PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A
MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Automatic Interpretation Strategies for Synthetic Aperture Radar Images [and Discussion]

S. Quegan, A. J. Rye, A. Hendry, J. Skingley, C. J. Oddy and D. Lane

Phil. Trans. R. Soc. Lond. A 1988 **324**, 409-421
doi: 10.1098/rsta.1988.0029

Phil. Trans. R. Soc. Lond. A **324**, 409–421 (1988)
Printed in Great Britain

Automatic interpretation strategies for synthetic aperture radar images

By S. Quegan[1], A. J. Rye[2], A. Hendry[2], J. Skingley[2] and C. J. Oddy[2]

[1] *Department of Applied and Computational Mathematics, The University, Sheffield S10 2TN, U.K.*
[2] *GEC Research Ltd, West Hanningfield Road, Gt. Baddow, Chelmsford, Essex CM2 8HN, U.K.*

[Plate 1]

Effective use of the large amounts of timely data to be provided by the coming generation of space-based radars will require automatic methods of image interpretation. The key to such interpretation is an image representation, based on low-level operations, which can support the introduction of high-level (rule-based) knowledge. This representation, described in this paper as a segmented image database (SID), is dependent on the performance of the low-level operations (segmentation, edge-detection and thin-line-detection) which generate it. Methods of quantifying the performance of these operations are described. Use of the SID to support classification based on context, and image–map matching that uses image structure, rather than geometrical matching, are demonstrated.

1. Introduction

The use of synthetic aperture radar (SAR) to investigate and monitor the Earth's surface is of increasing interest, because it can provide images with high spatial resolution unaffected by cloud or darkness, and it responds directly to motion in the observed scene. High resolution (of the order of metres) is important in many civilian and military applications.

Penetration of cloud is essential for routine land-use monitoring in, for example, Europe or the equatorial rain forests. The ability to operate through cloud and at night is particularly relevant at high latitudes, where there is a need to monitor ice and icebergs (Skriver & Gudmandsen 1986). The Doppler shift imposed on the return signal by a moving scatterer can be used to selectively image moving small objects (Raney 1971; Freeman 1984) or to image waves in the ocean (Rotheram 1983).

These attractive possibilities offered by SAR have provided the impetus to several major satellite projects for Earth observation, intended to be implemented before the end of the century. The European Space Agency, Japan and Canada are all developing free-flying SARS (*ERS*-1, *JERS*-1, and *Radarsat* respectively); the Space Shuttle has been used to carry SARS (the *SIR*-A and *SIR*-B missions), and should continue to do so when missions resume; and a SAR is seen as an essential element in the Space Station–Polar Platform.

The unique capabilities of SAR data unfortunately exact a price, whose value we do not fully know; this is learning how to interpret the data, which has special properties inherent in the use of coherent radiation with wavelengths of the order 1–30 cm. Tonal and textural properties of the images reflect volume and surface scattering in the elements of the scene. Image geometry is intimately connected with relative motion of the scene and SAR platform. Lastly, and of particular importance for the work presented below, SAR images are affected by a signal-dependent noise known as speckle (for a discussion of this phenomenon see, for example, Dainty

31-2

(1984)). Speckle is the first obstacle to be overcome in any attempt to interpret SAR images automatically.

Image interpretation has four basic elements:

(i) definition of objects of interest, e.g., fields, texture, areas of change, edges;

(ii) characterization of objects of interest *in image terms*;

(iii) use of (ii) to derive methods of seeking objects of interest;

(iv) referencing of the image and objects within it in a form suitable for comparison with other data.

Element (i) has the obvious implication that image interpretation is application-dependent. However, certain operations seem essential to many applications, notably image segmentation, edge-detection, line-detection, and post-segmentation operations, such as measurement of segment statistics and segment classification. Element (ii) can present great difficulties. Even such primitive concepts as lines, which are handled effortlessly by the eye and brain, seem to defy any simple description in terms of image statistics or elementary structures. Measurement of image properties is often not easy, because it is not clear what should be measured. However, there have been some successes; e.g., image texture has been successfully characterized by combining measurements with an image model (Oliver 1986). Element (iii) is often hindered by the difficulty in producing proper image characterization, and we are often obliged to use paradigms to drive algorithm development, with ensuing questions over the applicability of the results to real images.

Element (iv) is fundamental to use of the data. It is unlikely that the image is the only or even the primary data source of interest. Hence it is essential that the image structure be encoded in a form suitable for comparison and combination with other data. This encoding may be in terms of pixels (with ensuing requirement for geometric corrections) or in terms of a relational database. The generation and use of such a database is a central concern of this paper, because it provides the framework for introduction of knowledge-based rules into image interpretation. The low-level operations used to generate the database are discussed in the first half of the paper. The emphasis here is on quantifying and optimizing the performance of algorithms, thus defining the constraints in the data available to the database. The second half of the paper is concerned with the use of the database to compare images, and to use contextual information in classifying image segments.

2. Low-level image processing

By low-level image processing we will mean those operations necessary to generate a segmented image database, i.e. edge detection, thin-line detection and image segmentation. Other potentially important operations, such as statistical correlation algorithms, are not discussed (but see Deane *et al.* 1987). Although describing these operators as low level, we do not wish to imply that they are easy to develop or do not rely on sophisticated techniques. The opposite is in fact the case, and the search for reliable methods operating in speckle still continues. In the next section we summarize our approaches to quantitative evaluation of algorithms.

2.1. *Edge detection and segmentation*

As our working definition of an edge, we use a sharp discontinuity in scene reflectivity, which becomes smoothed as a result of the point spread function of the SAR imaging process

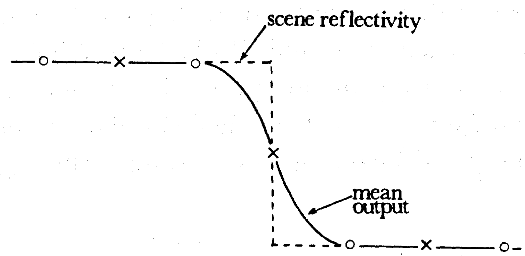

FIGURE 1. Mean output in the SAR image of a sharp reflectivity edge. Two possible digital samplings are noted by ○ and ×. The mean intensity indicated would be subject to speckle.

