal features are two dimensional objects which may be permanent features (such as a lake) or transient (for example the contents of a field or the extent of sea ice). The user first selects the band or composite bands they use for identifying a feature. The feature is then described by an expert pointing at it and setting pixel threshold values. KAGES allows this by the use of slider bars. All pixels within the thresholds which are contiguous are then grouped (Fig 2). If the grouping agrees with the expert's idea of the extent of the feature, it is named and information about thresholds and its minimum bounding rectangle (MBR) (Chang and Jungert, 1996) is stored.

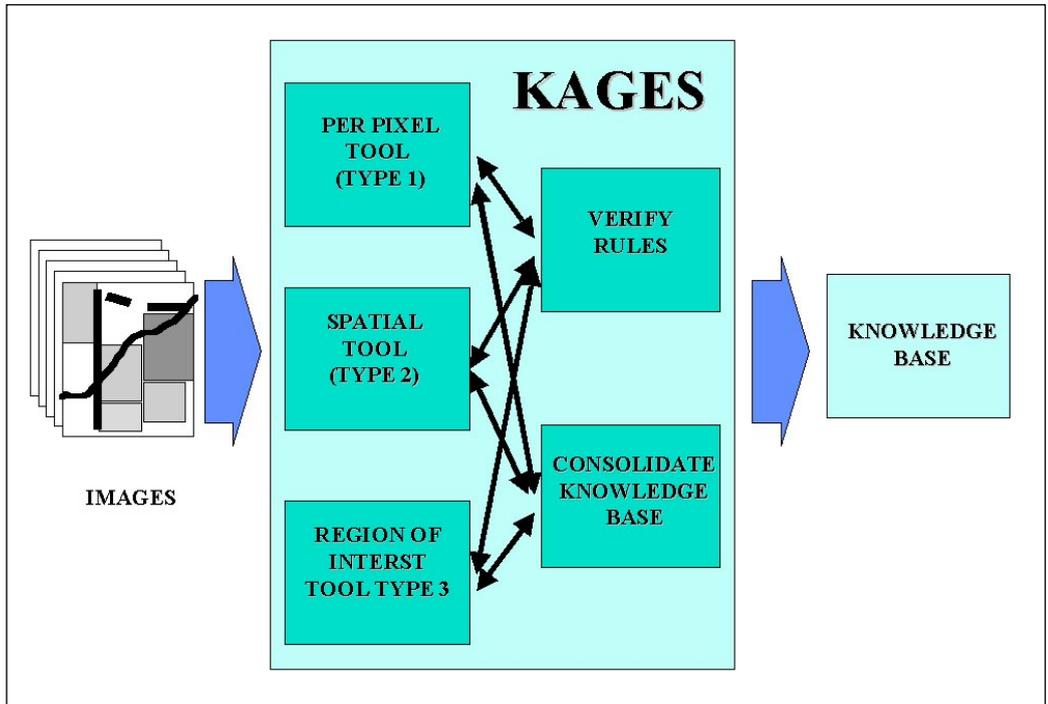


Figure 1: The KAGES system



The expert user can identify other occurrences of the object type on the training image (although it could be on other bands or band combinations). At any stage the user can ask for production rules defining object types to be generated.

Lines have proved to be the most difficult objects to deal with (Crowther and Hartnett, 1997). These one dimensional objects can be identified by either a line following algorithm or by line tracing. The later technique is necessary for lines such as municipal boundaries and other cadastral data. Information about lines is stored as either a set of raster points or as a vector, depending on how the line was acquired.

A user can choose as many examples of an object as they wish on whichever band or combination of bands they like. Once this has been done, the module which consoli-

date rules can be called. This filters the knowledge base combining rules with the same antecedents. The resultant rules can be fired individually with the results overlaid on band 1 (even if the objects were identified on other bands), or the entire Type 1 knowledge base can be applied.

