

**Biomedical Image Processing with
Morphology-Based Nonlinear
Filters**

**Approved by
Dissertation Committee:**

John A. Pearce, Supervisor

Jonathan W. Valvano

Rebecca Richards-Kortum

Ronald E. Barr

J. K. Aggarwal

Copyright
by
Mark Allen Schulze
1994

**Biomedical Image Processing with
Morphology-Based Nonlinear
Filters**

by

Mark Allen Schulze, B.A., B.S.E.E., M.S.E.

Dissertation

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

Doctor of Philosophy

The University of Texas at Austin

December 1994

Acknowledgments

Many people have helped me during my graduate studies, and I would like to thank all of them for their patience and support. I am particularly grateful to my supervising professor, Dr. John Pearce, who has given me so much guidance and insight over the past five years. Everyone in the Biomedical Engineering program has helped me at some point along the way, and I am especially grateful to its director, Dr. Lee Baker, for his leadership and for many kindnesses over the years. My dissertation committee members and my colleagues in Dr. Pearce's laboratories have also been of great assistance during my research. I would like to thank Dr. Gilbert Hillman and the Department of Radiology at the University of Texas Medical Branch at Galveston for providing the magnetic resonance images for this work, and Ranjit Desai for providing the ultrasound images.

Thanks also to the National Science Foundation and AT&T Bell Laboratories, whose fellowship support for my graduate studies was of more than just monetary value. I am especially grateful to Dr. Thrasyvoulos Pappas for providing me with valuable experience and guidance during my summer research assignment at Bell Labs. Finally, I would like to thank all of my family and friends in Texas who have helped me enjoy my time in Austin, especially my parents and my brother Keith. Your love and support (and patience!) has made it all possible.

Biomedical Image Processing with Morphology-Based Nonlinear Filters

Publication No. _____

Mark Allen Schulze, Ph.D.

The University of Texas at Austin, 1994

Supervisor: John A. Pearce

Nonlinear filtering techniques are becoming increasingly important in image processing applications, and are often better than linear filters at removing noise without distorting image features. However, design and analysis of nonlinear filters are much more difficult than for linear filters. One structure for designing nonlinear filters is mathematical morphology, which creates filters based on shape and size characteristics. Morphological filters are limited to minimum and maximum operations that introduce bias into images. This precludes the use of morphological filters in applications where accurate estimation of the true gray level is necessary.

This work develops two new filtering structures based on mathematical morphology that overcome the limitations of morphological filters while retaining their emphasis on shape. The linear combinations of morphological filters eliminate the bias of the standard filters, while the value-and-criterion filters allow a variety of linear and nonlinear operations to be used in the geometric structure of morphology. One important value-and-criterion filter is the Mean of Least Variance (MLV) filter, which sharpens edges and provides noise smoothing equivalent to linear filtering.

To help understand the behavior of the new filters, the deterministic and statistical properties of the filters are derived and compared to the properties of the standard morphological filters. In addition, new analysis techniques for nonlinear filters are introduced that describe the behavior of filters in the presence of rapidly fluctuating signals, impulsive noise, and corners. The corner response analysis is especially informative because it quantifies the degree to which a filter preserves corners of all angles.

Examples of the new nonlinear filtering techniques are given for a variety of medical images, including thermographic, magnetic resonance, and ultrasound images. The results of the filter analyses are important in deciding which filter to use for a particular application. For thermography, accurate gray level estimation is required, so linear combinations of morphological operators are appropriate. In magnetic resonance imaging (MRI), noise reduction and contrast enhancement are desired. The MLV filter performs these tasks well on MR images. The new filters perform as well or better than previously established techniques for biomedical image enhancement in these applications.

Table of Contents

List of Figures.....	x
List of Tables	xviii
1. Introduction.....	1
1.1. Overview.....	1
1.2. Organization	3
2. Mathematical Morphology	8
2.1. Definitions	8
2.1.1. Basic Morphological Operators.....	8
2.1.2. Compound Morphological Operators	14
2.1.3. Examples.....	15
2.2. Deterministic Properties	21
2.2.1. Median Filter Root Signals.....	21
2.2.2. Relationship of Morphological and Median Filtered Signals...25	
2.3. Statistical Properties	28
2.3.1. One-Dimensional Properties.....	28
2.3.2. Two-Dimensional Properties	33
2.4. Summary.....	38
3. New Filter Structures.....	40
3.1. Introduction.....	40
3.2. Linear Combinations of Morphological Operators.....	40
3.2.1. Definitions	41
3.2.2. Deterministic Properties	43
3.2.3. Statistical Properties	64
3.3. Value-and-Criterion Filter Structure.....	73
3.3.1. Introduction.....	73
3.3.2. Definition.....	77
3.3.3. The Mean of Least Variance (MLV) Filter	79

3.3.4. Deterministic Properties	81
3.3.5. Statistical Properties	86
4. New Analysis Techniques	98
4.1. Introduction.....	98
4.2. Continuous Time Analysis	99
4.2.1. Introduction and Definitions.....	99
4.2.2. Results.....	102
4.3. Breakdown Point	110
4.3.1. Definition and Analysis	110
4.3.2. Order Statistic Limited (α -Limited) Filters	113
4.3.3. Conclusions.....	118
4.4. Corner Response Analysis.....	118
4.4.1. Definitions	119
4.4.2. Results.....	122
4.5. Conclusions.....	142
5. Applications.....	143
5.1. Introduction.....	143
5.2. Thermography.....	145
5.2.1. Thermal Image Noise Processes.....	145
5.2.2. Filter Selection.....	146
5.2.3. Examples.....	149
5.3. Magnetic Resonance Imaging (MRI)	166
5.3.1. Introduction.....	166
5.3.2. Anisotropic Diffusion	168
5.3.3. Noise Estimation.....	176
5.3.4. Results.....	179
5.3.5. Evaluating Filter Performance Using a Phantom Image	190
5.3.6. Conclusions.....	199
5.4. Ultrasound.....	200
5.5. Conclusions.....	209

6. Conclusions.....	211
Appendix: Geometric Derivations of Corner Responses.....	214
A.1. Total Area of Corner.....	214
A.2. Corner Response of Morphological Opening	216
A.3. Corner Response of Erosion and the Midrange Filter	218
A.4. Corner Response of the Square Median Filter.....	222
A.5. Corner Response of the Averaging Filter	224
A.6. Corner Response of the Plus-Shaped Median Filter.....	230
Bibliography	233
Vita	239

List of Figures

Figure 2.1. (a) Erosion example. (b) Dilation example.	9
Figure 2.2. Example of threshold decomposition of a function into cross sections.	11
Figure 2.3. Illustration of an “effective” structuring element of opening and closing formed by rotation between successive erosions and dilations.....	15
Figure 2.4. Effect of erosion and opening on a 1-D signal.....	16
Figure 2.5. Effect of dilation and closing on a 1-D signal.....	17
Figure 2.6. Effect of open-closing (OC) on a 1-D signal.	17
Figure 2.7. Effect of close-opening (CO) on a 1-D signal.....	18
Figure 2.8. Original binary image.....	20
Figure 2.9. Image after erosion.....	20
Figure 2.10. Image after dilation.	20
Figure 2.11. Image after opening.....	20
Figure 2.12. Image after closing.	20
Figure 2.13. Image after OC.	20
Figure 2.14. Image after CO.	20
Figure 2.15. Oscillatory signal filtered by a 5-wide non-recursive median filter.....	23
Figure 2.16. Oscillatory signal filtered by a 5-wide recursive median filter.....	24
Figure 2.17. Density function of CO output with uniformly distributed input.....	31
Figure 2.18. Density function of OC output with uniformly distributed input.....	31

Figure 2.19. Density function of closing output with uniformly distributed input.	31
Figure 2.20. Density function of opening output with uniformly distributed input.	31
Figure 2.21. Density function of median filter output with uniformly distributed input.	32
Figure 2.22. Density functions of 1-D and 2-D closing output with uniformly distributed input.	36
Figure 2.23. Density functions of 1-D and 2-D closing output with uniformly distributed input.	37
Figure 3.1. Example of a signal that has a dilation with no constant neighborhoods.....	49
Figure 3.2. Illustration for Case III of the proof of Theorem 3.14.	55
Figure 3.3. Illustration for Case III of the proof of Theorem 3.15.	59
Figure 3.4. Density function of midrange filter output with uniformly distributed input.	65
Figure 3.5. Theoretical and simulated results for the midrange filter acting on uniformly distributed input.	67
Figure 3.6. Theoretical and simulated results for the pseudomedian filter acting on uniformly distributed input.	70
Figure 3.7. Theoretical and simulated results for the pseudomedian filter acting on normally distributed input.....	70
Figure 3.8. Theoretical and simulated results for the LOCO filter acting on uniformly distributed input.	72
Figure 3.9. Theoretical and simulated results for the LOCO filter acting on normally distributed input.....	73

Figure 3.10. Structure of morphological opening and closing (and value-and-criterion structure) with a 3×3 structuring element.	75
Figure 3.11. Example of the edge-enhancing characteristic of the MLV filter in one dimension.....	80
Figure 3.12. Example section of piecewise constant root signal of the MLV filter with a structuring element of width 3.	82
Figure 3.13. Example of oscillatory root of MLV filter with $ N = 3$ under previously-described tie-breaking scheme.	83
Figure 3.14. Response of MLV, median, and average filters to a 90° corner.	85
Figure 3.15. Theoretical and simulated results for the averaging filter acting on normally distributed input.....	88
Figure 3.16. Data for slope function $v(m)$, curve fit of the form cm^k , and the approximation $v(m) \approx \frac{4}{5}(1 + \sqrt{m})^{-1}$	92
Figure 3.17. Theoretical and simulated results for the 3-, 5-, and 7-wide MLV filter acting on a normally distributed input.	93
Figure 3.18. Theoretical and simulated results for the 25- and 51-wide MLV filter acting on normally distributed input.	94
Figure 4.1. Continuous time response of midrange, pseudomedian, and LOCO filters to sinusoidal input signals.	106
Figure 4.2. Continuous time response of averaging, MLV and median filters to sinusoidal input signals.....	107
Figure 4.3. Continuous time response of midrange, pseudomedian, and LOCO filters to square wave input signals.....	108
Figure 4.4. Continuous time response of averaging, MLV and median filters to square wave input signals.	109
Figure 4.5. Noisy input signal.....	116

Figure 4.6. Output of LOCO filter and λ -limited LOCO filter with 25-wide structuring element and $\lambda = 0.2$	117
Figure 4.7. Original image of corner of angle θ with rotation α	120
Figure 4.8. Example filtered image of corner.....	121
Figure 4.9. Fractional corner preservation $r(\theta, 0)$ of morphological opening, OC, CO, and the LOCO filter compared to median and pseudomedian filters.....	128
Figure 4.10. Corner attenuation $s(\theta, 0)$ of morphological opening, OC, CO, and the LOCO filter compared to median and pseudomedian filters.	130
Figure 4.11. Fractional corner preservation $r(\theta, 0)$ of the plus-shaped median filter compared to the square-shaped median filter.	131
Figure 4.12. Corner attenuation $s(\theta, 0)$ of the plus-shaped median filter compared to the square-shaped median filter.	132
Figure 4.13. Fractional corner preservation $r(\theta, 0)$ of the LOCO filter with a circular structuring element compared to the LOCO filter with a square structuring element.....	134
Figure 4.14. Comparison of fractional corner preservation for the square median filter at corner rotations $\theta = 0$ and $\theta = 45^\circ$	135
Figure 4.15. Comparison of fractional corner preservation for the square LOCO filter at corner rotations $\theta = 0$ and $\theta = 45^\circ$	136
Figure 4.16. Comparison of fractional corner preservation for the plus-shaped median filter at corner rotations $\theta = 0$ and $\theta = 45^\circ$	137
Figure 4.17. Comparison of fractional preservations of the square median filter determined with 63×63 and 5×5 square windows.	138

Figure 4.18. Comparison of fractional corner preservation of square-shaped median, LOCO, and averaging filters and the plus-shaped median filter at corner rotations $\theta = 0$	139
Figure 4.19. Comparison of fractional corner preservation of square-shaped median, LOCO, and averaging filters and the plus-shaped median filter at corner rotations $\theta = 45^\circ$	140
Figure 5.1. “Round” structuring element (5×5).....	150
Figure 5.2. Weighted 5×5 “round” structuring element.	151
Figure 5.3. Original image of laser heated porcine skin.....	153
Figure 5.4. Image filtered by 5-wide horizontal median filter and 3-wide vertical triangular filter.	153
Figure 5.5. Image filtered by 3×3 square LOCO filter.	153
Figure 5.6. Image filtered by 5×5 “round” LOCO filter.....	153
Figure 5.7. Image filtered by weighted 5×5 “round” LOCO filter.....	153
Figure 5.8. Image filtered by 5×5 square median filter.	153
Figure 5.9. Image filtered by 5-wide horizontal trimmed mean filter ($\alpha = 0.2$) and 3-wide vertical triangular filter.	154
Figure 5.10. Image filtered by 5-wide horizontal mean filter and 3-wide vertical triangular filter.	154
Figure 5.11. Phantom image (Gaussian profile).....	160
Figure 5.12. Noisy phantom image.....	160
Figure 5.13. Phantom filtered by 5-wide horizontal median filter and 3-wide vertical triangular filter.	160
Figure 5.14. Phantom filtered by 3×3 square LOCO filter.....	160
Figure 5.15. Phantom filtered by 5×5 “round” LOCO filter.	160
Figure 5.16. Phantom filtered by weighted 5×5 “round” LOCO filter.....	161

Figure 5.17. Phantom filtered by 5×5 square median filter.....	161
Figure 5.18. Phantom filtered by 5-wide horizontal trimmed mean filter ($\alpha=0.2$) and 3-wide vertical triangular filter.....	161
Figure 5.19. Phantom filtered by 5-wide horizontal mean filter and 3-wide vertical triangular filter.	161
Figure 5.20. Anisotropic diffusion of a signal with λ ranging from 1 to 6.	170
Figure 5.21. MLV filtering of a signal with $ N $ ranging from 3 to 11.	172
Figure 5.22 Comparison of anisotropic diffusion and MLV filtering for a two- dimensional rectangular pulse.	175
Figure 5.23. Scaled original MR image of human head (coronal section).	178
Figure 5.24. MR image with additive Gaussian noise.....	178
Figure 5.25. Original MR image of human head (coronal section).....	182
Figure 5.26. MLV filtered MR image ($N = 3 \times 3$).....	182
Figure 5.27. MLV filtered MR image ($N = 5 \times 5$).....	183
Figure 5.28. MLV filtered MR image ($N = 9 \times 9$).....	183
Figure 5.29. Anisotropic diffusion of MR image (3 iterations with $\lambda = 3.0$).	184
Figure 5.30. Anisotropic diffusion of MR image (3 iterations with $\lambda = 20.0$). ..	184
Figure 5.31. Anisotropic diffusion of MR image (15 iterations with $\lambda = 3.0$). ..	185
Figure 5.32. Anisotropic diffusion of MR image (15 iterations with $\lambda = 20$). ...	185
Figure 5.33. Median filtered MR image (5×5 square filter window).	186
Figure 5.34. Diagram of MRI head phantom.	193
Figure 5.35. Image of MRI head phantom.	193
Figure 5.36. Head phantom with additive Gaussian noise.	194
Figure 5.37. MLV filtered head phantom ($N = 3 \times 3$).	194

Figure 5.38. MLV filtered head phantom ($N = 9 \times 9$).	195
Figure 5.39. Anisotropic diffusion of head phantom (15 iterations with = 20).....	195
Figure 5.40. Anisotropic diffusion of head phantom (3 iterations with = 30).....	196
Figure 5.41. Anisotropic diffusion of head phantom (3 iterations with = 100).....	196
Figure 5.42. Original echocardiogram.....	206
Figure 5.43. Closing (3×3) of original image.	206
Figure 5.44. Median filtering (3×3) of closed image.....	206
Figure 5.45. Threshold of original echocardiogram.	207
Figure 5.46. Threshold of Fig. 5.43.	207
Figure 5.47. Threshold of Fig. 5.44	207
Figure 5.48. LOCO filtering (plus-shaped 3×3) of original image.....	208
Figure 5.49. MLV filtering (3×3) of original image.....	208
Figure 5.50. Threshold of Fig. 5.48.	208
Figure 5.51. Threshold of Fig. 5.49	208
Figure 5.52. Anisotropic diffusion (15 iterations, = 15) of original image.	209
Figure 5.53. Threshold of Fig. 5.52.	209
Figure A.1. Derivation of total area of corners with angles in the range 0	215
Figure A.2. Derivation of fractional corner preservation for morphological opening with a square structuring element.	217
Figure A.3. Derivation of fractional corner preservation for morphological erosion with a square structuring element (small angles).....	220

Figure A.4. Derivation of fractional corner preservation for morphological erosion with a square structuring element (larger angles).	221
Figure A.5. Derivation of fractional corner preservation for the square median filter (small angles).	223
Figure A.6. Derivation of fractional corner preservation for the square median filter for a right angle corner.....	224
Figure A.7. Derivation of fractional corner preservation for the square averaging filter (small angles).	227
Figure A.8. Derivation of fractional corner preservation for the square averaging filter for a right angle corner.....	229
Figure A.9. Derivation of fractional corner preservation for the plus-shaped median filter (small angles).	231
Figure A.10. Derivation of fractional corner preservation for circle-shaped morphological opening (small angles).....	232

List of Tables

Table 2.1. Median of filter output distributions with uniformly distributed input.	32
Table 3.1. Theoretical and simulated variance ratios for the averaging filter for odd window sizes 3 to 11 and 25, 51, and 101.	89
Table 3.2. Simulated and predicted variance ratios for the MLV filter for structuring elements of various sizes from 3 to 251.	95
Table 4.1. Summary of breakdown points of filters.	113
Table 4.2. Comparison of filter corner responses.....	140
Table 5.1. Mean square error (MSE) of filtered noisy signals.	147
Table 5.2. Mean percentage reduction of standard deviation of isothermal contour radius for various thermal image filtering techniques.....	156
Table 5.3. Convexity measure of isothermal contours for various thermal image filtering techniques.....	158
Table 5.4. Mean percentage increase in standard deviation of “isothermal” contour radius and mean convexity measure for various filtering techniques on the thermal image phantom.	163
Table 5.5. Mean percentage of false positive and negative identifications of points inside “isothermal” contour radius and mean convexity measure for various filtering techniques on the thermal image phantom.	165
Table 5.6. Noise estimation in filtered MR images.	187
Table 5.7. Image quality estimation in filtered MR images.	189
Table 5.8. Component ellipses of MRI head phantom.	192

Table 5.9. Percentage of points misclassified by various filters acting on phantom image.....	198
--	-----

1. Introduction

1.1. OVERVIEW

Nonlinear methods in signal and image processing have become increasingly popular over the past thirty years. There are two general families of nonlinear filters: the homomorphic and polynomial filters, and the order statistic and morphological filters [1]. Homomorphic filters were developed during the 1970's and obey a generalization of the superposition principle [2]. The polynomial filters are based on traditional nonlinear system theory and use Volterra series. Analysis and design of homomorphic and polynomial filters resemble traditional methods used for linear systems and filters in many ways. The order statistic and morphological filters, on the other hand, cannot be analyzed efficiently using generalizations of linear techniques. The median filter is an example of an order statistic filter, and is probably the oldest [3, 4] and most widely used order statistic filter. Morphological filters are based on a form of set algebra known as mathematical morphology. Most morphological filters use extreme order statistics (minimum and maximum values) within a filter window, so they are closely related to order statistic filters [5, 6].

While homomorphic and polynomial filters are designed and analyzed by the techniques used to define them, order statistic filters are often chosen by more heuristic methods. As a result, the behavior of the median filter and other related filters was poorly understood for many years. In the early 1980's, important results on the statistical behavior of the median filter were presented [7], and a

new technique was developed that defined the class of signals invariant to median filtering, the root signals [8, 9]. Morphological filters are derived from a more rigorous mathematical background [10-12], which provides an excellent basis for design but few tools for analysis. Statistical and deterministic analyses for the basic morphological filters were not published until 1987 [5, 6, 13]. The understanding of the filters' behavior achieved by these analyses is not complete, however, so further study may help determine when morphological filters are best applied.

This dissertation investigates the use of morphology-based nonlinear filters to enhance biomedical images. Specifically, new filters based on mathematical morphology are developed, analyzed, and applied to a variety of medical images. The behavior of the standard morphological filters is undesirable for certain applications, and the new filters are designed to overcome these weaknesses. Some new analysis techniques are introduced, including a method to quantify the response of filters to two-dimensional features. These new analysis methods and the basic statistical and deterministic analyses are used to compare the new filters with the standard filters. Finally, the new nonlinear filters are used to enhance magnetic resonance, thermographic, and ultrasound images and their performance is compared to established filtering techniques for each of the imaging modalities.

The accomplishments of this work include:

- Demonstrating the statistical and deterministic bias introduced by the standard morphological filters.

- Developing the linear combinations of morphological filters as a class of filters that overcome the bias problems of the standard morphological filters.
- Developing a new filter structure based on mathematical morphology that allows the use of both linear and nonlinear operations and that unifies a wide variety of filters into a single mathematical formalism.
- Performing deterministic and statistical analyses of the new nonlinear filters to illustrate their behavior.
- Developing new analysis techniques that assist in understanding the behavior of nonlinear filters, including techniques to determine responses to rapidly fluctuating signals, impulsive noise, and two-dimensional corners.
- Illustrating the use of linear combinations of morphological filters in noise reduction and isothermal contour estimation for thermography.
- Illustrating the use of a new nonlinear filter for contrast enhancement in magnetic resonance imaging.

1.2. ORGANIZATION

This dissertation begins with a review of mathematical morphology, including the statistical and deterministic properties of the morphological filters. These properties point out weaknesses (specifically, a bias problem) in the behavior of the standard morphological filters that motivate the development of new filters. Next, new filters that address the bias problem of the standard filters are introduced. Linear combinations of morphological operators are one of the new types of filters. This work develops the deterministic and statistical

properties of these filters and illustrates the potential advantages of these filters over the standard morphological filters.

Another new type of filter introduced in this work is the value-and-criterion filter. This filter structure uses the shape-based organization of morphology, but expands the operations used for the filtering beyond just the maximum and minimum operators. Thus, any linear or nonlinear function can be used to determine the output value from values in a window, and to determine which window to use to get the output value. A promising application of this new structure is for designing filters that sharpen edges and smooth noise simultaneously. One of these new filters is the “Mean of Least Variance” filter, or MLV filter, which is a significant improvement over previously defined edge-preserving smoothing filters. The deterministic and statistical properties of the MLV filter are also investigated to contrast its behavior with other morphology-based filters.

Since the usual statistical and deterministic analyses provide only an incomplete understanding of the behavior of nonlinear filters, new analysis methods are introduced here to gain further insight into the response of the filters. A technique to quantify the response of filters to periodic signals of various frequencies is outlined. This method is similar to Fourier analysis for linear filters, but is much more limited in scope because of the nonlinear nature of the filters examined. Nonetheless, this analysis gives valuable clues about the response of nonlinear filters to rapidly fluctuating signals. Another important property of many nonlinear filters is their resistance to outlying values and impulsive noise. The “breakdown point” is a measure of the robustness of filters

in the presence of outliers. This method is another way to help explain differences among filters.

The last analysis method developed in this dissertation furthers the understanding of the behavior of filters at two-dimensional structures. This technique, called “corner response analysis,” quantifies the percentage of binary corners of various angles that is preserved by a filter. By plotting this information in polar format, the change in the response of a filter to corners of various angles is easily visualized. This method is a major improvement over previous analyses that focused on general characteristics like noise reduction or one-dimensional characteristics like edge preservation. The response of the filter to different rotations of the same feature is also explored using corner response analysis, indicating whether a filter acts similarly to different rotations of 2-D objects.

The final portion of this work illustrates the use of the new nonlinear filters in biomedical image processing applications. The results for the various filters yield important information for selecting the proper filter for a given application. Among the considerations for selecting a filter are the signal and noise characteristics of the specific imaging modality and the type of information that is to be extracted from the data. The imaging modalities considered (thermography, magnetic resonance, and ultrasound) have a variety of different characteristics that call for different filters. The theoretical analyses in the earlier sections provide a solid basis for selecting appropriate filters for each modality.

Thermograms are very noisy, and often accurate temperatures need to be estimated from them. The goal in thermographic imaging, then, is to remove the noise without introducing any statistical bias that would affect the accuracy of

temperature readings inferred from the images. The linear combinations of morphological operators match this description and provide more control over shape than previous thermographic filtering techniques. Examples of the new filtering technique are given and compared to established filtering methods for thermography.

In magnetic resonance imaging (MRI), image processing problems of clinical interest include improving the contrast and reducing the noise in images, and segmenting images into regions corresponding to different tissue types. The gray levels in an MR image do not correspond to properties that need to be measured quantitatively, so filters that bias or otherwise alter the gray levels may be used. The MLV filter, one of the new value-and-criterion filters, provides excellent contrast enhancement in MRI by sharpening edges between homogeneous tissue regions and simultaneously smoothing the noise within these regions. The results of a single pass of the MLV filter compare favorably to many iterations of another emerging technique for MRI enhancement, anisotropic diffusion [14]. Examples of both techniques are shown and the noise levels for the filtered images are estimated and compared. The edge enhancing and noise smoothing properties of the MLV filter make it an excellent choice as a pre-processor for segmentation algorithms and contrast improvement schemes in MRI analysis.

In ultrasound images, preserving structures (edges and shapes) in the image is more important than determining accurate gray level values. For example, one might only wish to extract a region of interest from a particular ultrasound image to use for a three-dimensional reconstruction. This is a case

where the standard, biased morphological operators are expected to work well. This work shows that the standard morphological filters are more appropriate for ultrasound image processing than any of the newly defined unbiased filters. Examples of filtering by both the biased and unbiased filters are given and compared.

The new nonlinear filter structures and analysis techniques introduced here are useful in biomedical image processing applications because they expand the variety of available filters and tools. The characteristics of the imaging modality can be then used to help select a filter that achieves the desired results. These new filters and analysis tools improve the chances of finding a suitable filter for almost any image processing application. The three biomedical applications investigated here give examples of situations where the new filters are useful and situations where other techniques are preferred.

2. Mathematical Morphology

2.1. DEFINITIONS

Mathematical morphology is a set algebra used to process and analyze data based on geometric shapes. The theory of mathematical morphology was introduced by Matheron [10] in 1974 and refined by Serra [11, 12] in the 1980's. The basic morphological operations are erosion and dilation. For binary signals, erosion is a Minkowski set subtraction (an intersection of set translations), and dilation is a Minkowski set addition (a union of set translations). These operators were extended to operate on non-binary signals by Serra [11] and others [5, 15, 16]. There are two main types of morphological filters [5]: set processing and function processing filters. Set processing filters accept binary input signals and give binary output signals, while function processing filters accept binary or non-binary functions as input and yield non-binary functions as output. The interpretation of binary signals as sets and non-binary signals as functions is straightforward; more details are given in [5].

2.1.1. Basic Morphological Operators

Set Processing Operations

Let X denote an m -dimensional set and N denote a compact k -dimensional set ($k \leq m$), and let \mathbf{y} denote a point in \mathbb{R}^k and \mathbf{z} a point in \mathbb{R}^m . The set X is a binary signal or image to be filtered, and the set N is called the structuring element of the morphological filter. Define the symmetric set $\tilde{N} = \{-\mathbf{z} : \mathbf{z} \in N\}$, which is a reflection of N about the origin. The translation of a set to a point \mathbf{z} is

denoted by a subscript; for example, the set N translated to the point \mathbf{z} is $N_{\mathbf{z}}$. The set processing morphological erosion and dilation are defined by:

$$\text{Erosion: } X \mathbf{K} \tilde{N} = \{\mathbf{z}: N_{\mathbf{z}} \subseteq X\} = \bigcap_{\mathbf{y} \in N} X_{-\mathbf{y}} \quad (2.1)$$

$$\text{Dilation: } X \mathbf{D} \tilde{N} = \{\mathbf{z}: (N_{\mathbf{z}} \cap X) \neq \emptyset\} = \bigcup_{\mathbf{y} \in N} X_{\mathbf{y}} \quad (2.2)$$

The symbols \mathbf{K} and \mathbf{D} denote Minkowski subtraction and Minkowski addition, respectively [11, 12]. The erosion of a set X is then the set of points to which the structuring element N may be translated while remaining entirely within the original set X . The dilation of X is the set of points to which N may be translated and still intersect X with at least one point. Examples of erosion and dilation of some simple discrete sets are shown in Figure 2.1 below. Clearly, erosion shrinks a set while dilation expands a set.

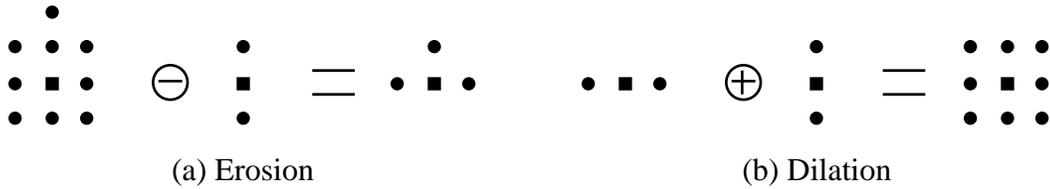


Figure 2.1. (a) Erosion example. (b) Dilation example. (Adapted from [5].) Squares (■) denote origin of the plane; circles (●) denote other points in the set.

In most applications of mathematical morphology, the structuring elements are symmetric about the origin, so that $N = \tilde{N}$. When this is the case, there is no distinction between erosion and Minkowski subtraction nor between dilation and Minkowski addition. Erosion and dilation also are duals of each other with respect to set complementation. If a superscript c denotes set complementation, then $X \mathbf{D} \tilde{N} = (X^c \mathbf{K} \tilde{N})^c$ and $X \mathbf{K} \tilde{N} = (X^c \mathbf{D} \tilde{N})^c$.

Function Processing Operations

Since most signal and image processing applications do not deal with binary data, mathematical morphology must extend to non-binary signals (functions) to be widely useful. This extension is performed by representing a function as an ordered set of binary signals [5, 11]. The cross section of a function at a particular level is a binary set, and the set of all such cross sections forms a complete representation of the function. This process of reducing a function to a set of binary signals is called threshold decomposition. The only restriction on threshold decomposition is that the function must be upper semicontinuous, which means that each cross section of the function must be a closed set. This is not a problem in most applications because all sampled functions are upper semicontinuous [5]. Sternberg [15, 16] used another technique to extend morphology to functions, but the resulting function processing filters are identical to those derived from threshold decomposition.

Threshold decomposition is illustrated in Figure 2.2, which shows three cross sections of a one-dimensional function $f(\mathbf{x})$. Given all the cross sections $X(f, t)$ of a function f , the original signal $f(\mathbf{x})$ may be uniquely reconstructed by simply “stacking” the cross sections. The value of the original signal at a location \mathbf{x} is equal to the highest value of t for which the location \mathbf{x} is included in the set $X(f, t)$. For a quantized signal, there are a finite number of levels (t) where the cross sections of the signal are taken.

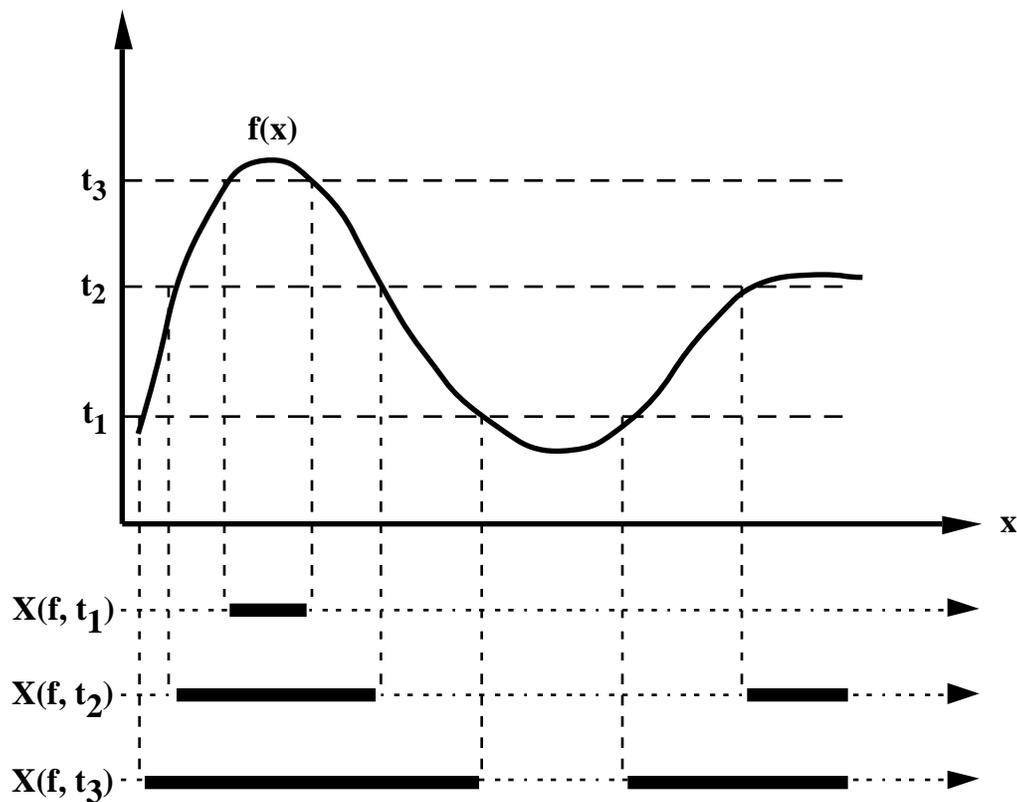


Figure 2.2. Example of threshold decomposition of a function into cross sections. (Adapted from [5].)

A set processing filter $f(X)$ is said to be “increasing” [5] if for any two sets A and B where $A \subseteq B$, the filtered sets maintain the same set relationship; that is, $f(A) \subseteq f(B)$. This property is also called the “stacking property” [17]. A discrete, binary set processing filter possesses the stacking property if and only if its output can be expressed as a Boolean function that does not contain the complement of any of the input variables [17]. Such expressions are called positive Boolean functions.

Set processing filters that are increasing (possess the stacking property) may be converted to function processing filters by performing the set processing operations on the individual cross sections of a signal; the filtered signal is found by stacking the filtered cross sections in the manner described previously. This process of converting the cross sections back into a function is a supremum operation. Examples of filters that are increasing and therefore may operate on a threshold decomposition of a signal are all order statistic filters [18, 19], including the median filter, and the morphological filters. Function processing filters that operate in this manner on the cross sections of a function are called stack filters [6, 17, 20]. Note that not all function processing filters obey the stacking property; those that do are part of a subset of function processing filters called function and set processing (FSP) filters [5, 13]. FSP filters are useful because they can accept either sets or functions as input, and give the same type of output as the input they receive.

The set processing morphological filters may be converted to a function processing operation by this stacking property. The resulting grayscale morphological filters are a subset of the stack filters. All stack filters are FSP filters, and so the morphological filters that use set structuring elements are FSP filters. The resulting expressions for the FSP morphological filters are [5, 13]:

$$\text{Erosion: } (f \mathbf{E} \tilde{N})(\mathbf{z}) = \inf \{f(\mathbf{y}): \mathbf{y} \in N_{\mathbf{x}}\} \quad (2.3)$$

$$\text{Dilation: } (f \mathbf{D} \tilde{N})(\mathbf{z}) = \sup \{f(\mathbf{y}): \mathbf{y} \in N_{\mathbf{x}}\} \quad (2.4)$$

where $f(\mathbf{y})$ denotes an m -dimensional upper semicontinuous function, N denotes a compact k -dimensional set ($k \leq m$), and \mathbf{y} and \mathbf{z} denote points in \mathbb{R}^k and \mathbb{R}^m ,

respectively. The infimum (inf) and supremum (sup) operations reduce to simple minimum and maximum operations, respectively, when acting on discrete signals.

