An edge-enhancing nonlinear filter for reducing multiplicative noise

Mark A. Schulze
Perceptive Scientific Instruments, Inc.
League City, Texas

ABSTRACT

This paper illustrates the design of a nonlinear filter for edge-enhancing smoothing of multiplicative noise using a morphology-based filter structure. This filter is called the Minimum Coefficient of Variation (MCV) filter. The coefficient of variation is the ratio of the standard deviation of a random process to its mean. For an image corrupted only by stationary multiplicative noise, the coefficient of variation is theoretically constant at every point. Estimates of the coefficient of variation indicate whether a region is approximately constant beneath the multiplicative noise or whether it contains significant image features. Regions containing edges or other image features yield higher estimates of the coefficient of variation than areas that are roughly constant. The MCV filter uses a morphological structure to direct low-pass filtering to act only over regions determined to be most nearly constant by measuring the coefficient of variation.

Examples of the use of the MCV filter are given on synthetic aperture radar (SAR) images of the earth. SAR images are corrupted by speckle, a predominantly multiplicative noise process. Therefore, the MCV filter is a good choice for reducing speckle without blurring edges. The MCV filter is useful for pre-processing in image analysis applications such as coastline detection in SAR images.

Keywords: nonlinear filtering, multiplicative noise, edge enhancement, synthetic aperture radar, mathematical morphology

1. INTRODUCTION

Nonlinear filters are widely used in image processing for reducing noise without blurring or distorting edges [1]. Median filters were among the first nonlinear filters used for this purpose [2, 3], but more recently a wider variety of options have become available, including mathematical morphology [4], anisotropic diffusion [5], and wavelet-based enhancement techniques [6]. Most of these methods focus on the problem of removing additive noise from images, since this is by far the most common type of corrupting noise. However, images acquired using a coherent illumination source such as lasers or synthetic aperture radar (SAR) exhibit speckle noise, which is a predominately multiplicative noise process. Multiplicative noise is, for most images, strikingly different from additive noise, and many of the nonlinear filtering techniques used to remove additive noise are not as effective at removing multiplicative noise.

The value-and-criterion filter structure [7] was developed to widen the variety of operations that can be used within a geometric filter structure defined by mathematical morphology. This structure has been used previously to develop a filter that reduces noise and enhances edges in images corrupted by additive noise [8]. This paper illustrates the use of the value-and-criterion filter structure to design a similar filter for use in multiplicative noise environments.

Further author information —
Email: schulze@persci.com; WWW: http://www.persci.com/~schulze/

2. FILTER DESIGN

2.1. Value-and-criterion filter structure

The value-and-criterion filter structure has been described previously [7–9]. It is based on the geometrical structure of mathematical morphology, but allows the use of a much wider variety of operations (both linear and nonlinear) than standard morphology. Morphological filters usually only use extreme order statistic operations (maximum and minimum). The new filter structure is much more flexible, but still retains the geometric characteristics of mathematical morphology.

The value-and-criterion filter structure is derived from the "compound" structure of the morphological operations opening and closing. Opening and closing are both sequential operations; opening consists of the successive application of the basic morphological operations of erosion and dilation, and closing is dilation followed by erosion. Like opening or closing, a value-and-criterion function has two distinct stages. However, a value-and-criterion filter can have two separate operations acting in parallel in the first stage, as opposed to a single operation (such as erosion in the case of opening). The second stage of a value-and-criterion function uses the output of one of these first stage operations to determine which output from the other first stage operation is selected as the final output.

In the first stage of a value-and-criterion filter, two operations act on the input image $f(\mathbf{x})$: the "value" function $\mathcal{V}$ and the "criterion" function $\mathcal{C}$. The window, or structuring element, over which these functions are defined is denoted N. The second stage of the filter is a "selection" operator $\mathcal{S}$ that acts on the outputs of the first stage. $\mathcal{S}$ operates over a 180° rotation of the structuring element N, which is denoted $\tilde{N}$. Often, N is symmetric about its center and therefore $\tilde{N} = N$. Letting $g(\mathbf{x})$ denote the output of a value-and-criterion function, and $v(\mathbf{x})$ and $c(\mathbf{x})$ respectively denote the outputs of the value function and the criterion function, the output of the filter is defined by equations (2.1)–(2.3) below [7–9].

$$v(\mathbf{x}) = \mathcal{V}\{f(\mathbf{x}); N\} \tag{2.1}$$

$$c(\mathbf{x}) = \mathcal{C}\{f(\mathbf{x}); N\} \tag{2.2}$$

$$g(\mathbf{x}) = v\left(\left\{\mathbf{x}' : \mathbf{x}' \in \tilde{N}_{\mathbf{x}};\, c(\mathbf{x}') = \mathcal{S}\left\{c(\mathbf{x});\, \tilde{N}\right\}\right\}\right) \tag{2.3}$$

$\tilde{N}_{\mathbf{x}}$ denotes the translation of $\tilde{N}$ such that it is centered at position $\mathbf{x}$.