(figure 1). Edges and segments are dual concepts, each giving rise to the other, and segmentation can be carried out by edge-detection, segment growing or a combination of both approaches (White 1985; Oddy & Rye 1983).

An obvious question when faced with the large variety of methods mooted for image segmentation concerns their relative merits and reliability. Because this is fundamental to operational use of these methods, we have sought to define procedures to compare different segmentation schemes. This has relied extensively on simulated images, because (*a*) the definition of the correct or desirable segmentation of a real image is not available, and (*b*) segmentation procedures normally involve nonlinear and/or adaptive operations which are not easy to investigate analytically. Though there are a number of properties of segmentation that may be regarded as desirable (e.g. area-preservation, edge position), we have concentrated on the two most fundamental properties, i.e. the generation of spurious segments (false alarms) and the merging of two distinct segments (edge-destruction). It has at present only proved possible to produce a method of comparing segmentation algorithms of the form shown in scheme 1.

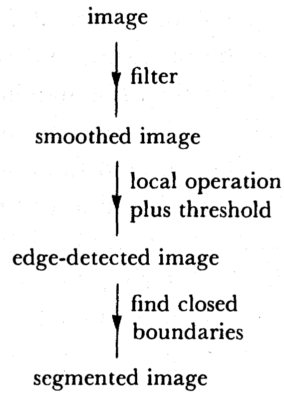

SCHEME 1. Edge-detection segmentation scheme.

This omits the important methods based on segment growing, and such possibly important techniques as iterated filtering before edge detection. At present we do not know how to extend our methods to these cases, except by a massive simulation programme.

The approach used is very simple in concept. Because edge-detection thresholds the output of some operation applied to the filtered image (see scheme 1), a false-alarm measure can be derived by examining the effects of applying the operation to filtered pure speckle images. This yields, for a particular algorithm, a relation between false-alarm rate and the threshold used.

Comparison of algorithms can then proceed on the basis of a comparable false-alarm rate (and hence fixed thresholds) for each algorithm. With the thresholds fixed, it is necessary to investigate the edge preservation or destruction properties of each algorithm. This is carried out by first generating simulated images which are dense in edges (figure 2), and then applying each algorithm to enough simulated images to ensure statistically significant results.

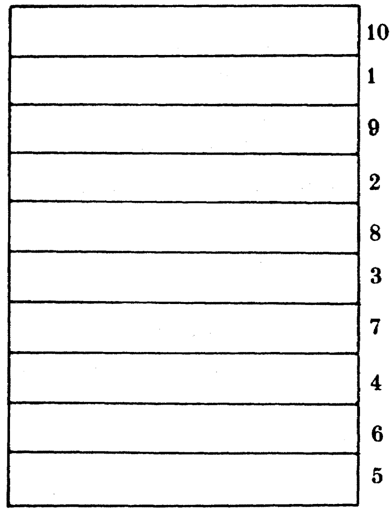

FIGURE 2. Edge image used for simulation. The numbers on the right give the relative brightnesses of the horizontal bands. Speckle is added before simulation is carried out.

These methods have been applied to a number of proposed schemes for edge-detection based segmentations, and have proved particularly useful for comparing schemes that have free parameters in the filtering stage of the operation. Of particular interest are results obtained with the Frost *et al.* (1982) filter combined with a matched edge detector (for a fuller discussion see Quegan *et al.* 1987). These indicate that the optimum performance is achieved (for any false alarm rate) by permitting the filter to tend to an average filter. A similar conclusion has been obtained from methods that use a filter due to Oddy & Rye (1983) and the modified Lee 2-σ filter (Lee 1983) (see Hendry *et al.* 1985), with the implications that (i) nothing is gained by greater sophistication in the filtering, hence simple and fast methods can be used, and (ii) the edge-smoothing inherent in the average filter does not prevent post-filtering edge-detection. These conclusions have been tested on real data, with results that confirm the findings from simulated data.

It must be stressed, however, that we cannot as yet define an optimum edge-detector or segmenter, only put limits on the reliability of one class of such algorithms. These algorithms have also been applied to only a limited class of data. The effects of oversampling the data have been investigated (the major conclusion being that any form of oversampling degrades the performance of the algorithms), but edge-orientation effects and the effects of resampling have not been considered. It must also be noted that the test pattern (figure 2) was designed to keep the results at different edges statistically independent, hence the results may not be applicable as the stripes become narrower. There is therefore a continuing need to develop the methods described above, to extend the range of image data for which they may be considered valid.

There is also a clear need to develop methods of assessing the performance of other classes of algorithms, if we are to provide the best techniques for applications purposes.

2.2. *Thin-line-detection*

Thin lines are important scene elements (roads, hedgerows, field boundaries, etc.), which contribute to human image segmentation. We characterize thin lines as features one to three pixels wide, which may separate regions of the same or different intensities, and which may traverse several different regions. Conceptually they could be regarded as segments, but in practice special-purpose algorithms are necessary to locate them, because their statistical separability from the background requires searching along particular directions (both because of their narrowness, and because they often do not differ greatly in intensity from the background in real SAR images).

Three types of line-detection methods have been considered, namely (1) local operators or masks (Vanderbrug 1976; Gurney 1980); (2) dynamic programming (Wood 1985); and (3) global transforms (Deans 1981; Duda & Hart 1972). Of these, the first method seems inapplicable to SAR images unless very large window sizes are used, and it does not perform well when the ratio of line to background intensity is near one (Hendry *et al.* 1985). The second method is applicable to SAR images, can cope readily with curved lines, is computationally expensive and needs further development to cope with linking of line segments (Wood 1985; Skingley 1986). The third method uses the Radon or Hough transforms, and has been successfully applied to SAR images (Murphy 1987; Hendry *et al.* 1985; Skingley 1986). The application of these two transforms to line-detection yields equivalent results (the Hough transform is a special case of the Radon transform). The Radon transform is faster (Murphy 1985), but the Hough transform is conceptually simpler, and can be readily analysed (Skingley & Rye 1987; Skingley 1986), as described below.

