METHODS FOR EVALUATING SPECKLE-SUPPRESSING FILTERS BASED ON EDGE DETECTION PERFORMANCE

by
M. Adair
Intera Technologies LId.
B. Guindon
Canada Centre for Remote Sensing

RÉSUMÉ

Des techniques fondées sur la détection des frontières comme mesure de performance pour l'évaluation de filtres destinés à réduire le chatoiement sont proposees dans le présent article. Étant donné que le chatoiement a pour effet non seulement de masquer l'information sur les frontières reelles, mais également de provoquer l'apparition de frontières artificielles, les filtres doivent être évalués en fonction de leur capacité à identifier les frontières réelles et à supprimer les frontières artificielles.

Les auteurs ont utilisé les frontières de champs agricoles observées sur des images RAS pour estimer le taux de frontières réelles identifiées avec deux modèles de filtres adaptatifs, soit le filtre Frost et le filtre Lee Sigma. L 'étude des frontières artificielles, d'un autre côté, s'avère difficile avec les images réelles en raison du fait que nous ne connaissons pas de façon précise les propriétés de rétrodiffusion des cibles en présence. Nous avons donc utilisé des images simulées pour analyser ce problème et, finalement, pour élaborer des modèles permettant d'évaluer la performance des deux filtres en question.

Les filtres Frost et Lee sont des fillres « adaptatifs » en ce sens qu'ils fixent le niveau de filtrage spatial en un point donné en fonction d'une mesure du contenu en frontières à l'intérieur de la fenêtre du filtre. Tout bruit résiduel est mulliplicatif: En conséquence, la densité et l'importance des frontières artificielles sont fonctions de l'information réelle, c'est-à-dire que des taux élevés de frontieres artificielles seront observés dans les regions où la rétrodiffusion radar est importante. En raison de cette relation, des precautions doivent être prises dans l'interprétation du contenu en information spatiale des images RAS, telle la texture d'image. Les auteurs proposent que les seuils délimitant l'importance des frontières soient déterminés à partir des modeles de performance et qu'ils soient d'abord appliqués à des images réelles pour éliminer les frontières qui sont potentiellement associées au bruit et non pas à la variabilité réelle de la rétrodiffusion des surfaces.

SUMMARY

Certain techniques are proposed to evaluate speckle-reducing fitters based on edge detection as a performance measure. Since speckle not only masks real edge information but also can trigger spurious edges (artefacts), it is argued that fitters must be gauged both in terms of real edge recovery and artefact suppression.

Agricultural field boundaries observed in airborne SAR images have been used to estimate real edge recovery rates for two adaptive fillers, the Frost and the Lee Sigma fitters. Studies of artefacts, on the other hand, are rendered difficult with real image data since a precise knowledge of the backscatter properties of surface targets is not available. Therefore, we have used simulated imagery to analyze this problem and, ultimately, to develop overall edge detection performance models for the two fitters in question.
Both the Frost and Lee fitters are "adaptive" in the sense that they determine the level of spatial averaging at a given image location according to the perceived edge content within the fitter window. Any residual noise remains multiplicative, and, hence, the density and magnitude of artefacts retain a signal dependence; that is, high artefact levels will be observed in regions of high radar backscatter. Because of this signal dependence, caution must be taken in interpreting the spatial information content of SAR images such as image texture. It is proposed that edge magnitude thresholds be determined from fitter performance models and first applied to real imagery to eliminate that portion of the overall edge content potentially associated with noise and not with real variability in surface backscatter.

INTRODUCTION

For many applications, the presence of speckle in synthetic aperture radar (SAR) images is detrimental to data interpretation. Numerous filters have been proposed to suppress speckle (e.g., Frost *et al.*, 1982; Lee, 1983, 1987; Crimmins, 1985). The utility of these filters has been gauged on the basis of both applications-specific criteria, such as observed improvement in agricultural field boundary definition (Goodenough *et al.*, 1984) and crop identification (Durand *et al.*, 1987), or on improvements in image primitive extraction, such as edge and linear feature detection (e.g., Modestino and Fries, 1977; Machuca and Gilbert, 1981; Wood, 1985; Touzi *et al.*, 1987).

The purpose of this paper is to extend the analysis of filter performance with regard to edge detection. The following three objectives have been established:

- to establish methods quantifying the improvement in edge extraction provided by a speckle-reducing filter . It is argued that filter performance must not only be gauged in terms of the resulting enhancement of real edge detectability, but also by the level of false edges (i.e., artefacts) triggered by residual noise;
- to demonstrate the utility of simulated imagery in assessing filter performance;
- to develop filter performance models and to demonstrate their use for confidence testing of the edge information extracted from real SAR images.

A comparison of the adaptive Frost filter (Frost *et al.* , 1982) and the multiplicative Sigma filter (Lee, 1983) has been undertaken to demonstrate the methodology set out in this paper. We have used the edge detection and thinning algorithms proposed by Nevatia and Babu (1980) as our edge primitive extraction standard.

TEST FILTERS

Both the Frost and Sigma filters are exhaustively described in the appropriate references; hence, only brief descriptions of their salient features are given.

The filters are similar in that they involve replacement of the grey level of each pixel, at location (i,j), and by the spatial average of some or all of the grey levels of pixels within a window of fixed size centred on (i,j). The filters differ in the pixel selection and weighting strategies used in their spatial averaging process. In addition, the Sigma filter can be applied recursively.

In the case of the Frost filter, spatial averaging is accomplished through the application of a radially symmetric, exponential weighting function. The "decay factor" of the exponential selected for each pixel location depends on the observed grey level mean square to variance ratio of pixels within the filter window. The closer this ratio is to the number of looks of the scene (assumed to be known *a priori),* the larger the decay factor and, hence, the greater the level of spatial averaging.

On the other hand, the Sigma filter uses radiometric criteria to select pixels for spatial averaging. The grey level of the central pixel in the window is compared with the grey levels of all other window pixels. Only those window pixels exhibiting a grey level within an acceptance range are selected for averaging. If the grey level of the central pixel is Z_{ij} and the number of looks of the scene is NL, the grey level acceptance range is selected to be

$$(1 - m\ N_L)z_{ij} \text{ to } (1 + m\ N_L)z_{ij},$$

where m is a constant. Adaptation of the filter to multiplicative noise is accounted for since the acceptance grey level range is proportional to the grey level of the central pixel.