This filter structure may be interpreted as having a set of subwindows within an overall filter window. Figure 1 below illustrates this interpretation. The subwindows are the same as the structuring element N, and the overall window is the union of N translated to each point in the rotated structuring element $\tilde{N}$. In Figure 1, N (and therefore $\tilde{N}$) has three points, and so there are three 3-point subwindows within the overall window (shown at the bottom of the figure) of seven points. Notice that because of the rotation of N to become $\tilde{N}$, the overall window is symmetric about its center.

The operation of a value-and-criterion filter at a point is equivalent to examining each subwindow within the overall window centered at that point and finding which subwindow has the optimal criterion function output (as defined by the selection function). The output of the value function over that subwindow then becomes the final filter output for the point under consideration. Note that all possible subwindows N within the overall window are examined because of the way the overall window is constructed from N itself.

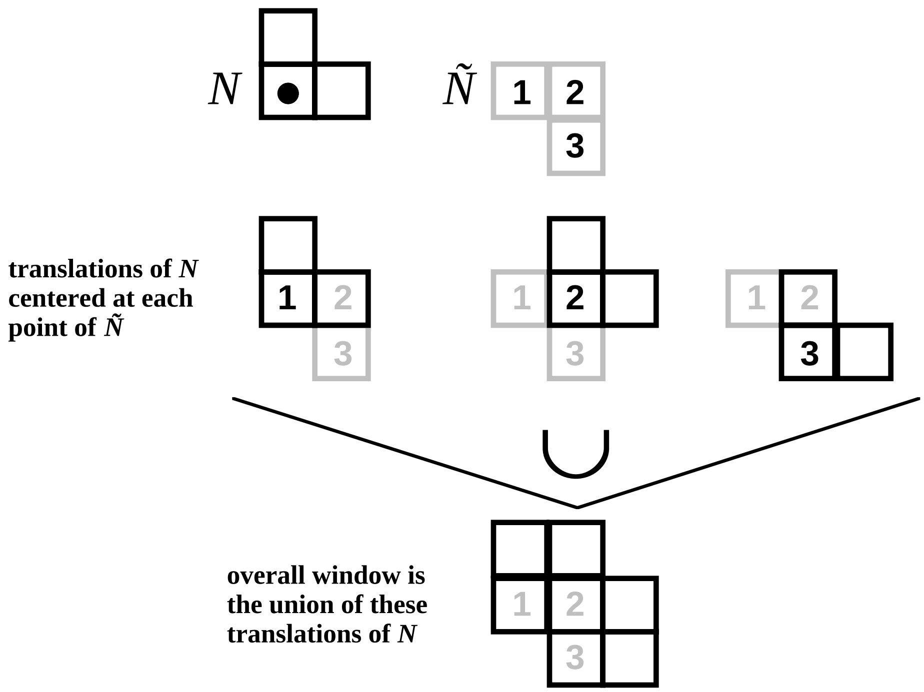

Figure 1. Formation of the overall window of a value-and-criterion filter with structuring element *N*.

The value-and-criterion filter lends itself well to the design of filters to smooth noise without blurring edges. This is accomplished by using a criterion function that detects the presence of image edges, and using a noise smoothing operation for the value function. The selection function should favor criterion function values that indicate that edges are not present within the structuring element. That way, the output of the value function is derived only from the one structuring element within the overall window that is least likely to contain an edge.

2.2. Multiplicative noise model

Multiplicative noise is generally more difficult to remove from images than additive noise, because the intensity of the noise varies with the signal intensity. A model for multiplicative noise is given in equation (2.4) below. The original (uncorrupted) signal is denoted by $s(\mathbf{x})$, the noise process by $n(\mathbf{x})$, and the noise-corrupted signal by $f(\mathbf{x})$. For most applications involving multiplicative noise, the noise process is assumed to be stationary with a mean of one ($\mu_n = 1$) and an unspecified variance σ_n^2.

