PROCEEDINGS

TWENTY-SIXTH ANNUAL ALLERTON CONFERENCE
ON COMMUNICATION, CONTROL, AND COMPUTING

VOLUME II

Mark W. Spong
Bruce Hajek
Conference Co-Chairmen

Conference held
September 28-30, 1988
Allerton House
Monticello, Illinois

Sponsored by
The Coordinated Science Laboratory
and
The Department of Electrical and Computer Engineering
of the
UNIVERSITY OF ILLINOIS
at
Urbana-Champaign

IMPROVING COVER TYPE IDENTIFICATION IN SPECKLED SAR IMAGES BY PREFILTERING AND SEQUENTIAL CLASSIFICATION *

QIAN LIN and JAN ALLEBACH
Laboratory for Applications of Remote Sensing and
School of Electrical Engineering
Purdue University, West Lafayette, IN 47907

ABSTRACT: Synthetic aperture radar utilizes coherent microwaves to produce images of the earth's surface. Due to the interference of coherent wavelets, the images appear speckled. This reduces the performance of per-pel classifiers. One way to increase the performance is to filter the image first, then classify the filtered image. For this purpose, we investigate several novel filters that have been reported in the literature. These are the geometric filter, adaptive LMMSE filter, and linear approximation filter. For comparison, we also consider conventional mean and median filters. In our experiments, we find that the mean filter with seven iterations gives the best result. The overall performance increased from 65.2% to 88.9%. We also assess the capability of these filters to preserve edges in the original image. We find that the geometric and median filters are the best in preserving edges, and that the linear approximation and adaptive LMMSE filters are the best in discriminating roads. Prefiltering the image effectively provides contextual information to the per-pel classifier. An alternate approach is to directly design a contextual classifier. We propose a new contextual classifier based on sequential decision theory. With this classifier, we find that the overall performance increases to 89.5%.

1. INTRODUCTION

Classification of cover type from remotely sensed images of the earth's surface provides an important means of gathering information about natural resources. Such images have traditionally been acquired with passive multispectral scanners operating in the visible and infrared; but more recently, active synthetic aperture radar (SAR) operating in the microwave region has also been used. In this case, images may be acquired at different incidence angles and with different transmitter/receiver polarizations to provide diversity corresponding to the different spectral bands of a multispectral scanner. Compared to multispectral scanners, SAR possesses the advantage of being able to operate day or night independent of weather or cloud conditions. However, SAR is a coherent system; and images acquired with it are degraded by speckle noise due to interference of waves reflected from multiple scatterers within each resolution cell. Speckle noise severely degrades the performance of per-pel classifiers which have been used successfully with multispectral images. In addition, speckle noise impedes identification of image detail, which is essential for correlating image information with that gathered from other sources including ground truth. Such information is, in fact, needed to identify the test samples on which the classifier design is based.

One approach to this problem is to filter the images to reduce speckle prior to classification. Mueller [1] has shown that conventional linear and median filters can improve classifier performance. However, these filters tend to degrade image sharpness to varying degrees. Recently, several novel filters have been developed with the objective of reducing speckle while maintaining image sharpness. These include the geometric filter [2] [3], an adaptive linear, minimum mean squared error (LMMSE) filter [4] [5], and a linear approximation approach in which the SAR image is expressed in terms of a basis set consisting of multispectral images [6]. The first objective of this paper is to evaluate the effect of these filters on classifier performance.

* Research supported by NASA under Contract No. NAGW-925.

Besides classification performance, another criteria for comparing filters is visual assessment. We define measures for edge spread and road contrast, and use them to evaluate the visual quality of SAR images processed with the various filters.

The filter-then-classify approach may be regarded as a rather indirect means of providing a per-pel classifier with contextual information. A second approach is to develop explicitly contextual classifiers that operate directly on the speckled image. Classifiers based on combining similar neighbors [7] and on compound decision theory [8] are examples. Image segmentation techniques [9] are also applicable. The second objective of this paper is to develop and evaluate a contextual classifier based on sequential decision theory [10].