3.2 TYPE 2 TOOL

The Type 2 tool is the spatial relationship tool which allows a user to determine the relationship between two objects (Fig 3). The two objects are shown named on the default band of an image set (usually band 1) with their MBRs. KAGES then determines the relationships between the two objects. Objects fall into the three scene primitive types and procedures have been developed to deal with relationships between them. These relationships fall into the following categories:

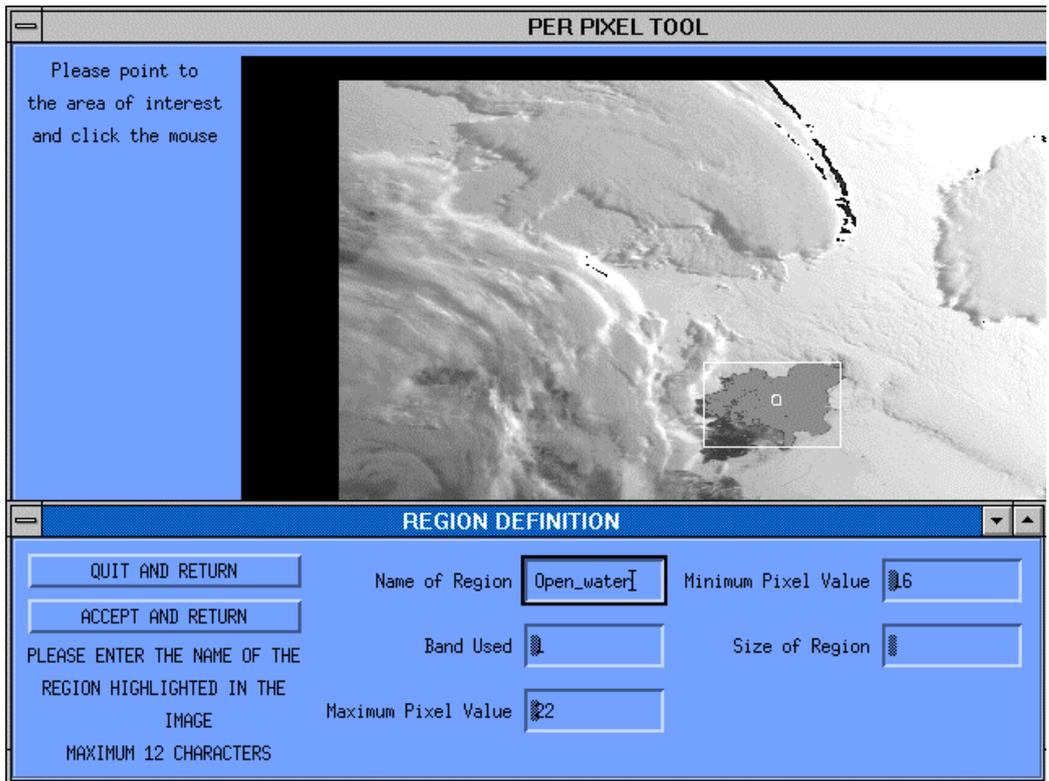


Figure 2. The per pixel Type 1 tool being used on an areal object. The object being defined in this Band 1 NOAA image of Vincennes Bay is open water. The objects Minimum Bounding Rectangle (MBR) is shown with the centroid of the MBR being marked by an O. Information on threshold values and the user entered feature name are shown in the dialogue window.



point : point	point : line	point : area
line : line	line : area	area : area

The type of relationships fall into three main categories based on Allen intervals as used by Egenhofer (1991). In all cases the following relations are calculated:

- Proximity
- Degree of overlap
- Orientation are calculated.

The spatial relationship module first determines what the two types of objects are it is dealing with. The system then determines all possible relationships between the objects. For example in the case of two lines being selected:

- Determine if the lines touch
- If they touch
 - Determine type of intersection
- Otherwise
 - Determine directional relationship
 - Determine degree of parallelism

Operationally, the user is required to select two objects, which are related, from a menu and then check the results of the relationships generated by the system. This is done by the use of a simple rule editor which allows a user to remove clauses which are chance relationships and not deterministic of the relationship between the two objects.

As a further example, if a point object is compared with an areal object, the system determines the distance of the