$$f(\mathbf{x}) = s(\mathbf{x}) \cdot n(\mathbf{x}) \tag{2.4}$$

The coefficient of variation of a random process, denoted γ, is defined as the ratio of the standard deviation to the mean of the random process, as shown in equation (2.5) below. For a stationary noise process with $\mu_n = 1$, the coefficient of variation is equal to the standard deviation ($\gamma_n = \sigma_n$).

$$\gamma = \frac{\sigma}{\mu} \tag{2.5}$$

Since the signal $s(\mathbf{x})$ in equation (2.4) is a constant at any given $\mathbf{x}$, and the coefficient of variation is not a function of $\mathbf{x}$, the coefficient of variation is theoretically constant at each point $\mathbf{x}$ in an image that is corrupted by stationary multiplicative noise $n(\mathbf{x})$ with a mean of one. That is, $\gamma_f = \gamma_n = \sigma_n$.

An estimate of the coefficient of variation of an image corrupted by multiplicative noise may be made by computing the sample mean and sample standard deviation over an area in the image known (or assumed) to be constant. Since the image signal $s(\mathbf{x})$ is presumed to be constant over such an area, the ratio of the measured mean to the measured standard deviation should be a good estimate of the coefficient of variation of the noise process, which is theoretically the same as the true coefficient of variation at each point in the image.

2.3. Minimum Coefficient of Variation (MCV) filter

The Minimum Coefficient of Variation (MCV) filter is a value-and-criterion filter designed to reduce multiplicative noise in images without blurring edges. It uses an estimate of the coefficient of variation as its criterion function. The coefficient of variation over a structuring element is estimated by taking the ratio of the sample mean to the sample standard deviation over the structuring element. If the image is constant everywhere within the structuring element, the estimated coefficient of variation should be close to the theoretical value (which is the same as the coefficient of variation of the noise alone). However, if the image has an edge or other features within the structuring element, the sample standard deviation will increase and the sample mean will change (either increase or decrease). It can be easily shown that the effects of a non-constant signal $s(\mathbf{x})$ within a structuring element on the sample mean and standard deviation are to increase the estimated coefficient of variation above the value that would be found for a constant $s(\mathbf{x})$. Therefore, structuring elements within an image that have a constant underlying signal will usually have a lower measured coefficient of variation than structuring elements that contain an edge or other image feature.

The selection function of the MCV filter is the minimum, so that the filtering operation (value function) effectively acts over the structuring element in the overall window that has the smallest coefficient of variation. This means that the noise smoothing function acts only over those structuring elements selected as being most nearly constant (by having the smallest coefficient of variation). The operation chosen to be the value function should be a filter appropriate for the distribution of the corrupting noise process. For example, if the noise distribution is known to be Gaussian, the best linear estimate is the sample mean and therefore the sample mean is an appropriate choice for the value function for an image corrupted by multiplicative Gaussian noise. The MCV uses the sample mean as its value function.

The MCV filter therefore uses the sample mean for the value function, the coefficient of variation as the criterion function, and the minimum as the selection function. It is a value-and-criterion filter specifically designed to reduce multiplicative noise. The equations for the MCV filter are given in (2.6)–(2.8) below.

$$v(\mathbf{x}) = \frac{1}{|N|} \sum_{\mathbf{y} \in N_{\mathbf{x}}} f(\mathbf{y}) \tag{2.6}$$

$$c(\mathbf{x}) = \frac{\left(\frac{1}{|N|} \sum_{\mathbf{y} \in N_{\mathbf{x}}} \left[f(\mathbf{y}) - v(\mathbf{x}) \right]^2 \right)^{1/2}}{v(\mathbf{x})} \tag{2.7}$$

$$\mathrm{MCV}\{ f(\mathbf{x}); N \} = v\left(\left\{ \mathbf{x}' : \mathbf{x}' \in \tilde{N}_{\mathbf{x}};\, c(\mathbf{x}') = \min\left[c(\mathbf{y});\, \mathbf{y} \in \tilde{N}_{\mathbf{x}} \right] \right\} \right) \tag{2.8}$$

The MCV filter is closely related to a value-and-criterion filter designed to remove additive noise known as the Mean of Least Variance (MLV) filter [7–9]. The MLV filter uses the sample variance as its criterion function rather than the coefficient of variation. In images corrupted by additive noise, the variance is theoretically minimum in structuring elements where the signal is constant.