Sharp spot noise (i.e., isolated bright or dark pixels) will not be eliminated with this algorithm alone. Hence, Lee has proposed an additional threshold, K, be included. If the number of pixels within the acceptance range is less than K, the radiometric criterion is replaced by a simple averaging of the central pixel's four nearest neighbours. K is selected to be much less than the total number of window pixels in order to prevent the obliteration of fine scale features. Lee has also proposed that a series of Sigma filters can be applied sequentially to the same image, each successive filter being of smaller size and lower threshold.

Given that the objective of this paper is the development of performance evaluation methodologies and that these filters are used primarily for illustrative purposes, we have not attempted to optimize the selection of filter parameters. For the Frost filter, we have selected a window size of 15 by 15 pixels, while for the Sigma case we have considered the effects of a 3 pass, m = 2 Sigma filter process whose filter window sizes and thresholds are 15 x 15, 7 x 7, 3 x 3, and K = 7, 3, 1, respectively. These are relatively large window sizes, but previous experience with automated matching of SAR edge image pairs indicates that extensive spatial averaging will be required for operational processing of low look SAR images (Guindon, 1985, 1986).

IMAGE DATA

Both real and simulated test scenes have been employed in our analyses. The real image data consist of airborne C-band amplitude scenes acquired over Melfort, Saskatchewan, on July 31 and August 13, 1983 (Brown *et a/.*, 1984). The images, with a resolution of 11.2 metres and a pixel spacing of 5.6 metres, are particularly well-suited to the present study since they contain straight edge and line features corresponding to field boundaries and the road network. Ancillary information in the form of colour infrared photography also was acquired and has been used to identify candidate edges for detailed analysis.

Simulated images have been used also to develop performance models and to estimate artefact generation. These images consist of an underlying control signal in the form of either checkerboard or line patterns. Multiplicative noise was then introduced to simulate the effects of multilooking in a power image using the analytic description of speckle given by Lee (1987). The mean grey level of the homogeneous regions is restricted to the 60-160 range (in an 8-bit image) to prevent saturation effects. The oversampling characteristics of the real imagery has also been emulated.

EDGE AND LINE FEATURE EXTRACTION ALGORITHM

We have opted to use the edge detection component of the linear feature extraction algorithm proposed by Nevatia and Babu (1980) as our reference edge operator. Our reasons for this selection are twofold: first, the algorithm is widely used as a primitive in computer vision work (Huertas and Nevatia, 1988: Medioni and Nevatia, 1984); and second, edges are characterized by direction as well as by magnitude, a feature not retained with classical operators such as the Sobel operator and that is necessary for edge thinning operations. A brief description of the algorithm follows.

A 5 × 5 neighbourhood around each image pixel is convolved with six different edge masks corresponding to edge directions of 0, 30, 60, 90, 120, and 150 degrees. The magnitude of the convolved output, hereafter referred to as the edge magnitude, and the direction of the mask giving the highest response at each pixel are recorded as edge data.

Edge thinning is accomplished by retaining, as edges, only those pixels that meet the following criteria:

- the edge magnitude of the pixel in question exceeds the edge magnitude of its two neighbours in directions normal to this edge;
- the edge directions of these two neighbours differ by no more than 30 degrees from the edge direction of the central pixel; and
- the edge magnitude exceeds a predefined threshold level. The threshold has been selected to be the magnitude for a random radiometric noise level of one grey level.

The major edge detection steps are illustrated in the sequence of example images in Figure 1.

METHODS FOR MEASURING AND PREDICTING FILTER PERFORMANCE

We propose to rate a given filter based on its ability to improve real edge detection and on the level of false edges (artefacts) that are detected arising from residual, unsuppressed noise.

Real Edge Recovery

Consider the case of a boundary between two homogeneous regions m and n on a SAR image. By homogeneous, we mean that the observed mean square to variance ratio of the region's pixel

