Rule-Based Classification of Water in Landsat MSS Images Using the Variance Filter

Paul A. Wilson

Abstract

The variance filter is a textural algorithm capable of distinguishing between flat, uniform water bodies and cloud or mountain shadows when applied to satellite imagery. The filter output forms the basis of rules used by a knowledge-based classifier, which segments water adaptively. The use of the filter in the unsupervised classification of water is demonstrated on two spectrally varied Landsat MSS images. The same images are segmented using a conventional thresholding algorithm. The two algorithms identify a similar proportion of the water pixels in both images; however, the rules-based algorithm does not generate any false positives, whereas the threshold algorithm misclassifies many shadow pixels as water. The rules-based algorithm is less efficient at finding small lagoons and swamps than at finding large water bodies.

Introduction

This paper describes one algorithm which has been developed as part of a long term project. The long term goal of the research is to devise methods of partially analyzing Landsat data without human intervention or guidance. The rationale for this is the need to solve what is sometimes referred to as the *terabyte problem*. There is a need to process, by computer, the very large volumes of data being received from Earth observation satellites, given that there is a shortage of trained analysts.

The short term goal of this work is to reliably identify (segment) pixels of water in a Landsat multispectral scanner (MSS) image, and without obtaining any false positive pixels. Cloud shadow and western facing mountainsides are examples of features which are often confused with water by classifiers.

The author has elected to use any algorithm which will provide useful results; the range of algorithms investigated in the project to date includes statistical methods, texture algorithms, contextual algorithms, rules, production systems, fuzzy logic, genetic algorithms, and neural networks. This paper is confined to a description of one texture algorithm, two contextual algorithms, and several rules. The rules were developed by trial and error in order to overcome several difficulties and to meet several criteria which are described in the next section.

The experimental work has been carried out on Landsat MSS images because the data sets are less costly than Thematic Mapper (TM) data, and MSS images continue to be widely used for land-use and resources studies. Preliminary tests with TM images suggest that the algorithms will be equally successful with that type of data.

Traditionally, segmentation of water is carried out under human supervision using spectral information. In studies of the Great Barrier Reef Marine Park, Australia, Jupp *et al.* (1985) state, "Using Band 7, each image is separated into water and other areas (such as land and clouds, etc.) by a simple mask. This operates on the basis that Band 7 is totally absorbed by water, providing a distinct separation." Jupp *et al.* comment on the difficulty of separating water from steep hillsides with a western aspect or from cloud shadows. The technique suggested by Jupp *et al.* is to identify deep water visually and then find the minimum values in all the bands and use these values to define the mask. Cloud shadows and steep mountain sides are also identified visually and digitized out of the image manually. Moller-Jensen (1990) classifies water by the following simple rule, using Thematic Mapper data:

- Band 4 (infrared, wavelength 0.76 to 0.9 μm) < 45 in value (digital number), and
- Band 5 (infrared, wavelength 1.55 to 1.75 μm) < 35 in value.

He claims, using this method, that "only a few non-water pixels are misclassified."

The cursory treatment given by these, and other, workers to the segmentation of water implies that the problem is straightforward to solve. Under supervision, this is so in many instances, using the methods already described. However, even visual classification can at times be difficult. For example, a small crater lake on the western side of a mountain ridge is virtually impossible to distinguish from deep shadow. One of the images used in this study was deliberately chosen to include such a lake, Lake Euramo in the Danbulla State Forest, Atherton.

Automatic Segmentation of Water

One objective of this project is to establish some broad principles which can be used for the fully automatic segmentation of water in MSS images which have not been pre-processed, or even de-striped. The images we used have only been radiometrically corrected by the supplier, The Australian Centre for Remote Sensing (ACRES).

Most of the commonplace techniques for image interpretation rely on the presence of a human analyst. Methods such as contrast enhancement, histogram equalization, and low pass filtering are representative of those used to assist visual interpretation (Campbell, 1987; Richards, 1986). For fully automatic processing, a change of paradigm is required, with less emphasis on traditional methods and more emphasis on artificial intelligence; hence, the use of rules in this study.

The design criteria for the algorithms developed in this research are that the processing of the images should have no human guidance and that the images should not be pre-

Photogrammetric Engineering & Remote Sensing,
Vol. 63, No. 5, May 1997, pp. 485–491.

0099-1112/97/6305-485$3.00/0

School of Electrical and Electronic Systems Engineering, Queensland University of Technology, P.O. Box 2434, Brisbane 4001, Australia (pa.wilson@qut.edu.au).

processed except for the radiometric correction as outlined previously. The algorithm should not classify any pixel erroneously; i.e., it is acceptable for the software to fail to classify any pixel, but false positives are unacceptable. Any unclassified pixels can be reconsidered later in the application of the algorithm. A misclassified pixel may give rise to incorrect knowledge about an image which could lead to false conclusions. Thus, the Jupp *et al.* and the Moller-Jensen methods described in the introduction are unacceptable under the last criterion.

Another important criterion is that the algorithm must be general enough to identify water under varied conditions as discussed below.

