

A Model to Support the Integration of Image Understanding Techniques within a GIS

Mark Gahegan and Julien Flack

Abstract

Traditionally, geographic information systems (GIS) have used data generated from remote sensing, but only after the data have been preprocessed in some way to provide a suitable classification. This leads to several weaknesses; the resulting framework is inflexible and cannot support multiple interpretations of the same area. A new model for a GIS is introduced that includes a set of image understanding methods, and expert knowledge governing the application of these methods. The methods employed are chosen automatically to best emphasize the types of features (rivers, fields, forests, etc.) that the user is currently investigating, and the type of imagery available. An implementation of the model is described along with details of the knowledge structures and image processing methods used. Examples are included to show the resulting adaptive nature of image interpretation and subsequent inclusion of feature descriptions into the GIS.

Introduction

The need to form an integrated environment for remote sensing data, to encompass both image interpretation and its subsequent use for geographic analysis and modeling, has received much interest of late (Zhou, 1989; Ehlers *et al.*, 1991). As a consequence, research has been directed towards the provision of a data model for this process, often by directing the output from the image interpretation to the input of the geographic database, by utilizing some feature extraction process.

The motivation for this work can be summarized by considering the broad spectrum of geographic information system (GIS) users. At one extreme, non-experts would like to be able to use remotely sensed data in GIS applications but do not understand the process involved in changing the image data into a more suitable form. Ideally, their needs should be supported, even if they do not wish to understand or be involved with the process. At the other extreme, expert users would like the ability to control the methods involved in extracting the spatial descriptions for use within GIS, and to experiment with different strategies and study their effect on the outcome.

Existing GIS incorporate remotely sensed data most usually in the form of a single classified *overlay* of a scene, formed according to established pixel-based spectral classification techniques. The overlay is constructed externally to the GIS, often with the aid of an image processing system. A good summary of this process is given by De Cola (1989). The overlay provides the description for the geographic objects of interest (referred to as *features*) after it is segmented into (geometrically) distinct regions. A single overlay often represents the *only* interpretation of the raw data available to the GIS and is therefore not ideally suited to some types of

problems. The classifications can suffer from being either too specific in their application, or too general. While it is relatively simple to produce more than one scene interpretation using existing image processing tools, it is a far from simple matter to import more than one classification into a GIS, because it causes a conflict over which one should be used. Each distinct classification can produce a different spatial description for each feature, in terms of both estimated land cover and shape. Existing GIS support only one spatial description of each feature; worse still, each feature actually draws its existence from this single description.

Proposed Enhancements

In a data modeling sense, the task of linking semantically the results of a scene interpretation process to geographic features (known to the GIS) is called *binding* (Korth and Silberschatz, 1991). Binding is the general process of joining together data from different levels of abstraction within the data model. The time at which binding occurs affects the flexibility of all database systems. Early bindings give rise to poorer flexibility whereas late bindings increase flexibility, usually to the detriment of performance¹. Here the binding process takes a set of pixels in the raw image data and casts it to somehow represent the spatial description of a feature. Depending on the application, this process may include a vectorizing stage where the boundary pixels are converted to vector form.

The main goal of this work is to allow a single feature in the GIS to have more than one geographic description, and, in addition, to put off the task of choosing a description until the current user task is known; that is, with late binding at execution time. By so doing, the description employed can be made highly suitable to the current task.

Example of Adaptive Feature Extraction

As an example of adaptive feature extraction, consider a GIS applied to the analysis of a remotely sensed Landsat TM scene containing farmland and a road network. In order to

¹The flexibility issue is similar to that of (say) a network database when compared to a relational database. In the relational model, the bindings between the various layers are not fixed until query execution time, giving rise to greater interpretative power because there are no restrictions on how attributes might be combined. In contrast, database models such as network or hierarchical support only fixed, pre-defined access paths, thus limiting the type of queries that can be solved.

Photogrammetric Engineering & Remote Sensing,
Vol. 62, No. 5, May 1996, pp. 483-490.

Department of Geographic Information Systems, Curtin University, P.O. Box U 1987, WA 6001, Australia.

0099-1112/96/6205-483\$3.00/0
© 1996 American Society for Photogrammetry
and Remote Sensing

identify the paddocks where a particular crop is grown, the image data could simply be classified, then individual paddocks could be extracted and passed to the GIS. Now instead consider a routing task involving the road network. The network is unlikely to show up well on the classified scene. The detection of features such as roads is problematic in lower resolution remotely sensed data; the roads often occupy a width of less than one pixel in the image; therefore, they do not form regions (groups of connected pixels). Their presence may be indicated in the raw image only as an aliasing or edge effect. Such features are here said to be *linear*. After performing a classification of the scene, linear features often become even more obscure and can be lost altogether, because they are difficult to characterize accurately with a training set. In order to emphasize them, we must apply an entirely different set of tools to the task, involving edge detection and line following. This in turn produces a different interpretation of the raw data where these features will be emphasized, and other features may be less well defined or even absent. The point here is that, until the nature of the current task is known, it is difficult to provide a suitable interpretation of the scene.