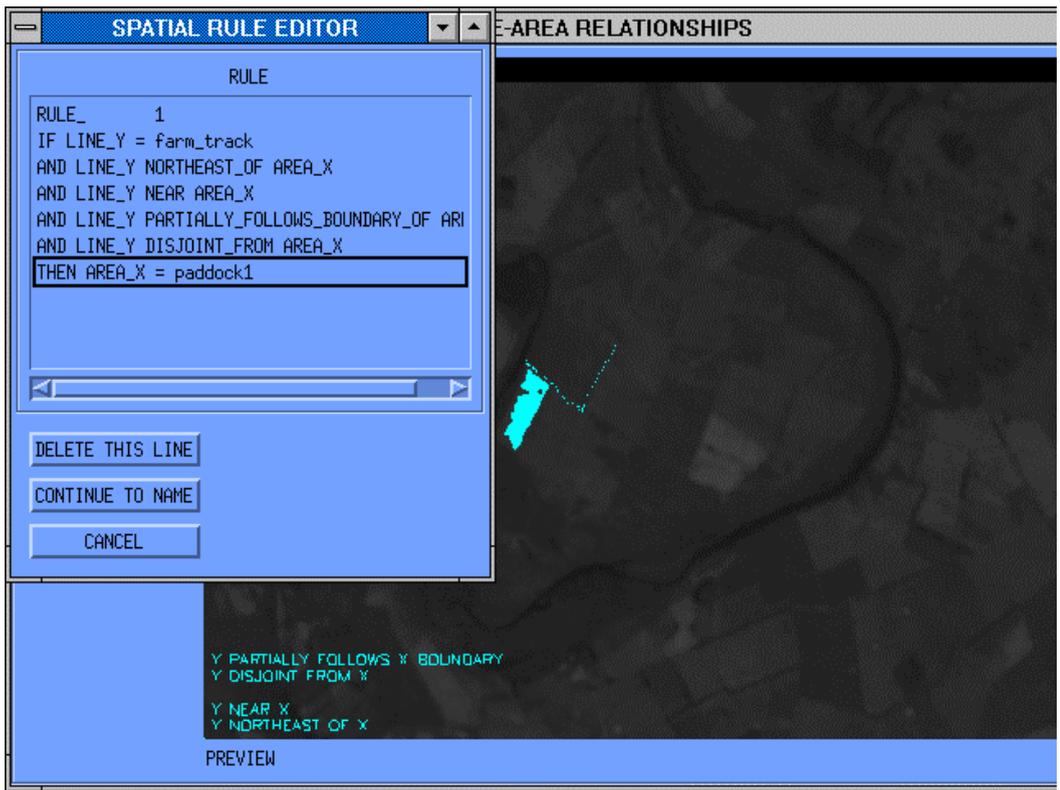


Figure 3: The spatial Type 2 tool being used to determine the relationships between a line object (a road) and an areal object (field) in a Landsat image of an area near Perth in Northern Tasmania. These are both functional primitives identified by the Type 1 tool. The relationships are shown on the image and as a rule in the rule editor window.



point from the centroid of the area, the direction of the point from the centroid (equivalent to orientation) and the relationship of the point to the boundary of the area. Possibilities are:

- Outside area and disjoint
- Outside area and touches boundary
- On the Boundary
- Inside area and touches boundary
- Inside area

3.3 TYPE 3 TOOL

The Type 3 tool allows a user to group objects defined by the Type 1 tool. Often these type of features have no natural boundary and are traced by an image identifier. This tool allows a user to trace a boundary directly on the image band they have chosen (Fig 4).

All objects which fall in the region of interest, whether fully or partially, are shown. Hence a river flowing through the region would be displayed, even though its points of termination may fall outside the region. Objects defined using other bands are also shown.

The results of the objects identified by the tool are displayed in a window labeled 'Check Membership' which can be manipulated by the user to remove objects which are not distinctive of the region. Once the user has completed this task a naming window is displayed. This name, the members of the region and the band used together with image identification are then stored.

4. TOWARDS A TYPE 4 TOOL

The aim of KAGES is to use an image as training data to create a rule base which can then be applied to other im-