Factors Which Cause Varied Reflectances from Water

The light reflected from water in an image can vary in intensity in several ways. The change of sun angle on swell, and whitecaps in choppy water result in specular reflection. The water may be turbid, or may contain vegetable matter such as algae or weed. The water may be shallow, giving rise to bottom reflectance (Campbell, 1987, Chapter 14; Lyzenga, 1981).

The data values (digital numbers) are also affected by the atmosphere, and the algorithm must cope with variations in atmospheric path radiance and scattered radiance from neighboring pixels. Variability caused by atmospheric radiance has been thoroughly discussed in the literature (Dave *et al.*, 1980; Switzer *et al.*, 1981; Kowalick *et al.*, 1983; Crippen, 1987; de Haan *et al.*, 1991).

The algorithm must accommodate noise caused by sensor gain variation (band striping) and quantization noise. Quantization noise is a result of converting a continuous or analog intensity value into a digital value, and the error is increased by expanding the range from the six bits captured by the camera to the eight bits of the data sold by ACRES. The error is increased again non-linearly by the radiometric correction algorithm employed by ACRES. Quantization noise can be greater than 50 percent for the low values of the infrared wavelengths for water, and it is particularly frustrating for developers of automatic algorithms.

The algorithm must also accommodate the difficulty of mixed pixels (mixels), both homogeneous (marshland) and edge mixels (shorelines). The author takes the view that the problem of mixels is being attacked in a different study, and classifying mixels, although an important consideration, is beyond the scope of this paper.

The Variance Filter Algorithm

A useful aspect of water (at least in oceans, wide rivers, and lakes) is that it is uniform and relatively flat, although some local variation in the brightness of water pixels is caused by all of the properties discussed in the previous section. It is the uniformity of water bodies that is exploited by the variance filter. The variance filter provides a measure of local homogeneity in an image. It can also be regarded as a non-linear non-directional edge detector. The filter emphasizes sudden changes in image brightness without any directional bias and is successful at identifying shorelines. High variance filter values at the shoreline are useful for limiting the region expansion algorithms (described in a later section). Mountain ridges cause very high variance filter values, and this fact is exploited in order to differentiate between mountain shadow and water.

The use of a texture channel in Landsat image interpretation is not new. The variance filter is one of a suite of filters based on straightforward statistical measures; mean, median, mode, variance, dispersion, etc. This set of filters, the so-called histogram filters, are used in many image processing applications. The variance filter has been mentioned briefly by Jain (1989, p. 344) as follows: "Variance can be used to measure local activity in the amplitudes." A similar type of filter has been used by Jupp *et al.* (1985) in the form of a texture channel previously proposed in the work of Haralick and Shanmugam (1974). The texture algorithm is the computed local root-mean-square (RMS) between the center pixel of a box and all the other eight pixels in the box. In the work of Jupp *et al.*, a 3- by 3-pixel box was used. Jupp *et al.* computed the texture only for Band 4 (MSS data) because the green band has the greatest water penetration. Haralick and Shanmugam (1974) specifically include variance in the appendix of their paper. The paper discusses the use of texture for automatic land-use classification, mentioning that water bodies display considerably more homogeneity than does grassland for example.

There have been other recent advances, in image classification, proposed by the statistical analysis community, where local homogeneity is used as one property in inferring class. These methods often rely on Bayesian computation and Markov random fields to represent local characteristics in an image (Besag, 1986; Smith and Roberts, 1993). The problems associated with Bayesian computation for Landsat image classification have been well documented in the literature (Skidmore and Turner, 1988). Skidmore and Turner, in their proposed non-parametric supervised classifier, have used a lookup table for modeling non-parametric search spaces. In a closely related approach, the use of rules in this project is a simple and effective way of modeling the non-linear and non-Gaussian search space of wind-ruffled, turbid, or shallow water. One purpose of this paper is to show that the variance filter is a straightforward way of using local homogeneity for image classification. The use of region-growing algorithms (described later) achieve similar ends to Besag's Markov random fields.

The variance filter algorithm consists of replacing a central pixel value with the variance of a specified set of pixel values surrounding it. The set need not be a square set of n by n pixels, where n is an odd number. Indeed, the set can be any specified group.

The variance of such a set is given as

$$S_x = \frac{1}{N-1} \sum_{i=1}^{N} (x_i - \overline{x})^2$$

where N is the number of pixels in the set, x_i is the value of pixel i, and x is the mean of pixel values in the set.

It makes little difference whether or not the unbiassed estimator definition of the variance is used. For a 5 by 5 set of pixels, N is 25 throughout, so the question of whether to use N or $N - 1$ is merely a question of scaling factor.

The computation is a two-pass one, because the mean of the set must be calculated on the first pass and the variance on the second.

Experiments were carried out on individual bands in order to determine an appropriate box size. A 3 by 3 variance filter on Band 7 was very successful in highlighting shorelines but did not smooth noise sufficiently. In particular, some uniform patches of cloud shadow were still indistinguishable from water. A 7 by 7 box was found to smooth edges too much so that some shorelines lost some definition. A 5 by 5 box, with the target pixel in the center, was found to give an optimum balance between definition and smoothing.