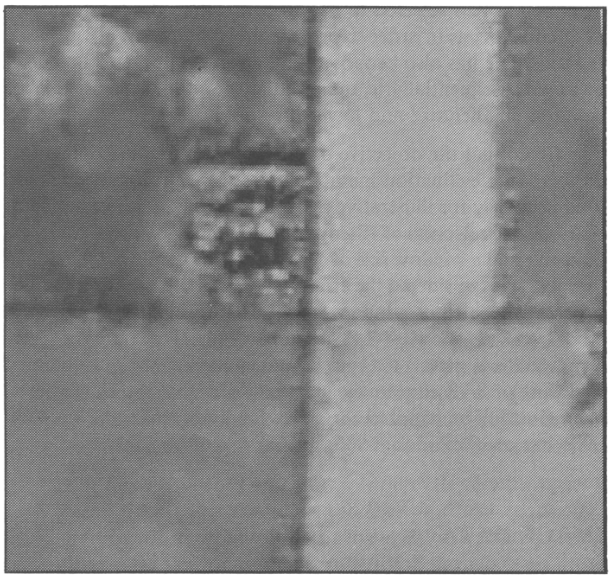

Figure 1a

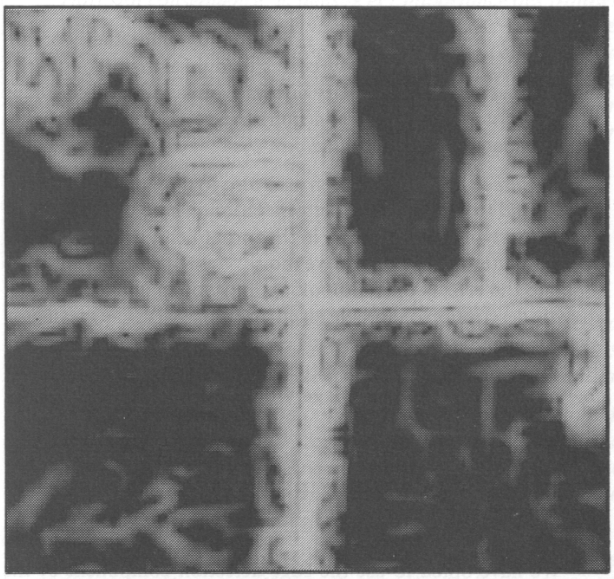

Figure 1b

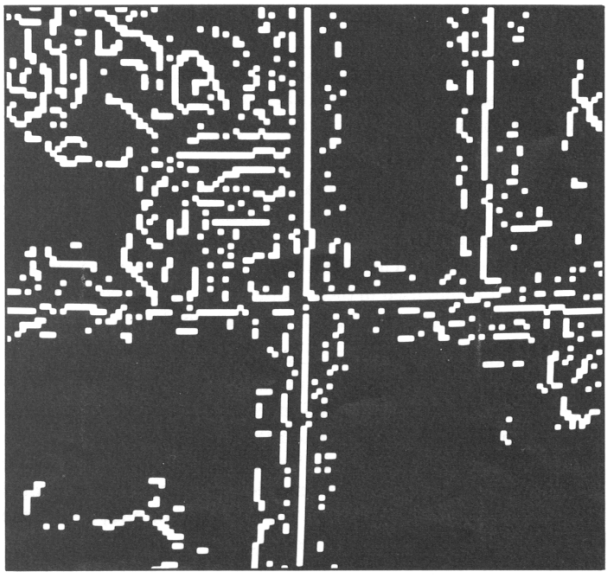

Figure 1c

Figure 1
An example of the application of the Nevatia and Babu edge detection algorithm to SAR imagery:
a) input filtered image, b) response from the convolution of edge masks, c) binary image following thinning.

grey levels is equal to the number of looks. If the scene is first filtered to reduce speckle and then processed with an edge detector, the performance of the filter can be gauged by the parameter F_{mn} the fraction of boundary edge pixels that are detected. Different filters can be compared through a statistical analysis of their resulting F_{mn} values using the classical binomial distribution theory (Snedcor and Cochran, 1976).

To use the above methodology with real imagery, one must first identify image segments containing regions that closely approximate the homogeneity condition and whose real edges can be identified by independent means. A set of potential candidates has been extracted from the Melfort SAR scene, which consists of field boundaries meeting the following criteria:

- The boundary is at least 30 pixels in length; and
- A visual inspection of aerial photography does not reveal the presence of surface cover variations within the bounding fields (i.e., the presence of standing water or regions of retarded crop growth.)

Apparent field homogeneity at optical wavelengths does not necessarily ensure homogeneity at radar wavelengths. To further identify potential problem cases, the grey level mean square to variance ratio (MSVR) for each bounding field was computed. Although we are using an amplitude image, the dynamic range of the fields in question is small enough that we may assume the MSVR for those fields that are "homogeneous" should be approximately the same and correspond to the effective number of looks of the scene. Therefore, we have computed the average value and standard deviation of the distribution of MSVR's for all candidate fields. All fields whose ratio differed from the mean by more than two standard deviations were then eliminated. The average MSVR of the remaining fields was computed to be 17.

As a precursor to estimating F_{mn} for a given candidate boundary, a best estimate of the true boundary location was made from a visual inspection of the raw SAR image. This best estimate will be referred to as the 'reference boundary'. Following filtering and edge detection, a real boundary pixel is deemed to be recovered if an edge pixel exists in the recovered edge map at a reference edge location, and the edge directions of the recovered and reference boundaries agree to within 30 degrees (i.e., the angular quantization the edge templates). If positional coincidence is not found for a given reference pixel, a search is also made in the two adjacent pixel positions perpendicular to the boundary direction for a recovered edge. This extended search area has been included to allow for possible errors in the manual positioning of the reference boundary.

A compilation of measured F_{mn} values is presented in Table 1 for the application of each of the filters on each of the test boundaries. We have applied a binomial analysis to determine the significance of the difference between the two F_{mn} values for the same boundary (Snedcor and Cochran, 1976). If the difference in F_{mn} values exceeds the 95 per cent limit in the uncertainty of the lowest recovery rate, we deem the recovery rate of the other filter to be significantly better.

An inspection of Table 1 indicates that, overall, in the cases of 33 of the 45 boundaries, application of the Frost filter results in a greater recovery rate than smoothing with the Sigma filter. In addition, of the 15 cases where there is a significant difference, 12 cases arise when the Frost filter outperforms the Sigma filter. We conclude that the Frost filter is preferred from the point of view of extracting real edge information when these filters are used in this configuration.

Table 1
A comparison of observed real edge recovery rates for the application of the Frost and Sigma filters to the 45 test boundaries

Test Boundary Number	Boundary Length (Pixels)	Frost Fmn	Sigma Fmn	Filter with Larger Fmn	Best Filter Where Difference is Significant
1	44	0.93	0.93		
2	62	0.94	0.87	F	
3	36	0.86	0.88	S	
4	55	0.91	0.93	S	
5	49	0.78	0.78		
6	45	0.87	0.84	F	
7	45	0.80	0.69	F	
8	53	0.92	0.96	S	
9	47	0.96	0.85	F	
10	71	0.80	0.54	F	F
11	54	0.93	0.83	F	
12	38	0.71	0.58	F	
13	52	0.92	0.87	F	
14	30	0.70	0.73	S	
15	60	0.82	0.87	S	
16	77	0.81	0.74	F	
17	49	0.69	0.60	F	
18	35	0.77	0.60	F	
19	59	0.90	0.49	F	F
20	101	0.65	0.53	F	F
21	61	0.92	0.95	S	
22	139	0.96	0.95	F	
23	139	0.94	0.92	F	
24	45	0.96	0.93	F	
25	91	0.93	0.89	F	
26	130	0.87	0.78	F	F
27	96	0.40	0.36	F	
28	84	0.88	0.75	F	F
29	104	0.75	0.69	F	
30	136	0.95	0.92	F	
31	136	0.62	0.53	F	F
32	101	0.64	0.54	F	
33	95	0.72	0.62	F	
34	94	0.70	0.86	S	S
35	86	0.22	0.40	S	S
36	78	0.64	0.50	F	F
37	119	0.78	0.73	F	
38	120	0.64	0.61	F	
39	125	0.87	0.74	F	F
40	106	0.44	0.52	S	S
41	67	0.84	0.73	F	
42	48	0.96	0.81	F	F
43	120	0.45	0.48	S	
44	125	0.84	0.74	F	F
45	176	0.87	0.77	F	F

MODEL BASED PREDICTION OF REAL EDGE RECOVERY RATE

Although the evaluation of filters using candidate edges from real scenes is desirable, use of simulation is attractive as well since it allows one to study a more diverse set of conditions affecting edge detectability and to classify detected edges as real or as artefacts since the underlying image signal is controlled and known precisely. In this and the following section, we derive predictive models for real edge recovery and artefact generation. These models are useful not only for further filter comparison but are of value for confidence testing of the edges found in real scenes.