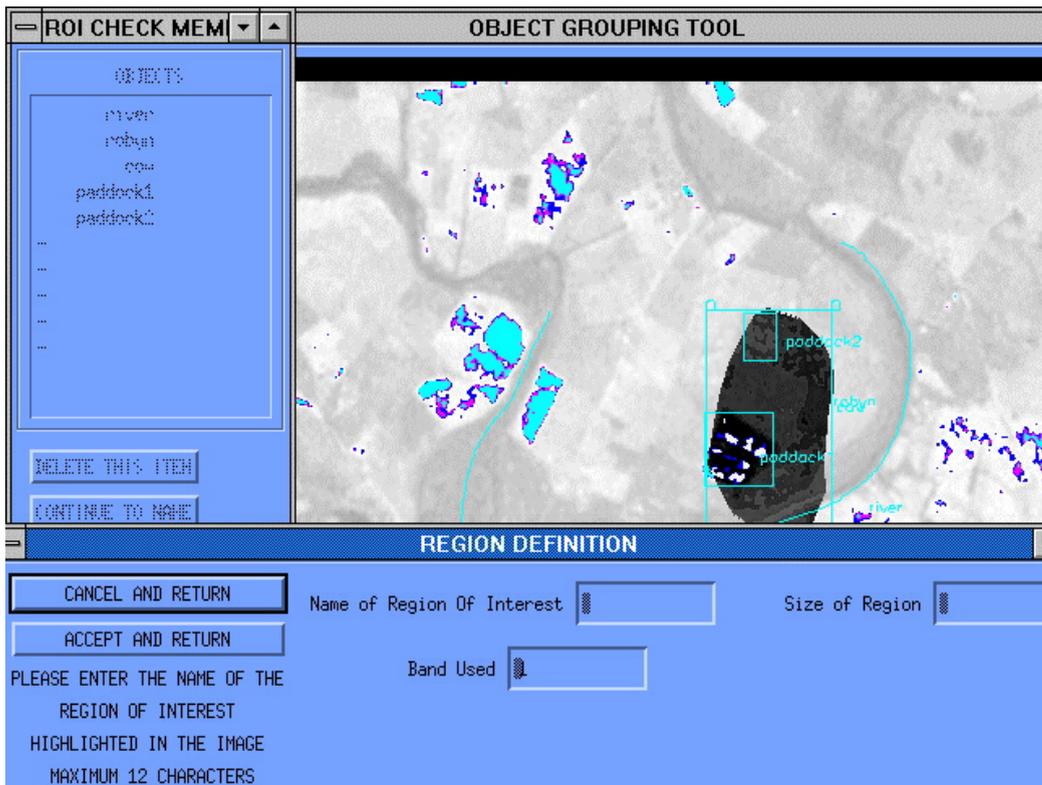


Figure 4: The Type 3 Region Of Interest tool being used to pick up point, area and line objects using band 1 of a Landsat image. The limits of the Region Of Interests MBR are marked with a (. Paddock1 and Paddock2 MBRs are drawn with their centroids named. Points (Robyn and the cow) are named with the location of the point at the left of the name. Line objects, in this case the river, are named near their centre point.





ages. It was not intended to be an image classifier. As part of the development, a verification tool which applies rules has been built and this will do classification. This tool operates in three stages:

- Apply Type 1 rules to the image
- Segment the image and label individual objects
- Apply spatial (Type 2) rules to the segmented image.

The result is a classified image which shows individual classified objects, areas for which no rule has fired (and hence are unclassified) and areas where conflicting rules have fired. These last two highlight image features which require further investigation by the KAGES tool.

A side effect of this development has been to allow this feature to be used as an image analyzer. As once the expert user is satisfied with the performance of the generated rules on the training image, those rules can then be applied to other images. Development of this tool turned KAGES from being just a knowledge acquisition tool into a tool with image analysis capabilities which will provide knowledge of complete scenes (Type 4 knowledge).

5. FULLY AUTOMATED OR HUMAN ASSISTED

Experience with Icemapper (Williams et al, 1997), which was developed using rules acquired by traditional interview techniques lead to the development of a system which provided a first best guess and which could then be adjusted by the image interpreter. Generally after three iterations a properly classified image was produced. This cut down the time of development of a classified image from about an hour (fully manual) to around ten minutes. This human directed system was preferred to a fully automated version based on neural networks (Kilpatrick and Williams, 1995).

KAGES was developed with this human assisted ethos in mind. The amount of time taken for a human to tell (in the form of interviews) and the knowledge engineer to interpret (into rules) took several months for Icemapper. This was a classic example of the "knowledge acquisition bottleneck" (Gaines, 1988). To speed the process up, the tool

that allowed an expert to show rather than tell their expertise was developed.

The tool also gets around the problem of feedback from the knowledge engineer in that by being able to apply the rules which they have developed, the expert user can see within four seconds, the results of applying those rules. If necessary, changes can then be made and the rules reapplied.

