INT. J. GEOGRAPHICAL INFORMATION SYSTEMS, 1993, VOL. 7, NO. 3, 247–262

An experimental system for the integration of GIS data in knowledge-based image analysis for remote sensing of agriculture

C. KONTOES, G. G. WILKINSON, A. BURRILL, S. GOFFREDO and J. MÉGIER

Institute for Remote Sensing Applications, Joint Research Centre, Commission of the European Communities, 21020 Ispra, Varese, Italy.

Abstract. This paper describes a knowledge-based system which ha developed for integrating easily-available geographical context information GIS in remotely-sensed image analysis. An experiment is described in wh maps and buffered road networks have been used as additional data la classifying single date SPOT images for estimates of crop acreages. Tl datasets have been digitised, co-registered to the satellite imagery, and mani using ARC/INFO. The knowledge base consists of both image context ru geographical context rules. Probabilistic information from the image classi from the rule base is combined using the Dempster-Shafer model of ev reasoning. Tests using ground data from the Département Loir-et-Cher, have shown that use of the knowledge-based system with GIS data g accuracy improvement of approximately 13 per cent compared to parametric image classifier alone.

1. Introduction

The classification of remotely-sensed imagery is normally performed usi multispectral information available for each pixel of an image. However, this pro does not always give high accuracy and the reliance on image data alone is not su for many applications. It is also generally acknowledged that human photointer use a considerable variety of contextual information and common-sense experi interpreting aerial or satellite imagery. Much of this contextual information and experience concerns factors which could be derived from maps or geogr information systems (GIS). Indeed, remote sensing and geographical infor systems are closely related in that not only does geographical information potential role in helping to interpret remotely-sensed data but also remote sensin major role as an information source for GIS (Trotter 1991). This relationship, of necessitates vigilance to avoid error-amplifying circularity.

At the Joint Research Centre work is underway on the development of aut satellite mapping methodologies that give high accuracy and which can be use continental scale in support of the operational requirements of European Com policy. One part of this activity currently concerns the integration of ancillary geographical information in the process of image interpretation. For this pur restriction is made to basic types of geographical data sets such as terrain mod maps, road networks and climatological zoning which are readily available thro the Community. These data can provide additional context information which c in satellite image analysis. The goal is to develop methods of image underst which use optimal classifiers exploiting background geographical imformation relevant to the imagery.

In view of the need to utilise both numerical and logical information in an integrated procedure for image analysis a knowledge-based system approach has been adopted. Such an approach has been proposed by various workers over the past ten years or so. McKeown (1987) and Goodenough *et al.* (1987) have both reported developments of knowledge-based or 'expert' systems for remote sensing, and several successful applications of rule-based methods for augmenting satellite data with ancillary topographic information have been reported. For example, Hutchinson (1982) used rules based on terrain slope to modify initial assignments of image classes concerning geological features such as dunes and alluvial fans in Landsat Multispectral Scanner images; Van Cleynenbreugel *et al.* (1990) used DTM information to assist in the extraction of road structures from SPOT imagery using knowledge such as 'mountain roads usually follow ground contours'. Cibula and Nyquist (1987) used both topographic and climatological data integrated in a GIS to refine the classification of Landsat Multispectral Scanner imagery. Slope data were used to discriminate water from forested areas in shadow. Also, aspect, elevation, slope and precipitation and climate regimes were used to discriminate specific forest types using boolean decision rules.

In this paper, we describe an experiment in which ancillary GIS information has been incorporated in a knowledge-based system with the aim of making estimates of agricultural production from satellite imagery. The twin requirements of regular estimates and of European scale coverage (at high latitudes with much cloud cover) often prevents the use of multi-temporal imagery (Kontoes and Stakenborg 1990). The experiment reported here thus concerns the use of GIS data to improve the classification of mono-temporal (single-data) imagery.

2. The knowledge-based system method

2.1. *Basic methodology*

In our methodology available geographical information is used in a post-classification process, i.e., to refine classes which have been derived initially solely from image information. This is done by combining the output from the image classifier with new class evidence coming from rules which are triggered by contextual information. This information may be derived both from the imagery itself and from ancillary geographical datasets. In our implementation a forward chaining control strategy is used in which the contextual information triggers rules leading to certain actions. The procedure is illustrated in figure 1. An ARC/INFO system is used to store and manipulate the ancillary geographical data.

The knowledge-based system uses a methodology of evidential reasoning for 'weighing up' class evidence which comes initially from the image classifier and subsequently from the rules. This procedure is based on the Dempster-Shafer (hereafter D-S) theory of evidence (Shafer 1976) and involves the computation of 'belief' values representing the combined degree of confidence in each class as a result of the combination of evidence from the various sources. This method was proposed for use in remote sensing by Lee *et al.* (1987). Recent applications of this technique for integrating multiple sources of data in remote sensing have been reported by Peacegood (1989) and by Srinivasan and Richards (1990). Its potential for integrating image processing and ancillary information in land cover classification was reported by Wilkinson and