To develop a predictive model, one must first identify those scene parameters likely to affect the edge recovery rate. Consider the case of a SAR scene of N_L looks and a boundary in that scene between two homogeneous regions whose mean grey levels are $\bar{g}_m$

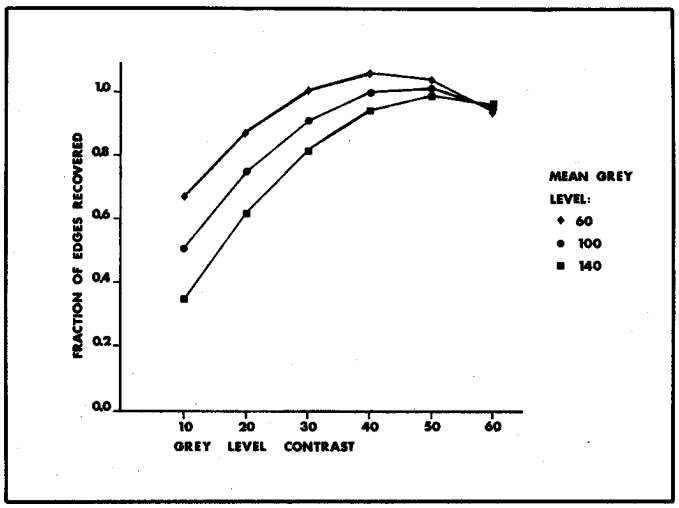

Figure 2a

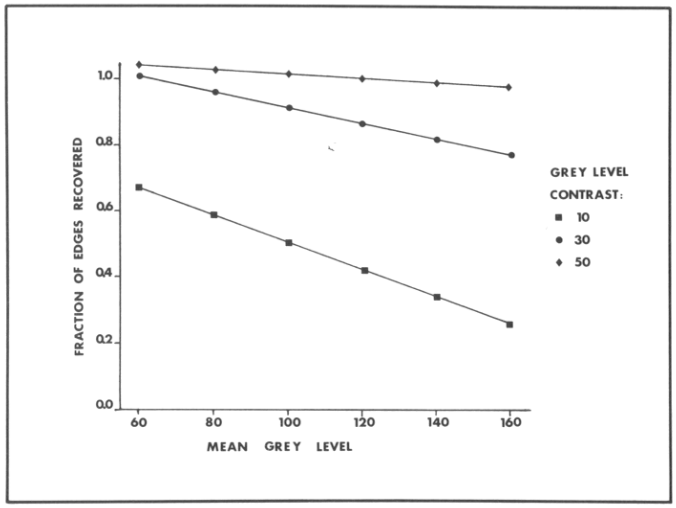

Figure 2b

Figure 2
Profiles of the polynomial fit of the real edge recovery rate as a function of (a) grey-level contrast and (b) mean grey level. The top curve in each case slightly exceeds $F_{mn} = 1.0$ since the polynomial fit was not restricted to lie below 1.0. In these cases the predicted F_{mn} is truncated to 1.0. The fit is valid only for $\Delta\bar{g}_{mn} < 60$ and $60 < \bar{g}_{mn} < 160$.

and $\bar{g}_n$. One obvious parameter affecting the detectability of the boundary will simply be its contrast, $\Delta\bar{g}$mn, where

$$\Delta\bar{g}_{mn} = |\bar{g}_m - \bar{g}_n|.$$

A second factor will be noise level characterized by N_L since noise can corrupt boundary edge magnitude or directional fidelity. Since residual noise will even persist following filtering, it is necessary to understand the nature of this residual noise.

In the case of both the Frost and Lee Sigma filters, the residual noise will be multiplicative. This arises because, in every region deemed to be homogeneous, the same level of spatial averaging is applied independent of the region's signal level. The effect of spatial averaging is a reduction of the noise in all homogeneous regions by the same scaling factor, but not to the same grey level sigma. In other words, the effect is equivalent to multilooking.

In conclusion, the so-called 'adaptiveness' of these filters applies only to their ability to adapt the level of spatial averaging to the perceived presence or absence of edge content within the filter window. Since the residual noise remains multiplicative, we predict that boundary detection will be a function of region brightness as well as boundary contrast. We quantify this brightness as the mean grey level of the two regions:

$$\bar{g}_{mn} = (\bar{g}_m + \bar{g})/2.$$

Our final model for a given filter will then be of the form:

$$F_{mn} = f(\bar{g}_{mn}.\ \Delta\bar{g}_{mn}, N_L).$$

For the purpose of comparison with the real candidate boundaries (i.e., $N_L = 17$ looks), one need only predict F_{mn} as a function of $\bar{g}_{mn}$ and $\Delta\bar{g}_{mn}$. For a given $(\bar{g}_m, \bar{g}_n)$ set, we can predict F_{mn} by generating a test pattern consisting of two blocks with the mean grey levels g_m and g_n, injecting 17 look multiplicative noise, filtering with the desired filter, applying edge extraction, and counting the number of recovered block boundary pixels. Since this procedure would be tedious for general purpose prediction, we have instead generated F_{mn} estimates on a regular ($\bar{g}_{mn}$, $\Delta\bar{g}$m) grid and approximated the recovery rate surface by a two-dimensional polynomial function of the form

$$F_{mn} = a + b*\bar{g}_{mn} + c*\Delta\bar{g}_{mn} + d*\bar{g}_{mn} + e*(\Delta\bar{g}_{mn})^2,$$

where a, b, c, d, and e are the least squares polynomial coefficients. The term in $\bar{g}_{mn}^2$ was discarded since its coefficient was found to be statistically insignificant. The polynomial provides a rapid method of performance prediction for any grey level combination.