3. EXAMPLES

3.1. Edge enhancement

The MCV filter not only preserves sharp edges between flat areas in images, but also sharpens edges that are not perfect step edges. This is illustrated in the one-dimensional example in Figure 2 below. The original signal has a ramp edge that extends over 15 points between constant areas. When an MCV filter with a structuring element of width 9 is applied to this signal, the edge is sharpened significantly. Note that the sharpened edge is shifted slightly to the left of the center of the original edge. This is because in the absence of noise, the coefficients of variation for the areas with higher signal values are relatively lower than those for areas with lower signal values. Therefore, near the center of the edge, structuring elements to the right of the point being filtered have a slightly lower coefficient of variation than those on the left and are thus favored by the selection operator. In the presence of multiplicative noise, this effect is lessened and edges are more likely to be sharpened near their true centers.

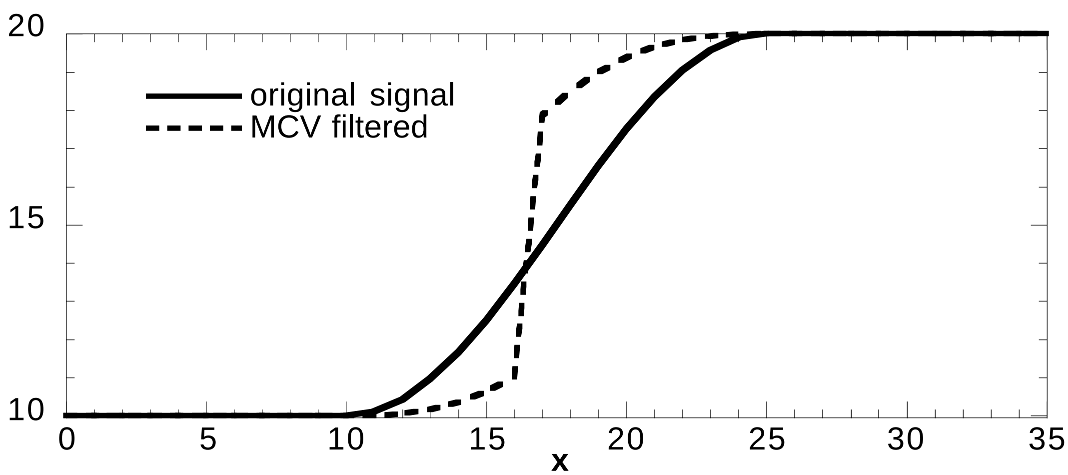

Figure 2. Edge enhancement example. Structuring element for the MCV filter is 9 points wide.

3.2. 1-D example

Figures 3 and 4 below illustrate the edge-preserving, noise-smoothing performance of the MCV filter on a 1-D signal. The original signal is a piecewise constant signal with constant areas 25 points wide. The baseline level is 10, and it has four evenly-spaced pulses of increasing height (at levels 25, 50, 100, and 200). This signal is corrupted by multiplicative Gaussian noise with a mean of 1 and a standard deviation of 1/3. Both the original and corrupted signals are shown in Figure 3 below. Figure 4 shows the results of filtering the noisy signal using an MCV filter with a structuring element 25 points wide, and also results of filtering using an MLV filter [7–9] with the same structuring element. The MCV filter is much more adept than the MLV filter at correctly identifying constant regions in multiplicative noise because it uses the coefficient of variation instead of the variance as its criterion function. The MCV filter correctly locates every edge, whereas the MLV filter consistently underestimates the width of the higher pulses. This demonstrates the value of using the MCV filter, which is specifically for multiplicative noise, over a filter designed for additive noise in this application.