In Sec. 2, we briefly describe the three new filters to be evaluated in the filter-then-classify approach. In Sec. 3, we also briefly describe the conventional per-pel Bayes classifier. We then present our new sequential classifier. In Sec. 4, we define the measures for assessment of visual quality. Finally, Sec. 5 contains the results from an experimental comparison of the various filter and classifier combinations.

2. FILTERS FOR SPECKLE REDUCTION

The geometric filter [2] [3] is essentially a morphological operator [11]. It may be explained as follows: Suppose a discrete two-dimensional gray level image is converted to a discrete three-dimensional volume with the additional z axis representing the gray value. Points below the gray level surface have value 1; and those above it have value 0. The volume is partitioned into parallel slices perpendicular to the x-y plane to obtain a set of binary images. Each binary image is filtered by running four 3×3 templates over it. Wherever a template matches, the center pixel is complemented; otherwise it remains the same. The filtered image is then complemented, and processed again in the same manner with four additional templates. Complementing the output of this step yields the final filtered output for the slice. The above procedure is applied sequentially to sets of slices taken at four different angles with respect to the x-axis. This constitutes one iteration. The final result is obtained by converting the binary images back to the original image form. It is important to note that this can always be done since the overall filtering operation obeys the stacking property [12].

The adaptive filter [4] [5] is based on a model in which the noise is multiplicative, and spatially uncorrelated with a unit-mean Gaussian distribution. Thus the (i,j)-th pixel $z_{i,j}$ in the SAR image is related to the terrain reflectance $x_{i,j}$ and speckle noise $v_{i,j}$ according to $z_{i,j} = x_{i,j}\, v_{i,j}$. Suppressing the subscripts (i,j), the LMMSE estimate of the terrain reflectance may be expressed as

$$\hat{x} = \bar{z} + k(z-\bar{z}),$$

where

$$k = \frac{\sigma_x^2}{\bar{v}^2\sigma_v^2 + \sigma_x^2},$$

and

$$\sigma_x^2 = \frac{\sigma_z^2 + \bar{z}^2}{\sigma_v^2 + \bar{v}^2} - \bar{x}^2.$$

The parameters σ_z^2 and $\bar{z}$ are calculated as sample averages within a neighborhood of the (i,j)-th pixel. Since the SAR image is an amplitude image preprocessed with a four-look average on intensity basis, $\sigma_v = 0.2536$ [5]. In addition, $\bar{v} = 1$; and this implies that $\bar{x} = \bar{z}$. In a flat region, $\sigma_x^2 \simeq 0$, and $\hat{x} \simeq \bar{x}$, which is the linear average of pixels in the neighborhood. In a high contrast region or edge area, σ_x^2 is large. In this case, $\hat{x} \simeq z$, and the value of the pixel is not changed.

Use of the approximation method [6] depends on the availability of multispectral images covering the same terrain as the SAR images, and which are registered with

them. The multispectral images are highly correlated with the SAR images, but are not speckled. Thus, they comprise a good basis set for a smooth, linear approximation to the SAR images.

To determine the smoothed image value $\hat{x}$ at a given pixel, we express the SAR image data $\vec{z}$ within a w-point window in terms of the p reference images $\vec{y}_i$, $i=1, \ldots, p$, as

$$\begin{bmatrix} z_1 \\ z_2 \\ \cdot \\ \cdot \\ z_w \end{bmatrix} = \begin{bmatrix} 1 & y_{11} & \cdot & \cdot & y_{1p} \\ 1 & y_{21} & \cdot & \cdot & y_{2p} \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ 1 & y_{w1} & \cdot & \cdot & y_{wp} \end{bmatrix} \begin{bmatrix} b_0 \\ b_1 \\ \cdot \\ \cdot \\ b_p \end{bmatrix} + \begin{bmatrix} e_0 \\ e_1 \\ \cdot \\ \cdot \\ e_p \end{bmatrix}$$

or in vector-matrix notation,

$$\vec{z} = Y\vec{b} + \vec{e},$$

where $\vec{e}$ is the approximation error. The coefficients $\vec{b}$ are chosen to give a least squares fit:

$$\vec{b} = (Y^TY)^{-1}(Y^T\vec{z}).$$

The smoothed SAR image value is then given by

$$\hat{x} = [1 \;\; y_{k1}, \ldots, y_{kp}]\,\vec{b},$$

where k is the index of the center pixel in the window.