To illustrate the results, Figures 2(a) and 2(b) show profiles of the surface in planes of constant $\bar{g}_{mn}$ and $\Delta\bar{g}_{mn}$, respectively. The recovery rate dependence on each parameter is evident.

To assess the accuracy of these performance models, predictions of F_{mn} values have been calculated for the 45 test boundaries compiled from the Melfort scene. Plots of observed versus predicted F_{mn} are presented in graphical form in Figures 3 and 4 for the Sigma and Frost filters, respectively. If the models are accurate, the relationship between observed and predicted F_{mn} in each case should be linear with a slope of 1. We have computed the least squares slope for each dataset using the constraint that the fitted line pass through the origin. The fitted slopes and the one-sigma uncertainties were found to be:

Frost filter: slope = 0.976 ± 0.096
Sigma filter: slope = 0.963 ± 0.069

We conclude that these simulation-based models provide good predictive capabilities for real edge recovery.

ARTEFACT LEVEL ESTIMATION AND COMPARISON

The analysis of artefacts is accomplished most easily by studying the derived edge content of filtered homogeneous regions. Since we do not have a precise knowledge of the scattering properties of the agricultural fields of the real image, we have generated simulated images of homogeneous regions composed of an underlying signal of a constant grey level with 17-look speckle noise added. Hence, any edges detected following filtering can only be artefacts arising from residual noise.

We quantify the artefact level by two parameters, namely, the fraction of pixels exhibiting an edge (F_A), and an edge magnitude

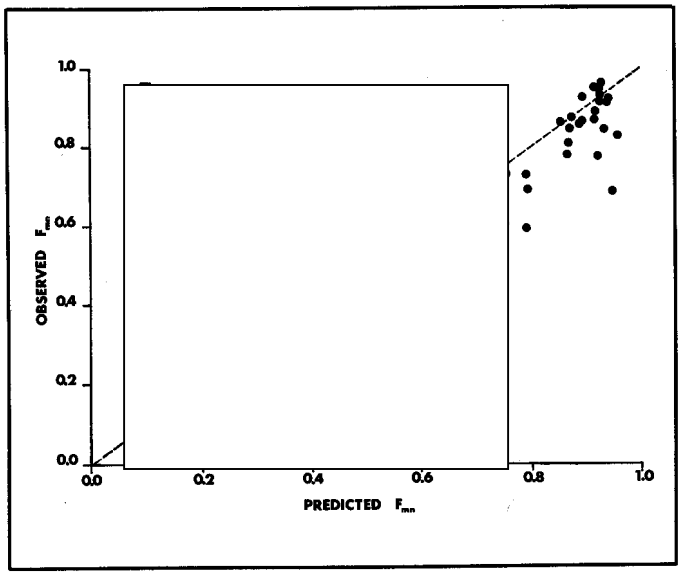

Figure 3
Comparison of predicted and observed values of F_{mn} for the 45 candidate field boundaries following the application of sigma filter. The error bar in the upper left-hand portion of the diagram indicates the average uncertainty in the observed F_{mn} values due to the finite length of the real field boundaries. The least squares linear fit to these data does not differ significantly from the unity line (i.e., the dashed line).

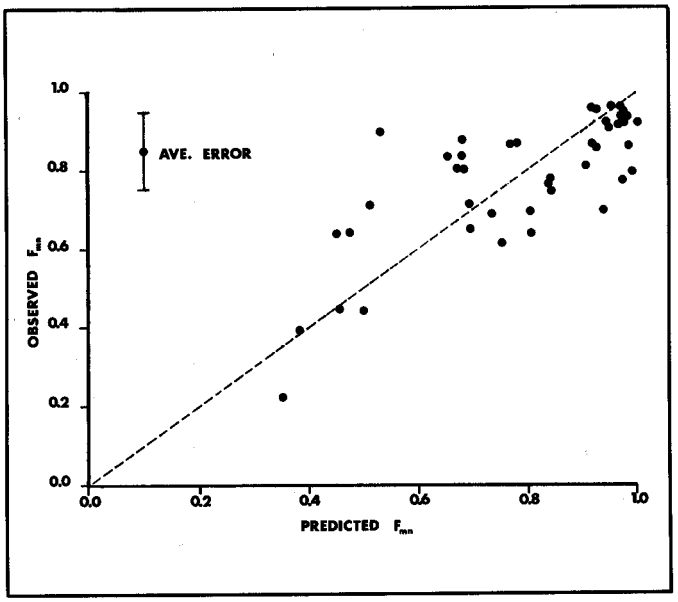

Figure 4
Comparison of predicted and observed values of F_{mn} for the 45 candidate field boundaries following the application of the Frost filter. The error bar in the upper left-hand portion of the diagram indicates the average uncertainty in the observed F_{mn} values due to the finite length of the real field boundaries. The least squares linear fit to these data does not differ significantly from the unity line (i.e., the dashed line).

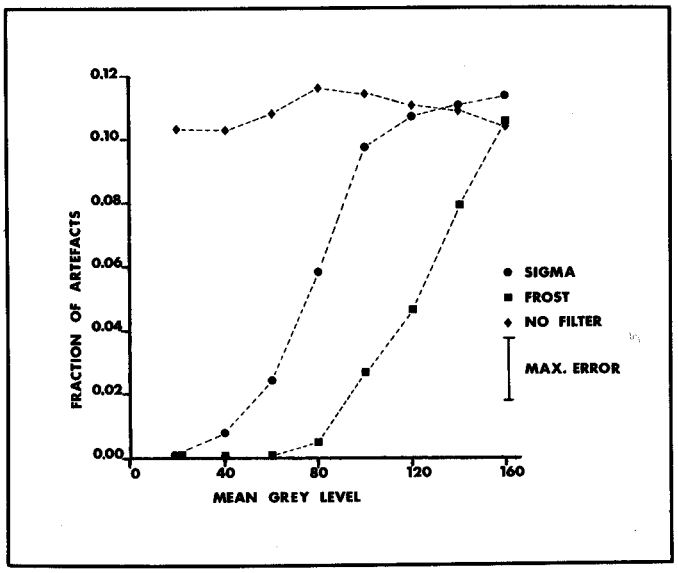

Figure 5
Fraction of pixels in homogeneous regions exhibiting an edge as a function of mean region grey level. The unfiltered image shows an artefact level that is independent of the mean grey level. Both filters show significant improvement in F_A, with the Frost filter being slightly better. The error bar is calculated for $F_A \sim 0.1$; as F_A approaches 0, the binomial error also approaches 0.