The Modulus Image

The variance filter can be used on any raster image, such as an individual band or any linear or non-linear combination of bands (one of the principle components for example). Very many variations on the algorithm are therefore possible.

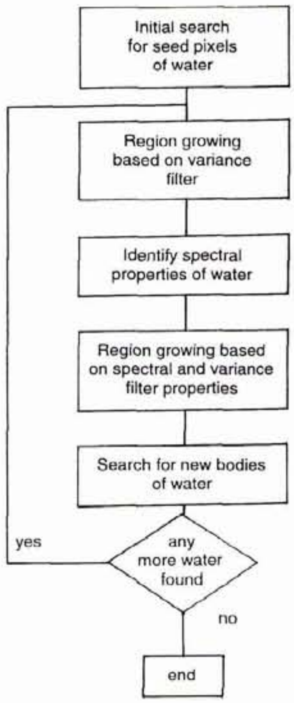

Figure 1. Flow diagram of the segmentation algorithm.

It is important to segment water with a range of spectral properties as previously discussed, and to distinguish boundaries between water and sand, soil, rock, and vegetation. It is essential also to distinguish water from cloud shadow over land and mountainsides. Notwithstanding the importance that Jupp *et al.* and others have placed on the infrared wavelengths, it was decided to use equal components of all the four wavelengths available with MSS data. One useful way of achieving this is to use the Euclidean distance of the pixel from the origin, which is referred to here as the *modulus* of the pixel vector. The modulus is also sometimes called the RMS value or the norm.

The pixel vector modulus is discussed by Pouch and Campagna (1990) in their paper on the calculation of the Hyperspherical Direction Cosine Transform, which is used for removal of brightness variation caused by the sun slope angle.

In this procedure, the variance filter of the modulus image is used as the basis for some of the rules. The modulus used in this work is calculated from

$$\text{modulus} = \sqrt{\frac{(w^2 + x^2 + y^2 + z^2)}{4}}$$

where w is the Band 4 intensity, x is the Band 5 intensity, y is the Band 6 intensity, and z is the Band 7 intensity.

The Rule-Based Classification Procedure

The procedure used here follows approximately the paradigm proposed by Ton *et al.* (1991) but with some important modifications. Their model proposes the use of a set of spectral rules to represent domain knowledge, and a set of spatial rules are used in region adjustment (*region growing* in this paper). It has already been mentioned that spectral rules alone are insufficient to distinguish water from other features such as cloud shadow and steep shaded hillsides. It is proposed here that domain knowledge is represented by both spectral and spatial information. Our region adjustment rules embody both spectral and spatial knowledge, because pixels already classified are used to define the likely spectral characteristics of water pixels yet to be found.

All the rules used here have been found by trial and error. The rules must be sufficiently general to find a reasonably sized body of water in the image if one exists and yet must not misclassify any pixel. The procedure must be reliable in spite of the considerable variation in brightness values for water from one image to another, caused by the factors discussed earlier. In particular, specular reflection and atmospheric path radiance values vary considerably between images.

Six images were used to develop the rules, these images representing a wide range of variation and a comprehensive sample of difficulties. Two of the images have been used for the comparison tests.

The program is iterative in nature; it continues to run until no more pixels of water can be classified. Within the iterative loop there are subroutines which also iterate until exhausted. A flow diagram of the procedure is shown in Figure 1. The rules which make up the subroutines are described in more detail in the following section. A key element of the procedure is the region-growing algorithm, of which two variations are used in the main iteration loop. Campbell (1987, Section 11.3) describes the region-growing algorithm as the *amoeba classifier*, and he commends it as an effective technique for large homogeneous regions.

Production Rules and Region Growing

The following section describes the rules which to date have given the most satisfactory results. The rules are under constant revision as subtle general properties of water emerge from the investigations.

The initial search for seed pixels was designed to establish the presence (or absence) of water in the image, and uses a very conservative set of rules. Values are digital values, also called the digital number.

Rules for finding seed pixels of water are as follows:

(1) the pixel value in Band 7 must be less than the image minimum value of Band 7 in the image plus 10: i.e., Band 7 $<$ (image minimum Band 7 + 10);
(2) the pixel value of the 5 by 5 variance filter of the modulus image must be 0 or 1;
(3) the Band 4 (green) value minus the Band 7 (infrared 2) must be greater than 15; and
(4) the pixel in question must be one of an unbroken row or unbroken column of at least 11 pixels, all of which must have a 5 by 5 variance filter (of modulus) value of less than ten.

These rules have successfully found a flat and uniform body of deep water about 750 metres across in every image tested so far. Rule 3 is designed to avoid large uniform patches of cloud shadow over forest or grassland. Rule 4 is very restricting as the algorithm will sometimes fail to find small or non-uniform water bodies. A recent refinement using a local response neural network (Wilson, 1996) has reduced the 750-metre restriction by a factor of three or so.