3. CLASSIFICATION

With per-pel classifiers, each image pixel is classified independently of the others in the image. Let $\vec{x}$ be a 3-tuple of the SAR image values for the three incidence angles at a particular pixel; and suppose that it belongs to one of M classes. Assuming that the distribution of $\vec{x}$ conditioned on it belonging to the i-th class is multivariate Gaussian with mean $\vec{m}_i$ and covariance Σ_i, the discriminant function for a Bayes classifier is given by

$$g_i(\vec{x}) = \ln(p_i) - \frac{1}{2}\ln\left(|\Sigma_i|\right) - \frac{1}{2}(\vec{x}-\vec{m}_i)^T\Sigma_i^{-1}(\vec{x}-\vec{m}_i).$$

Here p_i is the prior probability of the i-th class. The pixel is placed in that class i^* for which

$$g_{i^*}(\vec{x}) \geq g_i(\vec{x}), \quad i = 1,\ldots,M.$$

Since the signal to noise ratio of SAR images is very low, the per-pel classification rate is low compared to that of images obtained from noncoherent imaging systems. Filtering the speckled images reduces the variance of the data, making the gray value vector for each pixel more representative of the class to which it belongs. This results in an improvement in classifier performance.

Our contextual classifier examines the values of pixels in a neighborhood around the pixel in question to aid in its classification. The sequence in which these pixels are evaluated is indicated in Fig. 1. The pixel to be classified is evaluated first. At the n-th step in the algorithm, we attempt to classify this pixel based on its value $\vec{x}_1$ and that of its n-1 neighbors $\vec{x}_2, \ldots, \vec{x}_n$. If no decision can be made, we add the pixel $\vec{x}_{n+1}$ to the set and repeat the process. To reduce the chance of running across a boundary between two cover types, we limit the neighborhood to a 5×5 window.

We first describe the classification algorithm when there are only two classes H_0 and H_1. We then discuss the extension to the case of $M>2$ classes. Assuming neighboring pixels are independent and belong to the same class, we have, at the n-th stage of the algorithm, n i.i.d. vector observations $\vec{x}_1, \ldots, \vec{x}_n$. To test the hypothesis H_1 versus H_0, we compute the log-likelihood ratio

15	16	10	17	18
14	9	5	8	19
13	4	1	2	11
25	8	3	7	20
24	23	12	22	21

Fig. 1. Data sequence used in the sequential classifier.

$$s_n = \ln\left(\frac{f(\vec{x}_1, \ldots, \vec{x}_n \mid H_1)}{f(\vec{x}_1, \ldots, \vec{x}_n \mid H_0)}\right) = \sum_{k=1}^{n} \ln\left(\frac{f(\vec{x}_k \mid H_1)}{f(\vec{x}_k \mid H_0)}\right),$$

and compare it with two fixed thresholds $a < 0 < b$. If $s_n \leq a$, we choose H_0; and if $s_n \geq b$, we choose H_1. In this case, the algorithm terminates. Otherwise, $a < s_n < b$; so we get another data sample $\vec{x}_{n+1}$, and repeat the process.