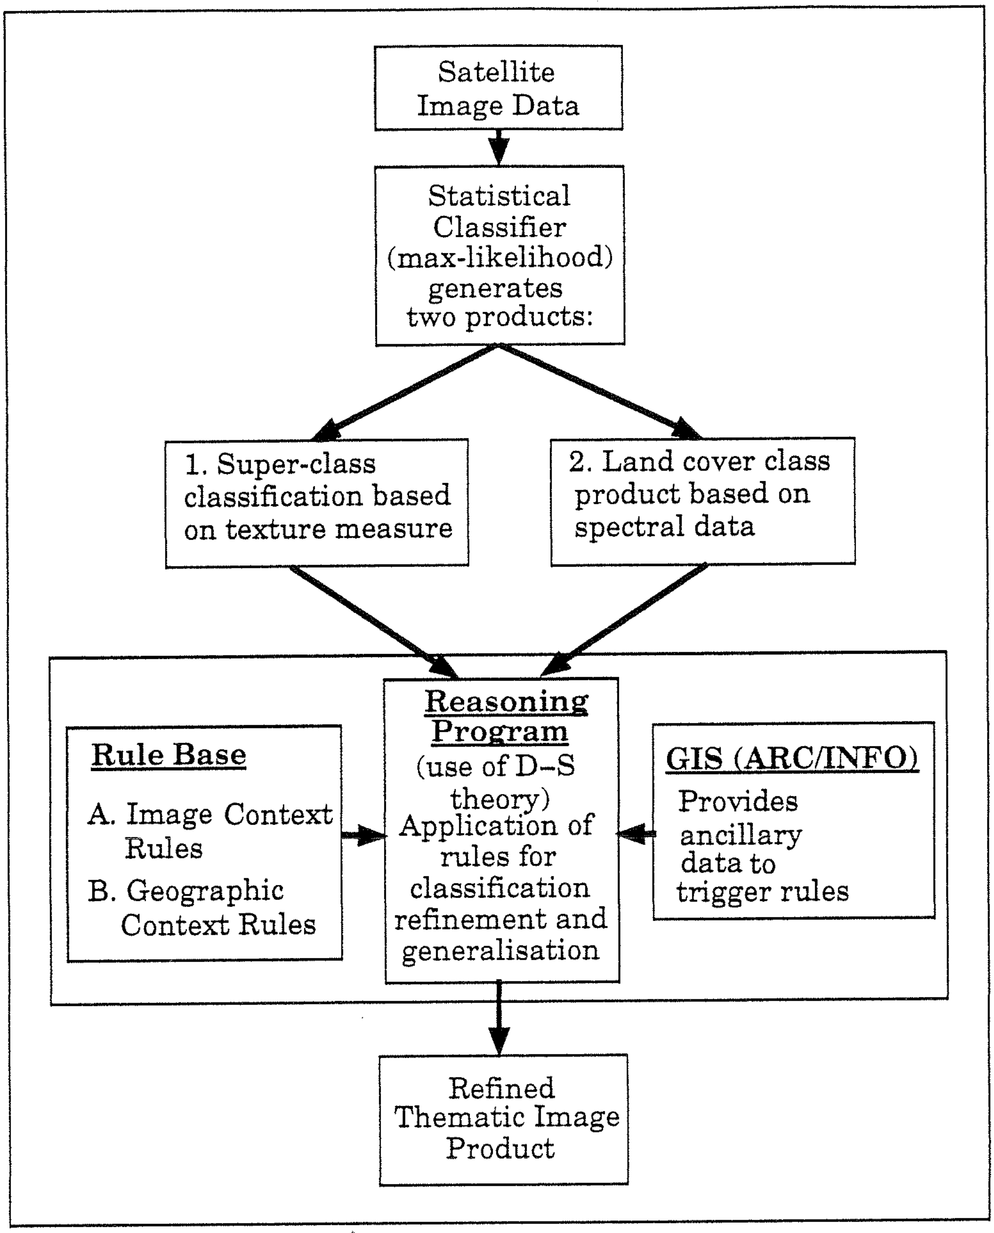

Figure 1. Knowledge-based system procedure.

Mégier (1990). A particular advantage of the D-S method is that it permits evidence to be associated not only with individual thematic classes but also with groups of classes (super-classes). This enables uncertainty to be handled. For example, if the net belief is low for individual classes but high for super-classes this implies that we have insufficient information to make an individual classification for an image pixel and that we should instead assign it to a super-class. This is useful in hierarchical classification schemes often used in land cover mapping.

As can be seen in figure 1, the ancillary information is brought into the process after the computation of a first classification product using the image classifier alone. The reasoning process is applied to each pixel separately. Rules in the knowledge base contain separate CONDITION and ACTION parts (see table 1). Initially the CONDITION parts are compared to the ancillary data applying to a pixel. If all of the CONDITION parts of a rule 'match' the available geographical data then the relevant rule is 'fired'. The rule then generates a piece of evidence (from its ACTION part) in the form of a 'support value' for a particular class or super-class (see table 1) which is put

Rule 391 Geographical context rule: soil-type

Conditions
- The dominant soil class within the immediate 4-connected window centred on this pixel is 'calcareous clay'
- The initial spectral class assignment for this pixel is 'maize'
- The dominant initial spectral classification for a 9 by 9 pixel window centred on this pixel is 'sunflower'

Actions:
- Refute the initial class assignment of this pixel with support value 0·9 [very strong disconfirming evidence]
- Confirm the class 'sunflower' as a candidate class for this pixel with support value 0·9 [very strong confirming evidence]

Explanation:
- There is considerable spectral confusion between 'maize' and 'sunflower' early in the growing cycle owing to background soil reflectance especially on calcareous clay soil with low water retention capacity. The rule aims to correct isolated misclassified 'sunflower' pixels which have been initially classified as 'maize' where they exist in a 'sunflower' dominated neighbourhood. The initial 'maize' class is thus strongly refuted and the 'sunflower' class is strongly supported.

Rule 650 Geographical context rule: access-route-proximity-type

Conditions
- The pixel is located more than 1000 m. from a road
- The dominant initial classification within the immediate 4-connected image window centred on this pixel is either 'forest' or 'pasture'
- The dominant initial classification in the 9 by 9 pixel window centred on this pixel is supported by the initial texture-based super-class assignment for this pixel

Actions:
- Confirm the dominant class of the neighbourhood as a candidate class for this pixel with support value 0·7 (strong confirming evidence)

Explanation:
- The pixel is located far from a road where we expect to find agricultural or natural vegetation classes. The dominant spectral class in the neighbourhood of this pixel is 'pasture' or 'forest' and there is also support for this at the super-class level from texture. The dominant class is thus strongly confirmed as a candidate class for this pixel.

Rule 290 Image context rule

Conditions:
- Three of the immediate 4-connected neighbours of this pixel have the same initial spectral class assignment as this pixel
- The spectral class assignments of this pixel and of its three matching 4-connected neighbours are not *all* supported by the super-class assignments derived from texture
- The super-class assignments of two out of these four pixels are consistent with their initial spectral class assignments

Action:
- Commit confirming evidence to the initial spectral class assignment for this pixel with support value 0·5 (medium confirming evidence)

Explanation:
- If three of the immediate 4-connected neighbours of this pixel have the same initial spectral class assignment as this pixel and the super-class information derived from texture is also consistent for two of the four 'matching' pixels, then the initial class assigned to this pixel is likely to be correct and should be supported, but only with medium strength because the texture-based evidence is not highly supportive in this region of the image

Table 1. Examples of rules

into a temporary storage area—a 'blackboard' (Hayes-Roth *et al.* 1983). After all rules have been checked for matching, the D-S method is used to combine all the support values from the matching rules to compute final belief values for all the classes and super-classes in the hierarchy. This has the desired effect of weighing up all the evidence. In this procedure class likelihood values generated by the image classifier are also used as support values so that there is a full integration of image information and ancillary geographical information. The optimum classification (taking account of both the image data and the ancillary geographical data) is indicated by the class or super-class with the maximum belief. The full list of procedural steps is outlined in figure 2.

2.2. *Evidence conbination procedure*

The support values generated by the rules indicate the strength of the evidence. The support values lie in the range $0 \rightarrow 1$ and can be used either as confirming evidence or disconfirming evidence. High values indicate a stronger degree of confirmation (or disconfirmation if the rule concerns disconfirming evidence). Rules which generate low support values can be regarded as contributors of weak evidence. Conversely rules which generate high support values contribute strong evidence. The example rules shown in table 1 have different levels of support. The possibility of expressing various degrees of certainty to confirm or disconfirm hypotheses about the true land-cover classes of pixels is a very useful technique. Typically the rules used in our system contain three or four conditions. In practice, however, more complex conditions based on coincident values of many different geographical parameters could be used depending on the availability and relevance of ancillary GIS datasets.