The main region growing algorithm follows this rule:

> The eight nearest neighbors to a pixel are examined and the number of these already classified as water is determined. To classify the new pixel as water, this number must be greater than half of the 5 by 5 variance filter of modulus value: i.e., if the variance filter of modulus value is 2 or 3, two near neighbors of water are required. If the variance filter of modulus value is 12 or 13, seven near neighbors of water are required.

Note that this algorithm has no rule regarding the spectral signature of the pixel; it is based on purely spatial data.

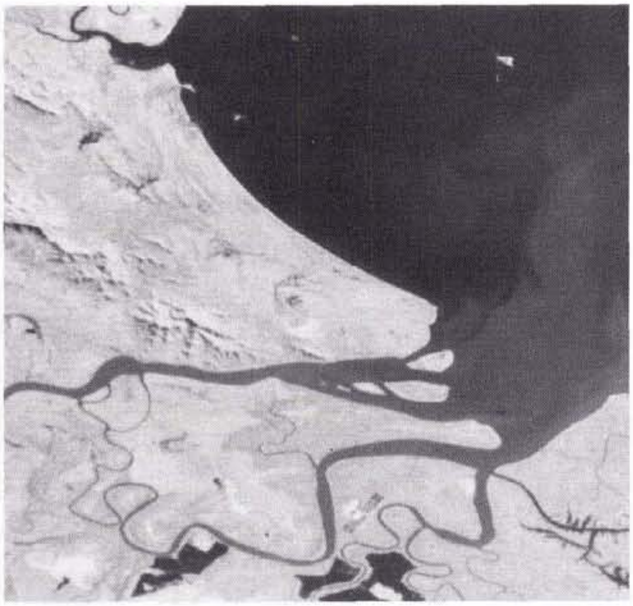

Figure 2. Landsat MSS image of the Fitzroy River Estuary in Band 6 (infrared 1).

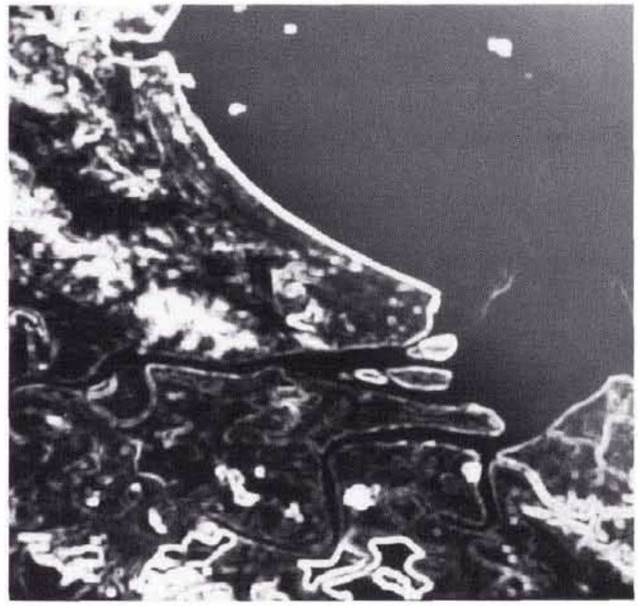

Figure 3. Variance filter of modulus image of the Fitzroy River Estuary.

However, the low variance filter of modulus value means that the pixel under examination must be spectrally similar to the surrounding pixels.

The second region growing algorithm has the following rules:

(1) The variance filter of modulus value must be less than ten,
(2) The pixel must be spectrally similar to pixels already identified as water, and
(3) At least one adjacent pixel must be already classified as water.

The spectral similarity algorithm originally employed a variation of the parallelepiped classifier (Campbell, 1986, Section 11.4) with the corners of the hyperbox truncated at 45 degrees to form octagonal rather than rectangular envelopes. Pixels of water already classified are used as the training set. The sides of the hyperbox are calculated to include all the pixels of water already classified. The octagonal box has now been replaced by a programmable local response neural network, which gives improved performance (Wilson, 1996).

The final algorithm in the iteration loop searches for other instances of water using information already obtained about the spectral characteristics of water in the image. This algorithm is a less constrained version of the original search for seed pixels of water. The rules for this are

(1) the value of the 5 by 5 variance filter of modulus image must be 0 or 1;
(2) the pixel must be spectrally similar to pixels already identified as water; and
(3) the pixel in question must be one of an unbroken row or unbroken column of at least 11 pixels, all of which must have a 5 by 5 variance filter of modulus value of less than ten.

This algorithm suffers from the same spatial limitations as the seed pixel finder, and the same improvement using the neural network algorithm is now being employed here.

The program described so far does not classify pixels close inshore because of the very high variance filter of modulus values generated by the shoreline. In future applications, pixels close to the shoreline may require a different procedure in order to classify the shoreline specifically. However, as a temporary measure in order to complete the classification to the shoreline, another region grower may be used after the main iteration loop. This region grower has the following rules:

(1) the new pixel must have one nearest neighbor as water, and
(2) the new pixel must be spectrally similar to water already found.