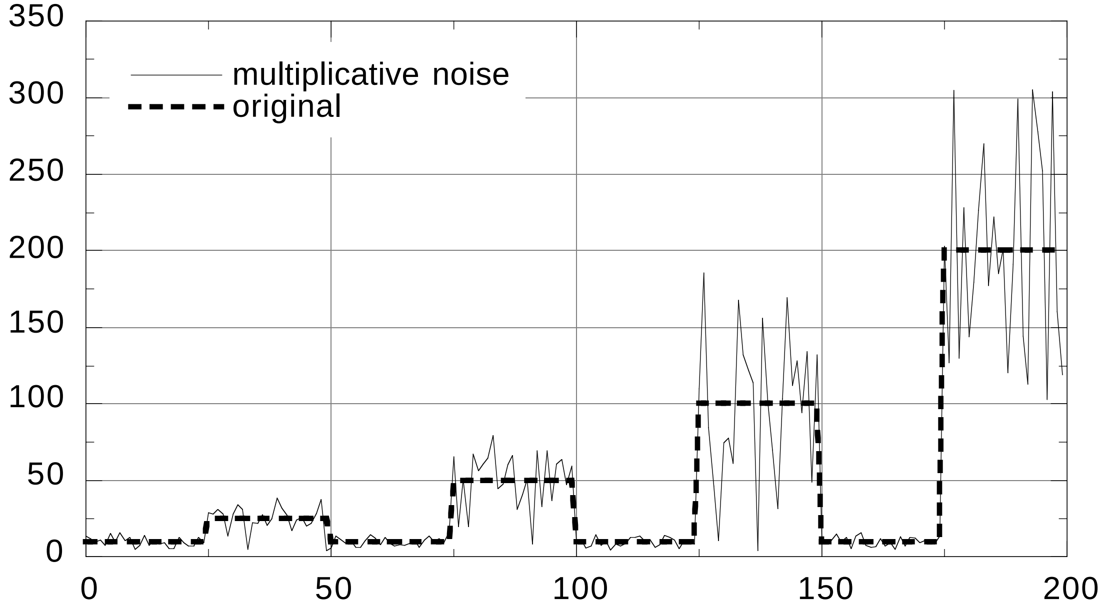

Figure 3. Original piecewise constant signal and signal corrupted by multiplicative noise.

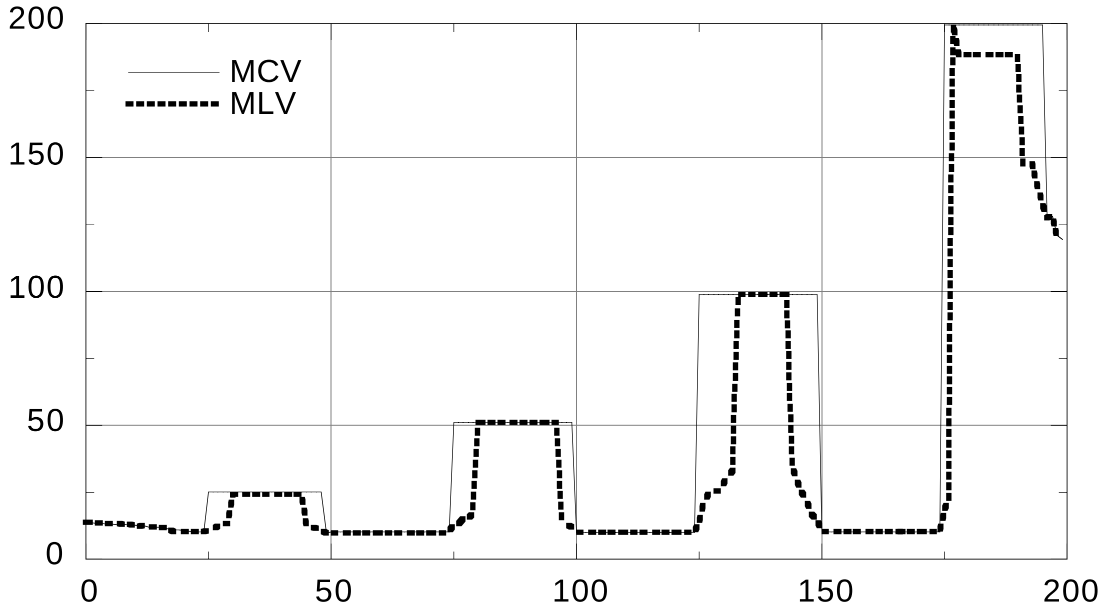

Figure 4. Noisy signal of Figure 3 filtered by MCV and MLV filters (width 25).

3.3. 2-D example

In this section, an image of several alphanumeric characters at various gray levels, shown in Figure 5(a), is used to demonstrate the behavior of the MCV filter on images corrupted by multiplicative noise. The image corrupted by multiplicative Gaussian noise with a mean of 1 and standard deviation of 1/5 is shown in Figure 5(b). The results of MCV filtering with a 3x3 structuring element are shown in Figure 5(c). Notice that the edges and corners of the characters are generally preserved correctly while the noise is much reduced. Figure 5(d) shows the result of MLV filtering with a 3x3 structuring element. The characters are distorted from the mislocation of edges due to the assumption of additive noise inherent in the MLV filter. The median filter is another nonlinear filter widely used for edge-preserving smoothing, and the results of a 3x3 median filter operating on the noisy image are given in Figure 5(e). Although it reduces the noise quite well, there are some distortions in the characters, especially at corners. Note, however, that the overall windows of the MCV and MLV filters used in Figures 5(c) and 5(d) are 5x5, and so these are perhaps more properly compared to the results of 5x5 median filtering, shown in Figure 5(h). This median filtered result shows much more distortion than the value-and-criterion filters in 5(c) and 5(d).