The structuring element of a morphological filter does not have to be a set. Like the signal the filter operates on, the structuring element may be a function. In this case, the output of the morphological filter is always a function, so such filters are not FSP, but are function processing. Instead of the structuring element set N , the structuring element is a k -dimensional upper semicontinuous function $g(\mathbf{z})$ that is defined over a compact region of support, S . As for the set structuring element case, the erosion and dilation are defined as the Minkowski subtraction and addition of the signal with the reflection of the structuring element about the origin. Let $\tilde{g}(\mathbf{z}) = g(-\mathbf{z})$ denote this reflection. The morphological erosion and dilation of a function by a function are given by:

$$\text{Erosion: } (f \ominus \tilde{g})(\mathbf{z}) = \inf \{ (f(\mathbf{z}+\mathbf{y}) - g(\mathbf{y})): \mathbf{y} \in S \} \quad (2.5)$$

$$\text{Dilation: } (f \oplus \tilde{g})(\mathbf{z}) = \sup \{ (f(\mathbf{z}+\mathbf{y}) + g(\mathbf{y})): \mathbf{y} \in S \} \quad (2.6)$$

Note that if the function structuring element $g(\mathbf{z})$ is zero over its entire region of support S , then the above expressions are equivalent to the FSP expressions with structuring element $N = S$. Set structuring elements are far more widely used than function structuring elements in applications of mathematical morphology. In theory, however, function structuring elements are a significant addition to morphological filtering because they are able to process signals and image based on a specific intensity profile over a certain shape. Set structuring elements assume a flat intensity profile over their shapes.

2.1.2. Compound Morphological Operators

Erosion and dilation are complementary operations, one shrinking the size of objects in an image and the other expanding them. However, erosion and dilation are not inverses of each other. Some objects are completely removed by erosion, and therefore cannot be restored by dilation. Likewise, dilation often joins nearby objects which erosion cannot then separate. The compound morphological operations formed by performing the complementary operators in sequence are the morphological operations “opening” and “closing.” Opening is defined as erosion followed by dilation, while closing is dilation followed by erosion. The structuring element used for the second operation is the reflection about the origin of the structuring element used for the first operation. The specific expressions for opening (denoted by a subscript) and closing (denoted by a superscript) for a set structuring element N acting on a function $f(\mathbf{z})$ are:

$$\text{Opening: } f_N(\mathbf{z}) = [(f \mathbf{\boxminus} \tilde{N}) \mathbf{\boxplus} N] (\mathbf{z}) \quad (2.7)$$

$$\text{Closing: } f^N(\mathbf{z}) = [(f \mathbf{\boxplus} \tilde{N}) \mathbf{\boxminus} N] (\mathbf{z}) \quad (2.8)$$

The structuring element rotation between the erosion and dilation creates an “effective” structuring element for opening and closing that is symmetric about the origin, even if the original structuring element is asymmetric. This is illustrated in Figure 2.3 below. In most applications, this rotation has no practical effect since the original structuring element is usually symmetric.

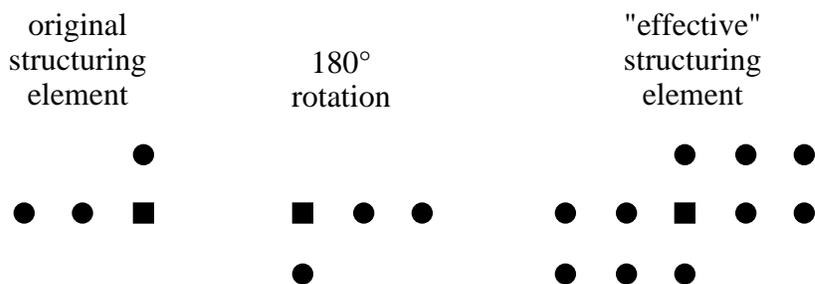


Figure 2.3. Illustration of an “effective” structuring element of opening and closing formed by rotation between successive erosions and dilations. Squares (■) denote the origin of the plane; circles (●) denote other points in the set.

The opening and closing operations are also complementary, and when applied in sequence, they form the doubly compound morphological operators open-close (OC) and close-open (CO):

$$\text{Open-Closing:} \quad \text{OC}(f; N) = (f_N)^N \quad (2.9)$$

$$\text{Close-Opening:} \quad \text{CO}(f; N) = (f^N)_N \quad (2.10)$$

2.1.3. Examples

1-D Examples

Figures 2.4–2.7 below illustrate the operation of the simple and compound morphological operators on a one-dimensional signal (frequency modulated in this example). For finite-length signals, there is a problem defining the output of filters near the ends of the signal; in the following examples, the output near the ends is found by “padding” the signal at either end with as many repetitions of the first or last value in the signal as necessary to define the output. Figure 2.4 demonstrates the effect of erosion and opening. Notice especially how opening preserves “negative” features of the signal, but cuts off positive impulses narrower than the structuring element. Figure 2.5 compares dilation and closing of the

same signal; the effect is exactly the opposite. Closing preserves positive features and cuts off negative impulses. The effects of cascaded operators, open-closing (OC) and close-opening (CO), are shown in Figures 2.6 and 2.7 respectively. Notice that as the signal frequency increases, the close-opening tends toward flat regions with high values, whereas the open-closing yields flat regions with low values.

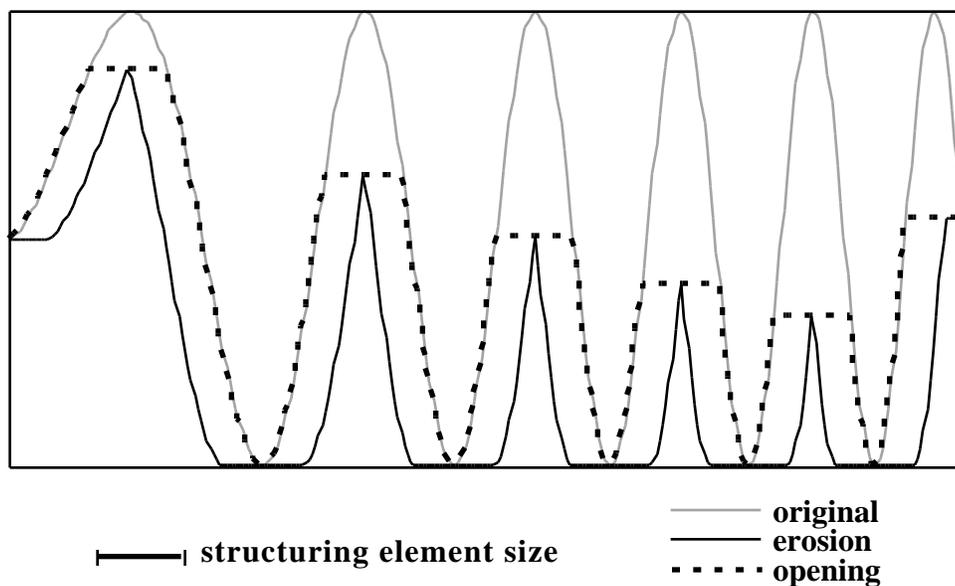


Figure 2.4. Effect of erosion and opening on a 1-D frequency-modulated signal.

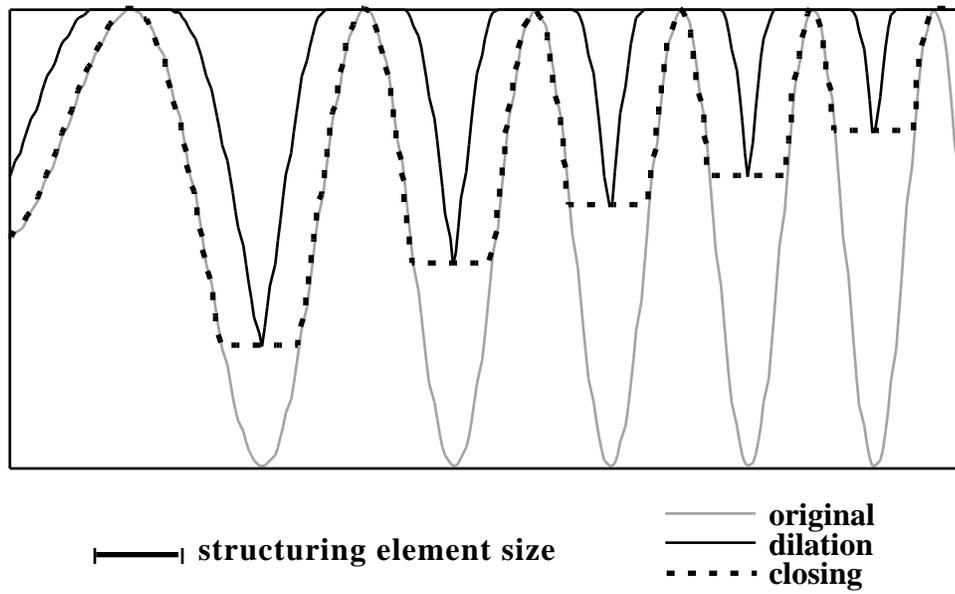


Figure 2.5. Effect of dilation and closing on a 1-D frequency-modulated signal.

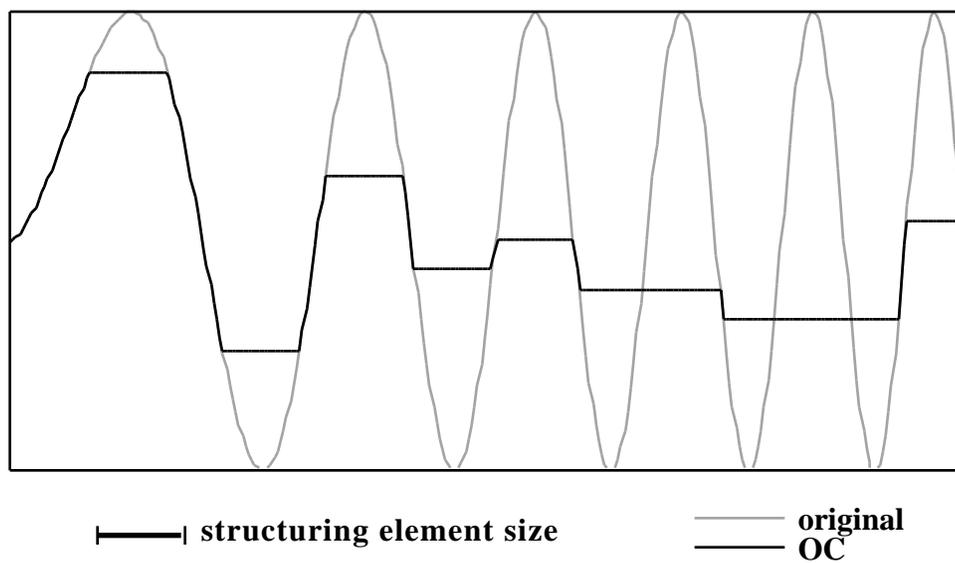


Figure 2.6. Effect of open-closing (OC) on a 1-D frequency-modulated signal.

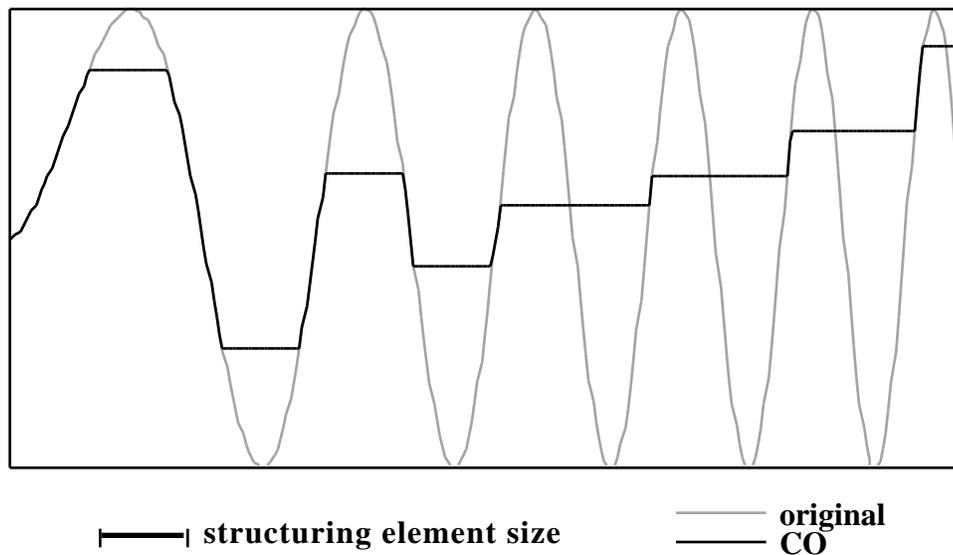


Figure 2.7. Effect of close-opening (CO) on a 1-D frequency-modulated signal.

2-D Examples

Figures 2.8–2.14 below illustrate the effect of the morphological filters in two dimensions with a square structuring element. Figure 2.8 is the original 110×110 pixel binary image of several characters. In the applications given in Chapter 5, image features are considered to have high gray level values and background to have low values; this means that dilation expands image features and erosion shrinks them. However, in the binary examples given below, the characters (features) are black, and so black will be considered to take on the value 1 and white the value 0. This is inverted from the convention of 0 for black and maximum range (often 255) for white which will be used later in this work.

Figures 2.9 and 2.10 show the erosion and dilation, respectively, of the original image in Figure 2.8 using a 5×5 square structuring element. Note that the thin parts of the characters disappear in the eroded image, while nearby portions of adjacent characters are merged in the dilated image. The opening and closing of the image with a 5×5 structuring element are illustrated in Figures 2.11 and 2.12, respectively. The opening, which expands the characters back from what was left after erosion, still shows many gaps where thin features were removed. The closing, however, preserves the thin features of the characters, but partially closes some of the enclosed white areas and leaves many of the characters touching. The doubly compound operators OC and CO, shown in Figures 2.13 and 2.14 respectively, are quite similar in appearance. There are only a few small differences between the OC and CO results, all of which are related to sharp corners in the original image (such as the angles in the “4”).

One important trait to notice in Figures 2.9–2.14 is the effect of the structuring element shape on the filtered images. In every case, the square shape of the structuring element has a noticeable effect, preserving straight edges and 90° corners and converting other features to a “more square” appearance. The fact that the structuring element shape greatly influences the output of the filter is one of the most important features of mathematical morphology.



Figure 2.8. Original binary image.

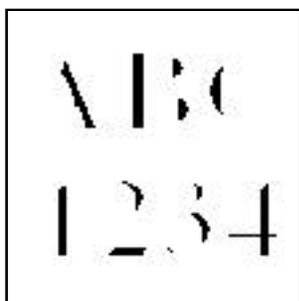


Figure 2.9. Image after erosion.



Figure 2.10. Image after dilation.

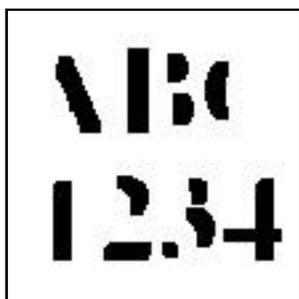


Figure 2.11. Image after opening.



Figure 2.12. Image after closing.

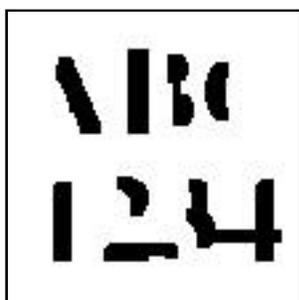


Figure 2.13. Image after OC.

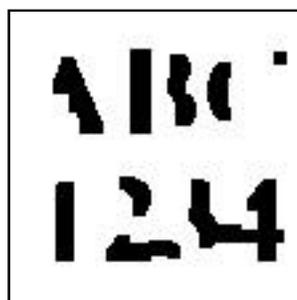


Figure 2.14. Image after CO.

2.2. DETERMINISTIC PROPERTIES

The deterministic properties of nonlinear filters have long been studied to gain some insight into their operation. Since standard frequency analysis techniques in the Fourier domain cannot be used with nonlinear filters, other methods have been developed to provide some insight into how these filters behave in a given situation. An important deterministic property of a nonlinear filter is its root signal set, the set of signals that are unchanged by the operation of the filter. Root signals are also called “fixed points” of a filter. The root signal set was specified for the median filter in 1981 [8, 9] and for the morphological filters in 1987 [6]. The relationship between the root signal sets of the median and morphological filters was also examined by Maragos in [6].

2.2.1. Median Filter Root Signals

The median filter is an order statistic (stack) filter that replaces the center value in the filter window with the median of the values in the window. If the values in the window are updated as the filter acts on the signal, it is called a recursive median filter. Non-recursive median filtering, which is far more common than recursive median filtering, always acts on the original values in the signal. For a signal $f(\mathbf{x})$ and a filter window W , the non-recursive median filter is denoted as shown in equation (2.11) below:

$$\text{med}(f; W)(\mathbf{x}) = \text{median} \{f(\mathbf{y}) : \mathbf{y} \in W_{\mathbf{x}}\} \quad (2.11)$$

where $W_{\mathbf{x}} = \{\mathbf{x} + \mathbf{a} : \mathbf{a} \in W\}$ is the window W centered at \mathbf{x} .

Repeated application of the median filter is denoted by a superscript; for example, $\text{med}^3(f; W)$ denotes the result of three iterations of the non-recursive median filter with window W over the signal f .

The root signal set of the 1-D median filter for finite-length signals consists only of signals that are everywhere locally monotonic of length $n + 2$, where W is $2n+1$ points long ($|W| = 2n + 1$) [8, 9]. This means that any section of a finite-length median root signal of at least $n + 2$ points is monotonic (nonincreasing or nondecreasing). This result assumes the signal is padded appropriately with constant regions to obtain the filter output near the ends, as described previously. Gallagher and Wise [9] stated the same result slightly differently: a finite-length median root signal consists only of constant neighborhoods and edges. A constant neighborhood is an area of constant value of at least length $n + 1$ (just over half the length of W) and an edge is a monotonic region of any length between two constant neighborhoods. This root signal set indicates that the median filter preserves slowly varying regions and sharp edges, but alters impulses and rapid oscillations.

For infinite-length signals, Tyan [8] showed that another type of root signal exists for the non-recursive median filter. These root signals, the “fast-fluctuating” roots, consist solely of rapid oscillations between two values. Nowhere in these signals is there even one monotonic region of length $n + 1$. For example, the infinite-length signal $\dots, 1, 0, 1, 0, 1, 0, 1, 0, \dots$ is a root of the non-recursive median filter with a window width of $4k+1$, where k is any positive integer. Although this type of signal is seldom encountered in practical applications, sections of a finite-length signal that fluctuate quickly (even if they

are not bi-valued) are often passed by the median filter without much smoothing. An example of this situation is shown in Figure 2.15. The original signal has an oscillation in it, and when filtered by a 5-wide median filter only the first and last two peaks in the oscillation are smoothed.

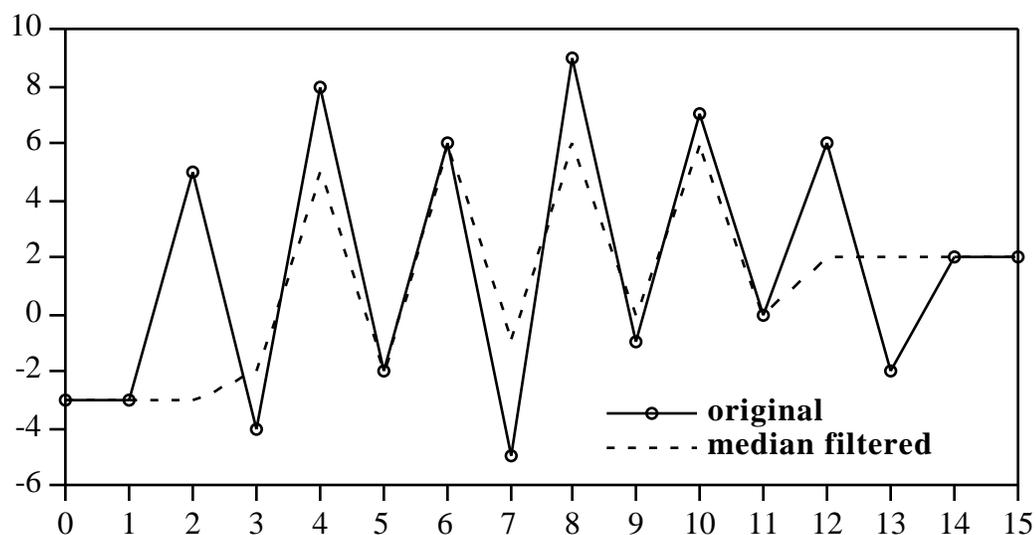


Figure 2.15. Oscillatory signal filtered by a 5-wide non-recursive median filter.

Gallagher and Wise [9] showed that repeated iteration of the non-recursive median filter on a signal of length L reduces the signal to a root signal in at most $(L-2)/2$ passes. This result comes from the observation that on the first pass, the first and last values must be unchanged; on succeeding passes, the area at the beginning and end of the signal that is unchanged by filtering is at least one point longer. Figure 2.15 also illustrates how oscillatory signals are reduced to a root a few points at a time from either end. The root signal resulting from repeated iteration of the non-recursive median filter is denoted $\text{med}(f; W)$. Although the

number of passes required to reduce most signals to a median root is fairly small, in some instances the number of passes required is very large.

The recursive median filter yields a root signal from any input signal in a single pass; however, the resulting root is usually not as good a representation of the original signal as $\text{med}(f; W)$. This is because the recursive filter allows signal values to propagate along the signal sequence, so that the output value at any point is not necessarily related to the values currently in the filter window. The recursive median filter is rarely used in practice for this reason. Figure 2.16 is an example of 5-wide recursive median filtering on the same signal as in Figure 2.15. The result is a root signal, but it is very different from the result of non-recursive median filtering.

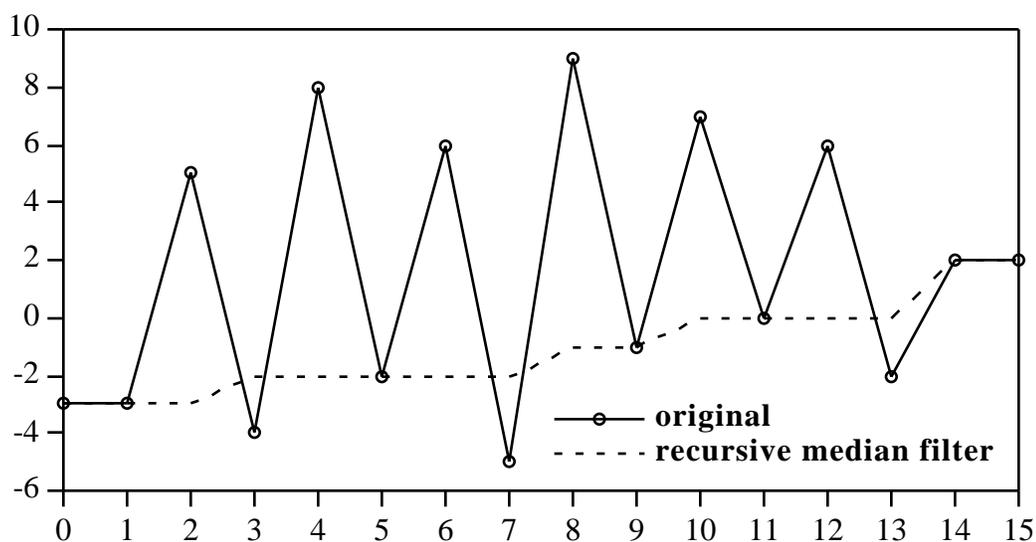


Figure 2.16. Oscillatory signal filtered by a 5-wide recursive median filter.

2.2.2. Relationship of Morphological and Median Filtered Signals

Ordering Relationships

The morphological filters have a specific ordering relationship with each other, with the original signal, and with the median filter root signal $\text{med}(f; W)$. Since erosion is a minimum operation, an eroded signal is less than or equal to the original signal at every point. A dilated signal is greater than or equal to the original signal at every point because it is a maximum operation. Morphological opening is the dilation of an eroded signal, so the opened signal is pointwise greater than or equal to the eroded signal. However, the opened signal is also pointwise less than or equal to the original signal. Similarly, the morphological closing of a signal is pointwise less than or equal to the dilated signal and greater than or equal to the original signal. These results are summarized in the inequality below.

$$(f \mathbf{E} \tilde{N}) \leq f_N \leq f \leq f^N \leq (f \mathbf{D} \tilde{N}) \quad (2.12)$$

Note that another immediate consequence of this result is that close-opening (CO) is pointwise not greater than closing, and open-closing (OC) is pointwise not less than the opening.

Maragos and Schafer [6] showed that the result of median filtering a signal with window W is bounded from below by the opening of the signal and bounded from above by the closing of the signal with a structuring element N :

$$f_N \leq \text{med}^k(f; W) \leq f^N$$

for $k = 1, 2, 3, \dots$ (2.13)

Likewise, open-closing and close-opening bound the root signal found by repeated median filtering:

$$\text{OC}(f; N) \text{ med } (f; W) \text{ CO}(f; N) \quad (2.14)$$

Together, the inequalities given above yield the following relationships:

$$(f \boxminus \tilde{N}) \text{ } f_N \text{ OC}(f; N) \text{ med } (f; W) \text{ CO}(f; N) \text{ } f^N \text{ } (f \boxplus \tilde{N}) \quad (2.15)$$

These results indicate that the morphological filters introduce a bias into a signal relative to the median filtered signal. The statistical properties of the morphological and median filters, as developed in Section 2.3, support this conclusion.

Root Signals

The root signal sets of nonlinear filters are important because they indicate the types of signals a nonlinear filter tends to produce. Erosion and dilation pass only constant signals unchanged; therefore, they tend to create flat spots in signals. These flat regions enlarge upon repeated erosion or dilation until a root signal of the operation is achieved: a constant signal equal to the global minimum or maximum value in the original signal. Opening and closing behave somewhat differently. For instance, both opening and closing are *idempotent*; that is, the first iteration of opening or closing results in a root signal for that operation. A signal that has been opened is invariant to further openings with the same structuring element; similarly, a closed signal is invariant to further closings. This is expressed by:

$$(f_N)_N = f_N \quad (2.16)$$

$$(f^N)^N = f^N \quad (2.17)$$

The root signal sets of opening and closing are quite different from each other: opening removes positive impulses and preserves negative impulses, while closing removes negative impulses and preserves positive impulses. Both operations preserve monotonic ramps and edges. However, the intersection of the root signal sets of opening and closing is equivalent to the non-fast fluctuating root signal set of the median filter with window W [6]. This relationship means that $f = \text{med}(f; W)$ if $f = f_N$ and $f = f^N$. A “constant neighborhoods and edges” median filter root signal is always a root signal of both opening and closing.

Opening followed by closing (OC) with structuring element N results in a signal that is invariant to further opening and closing, and is thus a root signal of the median filter with window W . Close-opening (CO) has the same property. Both OC and CO filtering therefore yield a root signal of the median filter in one pass. However, the root signals output by OC and CO are not necessarily equal to each other nor equal to the root signal achieved by repeated median filtering. One case of particular interest is that of the infinite-length fast-fluctuating bi-valued median filter root signals. These median roots are not root signals of OC or CO, as described below.

Recall that one of the conditions for the fast-fluctuating root signals was that there is no monotonic region of length $n + 1$ anywhere in the signal. Assume that the two values in the fast-fluctuating signal are a and b , and that $a < b$. Since $|N| = n + 1$, there is no location in the signal where the structuring element of the morphological operators does not contain both a and b . Therefore, erosion reduces the fast-fluctuating signal to a constant signal with value a , and dilation reduces it to a constant signal with value b . Erosion, opening, and OC all output

the constant value a , and dilation, closing, and CO output b . So, the fast-fluctuating root signals of the median filter are not roots of the morphological OC or CO operations; furthermore, the OC and CO reduce this type of signal to a constant in a single pass.

Although OC and CO seem to have some desirable deterministic properties similar to those of the median filter, there are some problems with the standard morphological operators. The dilation-based operators (closing and CO) always give an output that lies above or along the root signal found by median filtering, and the erosion-based operators (opening and OC) give an output along or below the median root. This indicates that the morphological filters introduce a deterministic bias to signals that they operate on. This makes sense because the filters are based on maximum and minimum operations. The statistical analyses of the next section also show that the morphological operators are biased with respect to the median filter.

2.3. STATISTICAL PROPERTIES

2.3.1. One-Dimensional Properties

The statistical properties of the standard one-dimensional morphological operators were derived by Stevenson and Arce [13]. They gave expressions for the output distribution function in terms of the input distribution function, when all points in the input signal are independent and identically distributed (i.i.d.). They determined the statistical properties by using threshold decomposition to find the cross sections of the input signal, then finding the probabilities of various

deterministic structures (constant neighborhoods, positive and negative impulses, and oscillations). The morphological operators map each of these structures deterministically to the same or another structure, and this knowledge combined with the stacking property allows the output statistics to be determined. For a 1-D discrete input signal $f(\mathbf{x})$ with values that are i.i.d. random variables with distribution function $G(\mathbf{z})$, the distribution function of the output of CO, denoted $H_{\text{co}}(\mathbf{z})$, with structuring element N of size $n + 1$ is [13]:

$$H_{\text{co}}(\mathbf{z}) = (n+1)G(\mathbf{z})^{n+1} - nG(\mathbf{z})^{n+2} + G(\mathbf{z})^{2n+2}(1-G(\mathbf{z})) + \frac{(n+2)(n+1)}{2}G(\mathbf{z})^{2n+2}(1-G(\mathbf{z}))^2 \quad (2.18)$$

Similarly, the output distribution for OC, denoted $H_{\text{oc}}(\mathbf{z})$, under the same assumptions is [13]:

$$H_{\text{oc}}(\mathbf{z}) = 1 - \left[(n+1)(1-G(\mathbf{z}))^{n+1} - n(1-G(\mathbf{z}))^{n+2} + (1-G(\mathbf{z}))^{2n+2}G(\mathbf{z}) + \frac{(n+2)(n+1)}{2}(1-G(\mathbf{z}))^{2n+2}G(\mathbf{z})^2 \right] \quad (2.19)$$

The probability density functions $h_{\text{co}}(\mathbf{z})$ and $h_{\text{oc}}(\mathbf{z})$ of the filter outputs are derived by differentiating the distribution functions with respect to \mathbf{z} . The input density function is $g(\mathbf{z}) = dG(\mathbf{z})/d\mathbf{z}$.

The last two terms in equation (2.18) result from the positive-impulse removing property of CO, whereas the first two terms result from the more basic properties of closing. Since closing alone does not remove positive impulses, its output distribution function is simply the first two terms of equation (2.18). Similarly, the output distribution function of opening is the same as that of OC with only the first two terms inside the bracket of equation (2.19). These distributions are shown in the equations below.

$$H_{\text{close}}(\mathbf{z}) = (n+1)G(\mathbf{z})^{n+1} - nG(\mathbf{z})^{n+2} \quad (2.20)$$

$$H_{\text{open}}(\mathbf{z}) = 1 - \left[(n+1)(1-G(\mathbf{z}))^{n+1} - n(1-G(\mathbf{z}))^{n+2} \right] \quad (2.21)$$

The output distributions for dilation and erosion are the distributions for maximum and minimum order statistics [21, 22]. The equations are:

$$H_{\text{dilate}}(\mathbf{z}) = G(\mathbf{z})^{n+1} \quad (2.22)$$

$$H_{\text{erode}}(\mathbf{z}) = 1 - [1 - G(\mathbf{z})]^{n+1} \quad (2.23)$$

The output statistics of the median filter are also found in order statistics research [7, 21]. For a 1-D median filter of window size $|W| = 2n+1$ with i.i.d. random input having a density function $g(\mathbf{z})$, the output density function is [7]:

$$h_{\text{med}}(\mathbf{z}) = (2n+1) \frac{2n}{n} g(\mathbf{z})G(\mathbf{z})^n [1 - G(\mathbf{z})]^n \quad (2.24)$$

Figures 2.17–2.21 below illustrate the output probability density functions of the CO, OC, closing, opening, and median filters for $n = 1, 3, 5,$ and 10 with i.i.d. random input uniformly distributed on $[0, 1]$. For the morphological filters, these values of n correspond to structuring elements of length $2, 4, 6,$ and $11,$ while for the median filter they correspond to window widths of $3, 7, 11,$ and $21.$ These plots show that the standard morphological operators are biased with respect to the median, and that the amount of bias increases as the structuring element size increases. The median values of the output distributions of the various operators with uniformly distributed input are shown in Table 2.1 for various filter sizes. Note that the complementary morphological operators are equally and oppositely biased, and that erosion and dilation are more biased than opening and closing, which are in turn more biased than OC and CO.

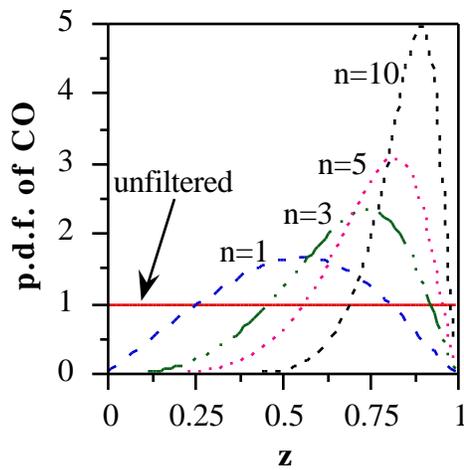


Figure 2.17. Density function of CO output with uniformly distributed input. (After [13].)

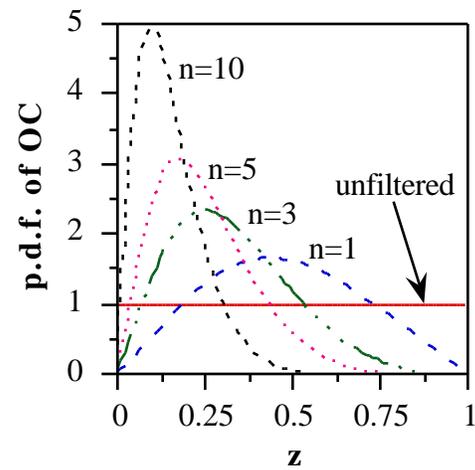


Figure 2.18. Density function of OC output with uniformly distributed input.

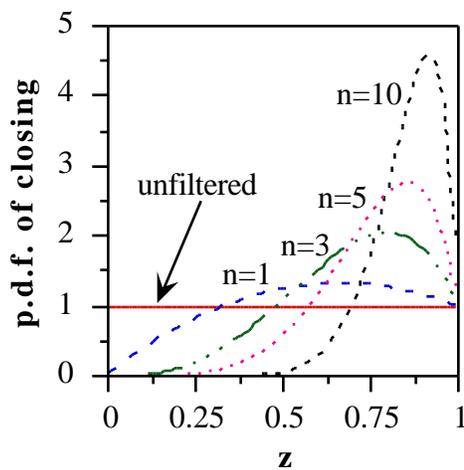


Figure 2.19. Density function of closing output with uniformly distributed input.

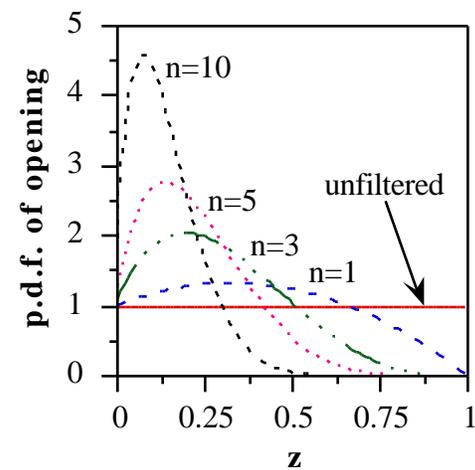


Figure 2.20. Density function of opening output with uniformly distributed input.

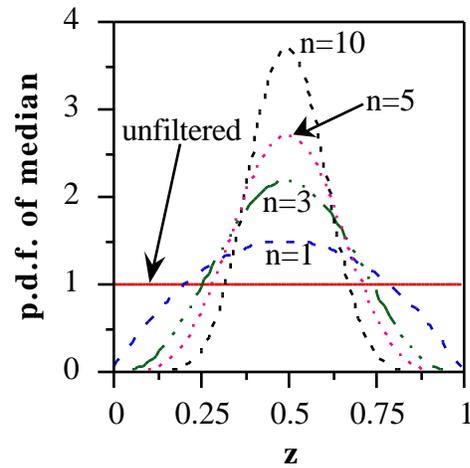


Figure 2.21. Density function of median filter output with uniformly distributed input.