Comparison Tests

The rule-based algorithm was tested against a conventional thresholding algorithm on two images. The images are subscenes (each 512 pixels by 480 pixels) and were selected because they present particularly difficult challenges. They contain a great diversity of spectral and spatial characteristics and a wide variety of water bodies including ocean, estuary, river, reservoirs, crater lakes, mangrove swamp, coastal lagoons, and saltworks evaporator ponds.

The locations of the scenes are

(1) The Fitzroy River Estuary	latitude 23°30'S	longitude 150°50'E
(2) Lake Tinaroo	latitude 17°15'S	longitude 145°40'E

The Fitzroy River Estuary is a delta of coastal swamp and flats, inlets, oxbow lakes, islands, and mangrove. The image contains a saltworks; a range of mountainous country; some beach, mudflats, and duneland; and some offshore islands. The river water is very turbid and is therefore bright in the red band. The ocean is shallow in places, and bottom reflectance causes high values in the green band. The image is cloud free and was recorded on 5 August 1986. Band 6 (infrared 1) is shown in Figure 2 and the variance filter of modulus image is shown in Figure 3.

Fitzroy River image minima and maxima are

	Band 4	Band 5	Band 6	Band 7
maximum values	139	158	179	124
minimum values	20	8	8	4
mean of the modulus of the image		34		

The Landsat 5 image recorded on 8 August 1986, of Tinaroo Dam (a very large irrigation reservoir) on the Atherton

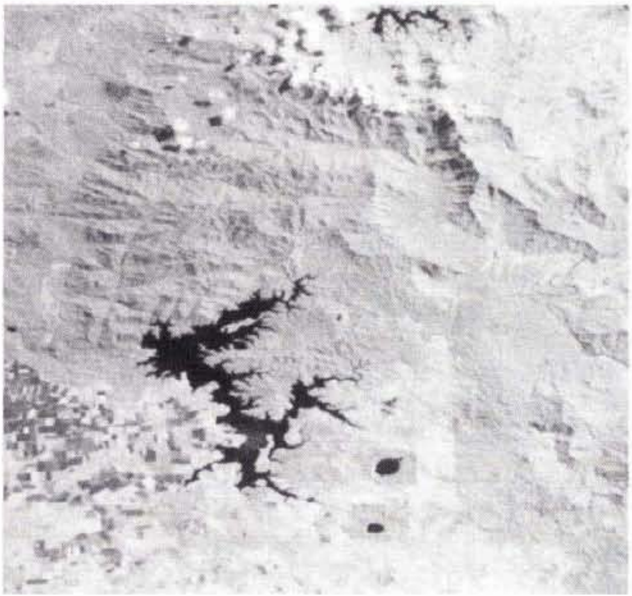

Figure 4. Landsat MSS image of the Tinaroo Dam Estuary in Band 6 (infrared 1).

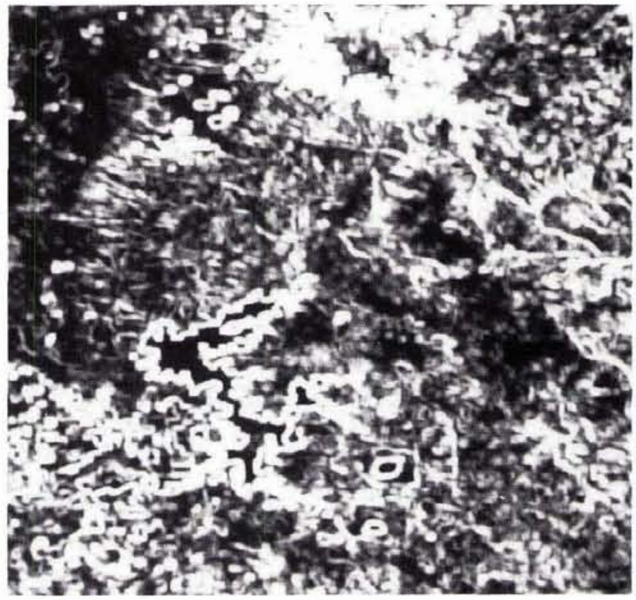

Figure 5. Variance filter of modulus image of the Tinaroo Dam area.

Tableland, is a bright picture. Band 7 values for the lake vary from a low of 1 to above 15. The image contains three crater lakes, Lake Barrine (1 square km, 204 pixels), Lake Eacham (0.4 square km, 84 pixels), and Lake Euramo (0.05 square km, 10 pixels). The western slopes of the Isley Range and the Little Mulgrave River have steep shaded gullies containing rainforest and there are patches of cloud in the north, the shadows of which are almost indistinguishable spectrally from the water in the lake. The Copperlode Falls Dam, in the north of the image, is a reservoir for the City of Cairns. This lake is long, winding, and very narrow. The Band 6 (infrared 1) image is shown in Figure 4 and the variance filter of modulus is shown in Figure 5.

The water in Lake Tinaroo is unusual in that, except for about six pixels, spectral values are higher than would be expected. The lake contains striations and filamentous structures radiating from a headland which are very bright in the green and red bands and are clearly visible in the infrared images. The structures are absent from an MSS image taken in August 1981 and from a Landsat Thematic Mapper image recorded in August 1992. They were also absent from the lake during a field trip in July 1993. It is assumed that the structures are mats of weed floating below the surface, and in a different project it has been possible to separate the spectral signature of the weed from that of the water (Wilson, 1994).