The combination of all the evidence, beginning with the support values, was performed according to the algorithm of Gordon and Shortliffe (1985), which is an approximation to the D-S approach. It is computationally faster than the full D-S method, although its use is restricted to hierarchical tree-structure classification

FOR EACH PIXEL:

STEP 1: Retrieve class likelihoods from statistical image classifier;

STEP 2: Retrieve ancillary GIS data;

STEP 3: Check if GIS data matches CONDITION PART of rules in the rule base;

STEP 4: Store all confirming or disconfirming support values from ACTION PARTS of all matching rules for all classes/super-classes on blackboard;

STEP 5: Merge confirming support values;

STEP 6: Merge disconfirming support values;

STEP 7: Combine results of STEP 5 and STEP 6, to compute 'belief' values for classes and super-classes;

STEP 8: Allocate the class/super-class of maximum belief.

[Steps 5, 6, and 7 are performed according to the D-S theory]

Figure 2. Algorithm for rule-based integration of GIS data in pixel analysis.

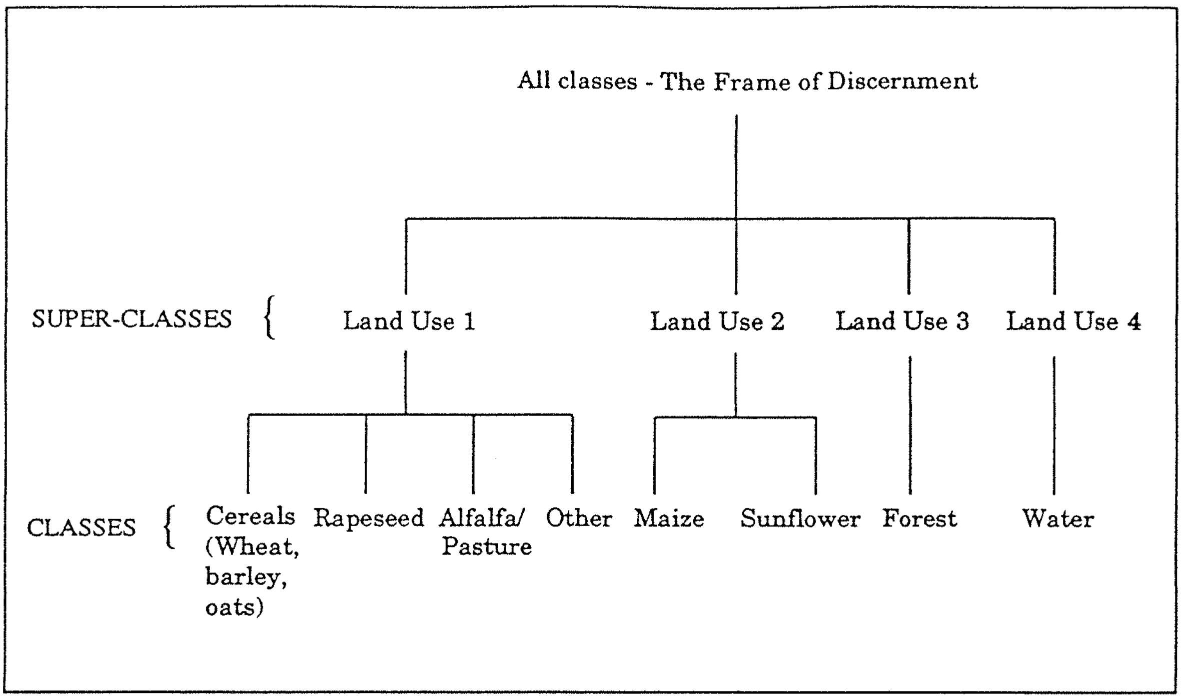

Figure 3. Hierarchical land cover classification scheme.

schemes, which is acceptable in this case (figure 3 shows the classification scheme which was used). In the evidence combination method each entry or node in the tree represents a set of one or more classes. The node at the top of the tree represents all classes and is usually called the frame of discernment or universe of discourse, and represents all classes under consideration. The evidence combination procedure then proceeds as follows: In steps 5 and 6 of figure 2 the support values from the rules both confirming and disconfirming a particular class or super-class have to be combined. If there are n pieces of evidence *confirming* a particular class or super-class X in the classification tree the individual support values S_i from each one are combined as follows:

$$m_x(X) = 1 - [(1 - S_1)(1 - S_2) \ldots\ldots (1 - S_n)] \tag{1}$$

where $m_x(X)$ denotes the 'basic probability assignment' (bpa) for class/super-class X derived from evidence relating solely to X itself. The remaining 'unassigned support' $(1 - m_x(X))$ is assigned as a 'null contribution' to the frame of discernment i.e. it is support which cannot be validly assigned to anything else since it represents lack of information. If there are some pieces of evidence refuting the class X the support values for these are combined in exactly the same way as above but the basic probability assignment calculated by this is assigned to the set complement X' of X.

After the combination of support values into basic probability assignments (bpa) it is necessary to propagate the effects of the evidence for individual classes or super-classes throughout the classification tree. This is necessary because, for example, strong support for a super-class should increase the level of confidence in its individual component classes. Thus a combined bpa value is computed known as 'belief' (step 7 of figure 2) by taking account of supporting parent-child class relationships:

$$m_T(X_i) = m_{x_1}(X_i) \oplus m_{x_2}(X_i) \oplus m_{x_3}(X_i) + \ldots \tag{2}$$

where $m_T(X_i)$ denotes the belief in class/super-class X_i arising from the combination of confirming evidence throughout the whole classification tree (T) and $\oplus$ denotes the

Dempster-Shafer combination method. The precise method of computing this will not be given here, but can be found in Gordon and Shortliffe (1985).

It is also necessary to modify the belief values just calculated to take into account all the disconfirming evidence. Formally this is done by computing for each class/super-class X_i in the classification tree:

$$m_T(X) \ominus m'_{x_m}(X'_1) \quad \text{and then:} \quad (m_T(X) \ominus m'_{x_m}(X'_1)) \ominus m'_{x_2}(X'_2) \ldots \text{etc.} \tag{3}$$

where $\ominus$ denotes the modification of belief for the effect of the disconfirming evidence. Again the details of the algorithm are beyond the scope of this paper but can be found in the original paper by Gordon and Shortliffe (op. cit.). However, it is important to realise that the calculation depends on the relationship in the tree between the class of interest X and other classes or super-classes X_1, $X_2 \ldots$ for which there exists disconfirming evidence. For this reason the algorithm involves checking the relationships by set operations. There are several possible relationships (and sub-cases of these) and they are all catered for in the algorithm and in our software implementation.