Table 2.1. Median of filter output distributions with uniformly distributed input.

size (n)	Median of Filter Output Distribution						
	Erosion	Opening	OC	Median	CO	Closing	Dilation
1	0.29	0.40	0.47	0.50	0.53	0.60	0.71
3	0.16	0.28	0.31	0.50	0.69	0.72	0.84
5	0.11	0.21	0.23	0.50	0.77	0.79	0.89
10	0.06	0.13	0.14	0.50	0.86	0.87	0.94

2.3.2. Two-Dimensional Properties

The statistical properties of the basic morphological and median filters in two dimensions are simply a direct extension of their one-dimensional properties. The general expressions for the output distributions of erosion and dilation with a structuring element N with $|N|$ pixels are:

$$H_{\text{dilate}}(\mathbf{z}) = G(\mathbf{z})^{|N|} \quad (2.25)$$

$$H_{\text{erode}}(\mathbf{z}) = 1 - [1 - G(\mathbf{z})]^{|N|} \quad (2.26)$$

Equations (2.25) and (2.26) are valid for N of any shape and size. For a square structuring element N of $(n+1) \times (n+1)$ pixels, the value of $|N|$ is $(n+1)^2$.

The 2-D median filter with the closest correspondence to the compound morphological filters has a filter window $W = (N \boxplus \tilde{N})$. For example, a square structuring element N of size $(n+1) \times (n+1)$ corresponds to a median filter window W that is a $(2n+1) \times (2n+1)$ square ($|W| = (2n+1)^2$). For any shape and size W , the output distribution of the 2-D median filter with window W is:

$$h_{\text{med}}(\mathbf{z}) = \binom{2w}{w} g(\mathbf{z})^w [1 - G(\mathbf{z})]^w \quad (2.27)$$

where $w = \frac{1}{2} (|W| - 1)$

Note that equations (2.25)–(2.27) are basically identical to the 1-D expressions given in equations (2.22)–(2.24). However, for the compound morphological operators, the output statistics are significantly more difficult to derive. They are not simple extensions of the one-dimensional cases because of different geometrical considerations that arise from the interaction of 2-D structuring elements between passes of erosion and dilation. Astola, Koskinen,

and Neuvo have developed the expressions for the 2-D output statistics of the compound morphological operators [23, 24]. Unfortunately, for all but very small and simple structuring elements, the expressions are intractable. The general expression given in [24] for the output distribution of any stack filter $h(\mathbf{x})$ is:

$$G(\mathbf{z}) = \sum_{\bar{\mathbf{x}} \in h^{-1}(0)} (1 - F(\mathbf{z}))^{w(\bar{\mathbf{x}})} F(\mathbf{z})^{m-w(\bar{\mathbf{x}})} \quad (2.28)$$

where $h^{-1}(0) = \{\bar{\mathbf{x}}: h(\mathbf{x}) = 0\}$; that is, $h^{-1}(0)$ is the set of all sets $\bar{\mathbf{x}}$ that result in a zero output from the filter h ; $w(\bar{\mathbf{x}})$ is the number of ones in $\bar{\mathbf{x}}$; and $m = |\bar{\mathbf{x}}|$. Note that $\bar{\mathbf{x}}$ must always be the same size and shape as the region of support of the filter; that is, $\bar{\mathbf{x}}$ includes all the points and only the points included in the positive Boolean expression for the stack filter $h(\mathbf{x})$. For erosion and dilation, $\bar{\mathbf{x}}$ must be the same size and shape as the structuring element N . For the median filter, $\bar{\mathbf{x}}$ conforms to the window W . For closing and opening, $\bar{\mathbf{x}}$ must have size and shape $(N \ominus \tilde{N})$, which is the region of support for those operations. Since $\bar{\mathbf{x}}$ is always the same size for a given filter, m is a constant.

To illustrate how equation (2.28) is simplified for a specific structuring element, consider the case of a 2×2 square structuring element. The “overall” window size (region of support) for the singly compound morphological operations (opening and closing) is a 3×3 square, so $m = 9$. For the closing operation, as many as five and as few as zero of the nine total pixels may equal 1 and still give an output of 0. Therefore, $w(\bar{\mathbf{x}})$ ranges from 0 to 5 and $m - w(\bar{\mathbf{x}})$ ranges from 9 to 4. The set where $\bar{\mathbf{x}}$ is all zeros gives a zero output, so one term of $G(\mathbf{z})$ is $F(\mathbf{z})^9$. There are 9 possible sets with only a single “1” in $\bar{\mathbf{x}}$; however, only eight of these sets give a zero output when closing is applied. Any set with a

zero in the central pixel gives a zero output to closing. Therefore, the corresponding term in equation (2.28) is $8(1-F(\mathbf{z}))F(\mathbf{z})^8$. Continuing in the same manner, we find that there are 26 sets with two 1's, 36 sets with three, 20 sets with four, and 4 sets with five 1's that give a zero output to closing. The four sets with five 1's are the sets that have a 2×2 block of 0's within the overall 3×3 window. The complete expression for the output distribution of closing with a 2×2 structuring element for i.i.d. input, derived from equation (2.28), is then:

$$\begin{aligned} G(\mathbf{z}) = & F(\mathbf{z})^9 + 8(1 - F(\mathbf{z}))F(\mathbf{z})^8 \\ & + 26(1 - F(\mathbf{z}))^2 F(\mathbf{z})^7 + 36(1 - F(\mathbf{z}))^3 F(\mathbf{z})^6 \\ & + 20(1 - F(\mathbf{z}))^4 F(\mathbf{z})^5 + 4(1 - F(\mathbf{z}))^5 F(\mathbf{z})^4 \end{aligned} \quad (2.29)$$

which simplifies to [24]:

$$G(\mathbf{z}) = -F(\mathbf{z})^9 + 4F(\mathbf{z})^8 - 2F(\mathbf{z})^7 - 4F(\mathbf{z})^6 + 4F(\mathbf{z})^4 \quad (2.30)$$

The difference between this expression for the 2-D closing and the corresponding 1-D closing is shown in Figure 2.22 below. The input points are i.i.d., uniformly distributed on $[0, 1]$. The output distributions for a 1-D closing with a structuring element of length 4 and for a 2-D closing with a 2×2 structuring element are similar, but the 2-D distribution is slightly less peaked, and therefore slightly less biased.

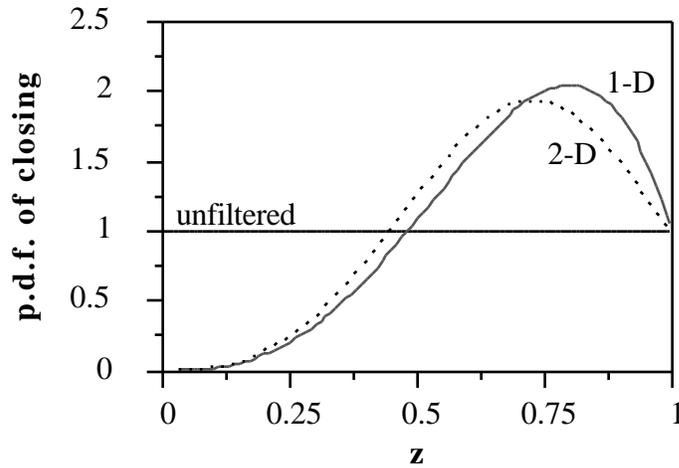


Figure 2.22. Density functions of 1-D and 2-D closing output with uniformly distributed input. 1-D structuring element length = 4. 2-D structuring element is 2×2 square.

Another structuring element for which equation (2.28) is tractable is a plus-shaped 3×3 element. The “overall” window for closing with this structuring element is a diamond-shaped region with 13 pixels. The resulting output distribution may have terms up to the 13th order. After simplification, the output distribution of closing with a 3×3 plus-shaped structuring element with i.i.d. random input with distribution $F(\mathbf{z})$ is:

$$G(\mathbf{z}) = -4F(\mathbf{z})^{12} + 6F(\mathbf{z})^{11} + 4F(\mathbf{z})^{10} - 2F(\mathbf{z})^9 - 8F(\mathbf{z})^8 + 5F(\mathbf{z})^5 \quad (2.31)$$

The output density function for the 3×3 plus-shaped structuring element with uniformly distributed input is shown in Figure 2.23 below. Also shown for comparison is the output density function of 1-D closing with a 5-wide structuring element. Again, the 2-D closing is less biased than the 1-D closing with the same number of points. One reason for this difference is that the overlap between

structuring elements is different in 2-D than in 1-D; there is much more overlap between adjacent structuring elements in one dimension.

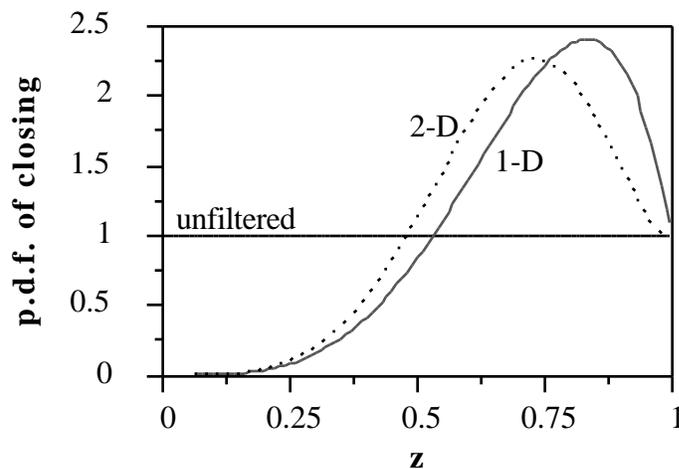


Figure 2.23. Density functions of 1-D and 2-D closing output with uniformly distributed input. 1-D structuring element length = 5. 2-D structuring element is 3×3 plus-shaped.

Although equation (2.28) gives us a way to compute 2-D output distributions, for large filtering windows the amount of computation required increases exponentially. Unless the combinations that produce zero output have a very simple form (as for the median filter, erosion, and dilation), it is nearly impossible to derive the coefficients for equation (2.28) for larger structuring elements. For example, to find the coefficients for the output distribution of closing with a 2×2 structuring element, only 512 different 3×3 sets must be examined to see if they produce zero output, and only 218 of those sets do not yield the answer trivially. If the structuring element size is increased to 3×3 , over 33 million different 5×5 sets must be checked, and over 16 million do not

have a trivial solution. For this reason, it is impractical to derive output distributions for 2-D filters that have a region of support larger than about five or six pixels. The CO and OC filters, even with a very small structuring element, have such large support regions (7×7 , or 49 pixels, for just a 2×2 structuring element) that finding their output distributions is not practical even on supercomputers.

Although the output distribution functions for most compound 2-D morphological filters cannot be found with current computer technology, it is clear from the above examples that the general behavior of the filters is similar to that of the corresponding 1-D filters. The differences arise from the fact that different geometric arrangements of the same number of 1's and 0's within a window give different output for the compound morphological operators. Filters that do not exhibit this geometric dependence, namely the order statistic filters, have 2-D output distributions that are identical to their 1-D distributions for the same number of pixels. For example, a 3×3 square median filter has the same output distribution as a 9-wide 1-D median filter.

2.4. SUMMARY

This chapter reviewed the definitions and many of the important properties of filters based on mathematical morphology. The stacking property and the distinctions between set and function processing filters are two concepts that are necessary to derive the deterministic and statistical properties of the morphological filters. The deterministic properties of the morphological filters,

especially the CO and OC operators, are similar to those of the median filter. Although these impulse-removing and edge-preserving properties are well known and often desired in the realm of nonlinear filtering, the statistical properties of the morphological operators reveal some potentially undesirable traits. All the morphological operators introduce a bias to the signals they filter. This bias becomes more severe as the size of the structuring element becomes larger. Although the shape control and deterministic properties of the morphological filters show great promise for a variety of applications, a way to eliminate the bias must be found to harness the full potential of the filters.

3. New Filter Structures

3.1. INTRODUCTION

This chapter introduces two new nonlinear filter structures based on the structure of mathematical morphology and derives statistical and deterministic properties for the new filters. The first set of new filters, linear combinations of morphological operators, eliminates the bias of standard morphological operators while retaining the shape control of morphology. The second set of new filters described here is a broader generalization of morphology called the value-and-criterion filter structure. This structure allows the use of both linear and nonlinear operations in a single filter. For example, one type of value-and-criterion filter applies a linear noise smoothing operation to an image while enhancing edges by using a nonlinear operation that selects the area where the linear operator acts. Using the standard terminology for nonlinear filter analysis, I prove 10 original theorems that describe the deterministic properties of the new filters, and I also develop approximations for the statistical properties of the filters.

3.2. LINEAR COMBINATIONS OF MORPHOLOGICAL OPERATORS

Since the complementary morphological operators (OC and CO; opening and closing; erosion and dilation) are equally and oppositely biased, an obvious way to try to remove the bias is by simply taking the average of the two operators. For symmetric input noise distributions, an evenly weighted average clearly should give an unbiased result; however, for asymmetric distributions an

unequally weighted average of the two operators may work better for removing the bias. This chapter describes the various types of linear combinations of morphological operators and illustrates how they alleviate some of the bias problems of standard morphological filters. Theorems describing the root signal sets of these linear combinations are presented, and approximations for the statistical properties of the filters are also given and compared to the properties of the standard morphological filters.

3.2.1. Definitions

Midrange Filter

The midrange filter is defined as the average of the maximum and the minimum values in a filter window, the midpoint of the range of values in the window. Since the erosion is a sliding minimum operation acting over the window N and dilation is a sliding maximum operation acting over the same window, another form of the midrange filter is the average of the two basic morphological operators, erosion and dilation. The midrange filter is a well-known estimator in the order statistics literature; see, for example, [25-27]. The midrange filter is optimal in the mean square sense among all filters that are linear combinations of order statistics for removing uniformly distributed noise from a constant signal [26]. The midrange filter is also the maximum likelihood estimator for uniformly distributed noise [26]. The notation for the midrange filter is given in equation (3.1) below.

$$\text{midr}(f; N) = \frac{1}{2} (f \mathbf{E} \tilde{N} + f \mathbf{D} \tilde{N}) \quad (3.1)$$

Pseudomedian Filter

The pseudomedian filter was originally defined in 1985 by Pratt, Cooper, and Kabir [28]. They defined the filter in one dimension to be the average of the maximum of the minima of $n+1$ subwindows within an overall window and the minimum of the maxima of the same subwindows. Each subwindow is $n+1$ points long, and is within an overall window of length $2n+1$. This structure corresponds to that of morphological opening and closing: the subwindows are the structuring elements N , and the overall window is W . Pratt recast the definition of the pseudomedian filter in his 1991 text [29] by forming the “maximin” and “minimax” functions, which he averages to find the pseudomedian. The maximin and minimax functions are equivalent to morphological opening and closing. The 1-D pseudomedian filter is therefore the average of the opening and closing, as noted previously by this author [30]. Pratt defined a two-dimensional pseudomedian filter in a manner that does not correspond to 2-D opening and closing; however, other work by the author [31, 32] generalized the pseudomedian filter to two dimensions in a manner corresponding to the definition of the pseudomedian as the average of opening and closing. The notation for pseudomedian filter is given in equation (3.2) below.

$$\text{pmed}(f; N) = \frac{1}{2} (f_N + f^N) \quad (3.2)$$

LOCO Filter

Following the example of the midrange and pseudomedian filters, another filter may be formed by taking the average of the OC and CO operators. In 1992

Song and Delp [33] suggested averaging cascaded opening and closing operations within the context of a more generalized filter structure; however, they did not study any of the properties of this filter. I call this filter a LOCO filter [34], because it is a Linear combination of OC and CO. The LOCO filter is given by equation (3.3) below.

$$\text{LOCO}(f; N) = \frac{1}{2} ((f_N)^N + (f^N)_N) \quad (3.3)$$

3.2.2. Deterministic Properties

Preliminaries

Although the properties of the standard morphological operators are understood well, the filters found by averaging these operators behave quite differently in many respects. For example, opening removes positive impulses and closing removes negative impulses, but their average (the pseudomedian filter) does not completely remove either positive or negative impulses [31, 32]. Nevertheless, the 1-D root signals of the midrange, pseudomedian, and LOCO filters are virtually the same as for their corresponding constituent filters, as will be shown subsequently.

To prove the root signal set results for the linear combination filters, several preliminary definitions and theorems must first be given. Most of these results are derived from earlier work on median and morphological filters. Every result in this section applies only to one-dimensional signals. For the median filter, the window W contains $2n+1$ points, and for the morphological filters the structuring element N contains $n+1$ points.

Definition 3.1: A root signal or fixed point of a filter is a signal invariant to the filter.

Definition 3.2 (Tyan [8]): A 1-D signal $f(\mathbf{x})$ is locally monotonic of length m , denoted LOMO(m), if every sequence $\{f(\mathbf{x}), f(\mathbf{x}+1), \dots, f(\mathbf{x}+m-1)\}$ in the signal is monotonic. A signal is nowhere LOMO(m) if there is no sequence $\{f(\mathbf{x}), f(\mathbf{x}+1), \dots, f(\mathbf{x}+m-1)\}$ in the signal that is monotonic.

Definition 3.3(Gallagher [9]): A constant neighborhood is a region of at least $n+1$ consecutive points in a signal that have the same value. An edge is a monotonic sequence between two constant neighborhoods such that the edge and constant neighborhoods taken together are monotonic.

Theorem 3.1(Maragos [6]): [See also equation (2.15).] For any signal $f(\mathbf{x})$,

$$(f \mathbf{K} \tilde{N}) f_N \text{ OC}(f; N) \text{ med } (f; W) \text{ CO}(f; N) f^N (f \mathbf{D} \tilde{N})$$

for all \mathbf{x} .

Theorem 3.2(Maragos [6]): [See also equation (2.12).] For any signal $f(\mathbf{x})$,

$$(f \mathbf{K} \tilde{N}) f_N f f^N (f \mathbf{D} \tilde{N})$$

for all \mathbf{x} .

Theorem 3.3: If $f(\mathbf{x})$ is a root signal of the LOCO filter, then

$$(f \mathbf{K} \tilde{N})(\mathbf{x}) f_N(\mathbf{x}) \text{ OC}(f(\mathbf{x}); N) f(\mathbf{x}) \text{ CO}(f(\mathbf{x}); N) f^N(\mathbf{x}) (f \mathbf{D} \tilde{N})(\mathbf{x})$$

Proof: Follows directly from the definition of the LOCO filter and Theorem 3.1.

Theorem 3.4 (Maragos [6]): If $f(\mathbf{x})$ is a root of both opening and closing by N , it is a root of median filtering by W , or

$$f(\mathbf{x}) = f_N(\mathbf{x}) = f^N(\mathbf{x}) \quad f(\mathbf{x}) = \text{med}\{f(\mathbf{x}); W\}$$

Theorem 3.5 (Maragos [6]): Open-closing or close-opening by N yields a root of median filtering by W in one pass:

$$\text{OC}(f(\mathbf{x}); N) = \text{med}\{\text{OC}(f(\mathbf{x}); N); W\}$$

$$\text{CO}(f(\mathbf{x}); N) = \text{med}\{\text{CO}(f(\mathbf{x}); N); W\}$$

for all \mathbf{x} .

Theorem 3.6 (Tyan [8]): If $f(\mathbf{x})$ is a root of the median filter with window W of size $|W| = 2n+1$ (that is, $f(\mathbf{x}) = \text{med}\{f(\mathbf{x}); W\}$), then $f(\mathbf{x})$ is either:

- (i) LOMO($n+2$), or
- (ii) infinitely long, nowhere LOMO($n+1$), and bi-valued.

These two cases may be restated as:

- (i) consists only of constant neighborhoods and edges [9], or
- (ii) infinitely long and takes on only two values, both of which are in every translation of the structuring element N in $f(\mathbf{x})$.

An important property that distinguishes the morphological filters and their linear combinations from the median filter is that they do not pass median root signals belonging to case (ii) of Theorem 3.6 above. In fact, root signals of this type are reduced to constant signals by any of the morphological operations and their combinations, as proven in Theorem 3.7 below.

Theorem 3.7: Median filter root signals that are infinitely long, nowhere LOMO($n+1$), and bi-valued are reduced to a constant signal by either erosion or dilation and therefore are not roots of any of the compound morphological operators (opening, closing, OC, and CO) or the midrange, pseudomedian, or LOCO filters.

Proof: A fast-fluctuating bi-valued median filter root signal $f(\mathbf{x})$ consists of only two values, a and b . Assume, without loss of generality, that $a < b$. Since each structuring element N in $f(\mathbf{x})$ has both a and b in it, $(f \ominus \tilde{N}) = a$ for all \mathbf{x} and $(f \oplus \tilde{N}) = b$ for all \mathbf{x} . Therefore, $f_N(\mathbf{x}) = (f_N)^N(\mathbf{x}) = a$ and $f^N(\mathbf{x}) = (f^N)_N(\mathbf{x}) = b$ for all \mathbf{x} and so $\text{midr}(f; N) = \text{pmed}(f; N) = \text{LOCO}(f; N) = \frac{1}{2}(a + b) f(\mathbf{x})$ for all \mathbf{x} . Therefore, $f(\mathbf{x})$ is not a root signal of any of the morphological operators or of the midrange, pseudomedian, or LOCO filters.

Q.E.D.

The results of Theorems 3.4, 3.6, and 3.7 together lead directly to the conclusion that if $f(\mathbf{x})$ is a root signal of both opening and closing by N , $f(\mathbf{x})$ consists of only constant neighborhoods and edges. A similar result, using Theorems 3.5 and 3.7, holds for the OC and CO filters.

Several additional preliminary results concerning the effects of erosion and dilation on a signal are required in order to prove the root signal sets of the linear combination filters. These results are given in Theorems 3.8–3.13 and their corollaries below. Definitions 3.4–3.6 will be used in these theorems.

Definition 3.4: A signal $f(\mathbf{x})$ is called n -increasing over some range of points R if for all points $\mathbf{x} \in R$ and $(\mathbf{x}+n) \in R$ there is a point $\mathbf{y} \in (\mathbf{x}, \mathbf{x}+n]$ such that $f(\mathbf{y}) > f(\mathbf{x})$.

Definition 3.5: A signal $f(\mathbf{x})$ is called n -decreasing over some range of points R if for all points \mathbf{x} where $\mathbf{x} \in R$ and $(\mathbf{x}+n) \in R$ there is a point $\mathbf{y} \in (\mathbf{x}, \mathbf{x}+n]$ such that $f(\mathbf{y}) < f(\mathbf{x})$.

Definition 3.6: A signal $f(\mathbf{x})$ is called n -monotonic over some range of points R if $f(\mathbf{x})$ is monotonic and n -increasing over R , or if $f(\mathbf{x})$ is monotonic and n -decreasing over R .

Theorem 3.8: For a signal dilation $(f \mathbf{D} \tilde{N})(\mathbf{x})$ not to contain a constant neighborhood, the original signal $f(\mathbf{x})$ must be infinitely long and at each point \mathbf{x} there must exist a point $\mathbf{y} \in [\mathbf{x}-n, \mathbf{x}+n]$ within n points in *either* direction where $f(\mathbf{y}) > f(\mathbf{x})$.

Proof: Finite length signals must have a global maximum, and every structuring element N that contains the global maximum has the same value for the dilation (the global maximum value itself). Therefore, the dilation of a finite length signal contains a constant neighborhood centered at the global maximum of the original signal. Thus, for a dilation without a constant neighborhood to exist, the original signal must be infinitely long.

At every point in the signal \mathbf{x} , for $f(\mathbf{x}) = (f \mathbf{D} \tilde{N})(\mathbf{x})$, there must be a point \mathbf{y} in within at least one of the structuring elements N containing \mathbf{x} where $f(\mathbf{y}) > f(\mathbf{x})$. Otherwise, each of the $n+1$ structuring elements that contain \mathbf{x} would have a maximum value of $f(\mathbf{x})$, and the dilation would have a constant neighborhood centered at \mathbf{x} . The $n+1$ structuring elements containing \mathbf{x} include all points from $\mathbf{x}-n$ to $\mathbf{x}+n$, so there must be a point \mathbf{y} within n points in either direction of \mathbf{x} where $f(\mathbf{y}) > f(\mathbf{x})$. *Q.E.D.*

Corollary: A signal dilation $(f \mathbf{D} \tilde{N})(\mathbf{x})$ does not contain any constant neighborhoods if and only if the original signal $f(\mathbf{x})$ is infinitely long and satisfies one of the following three conditions:

- 1) $f(\mathbf{x})$ is n -increasing everywhere;
- 2) $f(\mathbf{x})$ is n -decreasing everywhere; or
- 3) $f(\mathbf{x})$ is n -decreasing in the range $R_1 = (-\infty, \mathbf{x}_1)$ of \mathbf{x} and n -increasing in the range $R_2 = (\mathbf{x}_1, +\infty)$ of \mathbf{x} .

Justification: (See Figure 3.1 below.) The specification of $f(\mathbf{y}) > f(\mathbf{x})$ within $[\mathbf{x}-n, \mathbf{x}+n]$ from Theorem 3.8 means that $f(\mathbf{x})$ must be either n -increasing {if $\mathbf{y} \in (\mathbf{x}, \mathbf{x}+n]$ } or n -decreasing {if $\mathbf{y} \in (\mathbf{x}, \mathbf{x}-n]$ } at every point \mathbf{x} . Obviously, any signal that is everywhere n -increasing or everywhere n -decreasing satisfies these conditions. Now consider a signal where there are both n -increasing and n -decreasing regions. An n -decreasing region to the left of an n -increasing region means that there is a point \mathbf{x}_1 on the border between the two regions where there are larger signal values within n points to both the left and right of \mathbf{x}_1 . This is allowed by Theorem 3.8. However, an n -increasing region to the left of an n -decreasing region means that there exists a point \mathbf{x}_2 on the border between the two regions where there is no point in $[\mathbf{x}_2-n, \mathbf{x}_2+n]$ that has a larger value than $f(\mathbf{x}_2)$. This contradicts Theorem 3.8, so all signals with an n -increasing region to the left of an n -decreasing have a constant neighborhood when dilated. Therefore, since a transition from n -decreasing to n -increasing is allowed but from n -increasing to n -decreasing is not, an input signal that has a

dilation without any constant neighborhoods must satisfy one of the three conditions listed in the statement of the corollary.

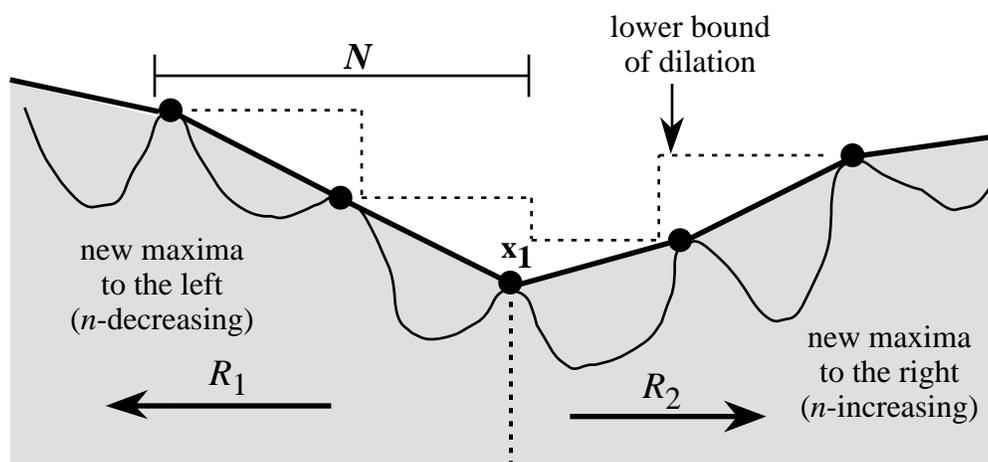


Figure 3.1. Example of a signal that has a dilation with no constant neighborhoods. Black dots represent “required” values; elsewhere the signal may take on any value in the shaded area.

Theorem 3.9: For a signal erosion $(f \boxminus \tilde{N})(\mathbf{x})$ not to contain a constant neighborhood, the original signal $f(\mathbf{x})$ must be infinitely long and at each point \mathbf{x} there must exist a point $\mathbf{y} \in [\mathbf{x}-n, \mathbf{x}+n]$ within n points in *either* direction where $f(\mathbf{y}) < f(\mathbf{x})$.

Proof: Exactly analogous to the proof of Theorem 3.8.

Corollary: The erosion of a signal $(f \boxminus \tilde{N})(\mathbf{x})$ does not contain any constant neighborhoods if and only if the original signal $f(\mathbf{x})$ is infinitely long and satisfies one of the following three conditions:

- 1) $f(\mathbf{x})$ is n -increasing everywhere;

- 2) $f(\mathbf{x})$ is n -decreasing everywhere; or
- 3) $f(\mathbf{x})$ consists of a single, contiguous n -increasing range R_1 immediately to the left of a single, contiguous n -decreasing range R_2 , and all points \mathbf{x} belong to either R_1 or R_2 .

Theorem 3.10: If a signal dilation $(f \mathbf{D} \tilde{N})(\mathbf{x})$ does not contain a constant neighborhood, it is infinitely long and either n -monotonic everywhere or monotonic n -decreasing to the left and monotonic n -increasing to the right of a set of n or fewer global minimum points. If a signal erosion $(f \mathbf{K} \tilde{N})(\mathbf{x})$ does not contain a constant neighborhood, it is infinitely long and either n -monotonic everywhere or monotonic n -increasing to the left and monotonic n -decreasing to the right of a set of n or fewer global maximum points.

Proof: Follows directly from Theorems 3.8 and 3.9 and the definitions of dilation and erosion.

Theorem 3.11: If a signal is nowhere equal to its dilation, $f(\mathbf{x}) < (f \mathbf{D} \tilde{N})(\mathbf{x})$ for all \mathbf{x} , then $f(\mathbf{x})$ is infinitely long and the dilation $(f \mathbf{D} \tilde{N})(\mathbf{x})$ does not contain a constant neighborhood. Similarly, if a signal is nowhere equal to its erosion, $f(\mathbf{x}) > (f \mathbf{K} \tilde{N})(\mathbf{x})$ for all \mathbf{x} , then $f(\mathbf{x})$ is infinitely long and the erosion $(f \mathbf{K} \tilde{N})(\mathbf{x})$ does not contain a constant neighborhood.

Proof: Finite length signals achieve a global maximum and minimum, and therefore must have points \mathbf{x}_1 and \mathbf{x}_2 where $f(\mathbf{x}_1) = (f \mathbf{D} \tilde{N})(\mathbf{x}_1)$ and $f(\mathbf{x}_2) =$

$(f \mathbf{E} \tilde{N})(\mathbf{x}_2)$. Therefore, if $f(\mathbf{x}) = (f \mathbf{D} \tilde{N})(\mathbf{x})$ for all \mathbf{x} , $f(\mathbf{x})$ is infinitely long, and if $f(\mathbf{x}) = (f \mathbf{E} \tilde{N})(\mathbf{x})$ for all \mathbf{x} , $f(\mathbf{x})$ is infinitely long. Suppose $(f \mathbf{D} \tilde{N})(\mathbf{x})$ contains a constant neighborhood centered at \mathbf{x}_1 . Each point \mathbf{y} in the structuring element N centered at \mathbf{x}_1 must satisfy $f(\mathbf{y}) = (f \mathbf{D} \tilde{N})(\mathbf{x}_1)$, but there must be at least one $f(\mathbf{y}) = (f \mathbf{D} \tilde{N})(\mathbf{x}_1)$ for the dilation to take on this value. So the dilation and the original signal are equal at this point, which contradicts the condition that $(f \mathbf{D} \tilde{N})(\mathbf{x}) = f(\mathbf{x})$. Therefore, a signal that is nowhere equal to its dilation cannot have a constant neighborhood in the dilation. A similar argument shows that a signal that is nowhere equal to its erosion cannot have a constant neighborhood in the erosion.

Q.E.D.

Theorem 3.12: If a signal is nowhere equal to its dilation and nowhere equal to its closing, then the closed signal must be infinitely long and n -monotonic everywhere.

Proof: By Theorem 3.11, a dilated signal that is nowhere equal to the original signal is infinitely long and does not contain a constant neighborhood. Then, by Theorem 3.10, the dilated signal must be either n -monotonic everywhere or have at least one global minimum point. The closed signal is the erosion of the dilation, and erosion of an n -monotonic signal yields a monotonic signal. Therefore, dilated signals that are n -monotonic everywhere lead to closed signals that are n -monotonic everywhere. Now assume the dilated signal has a global minimum value $(f \mathbf{D} \tilde{N})(\mathbf{x}_1)$ at the

point \mathbf{x}_1 . Every point \mathbf{y} in the structuring element N centered at \mathbf{x}_1 must satisfy $f(\mathbf{y}) \geq (f \mathbf{D} \tilde{N})(\mathbf{x}_1)$, but there must be at least one $f(\mathbf{y}) = (f \mathbf{D} \tilde{N})(\mathbf{x}_1)$ for the dilation to take on this value at \mathbf{x}_1 . The closed signal, however, has a constant neighborhood with value $(f \mathbf{D} \tilde{N})(\mathbf{x}_1)$ centered at \mathbf{x}_1 , because $(f \mathbf{D} \tilde{N})(\mathbf{x}_1)$ is a global minimum of the dilation, so the subsequent erosion expands it to a constant neighborhood. Therefore, the value of \mathbf{y} where $f(\mathbf{y}) = (f \mathbf{D} \tilde{N})(\mathbf{x}_1)$ is in this constant neighborhood and thus at that point the closed signal is equal to the original signal, which is a contradiction. Therefore if a signal is nowhere equal to its dilation and nowhere equal to its closing, the closed signal is infinitely long and n -monotonic everywhere.

Q.E.D.

Theorem 3.13: If a signal is nowhere equal to its erosion and nowhere equal to its opening, then the opened signal must be infinitely long and n -monotonic everywhere.

Proof: Analogous to the Proof of Theorem 3.12.

Midrange Filter

Using the above preliminary results, the complete root signal sets of the midrange, pseudomedian, and LOCO filters can now be found. The root signal set of the midrange filter for bounded signals consists only of constant signals. This is proven in Theorem 3.14 below.

Theorem 3.14: A bounded root signal of the midrange filter must be a constant signal.

Proof: Suppose $f(\mathbf{x})$ is a root signal of the midrange filter with structuring element N of length $|N| = n+1$; that is, $f(\mathbf{x}) = \text{midr}(f(\mathbf{x}); N)$. This proof will proceed by defining three possible cases for the root signals. The erosion and dilation of the signal may be equal at all points, or may be unequal at all points, or may be equal at some points and unequal at others. The case where erosion and dilation are equal at all points corresponds to the constant signal roots. The case where the erosion and dilation are unequal at all points can only give rise to infinite-length, unbounded root signals. All that is left to show is that there exist no root signals that have the erosion and dilation equal at some points and unequal at other points.

Case I: $(f \mathbf{E} \tilde{N})(\mathbf{x}) = (f \mathbf{D} \tilde{N})(\mathbf{x})$ for all \mathbf{x} .

The above assumption and the definition of the midrange filter together imply that $(f \mathbf{E} \tilde{N})(\mathbf{x}) = (f \mathbf{D} \tilde{N})(\mathbf{x}) = f(\mathbf{x})$ for all \mathbf{x} . Since the erosion and dilation are equal, the maximum and the minimum of $f(\mathbf{x})$ in the structuring element N are equal, and so all points in the structuring element must be equal. Since the erosion and dilation are equal for all \mathbf{x} , all points in $f(\mathbf{x})$ must then be equal. Therefore, $f(\mathbf{x})$ is a constant for all \mathbf{x} .

Case II: $(f \mathbf{E} \tilde{N})(\mathbf{x}) \neq (f \mathbf{D} \tilde{N})(\mathbf{x})$ for all \mathbf{x} .

Consider the case where there exists a point \mathbf{x}_1 such that $f(\mathbf{x}_1) = (f \mathbf{E} \tilde{N})(\mathbf{x}_1)$. Then, by the definition of the midrange filter, $(f \mathbf{D} \tilde{N})(\mathbf{x}_1) = f(\mathbf{x}_1)$, which is a contradiction of the assumption for Case

II. Similarly, the existence of a point \mathbf{x}_2 such that $f(\mathbf{x}_2) = (f \mathbf{D} \tilde{N})(\mathbf{x}_2)$ leads immediately to a contradiction.

Any root signal that satisfies Case II must then have no points where the original signal is equal to either the erosion or dilation. By Theorems 3.8 and 3.9 and their corollaries, $f(\mathbf{x})$ is then infinitely long and n -monotonic everywhere. For this to be true, the signal $f(\mathbf{x})$ cannot be bounded. Therefore, no bounded root signals of the midrange filter exist that satisfy $(f \mathbf{E} \tilde{N})(\mathbf{x}) = (f \mathbf{D} \tilde{N})(\mathbf{x})$ for all \mathbf{x} .