The result of applying scene understanding techniques to highlight a particular type of feature is here called an *image-view*, or *view* for short. Note that a view is not always the same as a traditional classification, but, rather, is more akin to the database notion, i.e., as a way of *viewing* the underlying data. It is from a view that the spatial description of certain individual features is extracted.

Multiple Spatial Representations

For many applications it is only necessary to provide a single spatial description for each feature of interest, but there are three special cases where multiple representations are appropriate.

The first is to support cartographic generalization so that the appearance of features can change as the viewing scale is changed. For example, an iconic representation of a building may be used until the real shape of the building is discernible, given the current display scale and resolution of the output device. A detailed description of this type of feature behavior within a GIS is given by Roberts *et al.* (1991). Support for scale-dependent behavior is becoming available in some GIS; for example, SMALLWORLD supports a change in the *presentation* of features through a range of display scales.

The second is in situations where the spatial descriptions of the features being considered are subject to uncertainty, either because the data are not accurate enough or the feature is by nature not well-defined. An example is where land-use change between forest and marginal land is gradual and the boundary identified for each individual feature is likely to be highly dependent on the different techniques used in its production. In existing GIS, the user is typically left with no choice as to how such arbitrary decisions are made regarding feature extraction, and worse still the users may not even be aware of the errors and uncertainties to which the methods used can lead. There is a worrying tendency among GIS users to believe that once a feature has been delineated by a series of vectors, then it is entirely correct, unambiguous, and reliable (Davis and Simonett, 1991). By supporting multiple spatial descriptions, expert users can be made aware of the margins of error involved and the effects on the raw data of different feature extraction policies. Hence, the effects of using different image interpretation methods on the same feature can be assessed. An expert user can choose particular strategies for feature formation, and experiment with the effects these have on the resulting spatial description.

The third use of multiple representations is where the behavior of features over time is to be studied. If some spatial or spectral property of a feature changes, the feature itself may still be the same (Langran, 1989); that is, its other attributes remain unaltered. An obvious example is the change in appearance of a paddock in which a harvested crop is grown, or the change in shape of a seasonal lake as the water level declines. In both cases there may be several possible geographic descriptions that could be applied to the feature depending on the time of year that the image data were captured. To the user the features themselves remain persistent, and such changes must somehow be included within the overall description of each feature.

Overview of the GIS

To support the required increase in functionality, the GIS must be extended in several ways. In general terms, a series of image analysis tools must be made available, and extra functionality must be added to support their use. It is not strictly necessary that these tools be physically located within the GIS, but rather that their behavior be under the control of the GIS. Specifically, the GIS must provide a level of functionality beyond that which is normally available, in the following areas:

- Support for multiple spatial descriptions of features:
 - to allow different interpretation strategies to be employed;
- Ability to create, store, access, and structure the required knowledge:
 - to manage the provision of expert knowledge for feature extraction;
- Ability to control the behavior of a set of image processing tools:
 - to allow feature extraction strategies to be under the control of the system or the expert user; and
- Runtime access to the query semantics and information about the domain:
 - to ascertain the subject of the query, and the domain information required to locate suitable feature extraction knowledge.

Ideally, these extensions would be incorporated into an existing, proven GIS. This is theoretically possible but requires an extremely complex shell to be built that surrounds the GIS to somehow include the missing functionality and storage structures. Most of this functionality is difficult or cumbersome to provide without access to the internal structure of a GIS, through source code or low level libraries, and consequently the solution adopted here uses a GIS testbed developed by the authors and others as a result of several research projects. Interested readers are directed towards Roberts *et al.* (1991) and Roberts and Gahegan (1993). To support the work described here, the GIS has been extended at the conceptual and physical layer. An additional new layer, the view layer, is added to manage the transition from raw image data to feature data (Gahegan, 1994). A description of the functionality of each layer is given below.

System Architecture

The system is built as a series of layers, each one being a further abstraction of the raw image data, in much the same way as is found with conventional databases. The structure is an expansion of the general ANSI/SPARC three-layer architecture for relational databases that is often used to describe conventional (non-spatial) database models. Note that other types of spatial data such as classified overlays and feature boundaries can also be included directly into the model as view data and feature data, respectively. The discussion below relates to the structure illustrated in Figure 1.

