Value-and-Criterion Filters: A new filter structure based upon morphological opening and closing

Mark A. Schulze and John A. Pearce

Department of Electrical and Computer Engineering and Biomedical Engineering Program
The University of Texas at Austin
Austin, Texas 78712

ABSTRACT

In this paper, we introduce the value-and-criterion filter structure and give an example of a filter with the structure. The value-and-criterion filter structure is based on morphological opening (or closing), which is actually two filters applied sequentially: the first assigns values based on the original image values, and the second assigns values based on the results of the first. The value-and-criterion structure is similar, but includes an additional step in parallel to the first that computes a different set of values to use as criteria for selecting a final value.

Value-and-criterion filters have a "value" function (V) and a "criterion" function (C), each operating separately on the original image, and a "selection" operator (S) acting on the output of C. The selection operator chooses a location from the output of C, and the output of V at that point is the output of the overall filter. The value-and-criterion structure allows the use of different linear and nonlinear elements in a single filter, but also provides the shape control of morphological filters.

An example of a value-and-criterion filter is the Mean of Least Variance (MLV) filter, which we define to have V = mean, C = variance, and S = minimum. The MLV filter resembles several earlier edge-preserving smoothing filters, but performs better and is more flexible and more efficient. The MLV filter smooths homogeneous regions and enhances edges, and is therefore useful in segmentation algorithms. We illustrate its response to various image features and compare it to the median filter on different biomedical images.

1. INTRODUCTION

The problem of smoothing noise without blurring edges is one of the oldest problems in image processing. Linear filters such as the average filter and other low-pass filters effectively smooth many types of noise, but also blur sharp edges in images. Nonlinear filters are often proposed as solutions to the problem of edge-preserving noise smoothing; however, they are not as easy to design and analyze as linear filters. One of the oldest and most common nonlinear filters for noise reduction in images is the median filter. Because of its widespread use and its relatively long history, the properties of the median filter are fairly well understood [1], [2], [3]. However, the response of the median filter to certain types of signals and noise is often not desirable; in particular, it is not effective in reducing high frequency periodic signals [3], [4]. Other nonlinear filters have been proposed to solve the problem of edge-preserving smoothing; many of these have a structure of small filter subwindows within an overall filter window. Typically, these filters find the most homogeneous subwindow and perform low-pass

Copyright 1993 Society of Photo-Optical Instrumentation Engineers.
This paper was published in E. R. Dougherty, J. Astola, and H. Longbotham, eds., *Nonlinear Image Processing IV, Proc. SPIE,* v. 1902, pp. 106-115, and is made available as an electronic reprint with permission of SPIE. One print or electronic copy may be made for personal use only. Systematic or multiple reproduction, distribution to multiple locations via electronic or other means, duplication of any material in this paper for a fee or for commercial purposes, or modification of the content of the paper are prohibited.

filtering (usually the arithmetic mean) on that subwindow to give the output for the overall window. A few examples of such filters are described in [5], [6], [7], [8]. In this paper, we develop a regular subwindow structure that can be used to design such nonlinear filters and we illustrate the response of a filter with this structure.

The idea of dividing a filter window into several smaller neighborhoods or subwindows is similar to the window structure that develops automatically in morphological opening or closing. Opening and closing are compound operators that are, respectively, morphological erosion followed by dilation, and dilation followed by erosion. Erosion is an operator that simply finds the minimum of the pixels in its "structuring element" (window), and dilation finds the maximum in its structuring element. The structuring elements of erosion and dilation compound when the two operations are performed consecutively, yielding an overall window for opening and closing which is almost twice as large. For example, if the structuring element of erosion and dilation is a 3x3 square, opening and closing draw their values from the 5x5 window formed by the union of translations of a 3x3 "subwindow" to each position of a 3x3 structuring element. There are then nine 3x3 subwindows in the 5x5 window, and these nine subwindows are all the possible different 3x3 subwindows in the window. This is illustrated in Figure 1 below. In most of the edge-preserving smoothing filter structures described previously, there are fewer subwindows, or the subwindows are irregularly shaped and/or differently sized.