6. KAGES AND THE NEED FOR SPEED

Speed is essential in two areas. First there needs to be sufficient processing power to handle the size and number of arrays containing the images. These are generally multi-dimensional with each band taking up a two dimensional array and the complete image set and any result images being other dimensions. For example a NOAA image of 1000 by 800 pixels with 5 bands and an array of a similar size for manipulating and displaying the results requires a minimum of 4.8 megabytes of memory. This is an underestimation of the actual memory requirement as composite bands will add further dimensions to the data structure as will other arrays which are used to store individual features and intermediate results of processing. A user requirement to compare the training image with other images will double this. The key factor is that the processing must take place sufficiently quickly for the user to respond to the feedback iteratively. Fortunately memory is cheap and the system has been run successfully on a 16 Mb machine.

The description of spatial relationships is one area where the implementation of parallel algorithms has potential to further speed up the system. We have identified the algorithms used in the current serial version of the system to have Multiple Input Multiple Processing (MIMP) (Ding and Densham, 1996) characteristics. Furthermore, to apply spatial rules for verification, all occurrences of one class of objects are compared with all occurrences of a second. In this case Multiple Input Single Process (MISP) could be used as the same procedure is used on multiple occurrences of an object type. This analysis provides exciting possibilities for future implementation.





For now, IDL routines, which allow manipulation of arrays, speed up data processing. A speed increase of about sixty (10 minutes to 10 seconds) was achieved when spatial relationship routines were converted from traditional loop processing.

7. APPLICATIONS

The system is currently being used to generate rules for comparison with those already manually determined and placed in the Icemapper sea ice identification system. The rules are comparable but are much quicker and easier to both generate and modify. Of interest, several of the rules in Icemapper can only be reproduced using the repertory grid interview tool as they are not visual in nature.

A second system currently under investigation and development involves crop identification in the North West of Tasmania. In this system the rules generated by KAGES are applied using the module currently used for verification. In this case the images are supplied by a GIS and the results will be ported back into the GIS. The results, which are so far incomplete, are being compared with the results gained from clustering techniques.

8. CONCLUSIONS

To capture visual knowledge for the use in an expert system coupled to a geographic information system, a primarily non text graphical tool is necessary. Such a visual knowledge acquisition tool has to have the following characteristics:

- It must have a graphical user interface
- It must be intuitive for the user to operate.
- It must acquire the user's knowledge by directly capturing their actions
- It must operate in real time and give instant feedback
- The rules generated must be visually verifiable

KAGES is designed to meet the user requirements but relies on powerful hardware to function successfully. The system is designed to be cross platform and work on both personal computers and workstations. Current generation personal computers with at least 16 Mb of RAM give sat-

isfactory performance. On workstations the system has given good results and has a high level of user satisfaction.

9. REFERENCES

Chang, S. and Jungert, E, (1996), Symbolic Projection for Image Information Retrieval and Spatial Reasoning, Academic Press, London, Ch. 2.

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

Crowther, P. and Hartnett, J. (1996), Using Repertory Grids for Knowledge Acquisition for Spatial Expert Systems, Proceedings of the 4th Australian and New Zealand Conference on Intelligent Information Systems. Adelaide, pp 14 - 17.

Crowther, P. and Hartnett, J. (1997), 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, pp 167 - 176.

Ding, Y. and Densham, P. J. (1996), Spatial Strategies for Parallel Spatial Modelling, International Journal of Geographic Information Systems, Vol. 10, No. 6, pp 669 - 698.

Egenhofer, M. J. (1991), Reasoning About Binary Topological Relations, in Advances in Spatial Databases, 2nd Symposium SSD '91, Gunther, O and Schek, H.J. (eds), Springer Verlag, pp. 143 - 159.

Gaines, B. R. (1988), An overview of Knowledge Acquisition and Transfer. In: Knowledge-based Systems Vol. 1., Gaines, B. R. and Boose, J. H. (eds), Academic Press, pp. 3 - 21

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



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

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

McKeown, D. M., Harvey, W. A. and Wixson, L. E. (1989), ëAutomating Knowledge Acquisition for Aerial Image Interpretationí, Computer Vision & Graphics, Vol. 46, pp. 37 - 81.

Nagao, M. (1990), Knowledge and Inference, Academic Press, San Diego Ch. 1.

Research Systems, Inc. (1994), Interactive Data Language Version 3.6.

10. ACKNOWLEDGMENTS

The expert system development work reported in this paper has been supported by an Antarctic Science Advisory Committee Research Grant.