To bound the error probabilities by α_0 and α_1

$$P(\text{choosing } H_1 \mid H_0) \leq \alpha_0,$$

$$P(\text{choosing } H_0 \mid H_1) \leq \alpha_1,$$

we let

$$a = \ln\left(\frac{\alpha_1}{1-\alpha_0}\right),$$

$$b = \ln\left(\frac{1-\alpha_1}{\alpha_0}\right).$$

For M classes, we calculate the log-likelihood functions of the data sequence for each class i

$$l_{i,n} = \ln(f(\vec{x}_1, \ldots, \vec{x}_n \mid H_i)),$$

and then calculate the log-likelihood ratio $s_n = l_{i1,n} - l_{i2,n}$ for the largest two of the log-likelihood functions. If we can make a decision based on s_n, we terminate. Otherwise, we add a sample, recalculate the likelihood functions for each class, and again attempt to make a decision as above. If a decision cannot be made with all 25 samples in the window, we choose the class with the largest likelihood function.

4. EDGE RENDITION AND ROAD DISCRIMINATION

To quantitatively describe the sharpness of edges, we define the edge spread S to be the distance between the two pixels which are the 90% transition points shown in Fig. 2. In digital imagery, it is easier to detect edges horizontally and vertically than at an arbitrary angle. Suppose we scan horizontally, and detect an edge. We calculate the spread S of the edge as if the edge were vertical. If the edge is at an angle α with the vertical direction, as shown in Fig. 3, we calculate a least squares estimate of the slope and scale the spread of the edge accordingly.

More precisely, let the x coordinates range over the integers from 0 to N-1, and the y coordinates be the distances to the edge, as shown in Fig. 3. The edge can be represented by the linear equation

$$y = mx + b,$$

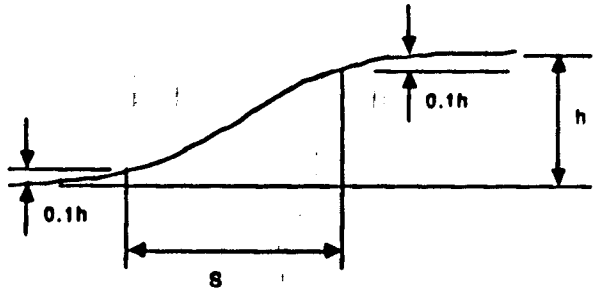

Fig. 2. Definition of edge spread.

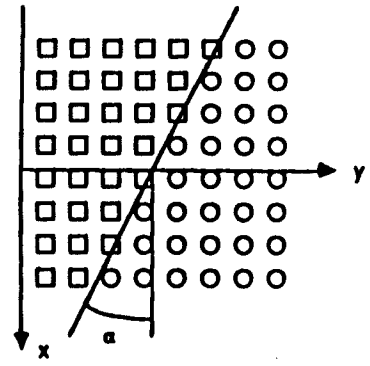

Fig. 3. Edge slope estimation.

where m is the slope we are trying to estimate. Let

$$A = \begin{bmatrix} x_1 & 1 \\ x_2 & 1 \\ \cdot & \cdot \\ \cdot & \cdot \\ \cdot & \cdot \\ x_{N-1} & 1 \end{bmatrix}, \quad X = \begin{bmatrix} m \\ b \end{bmatrix}, \quad B = \begin{bmatrix} y_1 \\ y_2 \\ \cdot \\ \cdot \\ \cdot \\ y_{N-1} \end{bmatrix};$$

then $AX \simeq B$. The least squares fit of the parameters m and b is given by

$$X = (A^TA)^{-1}A^TB.$$

Since the x_i are the integers from 0 to N-1, $(A^TA)^{-1}A^T$ can be precalculated. The corrected edge spread is then given by

$$S_c = S\cos(\alpha) = \frac{S}{\sqrt{1+m^2}}.$$

In SAR images, pixels representing roads have a smaller gray value than surrounding pixels. Suppose we use an array to represent a line segment across a road. Let a and c be the gray values of the two end points of the line segment, while b is

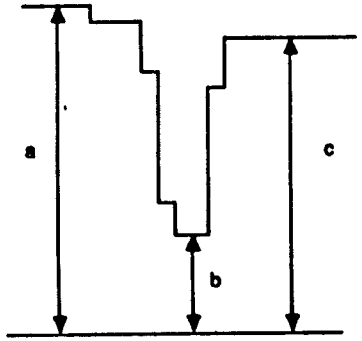

Fig. 4. Definition of road contrast.