At the end of the analysis the belief value calculated for each class or super-class indicates the combination of all available evidence impinging on it. The class or super-class with the maximum belief value after the final step can then be regarded as the optimum classification for the image pixel. However, in the original Dempster–Shafer theory other measures such as 'plausibility' and 'belief intervals' could be calculated. This is not, however, possible with the Gordon–Shortliffe approximation and the class or super-class of maximum belief is normally regarded as the optimum classification to be assigned to the pixel of interest. It is possible that the maximum belief value occurs for a super-class rather than for a single class. In this case it is to be understood that the super-class is the best classification to make. The fact that a single class has not emerged as the best classification implies that there is insufficient evidence to classify down to the single class level. If the maximum final belief occurs at the top of the tree this is an indication that there is insufficient information to make any useful classification at all.

3. Ancillary data sources

The method that has been developed is based on the classification of satellite images in raster format. Since our ancillary information originated from maps, it had to be digitised, pre-processed and converted into co-registered raster form in order to be integrated into the classification methodology. This was handled directly in ARC/INFO as described below. For testing the use of ancillary information two GIS data sets were prepared for an area in Région Centre, France, where we have been undertaking experiments on improving the accuracy of agricultural crop classification from satellites. These data sets were:

— *Soil type:* Soil type can have a strong influence on land-use suitability and is thus an important factor to consider in determining land-cover classes. Properties such as soil texture, depth, sub-surface drainage and water retention capability can have a strong influence on land-use suitability and can provide evidence in determining likely land-cover classes. The data for this study were derived from the Esquisse des Caratères Hydriques des Sols de la Région Centre from the Chambre Régionale d'Agriculture du Centre (1 : 250 000 scale) which is the only soil-related cartographic document existing for the whole region. Such a map is not ideal for interpreting satellite imagery but is typical of maps that would generally be available elsewhere for operational use.

— *Accessibility:* Land use is often also influenced by ease of access to networks for the transport of supplies or produce. This factor is usually most important in remote areas. The road network from the Institut Géographique National (IGN) sheet 2021 ouest (west) topographic map (1 : 25 000 scale) was digitised in this study to provide the necessary data layer.

4. Processing of ancillary data

4.1. *Preparation of the ancillary datasets*

The ancillary data sets were initially all computerised using the ARC/INFO Arcedit and Build modules on an IBM PC. The soil map was entered as a polygon coverage and the roads as a line coverage.

The road network was digitised into a coordinate system based on the IGN Lambert kilometric zone II étendu grid system which is the National Cartographic Projection system used in France. This coordinate system was chosen because the map documents had been produced with this projection and also the SPOT imagery used had been geo-referenced to these coordinates. Despite its much smaller scale and the absence of a graticule on the soil map, it was possible to digitise the soil zones into the same coordinate system, by using ground control points which could be identified on both the topographic map and the soil map.

The data sets were next exported and transferred to a VAXstation 3100 computer running Workstation ARC/INFO software. On the VAXstation, the road network was buffered at four distances and the resulting coverages combined to produce one polygon coverage of accessibility zones. This coverage and the map of soil zones were converted to a raster file format with 20 m. pixel size and stored in ERDAS image file format. In this format, they were ready for direct use as additional layers in the knowledge-based system (see figure 4). The ARC/INFO system allows easy input of map-derived data, for the combination and manipulation of such data (including coordinate transformation and surface modelling techniques), and for its conversion to raster format. While some of these initial manipulations could be accomplished in a raster system, in general, the processing is easier in vector format. Coordinate transformation in a vector-based system, for example, involves precise conversion of the individual points which represent the dataset, followed by rebuilding of topography. A raster-based system undertakes the transformation by a resampling of the raster elements: unless the raster is made very small (introducing a large requirement of disk space) a significant degree of error can be introduced. Similarly, the vector format software is more adept at surface modelling and interpolation. In fact, because of the constraints, many raster-based systems do not attempt to provide functions such as conversion between map projections or contour generation. ARC/INFO, however, performs these functions without problems and is thus clearly suited to pre-processing ancillary map data in preparation for its use in raster-based analysis with image data, including land-cover mapping using a knowledge-based system.

4.2. *Accuracy issues*

In combining map-derived and satellite-derived data, however, it is important to consider the inherent and operational errors associated with the original data sources and with the processing techniques respectively.

In our experiments SPOT-HRV satellite data of processing level 1B (radiometrically normalised and geometrically corrected for systematic deformations) were used. In view of (*a*) the good geometry of these data relating to the extremely small geometric

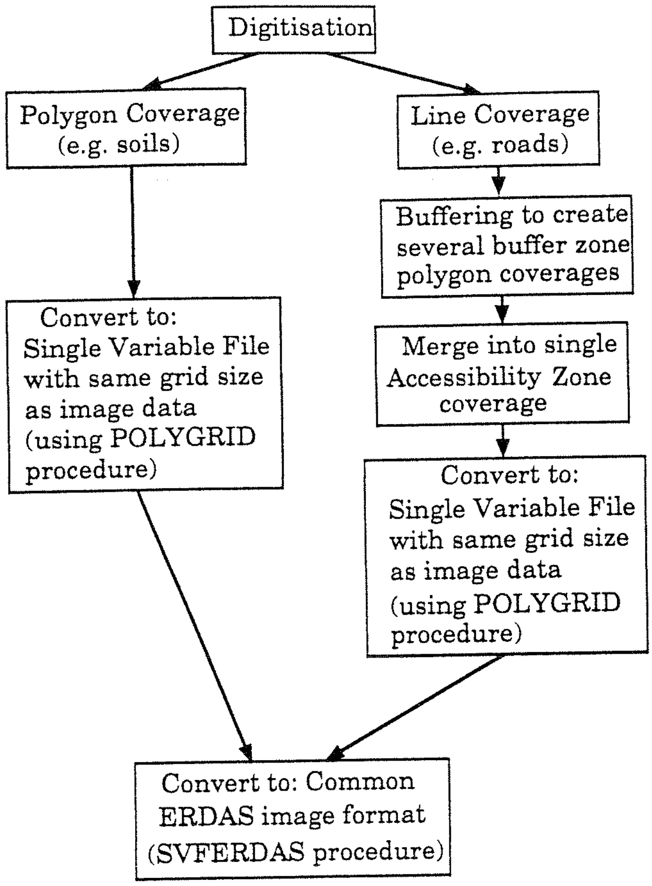

Figure 4. Processing of ancillary GIS datasets.

distortion in the along-track direction, (*b*) the relatively stable SPOT bus system (Boissin and Gardelle 1986) and (*c*) the relatively smooth terrain of the study area, the positional errors inherent to the satellite data are not considerable. Also the operational error (RMS error) introduced during the image to map registration process is only of 0·37 pixels in both directions. A first order polynomial function was used for the geometric correction and map registration of the satellite data. The satellite data were projected using the Lambert Conformal Cartographic Projection System at a scale of 1 : 25 000. The whole process was based on the use of seven ground control points distributed through the study area.