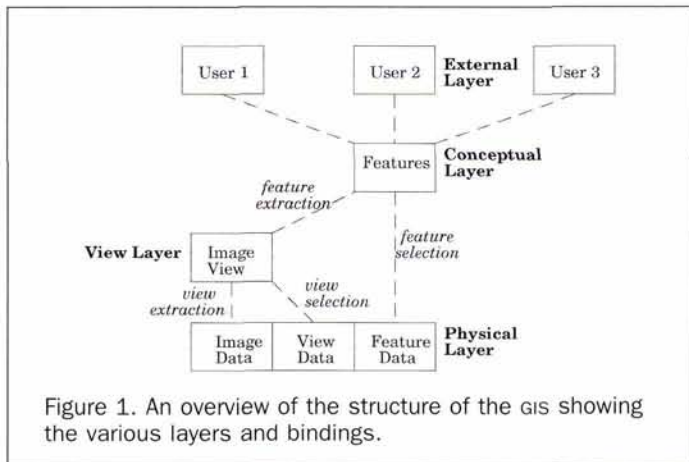


Figure 1. An overview of the structure of the GIS showing the various layers and bindings.

Physical Layer

The physical layer manages the storage and retrieval of all forms of data to and from secondary storage. It must contend with image data, view data, feature data, and attribute data (alpha-numeric data, often referred to as *aspatial*). Image data are always in raster format, while view data and feature data may be expressed in raster or vector form.

View Layer

The mappings between the physical data and the conceptual model of the domain are necessarily more complicated because the former contains image data and the latter contains features. Hence, the need for some further translation layer, the view layer, to bind them together. It is this layer and associated mappings that enable the non-expert user to obtain features from the image data without necessarily understanding all the processes and knowledge involved. The image processing component resides at this level. It is controlled by the knowledge base in the conceptual layer or optionally from the external layer in the case of expert user intervention. The detailed workings of this layer are described later under the heading "feature extraction."

Conceptual Layer

At the conceptual layer, the GIS is built around geographic features. Features are described by means of a specialization hierarchy within an object-oriented architecture (Smith and Smith, 1977). Various attributes, methods, and behavior are associated with each feature type, and together these form the class descriptions (Worboys *et al.*, 1991). Knowledge concerning feature extraction is also included as a series of frames (Minsky, 1975) associated with each feature type. Note that the spatial description is here just one more attribute of a feature, which can now assume more than one value. In fact, its value is not fixed until query time, when it is instantiated by analyzing the image data according to the stored knowledge.

External Layer

The external layer is where the user interacts with the GIS. At the external layer, various interfaces can be provided that form a *view* onto the conceptual layer, appropriate for a particular application. To the user or application program, the system contains a set of geographic features of certain predefined types which can be manipulated together to form maps, or produce other types of results. The user interface is independent of the type or source of the data. In other words, the non-expert user operates on features without

knowing whether they are defined as such in the physical data or are being formed by the system from image data.

How the System Operates

At execution time the subject of the current query or operation is found. Next, the system supplies an appropriate view of the image data, giving the best possible emphasis to that subject. To achieve this, the stored knowledge concerning feature-type occurrences is examined, and then is applied to form the image view before the query is processed. The information required to select the most suitable extraction frame is taken from the query itself, and also from the current domain (that is, the area under investigation, the time of year, the available sources of data) and is determined by matching with the frame header. The raw images are then manipulated according to the methods detailed in the frame. The end result is that a derived spatial representation becomes associated with, or causes the creation of, each feature instance. For regional features, this will most often be in the form of a set of pixels; however, the representation can be given as a set of vectors if this is more appropriate (for *linear* features, or where vectors are more desirable for any operations that are to follow). Views created for a particular task are usually saved for future reference. Consequently, before a new view is created, a check is made to see if the required feature data are already available.

Example of Operation

The user first selects a dataset, for example, the wheat belt in Western Australia (WA):

"dataset = wa-wheatbelt."

This determines the domain of study and is used to match against frame headers when searching for appropriate knowledge for feature extraction. The user continues by issuing a query such as

"display all wheat paddocks."

The subject of the query, "wheat paddocks," is determined and the corresponding feature type is located in the object hierarchy. The operation to be carried out (in this case, "display") can also place restrictions on the suitability of particular knowledge frames, depending on such factors as the current viewing scale and the preferred style of presentation (pixel or vector). From here a search is made through the frames, examining the headers to find a frame for extracting wheat paddocks in WA. Assuming that appropriate knowledge and data are available, a suitable frame will be found. If this frame has been previously executed, and the required features have already been extracted, then the frame simply points to the appropriate storage structure where the feature data are located. If not, the frame is passed to the feature extraction process. Here, the methods described in the frame are executed and feature descriptions are produced. These are then visualized using the appropriate drawing method inherited from the feature class. An example frame for the above query, using Landsat TM data, is described later.