Case III: $(f \mathbf{E} \tilde{N})(\mathbf{x}) = (f \mathbf{D} \tilde{N})(\mathbf{x})$ for some \mathbf{x} and $(f \mathbf{E} \tilde{N})(\mathbf{x}) < (f \mathbf{D} \tilde{N})(\mathbf{x})$ for some \mathbf{x} .

Choose \mathbf{x}_1 such that $(f \mathbf{E} \tilde{N})(\mathbf{x}_1) = (f \mathbf{D} \tilde{N})(\mathbf{x}_1) = f(\mathbf{x}_1)$. Referring to Case I above, this means that all points in the structuring element $N_{\mathbf{x}_1}$ centered at \mathbf{x}_1 are equal to $f(\mathbf{x}_1)$. Now assume, without loss of generality, that there is a point \mathbf{x}_3 immediately to the right of this constant area where $f(\mathbf{x}_3) < f(\mathbf{x}_1)$ and $f(\mathbf{x}_3) < (f \mathbf{E} \tilde{N})(\mathbf{x}_3) < (f \mathbf{D} \tilde{N})(\mathbf{x}_3)$. The point \mathbf{x}_2 that is immediately to the right of \mathbf{x}_1 is in the constant area defined by the structuring element $N_{\mathbf{x}_1}$ centered at \mathbf{x}_1 . However, the point \mathbf{x}_3 lies within the structuring element $N_{\mathbf{x}_2}$ centered at \mathbf{x}_2 . (See Figure 3.2.) This means that $(f \mathbf{E} \tilde{N})(\mathbf{x}_2) = f(\mathbf{x}_3)$, but since all other points in $N_{\mathbf{x}_2}$ are equal to $f(\mathbf{x}_1)$, $(f \mathbf{D} \tilde{N})(\mathbf{x}_2) = f(\mathbf{x}_1)$. For $f(\mathbf{x})$ to be a root signal at \mathbf{x}_2 , $f(\mathbf{x}_2) = \frac{1}{2}[f(\mathbf{x}_1) + f(\mathbf{x}_3)]$, which is a contradiction. Hence, there are no root signals of the midrange filter that have $(f \mathbf{E} \tilde{N})(\mathbf{x}) = (f \mathbf{D} \tilde{N})(\mathbf{x})$ for some \mathbf{x} and $(f \mathbf{E} \tilde{N})(\mathbf{x}) < (f \mathbf{D} \tilde{N})(\mathbf{x})$ for some \mathbf{x} .

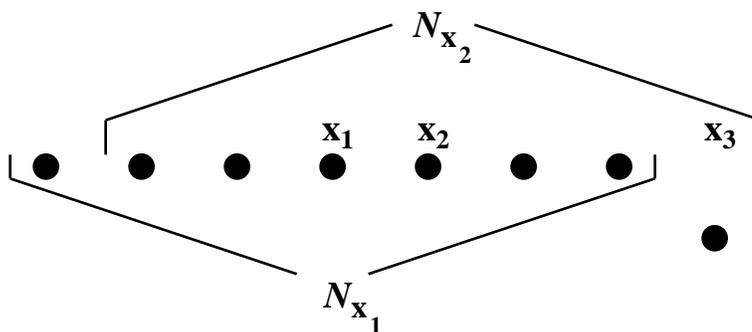


Figure 3.2. Illustration for Case III of the proof of Theorem 3.14.

Therefore, a bounded root signal $f(\mathbf{x})$ of the midrange filter must satisfy

$$(f \mathbf{K} \tilde{N})(\mathbf{x}) = (f \mathbf{D} \tilde{N})(\mathbf{x}) \text{ for all } \mathbf{x}, \text{ which means that it is a constant signal.}$$

Q.E.D.

Some infinite-length, unbounded signals are root signals of the midrange filter that satisfy the condition of Case II of the proof, that $(f \mathbf{K} \tilde{N})(\mathbf{x}) = (f \mathbf{D} \tilde{N})(\mathbf{x})$ for all \mathbf{x} . For instance, infinite-length signals with a constant rise over any run of just over half the length $\left(\frac{1}{2}(|N|+1) = \frac{1}{2}n+1 \right)$ of the structuring element are root signals of the midrange filter that satisfy Case II of the proof. It is physically impossible for an unbounded signal to exist, so these types of root signals are not important in real applications.

Unlike the compound morphological operators and their combinations, the midrange filter does not preserve edges. This behavior is indicated by the root signal set of the midrange filter, which does not include any signals with edges. The midrange filter severely distorts edges by taking the average of high and low points surrounding the edge.

Pseudomedian Filter

The deterministic behavior of the pseudomedian filter is quite different from that of the midrange filter. Theorem 3.15 below proves that the root signal set of the pseudomedian filter for bounded signals includes only signals that consist entirely of constant neighborhoods and edges.

Theorem 3.15: A bounded root signal of the pseudomedian filter consists only of constant neighborhoods and edges.

Proof: Suppose $f(\mathbf{x})$ is a root signal of the pseudomedian filter with structuring element N of length $|N| = n+1$; that is, $f(\mathbf{x}) = \text{pmed}(f(\mathbf{x}); N)$. This proof will proceed by defining three possible cases for the root signals. The opening and closing of the signal may be equal at all points, or may be unequal at all points, or may be equal at some points and unequal at others. The case where opening and closing are equal at all points corresponds to the root signals consisting of constant neighborhoods and edges. The other two cases show that there are not roots where the opening and closing are unequal at all points or where opening and closing are equal at some points and unequal at other points.

Case I: $f_N(\mathbf{x}) = f^N(\mathbf{x})$ for all \mathbf{x} .

Together, the definition of the pseudomedian filter and the above assumption imply that $f_N(\mathbf{x}) = f^N(\mathbf{x}) = f(\mathbf{x})$ for all \mathbf{x} . By Theorems 3.4, 3.6, and 3.7, this root signal $f(\mathbf{x})$ of the pseudomedian filter is everywhere

LOMO($n+2$); that is, $f(\mathbf{x})$ consists only of constant neighborhoods and edges.

Case II: $f_N(\mathbf{x}) = f^N(\mathbf{x})$ for all \mathbf{x} .

First, assume that there exists a point \mathbf{x}_1 such that $f(\mathbf{x}_1) = (f \boxtimes \tilde{N})(\mathbf{x}_1)$. By Theorem 3.2, $f_N(\mathbf{x}_1) = f(\mathbf{x}_1)$. Then, by the definition of the pseudomedian filter, $f_N(\mathbf{x}_1) = f^N(\mathbf{x}_1)$, which is a contradiction. Similarly, if there is a point \mathbf{x}_2 such that $f(\mathbf{x}_2) = (f \boxplus \tilde{N})(\mathbf{x}_2)$ then $f^N(\mathbf{x}_2) = f(\mathbf{x}_2)$. Again, $f_N(\mathbf{x}_2) = f^N(\mathbf{x}_2)$, which is a contradiction. Therefore, for Case II there can be no point \mathbf{x} where $f(\mathbf{x})$ is equal to the erosion, dilation, opening, or closing at \mathbf{x} . By Theorems 3.12 and 3.13, both $f_N(\mathbf{x})$ and $f^N(\mathbf{x})$ are then infinitely long and n -monotonic everywhere. By Theorem 3.2, $f^N(\mathbf{x}) \geq f_N(\mathbf{x})$, so the closing and opening must be both increasing or both decreasing, since any infinite n -monotonic increasing signal must be greater than any infinite n -monotonic decreasing signal as $\mathbf{x} \rightarrow +\infty$ and less than any infinite n -monotonic decreasing signal as $\mathbf{x} \rightarrow -\infty$. The average of two signals that are n -monotonic and either both increasing or both decreasing is a signal that is n -monotonic. Therefore, the pseudomedian of a signal $f(\mathbf{x})$ that satisfies $f_N(\mathbf{x}) = f^N(\mathbf{x})$ for all \mathbf{x} is an n -monotonic signal. However, if $f(\mathbf{x})$ is n -monotonic, then $f_N(\mathbf{x}) = f(\mathbf{x}) = f^N(\mathbf{x})$, which is a contradiction of $f_N(\mathbf{x}) = f^N(\mathbf{x})$. Therefore, there exist no root signals of the pseudomedian filter that satisfy Case II.

Case III: $f_N(\mathbf{x}) = f^N(\mathbf{x})$ for some \mathbf{x} and $f_N(\mathbf{x}) < f^N(\mathbf{x})$ for some \mathbf{x} .

For bounded signals, there exists a point \mathbf{x}_1 such that $f(\mathbf{x}_1) = (f \boxtimes W)(\mathbf{x}_1)$. This implies that $f(\mathbf{x}_1) = (f \boxtimes \tilde{N})(\mathbf{x}_1)$. By Theorem 3.2, $f_N(\mathbf{x}_1) = f(\mathbf{x}_1)$.

Then, by the definition of the pseudomedian filter, $f_N(\mathbf{x}_1) = f^N(\mathbf{x}_1)$. For this to be true, there must exist a point \mathbf{x}_2 within the structuring element $N_{\mathbf{x}_1}$ centered at \mathbf{x}_1 such that $f(\mathbf{x}_1) = (f \mathbf{E} \tilde{N})(\mathbf{x}_2)$. For this to be true without violating $f(\mathbf{x}_1) = (f \mathbf{E} W)(\mathbf{x}_1)$, all values of $f(\mathbf{x})$ within the structuring element $N_{\mathbf{x}_2}$ centered at \mathbf{x}_2 must be equal to $f(\mathbf{x}_1)$. Therefore, the values of $f(\mathbf{x})$ in $N_{\mathbf{x}_2}$ are a constant neighborhood of length at least $n+1$.

Within a constant neighborhood, $f_N(\mathbf{x}) = f^N(\mathbf{x})$. (See Case I.)

The constant neighborhood satisfies the first part of the assumption of Case III, so now an area where $f_N(\mathbf{x}) \neq f^N(\mathbf{x})$ is needed. But to achieve a root point where $f_N(\mathbf{x}) \neq f^N(\mathbf{x})$ means that $f_N(\mathbf{x}) < f(\mathbf{x}) < f^N(\mathbf{x})$ by Theorem 3.2. Choose a point \mathbf{x}_3 , where $f_N(\mathbf{x}_3) \neq f^N(\mathbf{x}_3)$, immediately to the right of an area of at least $n+1$ points where $f(\mathbf{x}) = f_N(\mathbf{x}) = f^N(\mathbf{x})$. Since $f_N(\mathbf{x}_3) < f(\mathbf{x}_3)$, the leftmost N containing \mathbf{x}_3 must contain a point with a value less than or equal to $f_N(\mathbf{x}_3)$. That is, there exists \mathbf{x}_L within the leftmost N containing \mathbf{x}_3 such that $f(\mathbf{x}_L) \leq f_N(\mathbf{x}_3) < f(\mathbf{x}_3)$. Similarly, there exists a point \mathbf{x}_h within the leftmost N containing \mathbf{x}_3 such that $f(\mathbf{x}_h) \geq f^N(\mathbf{x}_3) > f(\mathbf{x}_3)$. Note that \mathbf{x}_3 , \mathbf{x}_L , and \mathbf{x}_h are all within one structuring element N , and that $f_N(\mathbf{x}_h) = f^N(\mathbf{x}_h) = f(\mathbf{x}_h)$ and $f_N(\mathbf{x}_L) = f^N(\mathbf{x}_L) = f(\mathbf{x}_L)$.

1.) Assuming that \mathbf{x}_L is to the left of \mathbf{x}_h , all structuring elements N that contain \mathbf{x}_h also contain either \mathbf{x}_3 or \mathbf{x}_L or both. Since $f(\mathbf{x}_L) < f(\mathbf{x}_3)$, the highest value that $(f \mathbf{E} \tilde{N})(\mathbf{x}_h)$ can achieve is $f(\mathbf{x}_3)$. This means that $f_N(\mathbf{x}_h) \leq f(\mathbf{x}_3) < f(\mathbf{x}_h)$, which is a contradiction of $f_N(\mathbf{x}_h) = f^N(\mathbf{x}_h) = f(\mathbf{x}_h)$. (See Figure 3.3.)

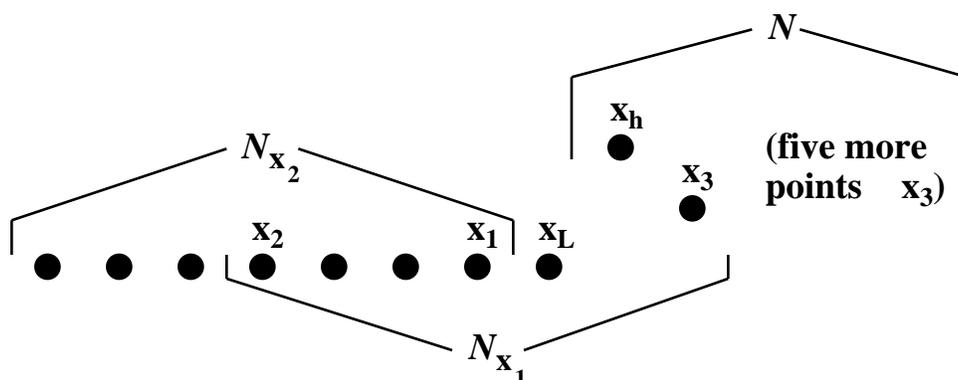


Figure 3.3. Illustration for Case III of the proof of Theorem 3.15.

2.) Assuming that \mathbf{x}_h is to the left of \mathbf{x}_L , all structuring elements N that contain \mathbf{x}_L also contain either \mathbf{x}_3 or \mathbf{x}_h or both. Since $f(\mathbf{x}_h) > f(\mathbf{x}_3)$, the lowest value that $(f \mathbf{D} \tilde{N})(\mathbf{x}_L)$ can achieve is $f(\mathbf{x}_3)$. This means that $f^N(\mathbf{x}_h) = f(\mathbf{x}_3) > f(\mathbf{x}_h)$, which is a contradiction of $f_N(\mathbf{x}_L) = f^N(\mathbf{x}_L) = f(\mathbf{x}_L)$. Therefore, a point satisfying the restrictions placed on \mathbf{x}_3 does not exist.

A similar argument may be made to show that a point \mathbf{x}_4 , with $f_N(\mathbf{x}_4) = f^N(\mathbf{x}_4)$, may not exist immediately to the *left* of an area of at least $n+1$ points where $f(\mathbf{x}) = f_N(\mathbf{x}) = f^N(\mathbf{x})$. Since a constant neighborhood is $n+1$ consecutive equal points that all have $f(\mathbf{x}) = f_N(\mathbf{x}) = f^N(\mathbf{x})$, the existence of a constant neighborhood in a signal means that the signal cannot have any points where $f_N(\mathbf{x}) \neq f^N(\mathbf{x})$ and therefore the signal must consist entirely of constant neighborhoods and edges. (See Case I.). If a root signal has a point where $f(\mathbf{x}_1) = (f \mathbf{K} W)(\mathbf{x}_1)$, then it must have a constant neighborhood and therefore must satisfy Case I only.

Therefore, a bounded root signal of the pseudomedian filter, $f(\mathbf{x})$, obeys the property $f_N(\mathbf{x}) = f^N(\mathbf{x})$ for all \mathbf{x} , and thus consists only of constant neighborhoods and edges.

Q.E.D.

Some infinite-length, unbounded signals do not satisfy the assumption in Case III of the proof that there exists a point \mathbf{x}_1 such that $f(\mathbf{x}_1) = (f \boxtimes W)(\mathbf{x}_1)$. For unbounded signals that do satisfy the assumption, the proof holds. The only obvious examples of signals that do not satisfy the assumption and are root signals of the pseudomedian filter are signals that are everywhere strictly increasing or everywhere strictly decreasing. It is physically impossible for an unbounded signal to exist, so these types of root signals are not important in real applications.

LOCO Filter

The root signal set of the LOCO filter is identical to that of the pseudomedian filter. This is shown in Theorem 3.16 below.

Theorem 3.16: A bounded root signal of the LOCO filter consists only of constant neighborhoods and edges.

Proof: Suppose $f(\mathbf{x})$ is a root signal of the LOCO filter with structuring element N of length $|N| = n+1$; that is, $f(\mathbf{x}) = \text{LOCO}(f(\mathbf{x}); N)$. This proof will proceed by defining three possible cases for the root signals. The OC and CO of the signal may be equal at all points, or may be unequal at all points, or

may be equal at some points and unequal at others. The case where OC and CO are equal at all points corresponds to the root signals of constant neighborhoods and edges. The other two cases show that there are not roots where OC and CO are unequal at all points or where OC and CO are equal at some points and unequal at other points.

Case I: $OC(f(\mathbf{x}); N) = CO(f(\mathbf{x}); N)$ for all \mathbf{x} .

The definition of the LOCO filter together with the above assumption implies that $OC(f(\mathbf{x}); N) = CO(f(\mathbf{x}); N) = f(\mathbf{x})$ for all \mathbf{x} . By Theorem 3.5, this signal must be a root signal of the median filter with window W . Then, by Theorems 3.6 and 3.7, this root signal $f(\mathbf{x})$ of the LOCO filter is everywhere LOMO($n+2$); that is, $f(\mathbf{x})$ consists only of constant neighborhoods and edges.

Case II: $OC(f(\mathbf{x}); N) \neq CO(f(\mathbf{x}); N)$ for all \mathbf{x} .

First, assume that there exists a point \mathbf{x}_1 such that $f(\mathbf{x}_1) = (f \boxminus \tilde{N})(\mathbf{x}_1)$. By Theorem 3.2, $OC(f; N)(\mathbf{x}_1) = f(\mathbf{x}_1)$. Then, by the definition of the LOCO filter, $OC(f; N)(\mathbf{x}_1) = CO(f; N)(\mathbf{x}_1)$, which is a contradiction. Similarly, if there is a point \mathbf{x}_2 such that $f(\mathbf{x}_2) = (f \boxplus \tilde{N})(\mathbf{x}_2)$ then $OC(f; N)(\mathbf{x}_2) = f(\mathbf{x}_2)$. Again, $OC(f; N)(\mathbf{x}_2) = CO(f; N)(\mathbf{x}_2)$, which is a contradiction. Therefore, for Case II there can be no point \mathbf{x} where $f(\mathbf{x})$ is equal to the erosion, dilation, OC, or CO at \mathbf{x} . By Theorems 3.12 and 3.13, both $f_N(\mathbf{x})$ and $f^N(\mathbf{x})$ are then infinitely long and n -monotonic everywhere, and so $OC(f(\mathbf{x}); N)$ and $CO(f(\mathbf{x}); N)$ are also infinitely long and n -monotonic everywhere. By Theorem 3.2, $CO(f(\mathbf{x}); N) \neq OC(f(\mathbf{x}); N)$, so the CO and OC must be both increasing or both decreasing, since any infinite n -

monotonic increasing signal must be greater than any infinite n -monotonic decreasing signal as $\mathbf{x} \rightarrow +\infty$ and less than any infinite n -monotonic decreasing signal as $\mathbf{x} \rightarrow -\infty$. The average of two signals that are n -monotonic and either both increasing or both decreasing is a signal that is n -monotonic. Therefore, the output of the LOCO filter on a signal $f(\mathbf{x})$ that satisfies $OC(f(\mathbf{x}); N) = CO(f(\mathbf{x}); N)$ for all \mathbf{x} is an n -monotonic signal. However, if $f(\mathbf{x})$ is n -monotonic, then $OC(f(\mathbf{x}); N) = f(\mathbf{x}) = CO(f(\mathbf{x}); N)$, which is a contradiction of $OC(f(\mathbf{x}); N) = CO(f(\mathbf{x}); N)$. Therefore, there exist no root signals of the LOCO filter that satisfy Case II.

Case III: $OC(f(\mathbf{x}); N) = CO(f(\mathbf{x}); N)$ for some \mathbf{x} and $OC(f(\mathbf{x}); N) < CO(f(\mathbf{x}); N)$ for some \mathbf{x} .

For bounded signals, there exists a point \mathbf{x}_1 such that $f(\mathbf{x}_1) = (f \boxminus W)(\mathbf{x}_1)$. By Theorem 3.3, $OC(f; N)(\mathbf{x}_1) = f(\mathbf{x}_1)$. Then, by the definition of the LOCO filter, $OC(f; N)(\mathbf{x}_1) = CO(f; N)(\mathbf{x}_1)$. For this to be true, there must exist a point \mathbf{x}_2 within the structuring element $N_{\mathbf{x}_1}$ centered at \mathbf{x}_1 such that $f(\mathbf{x}_1) = f^N(\mathbf{x}_2)$, and therefore there must exist a point \mathbf{x}_3 within the structuring element $N_{\mathbf{x}_2}$ centered at \mathbf{x}_2 such that $f(\mathbf{x}_1) = (f \boxminus \tilde{N})(\mathbf{x}_3)$. Note that $N_{\mathbf{x}_1}$ must overlap with $N_{\mathbf{x}_3}$. For $f(\mathbf{x}_1) = (f \boxminus \tilde{N})(\mathbf{x}_3)$ to be true without violating $f(\mathbf{x}_1) = (f \boxminus \tilde{N})(\mathbf{x}_1)$, all values of $f(\mathbf{x})$ within the structuring element $N_{\mathbf{x}_2}$ centered at \mathbf{x}_2 must be equal to $f(\mathbf{x}_1)$. Therefore, the values of $f(\mathbf{x})$ in $N_{\mathbf{x}_2}$ are a constant neighborhood of length at least $n+1$. Within a constant neighborhood, $OC(f(\mathbf{x}); N) = CO(f(\mathbf{x}); N)$. (See Case I.)

The constant neighborhood satisfies the first part of the assumption of Case III, so now an area where $OC(f(\mathbf{x}); N) = CO(f(\mathbf{x}); N)$ is needed. But to achieve a root point where $OC(f(\mathbf{x}); N) = CO(f(\mathbf{x}); N)$ means that $OC(f(\mathbf{x}); N) < CO(f(\mathbf{x}); N)$ by Theorems 3.1 and $f_N(\mathbf{x}) = OC(f(\mathbf{x}); N) < f(\mathbf{x}) < CO(f(\mathbf{x}); N) = f^N(\mathbf{x})$ by Theorem 3.3. The proof now proceeds exactly as for Case III in the proof of Theorem 3.15 for the pseudomedian filter. The existence of a constant neighborhood in a signal means that the signal cannot have any points where $f_N(\mathbf{x}) = f^N(\mathbf{x})$ and so there also cannot be any points where $OC(f(\mathbf{x}); N) = CO(f(\mathbf{x}); N)$. Therefore, the signal must consist entirely of constant neighborhoods and edges. (See Case I.) If a root signal has a point where $f(\mathbf{x}_1) = (f \boxtimes \tilde{N})(\mathbf{x}_1)$, then it must have a constant neighborhood and therefore must satisfy Case I only.

Therefore, a bounded root signal of the LOCO filter, $f(\mathbf{x})$, obeys the property $OC(f(\mathbf{x}); N) = CO(f(\mathbf{x}); N)$ for all \mathbf{x} , and thus consists only of constant neighborhoods and edges.

Q.E.D.

Despite the fact that the root signal sets of the pseudomedian and LOCO filters are identical, the results obtained from one iteration of these filters are usually quite different. Although both filters preserve edges and constant neighborhoods, the LOCO filter removes impulses completely, while the pseudomedian filter only reduces their amplitude by one-half. The root signal analysis does point out major differences among the midrange, pseudomedian, and LOCO filters, and indicates that the behavior of the pseudomedian and LOCO

filters in the presence of fast-fluctuating signals is much more stable than the behavior of the median filter. Although the deterministic properties of the filters give important insights into their behavior, finding the statistical properties will provide even more specific information about the similarities and differences among the filters.

3.2.3. Statistical Properties

The output statistics of the various linear combinations of morphological operators help describe the behavior of the filters in the presence of noisy signals. The statistical properties of the standard morphological operators were given in Chapter 2, and the biases that they introduce were demonstrated. This section shows that the averages of the complementary operators (the midrange, pseudomedian, and LOCO filters) do not introduce a statistical bias to noisy signals.

Midrange Filter

The statistical properties of the midrange filter are well-known in order statistics research [21, 22, 25, 35]. For a 1-D discrete input signal $f(\mathbf{x})$ with values that are i.i.d. random variables with distribution function $G(\mathbf{z})$, the distribution function $H_{\text{midr}}(\mathbf{z})$ of the output of the midrange filter with structuring element N of size $m = n+1$ is [22, 25, 35]:

$$H_{\text{midr}}(\mathbf{z}; m) = m \int_{-\infty}^{\mathbf{z}} [G(2\mathbf{z} - \mathbf{y}) - G(\mathbf{z})]^{m-1} g(\mathbf{y}) d\mathbf{y} \quad (3.4)$$

Differentiating equation (3.4) gives the probability density function $h_{\text{midr}}(\mathbf{z})$ of the output of the midrange filter:

$$h_{\text{midr}}(\mathbf{z}; m) = 2m(m-1) \int_0^1 g(\mathbf{z} + \mathbf{y})g(\mathbf{z} - \mathbf{y})|G(\mathbf{z} - \mathbf{y}) - G(\mathbf{z} + \mathbf{y})|^{m-2} d\mathbf{y} \quad (3.5)$$

The output p.d.f. of the midrange filter for uniformly distributed i.i.d. random input is shown in Figure 3.4 for structuring elements of length 3, 5, 7, and 11. Note that the output is very strongly weighted toward the center of the distribution, especially as the size of the structuring element increases. The median value of the input distribution and of all the midrange filter output distributions is 0.5, so the output of the midrange filter in this case is not biased with respect to the median. Since the mean of this distribution is equal to the median, the midrange filter output is also not biased with respect to the mean.

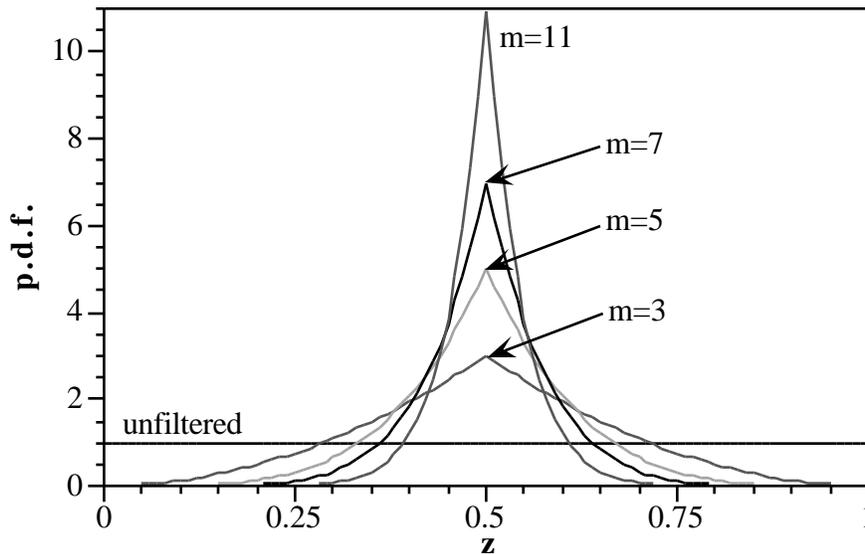


Figure 3.4. Density function of midrange filter output with uniformly distributed input.

Figure 3.5 below demonstrates the correspondence between the theoretical distributions and the actual results of the midrange filter applied to pseudorandom data generated by a computer. A signal containing 100,000 uniformly distributed pseudorandom numbers between 0 and 1 was generated using the pseudorandom number generator in the software package IDL, version 2.0.10 (Research Systems, Inc., Boulder, Colorado). The histogram of this signal using 100 bins is shown in Figure 3.5. The signal was then midrange filtered with structuring elements of lengths 3 and 7. The density functions of these filtered signals were then estimated from their histograms in the same manner, and the results were compared with histograms found from the theory in equation (3.5). The match between the theoretical and computer-simulated results, shown in Figure 3.5, is very good in this instance. Since the pseudorandom number generator used in these examples has a 32-bit seed number, the resulting sequence of pseudorandom numbers does not repeat for over 4 billion points (and therefore there is no repetition in the 100,000-point signal). The longer the (non-repeating) pseudorandom input signal is, the better the match to theory is. For this reason, the longest signal that could be filtered with a reasonable execution time was used for the simulation.

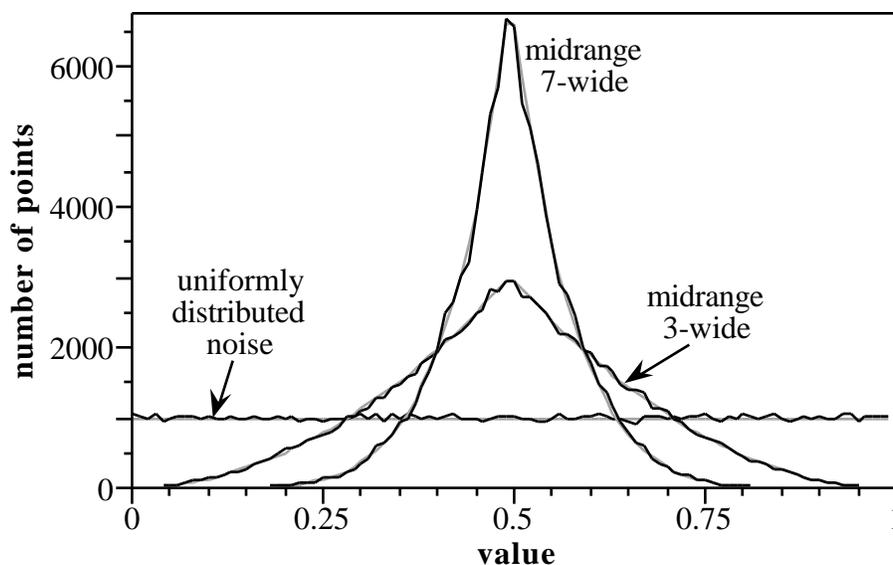


Figure 3.5. Theoretical and simulated results for the midrange filter acting on uniformly distributed input. 100,000-point simulation; 100 bins in histogram.

In two dimensions, the midrange filter output may be found in the same manner, since the midrange is a simple average of extreme order statistics and there are no geometrical effects to consider.

Pseudomedian Filter

The statistical properties of the pseudomedian filter are significantly more difficult to find than the properties of midrange filter. However, the properties of its constituent operators, opening and closing, are known [13]. Unfortunately, the opening and closing of the same signal are not independent, so finding an analytical expression for the output distribution of the pseudomedian filter for an arbitrary input distribution requires knowledge of the joint distribution of the opening and closing of the same signal.

However, a good match for the general shape of the output distribution of the pseudomedian filter is found using a formula that relies on the relationship between the distributions of dilation and closing and on the distribution of the midrange filter. The equation relating the output distributions of closing and dilation is:

$$h_{\text{close}}(\mathbf{z}; m) = m h_{\text{dilate}}(\mathbf{z}; m) - (m - 1)h_{\text{dilate}}(\mathbf{z}; m + 1) \quad (3.6)$$

I have found that a good empirical match for the general shape of output distribution of the pseudomedian filter can be achieved by using the same relationship as in (3.6) for pseudomedian and midrange distributions, and averaging the result with the output distribution of the midrange filter. This formula is given in equation (3.7) below.

$$\begin{aligned} h_p(\mathbf{z}; m) &= \frac{1}{2} [m h_{\text{midr}}(\mathbf{z}; m) - (m - 1)h_{\text{midr}}(\mathbf{z}; m + 1)] + \frac{1}{2} h_{\text{midr}}(\mathbf{z}; m) \\ &= \frac{1}{2} (m + 1)h_{\text{midr}}(\mathbf{z}; m) - \frac{1}{2} (m - 1)h_{\text{midr}}(\mathbf{z}; m + 1) \end{aligned} \quad (3.7)$$

The function h_p given in equation (3.7) is an excellent fit for the output distribution of the pseudomedian filter for uniformly distributed input. For other input distributions, such as the normal distribution, equation (3.7) does not exactly match the output distribution. However, a slight amplitude adjustment to this equation yields an excellent approximation of the output distribution. This adjustment is given in equation (3.8) below.

$$h_{\text{pmed}}(\mathbf{z}; m) = \frac{1}{p} (h_p(\mathbf{z}; m))^k \quad (3.8)$$

where k is a constant ($1 < k < 2$) which takes one of several forms according to the type of input distribution and may be a function of m , and p is a normalization constant defined by equation (3.9) below.

$$p = \int_{-}^{+} [h_p(\mathbf{z}; m)]^k d\mathbf{z} \quad (3.9)$$

For uniformly distributed input, $k = 1$ and therefore $p = 1$. This also holds for input density functions that are piecewise constant. For normally distributed input, $k = 2 - \frac{2}{m}$.

Figures 3.6 and 3.7 below compare the results of 100,000-point simulations of pseudomedian-filtered noise signals to the theoretical approximations given in equations (3.7) and (3.8). The histogram for the uniformly distributed input (Figure 3.6) has 100 bins between 0 and 1; the histogram for the normally distributed input (Figure 3.7) has 160 bins between -4 and 4 , or 20 bins per unit. The match between the data and theory is very close in all cases.

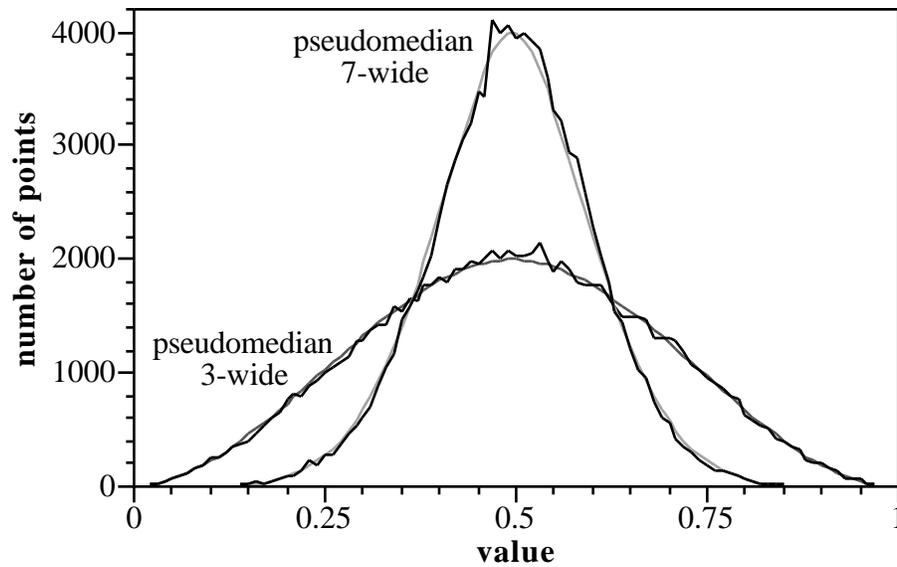


Figure 3.6. Theoretical and simulated results for the pseudomedian filter acting on uniformly distributed input. 100,000-point simulation; 100 bins in histogram.

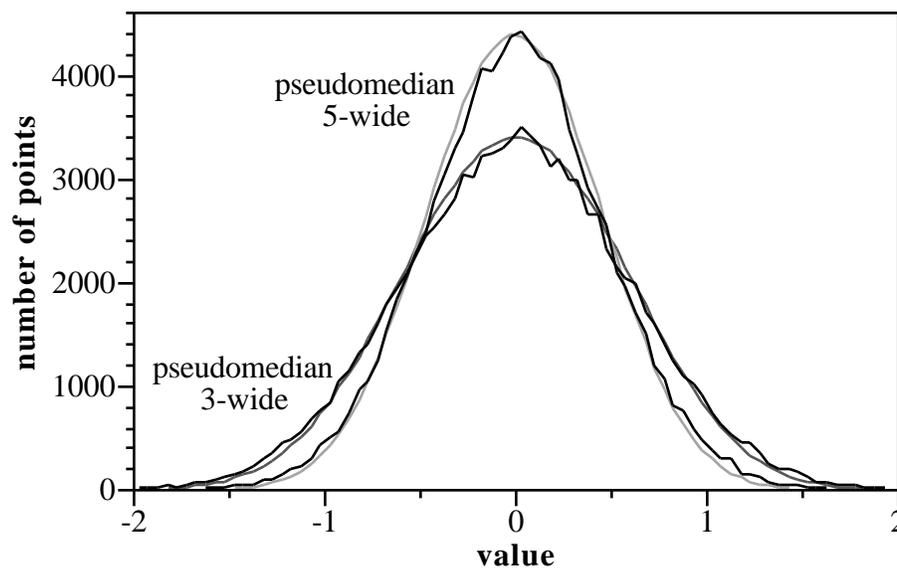


Figure 3.7. Theoretical and simulated results for the pseudomedian filter acting on normally distributed input. 100,000-point simulation; 20 bins per unit.