The results of MCV and MLV filtering using a 5x5 square structuring element are shown in Figures 5(f) and 5(g) respectively. Both show more distortion than the 3x3 filters, but the MCV filter preserves shapes in the image much more accurately. To illustrate the superior edge-preserving abilities of the nonlinear filters over linear filters, Figure 5(i) shows the result of mean filtering with a 3x3 window. As expected, the edges in the image are noticeably blurred. This shows that the MCV filter, even though it uses a mean filter as its value function, does not blur edges in an image corrupted by multiplicative noise because it uses the coefficient of variation to "guide" the filtering to areas that do not include edges.

One way to quantify the improvement made by filtering a noisy image is to compute the mean squared error (MSE) between the uncorrupted image and the filtered image. This is possible in this case because the image is synthetic and therefore the uncorrupted image is available. The MSEs over the entire image for the various filters shown in Figure 5 are given in the middle column of Table 1. The MCV filter with a 3x3 structuring element has the lowest MSE overall, and of the 5x5 filters, the MCV filter again has the lowest MSE. In both cases, the MSE of the MCV filter is less than half of the next-lowest error. This indicates the overall superiority of the MCV filter at smoothing noise and preserving edges in this image corrupted by multiplicative noise.

To compare the ability of the filters to reduce noise in regions without any edges, the MSE for the filters shown in Figure 5 was also measured only within the bright rectangle at the bottom of the image. To avoid any possible edge effects, the area taken was two pixels inside each boundary of the rectangle in the original image. The MSEs for the various filters within this constant area are given in the right column of Table 1. In this case, where no edges are present, the mean filter has the lowest MSE. This is expected, since it is the best linear estimate for Gaussian noise. Both the MCV and MLV filters use a mean filter as their noise smoothing function, but the criterion function of these filters reacts to some of the "structure" created by the noise and therefore the MSEs for these filters within this constant region are slightly higher than for the ordinary mean filter. The median filter has the highest MSE in this case. Notice however, that for the entire image, the MSE for the mean filter is much higher than for the other filters because of its blurring effect on edges.

(a) Original

(b) Noisy

(c) MCV 3x3

(d) MLV 3x3

(e) Median 3x3

(f) MCV 5x5

(g) MLV 5x5

(h) Median 5x5

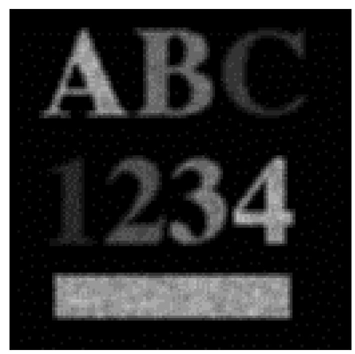

(i) Mean 3x3

Figure 5. Images demonstrating multiplicative noise filtering. (a) Original image; (b) corrupted by multiplicative noise; (c) MCV filtered [3x3]; (d) MLV filtered [3x3]; (e) median filtered [3x3]; (f) MCV filtered [5x5]; (g) MLV filtered [5x5]; (h) median filtered [5x5]; (i) mean filtered [3x3].

Table 1. MSE of filtered images compared to noiseless original.

Filter and size	MSE over entire image	MSE over constant area
(none - noisy image)	152	941
MCV 3x3	50	146
MLV 3x3	126	159
median 3x3	107	172
mean 3x3	264	110
MCV 5x5	125	41
MLV 5x5	302	51
median 5x5	308	64
mean 5x5	455	36

4. APPLICATION

4.1. Synthetic aperture radar (SAR) imaging

Synthetic aperture radar (SAR) is a coherently-illuminated imaging modality used in satellite remote sensing to image the earth's surface through cloudcover. SAR images are corrupted by speckle noise due to the coherent nature of the illumination [10]. This speckle noise can be accurately modeled as a purely multiplicative noise process [11]. A common image analysis problem in SAR remote sensing is to detect features and boundaries such as roadways and coastlines. The MCV filter is ideally suited for pre-processing SAR images to simplify these image analysis problems because it reduces the speckle noise while sharpening boundaries.

Figure 6 shows a section of a SAR image of a peninsula extending into Port Underwood in the Marlborough Sounds of New Zealand. The speckle noise is prominent on both the sea and land portions of the image. The results of MCV filtering this image with a 5x5 square structuring element are shown in Figure 7. The MCV filtered image has less noise and sharper edges than the original SAR image. Figure 9 shows the result of filtering the original image with a speckle noise reduction algorithm developed specifically for SAR by Durand [12]. The results of this so-called "local statistics" method are generally noisier than the MCV filtered results, and the edges are much less distinct.