Feature Extraction

We now turn our attention to describing the additional functionality required for the integration of feature extraction. From the standpoint of computer vision, the design of an interactive GIS incorporating image analysis tools provides a goal-oriented approach that is lacking where image understanding is seen as a separate preliminary process. A problem that is often encountered in image understanding is that of specifying accurately what constitutes a feature (Marr, 1982), or, in other words, "what exactly are we looking for?"

The fundamental difficulty with edge detections lies in the definition of an edge. The problem is identifying what constitutes a *significant* change in intensity. This involves specifying the structure of an edge either as a sharp step edge, or as a *gradual blurred intensity change*. Essentially, this is a problem of scale because a gradual blurred edge appears as a sharp edge at a greater scale and vice versa. In computer vision the problem may be addressed by scale space analysis (Sonka *et al.*, 1993) in which edges are marked only if they occur at several resolutions within the image. Essentially, the aim is to detect the same feature at a number of generalizations of the data. However, this simply shifts the problems because it is still necessary to specify a range of (theoretically infinite) resolutions and to state how the feature is reconstructed from its multiple detections. We approach the problem of specifying significance in a more simplistic manner. We regard the knowledge concerning the resolution of the image and the scale of a particular feature type to be contextual, prerequisite knowledge for feature extraction which is attached to the relevant frame. For practical purposes, this means that a frame contains the processes and parameter values which are necessary to extract the given feature type. For example, we know that (as a function of scale) roads may appear in Landsat TM imagery as small-scale (aliased) edges; therefore, the smoothing parameter for the Canny edge detection (Canny, 1986) is set to emphasize edges of this type. Although this simple technique performs quite well, it is often true that there exists no direct relationship between a feature type and the most appropriate parameter values to extract it from its surroundings. We are currently investigating a *more robust, iterative* technique in which parameter values are varied strategically based on an evaluation of the results obtained.

Feature-Based Processing

Feature-based methods use the properties and relationships of features as their basis for grouping together the graphical elements within a view into spatial descriptions of features. Hence, the feature-based descriptions abstract from the pixel-based view of the scene and enable structural characteristics of the data to be established. Attributes such as size, shape, and relationships with other features provide useful cues for the purposes of interpretation. Recently, a number of researchers have realized the potential of using, in particular, region-based techniques in the analysis of remotely sensed data. Nichol (1990) described some basic algorithms and demonstrated the usefulness of this approach. Bartl *et al.* (1993) showed how *region-based analysis* can be useful for fusion of information from multiple platforms. Taylor *et al.* (1986) used region-based techniques as the basic reasoning element in their knowledge-based segmentation system. Similarly, line-based techniques have been used in the automatic extraction of linear features such as roads (Cleyenbreugel *et al.*, 1990) and lineaments (Wang and Howarth, 1990).

Once detected, feature attributes such as size, shape, and minimum bounding box are calculated and the feature description is stored within a persistent feature store, known as a feature database (FDB). In essence, this is the "spatial database" component of the GIS. The mapping from an image view to features is not straightforward due to the presence of uncertainty. There is a decision point where the spatial extent of a feature must be determined from the "fuzzy" image view. Although uncertainty is still required, it must become feature based, as opposed to pixel based. For example, consider the mapping of a four-band view representing the *a posteriori* probabilities of a maximum-likelihood classification onto a database of features. To form regions, a single label must be assigned to each pixel (in this case, the maximum *a posteriori* value of the four bands). One method to

evaluate the uncertainty of such a feature is to simply calculate an average of the pixel-based uncertainty values that constitute the feature. In the previous example, the certainty of a region would be an average of the *a posteriori* probabilities of the defining pixels in the view.

Combining Evidence

The combination of evidence from a number of processes offers two major advantages. First, reliance of the interpretation on a single property of a feature type is reduced, and second, a variety of properties can also be considered. For example, many native forests have a relatively poor spectral discrimination (caused by irregular tree density), yet may still be identified by virtue of their boundary definition with the surrounding areas. Such concepts form the basis of recent work in the attempt to unify results from low-level process to form a more accurate image segmentation (Pavlidis and Liow, 1990). Additionally, by weighting the importance of one attribute over another, we may successfully apply the same extraction technique to a range of dissimilar feature types. The problem presented by this approach is exactly how the various sources of information (or intermediate views) may be combined. The approach taken here is to defer the formation of spatial descriptions until all available information has been considered, as opposed to forming the descriptions during low-level analysis and subsequently splitting, merging, or discarding them. In this situation the low level processing is simply used as a focusing mechanism to identify areas of the image which may merit further investigation.