LOCO Filter

The statistical properties of the LOCO filter are also difficult to describe analytically, since the OC and CO are not independent and their joint distribution is not known. However, the output p.d.f. of the LOCO filter has the same basic shape as that of the pseudomedian filter for any input; the LOCO filter output is usually just slightly more peaked. The proposed estimation of the output density function of the LOCO filter for random i.i.d. input is:

$$h_{\text{loco}}(\mathbf{z}; m) = \frac{1}{p_{\text{med}}} [h_{\text{pmed}}(\mathbf{z}; m)]^p \quad (3.10)$$

where p is a constant ($1 < p < 2$) which depends upon the input distribution and may be a function of m , and p_{med} is a normalization constant defined by equation (3.11) below.

$$p_{\text{med}} = \int [h_{\text{pmed}}(\mathbf{z}; m)]^p d\mathbf{z} \quad (3.11)$$

For large structuring element sizes, the output distributions of the pseudomedian and LOCO filter for uniformly distributed input are virtually identical. The best fit to the simulations for uniformly distributed input are found when $p = \frac{m+1}{m}$. For normally distributed input, the relative difference between the pseudomedian and LOCO filters is constant for all structuring element sizes, with $p = \frac{4}{3}$.

Figures 3.8 and 3.9 below compare the results of 100,000-point simulations of LOCO-filtered noise signals to the theoretical approximation given in equation (3.10). The histogram for the uniformly distributed input (Figure 3.8) has 100 bins between 0 and 1; the histogram for the normally distributed input

(Figure 3.9) has 160 bins between -4 and 4 , or 20 bins per unit. The match between the data and theory is very close in all cases.

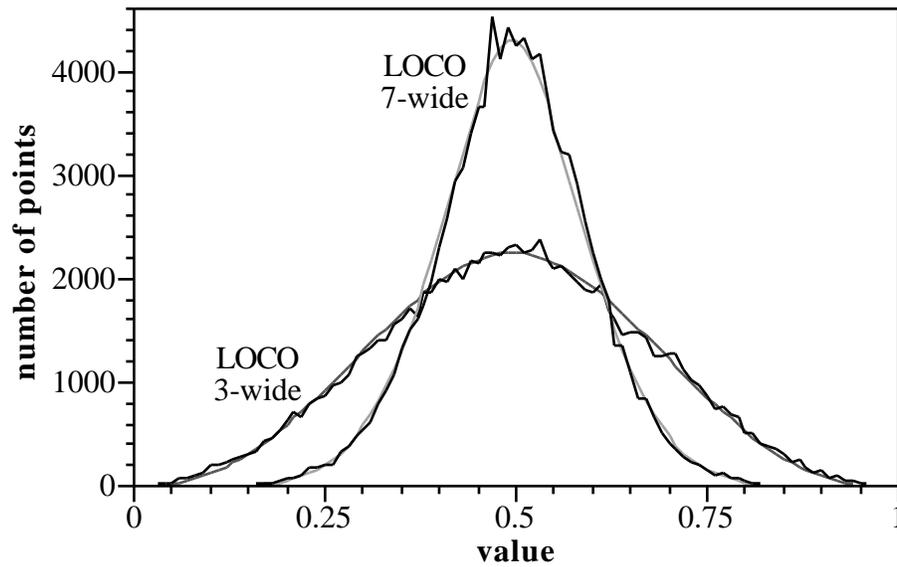


Figure 3.8. Theoretical and simulated results for the LOCO filter acting on uniformly distributed input. 100,000-point simulation; 100 bins in histogram.

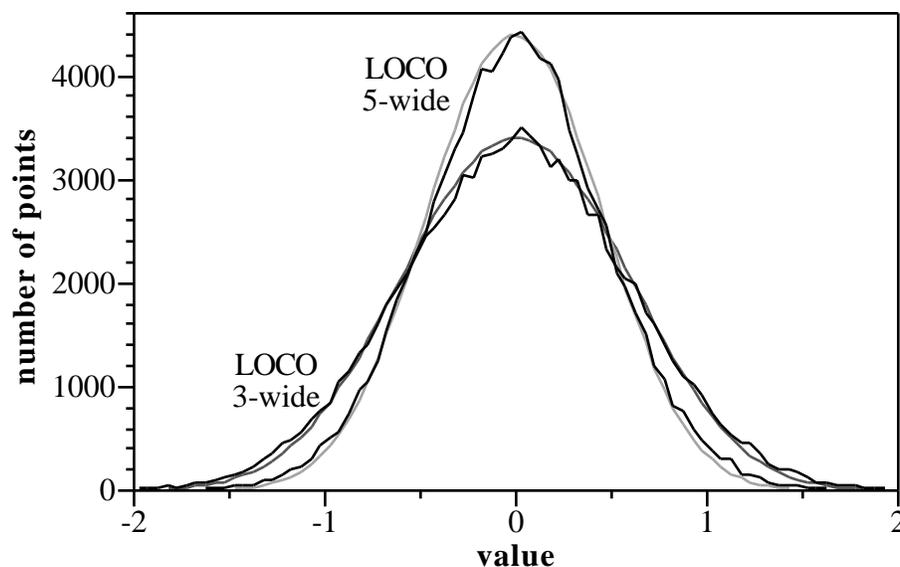


Figure 3.9. Theoretical and simulated results for the LOCO filter acting on normally distributed input. 100,000-point simulation; 20 bins per unit.

3.3. VALUE-AND-CRITERION FILTER STRUCTURE

3.3.1. Introduction

The value-and-criterion filter structure, defined here for the first time, is derived from the structure of the singly-compound morphological operations (opening and closing). This new structure is much more flexible than the morphological structure, because it allows operations other than just the minimum and maximum to be performed on the data. Another advantage of the value-and-criterion structure is that it is not just a simple sequential structure where the output of the second stage of the filter is based only upon the output of the first

stage. A value-and-criterion filter has two parallel operations in its first stage, one of which is used as a criterion for selecting the final output value.

The value-and-criterion filter structure is related to several nonlinear filters proposed previously to solve the problem of edge-preserving smoothing [36-41]. These filters have a set of small neighborhoods or “subwindows” within an overall filter window. Typically, the filters find the most homogeneous subwindow (by various measures of homogeneity) within the overall window and perform low-pass filtering (usually the arithmetic mean) on that subwindow to give the output for the overall window. The value-and-criterion structure generalizes these filters by providing a complete, regular subwindow structure that actually defines the overall window.

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. 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 3×3 square, opening and closing draw their values from the 5×5 window formed by the union of translations of a 3×3 subwindow to each position of a 3×3 structuring element. There are then nine 3×3 subwindows in the 5×5 window, and these nine subwindows are all the possible different 3×3 subwindows in the window. This is illustrated in Figure 3.10 below. In most of the edge-preserving smoothing filter structures described previously, there are fewer subwindows, and the subwindows are often irregularly shaped or differently sized.

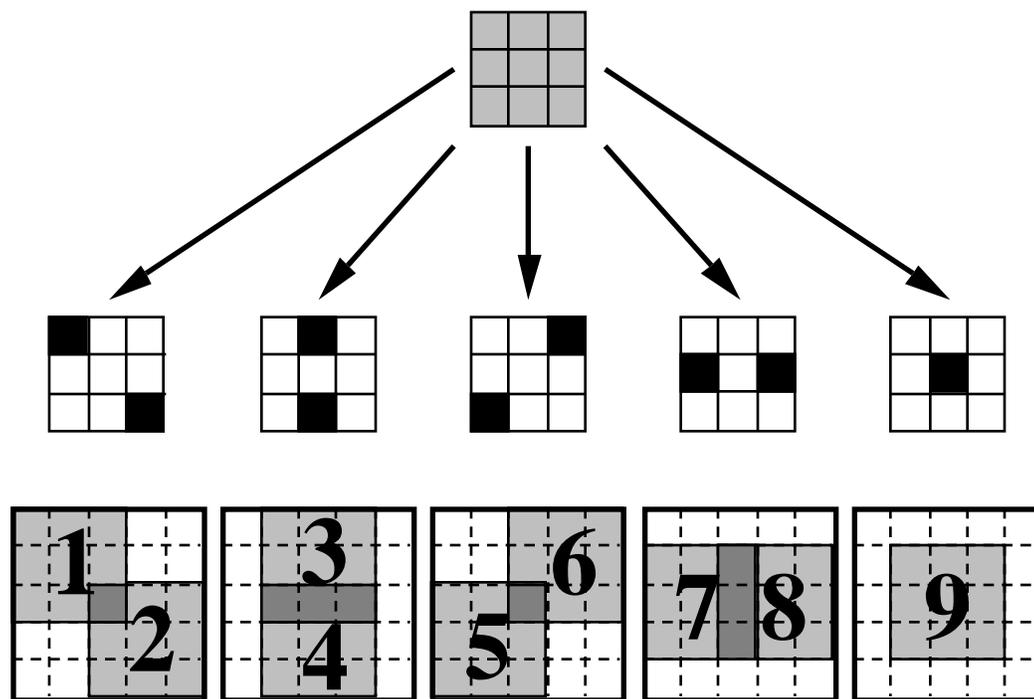


Figure 3.10. Structure of morphological opening and closing (and value-and-criterion structure) with a 3×3 structuring element.

The number of “subwindows” within the overall window grows quickly as the size of the structuring element increases. For an $n \times n$ square structuring element, there are n^2 subwindows within a square overall window of size $(2n-1) \times (2n-1)$. This tremendous growth in the number of subwindows is a key difference between the previous edge-preserving smoothing filter structures and the morphology-based value-and-criterion filter structure. Because the previous filter structures were not as computationally efficient as the value-and-criterion structure, they typically only considered several (often four, but sometimes as many as six or seven) subwindows, and the subwindows were only defined for a particular window shape. The value-and-criterion structure, however, identifies

all the subwindows of the same shape and size within an overall window. (In fact, the overall window of a value-and-criterion filter is actually defined by the subwindows; the “effective” window examples for the compound morphological operations in Chapter 2 illustrate this process.) Yet because of their more efficient definition, value-and-criterion filters execute considerably faster than most implementations of related smoothing filters.

The class of edge-preserving smoothing filters considered here first computes some measure of homogeneity on the subwindows to use as a criterion for selecting a subwindow, and then computes 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 key differences between the value-and-criterion filter structure and the morphological structure are the number and type of operations allowed in the various filtering stages. The value-and-criterion filter structure uses the complete, regular subwindow structure of opening and closing, but the first filtering stage performed on those subwindows is a set of two operations in parallel. 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. In addition to providing a better design structure for edge-preserving filters, the value-and-criterion structure is a uniform method for describing diverse types of nonlinear filters using a single mathematical formalism.

3.3.2. Definition

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 signal or image $f(\mathbf{x})$. Both V and C operate over the same structuring element N . Another operator, the “selection” function S , acts on the output of C . The selection operator acts over \tilde{N} , the 180° rotation of structuring element N . (Recall 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 used 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 signal or 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:

$$v(\mathbf{x}) = V[f(\mathbf{x}); N] \quad (3.12)$$

$$c(\mathbf{x}) = C[f(\mathbf{x}); N] \quad (3.13)$$

$$g(\mathbf{x}) = v\left(\left\{\mathbf{x} : \mathbf{x} \in \tilde{N}_{\mathbf{x}}; c(\mathbf{x}) = S[c(\mathbf{x}); \tilde{N}]\right\}\right) \quad (3.14)$$

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} may be chosen by the selection operator if two or more values of the criterion function in the window are equal. 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. The examples in this work 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 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 $(N \ominus \tilde{N})$ which has the smallest variance. (See Figure 3.10.) This is the Mean of Least Variance (MLV) filter, which I have developed as a significant improvement over previous edge-preserving smoothing filters. The properties of the MLV filter are explored in detail in the next section.

The value-and-criterion filter structure is powerful 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 common technique 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 provides an efficient implementation that avoids the redundant computation of the earlier edge-preserving filters. The structure also outlines a potentially parallel implementation of the filters.

3.3.3. 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:

$$v(\mathbf{x}) = \frac{1}{|N|} \sum_{\mathbf{y} \in N_{\mathbf{x}}} f(\mathbf{y}) \quad (3.15)$$

$$c(\mathbf{x}) = \frac{1}{|N|} \sum_{\mathbf{y} \in N_{\mathbf{x}}} |f(\mathbf{y}) - v(\mathbf{x})|^2 \quad (3.16)$$

$$\text{MLV}[f(\mathbf{x}); N] = v\left(\left\{\mathbf{x} : \mathbf{x} \in \tilde{N}_{\mathbf{x}}; c(\mathbf{x}) = \min[c(\mathbf{y}) : \mathbf{y} \in \tilde{N}_{\mathbf{x}}]\right\}\right) \quad (3.17)$$

Note that the $\frac{1}{|N|}$ term in the definition of $c(\mathbf{x})$ 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, the value of $v(\mathbf{x})$ closest to $f(\mathbf{x})$ is chosen, with ties settled 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 good for restoring exact gray levels of the original image from a noisy or distorted version. An example of the edge-enhancing characteristic of the MLV filter is shown in Figure 3.11 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 ($|N| = 5$).

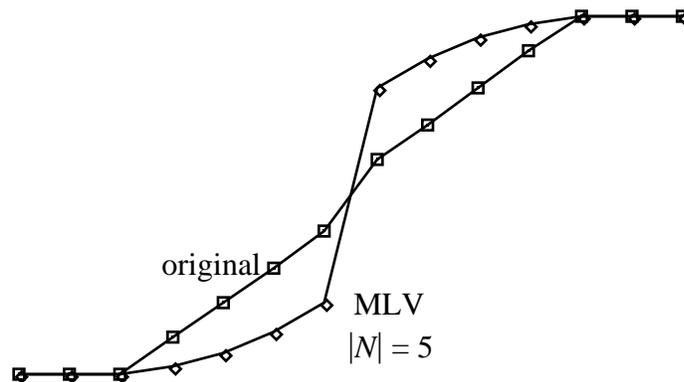


Figure 3.11. Example of the edge-enhancing characteristic of the MLV filter in one dimension. Original signal (\square); MLV filtered with $|N| = 5$ (\diamond).

3.3.4. Deterministic Properties

The value-and-criterion structure embraces such a wide variety of nonlinear filters that little can be said about the deterministic properties of value-and-criterion filters in general. However, the same types of deterministic analyses performed for the morphological and median filters can be performed for particular value-and-criterion filters. The deterministic properties of the MLV filter give some indication of how value-and-criterion filters that perform smoothing over homogeneous neighborhoods behave.

Since the MLV filter only alters signal values by averaging, the root signal set of the averaging filter should be a subset of the root signal set of the MLV filter. The root signal set of a 1-D averaging filter with window N consists only of constant signals and infinite-length constant-slope signals, and these are trivially roots of the 1-D MLV filter also. However, the MLV filter is also able to preserve perfect step edges between long 1-D constant regions, since it will take the average of one of its constant subwindows (zero variance) and not of the subwindows with the step. The constant regions must be at least as long as one subwindow (structuring element), $|N|$. An example illustrating how the MLV filter preserves perfect edges between constant neighborhoods is shown in Figure 3.12 below. The structuring element width is 3 in this 1-D example, and so the overall window width is 5. At position \hat{A} , the three subwindows within the overall window are A , B , and C . The variance (σ^2) of structuring element A is zero, but the variances of B and C are 2. Therefore, the subwindow of least variance is A , and the mean (μ) of A is zero, so the value at point \hat{A} is unchanged.

Similarly, at position ζ , structuring element D has zero variance, and so its mean is the output of the MLV filter. The value at point ζ is therefore also unchanged.

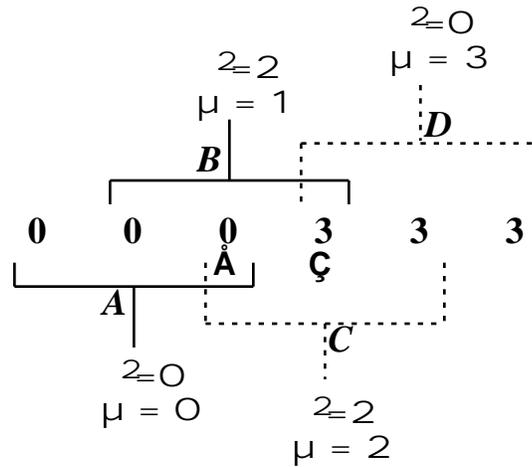


Figure 3.12. Example section of piecewise constant root signal of the MLV filter with a structuring element of width 3.

Signals where every point belongs to a constant neighborhood are called piecewise constant (PICO) signals. PICO signals and infinite-length constant ramp signals are roots of the MLV filter. No other types of edges or ramps are left unchanged by the MLV filter; however, a particular type of oscillation can be a root of the MLV filter under the tie-breaking scheme defined earlier. An example for $|\mathcal{N}| = 3$ is illustrated in Figure 3.13 below.

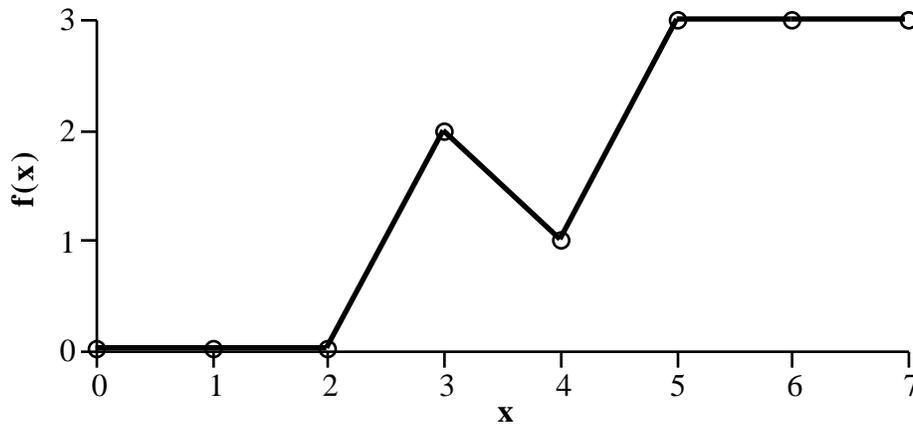


Figure 3.13. Example of oscillatory root of MLV filter with $|N| = 3$ under previously-described tie-breaking scheme.

In this example, the signal at $x = 3$ and at $x = 4$ has sample variance of 2, while the mean at $x = 3$ is 1 and the mean at $x = 4$ is 2. Since the variance is greater than 2 at $x = 2$ and $x = 5$, there is a tie for the least variance at both $x = 3$ and $x = 4$. Since ties are settled in favor of the value nearest the original value, this signal is a root of the MLV filter. Other tie-breaking schemes, however, would likely prevent this signal from being a root signal.

It is important to note that this signal relies on an exact tie for the least variance at two or more points in the signal. Small perturbations in the signal would cause these variances not to be equal, and the oscillation would be removed by the MLV filter. This type of root with an oscillation is an *unstable* root of the MLV filter under the given tie-breaking conditions; input signals within a small distance of an oscillatory root map to output signals much further away from the original signal after MLV filtering. By contrast, the piecewise constant roots are

stable; small changes in them lead to an MLV output that is very near the original signal. A constant ramp signal is a root signal under most tie-breaking schemes, but perturbations cause the MLV filter to break down such a signal into “stair steps.” Therefore, the root signal set of the MLV filter consists of a set of stable, universal roots, piecewise constant signals, and a set of unstable roots that depend on the tie-breaking scheme that is adopted. Under some schemes, the unstable roots may consist only of constant-slope signals, while in other cases certain types of oscillations may also be roots.

In addition to enhancing blurred edges, the MLV filter tends to sharpen any features in an image that conform to its structuring element shape. This is similar to the shape behavior of the morphological filters, except that the MLV filter actually enhances blurred features instead of merely preserving them. This is illustrated in Figure 3.14 below. The original image is a 90° corner (black, value 0) that is surrounded by a band of gray (value 127) on a background of white (value 255). Since the gray is slightly closer to black than white, the MLV filter with a 3×3 structuring element sharpens the edges of the corner by making the gray band much darker. The original image is a root of morphological opening and closing, and so the pseudomedian, OC, CO, and LOCO filters all do not change the image at all. For comparison, the results of 5×5 median filtering and 5×5 average filtering are also shown. The median filter rounds off the corner and maintains the gray band, while the average filter blurs the corner.

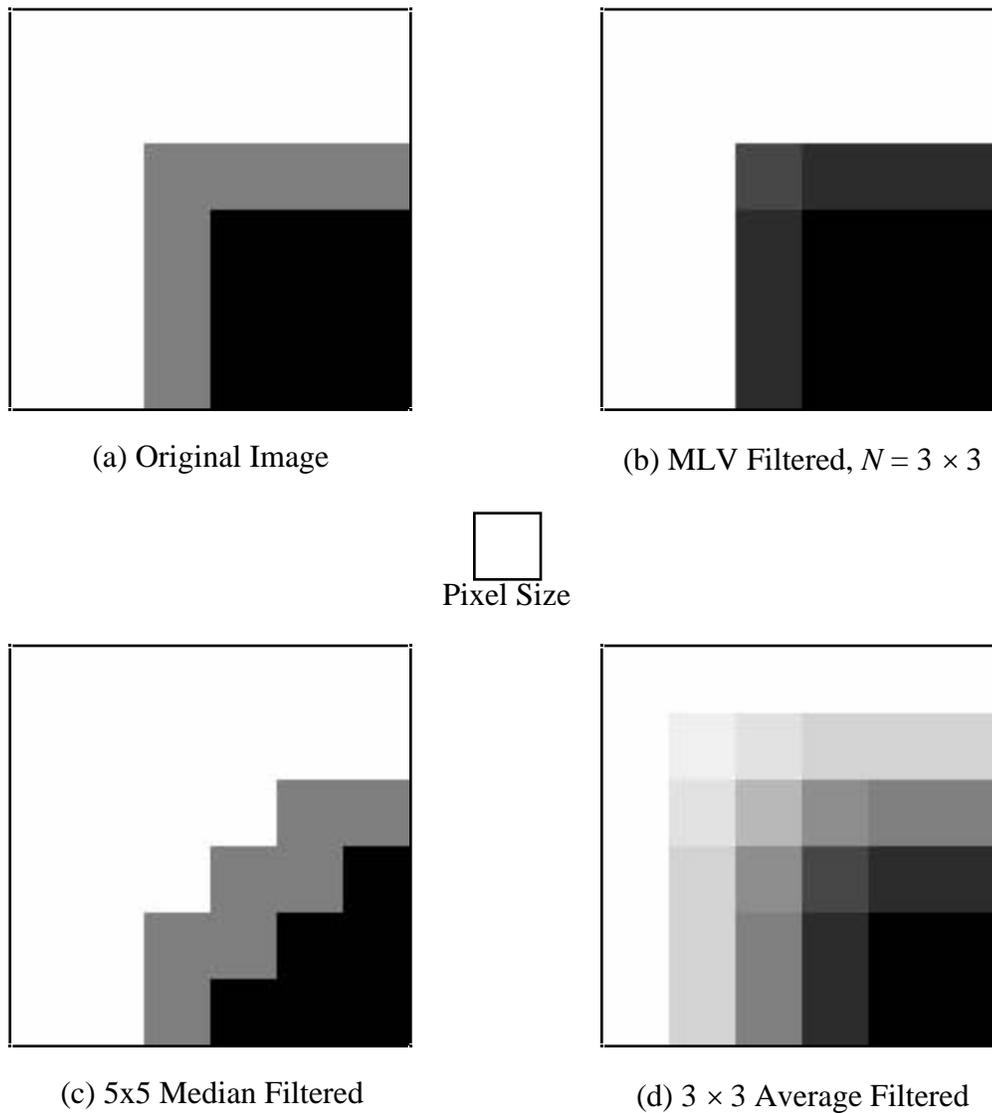


Figure 3.14. Response of MLV (b), median (c), and average (d) filters to a 90° corner (a). Each image section is 6×6 pixels square.

The main difference between the MLV filter and the morphological filters illustrated by the deterministic properties is that the MLV filter sharpens ramp edges while the morphological filters preserve them without sharpening them. The “stable” roots of the MLV filter are piecewise constant, and so they have

only perfect edges. However, since the main filtering operation of the MLV filter is averaging, it is effective at removing noise in signals and images also. In fact, the next section gives the statistical properties of the MLV filter, which are very similar to a simple averaging filter. This combination of well-understood linear-type noise smoothing with edge enhancement is potentially of great use in many applications.

3.3.5. Statistical Properties

As is true for the deterministic properties, little can be said about the statistical properties of the value-and-criterion filters in general. Once again, the MLV filter will be used as an instructive example. Since its main filtering operation is simple averaging, its statistical properties are likely similar to those of the averaging filter.

The averaging filter is simply a sliding mean operation, so for i.i.d. random input its properties are the same as the sample mean. Although the properties of the sample mean have been studied extensively, there is no simple expression for the output distribution for an arbitrary input distribution. Closed-form results for the sample mean are known for a normal input distribution; for other distributions, the output is approximated for larger sample sizes by the Central Limit Theorem [22]. The MLV filter takes the mean over the subwindow (structuring element) with the smallest sample variance; however, despite the fact that the variance is “small,” the points in this structuring element all come from the same distribution as every other point in the signal.

Consider an i.i.d. random input signal with a normal distribution of mean μ and variance σ^2 . This distribution is denoted by (μ, σ^2) . The distribution of the sample mean of n samples of an (μ, σ^2) distribution is normally distributed with mean μ and variance $\frac{\sigma^2}{n}$ [22]. Figure 3.15 shows the theoretical and experimental results for averaging filtering acting on normally distributed input. The input probability density function is shown, both in theory and for a 100,000-point simulation with 100 histogram bins. The theoretical and simulation results for averaging filtering with 3-, 5-, and 7-wide windows are also shown. This figure demonstrates the quality of the match between the theory and simulation for the derived relationships. The numerical results for both the theory and the simulation are tabulated in Table 3.1.

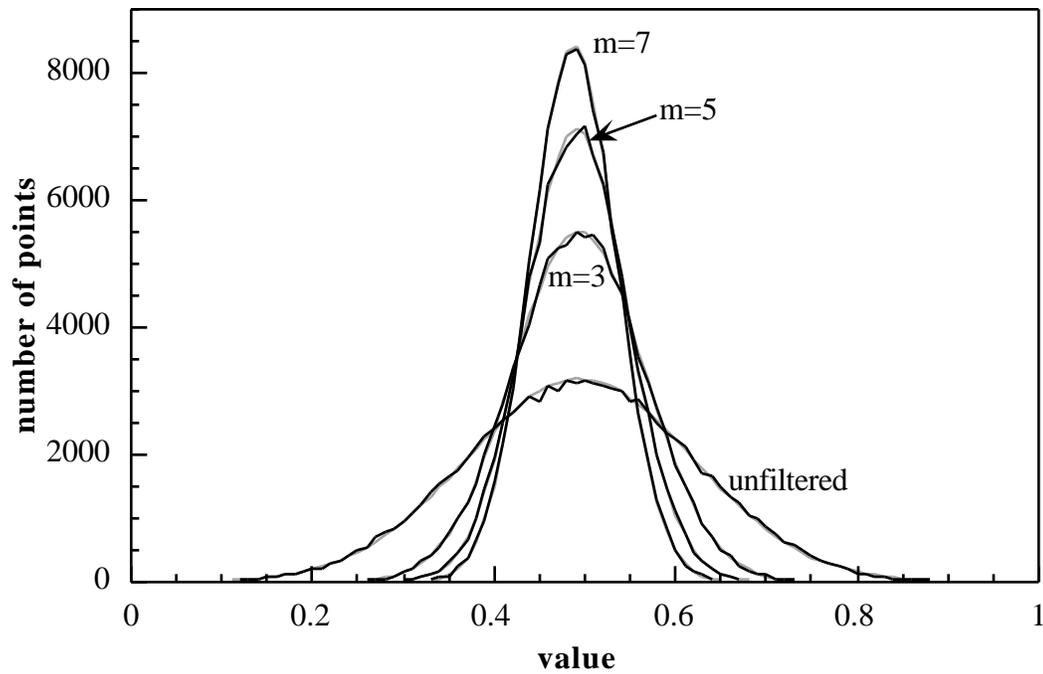


Figure 3.15. Theoretical and simulated results for the averaging filter acting on normally distributed input. 100,000-point simulation; 100-bin histogram.

Table 3.1. Theoretical and simulated variance ratios for the averaging filter for odd window sizes 3 to 11 and 25, 51, and 101.

size (m)	$\frac{\sigma^2}{m}$ (observed)	$\frac{\sigma^2}{\frac{m}{2}}$ (observed)	$\frac{\sigma^2}{\frac{m}{2}}$ (theory)	percentage deviation
1	1.001	(1.000)	1.000	—
3	0.333	3.01	3.00	+0.20%
5	0.201	4.99	5.00	-0.25%
7	0.1440	6.96	7.00	-0.66%
9	0.1120	8.94	9.00	-0.70%
11	0.0917	10.91	11.00	-0.79%
25	0.0401	25.0	25.0	-0.11%
51	0.01914	52.3	51.0	+2.48%
101	0.00958	104.5	101.0	+3.46%

For the 1-D MLV filter with structuring element length $|N| = m$, the overall window size of the filter is $2m-1$, but the mean is taken only over the m contiguous elements in the structuring element with the smallest variance. If this selection based on the variance did not occur, the properties would be the same as for the sample mean with m samples. However, since the MLV filter tends to create sharp “jumps” in a signal, one might expect that the MLV filter does not reduce the variance of the input quite as much as the averaging filter does for m samples.

Two limiting cases will help determine an appropriate relationship between the statistical properties of the MLV filter and the averaging filter. The first is that for $m = 1$, the input distribution is unchanged by filtering because the sample mean of a single value is the value itself. The second limiting case is that as $m \rightarrow \infty$, the properties of the MLV filter should approach those of the averaging filter of width m . This is because the sample variances of the various subwindows should be approximately equal for a large number of samples, and the sample means of the subwindows should also be very close to each other.

The relationship between the MLV filter and averaging filter for a normally distributed input is explored in Tables 3.1 and 3.2. Table 3.1 gives the sample variance of 100,000-point simulation results for the averaging filter of odd window sizes (m) from 3 to 11 and 25, 51, and 101. The ratio of the variance for each filter size to the original variance is computed and compared to the variance ratios predicted by the theory. The agreement is excellent, with less than a 4% error in the variance ratios for all filter sizes. Table 3.2 shows the same data for the MLV filter of various structuring element sizes from 3 to 251. These data will be used to derive an approximate theoretical relationship for the MLV filter, and the values computed from this theory are shown in Table 3.2 also.

Since the primary operation of the MLV filter is a sample mean, it is reasonable to postulate that for normally distributed i.i.d. input (μ, σ^2) , the output distribution of the MLV filter is normally distributed with the same mean μ and a smaller variance $\frac{\sigma^2}{m}$, where m is a function of the structuring element size m . Since it often preserves abrupt changes in the signal, an MLV filter is expected to reduce the variance less than the averaging filter of the same size as

the structuring element. Therefore, the relation $m \approx m$ is expected. The earlier two observations about the MLV filter indicate that $m \approx 1$ when $m = 1$ and that $m \approx m$ as $m \rightarrow \infty$. One way to achieve these constraints is to postulate that $m \approx [1 - v(m)](m - 1) + 1$, where $[1 - v(m)]$ represents a slope function with $v(m) \rightarrow 0$ as $m \rightarrow \infty$ and the slope is taken from $m = 1$ because that is the degenerate case rather than $m = 0$.

To find an approximation for $v(m)$, a 100,000-point signal of normally distributed pseudorandom noise was MLV filtered with various structuring element widths, and the corresponding estimate of m was recorded by dividing the sample variance of the unfiltered signal by that of the filtered signal (see Table 3.2). From this estimate and the known value of m , the value of $v(m)$ is found, and the plot of $v(m)$ for all the values of m in Table 3.2 is shown in Figure 3.16 below. Since $v(m) \rightarrow 0$ for all m , and $v(m)$ is asymptotically zero as m becomes very large, a likely candidate to fit this curve is a power function; that is, $v(m) \approx cm^k$, where c and k are constants. The least squares fit for this function, as shown in Figure 3.16, is for $c \approx \frac{2}{3}$ and $k \approx -\frac{1}{2}$. However, this fit is significantly different from the simulation for small m , ($m < 10$) which is troublesome since these are the most common sizes for the filter. A better fit, especially for smaller window sizes, is found with an approximation of the form $v(m) \approx c(1 + m^k)^{-1}$.

The least squares fit for this form is for $c = 0.8754$ and $k = 0.5357$, with a squared correlation coefficient $r^2 = 0.9674$. This can be approximated by the function $v(m) \approx \frac{7}{8}(1 + \sqrt{m})^{-1}$, although for the smallest (and most common) values of m , $v(m) \approx \frac{4}{5}(1 + \sqrt{m})^{-1}$ is an even better approximation. This approximation is plotted in Figure 3.16 along with the experimental data and the original power

function fit. This final expression for $v(m)$ results in the following approximation for m :

$$m = m - \frac{4}{5}(1 - \sqrt{m}) \quad (3.18)$$

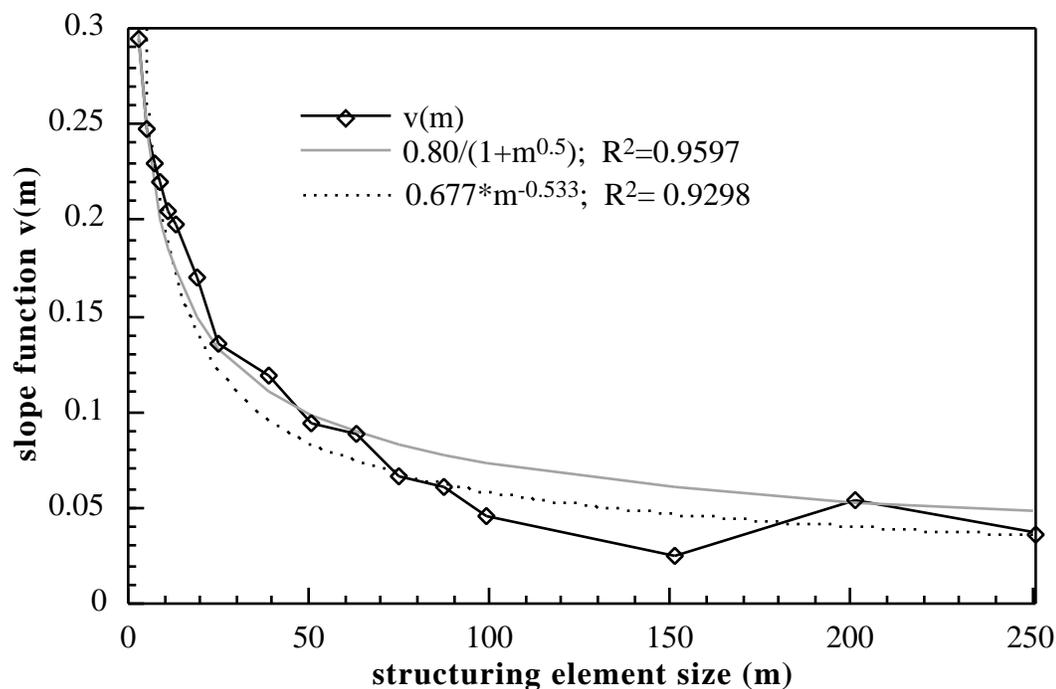


Figure 3.16. Data for slope function $v(m)$, curve fit of the form cm^k , and the approximation $v(m) = \frac{4}{5}(1 + \sqrt{m})^{-1}$.

This expression (3.18) is used to determine the “predicted” values given in Table 3.2, and the percentage difference between the observed value of m and the predicted value from (3.18) is less than $\pm 4\%$ in all cases. This closeness of fit is similar to that of the averaging filter, as shown in Table 3.1.

Histogram plots for the MLV filter with various structuring element sizes for the 100,000-point simulation are shown in Figures 3.17 and 3.18 below. Also

plotted are the expected histograms using the approximation of equation (3.18). There is very close agreement between the predicted and simulated histograms, indicating that the relationship given in (3.18) for the MLV filter is a good estimate.