In contrast to the satellite image, the soil data introduced larger inherent and operational errors into the analysis. The inherent errors relate mainly to the map scale, which leads to generalisations and thus to mis-labelling errors, especially in cases of small soil inclusions in large uniform soil types. The operational errors are also very important and relate not only to the map scale and human errors in digitising but also to the procedure followed for the map construction. Because of the scale we expect high uncertainty in positioning the boundary lines between different soil types. The level of these 'sliver-type' errors becomes even higher if we consider that usually in soil maps the boundaries between different soil types are lines of transition (probability, confidence etc.) rather than the abrupt limits used in conventional cartography. Also, significant errors of identification are going to be introduced because the map production was based on a very small set of sampling points and mainly by exploiting geological maps. A true pedological map in the appropriate scale of 1 : 50 000 existed only in some specific areas of the Loir-et-Cher département. Inside our study area, which extends westward to the city of Blois, there exist only two sampling points where properties such as water capacity, useful water reserve, soil texture and structure have been estimated. The map production in that area was originally based mainly on the use of two geological maps (1 : 50 000 and 1 : 80 000) and some related local studies made by the 'Institut National de Recherches Agronomiques'. The quality of the input data for the construction of the soil map is clearly not ideal. But in view of the final map scale this product was accepted.

Another kind of operational error introduced into the study comes from the imprecision of manual digitising and from the warping of paper maps. In our experiment one image pixel represents an area of 20 m by 20 m on the ground and thus the width of the pixel is equivalent to 0·08 mm on the soil map and 0·8 mm on the topographic map. It thus becomes clear that digitisation and map warping may result in errors which amount to several pixels.

While the above remarks are not meant to imply that the data sets produced were necessarily unsuitable for use in combination with raster data, the accuracy of the results of the knowledge-based system is clearly influenced by the propagation of errors from the map data sets. Also the error will increase as more data sets are used. Error propagation is, of course, an important research topic in its own right and has received much attention (e.g., Heuvelink *et al.* 1989, Lunetta *et al.* 1991, Walsh *et al.* 1987). However, the propagation of error into and through a knowledge-based system using GIS data is not well-understood. Our approach has been to assign support values to rules taking account of the estimated reliability of the data set and ultimately to quantify the accuracy of the final image-derived product. It would be interesting, as a future exercise, to examine the sensitivity of this accuracy as a function of the error level in the ancillary map data sets used to trigger rules.

5. Development of an interactive front-end system

One of the secondary objectives of this work was to provide the possibility of training users in how GIS data sets can be used effectively to aid image analysis. For this purpose an interactive window-based user interface has been created for the computer (VAXstation 3100) on which the system has been developed. The purpose of the interface is to allow a user to set up a land-cover classification scheme, view a classified image, control the execution of the knowledge-based system to improve the classification and then select individual pixels to see which rules from the knowledge base have been used in this process, i.e., direct interaction with the blackboard. The user is then provided with the option of seeing details of the rules (in English text form) and if necessary requesting a full explanation of the purpose of the rule. All of this is accomplished using pull-down menus and mouse buttons. This interface is of value in assisting less experienced users and has a potential role in training. It is planned as a future enhancement to this interface to allow users cyclically to 'tune' rules or their support values on-line if the results are perceived to be unsatisfactory as follows:

$$
\begin{tikzcd}
\text{Classify} \arrow[d] & \\
\text{Apply knowledge-base system} \arrow[d] & \text{Modify rule-base} \arrow[l] \\
\text{View and evaluate results} \arrow[d] \arrow[ur] & \\
\text{Exit} &
\end{tikzcd}
$$

6. Experimental test

6.1. *Initial image classifications*

An experimental test of the system was performed for an area of $280\,km^2$ in the Département Loir-et-Cher of Région Centre using single-date multispectral SPOT-HRV imagery from 1 May 1986. This region of France is primarily a mixture of agricultural land and forests. Our experiments were aimed at both improving the accuracy of the inventory of crop classes and creating a thematically-smoothed image product for use by decision-makers. The classification scheme used in the experiment is as shown in figure 3. Before using the knowledge-based system the SPOT image data were classified using a standard parametric procedure (maximum *a posteriori* probability). However, two products were generated one giving the individual land-cover classes based on spectral classification and the other giving the land-cover super-classes in the hierarchical classification scheme using a textual classification of the image. The textural classification was based on the use of four of the twenty-eight computable features of local image texture suggested by Haralick *et al.* (1973). The four chosen features were angular second moment, contrast, correlation and variance. They are related to the variability and local correlation of the pixel intensity values in the image and appear to be useful in separating the image data into the land use super-classes. The full mathematical definitions of the texture terms can be found in Haralick *et al.* (1973). The possibility of using class information derived from spectral characteristics and super-class information derived from textural information provides additional flexibility for creating rule conditions.

6.2. *Selection of rules*

In the experiment two separate categories of rules were used in the knowledge-based system. The first set concerned only spatial context in the image. These were used to enforce spatial consistency in the image product—such as approach has been used in the past by Wu *et al.* (1988). These rules provide evidence either confirming or disconfirming particular classes for an image pixel depending on the classes and super-classes initially assigned to its neighbours (i.e., the rules tend to help a pixel to join the class of a majority of its neighbours if there is sufficient evidence for doing so). This has an important noise-removing and generalising effect.

The second set of rules directly concerned geographical context and used the GIS map-derived data sets as their basis. These rules concerned the effect of the soil type in supporting or suppressing the possibility of different agricultural classes, and the effect of distance from communications as a factor in affecting the likelihood of the existence of agricultural classes. In the test area chosen there is a tendency for the agricultural classes to be located nearer the roads—the forested areas are usually more extensive and more remote. The buffered road network dataset was used to trigger the rules concerning this effect. Typical rules of both categories are shown in table 1. The full final rule base consisted of a total of 52 rules.

The choice of rules for a knowledge-based system is in practice very difficult but several useful texts have been written on this subject, such as Hart (1986) and Keravnou and Johnson (1986). It is normally expected that a 'knowledge engineer' will elicit information relevant to the problem under consideration from an expert in the field (the so-called domain expert). This process of eliciting knowledge can be performed using models such as the repertory grid described in Hart (1986) in which the domain expert selects the most important variables in the analysis procedures and gives subjective ratings of their importance. These are then used to construct tables of potential rules.

This kind of approach is designed primarily for the situation in which the person building the system (the knowledge engineer) has no familiarity with the subject and the domain expert has little or no familiarity with the technique used to program the knowledge.

In the experiments reported here the domain experts and knowledge engineers were effectively the same people. The choice of variables for use in the rules was based primarily not on any consideration of what would be the most useful GIS variables to use in helping to improve the analysis of remotely-sensed imagery, but on which variables were actually available from the GIS data sets and which could be available more generally for most regions of Europe—this is far removed from the ideal situation often envisaged by knowledge-based system builders. The subjective importance attached by the team to the variables, such as soil type, were then used to assign support values to the rules for use in the evidential reasoning procedure.