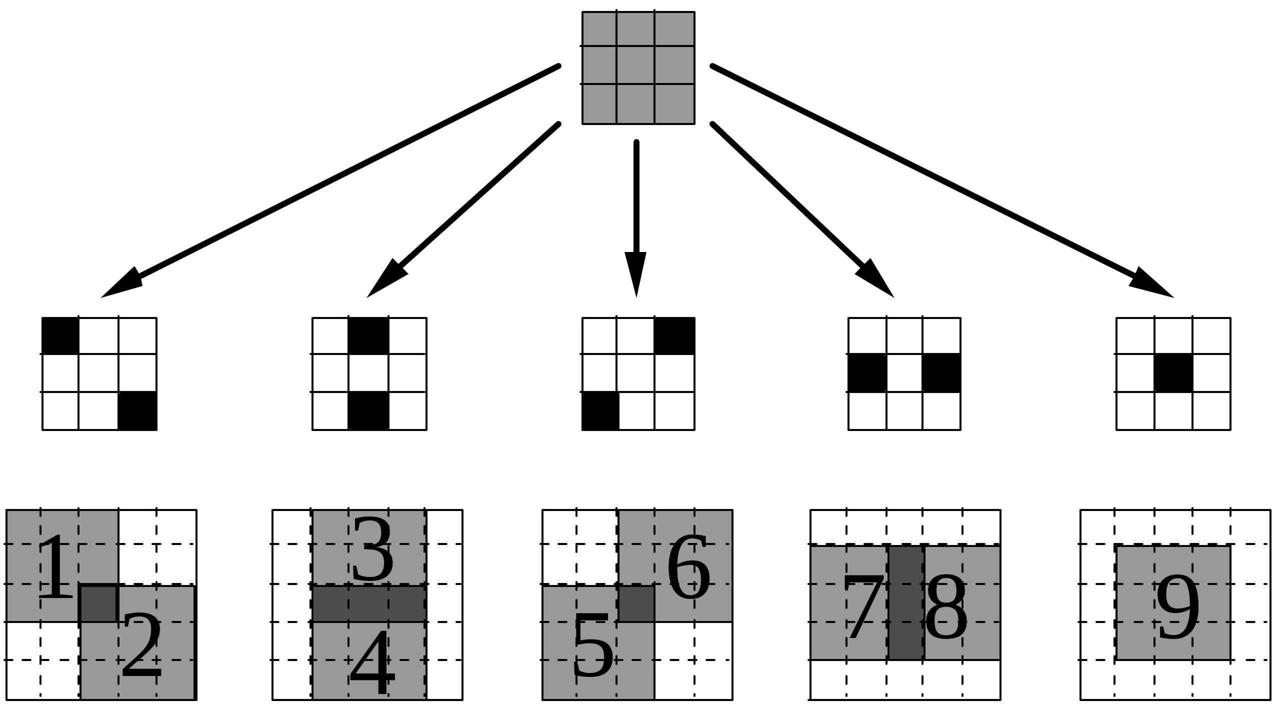

Figure 1. Structure of opening/closing with 3x3 structuring elements.

There are two key differences between the edge-preserving smoothing filter structures and the opening/closing structure. As previously noted, the opening/closing structure identifies *all* the subwindows of the same shape and size within an overall window (indeed, it defines the overall window by using the subwindows). The earlier smoothing filter structures are typically neither as complete nor as regular in their definitions of the subwindows. However, the edge-preserving smoothing filters first compute some measure of homogeneity on the subwindows to use as a criterion for selecting a subwindow, and then compute the output of the overall window using a different function on the original values in the subwindow. These two operations are essentially in parallel, because they both rely only on the original image values for their basic computations. Opening and closing, by contrast, result from consecutive filtering operations; that is, the values output by the first operation are used to compute the final output. Any information destroyed by the first operation of opening or closing cannot be recovered by the second operation.

The value-and-criterion filter structure uses the complete, regular subwindow structure of opening and closing, but changes the operations performed on those subwindows to be in parallel instead of in series. The output of one operation, the "criterion" function, is used to select one of the subwindows in the overall window (for example, the subwindow with the smallest sample variance). The output of another operation, the "value" function, on the selected subwindow is the final output of the filter. The value-and-criterion filter structure adds the regularity of morphological opening and closing to the earlier types of edge-preserving smoothing filters and yields a significant improvement in computational efficiency over the earlier filters. It also allows us to use any type of linear or nonlinear function in a filter while still retaining much of the shape control of the standard morphological operators.