The Hough transform parametrizes all possible lines in the image in terms of R, the distance from the centre of the image, and θ, the orientation (see figure 3), and sums pixel values along

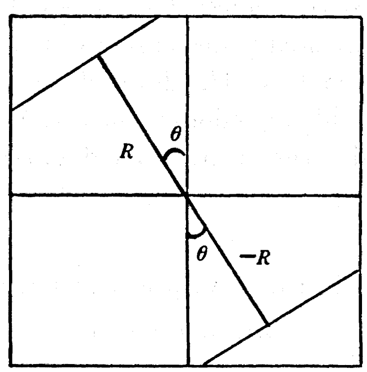

FIGURE 3. The Hough transform coordinate scheme.

these lines, assigning mean intensity of the line to each (R, θ) point. Bright and dark lines in the image therefore transform to peaks and troughs in the (R, θ) domain. Detection probabilities for such peaks and troughs can be derived analytically (see Skingley & Rye 1987). Figure 4 shows calculated detection probabilities for lines of various lengths and brightnesses. Comparison

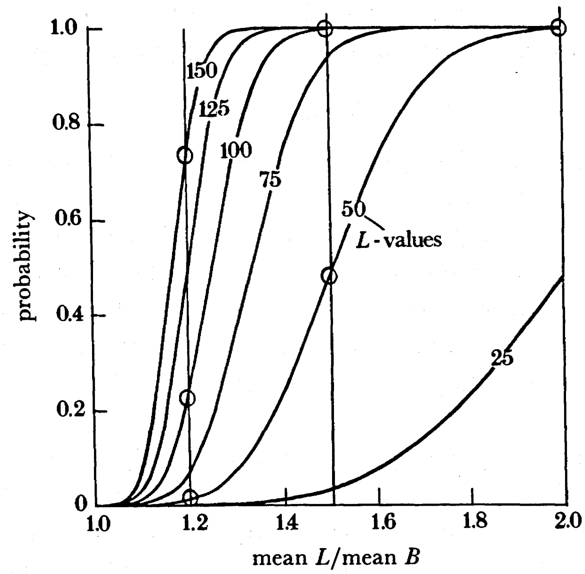

FIGURE 4. Calculated detection probabilities for lines of various lengths (indicated on each curve) embedded in a uniform background, with Rayleigh distributed speckle. The abscissa gives the ratio of line, L, to background brightness, B. The maximum possible line length is taken to be 150 pixels.

with performance on simulated images confirmed these theoretical results, though an excessive number of spurious detections were noted. These can be understood in terms of correlation between samples in the transform space (Skingley 1986).

Detection alone is not enough, because it is clear from the formulation of the Hough transform that short bright lines yield the same signature as long, less bright lines. Work on end-point detectors and removal of false alarms is still continuing.

The applicability of results based on simulation to real data must be carefully considered, because at least some lines observed in real images do not conform to the simple speckle model adopted in the simulation (Quegan *et al.* 1988). The Hough–Radon transform has been applied to real SAR images, and successfully used to detect ship wakes and quite subtle patterns of lines corresponding to rides through woodland (Hendry *et al.* 1985; Murphy 1987). Whether these results are consistent with theoretical predictions is unclear. At present, also, thin-line-detection is not properly incorporated in the segmented image database, and more work is required to accomplish this.

3. High-level processing

3.1. *Introduction*

Segmentation techniques, such as those described above, generate an initial spatial description of a scene in a form suitable for further analysis of the image data. The development of techniques that manipulate and process image regions has been supported by the construction of a standard database of segmented images. This database is described in the next section.

The spatial representation of a scene may be matched to other geographic data, such as maps, to generate a mapping function between the data, as a precursor to the fusion of multisensor and multitemporal images. One approach to matching segmented images is described below.

It has been recognized, for many years, that contextual information is a valuable tool for producing land-use classification from remote-sensing images. However, until recently many classification methods used only the spectral response of single pixels and their neighbours. Segmentation provides the necessary primitives for the development of classifiers that use the contextual information present in a scene and a simple demonstration system for woodland is presented.

3.2. *A segmented image database* (SID)

A database for storing region information from an image after it has been segmented has been designed by GEC Research (Cruse *et al.* 1984). This Segmented Image Database (SID) stores information on regions, region boundaries, adjacencies and attributes in a flexible manner, thus allowing the database to be accessed in a number of ways for different analysis operations. The SID can be easily updated after these operations, and can store images from any source.

Two implementations of the SID have been designed in FORTRAN and PROLOG. The FORTRAN implementation is suitable for low-level image-analysis operations and the PROLOG implementation is suitable for high-level tasks.

In FORTRAN, the database is basically a tree structure, with additional pointers between branches, and uses a dynamic memory management system. The data are contained in three main indexes: a region index, containing all the regions in the image, a line segment index, containing all the line segments that make up the boundaries of the regions, and an attribute index, containing information about all attributes that have been calculated for any of the regions. These attributes may include mean grey level, texture measures, shape measures, polygonal approximations, symbolic labels and any string or vector the user wishes to associate with a region. A suite of subroutines has been written to interact with the SID to create, read, delete or update records.

The network structure of the FORTRAN implementation may be converted to a relational database in PROLOG. This allows very simple facilities for interrogating the database and rapid assessment of new algorithms.

Many different image analysis operations may be performed with the SID, e.g. attribute extraction, region merging, region classification, and shape recognition. It is also easy to create graphical descriptions of the segmented image from SID, such as a region adjacency graph, which can be used for structural matching techniques.

3.3. *Structural matching*

Structural matching is a technique which matches images or images and maps after they have been segmented. The segmented images are represented by attributed region adjacency graphs (RAGS), in which the nodes of the graph represent the regions, and the branches represent the adjacency relations. Both the nodes and the branches of the graph may have attributes associated with them. These may be the size, shape or texture of regions and the direction of adjacency. The images are then matched by matching their RAGS.

This region-based approach to matching will be less sensitive to speckle than methods relying on local correlation, as long as the RAG is of 'sufficient' accuracy. It is not yet possible to relate the performance measures discussed in §2 to a proper description of the accuracy of the RAG. However, the method of matching described below seems capable of coping with highly distorted RAGS, as will be clear. This method may also be used for multisensor and multi-temporal image matching when pixel-by-pixel comparisons are not sufficiently robust.

The matching strategy developed by GEC Research (Oddy 1986; Oddy *et al.* 1988) involves first matching a number of control regions from the images. These should be distinctive in some way (for example in size) and ideally should be well distributed throughout the images. The second stage then attempts to match regions adjacent to those already matched, until no more regions can be matched. Thus the matching proceeds outwards from the control regions.