In order to combine evidence, it is either necessary to define the relationship between each of the sources or to map each source into the same domain. If we define the domain into which each data set is mapped as the range of probabilities representing a particular hypothesis or goal, then each contributing data source must only define its behavior with respect to the goal and not with each of the other contributing sources. This concept is similar to the multisource data analysis approach taken by Lee *et al.* (1987) in so far as data are mapped into "data classes" as a bridge between the data and user defined "information classes."

Consider the example of using two sources of information, an edge detection and a classification. Given a specific goal, such as "identify wheat paddocks," we must be able to create a mapping from each data set to the goal. Without loss of generality, we may specify the task as the process of extracting wheat paddock boundaries. The *a posteriori* probabilities of a wheat classification are shown in Figure 3a. By thresholding these *a posteriori* probabilities, we obtain regions whose boundary may be determined, and whose probabilities are a function of the pixels they encompass. By assuming that edge strength represents the probability of wheat paddock boundaries, we have mapped both data sets into the same domain and may now combine them. Figure 3b shows an edge detection of the same area. Figure 3c shows the combined view of spectral and edge evidence. In this figure, the darker pixels indicate a high probability of a wheat paddock boundary. In essence, we are using Bayesian evidence pooling (Pearl, 1988). The domain represents the *a priori* probability of the hypothesis (goal) and evidence from various sources are used to update the hypothesis: i.e.,

$$O(H|e^1, e^2, \dots, e^N) = O(H) \prod_{k=1}^N L(e^k|H) \quad (1)$$

where $O(H)$ is the *a priori* odds; $L(e^k|H)$ is the *likelihood* of evidence e^k given Hypothesis H ; and $O(H|e^1, e^2, e^3, \dots, e^N)$ is the *a posteriori* odds, given the sources of evidence e^1, e^2, \dots, e^N . The odds of a hypothesis H are defined as $P(H)/P(\neg H)$,

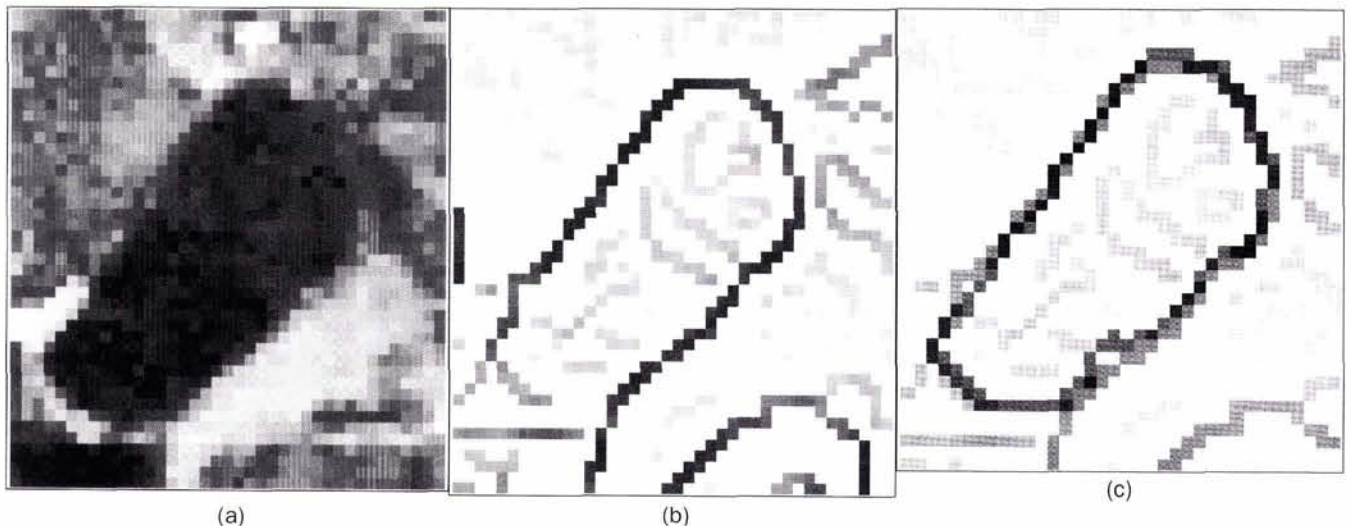


Figure 3. (a) An extract showing the *a posteriori* probabilities of wheat. The darker the pixel, the greater the probability of wheat. (b) An extract of the same area showing edge strength. The darker the pixel, the greater the edge strength. (c) Combined evidence created by the intersection of two probabilistic sources of information: in this case, edge and region boundary evidence.

and the likelihood of a hypothesis H given evidence e is defined as $P(e|H)/P(e|\neg H)$. This principle is used in an algorithm similar in nature to Milgram's convergent evidence technique (Milgram and Kahl, 1979). The algorithm can be summarized as follows: (1) threshold spectral *a posteriori* probabilities at T , (2) group pixels into regions, (3) determine the region boundary, (4) evaluate the intersection of the region boundary and the edge evidence using Equation 1 above, (5) retain region and exit if the *a posteriori* odds are sufficiently high, and (6) increment T and repeat.