2. THE VALUE-AND-CRITERION FILTER STRUCTURE

The value-and-criterion filter structure has two functions, the "value" function V and the criterion function C, which operate independently (and may operate in parallel) on the original image $f(\mathbf{x})$. Both V and C operate over the same window, or structuring element, N. Another operator, the "selection" function S, acts on the output of C. The selection operator uses a window, $\tilde{N}$, that is a 180° rotation of structuring element N. (Note that in morphological opening by N, erosion acts on N and the subsequent dilation acts on $\tilde{N}$.) Although S acts on the output of C, the information that we use from S is not a value, but rather the location $\mathbf{x}$ where C takes on the value selected by S. Letting $g(\mathbf{x})$ denote the filtered image and $v(\mathbf{x})$ and $c(\mathbf{x})$ denote the output of the value and criterion functions, respectively, the general value-and-criterion filter is defined by:

$$v(\mathbf{x}) = V\{f(\mathbf{x}); N\}$$
$$c(\mathbf{x}) = C\{f(\mathbf{x}); N\}$$
$$g(\mathbf{x}) = v\Big(\big\{\mathbf{x}' : \ \mathbf{x}' \in \tilde{N}_{\mathbf{x}}; c(\mathbf{x}') = S\{c(\mathbf{x}); \tilde{N}\} \big\}\Big)$$

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

Note that more than one value of $\mathbf{x}'$ is chosen by the selection operator if two or more values of the criterion function in the window are equally minimal. In this case, the above definition may give more than one value of $v(\mathbf{x})$ for output. Some method of deciding among these values is required. Two potential solutions are: (1) to average all the selected values of $v(\mathbf{x}')$ to yield the final output, and (2) to choose the value of $v(\mathbf{x}')$ closest to the value of $f(\mathbf{x})$ (that is, such that $|f(\mathbf{x})-v(\mathbf{x}')|$ is minimum) as the final output, settling ties consistently in favor of either the higher [$v(\mathbf{x}')>f(\mathbf{x})$] or lower [$v(\mathbf{x}')<f(\mathbf{x})$] value. In this paper, we use the second solution, and settle ties in favor of the higher value. This solution is less resistant to noise at $f(\mathbf{x})$ than the first solution, but is easier to implement and performs better at sharp, noiseless edges (where ties are likely). The first solution provides better noise reduction in some situations, but sometimes blurs sharp edges slightly.

Morphological opening and closing can be easily expressed in terms of the value-and-criterion structure. Opening results if both V and C are the minimum (or infimum) operator and S is the maximum (or supremum) operator. Closing is the case when both V and C are the maximum operator and S is the minimum operator. More interesting possibilities arise when V and C are not the same. For example, if V is the sample mean, C is the sample variance, and S is the minimum, the resulting filter outputs the mean of the neighborhood N in the composite window ($N \oplus N$, where the symbol $\oplus$ denotes morphological dilation) which has the smallest variance. (If N is a 3x3 square structuring element, then the composite window $N \oplus N$ is 5x5 square. See Figure 1.) We call this filter the Mean of Least Variance (MLV) filter. It will be discussed in detail in the next section.

The value-and-criterion filter structure is a powerful idea because it allows one to define the shape (structuring element) over which the filter will operate and to set a criterion with which to choose the region to filter. In edge-preserving noise smoothing, one of the most common techniques is to selectively smooth homogeneous regions, leaving regions with edges, lines, or other details unchanged. The value-and-criterion structure is well suited to designing filters to do this, and it suggests an efficient implementation that avoids the redundant computation of the earlier edge-preserving filters. The structure also outlines a potential parallel implementation of the filters.

3. THE MEAN OF LEAST VARIANCE (MLV) FILTER

The Mean of Least Variance (MLV) filter is a value-and-criterion filter that uses the sample variance as the criterion for selecting the most homogeneous neighborhood and the sample mean for computing a final output value from that neighborhood. Since the sample variance requires the sample mean for its computation, the V and C functions are not really independent in this case. The MLV filter is described by:

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

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

$$\mathrm{MLV}\{\mathrm{f}(\mathbf{x}); N\} = \mathrm{v}\Big(\big\{ \mathbf{x}' : \mathbf{x}' \in \tilde{N}_{\mathbf{x}};\, \mathrm{c}(\mathbf{x}') = \min\{ \mathrm{c}(\mathbf{y}) : \mathbf{y} \in \tilde{N}_{\mathbf{x}} \} \big\} \Big)$$

where $N_{\mathbf{x}}$ and $\tilde{N}_{\mathbf{x}}$ denote, respectively, the structuring elements N and $\tilde{N}$ centered at position $\mathbf{x}$. Note that the $1/|N|$ term in the definition of $\mathrm{c}(\mathbf{x})$ is a normalization constant and does not affect the selection operator (minimum); therefore, it may be removed without changing the operation of the filter. Also, recall that if several values of $\mathbf{x}'$ satisfy the selection criterion, we have chosen to select the value of $\mathrm{v}(\mathbf{x}')$ closest to $\mathrm{f}(\mathbf{x})$, settling ties in favor of higher values.