Figures 9–11 show the results of the Canny edge detection algorithm [13] on the images in Figures 6–8. These images show that the important boundaries in the image (coastlines and ridgelines) are much more accurately found in the MCV filtered image than in either the unfiltered or local statistics filtered images. This shows that the MCV filter can be used to perform useful pre-filtering on SAR images to reduce speckle noise and enhance edges.

5. CONCLUSIONS

This paper demonstrates the design of an edge-enhancing noise smoothing filter for multiplicative noise using the value-and-criterion filter structure. The morphological basis of this filter structure means that this filter has some well-known geometrical properties. The ability of this filter to accurately locate edges in multiplicative noise and sharpen them is illustrated by one-dimensional examples. The superiority of this filter at reducing multiplicative noise on a synthetic image is also demonstrated. One application of this filter that shows much promise is as a pre-filter for reducing noise in SAR images before performing image analysis and feature detection.

ACKNOWLEDGMENTS

The author would like to thank his former colleagues at Landcare Research New Zealand Ltd., especially Ms. Stella Belliss, Dr. David Pairman, Dr. Qing Wu, and Dr. Stephen McNeill, for their assistance with the SAR imagery for this work.

This work was partially funded by the New Zealand Foundation for Research, Science and Technology under contract C09506.

REFERENCES

1. I. Pitas and A.N. Venetsanopoulos, "Edge detectors based on nonlinear filters," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, v. 8, n. 4, pp. 538-550, 1986.
2. J.W. Tukey, *Exploratory Data Analysis*, Addison-Wesley, Reading, Massachusetts, 1971.
3. N.C. Gallagher Jr. and G.L. Wise, "A theoretical analysis of the properties of the median filter," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, v. 29, n. 6, pp. 1136-1141, 1981.
4. J. Serra, *Image Analysis and Mathematical Morphology*, Academic Press, London, 1982.
5. P. Perona and J. Malik. "Scale space and edge detection using anisotropic diffusion," *IEEE Workshop on Computer Vision*, pp. 16-22, Miami, 1987.
6. J. Lu, D.M. Healy Jr., and J.B. Weaver, "Contrast enhancement of medical images using multiscale edge representation," *Optical Engineering*, v. 33, n. 7, pp. 2151-2161, 1994.
7. M.A. Schulze and J.A. Pearce. "Value-and-criterion filters: A new filter structure based upon morphological opening and closing," *Nonlinear Image Processing IV*, E.R. Dougherty, J. Astola, and H. Longbotham, editors, SPIE v. 1902, pp. 106-115, San Jose, California, 1993.
8. M.A. Schulze and J.A. Pearce. "A morphology-based filter structure for edge-enhancing smoothing," *1994 IEEE International Conference on Image Processing (ICIP-94)*, v. 1, pp. 530-534, Austin, Texas, 1994.
9. M.A. Schulze, *Biomedical Image Processing with Morphology-Based Nonlinear Filters*, Ph.D Dissertation, The University of Texas at Austin, 1994.
10. J.C. Curlander and R.N. McDonough, *Synthetic Aperture Radar: Systems and Signal Processing*, Wiley, New York, 1991.
11. J.W. Goodman, "Statistical properties of laser speckle patterns", in *Laser Speckle and Related Phenomena*, J.C. Dainty, editor , pp. 9-75, Spring-Verlag, Berlin, 1975.
12. J.M. Durand, B.J. Gimonet, and J.R. Perbos, "SAR data filtering for classification," *IEEE Transactions on Geoscience and Remote Sensing*, v. 25, n. 5, pp. 629-637, 1987.
13. J. Canny, "A computational approach to edge detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, v. 8, n. 6, pp. 679-698, 1986.