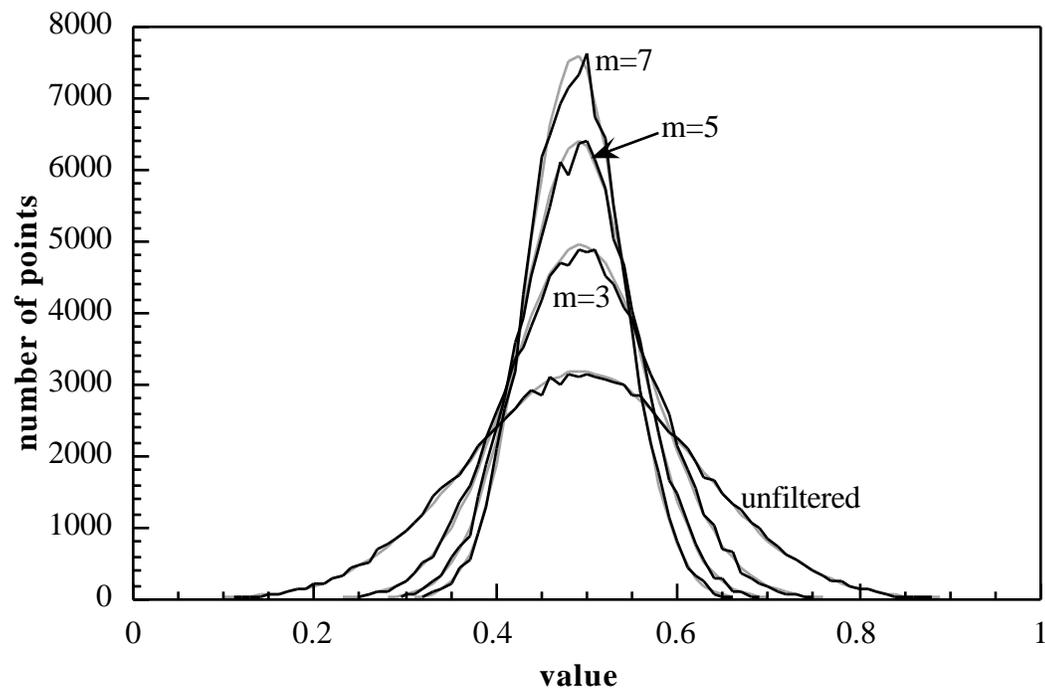


Figure 3.17. Theoretical and simulated results for the 3-, 5-, and 7-wide MLV filter acting on a normally distributed input. 100,000-point simulation; 100-bin histogram. Theoretical distributions based on the relationship given in the text.

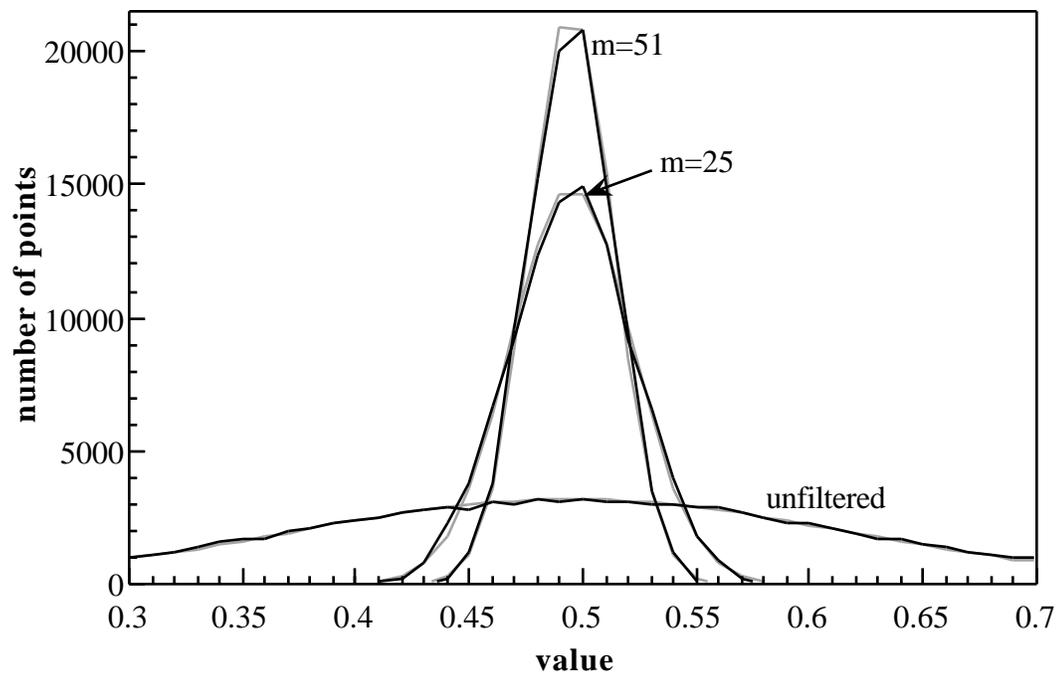


Figure 3.18. Theoretical and simulated results for the 25- and 51-wide MLV filter acting on normally distributed input. 100,000-point simulation; 100-bin histogram. Theoretical distributions are based on the relationship given in the text. Note that the value scale is truncated (0.3–0.7) to show detail.

Table 3.2. Simulated and predicted variance ratios for the MLV filter for structuring elements of various sizes from 3 to 251.

size (m)	$\frac{\sigma^2}{m}$ (observed)	$\frac{\sigma^2}{m}$ (observed)	$\frac{\sigma^2}{m}$ (predicted)	percentage deviation
1	1.001	(1.000)	1.000	—
3	0.415	2.41	2.41	0.12%
5	0.250	4.01	4.01	0.03%
7	0.178	5.62	5.68	1.08%
9	0.1381	7.25	7.40	2.12%
11	0.1118	8.95	9.15	2.19%
13	0.0941	10.63	10.92	2.64%
19	0.0628	15.94	16.31	2.33%
25	0.0460	21.7	21.8	0.28%
39	0.0290	34.5	34.8	0.91%
51	0.0216	46.3	46.1	-0.49%
63	0.01740	57.5	57.5	-0.10%
75	0.01427	70.1	68.9	-1.76%
87	0.01224	81.7	80.3	-1.72%
99	0.01059	94.5	91.8	-2.83%
151	0.00679	147.3	142.0	-3.64%
201	0.00526	190.2	190.5	0.14%
251	0.00413	242.	239.	-1.21%

For input distributions other than the normal distribution, the properties of both the averaging filter and the MLV filter are much more difficult to describe in detail. For the averaging filter, an approximation for the output distribution for any input distribution can be found using the Central Limit Theorem. This approximation becomes increasingly good as the size of the filter increases. A direct result of the Central Limit Theorem is that for a distribution with mean μ and variance σ^2 , the sample mean with n samples has a limiting distribution (as $n \rightarrow \infty$) that is normal with mean μ and variance $\frac{\sigma^2}{n}$. Thus the output distribution of an n -wide averaging filter for any input distribution with mean μ and variance σ^2 can be approximated by an $(\mu, \frac{\sigma^2}{n})$ distribution. As n increases, this approximation improves. For the MLV filter, the Central Limit Theorem also holds, with the effective filter size m replacing n in the expressions. For large structuring elements (large m), $m \gg \sigma$ and the MLV filter properties are roughly the same as the averaging filter with a window of size m . For smaller structuring elements, the statistical properties of the MLV filter are similar to those of the averaging filter with a window size slightly smaller than m .

Although the behavior of the MLV filter is quite different qualitatively from that of the averaging filter, the statistical properties of the filters are quite similar. Both filters reduce the variance of normally distributed input noise processes. The main difference is the factor by which the noise is reduced; for the averaging filter it is linear with the size of the filter, but for the MLV filter it is slightly less for all filter sizes. Although the smoothing operation of the MLV filter is simple averaging over a contiguous neighborhood, the selection process ensures that some abrupt changes in the signal will be preserved, so the output

variance of the MLV filter is slightly greater than that of the averaging filter for the same filter sizes. In contrast, the difference in the deterministic properties of the two filters is very significant. The fact that the MLV filter preserves only perfect edges and constant neighborhoods indicates its tendency to preserve and enhance sharp changes in a signal. This is in stark contrast to the averaging filter, which blurs sharp edges. Taking both the deterministic and statistical properties into consideration, one sees that the MLV filter removes noise from signals nearly as well as the averaging filter, but preserves and enhances edges in the underlying signal.

4. New Analysis Techniques

4.1. INTRODUCTION

Most of the techniques used to describe the behavior linear filters, such as Fourier analysis, do not work for nonlinear filters. This is because nonlinear filters do not obey the superposition principle; that is, the output of a nonlinear filter acting on the sum of two signals is not the same as the sum of the outputs of the same filter acting on the two signals individually. Although finding the root signals and statistical properties of nonlinear filters does shed some light on how the filters behave, this picture of filter performance is far from complete. In this chapter, three new techniques for describing the behavior of nonlinear filters will be introduced. Each technique demonstrates some aspect of how the actions of the morphological and MLV filters differ from those of the median and averaging filters.

The first technique, continuous time analysis, is similar to Fourier analysis for nonlinear filters. This analysis finds the response of filters to periodic signals of various frequencies. Since superposition does not apply, the responses for different frequencies cannot be combined, so the continuous time analysis is much more limited than Fourier analysis. The second analysis technique indicates the resistance of a filter to large outliers in its input. This robustness is measured by the breakdown point of the filter, which is the percentage of input points in a filter window that can be unbounded without forcing the filter to produce unbounded output. The breakdown point of almost any filter can be adjusted by limiting the

range of outputs with order statistics. The last analysis technique in this chapter, two-dimensional corner response analysis, quantifies the behavior of 2-D filters at corners of all angles. Filters that preserve corners with small angles are expected to preserve more detail and sharpness in an image, while filters that remove small angles generally give smoother, more rounded results.

4.2. CONTINUOUS TIME ANALYSIS

4.2.1. Introduction and Definitions

Continuous time analysis of a filter indicates the amplitude of the filter output when its input is a periodic signal of constant amplitude and frequency. This analysis is limited by several assumptions which are necessary to describe the responses of nonlinear filters. However, since the analysis is performed in continuous time, the results may be given in closed form for any signal frequency and filter size.

The continuous time response of a filter to a particular periodic waveform of a given frequency is defined to be the ratio of the peak amplitude of the filtered waveform to the peak amplitude of the original waveform. This concept of a “response” is different from the usual meaning of “response” in Fourier analysis because it is merely a peak amplitude comparison. Nonlinear filters often distort periodic signals in a manner very different from linear filters, so comparing peak amplitudes gives an indication of the relative intensities of the filter input and output; however, the frequency distortions introduced by nonlinear filters cannot be measured directly, as they can for linear filters using Fourier analysis. Despite

this fact, there have been attempts to measure the frequency responses of the median and morphological filters using Fourier analysis [42, 43]. A potential advantage of comparing peak amplitudes is that the result is a single value that accounts for the entire response of the filter at a given frequency, as opposed to spreading the effect of the filter over the entire frequency spectrum. Therefore, the effects of the filter at a given frequency are condensed into a single value, and the values for a variety of frequencies can be plotted to give an overall impression of the behavior of the filter. Comparing the power of the filtered signal to the power of the periodic input signal is another potential way to quantify the changes caused by filtering. However, it is much more difficult to compute the power of the filtered signals in continuous time for all frequencies than to compute the ratio of peak amplitudes, and in image processing applications the amplitude of an oscillation is often more representative of its visual impact than its power.

In the strictly limited case of single frequencies, the output of virtually all linear and nonlinear filters is either exactly in phase or 180° out of phase with the original signal. For signals that are symmetric in amplitude, a 180° phase shift corresponds simply to a negative ratio of peak amplitudes; that is, the sign of the output signal is opposite that of the input signal at the same time index. This means that the amplitude and phase responses can be combined into a single set of values where negative amplitude responses indicate a 180° phase shift.

Consider a filter H acting on a periodic signal $g(t)$ of period T . Assume that $g(t)$ is symmetric both in time and in amplitude; that is, $g(T - t) = -g(t)$ for all t . The filtered signal is denoted $h(t) = H\{g(t)\}$. The window size of the filter is w . The most appropriate units for the frequency f of the signal $g(t)$ are cycles per

filter window, or $f = \frac{w}{T}$ (w is in units of time per window, and T is in units of time per cycle). Assume that g and h are infinitely long. For continuous time and a time-invariant filter H , the output $h(t)$ over any cycle $[t_0, t_0 + T]$ is the same as for any translation $[t_0 + kT, t_0 + (k+1)T]$ (for any integer k). Therefore, the continuous time response of a filter H with window size w to a periodic signal $g(t)$ of frequency $f = \frac{w}{T}$ is denoted $R_H(f)$ and given by equation (4.1) below.

$$R_H(f) = \frac{h(t_p)}{g(t_p)} \quad (4.1)$$

$$\text{where } t_p = t: g(t) = \max_t \{g(t)\}_{[t_0, t_0+T)}$$

A minor issue in determining these responses in continuous time is generalizing the inherently discrete median and averaging filters to operate in continuous time. These generalizations are fairly straightforward [30, 31, 44]. The continuous time averaging filter of width w , $A\{g(t); w\}$, is defined by:

$$A\{g(t); w\} = \frac{1}{w} \int_{t-\frac{w}{2}}^{t+\frac{w}{2}} g(t) dt \quad (4.2)$$

The continuous time median filter of width w , $M\{g(t); w\}$, is defined by:

$$M\{g(t); w\} = m(t)$$

where $m(t)$ minimizes the expression

$$\left| \int_{t-\frac{w}{2}}^{t+\frac{w}{2}} \text{sgn}[g(t) - m(t)] dt \right| \quad (4.3)$$

and $\text{sgn}[\cdot]$ is the signum function

For an input signal $g(t)$ with jump discontinuities, there is the possibility of more than one value of $m(t)$ satisfying equation (4.3) for some t . This indeterminacy is resolved in different ways by Fitch, Coyle, and Gallagher [44] and in my work [30], but the specific way of handling this situation does not affect the results of the continuous time analysis.

4.2.2. Results

I have reported the results of the continuous time analysis for the median, pseudomedian, midrange, and averaging filters previously [30, 31] for sinusoidal, triangle, and square waves. These results will be summarized below, along with additional results for the LOCO and MLV filter. The individual morphological operators are not well-suited to this type of analysis, because their output is not symmetric in amplitude. If the response of the standard morphological filters is considered to be the average of the peak responses at the input signal maximum and minimum, then the response of each operator is identical to that of its complement, and identical to the filter that is the average of the complementary operators. This means that the responses of erosion and dilation are the same as the response of the midrange filter; closing and opening have the same response as the pseudomedian filter; and CO and OC have the same response as the LOCO filter. Recall also that for the compound morphological operators and the pseudomedian and LOCO filters, the effective window size w is (in continuous time) exactly twice the length of the structuring element ($w = 2n$).

Morphological Filters

The LOCO filter removes both positive and negative impulses in a signal, whereas the pseudomedian filter only reduces their amplitude. Because of this, the LOCO filter “cuts off” at a much lower frequency than the pseudomedian filter. At an input signal frequency of one cycle per window, the output of CO is a constant equal to the maximum of the signal, and the output of OC is a constant equal to the minimum. Therefore, at this frequency and all higher frequencies, the LOCO filter has a response of zero. Opening, however, removes only the positive part of a signal of with $f = 1$ cycle/window, and closing removes only the negative part. Their average (the pseudomedian) therefore has a response exactly halfway between perfect preservation and the response of the LOCO filter. This gives a general relationship between the continuous time response of the pseudomedian filter, $R_{\text{pmed}}(f)$, and the response of the LOCO filter, $R_{\text{LOCO}}(f)$, for the region $0 \leq f \leq 1$ cycles per windows:

$$\begin{aligned} R_{\text{pmed}}(f) &= \frac{1}{2} (1 + R_{\text{LOCO}}(f)) \\ R_{\text{LOCO}}(f) &= 2R_{\text{pmed}}(f) - 1 \end{aligned} \tag{4.4}$$

Averaging and MLV Filters

The basic filtering operation of the MLV filter is averaging, so its continuous time response is expected to have some resemblance to the response of the averaging filter. Consider a periodic signal that is monotonic between adjacent maxima and minima. Of the structuring elements that contain the peak of the signal, the one centered at the peak must have the lowest variance. The output when the MLV filter is centered at the peak of the input signal is then the

average of the structuring element centered at the peak. This is exactly the same value as the averaging filter with window size equal to that of the MLV structuring element. This is true regardless of the frequency of the input signal. Therefore, the continuous time response of the MLV filter with structuring element size n is the same as the response of the averaging filter with filter window width n . However, the shapes of the MLV filter and averaging filter outputs are very different; in general, only the peak output values of the two filters are the same.

Sinusoidal Input

For sinusoidal input signals, analytic expressions for the continuous time responses of the various filters mentioned above have been found and are given below. Note especially that for $0 \leq f \leq 1$ cycles per window, the expressions for the median and LOCO filters are equal.

For the LOCO filter:

$$R_{\text{LOCO}}(f) = \begin{cases} \cos\left(\frac{f}{2}\right), & 0 \leq f \leq 1 \\ 0, & f > 1 \end{cases} \quad (4.5)$$

For the pseudomedian filter [30, 31]:

$$R_{\text{pmed}}(f) = \begin{cases} \frac{1}{2} \left(1 + \cos\left(\frac{f}{2}\right)\right), & 0 \leq f \leq 2 \\ 0, & f > 2 \end{cases} \quad (4.6)$$

For the midrange filter [30, 31]:

$$R_{\text{midr}}(f) = \begin{cases} \frac{1}{2} (1 + \cos(f)), & 0 \leq f \leq 1 \\ 0, & f > 1 \end{cases} \quad (4.7)$$

For the averaging filter [30, 31] and MLV filter:

$$R_{\text{avg}}(f) = \text{sinc}(f) \quad (4.8)$$

For the median filter [30, 31]:

$$R_{\text{med}}(f) = (-1)^k \cos\left(\frac{f}{2k+2}\right), \quad k + \frac{k}{2k+1} \leq f < k + 1 + \frac{k+1}{2k+3} \quad (4.9)$$

for $k = 0, 1, 2, \dots, K$

These functions are plotted in Figures 4.1 and 4.2 below. The morphological filters have zero response (constant output signal) for all inputs above a particular frequency (1 cycle per window for the LOCO and midrange filters, 2 cycles per window for the pseudomedian filter). However, the median, averaging, and MLV filters do not have zero response except for frequencies with an exact integer number of cycles per window. These filters also have negative responses for certain frequency ranges, which indicate that they cause a 180° phase shift in the signals. This effect is related to the problem of pseudoresolution in imaging [31], and also helps explain the phenomenon of fast-fluctuating and recurrent roots of the median filter.

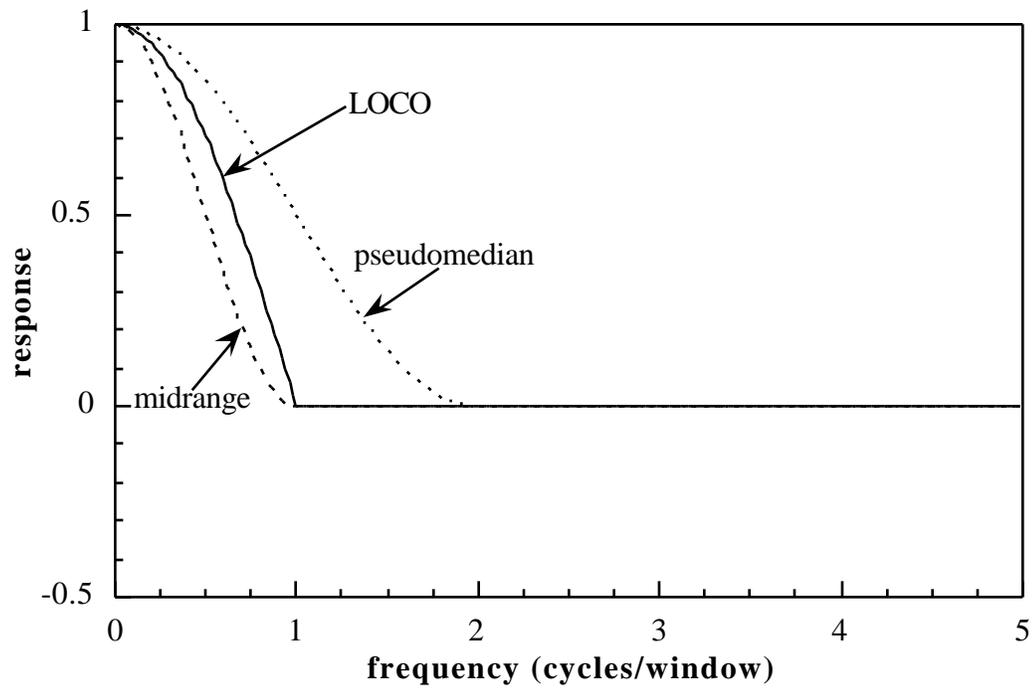


Figure 4.1. Continuous time response of midrange, pseudomedian, and LOCO filters to sinusoidal input signals.

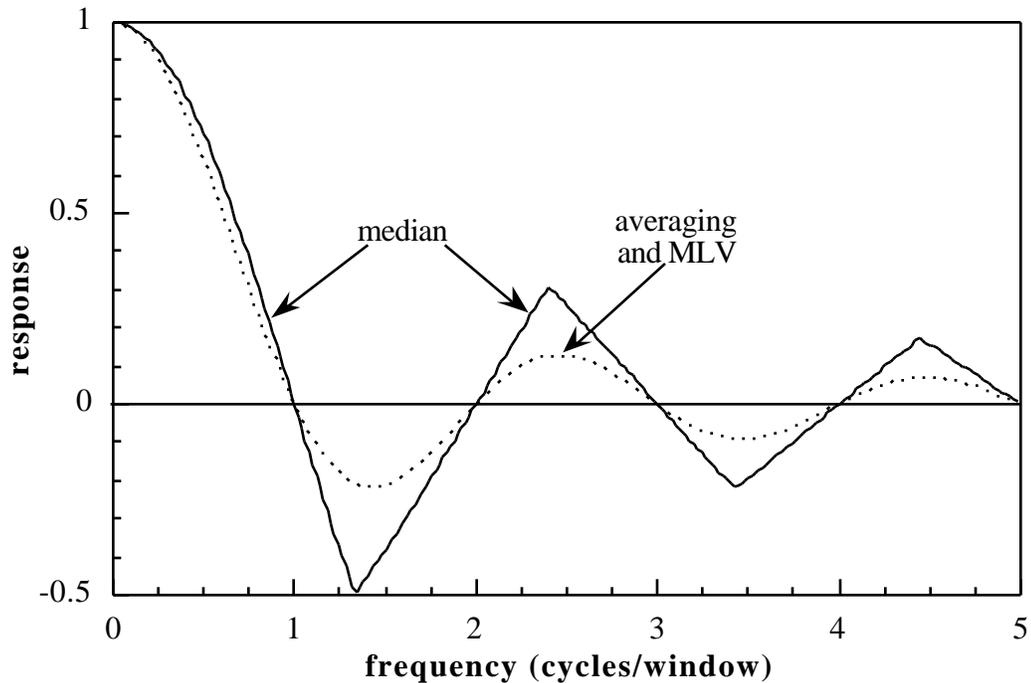


Figure 4.2. Continuous time response of averaging, MLV and median filters to sinusoidal input signals.

Square Wave Input

The existence of the fast-fluctuating and recurrent roots of the median filter is further clarified by the results for square wave input, shown in Figures 4.3 and 4.4 below. The morphological filters cut off all frequencies above 1 cycle per window (1/2 cycle per window for the midrange filter), but the median, averaging, and MLV filters have nonzero responses for all non-integer frequencies. The median filter response, shown in Figure 4.4, demonstrates the three types of median filter root signals: locally monotonic roots for frequencies between 0 and 1 cycle per window; recurrent roots for frequencies from 1 to 2, 3 to 4, ... cycles

per window; and fast-fluctuating bi-valued roots for frequencies from 2 to 3, 4 to 5, ... cycles per window.

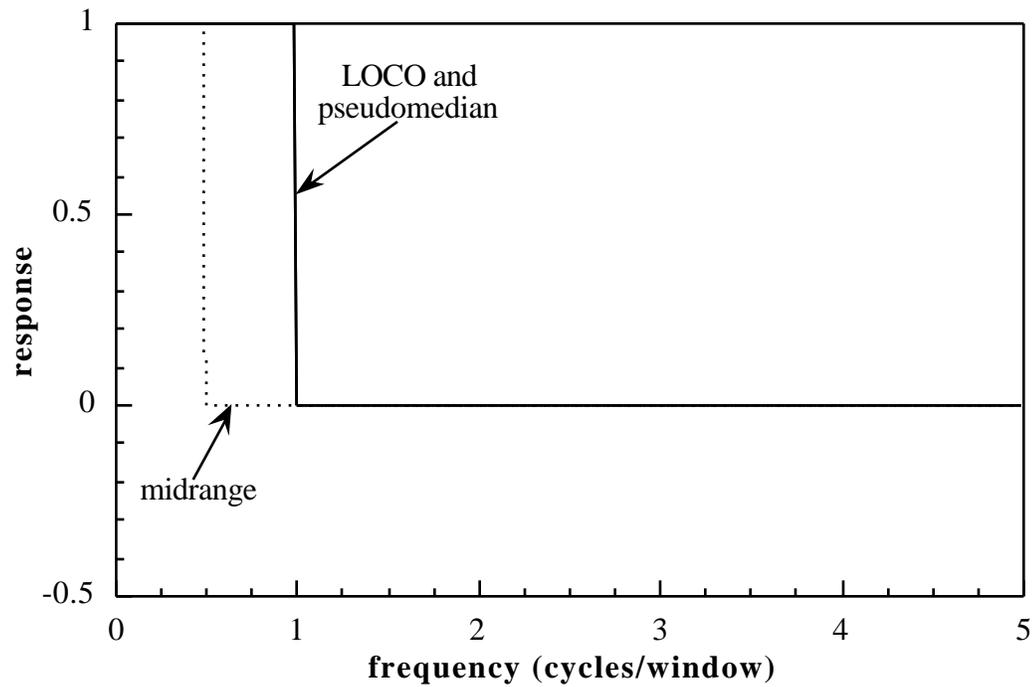


Figure 4.3. Continuous time response of midrange, pseudomedian, and LOCO filters to square wave input signals.

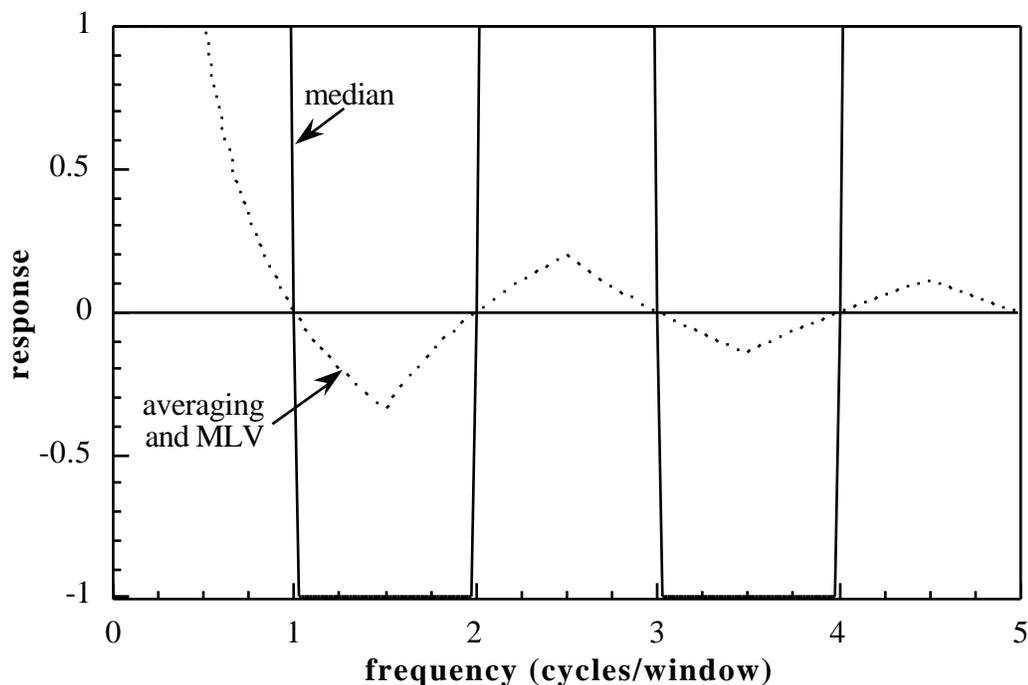


Figure 4.4. Continuous time response of averaging, MLV and median filters to square wave input signals.

The continuous time response of the morphological and MLV filters illustrates their peak response to periodic signals of various frequencies. Although the results for different frequencies cannot be combined to draw any conclusions about the behavior of the filters for multi-frequency input, this analysis does provide some help in understanding what types of fluctuations are preserved and removed by the different filters. A particularly interesting result is that for both sinusoidal and square wave input, the continuous time peak responses of the LOCO and median filter are identical for frequencies between 0 and 1 cycle per window. For higher frequencies, however, the LOCO filter has zero response (constant output), while the median filter has a nonzero response.

This indicates that the LOCO filter and the other morphological filters are better at suppressing rapid fluctuations in a signal or image than the median filter is. The continuous time analysis of nonlinear filters is thus a useful tool for understanding the behavior of filters in the presence of high frequency periodic noise.

4.3. BREAKDOWN POINT

4.3.1. Definition and Analysis

The continuous time analysis of a filter indicates its susceptibility to rapidly fluctuating noise in a signal. Another analysis technique that measures the susceptibility of a filter to outliers (impulses) in a signal is the breakdown point. The breakdown point is the smallest percentage of points in a filter window that, when chosen arbitrarily large (positive or negative), can create an arbitrarily large filter output [45-47]. The breakdown point, denoted b^* , has a specific value that varies with the window size of the filter; however, as the window size becomes very large the breakdown point usually approaches an asymptotic value, b^* . For example, the median filter with a window W , where $|W| = 2n+1$, has a breakdown point of $b^* = \frac{n+1}{2n+1}$ and an asymptotic breakdown point of $b^* = 0.50$ (or 50%). The averaging filter, on the other hand, has $b^* = 0$, because $b^* = \frac{1}{|W|}$. The breakdown point is an indication of the percentage of points in a filter window that may be outliers (unexpectedly large or small values) without ever causing the filter output itself to become an outlier. A filter with a high breakdown point is considered to

be robust to outliers, while a low breakdown point indicates that a filter is not good at resisting impulse noise.

Erosion and opening are less than or equal to their input signal at all points, so they break down in the presence of even a single negative outlier. Similarly, dilation and closing break down in the presence of one positive outlier. Therefore, all these operators have the same breakdown point, $\frac{1}{|W|}$ and $\tau^* = 0$, where W is the effective window size of the operations. The doubly compound operators OC and CO, however, reject single outliers but still can give aberrant output in one dimension with as few as two outliers in the filter window. Therefore, the 1-D OC and CO have $\frac{2}{|W|}$, which is twice as high as for the other morphological operators. In two dimensions, four points in the filter window are required to cause OC or CO to break down, so $\frac{4}{|W|}$. However, the asymptotic breakdown point for both the 1-D and 2-D OC and CO is the same as for the other operators, $\tau^* = 0$. Since the breakdown of one of the constituent operators of the midrange, pseudomedian, and LOCO filters causes the overall filter to break down as well, these filters each have the same breakdown point as their constituent filters.

The MLV filter tends to take its averages over structuring elements that do not include outliers, since outliers greatly increase the variance of a subwindow. However, when an outlier lies directly in the center of the overall window of the MLV filter, that outlier is included in every structuring element considered by the filter. Since the filter takes the average of one of these structuring elements, a single outlier can cause the MLV filter to break down: $\frac{1}{|W|}$ and $\tau^* = 0$. This is the same breakdown point as for the averaging filter, but the MLV performs

significantly better in the presence of single outliers. The averaging filter includes the effect of a single outlier in its output over an area the size of the filter window, while the MLV filter generally limits the effect of the outlier to a single point in the output. Other filters using the value-and-criterion filter structure may be designed to be more robust to impulses; for example, the averaging operation can be replaced with a median operation. In an environment where impulse noise is routinely encountered, there may be an advantage to replacing the averaging operation with the median (and perhaps also the variance operation with an appropriate scale measure) to improve the resistance of the filter to outlying values.

A summary of the breakdown points of the various filters discussed in this section is shown in Table 4.1 below. The LOCO filter and the OC and CO have different values for filters of different dimensions; for the other filters listed, these values are valid for any size and any dimension. The breakdown point gives an objective criterion on which to evaluate the performance of a filter in the presence of impulse-type noise. It can also be used as a tool to assist the design of filters that are robust to large outliers, as will be illustrated in the next section.

Table 4.1. Summary of breakdown points of filters.

Filter(s)	Breakdown Point	Asymptotic Breakdown Point *
Erosion, Dilation, Midrange	$\frac{1}{ W }$	0
Opening, Closing, Pseudomedian	$\frac{1}{ W }$	0
OC, CO, LOCO (1-D)	$\frac{2}{ W }$	0
OC, CO, LOCO (2-D)	$\frac{4}{ W }$	0
Median	$\frac{ W + 1}{2 W }$	0.5
Averaging, MLV	$\frac{1}{ W }$	0

4.3.2. Order Statistic Limited (α -Limited) Filters

One way to increase the breakdown point of a filter is to limit the range of filter outputs to a range of the order statistics of the filter window [27, 46-49]. For example, if a certain fraction α (where $0 < \alpha < 0.5$) of both the highest and lowest order statistics of a filter window are excluded from the range of possible outputs, the breakdown point is at least $\alpha^* = \frac{1}{2}$, since outliers must fill the excluded order statistics before the output range extends to the level of the outliers. Examples of this technique include the α -trimmed means [27, 46-48] and the LUM (Lower order statistic, Upper order statistic, Middle sample) filter [49]. An α -trimmed mean filter takes the average of the values remaining in the filter window after the extreme order statistics are removed; when $\alpha = 0$, the averaging filter results, and

$\alpha = 0.50$ corresponds to the median filter [48]. All other values of α result in an intermediate trimmed mean filter. The asymptotic breakdown point of an α -trimmed mean filter is $b^* = \alpha$.

The LUM filter for smoothing [49] changes the value in the middle of the filter window only if it is one of the extreme order statistics removed from the range; the output is then the value in the allowed range nearest to the removed value. This is exactly equivalent to a center-weighted median filter, which is simply the median filter taken on a window which has the center sample repeated a certain number of times. In this case, if the filter window size is originally $|W|$, the center point is repeated $(1-2\alpha)|W|$ times, and the median is taken over all $|W|+(1-2\alpha)|W|$ points. The smoothing LUM filter (or center-weighted median filter) with parameter α also has an asymptotic breakdown point of $b^* = \alpha$.

Limiting the output range of a filter in this way creates a parameterized spectrum of filters based on almost any original filter. This spectrum of filters ranges from the original filter at $\alpha = 0$ to the median filter at $\alpha = 0.5$. As α increases, the resistance of the filter to impulse noise improves, as evidenced by the breakdown point of the filter. All of the nonlinear filters discussed previously may be “ α -limited” in this manner to improve their robustness in the presence of outliers. For the linear combinations of morphological filters, better results are obtained by α -limiting the constituent morphological filters individually, since the averaging operation sometimes allows the effects of outliers to appear in the output without violating the range constraints placed on the filter by the α -limiting operation.

An example of how α -limiting can improve the behavior of a filter in the presence of impulse noise is shown in Figures 4.5 and 4.6 below. The input signal, plotted in Figure 4.5, is four cycles of a sine wave sampled at 250 points per cycle and corrupted by impulse noise at 10% of these points. The noise process used in this case adds the value one to 5% of the points and subtracts the value one from 5% of the points. This is a type of “salt-and-pepper” noise [50]. The output of the standard LOCO filter and the LOCO filter with α -limited OC and CO operations is shown in Figure 4.6. The structuring element size for both cases was 25. The impulse noise in the input causes severe distortion in the output of the standard LOCO filter. However, by limiting the output of the OC and CO operations to the central 60% of the order statistics ($\alpha = 0.2$), the output of the LOCO filter is greatly improved.