Since structuring elements containing part of an edge typically have a relatively large variance, the output of the MLV filter usually will not be an average of values across an edge. Ramp-like edges between homogeneous regions are sharpened by the MLV filter, since it averages the "flattest" neighborhood around intermediate points in the ramp. Therefore, the MLV filter is well suited to edge detection or segmentation applications, but is not a good choice when exact restoration of the original image from a noisy or distorted version is the goal. An example of the edge-enhancing characteristic of the MLV filter is shown in Figure 2 below. A ramp-like edge in a one-dimensional signal becomes much sharper after processing by an MLV filter with a 5-wide one-dimensional structuring element ($|N| = 5$).

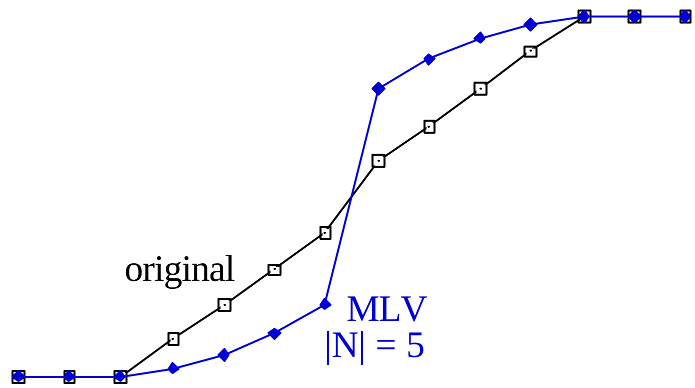

Figure 2. Example of edge-enhancing characteristic of the MLV filter in one dimension. Original signal (▫); MLV filtered with $|N| = 5$ (◆).

Finding the root signal set of the one-dimensional MLV filter also yields important insights into the behavior of the filter. A root signal of a filter is a signal that is unchanged by the operation of the filter on it. Root signals of the 1-D median filter with window width $|W| = 2|N|+1$ are such that each set of $|N|+2$ consecutive points in the signal is monotonic; that is between increasing and decreasing sections there must be a *constant neighborhood*, an area of at least $|N|+1$ consecutive points with the same value. Root signals of the 1-D average filter, on the other hand, are only constant signals or constant-slope infinite-length ramp signals. Any signal that is a root of the average filter must also be a root of the MLV filter, since it only alters values by averaging. However, the 1-D MLV filter is also able to preserve a step between long constant regions, since it will take the average of one of its constant subwindows and not of the subwindows with the step. The constant regions must be at least as long as one subwindow, $|N|$, and no other types of edges or ramps are left unchanged by the MLV filter. Therefore, the root signal set of the 1-D MLV filter consists of piecewise constant signals with constant regions at least $|N|$ points long, and constant-slope infinite-length ramp signals. This root signal set illustrates the edge-enhancing and noise-reducing properties of the MLV filter, since signals that do not have constant (homogeneous) regions separated by step edges are modified by the MLV filter. The choice of structuring element size is important for the behavior of the MLV filter at ramp edges, since with a small structuring element the MLV filter sometimes creates "stairsteps" (small flattened regions the same size as the structuring element separated by sharp steps) on long ramps.

In two dimensions, the shape of the structuring element plays an important role in the response of the MLV filter. Since the output of the filter at each point is the average over a structuring element in the image, features that are compatible with the structuring element are more readily preserved than those that are not.