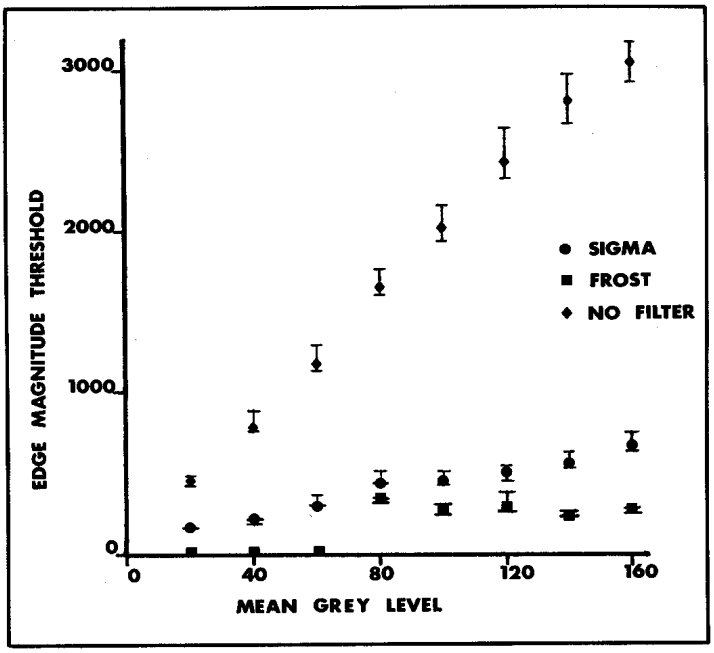

Figure 6
Edge magnitude threshold encompasses 95 per cent of the artefacts in a homogeneous region, as a function of the mean region grey level. Filtering significantly reduces artefact edge magnitude; however, the residual noise in homogeneous regions still exhibits a multiplicative nature following filtering (i.e., a dependence on the mean grey level).

threshold that encompasses 95 per cent of the detected edges (T_{95}). Figures 5 and 6 illustrate the dependencies of F_A and T_{95} on the homogeneous region signal level (i.e., mean grey level) for 17-look speckle. Besides the test filters, artefact levels for the case of no filtering prior to edge detection are included for comparison. The principal results of the comparisons are summarized as follows:

- If no filtering is applied, significant levels of false edges will be found. Both filters are found to be useful in significantly reducing artefact levels, particularly in low signal level regions.

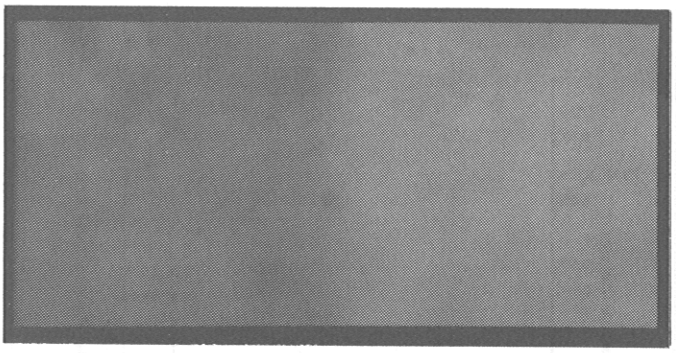

Figure 7a

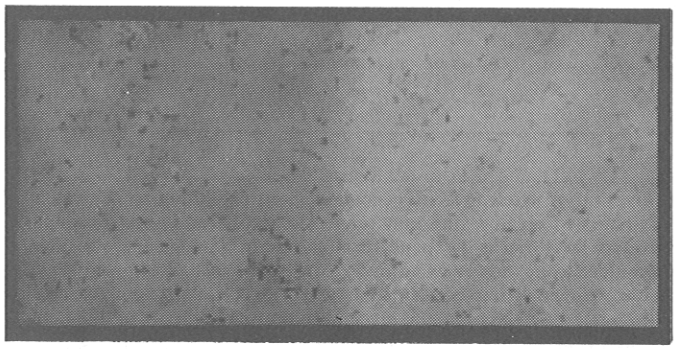

Figure 7b

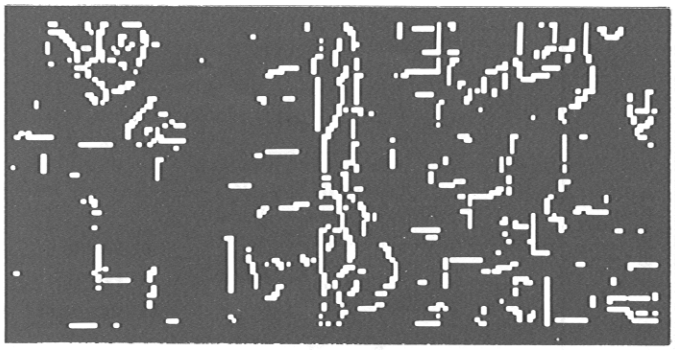

Figure 7c

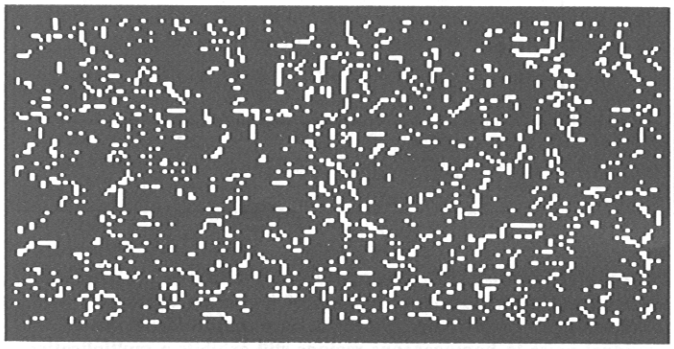

Figure 7d

Figure 7
Comparison of filter performance in homogeneous regions: (a) Frost-filtered image, (b) Sigma-filtered image, (c) artefact edges remaining following Frost filtering, and (d) artefact edges remaining following Sigma filtering. Note the many dark spots in 7(b), which are absent in 7(a) due to limited averaging of deviant pixels.