Feature attributes such as shape and size can be used as an additional source of (derived) evidence. For example, consider a region with an attached probability P (wheat paddock). A size descriptor for wheat paddocks can be used to obtain the conditional probability of wheat paddocks given the size of the current region. This conditional evidence is viewed as an independent source of information and is used to update the probability of the region.

Figure 4 shows a frame describing the extraction techniques used for identifying wheat paddocks from Landsat TM imagery. The frame header indicates that this particular sequence of extraction techniques is relevant only within a limited context as specified by the query and the domain. A Canny edge detection of band 4 provides the necessary edge evidence. (Band 4 was selected because it is the best discriminant for vegetation and because a multispectral edge detection was not available.) A Mahalanobis distance metric and a training set for wheat paddocks, stored in a descriptor slot of the current frame, are used to map the multispectral data onto *a posteriori* probabilities. This information is combined using the algorithm described above, and the resulting regions are stored in the feature database.

Results

Figure 5 shows the results of applying the frame shown in Figure 4 to a TM scene of a section of the West Australian wheat belt. Wheat regions are shown hatched and overlaid on a false color composite of bands 3, 4, and 5. Much of the area (particularly in the top and bottom right part of the extract) suffers from salt damage, resulting in a reduced crop, especially in the valley floors.

Figure 6 demonstrates how a number of spatial representations of the same feature can be obtained by manipulating formation constraints. Figure 6a shows a close up of region "A" in Figure 5, which was formed by identifying regions with a probability of wheat higher than 0.7 as hatched. In Figure 6b, this constraint has been increased to 0.8. This form of analysis can be used to determine upper and lower bounds for the spatial extents of a feature. The difference between the two, in this case, indicates areas of crop damage caused by salination. By using a fairly unconstrained approach, the regions produced approximate to the paddock boundaries, and hence are useful for *cadastral* type operations. By tightening the constraints, the regions give a more true indication of the actual area of cover attributable to the various ground conditions, and are more useful, for example, in determining estimates of crop yield. At run time, one is selected in preference based on the type of analysis to be performed.

Figure 7 again shows a situation where two distinct spatial representations have been formed from a single image. Figure 7a shows how a classification misrepresents roads (shown as a series of small fragmented dark regions running through the largest gray region). In some cases, the represented width of the road is three pixels wide, or 90 metres at the scale of Landsat TM. The actual width of roads in this area is around 10 to 15 metres. In Figure 7b, pixels lying on or near the road that are not part of a neighboring feature are reclassified, and the road is represented more faithfully as a series of vectors, shown as a pair of parallel dark lines. A more detailed discussion of this work is presented in Flack *et al.* (1994). Both the pixels in Figure 7a and the arcs and nodes in Figure 7b are part of the description of the same road features. For most work within a GIS, the representation shown by Figure 7b will be chosen over Figure 7a because it has the property of connectedness.

Conclusions

An extended GIS has been developed that encompasses an image interpretation system. The binding between the images and features of interest is delayed until run time, and this approach gives more flexibility in scene interpretation than

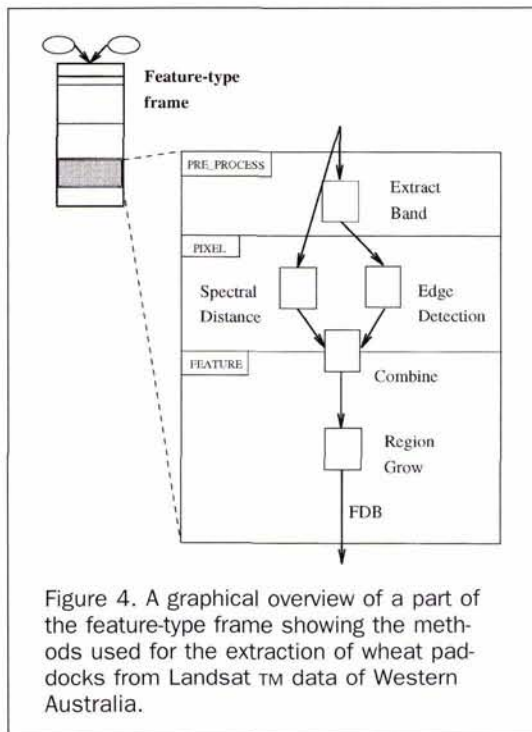


Figure 4. A graphical overview of a part of the feature-type frame showing the methods used for the extraction of wheat paddocks from Landsat TM data of Western Australia.