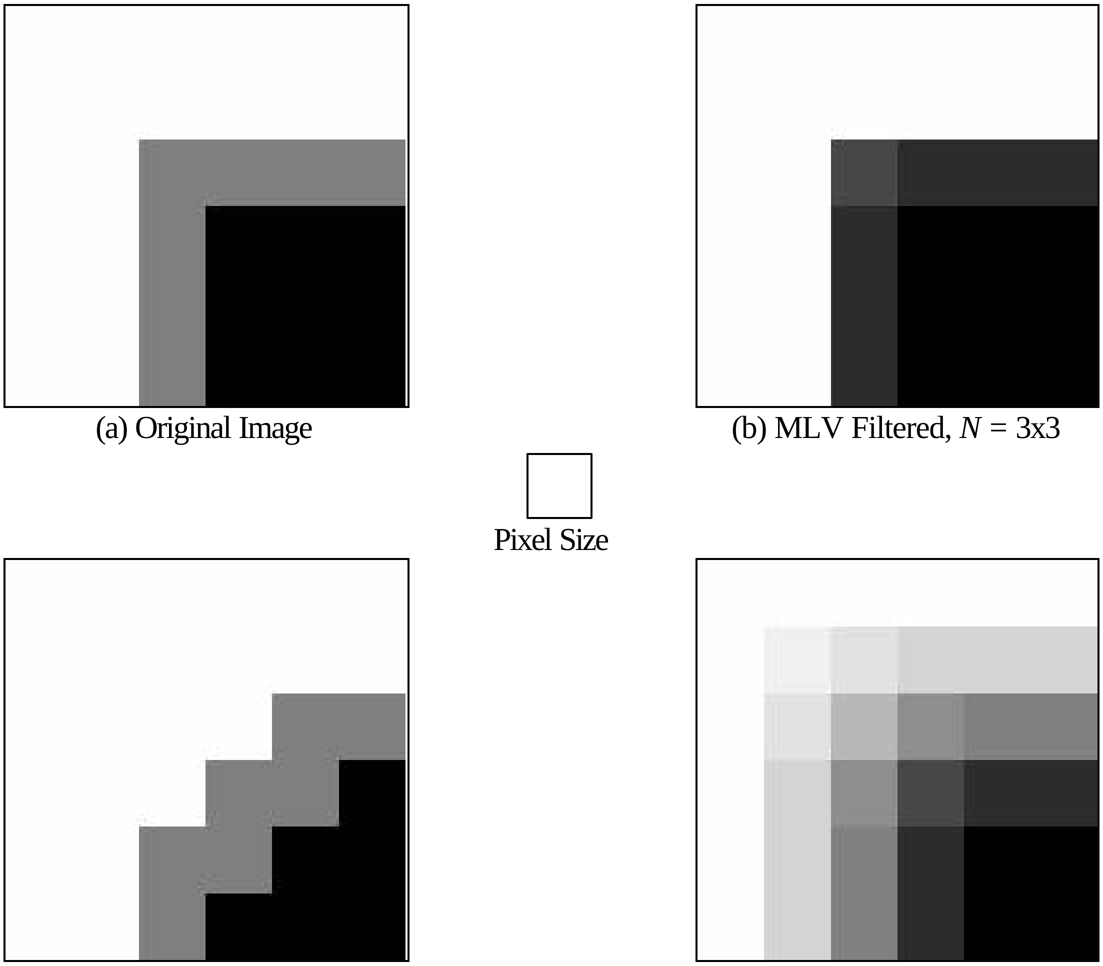

(a) Original Image

(b) MLV Filtered, N = 3x3

(c) 5x5 Median Filtered

(d) 3x3 Average Filtered

Figure 3. Response of MLV, median, and average filters to a 90° corner.
Each image section is 6x6 pixels square.

For example, if the structuring element is a square, one expects the MLV filter to preserve straight edges and 90° corners since these features are characteristic of a square. Figure 3 illustrates these properties of the MLV filter. Figure 3(a) is a 6x6 section of a 90° corner in an image. Figure 3(b) is the result of MLV filtering 3(a) with a 3x3 square structuring element. Figures 3(c) and 3(d) are, respectively, the results of median filtering with a 5x5 square window, and mean (or average) filtering with a 3x3 square window. This example demonstrates that the MLV filter with a 3x3 square structuring element enhances straight edges and 90° corners. The median filter preserves the straight edges, but rounds the corner off somewhat. The average filter blurs both edges and corners in an image.

Figure 4 is an example of the behavior of the MLV filter at a noisy edge. The sections shown in Figure 4 are 6x6 pixels square, and there are two intermediate values at the edge between constant regions on the left and right. The original image is shown in Figure 4(a). Uniformly distributed random noise with a range of one-fifth the difference between the constant regions is added to every pixel in the image. Figure 4(b) is the noisy edge. The results of MLV filtering with a 3x3 structuring element, 5x5 median filtering, and 3x3 average filtering the noisy image are shown in Figures 4(c), 4(d), and 4(e) respectively. Again, the edge-enhancing property of the MLV filter is observed. The median and average filters both blur the noisy edge, the median filter only slightly and the average filter quite noticeably. All three filters are effective at reducing the noise. The output of the median filter is visually most like the original, while the MLV-filtered result has a clearer transition between the two constant regions. Thus, the MLV filter is better suited to segmentation applications and the median filter to restoration applications.