the minimum gray value of the line segment, as shown in Fig. 4. We define the contrast of the road to be

$$C = \frac{\frac{1}{2}(a+c) - b}{\frac{1}{2}(a+c)}.$$

5. EXPERIMENTAL RESULTS

The images used in the experiment were obtained by the Shuttle Imaging Radar-B (SIR-B) during Space Shuttle Flight 41-G in October, 1984. Three sets of data were collected over a forested area in northern Florida on Oct. 9, 10, and 11, at incidence angles 58°, 45°, and 28°, respectively. The images were digitally processed by the Jet Propulsion Laboratory (JPL) in Pasadena, CA [13]. A set of Thematic Mapper (TM) images over the same area was obtained on Oct. 12, 1984; and they were registered with the SIR-B images by JPL. The images we used here have 512×512 pixels. Each pixel covers a 28.5×28.5 m^2 area on the ground. The SAR images with incidence angles 28°, 45°, and 58° are referred to as bands 1, 2, and 3, respectively. The TM images have seven bands, which are used as the reference images in the linear approximation filter.

By calculating the statistics of several sample fields, it was found that there is not a great difference in correlation between different cover types, which means that the "texture" we see in the image does not contain much information about cover type. In fact, if we subtract the local mean and scale by the local standard deviation making the transformation

$$\tilde{x}_{ij} = (x_{ij} - \bar{x}) \,/\, \sigma_x + 128$$

for every pixel in the image, the interior region of different ground cover types cannot be distinguished in the resulting image, at least not visually, although the roads and sharp edges are still present. On the other hand, if we only subtract the local mean, or only scale by the local standard deviation, then there still exist visible differences between ground cover types. This suggests that most of the information about ground cover type is contained in the local mean and variance. The speckle index $\sigma_x/\bar{x} \simeq 0.26$ in most of the fields. This is consistent with the expected value of 0.2536 discussed in Sec. 2.

The images were filtered with the geometric filter, the adaptive LMMSE filter, the linear approximation filter, and conventional 3×3 mean and 3×3 median filters.

The geometric filter was run for four iterations, as is suggested in [3]. The mean and median filters were run for seven iterations to reduce the speckle index to about the same level as that obtained with the geometric filter. The adaptive LMMSE filter and the linear approximation filter were run for only one iteration, since they are based on models of the original speckled image. The window size of the adaptive LMMSE filter was chosen to be 5×5. With the linear approximation filter, singularity is possible when inverting the reference image matrix since the TM images are so smooth. This problem was circumvented when necessary by perturbing the matrix Y.

To gain a quantitative idea of how the filters perform, we concentrate on a 32×32 pixel water region of band 1. Since this region is homogeneous, the variations are due to speckles only. The speckle indices for the unfiltered and filtered images are shown in Table 1.

Table 1. Effect of Filters on Speckle Index

Filter Type	No. Iterations	Speckle Index
original image	0	0.2627
3×3 mean	7	0.1025
3×3 median	7	0.1000
geometric	4	0.0832
adaptive LMMSE	1	0.1376
linear approx.	1	0.1717

We see that the geometric filter with four iterations is most effective in reducing speckle noise.

Eighteen training fields and 42 test fields were selected based on local lumber companies' inventories and maps. These fields were grouped into six different classes: water, clear cut with trees aged 1 to 6 years, pine aged 6 to 36 years, pine older than 36 years, mixed swamp, and cypress swamp. The prior probabilities of each class were obtained by analyzing the available ground truth, and were set as in Table 2.