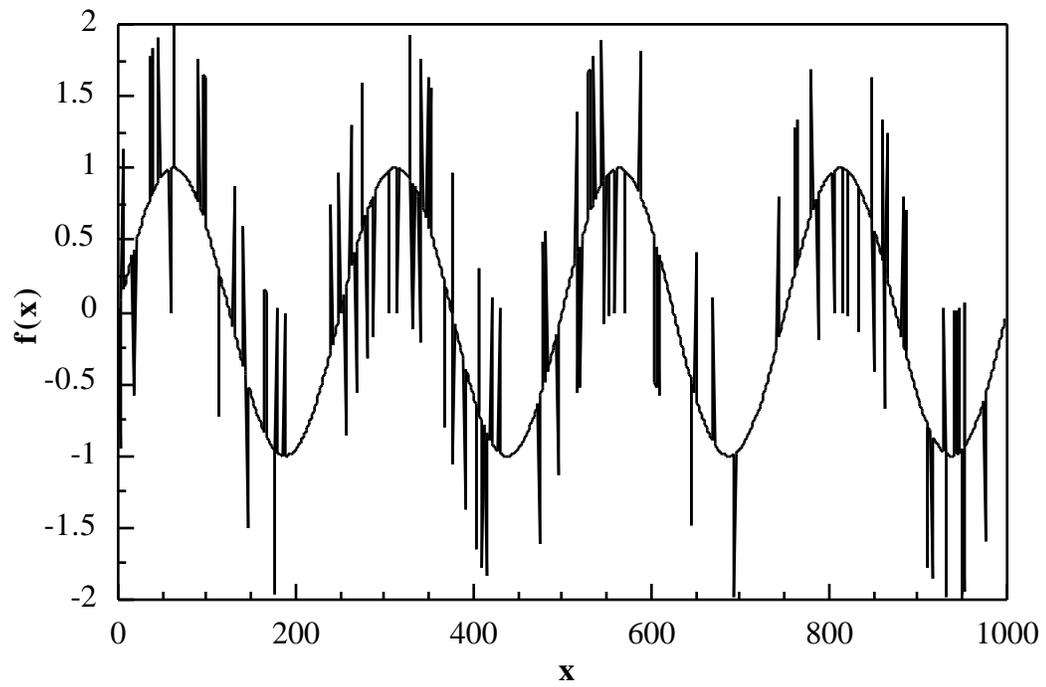


Figure 4.5. Noisy input signal. 1000-point sinusoid with 5% positive impulses and 5% negative impulses.

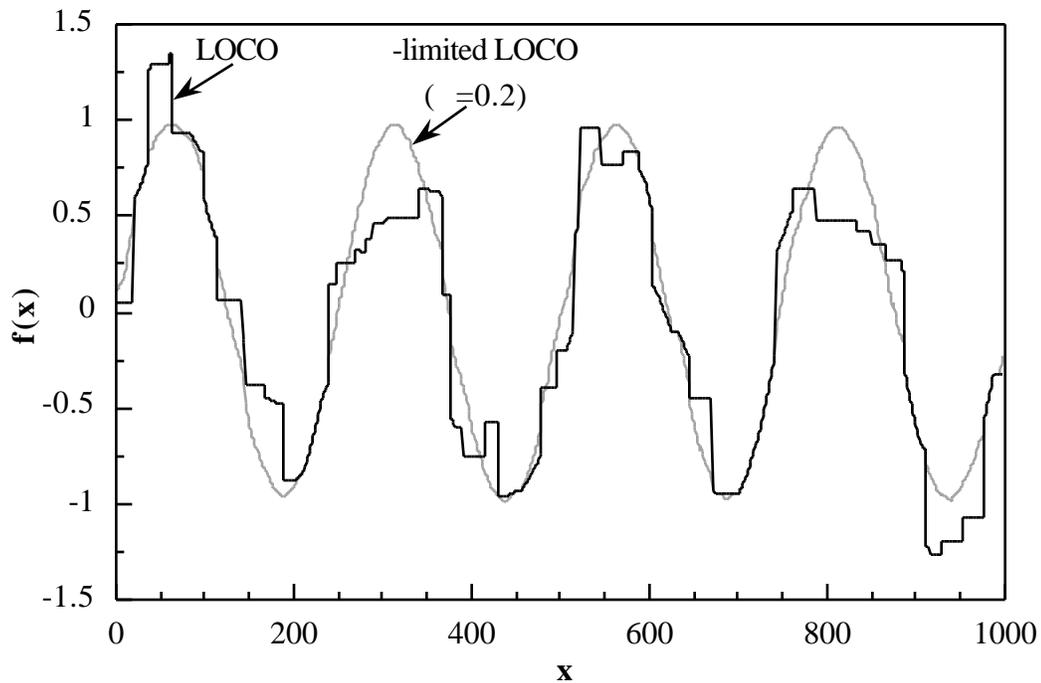


Figure 4.6. Output of LOCO filter and λ -limited LOCO filter with 25-wide structuring element and $\lambda = 0.2$.

The method of λ -limiting the output range of a filter increases the robustness of a filter to outliers (as measured by the breakdown point) to any desired value up to $\lambda^* = 50\%$. At $\lambda = 0$, the filter is unchanged, and at $\lambda = 0.5$, the filter output is identical to that of the median filter. However, intermediate values of λ result in filters that retain many of the features of the original filter, but achieve some of the impulse rejection capability of the median filter. For instance, the λ -limited LOCO filter retains the LOCO filter's ability to resist high frequency periodic noise but also rejects impulses quite well when λ is chosen appropriately. (See Figure 4.6.) This ability to form intermediate filters by λ -limiting is also important in two dimensions; for example, the corner-preserving

properties of the LOCO filter discussed in the next section can be preserved while adding a measure of impulse rejection.

4.3.3. Conclusions

The breakdown point provides another way to quantify the susceptibility of nonlinear filters to impulse noise, and illustrates weaknesses of the pseudomedian, LOCO, and MLV filters in the presence of impulse noise. Order statistic limiting, or α -limiting, is one way to modify almost any filter to resist impulse noise better. For example, suppose an application is well suited to LOCO or MLV filtering except that a significant amount of impulse noise is present. The shape control and other assets of these filters can be preserved while improving the impulse resistance of the filter by using the α -limiting technique. Thus, using the breakdown point to analyze filters and α -limiting to alter their responses increases the variety of filtering methods available for use in environments with significant impulse noise, such as thermography.

4.4. CORNER RESPONSE ANALYSIS

The response of one- and two-dimensional filters to edges has been a topic of much research [9, 51, 52]. The edge preserving properties of nonlinear filters such as the median filter is one reason that such filters have become important in signal and image processing applications. However, the edge response of a filter gives little information about its response to two-dimensional structures such as corners. For example, the median filter preserves edges and eliminates impulse noise, but it noticeably rounds off sharp two-dimensional corners. This section

introduces a new technique to measure the degree of preservation of corners of all angles by a filter with a given window shape. This technique may be applied to discrete filter windows acting on digital images, or may be used to derive theoretical results in continuous space that yield more general information about filters and window shapes. The corner responses may also be derived for corners oriented at any angle relative to the image axes.

4.4.1. Definitions

Fractional Corner Preservation

A corner of angle θ is defined in 2-D space at the origin of the x - y plane, as shown in Figure 4.7. The corner is rotated with respect to the x -axis by an angle θ . The original image $f(x,y)$ of the corner consists of only two values; in Figure 4.7, white is assumed to correspond to the value 0 and black to correspond to 1. Assume that $f(x,y)$ takes on values in the range $0 < f(x,y) < 1$. The regions I, II, III, and IV in Figure 4.7 are each the size of the smallest square into which the overall filter window fits. The total area of the corner defined by θ and r is defined by the area within the angle θ that is also within the regions I, II, III, and IV. Thus in Figure 4.7, the total corner area, A , is the black area inside region I. Since $f(x,y) = 1$ everywhere in A ,

$$\int_A f(x,y) dx dy = A \quad (4.10)$$

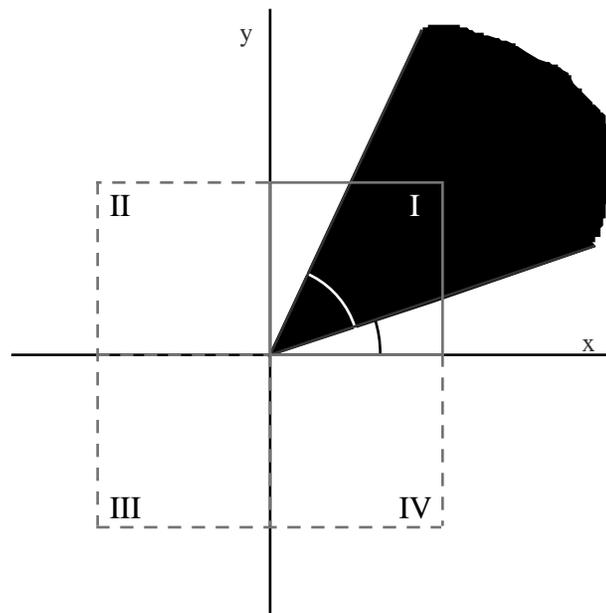


Figure 4.7. Original image of corner of angle with rotation .

Filtering usually changes parts of the corner from black (1) to gray or white (values less than 1). This is illustrated in Figure 4.8. If $g(x,y)$ denotes the filtered image, then we define the preserved corner area P as:

$$P = \int_A g(x,y) dx dy \quad (4.11)$$

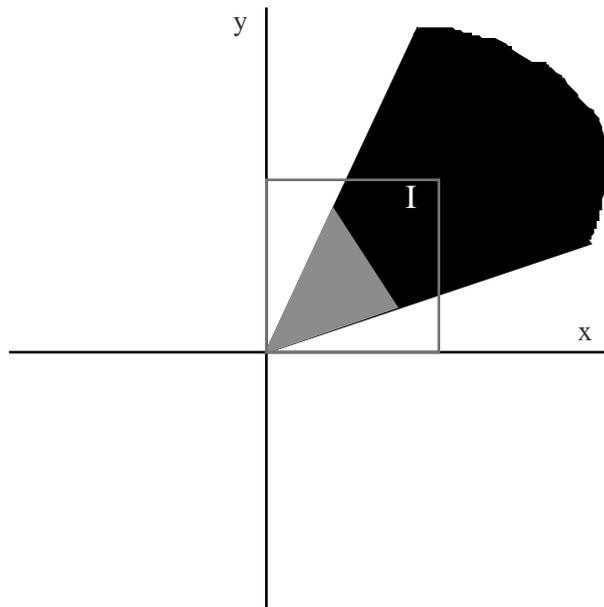


Figure 4.8. Example filtered image of corner.

For the median filter and other filters that do not return intermediate values for a bi-valued input, P is the area of A that the filter preserves as black (pixel value 1). For filters that do yield intermediate values, such as the averaging filter, P corresponds to the relative amplitude preservation of the corner only in the area defined by the original corner.

The fractional preservation of a corner of angle θ oriented at an angle ϕ with respect to the x-axis, denoted $r(\theta, \phi)$, is defined as:

$$r(\theta, \phi) = \frac{P}{A} \quad (4.12)$$

The response of a filter to corners of all angles at a particular orientation ϕ_0 is therefore only a function of the angle θ , denoted $r(\theta, \phi_0)$. Note that r ranges in value from 1 (perfect corner preservation) to 0 (complete removal of the corner in the region of interest).

Corner Attenuation

The fractional preservation can also be expressed as an attenuation in decibels. The corner attenuation in dB, $s(\theta, \alpha)$, is defined as:

$$s(\theta, \alpha) = 20 \log[r(\theta, \alpha)] = 20 \log\left[\frac{P}{A}\right] \quad (4.13)$$

From this definition, $s(\theta, \alpha) = 0$ dB corresponds to perfect corner preservation, and $s(\theta, \alpha) = -3$ dB corresponds to an average fractional corner preservation of $\frac{1}{\sqrt{2}}$ (that is, about 71% of the area is preserved intact, or the entire area is preserved at about 71% amplitude).

Continuous or Discrete Space Representation

The responses $r(\theta, \alpha)$ and $s(\theta, \alpha)$ may be determined either for discrete filters acting on a discrete lattice (that is, x and y are integers) or for continuous filters in continuous space (x and y are real). In the discrete case, each point (x, y) in the lattice corresponds to an angle θ that is the arctangent of $\left(\frac{y}{x}\right)$, and the total corner area A of an angle θ at orientation α (as in Figure 4.7) is assumed to be all points (x, y) in the regions I, II, III and IV where $\theta < \alpha < \theta + \pi$.

For continuous space, equations (4.10)–(4.12) apply directly, using the continuous space definitions of the filters. The response of a discrete filter is usually very similar to that of its corresponding continuous filter, especially for large window sizes. Very large discrete windows give a numerical approximation of the continuous space integrals in equations (4.10) and (4.11).

4.4.2. Results

For many filters, such as the averaging and median filter, closed form solutions for $r(\theta, \alpha)$ and $s(\theta, \alpha)$ are difficult to derive, even for $\alpha = 0$. However,

values of r and s at particular values of θ and ϕ often are readily found, such as $\theta = 45^\circ, 90^\circ, 135^\circ$, and 180° . These values may then be compared with the results of discrete simulations with large window sizes to verify the accuracy of the numerical approximations. The results given in this section are derived in detail in the Appendix.

Morphological Filters

Closed form solutions for $r(\theta, \phi)$ and $s(\theta, \phi)$ can be derived in continuous space for the morphological filters. This is because the binary morphological filters are geometrical operations, so their response can usually be determined by the interaction of the filter window shape with the corner. As with all comparisons between morphological and non-morphological filters, the appropriate window to consider the morphological filters acting over is the compounded window (dilation of the structuring element with itself).

For a square structuring element of arbitrary size in continuous space, the fractional corner preservation for the morphological closing is unity for all θ and ϕ ; that is, $r(\theta, \phi) = 1$. The corner attenuation is therefore $s(\theta, \phi) = 0$ for all angles θ and rotations ϕ . This result is found by noting that even for the sharpest corners, the square dilation expands them so that the subsequent erosion returns them to their original shape.

However, the morphological opening tends to clip off sharp corners because its first operation is erosion. Opening preserves high values only at points where an entire structuring element including that point is filled with high values. By simple geometry, closed form solutions for $r(\theta, \phi)$ and $s(\theta, \phi)$ may be

obtained for opening. For $\theta = 0$ and $\alpha < \tan^{-1}(0.5)$, roughly 27° , the square structuring element does not fit even partially within both the corner and region I, so $r(\alpha) = 0$. (Recall that since the structuring element is compounded to form the overall window, the structuring element is half as wide and half as tall as region I.) The complete expression for $r_{\text{open}}(\alpha, 0)$ when $\theta = 0$ for opening with a square structuring element is given in equation (4.14) below, and the complete derivation is given in the Appendix.

$$r_{\text{open}}(\alpha, 0) = \begin{cases} 0, & 0 \leq \alpha < \tan^{-1}\left(\frac{1}{2}\right) \\ 1 - \frac{1}{4} \cot^2 \alpha, & \tan^{-1}\left(\frac{1}{2}\right) \leq \alpha < \frac{\pi}{4} \\ \frac{1 - \frac{5}{8} \cot^2 \alpha}{1 - \frac{1}{2} \cot^2 \alpha}, & \frac{\pi}{4} \leq \alpha < \frac{\pi}{2} \\ 1, & \frac{\pi}{2} \leq \alpha \end{cases} \quad (4.14)$$

As expected, opening with a square structuring element perfectly preserves all obtuse and right angles when $\theta = 0$. For angles $\alpha < \frac{\pi}{2}$, the corner analysis is complementary to the case for $0 < \theta < \frac{\pi}{2}$. The effective corner angle is then $\theta = \frac{\pi}{2} - \alpha$, and the corner consists of an area of zero values in a background of value one. The total corner area would be restricted to the area of the 0's (instead of the area of the 1's) within regions III and IV of Figure 4.7. In this chapter, results will be given for $\frac{\pi}{2} < \theta < \pi$, and the results for $\theta < \frac{\pi}{2}$ follow by using converse logic.

Close-opening (CO) has exactly the same fractional corner preservation as opening, since the result of the first operation (closing) is the same as the original corner image for any α and θ . Similarly, open-closing (OC) has the same $r(\alpha, \theta)$

as opening and CO because the second operation (closing) does not alter the result of the first operation (opening) in the case of a simple binary corner. Therefore, the fractional corner preservation in equation (4.14) is valid for opening, OC, and CO.

Since the outputs of OC and CO are identical for the original corner images, the LOCO filter has the same response, as given in equation (4.14). Therefore, the opening, OC, CO and LOCO filters all have the same fractional corner preservation, as indicated in equation (4.15) below.

$$r_{\text{LOCO}}(f, g) = r_{\text{OC}}(f, g) = r_{\text{CO}}(f, g) = r_{\text{open}}(f, g) \quad (4.15)$$

The pseudomedian filter, however, is the average of two different outputs, those of opening and closing. Closing has fractional corner preservation that is always 1, and the expression for the corner preservation of opening is given in equation (4.14). The fractional corner preservation of the pseudomedian filter is therefore given by:

$$r_{\text{pmed}}(f, g) = \frac{1}{2} [1 + r_{\text{open}}(f, g)] \quad (4.16)$$

The fractional corner preservation for the midrange filter is given in equation (4.17) below, and its derivation is given in the Appendix. The relationship between the corner response of the midrange filter and dilation is the same as between the pseudomedian filter and opening: $r_{\text{midr}}(f, g) = \frac{1}{2} [1 + r_{\text{erode}}(f, g)]$. The midrange filter is perhaps not as well suited to corner response analysis as are the other filters discussed here, since it tends to severely distort edges and corners. After midrange filtering, corners and edges extend far beyond their original borders, and the unusually small values for the fractional corner preservation in equation (4.17) reflect this spreading caused by the filter.

$$\begin{aligned}
r_{\text{midr}}(\alpha, 0) = & \begin{cases} \frac{1}{2}, & |\alpha| \leq \tan^{-1}(2) \\ \frac{\tan^2|\alpha| + 4\tan|\alpha|}{16\tan|\alpha| - 8}, & \tan^{-1}(2) < |\alpha| \leq \tan^{-1}(3) \\ \frac{10\tan|\alpha| - 9}{16\tan|\alpha| - 8}, & \tan^{-1}(3) < |\alpha| < \frac{\pi}{2} \\ \frac{5}{8}, & \frac{\pi}{2} \leq |\alpha| \leq \frac{3\pi}{4} \\ \frac{4\tan|\alpha| + 18 - \cot|\alpha|}{8\tan|\alpha| + 32}, & \frac{3\pi}{4} < |\alpha| \leq \tan^{-1}\left(-\frac{1}{3}\right) \\ \frac{13\tan|\alpha| + 24}{8\tan|\alpha| + 32}, & \tan^{-1}\left(-\frac{1}{3}\right) < |\alpha| \leq \pi \end{cases}
\end{aligned} \tag{4.17}$$

Median and Averaging Filters

The fractional preservation for the median and averaging filters have been determined numerically (in discrete space), using 63×63 square filter windows. The corner responses of the median and averaging filters can be determined in continuous space for certain α and β , and these numbers may be used to check the simulation to see how close to the continuous case it is. Some of the known values of the median filter fractional preservation in continuous space with $\beta = 0$ are given in equation (4.18) below. Derivations are given in the Appendix.

$$\begin{aligned}
r_{\text{med}}(\alpha, 0) = & \begin{cases} 0, & 0 \leq |\alpha| < \tan^{-1}\left(\frac{1}{2}\right) \\ \frac{2}{3}, & |\alpha| = \frac{\pi}{4} \\ \frac{5}{4} - \frac{1}{2}\ln 4 + \frac{1}{2}\ln 2, & |\alpha| = \frac{\pi}{2} \\ 1, & |\alpha| = \pi \end{cases}
\end{aligned} \tag{4.18}$$

Some of the known values of the averaging filter fractional preservation in continuous space with $\beta = 0$ are given in equation (4.19) below. These values for

both the median and averaging filters match the 63×63 simulations extremely well.

$$r_{\text{avg}}(\theta, 0) = \begin{cases} \frac{43}{64} \tan \theta, & 0 \leq \theta < \tan^{-1}\left(\frac{2}{3}\right) \\ \frac{49}{64}, & \theta = \frac{\pi}{2} \\ \frac{7}{8}, & \theta = \pi \end{cases} \quad (4.19)$$

Comparison of Filter Corner Responses

Figure 4.9 is a plot of the fractional corner preservation of the opening, OC, CO, and LOCO operations compared with those of the pseudomedian and median filters for square structuring elements. The corner response functions are plotted in polar coordinates to give a visual indication of the angles of corners that are preserved and those that are destroyed. All curves in the figure are for corners aligned with the x -axis ($\phi = 0$) and π . Note that for every angle $\theta > \tan^{-1}(0.5)$, the LOCO filter preserves more of the corner than the median filter. The median filter does not preserve any corner exactly except the null case of $\theta = \pi$ (180°), where the corner becomes a straight edge. The square LOCO filter, on the other hand, perfectly preserves corners with $\theta = \frac{\pi}{2}$ (90°) when $\phi = 0$.

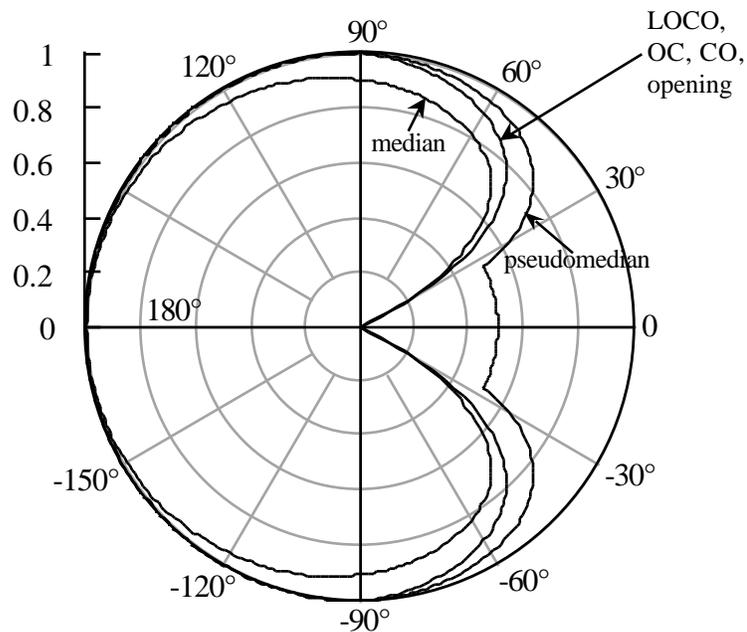


Figure 4.9. Fractional corner preservation $r(\theta, 0)$ of morphological opening, OC, CO, and the LOCO filter compared to median and pseudomedian filters (square structuring elements/filter windows).

Figure 4.10 is another illustration of the same information. In this case, the attenuation of the corner is plotted in rectangular coordinates, with the angle of the corner (in degrees) as the x -axis and the attenuation $s(\theta, 0)$ in decibels as the y -axis. This way of plotting the information puts it in a form similar to the forms used for analog filters, and by analogy one may define a “passband” and a “stopband” for the filters. For example, if the corner “passband” of a filter is defined to be the region where the attenuation is between 0 and -3 dB, the corner “passband” of the LOCO filter is approximately 43° . The median filter's corner “passband,” in contrast, begins at about $\theta = 48^\circ$. If the corner “stopband” is defined as the region where the attenuation is greater than -10.67 dB, then the

corner “stopband” of the LOCO filter runs from 0 to about 30° , whereas for the median filter it runs from 0 to roughly 31° . The LOCO filter therefore has a wider corner “passband” and a sharper cutoff than the median filter, while the two filters have approximately the same corner “stopband.” The pseudomedian filter has a corner “passband” of approximately 34° , but does not have a corner “stopband” according to the definition given above, since its maximum attenuation is -6 dB. This analysis shows that these filters are analogous to low pass filters, because they “pass” corners with angles above a certain value (the “cutoff” angle) and “stop” corners sharper than a given angle. (Sharp corners, like sharp edges, have higher frequency content than more obtuse corners.)

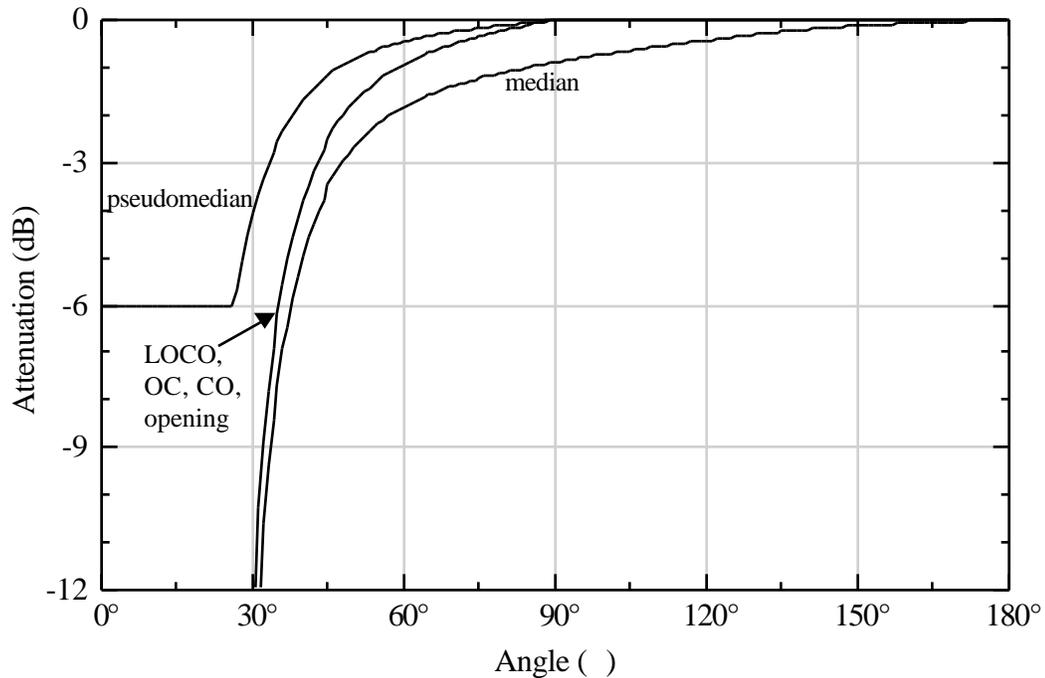


Figure 4.10. Corner attenuation $s(\theta, 0)$ of morphological opening, OC, CO, and the LOCO filter compared to median and pseudomedian filters (square structuring elements/filter windows).

Effect of Filter Window Shape

Different window shapes may be used to alter the response of filters to edges and corners. For instance, a plus-shaped window is often used with the median filter to preserve 90° corners aligned with the coordinate axes ($\theta = 0$). Figure 4.11 illustrates the corner response of the plus-shaped median filter compared to the square-shaped median filter. A closed form solution for this corner response, $r_{psm}(\theta, 0)$, is derived using geometry in continuous space in the Appendix, and is shown in equation (4.20) below. (Note that in continuous space a plus-shaped window has infinitesimally thin “arms,” but the amount of the

window that must be covered by 1's to result in an output of 1 is still easily determined.)

$$r_{\text{psm}}(\theta, 0) = \begin{cases} \tan^{-1} \left| \frac{1}{4} \frac{1}{1+\tan|\theta|} \right| & 0 \leq |\theta| < \tan^{-1} \left(\frac{1}{2} \right) \\ 1 - \frac{1}{4} \frac{1}{\tan|\theta| \left| (1+\tan|\theta|) \right|} & \tan^{-1} \left(\frac{1}{2} \right) \leq |\theta| < \frac{\pi}{4} \\ \frac{7}{8} - \frac{1}{2} \cot|\theta| + \frac{1}{8} \frac{\tan|\theta|}{1+\tan|\theta|} - \frac{2 \tan|\theta|}{2 \tan|\theta| - 1} & \frac{\pi}{4} \leq |\theta| < \frac{\pi}{2} \\ 1 & \frac{\pi}{2} \leq |\theta| \end{cases} \quad (4.20)$$

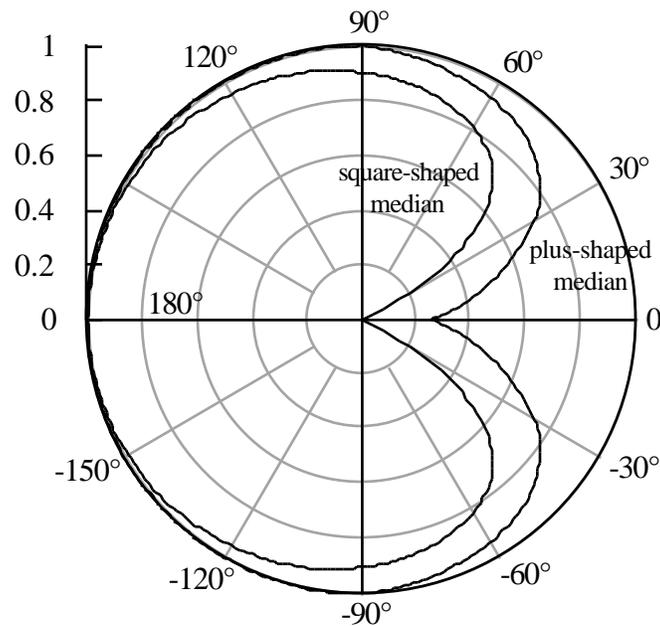


Figure 4.11. Fractional corner preservation $r(\theta, 0)$ of the plus-shaped median filter compared to the square-shaped median filter.

Figure 4.12 shows the attenuation of the plus-shaped median filter in decibels. Its corner “passband” (as defined above) is for angles greater than about 29°, while its “stopband” is only 0 to 4°. The plus-shaped median filter therefore

preserves corners significantly better than the square-shaped median filter when $\alpha = 0$, but its cutoff is not nearly as sharp, and it does not completely remove very acute corners. This may sometimes be an advantage, since the plus-shaped median filter does preserve straight, thin lines aligned with either of its axes. However, the plus-shaped median filter is very sensitive to corner orientation, as will be shown in the next section.

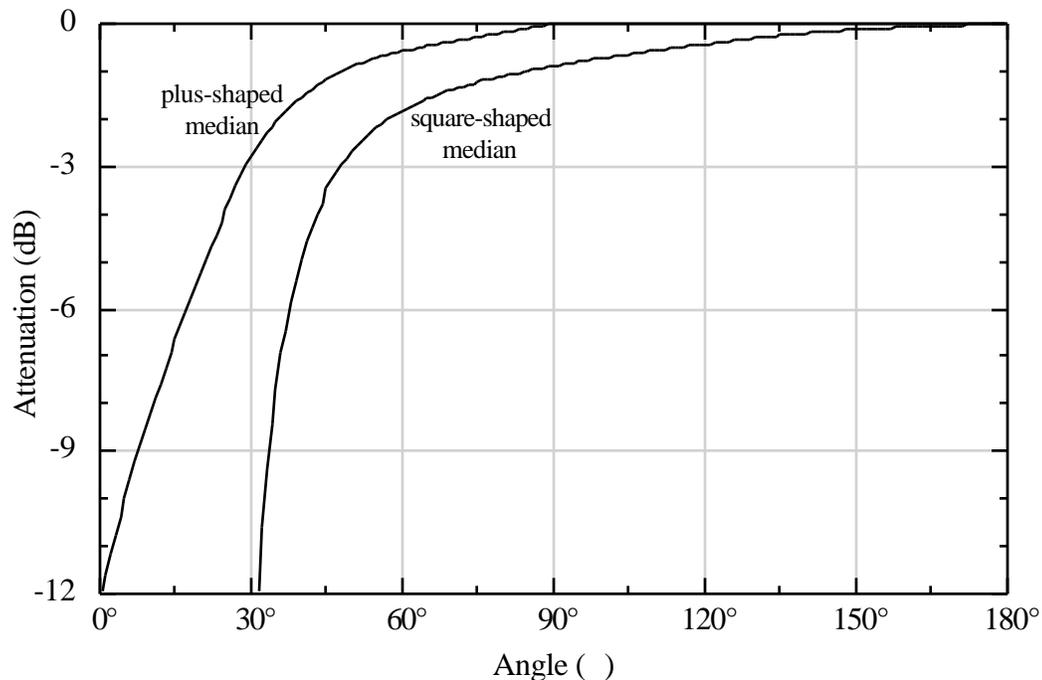


Figure 4.12. Corner attenuation $s(\alpha, 0)$ of the plus-shaped median filter compared to the square-shaped median filter.

Morphological filters are very sensitive to structuring element shape. Square structuring elements are used widely and are very good at preserving obtuse corners, but the sharp 90° corners of square structuring elements can sometimes create a “stairstep” effect at diagonal edges in noisy images.

Therefore, circular structuring elements are sometimes used to avoid this effect. In discrete space, the circular shape can only be approximated, but in the continuous space an exact circle can be used to perform the geometric calculations leading to a closed form solution for the corner response of opening, OC, CO, and the LOCO filters. (See the Appendix.) However, the form of the solution is difficult, especially in the few degrees just above the angle where it exhibits complete cutoff. The tractable portions of the response $r_{\text{open}}(\theta, 0)$ for a circular structuring element are:

$$r_{\text{open}}(\theta, 0) = \begin{cases} 0, & 0 \leq \theta < \frac{\pi}{4} \\ 1 - 2\cot|\theta| f(\theta), & \frac{\pi}{4} \leq \theta < \frac{3\pi}{4} \\ 1 - \frac{1}{1 - \frac{1}{2}\cot|\theta|} f(\theta), & \frac{3\pi}{4} \leq \theta < \frac{5\pi}{4} \\ 1 - \frac{1}{2 + \frac{1}{2}\tan|\theta|} f(\theta), & \frac{5\pi}{4} \leq \theta < \frac{3\pi}{2} \end{cases}$$

$$\text{where } \begin{cases} 5\sin\theta = 1 + \cos\theta, \text{ or } 22.6^\circ \\ 4\sin\theta = 1 + \cos\theta, \text{ or } 28.1^\circ \\ f(\theta) = \frac{1}{16} \frac{|\sin\theta|}{|1 - \cos\theta|} - \frac{1}{32} (|\theta| - \frac{\pi}{2}) \end{cases} \quad (4.21)$$

The fractional preservation of the LOCO filter (and opening, OC, and CO) with a round structuring element is compared to that with a square structuring element in Figure 4.13. Values for the 6° of angle missing in equation (4.21) were approximated by interpolation and verified by a discrete simulation. The morphological filters with circular structuring elements preserve more acute corners than for square structuring elements, but do not perfectly preserve obtuse and right angles. The corner “passband” is roughly $\theta > 37^\circ$ and “stopband”

approximately $< 27^\circ$ for the circular structuring element filters. Therefore, the transition band (range of angles between the “passband” and “stopband”) of only about 10° for the LOCO filter with a round structuring element is smaller than the transition bands for the LOCO and median filters with square windows.

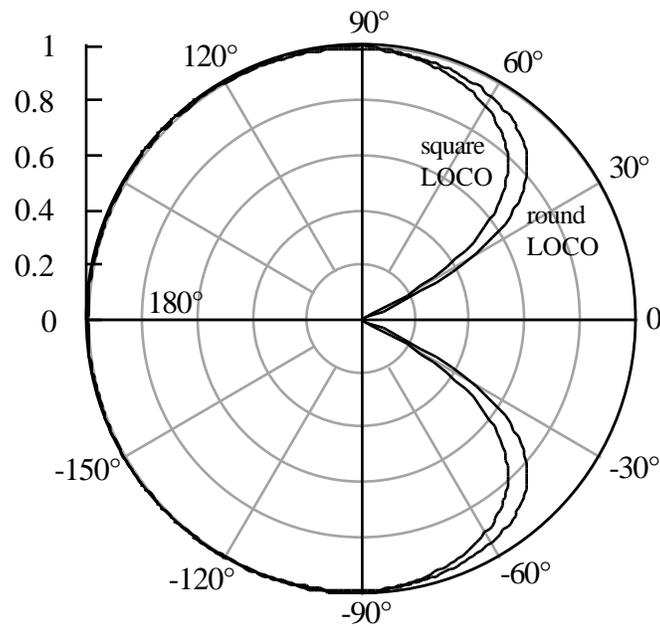


Figure 4.13. Fractional corner preservation $r(\theta, 0)$ of the LOCO filter with a circular structuring element compared to the LOCO filter with a square structuring element.