One potential disadvantage of the MLV filter is that its value function is the average, which makes it somewhat susceptible to impulsive noise. Although the selection of the structuring element with the

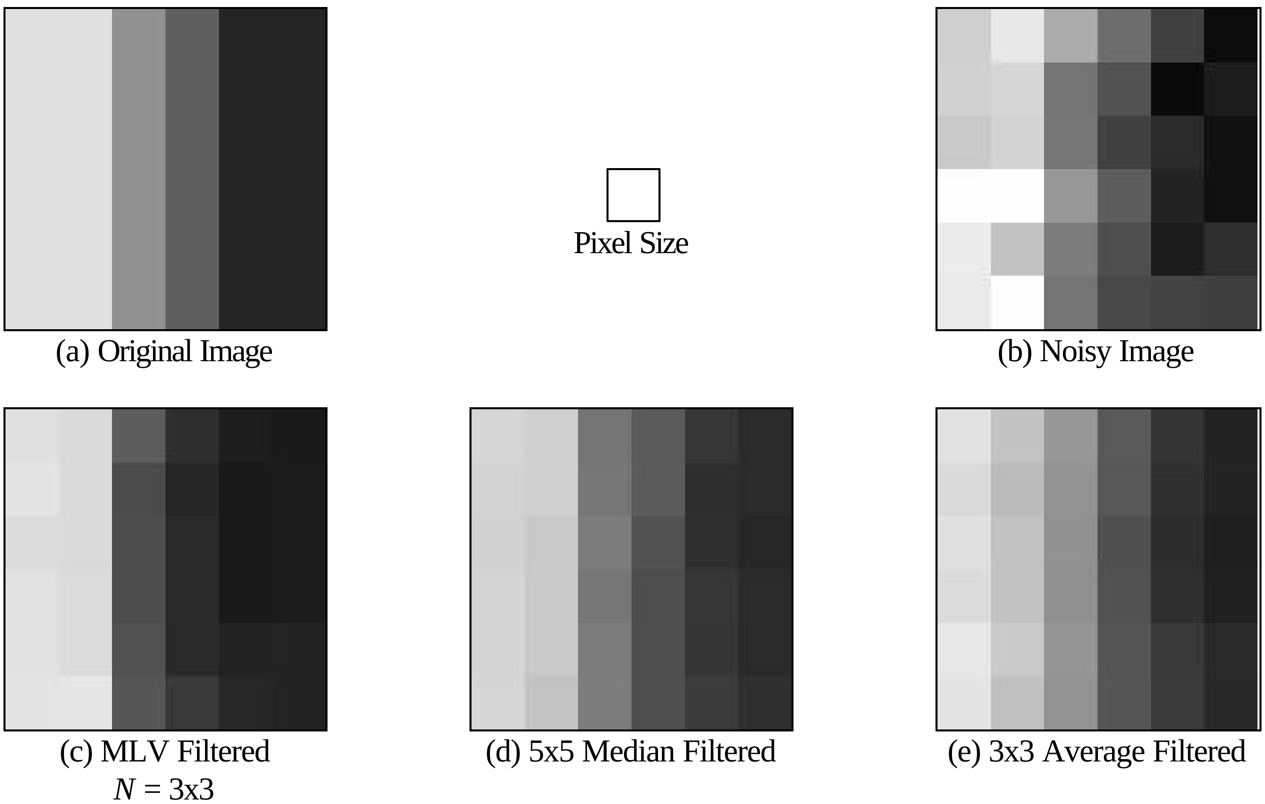

Figure 4. Response of MLV, median, and average filters to a noisy edge.
Noise is uniformly distributed, range = 1/5 range of original image.
Each image section is 6x6 pixels square.

smallest variance helps to minimize the effects of impulses, when the overall window is centered at an impulse, that impulse appears in every structuring element. This means that the output of the MLV filter will be slightly corrupted by the impulse. We can modify the MLV filter using the value-and-criterion structure to provide better resistance to impulsive noise by changing the value function from the sample mean to the median and making a corresponding change in the criterion function. (Use of the sample variance no longer makes sense, since it is the appropriate scale estimate only if the sample mean if the location estimate [9]. Changing the location estimate to the median requires a change in the scale estimate, which should be used for the criterion function.) Although such a filter may be more resistant to impulsive noise, we have found that the filters with the median for the value function do not perform as well as the MLV filter in most image processing situations. In addition, since the behavior of the median and its scale estimator are not as well understood as the behavior of the sample mean and variance, the properties of the median-based filters are not as readily described as those of the MLV filter.

6. APPLICATIONS OF THE MLV FILTER

As noted previously, the MLV filter seems well suited to applications in image segmentation algorithms. The tendency of the MLV filter to sharpen edges and smooth homogeneous regions is useful when segmenting images based on differences in gray levels when the images are noisy or otherwise degraded. Two examples using biomedical images are given in this section.