Image minima and maxima for Lake Tinaroo are

	Band 4	Band 5	Band 6	Band 7
maximum values	255	255	255	248
minimum values	21	10	11	1
mean of the modulus of the image	51			

The high maximum values are caused by the clouds in the north.

The conventional method of segmenting water is to classify as water any pixel having a value in Band 7 (infrared 2) of less than a specified threshold (Jupp *et al.*, 1985). In the Fitzroy River image, the thresholds used were 12, 17, and 27. In the Tinaroo Dam image, the thresholds were 12, 17, 21, and 27. For both images, the rules-based algorithm was exactly as described in the earlier sections. All programs were executed on a DEC 5000 workstation at the Queensland University of Technology (QUT). All programs have been written in ANSI C. The graphical side of the project uses a proprietary satellite imagery package written for personal computers at QUT.

The Fitzroy River image contains approximately 114,000 pixels of water, or about 46 percent of the image. Exactly how many pixels of water there are depends on the percentage cut off for mixels (here, approximately 80 percent of water in a mixel). The results for the Fitzroy River image are given in Table 1. The number of *water pixels* is the total number of pixels classified by the algorithm, including those wrongly classified. The *failed* column represents the number of true water pixels not found by the algorithm, and the *misclassified* column represents the number of pixels classified as water which were, in fact, mountain shadow or cloud shadow. Figure 6 shows the water classified by the rules-based algorithm, with classified water shown as uniform pale gray. Figure 7 shows the water classified by a threshold of less than 17. Misclassified pixels to the west of the mountain ridge can be seen in Figure 5. Correctly classified swamplands are also visible to the east of the mountain range.

The Tinaroo Dam image contains 6589 pixels of water, about 2.7 percent of the image. The results are given in Table 2. The columns are similar to those in Table 1 but an additional column, *failed weed*, records the number of pixels of weed infested water which the algorithm failed to classify. Figure 8 shows the Tinaroo Dam area classified by the rules-based algorithm. Figure 9 shows the same area classified by a threshold of less than 21. Very many misclassified pixels can be seen in the cloud and mountain shadows.

In terms of the number of misclassified pixels, the rules-based algorithm outperforms the threshold one, because it does not misclassify a single pixel in either image. The small number of pixels misclassified by the threshold algorithm in the Fitzroy River image is a function of the small number of deep shadow pixels. The larger proportion of misclassified pixels in the Lake Tinaroo image is caused by the longer mountain ranges and the patches of deep cloud shadow in the north. One of the misclassified pixels in the Lake Tina-

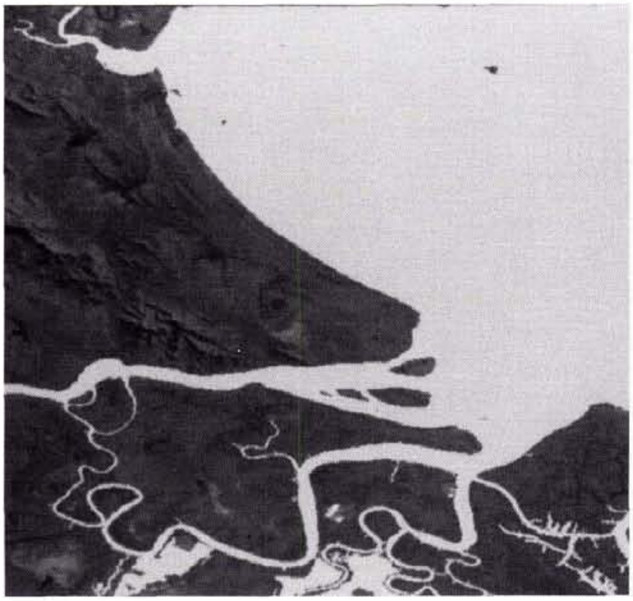

Figure 6. Water in the Fitzroy River Estuary area classified by the rules-based algorithm.

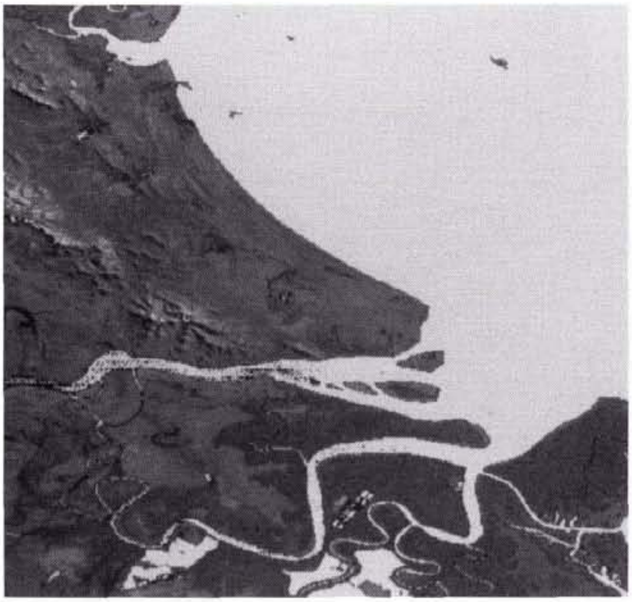

Figure 7. Water in the Fitzroy River Estuary classified by a threshold < 17.

roo image for a threshold of 27 is part of a ploughed field on dark red basaltic soil.