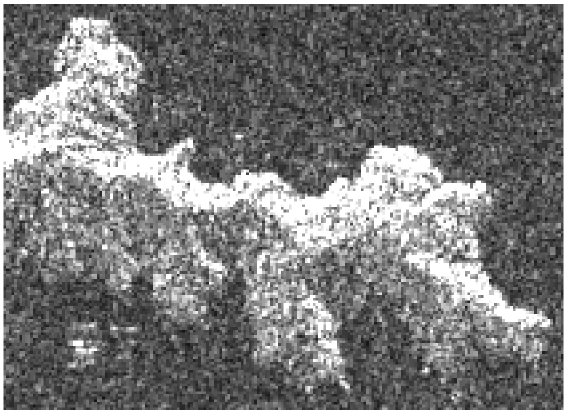

Figure 6. Original SAR image

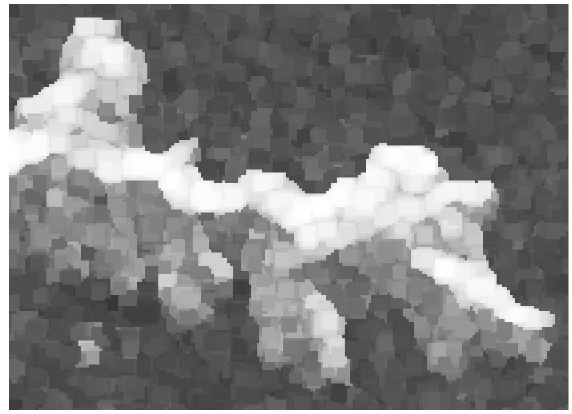

Figure 7. MCV filtering (5×5).

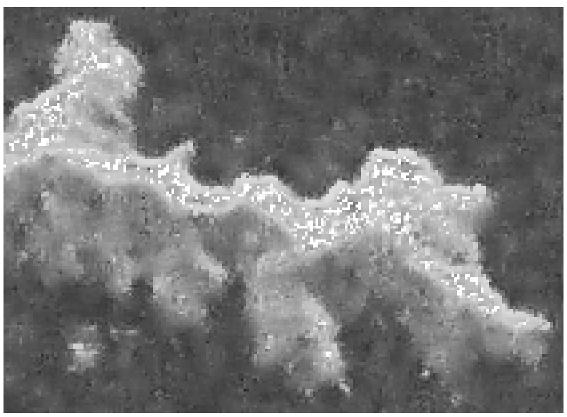

Figure 8. Local statistics filtering (7x7).

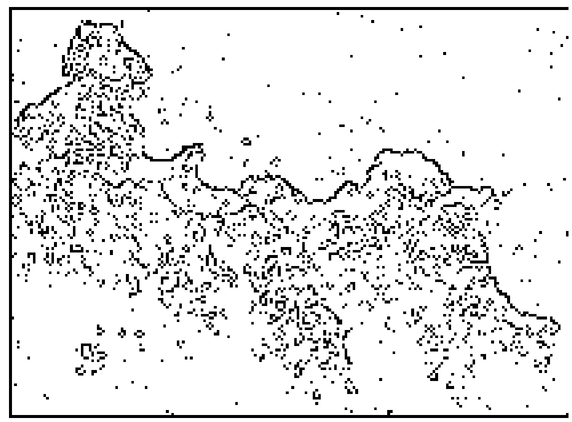

Figure 9. Edge detection of Figure 6.

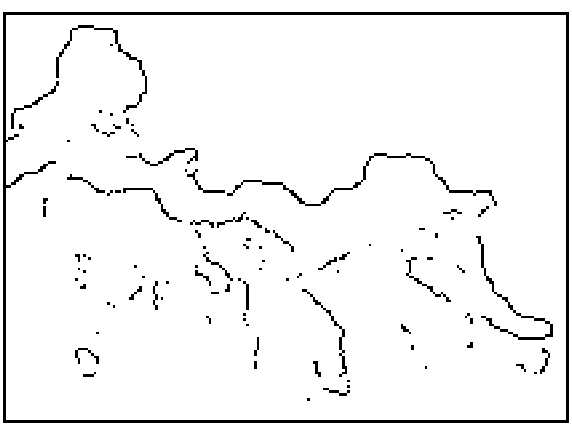

Figure 10. Edge detection of Figure 7.

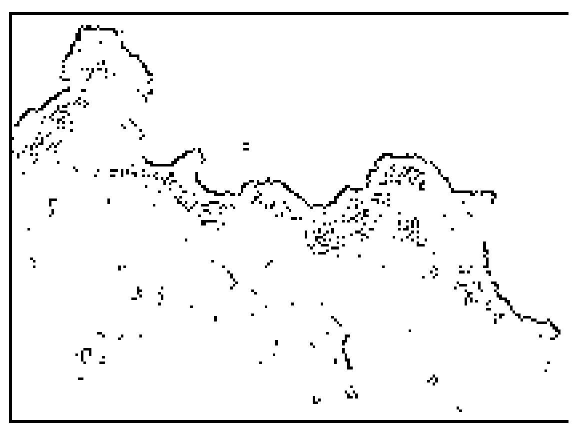

Figure 11. Edge detection of Figure 8.