A transverse cross-sectional image of a human brain taken by magnetic resonance (MR) imaging is shown in Figure 5(a). MR images are typically segmented into regions of white matter, gray matter, ventricles, and other (presumably abnormal) tissue. The results of MLV filtering with a 3x3 square structuring element are shown in Figure 5(b). The boundaries between tissue types are more distinct than in the original, and noise in the image is reduced. Figure 5(c) is the result of 5x5 median filtering the original image. Noise is also reduced in this image, but the boundaries seem blurred.

Each iteration of the MLV filter brings the signal or image closer to a root of the filter; that is, the “flat” regions become increasingly homogeneous and the edges become sharper. This fact can be used to our advantage in segmentation, since we may classify those points that are obviously within a particular region, while delaying classification of points that lie near a threshold. Iterating the MLV filter then brings some of the previously indeterminate points closer to one of the segmentation classes, and they can then be classified. For MR images, better results are obtained by varying the threshold with position in the image, since the images are typically brighter in some areas than in others. For example, the lower left corner in Figure 5(a) is significantly brighter than the upper right corner. Using the MLV filter iteratively with variable thresholding, we have obtained satisfactory segmentations of magnetic resonance images.

Another segmentation problem we have considered is thermographic venography of the peripheral vascular system. A thermal camera is used to image arms or legs while the hands or feet are immersed in warm water. The warm blood returning from the extremities creates a temperature difference between the subcutaneous veins of the forearms or legs and the surrounding tissue. Veins that are near the surface of the skin are then visible in the thermal image. Thermographic venograms help in diagnosing peripheral vascular disorders [10].

An example of a thermogram of a forearm after immersing the hands in warm water for several minutes is shown in Figure 6(a). The results of MLV filtering with N = 3x3 square and of median filtering with a 5x5 square window are shown in Figures 6(b) and 6(c) respectively. Once again, notice how much sharper the edges are in the MLV-filtered image, and that the median-filtered image is somewhat blurry.

The MLV filter thus enhances the contrast between the veins and background and sharpens edges while reducing noise in thermographic venography. The median filter is effective in reducing noise, but does not result in edges as sharp as the MLV filter. As a pre-filter for automatic vessel tracing, the MLV filter is preferred over the median filter.

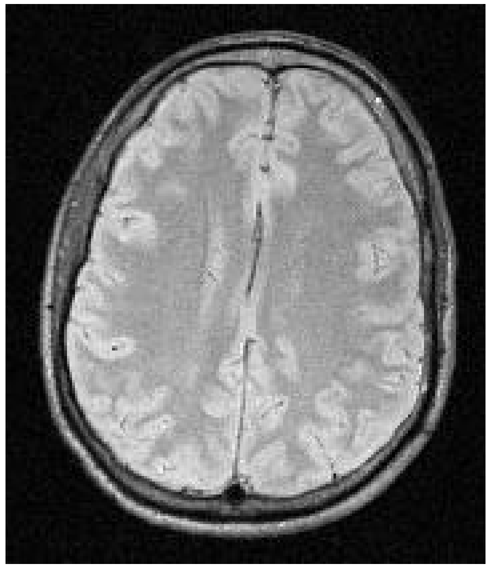

(a) Original MR image of human brain
(transverse section)