The rules for matching two regions are based on a comparison of region attributes and region shape. Region shape is determined by region vectors from a polygonal approximation of the boundaries.

The matching process allows for partial matching of regions. This is necessary because a complete match between two RAGS is very unlikely, due to genuine differences on the ground between images taken at different times, to different sensors discriminating between different surface types and to imperfect segmentations.

Output from the matching scheme is initially a list of matching region pairs. These can then be used to generate automatically lists of control points which may be used to determine a mapping function from one image to the other image or map. These control points may be the midpoints of matching vectors, or the corners of matching fields.

The above scheme has been implemented in a mixture of FORTRAN and PROLOG on a VAX 11/750. The programs used to generate the SID files containing the attributed RAG are written in FORTRAN. These files are then converted to a PROLOG relational database which is used as input to the structural matching. The matching algorithm is written in PROLOG and runs in the POPLOG system.

The technique described above has been used to match an airborne SAR-580 image (figure 5, plate 1) to a digital map (not shown) of an area near Feltwell in Norfolk. The map was digitized at GEC Research from an ordnance survey 1:25000 map. The map area was originally surveyed in the 1920s, and the SAR image was acquired in 1982; several differences are evident between the two datasets.

The map and the image were segmented using techniques developed at GEC Research (Keen 1986; Cruse *et al.* 1986). Figures 6 and 7 show line drawings of these segmentations.

Matching was performed with the map as the reference image. Rough orientation and scale information were input to the process. As many control regions as possible were then matched, followed by matching of adjacent regions, with gradual relaxation of the shape matching thresholds. The results of this matching are shown in figures 6 and 7. Matching regions are assigned the same letter. In spite of the fact that only a partial segmentation could be generated the matching performed well for regions which were extracted. In particular, pairs A, V and F match well. A number of partial matches are also obtained, in areas that can easily be matched manually (e.g. pairs y, X, N and T) and in more complicated areas (e.g. pairs w, n, u and r).

A set of control points from the midpoints of matching vectors have been generated and used to define an affine transformation from image to map. The outcome was in exact coincidence with the transformation obtained by manual selection of control points, an excellent result in view of the poor segmentation of the SAR-580 image that was provided for the matching process.

FIGURE 5. SAR-580 image near Feltwell, Norfolk.

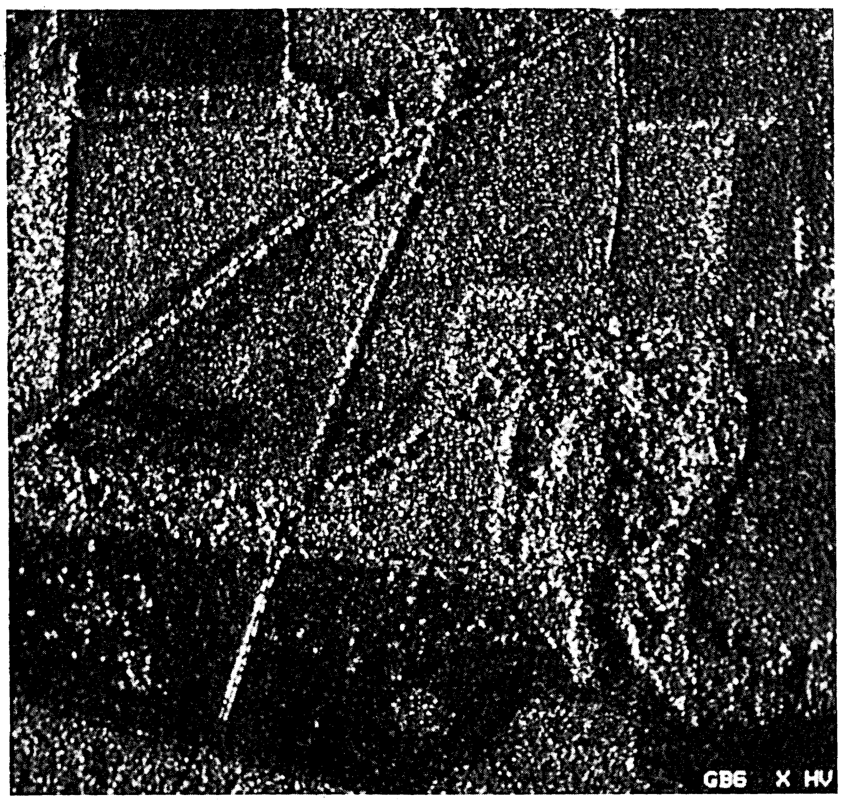

FIGURE 8. SAR-580 image of woodland area.

(*Facing p.* 416)

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

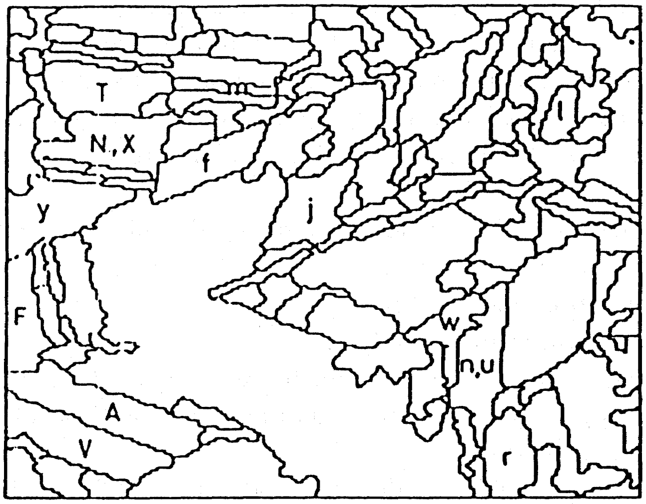

FIGURE 6. Segmentation of lower left section of figure 5 (the sharp river bend is in the top right of the segmentation).