- Of the two filters, the Frost filter is significantly better at reducing both the density and magnitude of artefacts. A quantitative understanding of this result can be gained by comparing Frost and Sigma filtered images of the same parent noisy region. An example case of two regions of constant grey level is shown in Figure 7. Figures 7(a) and 7(b) show the Frost filtered and Sigma filtered images, respectively, while Figures 7(c) and 7(d) show their respective edge images. Since the Frost filter selects the level of spatial averaging based on a statistic derived from all pixels within the filter window, and since the window is large, this statistic, and hence the averaging process, does not vary significantly on a pixel to pixel scale. The residual grey level fluctuations following Frost filtering are therefore of a low spatial frequency.

  On the other hand, the specific pixels selected for averaging with the Sigma filter can vary dramatically from anyone pixel to its neighbour since the grey-level selection range is governed by the central pixel's grey level. In addition, averaging can be limited to four pixels, independent of the filter window size, if the central pixel is particularly bright or dark. This limited averaging tends to result in a grey-level segregation between these deviant pixels and the rest of the population. The ultimate result is a residual high frequency noise, which triggers artefacts of higher density and magnitude than is the case for the Frost filter.

- For null filtering, T_{95} increases monotonically with region mean brightness, but F_A remains approximately constant at about 0.10. This constancy reflects the fact that the principal edge thinning criteria (i.e., comparative edge magnitude dominance and directional consistency) are independent of absolute edge magnitude. Although bright homogeneous regions will exhibit brighter edges, the spatial randomness of the parent noise and subsequently the thinned edge density should be brightness-independent.

From this discussion, we conclude that artefacts will be present even following filtering and that these should be excluded by edge magnitude thresholding. To illustrate this point, we consider the example problem of determining whether a set of fields, extracted from the Melfort SAR scene, exhibit significant within-field texture. Sequences of images of a proposed analysis procedure are shown in Figures 8 and 9 for two example cases. In each sequence, images (a) and (b) illustrate the Frost filtered image and all recovered edges, respectively. Since the level of artefact generation within any field will be a function of the field's mean grey level, the magnitude of the artefact threshold (i.e., T_{95}) must be selected and applied on a per field basis. The third image in each sequence shows the remaining edges following thresholding of the central portions of each field. Two conclusions can be drawn from these figures:

- most within-field edges observed in Figures 8(b) and 9(b) exhibit magnitudes at levels expected from residual noise; and
- the apparent relationship of high-edge density with high mean field brightness is accountable by the multiplicative nature of the residual unfiltered noise.

The implications of this analysis are equally applicable to SAR imagery of any terrain type, namely, that the effects of signal

dependence and spurious texture levels expected for a nominally 'homogeneous' region should be determined and removed before any interpretation of image tonal variations.

CONCLUSIONS AND DISCUSSION

Methods have been developed to evaluate the relative and absolute effectiveness of speckle-reducing filters. The performance of a given filter is gauged from the characteristics of edges extracted from the resulting filtered image. Edge extraction has been selected since it constitutes a basic operation in image segmentation, texture, and linear feature extraction. Two adaptive filters (Frost and Lee Sigma) and the edge extraction algorithm of Nevatia and Babu (1980) have been employed to illustrate the proposed evaluation methodologies. The important issues noted in this paper are summarized as follows:

- The presence of speckle not only masks real edge information but also triggers spurious edges (artefacts) unrelated to variations in the backscatter properties of the reflecting surface. Therefore, filters must be gauged in terms of both real edge recovery and artefact suppression.
- Adaptive filters such as the Frost and Lee Sigma filters are adaptive in the sense they vary the level of spatial averaging according to the perceived edge content within the filter window. Residual speckle, following filtering, retains its multiplicative character.
- Filter performance with regard to real edges can be quantified by a real edge recovery rate F_{mn}. This parameter is simply the fraction of boundary edges between two homogeneous regions (mean grey levels $\bar{g}_m$ and $\bar{g}_n$, which are detected following filtering. Given the multiplicative nature of residual noise, F_{mn} can be expected to be a function of absolute region grey level, as well as boundary grey level, contrast. This prediction has been confirmed with the aid of simulated imagery. Models of real edge recovery performance can also be derived using simulation.
- Artefact generation also has been explored with the aid of simulation. Because of the multiplicative nature of residual noise, the level of artefact generation (both magnitude and spatial density) increases with increasing region brightness.
- Comparative performance studies with the proposed measures are easily achieved since the statistical signif - icance of observed differences can be estimated with binomial statistical theory.
- Based on tests with 17-look imagery, the Frost filter outperforms the Lee Sigma filter in this configuration.

The existence of artefacts must be accounted for when an interpretation is made of SAR image parameters derived from local grey level variability. A principal example is the case of texture measures, such as those proposed by Haralick (1979). Our proposed method consists of applying the sequence of filter and texture algorithms to simulated imagery to estimate the magnitude and signal-dependence of the spurious texture generated by speckle alone. For any image sub-area of interest (e.g., an agricultural field), one can compare the observed texture magnitude and density distribution with the expected artefact levels to test for the existence of additional texture related to surface cover variability, and apply thresholding to minimize the artefact contribution.