The final rule base effectively evolved from a number of trial experiments which began by trying to list any possible relationships between the GIS data and the possible land cover classes. In some cases the relationships were relatively clear, in others they were believed to be rather tenuous. The support values were chosen accordingly.

During the development phase, however, the number of times that each rule was triggered during a test run with the data sets was monitored. This was useful in indicating which kinds of rules had most effect and also which were either not used at all or had very little use. This procedure was useful in helping to look for new possible rules which might be important and also helped in 'thinning' the rule base to remove rules which were effectively not needed and whose inclusion only slowed the execution.

6.3. *Test results*

The ancillary soil and buffered road network datasets are illustrated for the test area in figures 5 (*a*) and 5 (*b*). The effect of using the knowledge-base is illustrated in figures 6 (*a*) and 6 (*b*). Figure 6 (*a*) shows the initial product of image classification after the use of only the parametric image classifier on the satellite data alone. This product is noisy and contains errors. The accuracy of this product was evaluated by using control data from six ground segments providing ground information for 4601 image pixels. The overall accuracy achieved at this stage (averaged over all the classes) was 64·5 per cent. Figure 6 (*b*) shows the final classification product after the use of the knowledge-based system. Here the product is much less noisy and the accuracy has improved to 77·3 per cent (again averaged over all the classes). The overall improvement in accuracy by using the knowledge-based system with the map data is thus 12·8 per cent. This improvement in the mapping of the land cover classes in such agricultural regions is useful and leads to the possibility of more accurate estimations of the total crop areas. This will ultimately assist in attaining the more distant goal of operationally monitoring potential crop yields. Although this result is encouraging it is hoped that the inclusion of more rules and the tuning of the rule base support values could lead to even greater improvement in the image product, and there is much scope for further experimentation with this.

7. Discussion and conclusions

It is clear that there is much to be achieved from the use of ancillary geographical information in satellite image classification, although, the methods for integrating the different kinds of data in one thematic mapping procedure are complex and non-trivial to develop. We are confident, however, that the knowledge-based system approach,

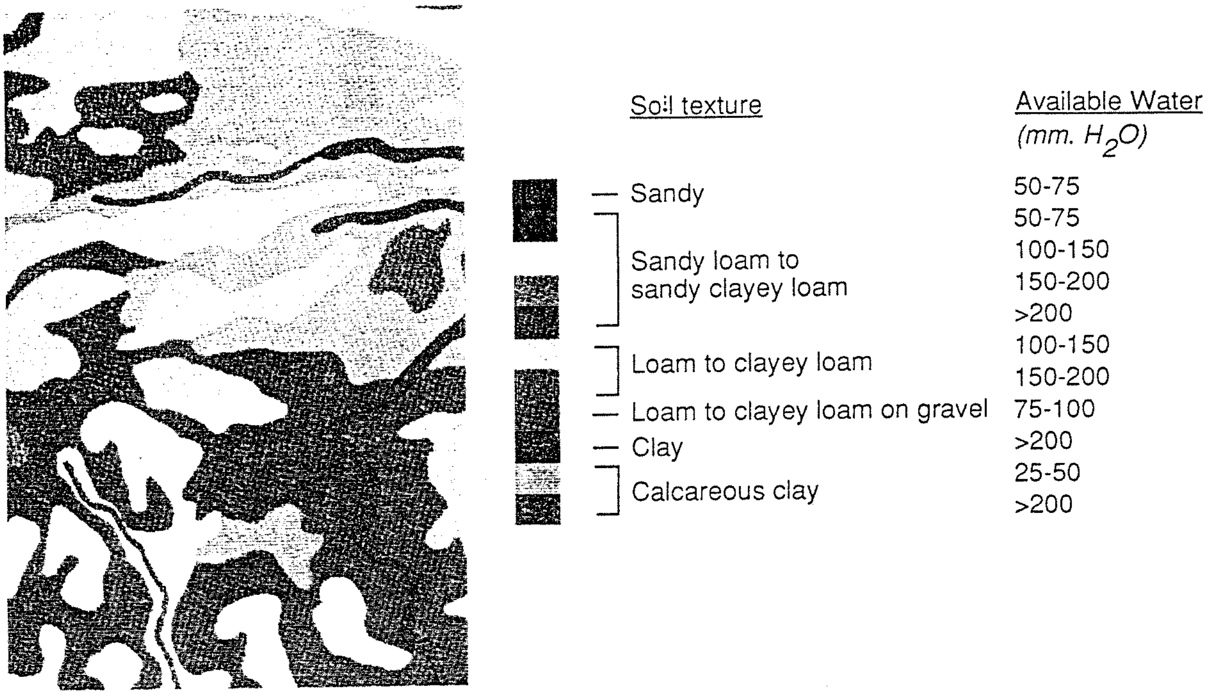

(*a*)

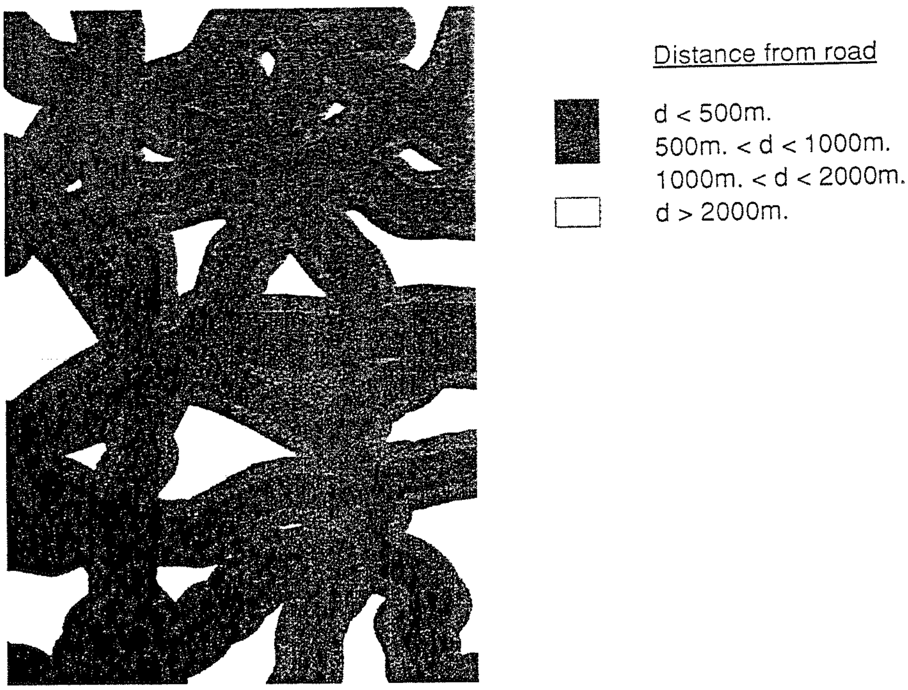

(*b*)

Figure 5. Ancillary GIS datasets used for triggering rules—20 km by 14 km experimental zone from Départment Loir-et-Cher, Région Centre, France. (*a*) Rasterised soil map. (*b*) Rasterised buffered road network.