The rules-based algorithm also outperforms the thresholding algorithm because it operates reliably without human guidance and hence meets one of the stated objectives. For optimum operation, the threshold value must by selected by an analyst from histogram data.

The threshold algorithm outperforms the rules-based one by classifying very small bodies of water. The failure of the rules-based algorithm to find the oxbow lake and several areas of lagoons (known locally as billabongs) and swamps is disappointing. The disappointment is tempered by the knowledge that most of the swamp and lagoon pixels contain

Table 1. Comparison Results for the Fitzroy River Estuary Image

Algorithm	water pixels	failed to find	misclassified
water pixels in image	~114000		
threshold<12	97051	~17000	26
threshold<17	108576	~5500	75
threshold<27	113252	~1000	356
rules-based	113870	~150	0

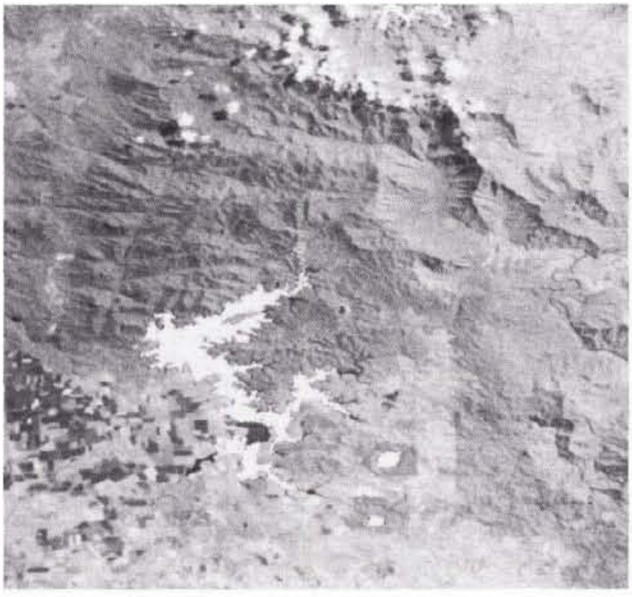

Figure 8. Water in the Tinaroo Dam area classified by the rules-based algorithm.

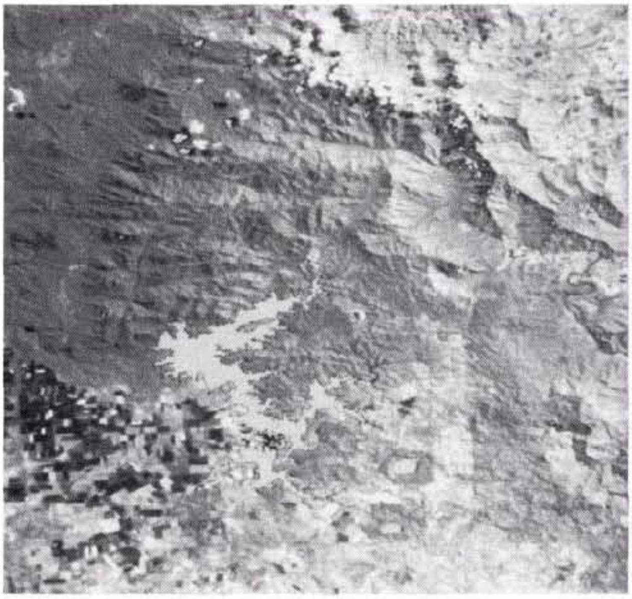

Figure 9. Water in the Tinaroo Dam area classified by a threshold < 21.

TABLE 2. COMPARISON RESULTS FOR THE TINAROO DAM IMAGE

Algorithm	water pixels	failed to find	failed (weedy)	misclassified
water pixels in image	6589			
threshold<12	1858	4731	n.a.	1
threshold<17	5804	785	n.a.	22
threshold<21	6753	190	190	354
threshold<27	8111	44	44	1566
rules-based	5734	855	339	0

much vegetation and are therefore mixels, and the algorithm does leave the pixels as unclassified. The algorithm failed to find Lake Euramo in the Tinaroo Dam image, but this was expected because only two pixels are 100 percent water. Again, the failure of the algorithm to find the weed infested part of the lake is disappointing in exactly the same way as for the swampy country. The pixels of weed are so patchy that they give rise to high variance filter of modulus values as well as high reflectance values in all bands.