Table 2. Prior Probabilities for Six Ground Cover Types

Cover Type	Prior Probability
Water	0.03
Clear cut and pine (1 - 3 yrs.)	0.16
Pine (6 - 36 yrs.)	0.35
Pine (> 36 yrs.)	0.20
Mixed swamp	0.20
Cypress swamp	0.06

Classification was performed on the original image, on images filtered with the geometric filter, the adaptive LMMSE filter, and the linear approximation filter, as well as the 3×3 mean and median filters. The classifier's performance was estimated using the holdout method [14]. The results are shown in Fig. 5. We see that using filtered images greatly improves the classification rate. The mean filter with seven iterations gave the best results.

For the sequential classifier, we chose the error probabilities to be $\alpha_0 = \alpha_1 = 0.01$. We did the sequential probability ratio test on the original and filtered images. The classification rate is also shown in Fig. 5. We can see that the classification rate increased in all cases by varying degrees. The largest increase occurred when sequential classification was done on the raw data. Note that the classification rate did not reach the theoretical error probability. This is mainly because the theoretical value is for a binary hypothesis test. Other reasons include truncation of the data sequence and statistical inhomogeneity of the images.

To gain a qualitative idea of the characteristics of the different filters, we show in Fig. 6, portions of the original image, the images filtered by the geometric filter, the mean filter, and the median filter. These portions each contain 256×256 pixels. We see that the geometric filter preserved edges very well, while effectively reducing fluctuations due to the speckle noise. Compared with the geometric filter, the mean

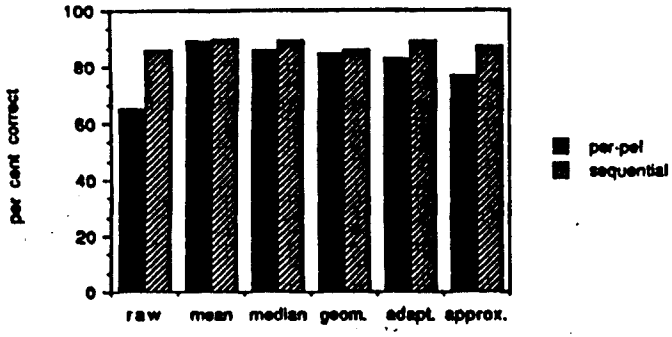

Fig. 5. Performance of per-pel and sequential classifiers.

Fig. 6. Portion of original image and images filtered with the mean, median and geometric filters (left to right and top to bottom).

filter smeared edges. Edge rendition with the median filter is similar to that of the geometric filter, although the median filter obliterated some distinct features present in the original image, which were preserved by the geometric filter. Figure 7 shows portions of the original image, the image filtered by the linear approximation filter, the TM image (band 5) and the image filtered by the adaptive LMMSE filter. The linear approximation filter has the special characteristic of enhancing the edges which are blurred and hardly distinguishable in SIR-B images, but which are very clear in the TM images, since it uses the TM images to approximate the SIR-B images. As expected, the adaptive filter passed edges in the original image with little change.

We calculated the edge spread and road contrast for the various filtered images. To detect edges, we used the boundary information of a classified image, because the original SIR-B image are severely speckled. For the same reason, we used a TM image to detect roads. The result is shown in Table 3.

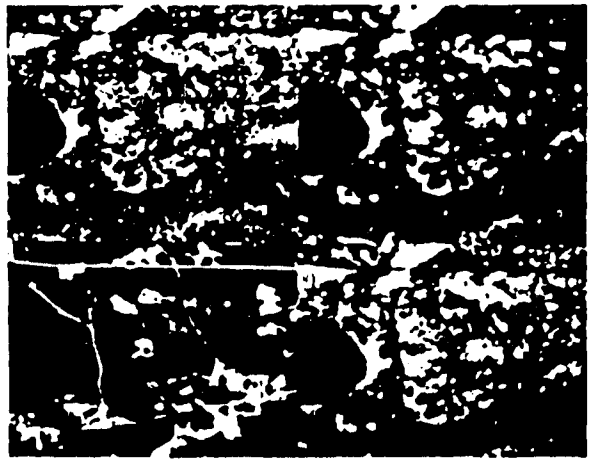

Fig. 7. Portion of original image, image filtered with the adaptive LMMSE filter, TM image (band 5), and image filtered with the linear approximation filter (left to right and top to bottom).