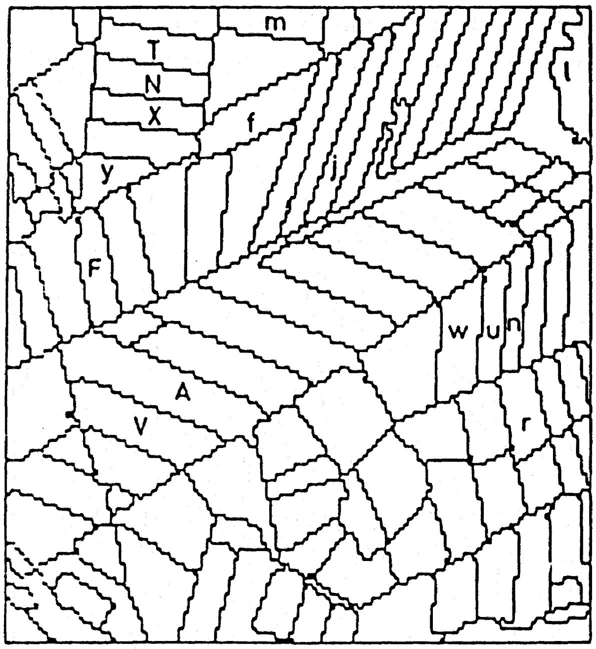

FIGURE 7. Segmentation of the ordnance survey digital map corresponding to figure 6.

3.4. *A simple contextual expert system for woodland*

This section describes a simple demonstration expert system which uses contextual and structural information to identify hedges and woodland in a high-resolution SAR-580 image.

High-resolution (3 m) airborne SAR images often contain shadow and highlight regions when imaging woodland and hedgerows. These linear features are typically 2–3 pixels wide and are often partly detected by segmentation algorithms as elongated regions. A system that uses a number of simple rules to interpret a scene has been implemented in PROLOG by using POPLOG on a VAX11/750.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

The input to the woodland expert system consists of a segmented image which includes the region mean grey level, region adjacencies, elongation and direction. The region adjacency directions are defined by the mean direction of the common edge(s) between two regions. The elongation measure used here is based on the minimum spanning rectangle of a region; the ratio of length to width gives the elongation and the region direction is given by the major axis of the rectangle.

The expert system developed is very simple and uses the following basic rules:

1. shadows are long, thin and dark;
2. highlights are long, thin and bright;
3. hedges are highlights in front of shadows;
4. woods are regions in front of shadows which are not part of a hedge.

These basic rules are supplemented by other rules to define 'in front of' for example. Thus:

a wood is IN FRONT OF a shadow

IF the adjacency direction of the shadow to wood (corrected for the flight direction) is less than 75°

AND the azimuth length of the shadow overlaps at least 30% of the azimuth length of the wood.

The test image used in this study was a multipolarization image (figure 8, plate 1) of a rough textured woodland area with tree-lined roads and heathland. The initial segmentation is shown in figure 9.

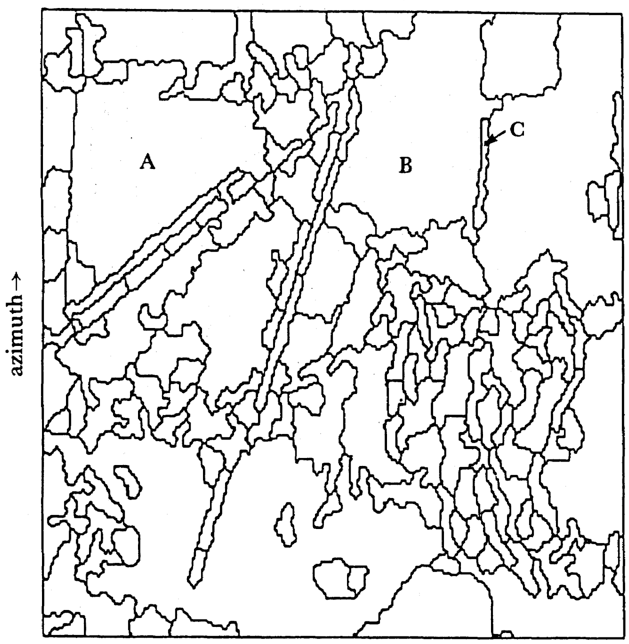

FIGURE 9. Segmentation of figure 8.

First results from the system are shown in figure 10. The hedges are clearly delineated. It is interesting to note that one road is imaged as a shadow in front of a highlight whereas the other road is a highlight in front of a shadow. The first shadow implies a wood in region A but a map of the area suggests that this is not correct. Region B is shown as a wood, because shadow C

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

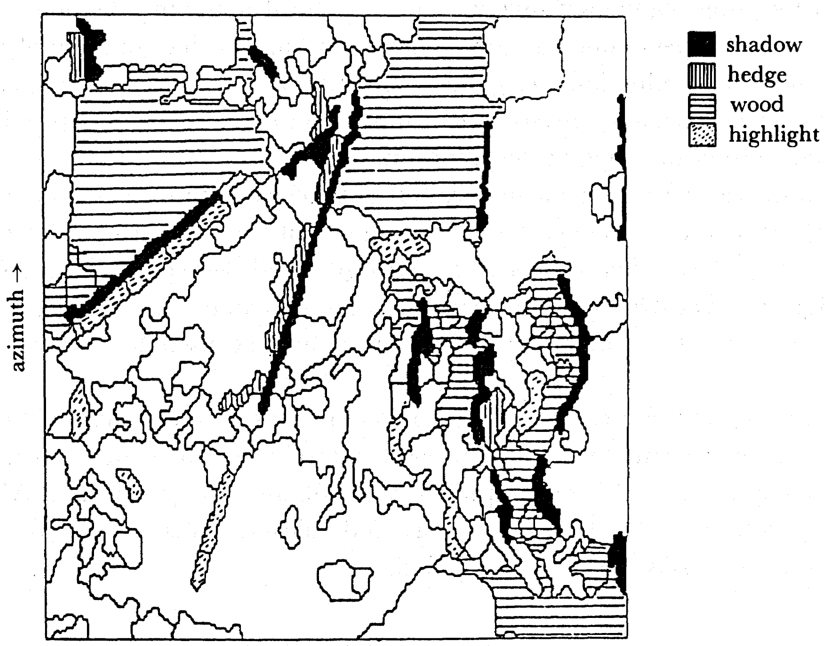

FIGURE 10. Classification of figure 8 using context.

is detected by the segmentation, but the highlight or hedge next to it in the image is missed by the segmentation and region B is incorrectly classified. The interpretation of the textured woodland area, in centre right of the image, is promising.