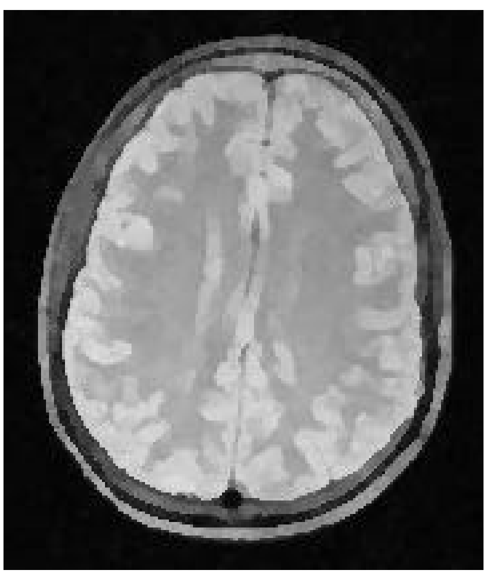

(b) MLV Filtered, N = 3x3

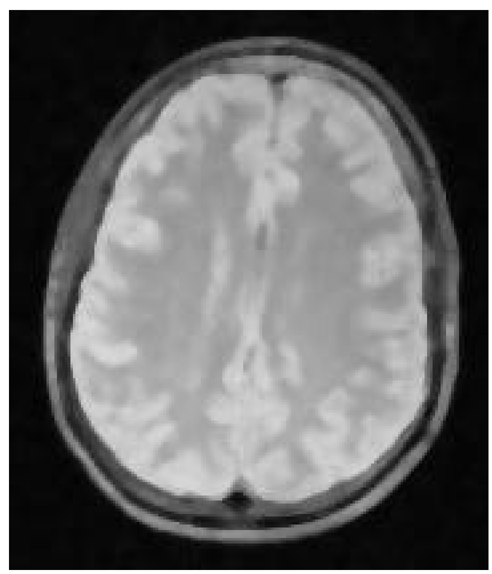

(c) 5x5 Median Filtered

Figure 5. Comparison of MLV and median filtering on a magnetic resonance image.

(a) Original thermographic image of forearms
fingers immersed in warm water

(b) MLV Filtered, N = 3x3

(c) 5x5 Median Filtered

Figure 6. Comparison of MLV and median filtering in thermographic venography.

7. CONCLUSIONS

The value-and-criterion filter structure is a new framework for designing filters with the shape control of mathematical morphology but with a wider variety of operators. The criterion function of a value-and-criterion filter determines which of the neighborhoods near a pixel will be used to determine the output. The value function determines what the output will be based on the pixels in the chosen neighborhood. This technique allows the value function to avoid operating on features such as edges if an appropriate criterion function is chosen. An example of a value-and-criterion filter is the MLV (Mean of Least Variance) filter, which takes the average of the neighborhood with the smallest sample variance as its output. We have shown that the MLV filter reduces noise, sharpens edges, and preserves shapes similar to those of its structuring element. The properties of the MLV filter are very desirable in segmentation algorithms, especially simple thresholding algorithms. In addition, the value-and-criterion structure illustrates an efficient, partially parallel implementation of the filters. The MLV filter executes more quickly than many of the older edge-preserving smoothing schemes with irregularly shaped and sized neighborhoods. The value-and-criterion filter structure thus is an efficient and complete scheme for designing filters that supports the development of interesting new filters such as the MLV filter.

8. ACKNOWLEDGMENTS

This material is based upon work supported under a National Science Foundation Graduate Fellowship and an AT&T Bell Laboratories Ph.D. Scholarship. M. A. Schulze is grateful for the support from both of these sources. The authors also thank Eric Chan for providing thermographic images for this research.

9. REFERENCES

[1] N. C. Gallagher, Jr. and G. L. Wise. "A theoretical analysis of the properties of the median filter." *IEEE Trans Acoust Speech Signal Process*, v. 29 n. 12 (1981) pp. 1136-1141.

[2] B. I. Justusson. "Median filtering: statistical properties." In T. S. Huang, ed., *Two-Dimensional Signal Processing II*. Berlin: Springer-Verlag, 1981, pp. 161-196.

[3] S. G. Tyan. "Median filtering: deterministic properties." In T. S. Huang, ed., *Two-Dimensional Signal Processing II*. Berlin: Springer-Verlag, 1981, pp. 197-217.

[4] J. Astola, P. Heinonen, and Y. Neuvo. "On root structures of median and median-type filters." *IEEE Trans Acoust Speech Signal Process*, v. 35 n. 8 (1987) pp. 1199-1201.

[5] J. Y. Hsiao and A. A. Sawchuk. "Supervised textured image segmentation using feature smoothing and probabilistic relaxation techniques." *IEEE Trans Patt Anal Mach Intell*, v. 11 n. 12 (1989) pp. 1279-1292.

[6] F. Tomita and S. Tsuji. "Extraction of multiple regions by smoothing in selected neighborhoods." *IEEE Trans Syst Man Cyber*, v. 7 (1977) pp. 107-109.

[7] M. Nagao and T. Matsuyama. "Edge preserving smoothing." *Comput Graphics Image Process*, v. 9 (1979) pp. 394-407.

[8] J.-S. Lee. "Refined filtering of image noise using local statistics." *Comput Graphics Image Process*, v. 15 (1981) pp. 380-389.

[9] I. Pitas and A. N Venetsanopoulos. *Nonlinear Digital Filters: Principles and Applications*. Boston: Kluwer Academic, 1990.

[10] E. K. Y. Chan. "Vessel extraction, visualization and tracing in themographic venography." Ph.D. Dissertation, University of Texas at Austin, August 1991.