Table 3. Edge Spread and Road Contrast Measures

Filter Type	No. Iterations	Edge Spread S	Road Contrast C
original image	0	2.794	0.289
3×3 mean	7	5.317	0.077
3×3 median	7	3.767	0.105
geometric	4	3.500	0.103
adaptive LMMSE	1	3.473	0.157
linear approx.	1	3.114	0.221

The linear approximation and adaptive LMMSE filters result in the smallest edge spread and the largest road contrast. With approximately the same speckle index, the geometric filter preserves edges and roads better than the mean and median filters.

6. CONCLUSION

Prefiltering the images improved the performance of the per-pel Bayes classifier by up to 36.3%. The mean filter gave the best classification rate, but was inferior to the other filters in edge rendition and road discrimination. The linear approximation and adaptive LMMSE filters yielded the best edge spread and road contrast. The contextual sequential classifier increased the classification rate by up to 37.3% compared to per-pel classification of the original image. With this classifier, performance did not vary much between the original image and images processed by the various filters.

ACKNOWLEDGMENT

The authors wish to thank Prof. R. M. Hoffer, Dr. D. F. Lozano, and Mr. K. Lee for their assistance in providing the SAR and TM images and the ground truth information, and for many helpful discussions. One of the authors (J. A.) also wishes to thank Prof. D. C. Munson, Jr. for pointing out the geometric filter.

REFERENCES

[1] P. W. Mueller, *Spatial Filtering of Shuttle Imaging Radar-B Data*, M.S. Thesis, Purdue Univ., 1987.

[2] T. R. Crimmins, "Geometric Filter for Speckle Reduction," *Appl. Opt.*, Vol. 24, pp. 1438-1443, May 1985.

[3] T. R. Crimmins, "Geometric Filter for Reducing Speckle," *Opt. Eng.*, Vol. 25, pp. 651-654, May 1986.

[4] J. S. Lee, "Speckle Suppression and Analysis for Synthetic Aperture Radar," *Opt. Eng.*, Vol. 25, pp. 636-643, May 1986.

[5] J. S. Lee, "Statistical Modelling and Suppression of Speckle in Synthetic Radar Images," *Proc. of IGARSS '87 Symposium*, Ann Arbor, MI, 18-21 May 1987, pp. 1331-1339.

[6] V. T. Tom, M. J. Carlotto, "Adaptive Least-Squares Technique for Multi-band Image Enhancement," *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 704-707, March 1985.

[7] R. L. Kettig and D. A. Landgrebe, "Classification of Multispectral Image Data by Extraction and Classification of Homogeneous Objects," *IEEE Trans. Geosci Electron.*, Vol. GE-14, pp. 19-26, Jan. 1976.

[8] P. H. Swain, S. B. Vardeman, and J. C. Tilton, "Contextual Classification of multispectral image data," *Pattern Recognition*, Vol. 13, No. 6, pp. 429-441, 1981.

[9] A. Rosenfeld and A. C. Kak, *Digital Picture Processing*, Vol. 2, 2nd Ed., New York: Academic Press, 1982.

[10] A. Wald, *Sequential Analysis*, New York: Dover, 1947.

[11] J. Serra, *Image Analysis and Mathematical Morphology*, New York: Academic Press, 1982.

[12] P. D. Wendt, E. J. Coyle, and N. C. Gallagher, "Stack Filters", *IEEE Trans. Acoust. Speech Signal Process.*, Vol. ASSP-34, pp. 898, August 1986.

[13] J. C. Curlander, "Performance of the SIR-B Digital Imaging Processing Subsystem," *IEEE Trans. Geosci. Remote Sensing*, Vol. GE-24, pp. 649-652, July 1986.

[14] P. A. Devijver and J. Kittler, *Pattern Recognition: A Statistical Approach*, Englewood Cliffs, NJ: Prentice Hall, 1982.