This woodland expert system shows potential, is very simple and demonstrates how even a small number of rules can be used to assist the interpretation of an image. The interpretation given here relies totally on the structural information in a segmented image and its performance is dependent on the quality of the initial segmentation. Most of the failures of the system are due to a failure by the segmentation scheme to detect some very weak features (highlights and shadows) in the image. In some cases it is very difficult to interpret the original image manually without prior knowledge of the test site or a map of the area.

4. Conclusions

Automatic interpretation of SAR images requires an image representation to support the introduction of knowledge into the interpretation. The segmented image database (SID) provides such a representation. Its formation relies on low-level operations, whose performance imposes constraints on the inferences possible from the SID. Quantitative performance measures for some of the necessary low-level operations are possible. In the case of segmentation based on edge detection, these measures indicate how to optimize the behaviour of the algorithms considered. The methods adopted do not apply to all segmentation strategies, and further work is required. Similarly, the methods allowing a theoretical analysis of the Hough transform as a line-detector cannot obviously be applied to other detection algorithms.

The use of the SID as a means of introducing knowledge into the interpretation of SAR images is at a very early stage. However, the work that has been carried out demonstrates the potential offered by this technology. It has been shown that the matching of image and maps may be achieved by using segmented images. The structural matching methods appear to be robust and not entirely dependent on a 'good' segmentation. The development of contextual classifiers *does* require the production of 'good' segmentations, and work to quantify and improve segmentation performance needs to continue. Contextual classifiers provide important additional information for image interpretation.

The SID provides a means for integration of knowledge into the segmentation process, i.e. it can check its own consistency. This is an important area requiring a lot more work. Nazif & Levine (1984) have demonstrated knowledge-based segmentation methods using general knowledge of object structures. In remote sensing we can obtain specific, if noisy, knowledge of ground structures from map data, and this should be used to reinforce and support edge detection.

The development of a comprehensive image interpretation system will require a number of different expert systems that can cope with many different scenes and image interpretation tasks. Such systems have been proposed by Goldberg *et al.* (1985) and ARTS-IP (1986) but are clearly long-term goals. In the short and midterm special purpose systems for specific applications and tasks may be anticipated.

References

ARTS-IP 1986 Adaptive real time strategies for image processing. Report on Esprit Project 867.

Cruse, D., Oddy, C. J. & Wright, A. 1984 In *Proc. 7th ICPR*, pp. 494–496. IEEE Computer Society Press.

Cruse, D., Hendry, A., Rye, A. J. & Barnes, I. 1986 Marconi Tech. Rpt. MTR 86/62.

Dainty, J. C. (ed.) 1984 *Laser speckle and related phenomena*. Berlin: Springer-Verlag.

Deane, G. C., Quegan, S., Churchill, P. N., Hendry, A., Evans, R. & Trevett, J. W. 1987 In *Proc. ESA Workshop on SAR Applications*, *Frascati*, SP-264, pp. 95–102.

Deans, S. R. 1981 *IEEE Trans.* **PAMI-3**, 185–188.

Duda, R. O. & Hart, P. E. 1972 *Communs Ass. comput. Mach.* **15**, 11–15.

Frost, V. S., Stiles, J. A., Shanmugan, K. S. & Holtzmann, J. C. 1982 *IEE Trans.* **PAMI-4**, 157–165

Freeman, A. 1984 Proc. IGARSS-84, ESA Pubn. SP-215, pp. 65–70.

Goldberg, M., Goodenough, D. G., Alvo, M. & Karan, G. M. 1985 *Proc. IEEE* **73**, 1054–1063.

Gurney, C. M. 1980 *IEE Trans.* **GE-18**, 204–211.

Hendry, A., Quegan, S., Rye, A. J., Skingley, J. & Churchill, P. 1985 Marconi Tech. Rpt MTR 84/59.

Keen, A. 1986 GEC Technical Note PALS/WN 86.

Lee, J. S. 1983 *IEEE Trans.* **SMC-12**, 83–89.

Murphy, L. M. 1987 *Patt. Recog. Lett.* **4**, 279–284.

Nazif, A. M. & Levine, M. D. 1984 *IEEE* **PAMI-6**, 555–577

Oddy, C. J. 1986 Post segmentation image–image matching. *Marconi Tech. Memo* ITM 85/119.

Oddy, C. J. & Rye, A. J. 1983 *Patt. Recog. Lett.* **1**, 443–449.

Oddy, C. J., Skingley, J. & Howard, P. D. 1988 In *Proc. ESA workshop on Image rectification techniques for space-borne SAR*. (In the press.)

Oliver, C. J. 1986 *Inverse Problems* **2**, 481–518.

Quegan, S., Lamont, J. M., Miller, R. J., Meadows, P. J., Wright, A., Veck, N. J., Churchill, P. N., Holmes, G., Trevett, J. W., Sloggett, D. & Hamer, P. 1987 In *Proc. ESA Workshop on SAR Applications*, *Frascati*, SP-264, pp. 37–45.

Quegan, S., Hendry, A. & Skingley, J. 1988 In *Proc. IMA Conf. Mathematics and its Applications in Remote Sensing*. (In the press.)

Raney, R. K. 1971 *IEE Trans.* **AES-7**, **3**, 499–505.

Rotheram, S. 1983 In *Satellite microwave remote sensing* (ed. T. Allen), ch. 10. New York: Wiley.

Skingley, J. 1986 Marconi Tech. Rpt. MTR 86/72.

Skingley, J. & Rye, A. J. 1987 *Patt. Recog. Lett.* **6**, 61–67.

Skriver, H. & Gudmandsen, P. 1986 Final Rpt to ESA, Contract No. 5442/83/D/IM(SC).

Vanderbrug, G. J. 1976 *IEEE Trans.* **GE-14**, 37–44.

White, R. G. 1985 RSRE Memo no. 3900.

Wood, J. W. 1985 RSRE Memo no. 3841.

Discussion

D. Lane (*Intelligent Automation Laboratory, Department of Electrical and Electronic Engineering, Heriot-Watt University, Edinburgh, U.K.*). The effectiveness of a segmentation strategy usually depends upon the appropriate selection of one or a number of parameters, such as thresholds, for a given class of data. A truly automatic interpretation strategy should not rely on any manual input of these parameters. What prospect is there, therefore, for automatically selecting processing parameters, by using perhaps high-level knowledge?