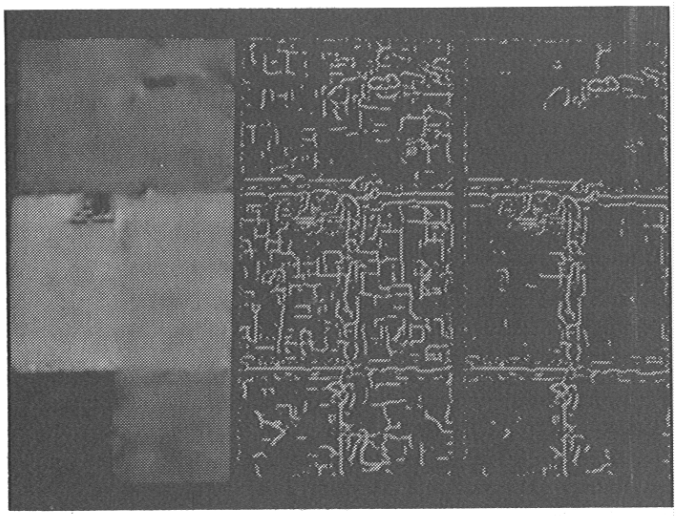

Figure 8

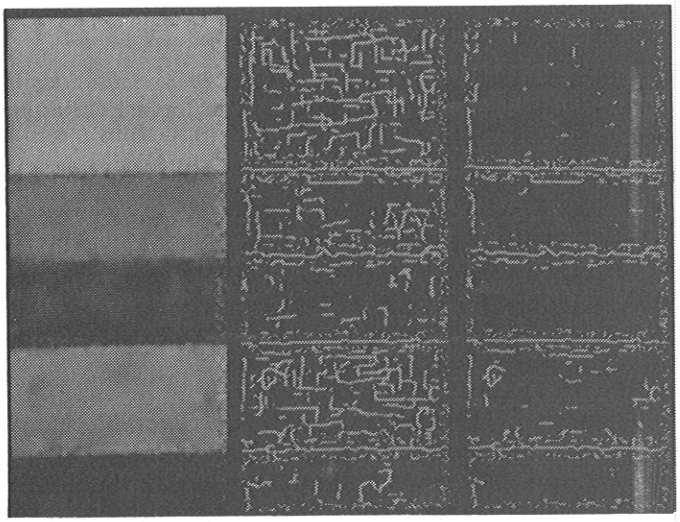

Figure 9

Figures 8, 9
Examples of a possible adaptive thresholding technique: (a) a section of a real SAR scene, Frost filtered, (b) binary edge image from (a), and (c) the result of applying a threshold that varies according to the mean grey level of the region. The edge pixels remaining in the homogeneous regions are most likely due to within-field texture.

The signal dependence of residual noise is a significant drawback for the adaptive filters described here. In theory, this problem can be eliminated by converting to a log power scale, thereby transforming speckle to an additive noise. Alternate, nonparametric measures of spectral homogeneity, which exhibit signal independent characteristics, are currently being investigated and will be presented in a future publication.

REFERENCES

Brown, R.J., Guindon, B., Teillet, P.M. and Goodenough, D.G. (1984) -*Crop Type Determination from Multitemporal SAR Imagery*, **Proceedings, Ninth Canadian Symposium on Remote Sensing, pp. 683-691.**

Crimmins, T .R. (1985) -*Geometric Filter for Speckle Reduction,* **Applied Optics, Vol. 24, pp. 1438-1443.**

Durand, J.M., Gimonet, B.J. and Perbos, J.R. (1987) -*SAR Data Filtering for Classification,* **IEEE Transactions on Geoscience and Remote Sensing, Vol. GE-25, pp. 629-637.**

Frost, V .S., Stiles, J.A., Shanmugan, K.S. and Holtzman, J.C. (1982) -*A Model for Radar Images and Its Application to Adaptive Digital Filtering of Multiplicative Noise,* **IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. P AMI-4, pp. 157-166.**

Goodenough, D.G., Guindon, B., Meunier, J.-F. and Swanberg, N.A. (1984) -*Adaptive Filtering and Image Segmentation for SAR Analysis,* **Proceedings Tenth International Symposium on Machine Processing of Remotely Sensed Data, pp. 315-324.**

Guindon, B. (1985) -*Automated Control Point Acquisition in Radar-Optical Image Registration,* **Canadian Journal of Remote Sensing, Vol. 11, pp. 103-112.**

Guindon, B. and Maruyama, H~(1986) -*Automated Matching of Real and Simulated SAR Imagery as a Tool for Ground Control Point Acquisition,* **Canadian Journal of Remote Sensing, Vol. 12, pp.149-158.**

Haralick, R.M. (1979) -*Statistical and Structural Approaches to Texture,* **Proceedings, IEEE, Vol. 67, pp. 786-804**.

Huertas, A. and Nevatia, R. (1988) -*Detecting Buildings in Aerial Images*, **Computer Vision and Image Processing, Vol. 41, pp.131-152.**

Lee, J.-S. (1983) -*A Simple Speckle Smoothing Algorithm for Synthetic Aperture Radar Images,* **IEEE Transactions on Systems, Man and Cybernetics, Vol. SMC-13, pp. 85-89.**

Lee, J .-S. (1987) -*Statistical Modelling and Suppression of Speckle in Synthetic Aperture Radar Images,* **1987 IEEE International Geoscience and Remote Sensing Symposium Digest, pp. 1331-1339.**

Machuca, R. and Gilbert, A.L. (1981) -*Finding Edges in Noisy Scenes,* **IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. PAMI-3, pp. 103-111.**

Medioni, G. and Nevatia, R. (1984) -*Segment-Based Stereo Matching,* **IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. PAMI-6, pp. 675-685.**

Modestino, J.W. and Fries, R.W. (1977) -*Edge Detection in Noisy Images Using Recursive Digital Filtering*, **Computer Graphics and Image Processing, Vol. 6, pp. 409-433.**

Nevatia, R. and Babu, K.R. (1980) -*Linear Feature Extraction and Description,* **Computer Graphics and Image Processing, Vol. 13, pp. 257-269.**

Snedcor, G. W. and Cochran, G.M. (1976) -*Statistical Methods,* **Iowa State University Press.**

Touzi, R., Lopes, A. and Bousquet, P. (1987) -*A Statistical and Geometrical Edge Detector for SAR Image Segmentation,* **IEEE International Geoscience and Remote Sensing Symposium Digest, pp. 1469-1474.**

Wood, J. W. (1985) -*Line Finding Algorithms for SAR,* **Royal Signals and Radar Establishment Memorandum 3841.**