Figure 5. Wheat coverage overlaid on a false color composite of bands 3, 4, and 5 Landsat TM of the West Australian wheat belt.

is currently available when using image processing as a precursor to the GIS. The results presented demonstrate the adaptable behavior of the system when presented with different types of queries. By using a variety of independent image processing tools and combining the results, the accuracy (and usefulness) of the resulting interpretation is improved.

One aspect of knowledge still missing from the system described is that of image selection. Various features are differentiated most accurately from specific image sources or at specific times in the year. For example, images captured at full crop may give the best results when looking for wheat paddocks. Knowledge of this type should be used to impose

an order on the frames so that the images most likely to give good results are used first (and as the default).

The use of a knowledge base presupposes that expert knowledge is available and can be gathered at an appropriate level of detail for the tasks to be carried out. The initial costs involved in configuring a knowledge base are high, and this must be justified by considering the number of non-expert users to whom the system becomes useful as a result of this effort.

The use of GIS for all manner of spatially oriented problems is increasing, as is the desire to have access to up-to-date remote sensing data. It is this reason, along with the trend for GIS to be used by non-experts, that motivates the work described here.

There is, of course, more work to do. Specifically, we are keen to build some automated means of query evaluation into the system that can modify expert rules in the light of user feedback. We are presently investigating the tracking of recognized features between temporally sequenced images as a means of monitoring land-use change.

Acknowledgments

Our thanks go to CSIRO Division of Mathematics and Statistics in Perth, with whom we are collaborating, for helping with some of the data requirements for this project.

References

- Argialas, D., and A. Harlow, 1990. Computational Image Interpretation Models: An Overview and a Perspective, *Photogrammetric Engineering & Remote Sensing*, 56(6):871-886.
- Bartl, R., W. Schneider, and A. Pinz, 1993. Information Fusion in Remote Sensing: Combining Austromir data and Landsat TM Data, *Proceedings, 25th International Symposium, Remote Sensing and Global Environmental Change*, 2:163-172.
- Canny, J., 1986. A Computational Approach to Edge Detection, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(6): 679-698.
- Davis, F.W., and D.S. Simonett, 1991. GIS and Remote Sensing, *Geo-*

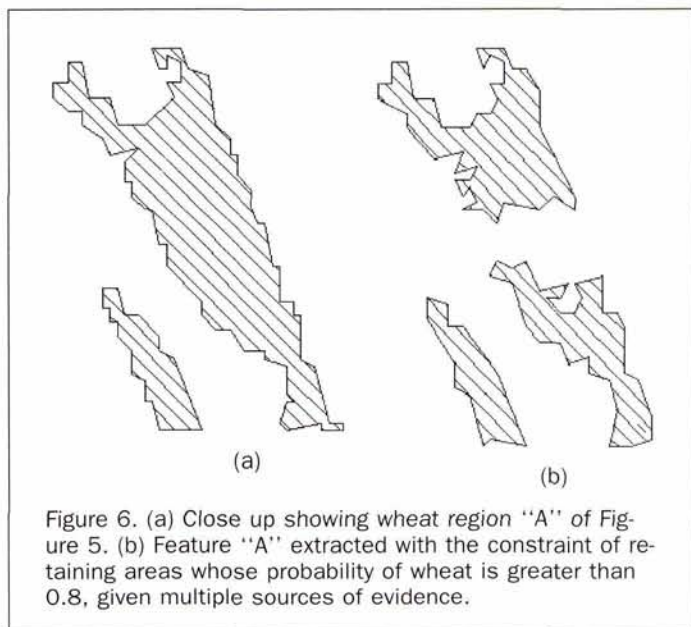




Figure 6. (a) Close up showing wheat region "A" of Figure 5. (b) Feature "A" extracted with the constraint of retaining areas whose probability of wheat is greater than 0.8, given multiple sources of evidence.



(a)



(b)

Figure 7. (a) Extract showing roads represented as pixels within a classification of the scene. This representation is inaccurate in both geometry and topology. (b) Roads represented by vectors provide a far more accurate and useful representation of the feature type.