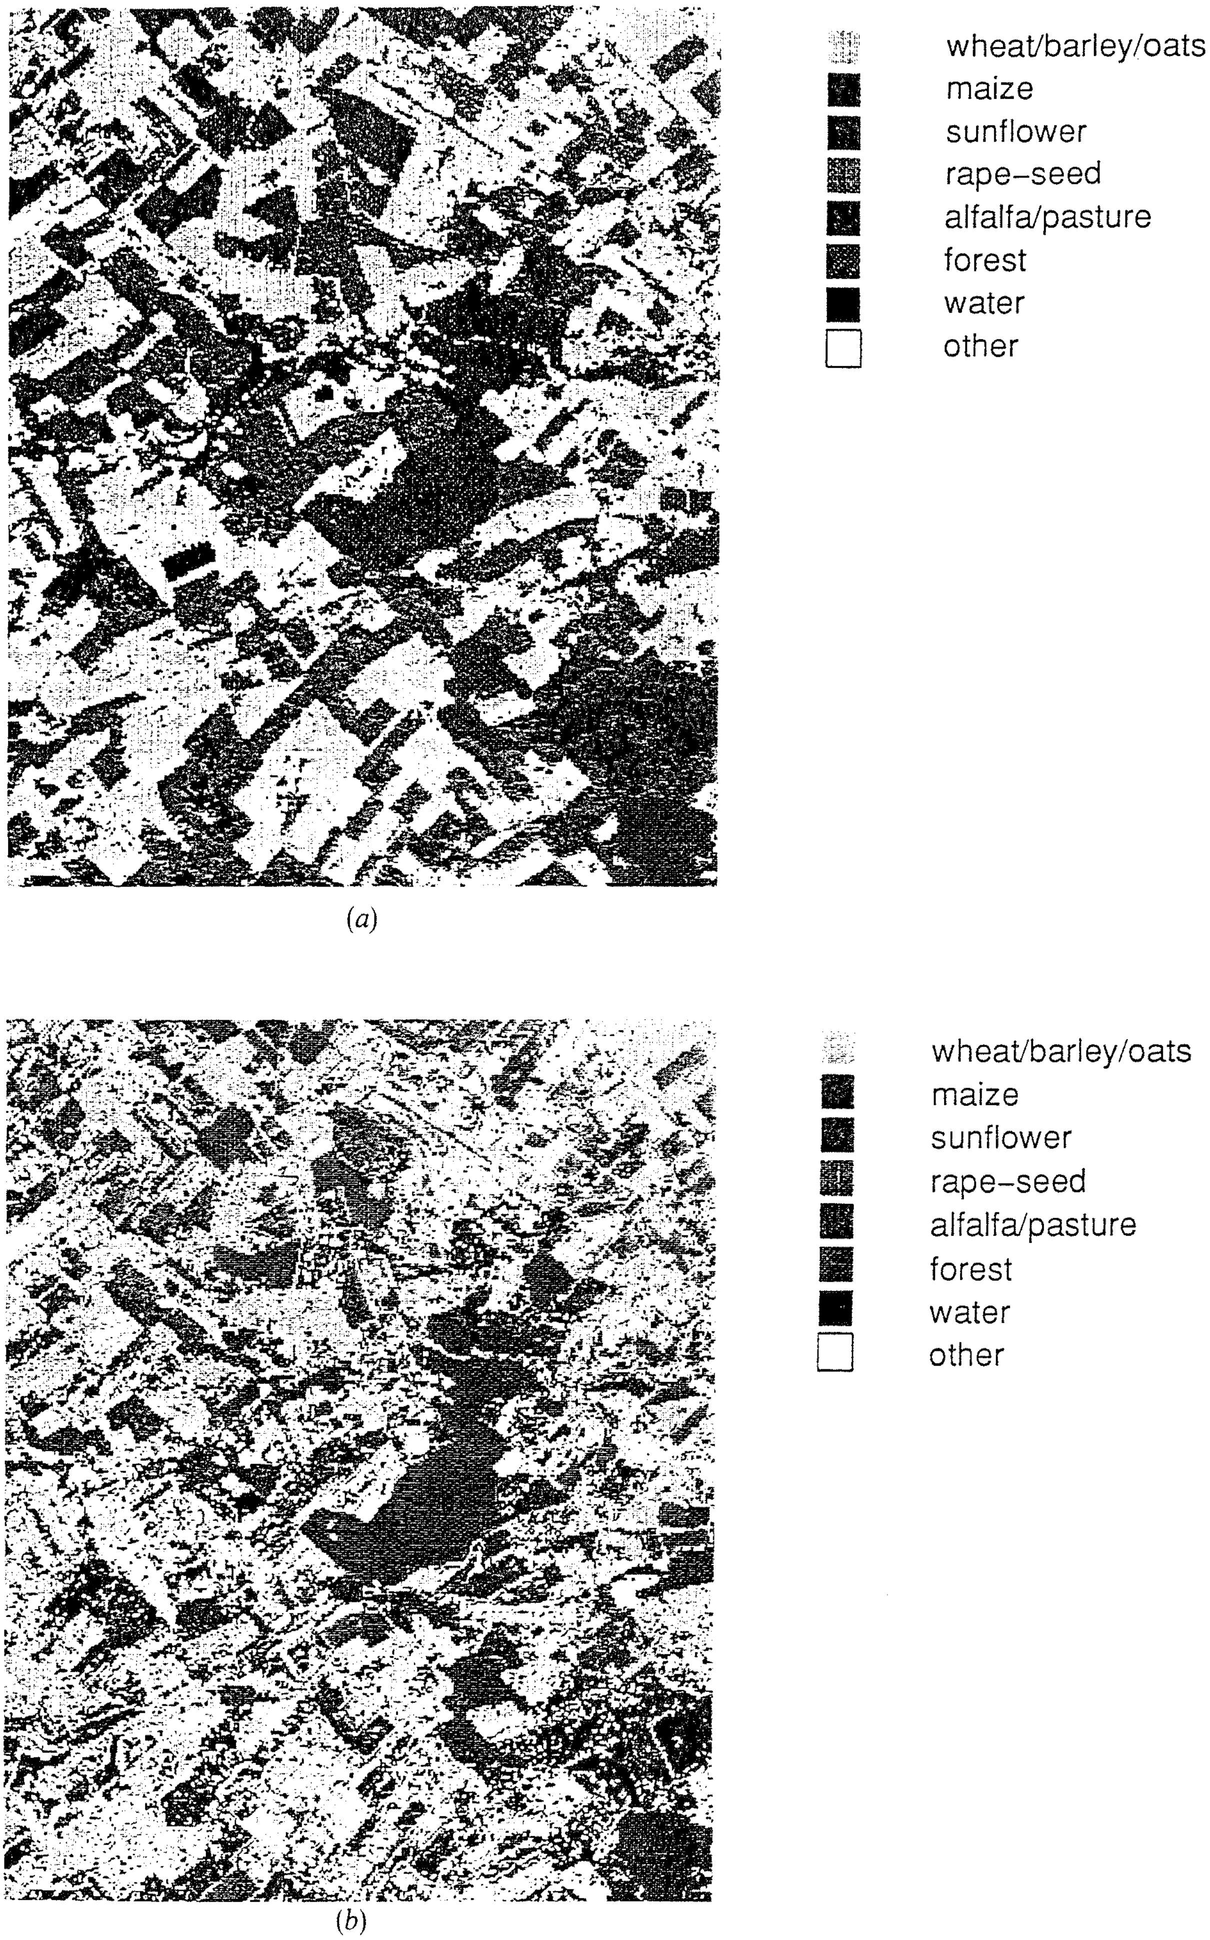

Figure 6. Test SPOT Image, Département Loir-et-Cher, Région Centre, France. Sub-window of 377 by 300 pixels (approx. 7·5 km by 6 km). (*a*) Initial single-date classification product. (*b*) Final product after use of knowledge-based system with GIS datasets.

together with evidential reasoning based on D-S theory, offers a robust method for making this integration and for combining mathematically all the information relevant to a given pixel. In practice this method could make use of extremely complex geographical relationships if these were necessary and were made available by making composite data sets in the GIS.

Also, since it is possible to construct very complex rules if necessary and, moreover, to assign a range of support values, it is also possible to use the ancillary data in a very flexible way. Some pieces of data can be given more emphasis than others in influencing the output of the classifier. In fact, the importance of each type of data is stated explicitly in the rules. Overall, by comparison to conventional purely numerical classifiers, the knowledge-based system method is more subjective because estimates of support values have to be made by the developer. The value of the technique must be proven by using ground truth data to verify the accuracy of the product. In our case this has been done and has demonstrated a significant improvement which ultimately will be useful from an operational point of view.

References

BOISSIN, B., and GARDELLE, J. P., 1986, SPOT localization accuracy and geometric image quality, Earth remote sensing using the Landsat Thematic Mapper and the SPOT sensor systems. *Proceedings of S.P.I.E.*, **660**, 83–92.

CIBULA, G. C., and NYQUIST, M. O., 1987, Use of topographic and climatological models in a geographical data base to improve Landsat MSS classification for Olympic National Park. *Photogrammetric Engineering and Remote Sensing*, **53**, 67–75.

GOODENOUGH, D. G., GOLDBERG, M., PLUNKETT, G., and ZELEK, J., 1987, An expert system for remote sensing. *I.E.E.E. Transactions on Geoscience and Remote Sensing*, **GE-25**, 349–360.

GORDON, J., and SHORTLIFFE, E. H., 1985, A method for managing evidential reasoning in a hierarchical hypothesis space. *Artificial Intelligence*, **26**, 323–357.

HARALICK, R. M., SHANMUGAM, K., and DINSTEIN, I., 1973, Textural features for image classification. *I.E.E.E. Transactions on Systems, Man and Cybernetics*, **SNC-3**, 610–621.

HART, A., 1986, *Knowledge Acquisition for Expert Systems* (London: Kogan Page).

HAYES-ROTH, F., WATERMAN, D. A., and LENAT, D. B., 1983, *Building Expert Systems* (Reading, Mass: Addison-Wesley).

HEUVELINK, G. B. M., BURROUGH, P. A., and STEIN, A., 1989, Propagation of errors in spatial modelling with GIS. *International Journal of Geographical Information Systems*, **3**, 303–322.