Conclusion

The variance filter is not a new algorithm, and some previous work has been discussed. Its ability to distinguish between flat, uniform terrain and rugged or varying country when used with satellite imagery renders it useful for segmenting water, and, where artificial intelligence methods are used for the segmentation, it can be used as an extra channel of information. A region growing algorithm which includes a variance filter of modulus channel as well as spectral data has been shown to be relatively immune to variations in the spectral characteristics of the imagery. Successful tests on two Landsat MSS images have been described. When applied with conservative rules, the algorithm ignores regions of cloud shadow and shaded mountainside which are spectrally similar to water and which are classified as water by conventional threshold algorithms. The variance filter of modulus has been shown to be particularly successful at highlighting the shorelines of large water bodies. Further research is aimed at developing the usefulness of this filter in the classification of water by artificial intelligence methods.

Acknowledgments

I am indebted to the Australian Key Centre for Land Information Studies (AKCLIS), for supporting the project, and to the Queensland Government, Department of Lands for the use of the images for the tests. I should like to thank Associate Professor Greg Allen and Pauline Catt, Director of the Centre for Remote Sensing, both at James Cook University, Townsville, for their encouragement and advice.

References

Besag, J., 1986. On the Statistical Analysis of Dirty Pictures, *Journal of the Royal Statistical Society* B, 48(3):259–302.

Campbell, J.B., 1987. *Introduction to Remote Sensing*, The Guildford Press, New York.

Crippen, R.E., 1987. The Regression Intersection Method of Adjusting Image Data for Band Ratioing, *International Journal of Remote Sensing*, 8(1):137–155.

Dave, J.V., 1980. Effect of Atmospheric Conditions on Remote Sensing of a Surface Nonhomogeneity, *Photogrammetric Engineering & Remote Sensing*, 46(9):1173–1180.

de Haan, J.F., J.W. Hovenier, J.M.M. Kokke, and H.T.C. van Stokkom, 1991. Removal of Atmospheric Influences on Satellite-Borne Imagery: A Radiative Transfer Approach, *Remote Sensing of the Environment*, 37:1–21.

Haralick, R.M., and K.S. Shanmugam, 1974. Combined Spectral and Spatial Processing of ERTS Imagery Data, *Remote Sensing of Environment*, 3:3–13.

Jain, A.K., 1989. *Fundamentals of Digital Image Processing*, Prentice Hall, Englewood Cliffs, New Jersey.

Jupp, D.L.B., K.K. Mayo, D.A. Kuchler, S.J. Heggen, S.W. Kendall, B.M. Radke, and T. Ayling, 1985. *Landsat Based Interpretation of the Cairns Section of the Great Barrier Reef Marine Park*, The C.S.I.R.O. Natural Resources Series No. 4, C.S.I.R.O. Division of Water and Land Resources, Australia.

Kowalik, W.S., R.J.P. Lyon, and P. Switzer, 1983. The Effects of Additive Radiance Terms on Ratios of Landsat Data, *Photogrammetric Engineering & Remote Sensing*, 49(5):659–669.

Lyzenga, D.R., 1981. Remote Sensing of Bottom Reflectance and Water Attenuation Parameters in Shallow Water Using Aircraft and Landsat Data, *International Journal of Remote Sensing*, 2(1):71–82.

Moller-Jensen, L., 1990. Knowledge-Based Classification of an Urban Area Using Texture and Context Information in Landsat-TM Imagery, *Photogrammetric Engineering & Remote Sensing*, 56(6): 899–904.

Pouch, G.W., and D.J. Campagna, 1990. Hyperspherical Direction Cosine Transformation for Separation of Spectral and Illumination Information in Digital Scanner Data, *Photogrammetric Engineering & Remote Sensing*, 56(4):475–479.

Richards, J.A., 1986. *Remote Sensing Digital Image Analysis: An Introduction*, Springer-Verlag.

Skidmore, A.K., and B.J. Turner, 1988. Forest Mapping Accuracies Are Improved Using a Supervised Nonparametric Classifier With SPOT Data, *Photogrammetric Engineering & Remote Sensing*, 54(10):1415–1421.

Smith, A.F.M., and G.O. Roberts, 1993. Bayesian Computation Via the Gibbs Sampler and Related Markov Chain Monte Carlo Methods, *Journal of the Royal Statistical Society*, B, 55(1):3–23.

Switzer, P., W.S. Kowalik, and R.J.P. Lyon, 1981. Estimation of Atmospheric Path-Radiance by the Covariance Matrix Method, *Photogrammetric Engineering & Remote Sensing*, 47(10):1469–1476.

Ton, J., J. Sticklen, and A.K. Jain, 1991. Knowledge-Based Segmentation of Landsat Images, *IEEE Transactions on Geoscience and Remote Sensing*, 29(2):222–232.

Wilson, P.A., 1994. The Automatic Segmentation of Water, with Reference to the Tinaroo Falls Dam, North Queensland, *7th Australasian Remote Sensing Conference*, Melbourne, Australia, March.

———, 1996. Local Response Neural Networks for Experimental Classification of Remotely Sensed Images, *8th Australasian Remote Sensing Conference*, Canberra, Australia.

(Received 27 May 1993; revised and accepted 1 July 1994; revised 26 February 1996)