- graphical Information Systems*, Vol. 1 (D.J. Maguire, M.F. Goodchild, and D.W. Rhind, editors), Longman, England, pp. 191–213.
- De Cola, L., 1989. Fractal Analysis of a Landsat Scene, *Photogrammetric Engineering & Remote Sensing*, 55(5):601–610.
- Duda, R.O., and P.E. Hart, 1973. *Pattern Classification and Scene Analysis*. Wiley, New York.
- Ehlers, M., D. Greenlee, T. Smith, and J. Star, 1991. Integration of Remote Sensing and GIS: Data and Data Access, *Photogrammetric Engineering & Remote Sensing*, 57(6):669–675.
- Flack, J.C., M.N. Gahegan, and G.A.W. West, 1994. The Use of Sub-Pixel Measures to Improve the Classification of Remotely-Sensed Imagery of Agricultural Land, *Proceedings, Seventh Australasian Remote Sensing Conference*, Melbourne, Australia, pp. 534–543.
- Gahegan, M.N., 1994. A Consistent User-Model for a GIS Incorporating Remotely Sensed Data. *Innovations in GIS* (M.F. Worboys, editor), Taylor and Francis, London.
- Gong, P., and P.J. Howarth, 1990. The Use of Structural Information for Improving Land-Cover Classification Accuracies at the Rural-Urban Fringe, *Photogrammetric Engineering & Remote Sensing*, 56(1):67–73.
- Korth, H.F., and A. Silberschatz, 1991. *Database System Concepts (Second Edition)*, McGraw-Hill, Chapter 1.
- Langran, G., 1989. A Review of Temporal Database Research and Its Use in GIS Applications, *International Journal of GIS*, 3(3):215–232.
- Lee, T., J.A. Richards, and P.H. Swain, 1987. Probabilistic and Evidential Approaches for Multisource Data Analysis, *IEEE Transactions on Geoscience and Remote Sensing*, 25(3):283–293.
- Marr, D., 1982. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*, W.H. Freeman and Company, 270 p.
- Milgram, D.L., and D.J. Kahl, 1979. Region Extraction Using Convergent Evidence, *Computer Graphics and Image Processing*, 11:1–12.
- Minsky, M., 1975. A Framework for Representing Knowledge, *The Psychology of Computer Vision* (P. Winston, editor), McGraw-Hill, New York.
- Nichol, D.G., 1990. Region Adjacency Analysis of Remotely-Sensed Imagery, *International Journal of Remote Sensing*, 11(11):2089–2101.
- Pavlidis, T., and Y.T. Liow, 1990. Integrating Region Growing and Edge Detection, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(3):225–233.
- Pearl, J., 1988. Probabilistic Reasoning in Intelligent Systems, *Networks of Plausible Inference*, Morgan Kaufmann, California, Chapter 2.
- Roberts, S.A., M.N. Gahegan, J. Hogg, and B.S. Hoyle, 1991. Application of Object-Oriented Databases to Geographic Information Systems, *Information and Software Technology*, 33(1):38–46.
- Roberts, S.A., and M.N. Gahegan, 1993. An Object-Oriented Geographic Information System Shell, *Information and Software Technology*, 35(10):561–572.
- Smith, J.M., and D.C.P. Smith, 1977. Database Abstractions: Aggregation and Generalisation, *ACM Transactions on Database Systems*, 2:105–133.
- Sonka, M., V. Hlavac, and R. Boyle, 1993. *Image Processing, Analysis and Machine Vision*, Chapman and Hall.
- Ton, J., J. Sticklen, and A.K. Jain, 1991. Knowledge-Based Segmentation of Landsat Images, *IEEE Transactions on Geoscience and Remote Sensing*, 29(2):223–231.
- Torre, V., and T. Poggio, 1986. On Edge Detection, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-8(2):147–163.
- Van Cleynenbreugel, J., F. Fierens, P. Suetens, and A. Oosterlinck, 1990. Delineating Road Structures on Satellite Imagery by a GIS-Guided Technique, *Photogrammetric Engineering & Remote Sensing*, 56(6):893–898.
- Wang, J., and P.J. Howarth, 1990. Use of the Hough Transform in Automated Lineament Detection, *IEEE Transactions on Geoscience and Remote Sensing*, 28(4):561–566.
- Wharton, S.W., 1987. A Spectral-Knowledge-Base Approach for Urban Land-Cover Discrimination, *IEEE Transactions on Geoscience and Remote Sensing*, 25(3):273–283.
- Worboys, M.F., H.M. Hearnshaw, and D.J. Maguire, 1990. Object-Oriented Data Modelling for Spatial Databases, *International Journal of GIS*, 4(4):369–383.
- Zhou, Q., 1989. A Method for Integrating Remote Sensing and Geographic Information Systems, *Photogrammetric Engineering & Remote Sensing*, 55(5):591–596.

(Received 28 September 1993; revised and accepted 1 November 1994; revised 17 February 1995)