HUTCHINSON, C. F., 1982, Techniques for combining Landsat and ancillary data for digital classification improvement. *Photogrammetric Engineering and Remote Sensing*, **48**, 123–130.

KERAVNOU, E. T., and JOHNSON, L., 1986, *Competent Expert Systems* (London: Kogan Page).

KONTOES, C., and STAKENBORG, J., 1990, Availability of cloud-free Landsat images for operational projects. The analysis of cloud-cover figures over the countries of the European Community. *International Journal of Remote Sensing*, **11**, 1599–1608.

LEE, T., RICHARDS, J. A., and SWAIN, P. H., 1987, Probabilistic and evidential approaches for multisource data analysis. *I.E.E.E. Transactions on Geoscience and Remote Sensing*, **GE-25**, 283–293.

LUNETTA, R. S., CONGALTON, R. G., FENSTERMAKER, L. K., JENSEN, J. R., MCGWIRE, K. C., and TINNEY, L. R., 1991, Remote sensing and geographic information system data integration: error sources and research issues. *Photogrammetric Engineering and Remote Sensing*, **57**, 677–687.

MCKEOWN, D. M. JR., 1987, The role of artificial intelligence in the integration of remotely sensed data with geographic information systems. *I.E.E.E. Transactions on Geoscience and Remote Sensing*, **GE-25**, 330–347.

PEACEGOOD, G., 1989, A Knowledge-Based System for Extraction and Recognition of Linear Features in High Resolution Remotely-Sensed Imagery. Unpublished Ph. D. Thesis, Council for National Academic Awards/Kingston Polytechnic UK.

SHAFER, G., 1976, *A Mathematical Theory of Evidence* (Princeton: Princeton University Press).

SRINIVASAN, A., and RICHARDS, J. A., 1990, Knowledge-based techniques for multi-source classification. *International Journal of Remote Sensing*, **11**, 505–525.

TROTTER, C. M., 1991, Remotely-sensed data as an information source for geographical information systems in natural resource management: a review. *International Journal of Geographical Information Systems*, **5**, 225–239.

VAN CLEYNENBREUGEL, J., FIERENS, F., SUETENS, P., and OOSTERLINCK, A., 1990, Delineating road structures on satellite imagery by a GIS-guided technique. *Photogrammetric Engineering and Remote Sensing*, **56**, 893–898.

WALSH, S. J., LIGHTFOOT, D. R., and BUTLER, D. R., 1987, Recognition and assessment of error in geographic information systems. *Photogrammetric Engineering and Remote Sensing*, **53**, 1423–1430.

WILKINSON, G. G., and MÉGIER, J., 1990, Evidential reasoning in a pixel classification hierarchy—a potential method for integrating image classifiers and expert system rules based on geographic context. *International Journal of Remote Sensing*, **11**, 1963–1968.

WU, J.-K., CHENG, D.-S., WANG, W.-T., and CAI, D.-L., 1988, Model-based remotely-sensed imagery interpretation. *International Journal of Remote Sensing*, **9**, 1347–1356.