S. Quegan. The first part of the paper described an approach to defining optimal parameters on purely statistical grounds, making use only of general properties of the data. For an initial segmentation such optimized procedures seem appropriate; they should not appeal to high-level knowledge, but to known properties of the data, based on the features of the sensor. The methods described offer hope that optimal parameters can be defined, though at present we can only cope with a limited class of algorithms and of data. Where high-level knowledge should surely be introduced is in assessing and perhaps adjusting the initial segmentation. This may involve reprocessing with altered parameters, but may instead involve closer scrutiny of selected areas, or imposition of rules to refine the segmentation. That is an interesting and difficult area of study.

MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES
THE ROYAL SOCIETY A
PHILOSOPHICAL TRANSACTIONS OF

FIGURE 5. SAR-580 image near Feltwell, Norfolk.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

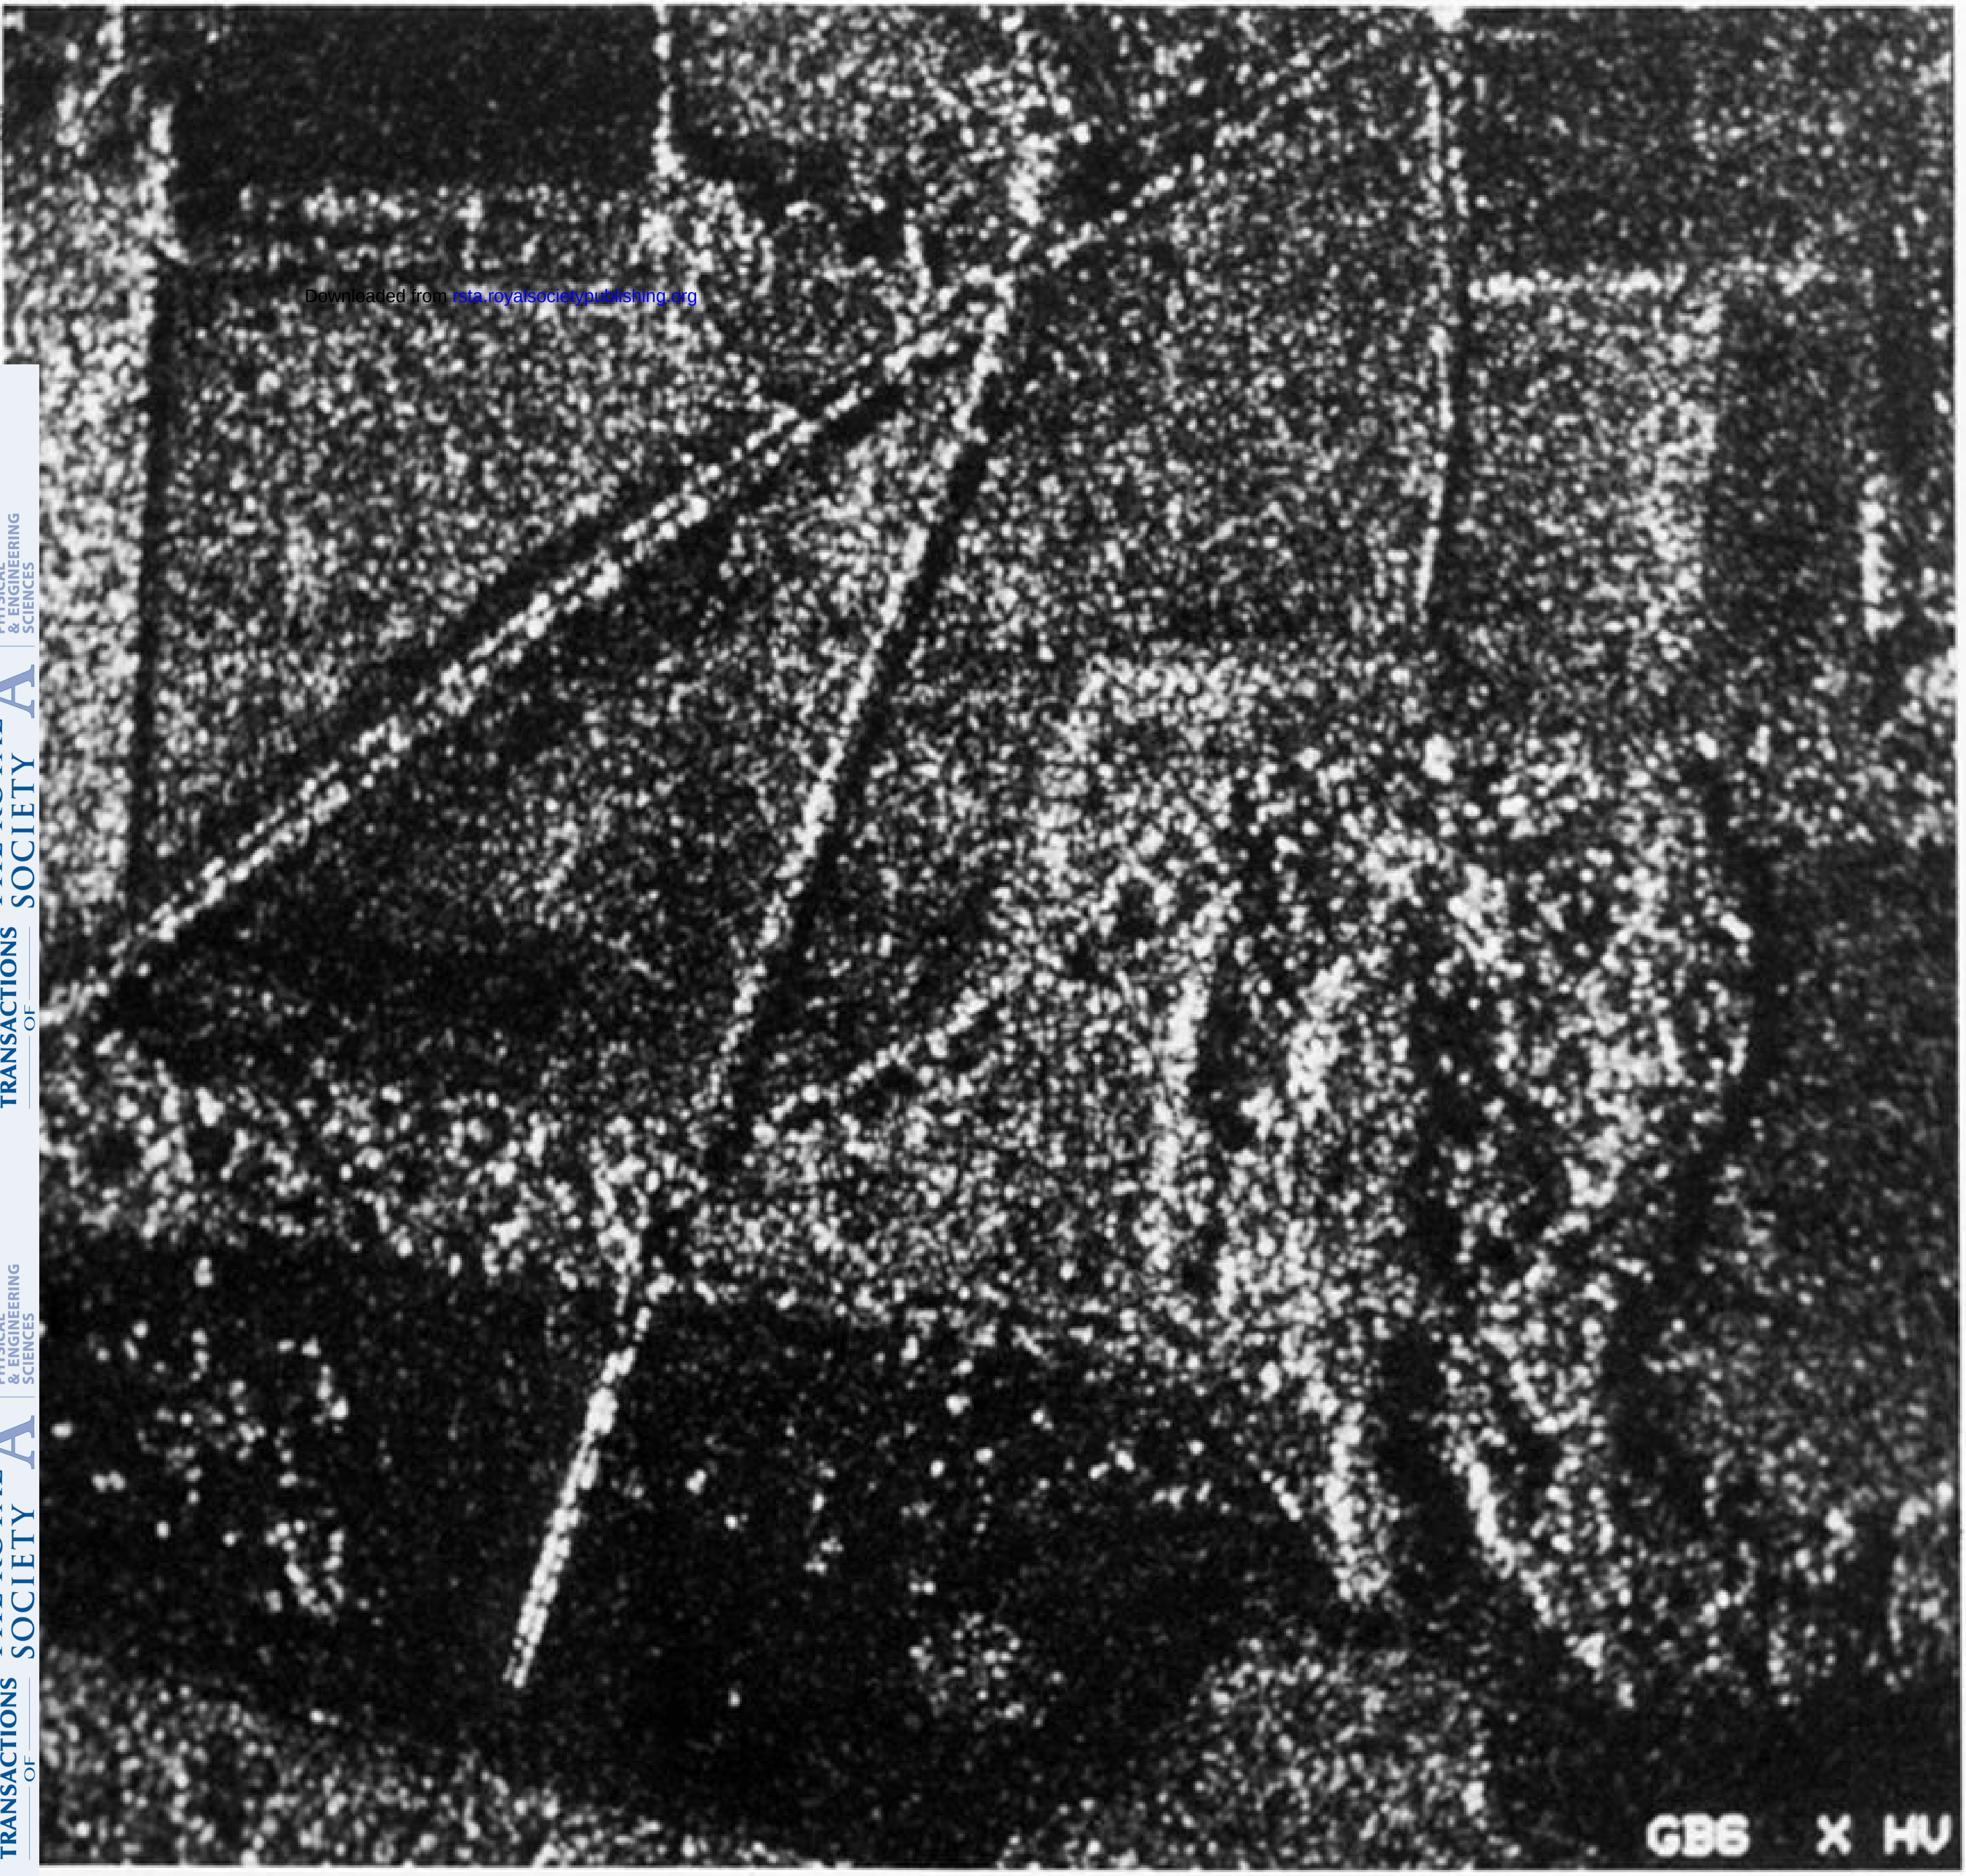

FIGURE 8. SAR-580 image of woodland area.