Effect of Corner Orientation Angle

Another factor that can affect the corner response of a filter is the orientation angle (θ) of the corner with respect to the filter window. For some filters, such as the plus-shaped median filter, this is a very important factor, whereas for others, such as the circular LOCO filter, it is virtually insignificant. Typically, a corner orientation angle of $\theta = 45^\circ$ produces the greatest difference in

the filter response from the $\theta = 0$ response. The difference between the response of the square median filter at $\theta = 0$ and $\theta = 45^\circ$ is shown in Figure 4.14. Most of the differences are very minor, and the corner “passband” and “stopband” are virtually unchanged. The $\theta = 45^\circ$ response was simulated the same way as the $\theta = 0$ response, with a 63×63 square window on a discrete lattice.

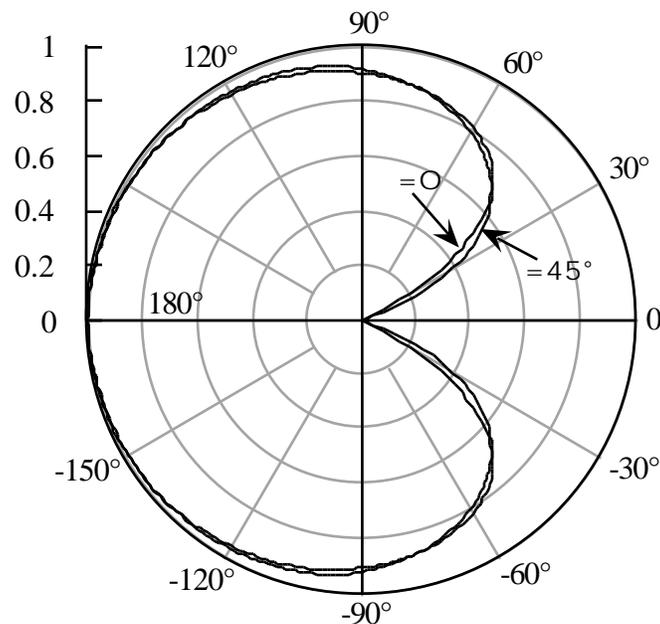


Figure 4.14. Comparison of fractional corner preservation for the square median filter at corner rotations $\theta = 0$ and $\theta = 45^\circ$.

The LOCO filter (and opening, OC, and CO) shows slightly more dependence on θ than the median filter. This is expected because of the importance of structuring element shape in morphological filtering. Figure 4.15 illustrates the change in the response of the LOCO filter with a square structuring element when θ changes from 0 to 45° . The transition and stop bands are very similar, but the “passband” of the LOCO filter shows significantly less

preservation in the region $45^\circ < \theta < 135^\circ$ for $\alpha = 45^\circ$ than for $\alpha = 0$. The “passband” of the LOCO filter at $\alpha = 45^\circ$ is approximately 44° , while the “stopband” is about 32° ; both bands are within about 1° of the LOCO filter bands for $\alpha = 0$.

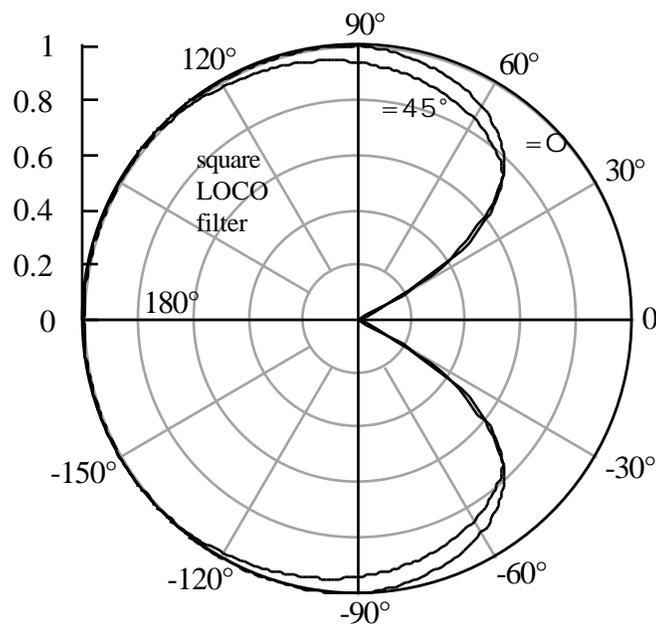


Figure 4.15. Comparison of fractional corner preservation for the square LOCO filter at corner rotations $\alpha = 0$ and $\alpha = 45^\circ$.

The plus-shaped median filter shows even more extreme changes in its response when the corner orientation angle changes from 0 to 45° . (At $\alpha = 45^\circ$, the plus-shaped median is basically equivalent to a “cross-shaped” or “X-shaped” median with $\alpha = 0$.) At $\alpha = 45^\circ$, the plus-shaped median preserves less of the corner at nearly every angle, as shown in Figure 4.16. The fractional preservation for $\alpha = 45^\circ$ was determined numerically using a 25×25 window on a discrete lattice. The perfect preservation band for this filter changes from 90° at $\alpha = 0$

to 135° at $\theta = 45^\circ$, just as for the LOCO filter. The corner “passband” for the plus-shaped median at $\theta = 45^\circ$ is about 31° , and the “stopband” is roughly 19° . The corner “passband” is similar for the two rotations, but the corner “stopband” of the plus-shaped median filter is much wider at $\theta = 45^\circ$ than at $\theta = 0$. The transition band has shrunk by more than half from 25° at $\theta = 0$ to only about 12° at $\theta = 45^\circ$.

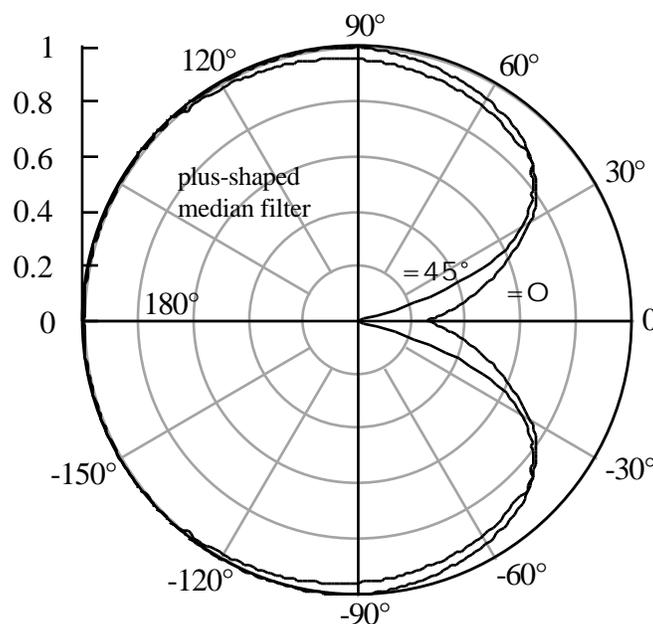


Figure 4.16. Comparison of fractional corner preservation for the plus-shaped median filter at corner rotations $\theta = 0$ and $\theta = 45^\circ$.

Effect of Filter Window Size in Discrete Space

Many of the above results were given for the “ideal” situation of continuous space or for very large discrete window sizes. Since most image processing filters use relatively small windows, it is important to investigate how

the continuous corner response curves relate to those derived for more realistic filter sizes. Figure 4.17 shows the corner response of a 5×5 square median filter compared to the 63×63 square result given earlier. Most of the differences between these two responses are discretization effects; that is, there are only a relatively small number of distinguishable angles in a 5×5 window. The overall trends and shape of the response are basically the same for small filter sizes as for large windows or for continuous space. The theoretical values for the median filter corner response are also shown (with the symbol \blacklozenge) in Figure 4.17. These values agree almost exactly with both the 63×63 and 5×5 simulations.

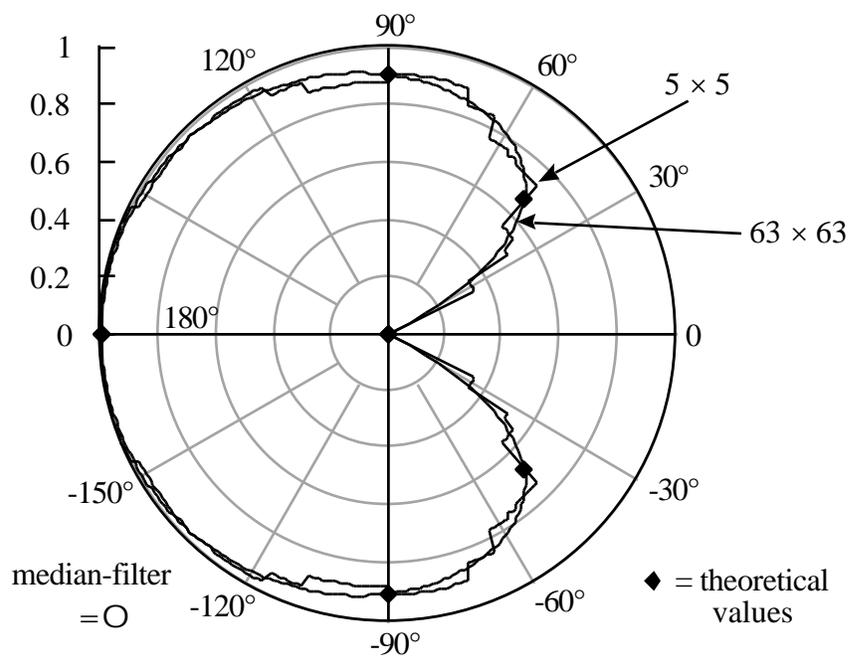


Figure 4.17. Comparison of fractional preservations of the square median filter determined with 63×63 and 5×5 square windows. Known theoretical values for the median filter shown by \blacklozenge .

Summary of Results

Comparisons of four different filters at $\theta = 0$ and at $\theta = 45^\circ$ are shown in Figures 4.18 and 4.19, respectively. The four filters are the square-shaped median, LOCO, and averaging filters and the plus-shaped median filter. Significant differences among the responses of these filters are easily observed. The corner “passband” and “stopband” information for these filters are collected in Table 4.2 below. The “perfect rejection” and “perfect preservation” bands correspond to regions of θ where $r(\theta, \omega) = 0$ and $r(\theta, \omega) = 1$ respectively.

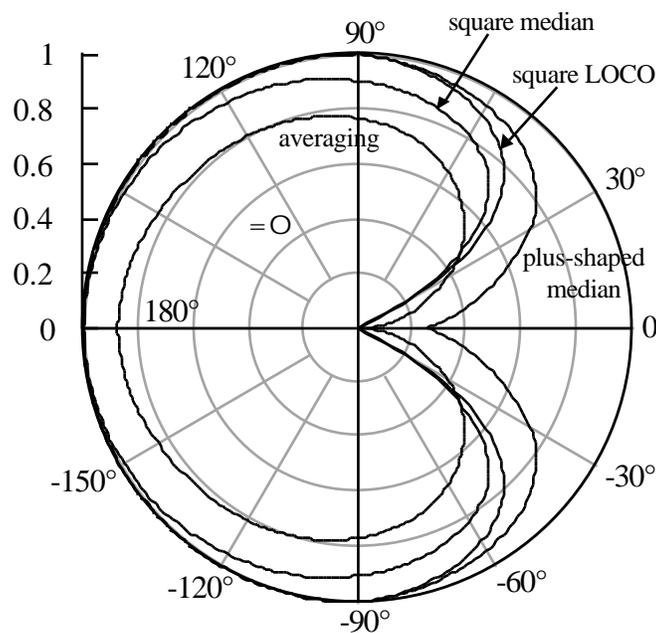


Figure 4.18. Comparison of fractional corner preservation of square-shaped median, LOCO, and averaging filters and the plus-shaped median filter at corner rotations $\theta = 0$.

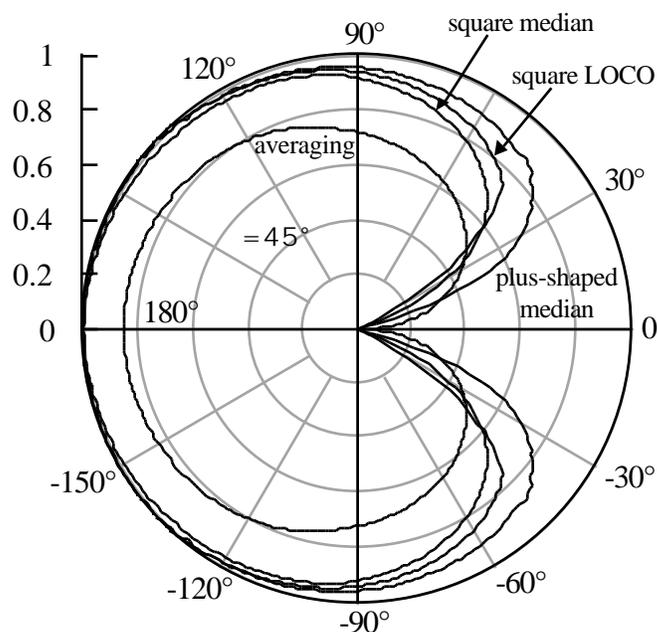


Figure 4.19. Comparison of fractional corner preservation of square-shaped median, LOCO, and averaging filters and the plus-shaped median filter at corner rotations $\theta = 45^\circ$.

Table 4.2. Comparison of filter corner responses.

Filter (Shape)	Rotation	Perfect Rejection ($r = 0$)	Stopband $s < -10.7\text{dB}$ ($r < .293$)	Transition Band Width	Passband $s > -3\text{dB}$ ($r > .707$)	Perfect Preservation ($r = 1$)
Median (Square)	$\theta = 0$	27°	31°	17°	48°	180°
	$\theta = 45^\circ$	19°	27°	22°	49°	180°
LOCO (Square)	$\theta = 0$	27°	31°	12°	43°	90°
	$\theta = 45^\circ$	27°	32°	12°	44°	135°
Averaging (Square)	$\theta = 0$	0^*	24°	47°	71°	—
	$\theta = 45^\circ$	0	17°	69°	86°	—
Median (Plus)	$\theta = 0$	—*	4°	25°	29°	90°
	$\theta = 45^\circ$	14°	19°	12°	31°	135°

* $r(\theta, 0) = 0$ as 0 for the averaging filter, but not for the plus-shaped median filter.

The values in Table 4.2 above indicate that the square-shaped LOCO filter has the sharpest transition between corner removal and corner preservation for $\theta = 0$, and that its response changes the least when θ changes from 0 to 45° . Thus, the LOCO filter is the closest of these four filters to an “ideal” filter with a very small transition band and a corner response that changes little for different corner orientations. However, there are usually many considerations taken into account other than just the response at corners when selecting a two-dimensional filter. For example, the degree of noise reduction offered by the filters in Table 4.2 varies, as does the type of noise they perform best against. Also, the corner response characteristics of opening, OC, CO, and the LOCO filter are all identical, despite obvious visual differences among the results of the filters in real image processing applications. Clearly, the corner response is not a complete description of the behavior of a filter.

Although the corner response technique is a valuable analysis tool, it does not indicate how to design a filter with given corner response characteristics. The cutoff angles of filters vary with the angle of rotation θ and the shape of the filter, but usually the cutoff angle cannot be drastically changed for a given filter. Also, the fractional preservation $r(\theta, \theta_0)$ only indicates the amount of area preserved, not any change in the shape of that area. Some filters, especially the median and morphological filters, achieve partial preservation of corners by changing the shape of the corners: the median filter rounds off sharp corners, while morphological operators tend to make corners look more like their structuring element. Thus a 45° corner filtered by open-closing with a square structuring element is 75% preserved [$r(45^\circ, 0) = 0.75$], but the result may look like a more

obtuse angle because this operation tends to clip corners to make them as close to right angles as possible.

4.5. CONCLUSIONS

This chapter introduced three new methods for analyzing the behavior of nonlinear filters. These analysis techniques demonstrate some important differences among the morphological, median, averaging, and MLV filters. Continuous time analysis yields the peak response of filters to periodic signals of different frequencies. This analysis highlights the differences between the morphological and non-morphological filters in the presence of rapidly fluctuating periodic signals or noise. The breakdown point gives a measure of the resistance of a filter to large outliers in its input. Using order statistics to limit the range of outputs of a filter allows custom filters with almost any breakdown point to be designed.

The last technique developed in this chapter analyzes the behavior of two-dimensional filters at corners. This method finds the fraction of corners of different angles and orientations preserved by a filter, and is important because it is one of very few techniques available to analyze the behavior of filters on two-dimensional structures. The corner response of a filter gives an indication of whether the filter is more likely to preserve details and sharp features in an image or remove them. The analysis techniques introduced in this chapter provide a significant improvement in understanding the behavior of nonlinear filters, especially their response to high frequency periodic signals, impulse noise, and two-dimensional corners.

5. Applications

5.1. INTRODUCTION

Nonlinear filtering has become increasingly important in many image processing applications. Initially, the attraction to nonlinear filters was mostly limited to the impulse-removing and edge-preserving qualities of the median filter [4, 53]. However, as the number and sophistication of nonlinear filters have increased, so has the variety of applications for these filters. The shape-based methods of mathematical morphology, in particular, are now used in a wide variety of medical applications, including electrocardiography [54, 55], ultrasound imaging [56-58], radiology [59-61], and histologic image analysis [62, 63].

For some biomedical imaging applications, the bias introduced by the morphological operations is not a problem. For example, most segmentation and feature extraction problems do not require accurate estimation of the original gray levels in an image, and some imaging modalities (such as ultrasound) are not adversely affected by the bias introduced by morphological filtering [56, 57]. However, the true gray levels of certain types of biomedical images are important. Thermography is one example of this type of application; the gray levels in thermal images can be used to find temperatures in an image. This is particularly important in images used to analyze thermal effects of electromagnetic irradiation (by laser or radio frequency electrosurgery) on tissue. For these images, even slight distortions of the actual gray levels by filtering to remove noise and artifacts are unacceptable. When the shape control of mathematical morphology is desired

in these applications, the linear combinations of morphological operators largely overcome the problems of the standard filters.

For segmentation problems, the shape control of mathematical morphology also holds much promise. Many segmentation and feature extraction problems involve finding particular shapes within images. However, if an image is noisy and the edges between regions are not distinct, standard morphological filters do not enhance the image in a way that improves segmentation performance. The deterministic and statistical properties of these filters described in previous chapters indicate that the morphological operations do not sharpen edges and are not particularly good at removing noise, either. These limitations are overcome by the value-and-criterion filter structure, which includes filters (such as the MLV filter) that sharpen edges and remove noise quite well.

This chapter shows how the nonlinear filters developed in the previous chapters may be used for restoration of noisy thermal images and for enhancement of magnetic resonance images. Linear combinations of morphological operators are useful when shape-based filtering is needed for thermal images, since the bias of the single operators alone cannot be tolerated in calibrated thermal images. The value-and-criterion filter structure can be used to design edge-enhancing filters that increase the sharpness of magnetic resonance images while reducing noise. The use of these filters is also compared to conventional morphology for ultrasound imaging. These applications also illustrate the potential utility of shape-based nonlinear filters in more complex algorithms for biomedical image analysis.

5.2. THERMOGRAPHY

5.2.1. Thermal Image Noise Processes

Thermal images are subject to a variety of noise processes due to the photodetector and the nature of thermal radiation [53]. The resulting overall noise process is often described as “heavy-tailed Gaussian,” meaning that the noise is roughly additive Gaussian noise with some additional impulses present. However, the noise present in thermal images is not overwhelmingly impulsive, as will be shown below. Hogg [64] and David [25] defined a statistic q that is a good indicator of the tail length of a distribution. This statistic is given in equation (5.1) below [25].

$$q = \frac{\bar{u}(0.05) - \bar{l}(0.05)}{\bar{u}(0.5) - \bar{l}(0.5)} \quad (5.1)$$

where $\bar{u}(\)$ = mean of the largest n order statistics of the distribution
 $\bar{l}(\)$ = mean of the smallest n order statistics of the distribution
 n = number of samples taken from the distribution

The statistic q is proposed by Hogg as a substitute for the sample kurtosis as a measure of impulsivity: a high value of q indicates that the noise process is impulsive.

Ryu estimated the value of q for the noise process in thermal images to be $q = 2.32$, compared to a value for additive Gaussian noise of $q = 2.19$ [53]. However, Hogg quotes a value of $q = 2.58$ for the normal (Gaussian) distribution [64]. Ryu indicated that the thermal image noise process is slightly more impulsive than Gaussian noise, but the results of Hogg indicate that it is slightly *less* impulsive. Regardless, the thermal image noise process is not very impulsive, but is actually very similar to Gaussian noise. The double-exponential

distribution, a standard model for very heavy-tailed (impulsive) noise, has $q = 3.30$.

5.2.2. Filter Selection

Hogg [64] suggests using q -trimmed means for estimating the center of symmetric distributions based upon the value of q . This robust estimation technique can be used to remove noise from thermal images in areas where the temperature is constant, since the noise is symmetric and the center of the distribution corresponds to the true temperature. For $2.0 < q < 2.6$, Hogg recommends the standard sample mean (that is, the averaging filter) as the best estimator of the center of the distribution. Only for very large q (greater than 3.2) does he recommend severely trimmed means or the median filter. However, there are other considerations to take into account in thermal images that affect the filtering choice. The averaging filter and (to a lesser extent) the q -trimmed means blur step edges in images; therefore, in regions of rapid temperature change, these filters introduce unacceptable distortion. In these instances, the median filter may be preferred because it is less susceptible to distortion introduced by temperature differences among nearby pixels.

To illustrate the various signal and noise types for which the averaging, median, and trimmed mean filters work best, consider three different signals of 1000 points: a constant signal c with additive white Gaussian noise with zero mean and unit variance $(0,1)$, denoted $f_1 = c + (0,1)$; a constant signal c with additive Gaussian noise $(0,1)$ and 5% (that is, 50 points) impulsive noise $i(0.05)$, denoted $f_2 = c + (0,1) + i(0.05)$; and a signal e that has 19 perfect step edges with

50 points separating each edge with additive Gaussian noise $(0,1)$, denoted $f_3 = e + (0,1)$. The impulses i have amplitude 5, and the step edges in e have amplitude 10. The noisy signals f_1 , f_2 , and f_3 are filtered by 5-wide averaging, median, trimmed mean ($\alpha = 0.2$), and LOCO filters, and the mean square error (MSE) between the filtered noisy signals and the original (noiseless) signal is computed. The LOCO filter with a 3-wide structuring element is included for comparison purposes, as the morphological filters are not a standard robust estimation technique. The results are summarized in Table 5.1 below.

Table 5.1. Mean square error (MSE) of filtered noisy signals (1000 points).

Original Signal	Noisy Signal	MSE of Noisy Signal	MSE of Averaging Filtered Noisy Signal	MSE of Median Filtered Noisy Signal	MSE of Trimmed Mean Filtered Noisy Signal ($\alpha = 0.2$)	MSE of LOCO Filtered Noisy Signal
c (const)	$f_1 = c +$	1.062	0.175	0.302	0.212	0.240
c (const)	$f_2 = c + i +$	2.406	0.444	0.343	0.304	0.358
e (edges)	$f_3 = e +$	1.062	0.921	0.338	0.660	0.259

These results indicate the situations under which each of these filters performs best. The averaging filter is optimal in the least squares sense for a constant signal with additive Gaussian noise, and its MSE is smallest of the filters in that particular case. However, when 5% impulsive noise is also added to the constant signal, the averaging filter does not perform as well as the other filters.

The trimmed mean filter is better than the median filter in this case because the trimming is enough to remove the impulses, but the mean operation removes the Gaussian noise better than the median operation. When the signal has perfect step edges, though, the median filter performs better because both the trimmed mean and averaging filters blur the edges slightly. Note that the performance of the median filter is good in all three cases; however, the benefits of the median are not apparent unless sharp edges or other such features are present in the input. For slowly varying signals, the standard averaging filter removes non-impulsive noise well, and the trimmed mean filter is better for removing more impulsive noise. The LOCO filter performs surprisingly well for all three signal and noise types shown above, and its results are the best of any of the filters for the signal with sharp edges. This is because the shape of the signal conforms with the shape of the structuring element. Note that the LOCO filter is better than the median filter at removing Gaussian noise (but not as good as the averaging filter), and worse than the median filter at removing impulsive noise (but better than the averaging filter).

Ryu [53] found that 5-wide median filtering on the horizontal lines of thermal images removed the noise in the images effectively, and that a 3-wide linear filter with triangular weights (0.25, 0.50, and 0.25) acting vertically smoothed the edge jitter associated with the images. This technique does not introduce any bias to the image, which is important in quantitative applications of thermography. Another potential filtering technique for thermal images is LOCO filtering, since it also does not introduce any bias, but does remove Gaussian and impulse noise while preserving sharp edges. Morphological filtering is also

promising for thermal images because the shape of the structuring element can be selected to specify what types of shapes should be preserved in the images, thus removing noise and vertical edge jitter at the same time.

Since the noise process in thermal images is not overwhelmingly impulsive (as demonstrated by the values of q given in Section 5.2.1 above), it may be worthwhile to examine the use of the standard and trimmed means instead of the median in Ryu's filtering technique. Also, the effect of LOCO filtering with various structuring element shapes and weights corresponding to thermal image features will be explored.

5.2.3. Examples

Filtering

An unfiltered thermal image of a portion of porcine skin heated by laser illumination is shown in Figure 5.3 below. This image is shown in pseudocolor of 20 levels to improve the visibility of the isothermal contours in the image. Noise is very prominent in the image, and edge jitter is visible along the isothermal contours between pseudocolor levels. Figures 5.4–5.10 show the image after filtering by various schemes outlined above. The result of filtering by Ryu's combined median and linear filtering scheme is shown in Figure 5.4. The noise is reduced quite well in this scheme, but the edge jitter is still visible along some of the isothermal contours. The values at the peak of the laser spot are somewhat reduced, but this filtering method does not severely distort the values at the peak.

Figures 5.5 and 5.6 illustrate the results of morphological (LOCO) filtering with set (flat) structuring elements. A square structuring element of size

3×3 was used to filter the image in Figure 5.5. This structuring element is not a good match for the features in the original image, which are roughly circular. Square LOCO filtering produces a “blocky” result; the square structuring element tends to produce 90° corners along the isothermal contours and squared-off features instead of rounded ones. A logical choice for the structuring element shape for this image is a circle. Unfortunately, it is difficult to approximate a circle with a small structuring element in discrete space. The best approximation of a reasonably small size is a 5×5 square with the four corner pixels removed, as illustrated in Figure 5.1 below. The result of LOCO filtering the original image with this structuring element is shown in Figure 5.6. The isothermal contours are smooth and the noise has been removed very effectively, but the values at the peak of the laser spot are attenuated significantly. This is because the structuring element is a set, and therefore tries to produce flat regions in the image that are the same shape as the structuring element. This is exactly what is seen in the center of the spot in Figure 5.6.

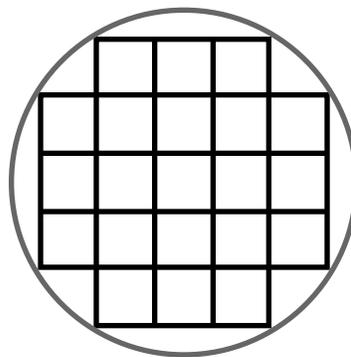


Figure 5.1. “Round” structuring element (5×5).

One way to overcome this limitation of set structuring elements is to use a function, or “weighted,” structuring element. By selecting weights for each pixel in a structuring element, amplitude features such as the peak in the laser spot image can be preserved. Figure 5.7 illustrates the result of LOCO filtering using a 5×5 “round” structuring element with weights designed to preserve the peak of the spot. These weights are given in Figure 5.2 below. Note that a weighted structuring element tries to create areas with the same absolute difference in amplitudes as its weights. The result of this weighted LOCO filtering (Figure 5.7) shows much better preservation of values at the peak of the signal, but noise reduction and edge jitter are more visible as well. The structuring element weights allow impulses to be partially preserved in an image, since the prototype provided by the structuring element is not flat.

	0	2	0	
0	4	5	4	0
2	5	6	5	2
0	4	5	4	0
	0	2	0	

Figure 5.2. Weighted 5×5 “round” structuring element.

Two-dimensional median filters are usually not used on thermal images despite their robust noise removal and edge-preserving properties, because they tend to over-smooth the image. For comparison purposes, the result of median filtering the laser spot image using a 5×5 square filter window is shown in Figure 5.8. Although the result is very rounded, this may not be the most accurate

representation of the shapes in the original image. The tendency of the 2-D median filter to round off corners was analyzed earlier (see Chapter 4) and this effect is clearly visible in Figure 5.8.

Finally, Figures 5.9 and 5.10 show the effect of replacing the median filter in the filtering technique of Ryu with a trimmed mean and a simple mean, respectively. As shown in Section 5.2.1, the overall noise process of thermal images is not as impulsive as perhaps it was thought to be previously. Hogg [64] and other robust estimation researchers have suggested that a trimmed or a standard mean may be better than the median for estimating the center of a noise process with these characteristics. The result of using a horizontal trimmed mean filter of width 5 with $\alpha = 0.2$ (Figure 5.9) seems quite smooth, with excellent noise removal and reasonable rejection of edge jitter. However, using a horizontal standard mean (averaging) filter of width 5 (Figure 5.10) gives a result that does not suppress the vertical edge jitter as well as most of the other filters examined.

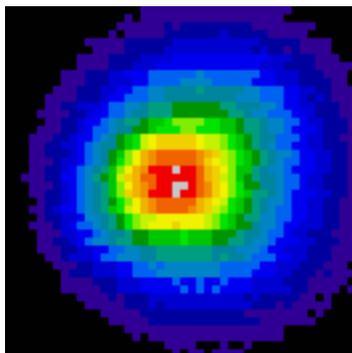


Figure 5.3. Original image of laser heated porcine skin.

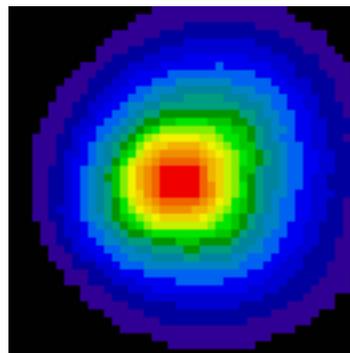


Figure 5.4. Image filtered by 5-wide horizontal median filter and 3-wide vertical triangular filter (Ryu [53]).

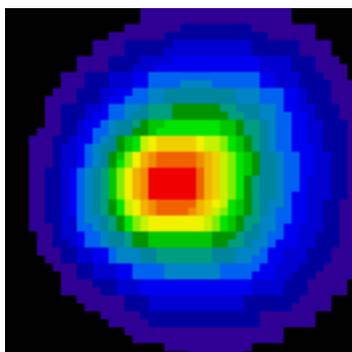


Figure 5.5. Image filtered by 3×3 square LOCO filter.

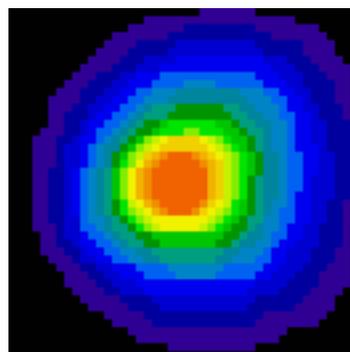


Figure 5.6. Image filtered by 5×5 "round" LOCO filter.

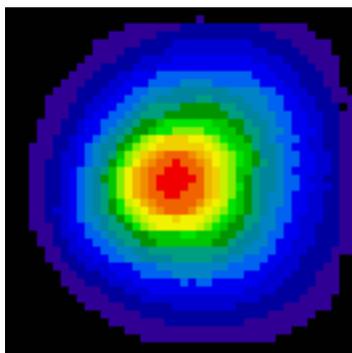


Figure 5.7. Image filtered by weighted 5×5 "round" LOCO filter.

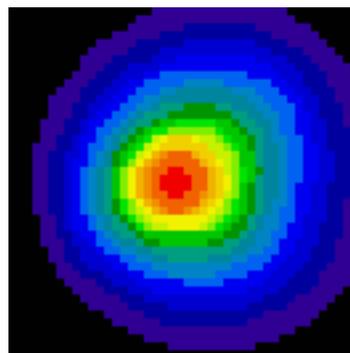


Figure 5.8. Image filtered by 5×5 square median filter.

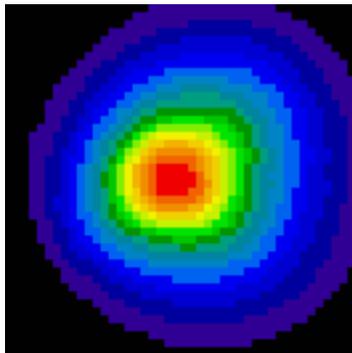


Figure 5.9. Image filtered by 5-wide horizontal trimmed mean filter ($\alpha=0.2$) and 3-wide vertical triangular filter.

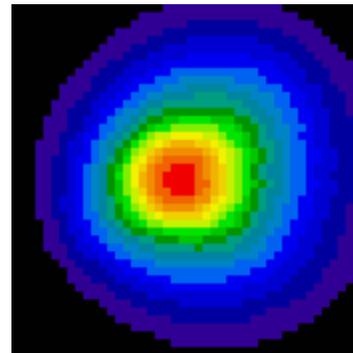


Figure 5.10. Image filtered by 5-wide horizontal mean filter and 3-wide vertical triangular filter.

Evaluating Filter Performance

Although the differences among the results of the seven different filtering techniques are readily visible, there is still a need to determine how well each of the various filtered images represents the “truth” image. As for most biomedical images, the “truth” image in this case is unknown and unknowable. However, some reasonable assumptions can be made about the temperature distribution imaged here that make it possible to quantify the performance of the various filters on this image. These simplifying assumptions are that the laser irradiation is normal to the skin surface, that the scene is viewed by the camera at a 90° angle, and that the thermal properties of the skin are relatively isotropic. Note that as with any biomedical experiment, these assumptions are not strictly true and the “truth” of the experiment will vary significantly from that deduced from the assumptions.

Using the above simplifying assumptions, the isothermal contours in the “truth” image are assumed to be circular, with each centered at the point of

incidence of the laser irradiation. Therefore, one way of quantifying the performance of the various filters is to measure the degree to which various isothermal contours in the filtered images deviate from a circle. To measure this deviation, the image is thresholded at a given gray level, and the distances from the centroid of the resulting shape to its boundary are measured. The standard deviation of these distances is measured and reported as a percentage of the mean distance, so that the measure is consistent for different circle sizes. For the original image, the standard deviation of the radius ranged from 6.45% at a gray level of 70 (near the edge of the spot) to 17.87% at a gray level of 195 (near the peak of the spot). The range of gray levels for the original image shown in Figure 5.1 is 20 to 221. To quantify the improvement made by a particular filter, the standard deviation of the radius is measured for gray levels from 70 to 195 in increments of five levels, for a total of 26 measurements. Since the standard deviation of the radius increases as the gray level gets higher, the percentage that each filter reduces the standard deviation is computed for each of the 26 measurements, and the average reduction for each filter is then computed. The mean extent to which a filter reduces the standard deviation of the radius indicates how much the filter improves the circularity of the isothermal contours in the image.

The results of this analysis are summarized in Table 5.2 below. The combined linear and nonlinear filtering technique of Ryu shows the largest reduction in the standard deviation of the radius, indicating that it generally gives the most circular results. This is somewhat surprising, considering that the two-dimensional median filter and the round LOCO filter both have a strong tendency

to produce smooth, circular contours. Much of the credit for the superior performance of Ryu's technique belongs with the vertical linear filter, since similar reductions in the standard deviation are observed when the horizontal median filter is replaced with the trimmed mean or simple averaging filter. The square LOCO filter does not do a good job of improving the circularity of the isothermal contours; this is confirmed by its small reduction in the standard deviation of the radius.

Table 5.2. Mean percentage reduction of standard deviation of isothermal contour radius for various thermal image filtering techniques.

Filter	Standard deviation of radius at gray level 70 (as a % of mean radius)	Standard deviation of radius at gray level 195 (as a % of mean radius)	Mean percentage reduction of standard deviation of radius
(none)	6.45%	17.87%	—
Ryu technique (5-wide horizontal median + 3-wide vertical triangular)	5.87%	15.16%	9.48%
Square LOCO 3 × 3	5.96%	15.16%	1.78%
Round LOCO 5 × 5	5.70%	14.43%	5.69%
Weighted round LOCO 5 × 5	5.77%	14.43%	6.95%
Median 5 × 5	6.21%	15.75%	7.36%
5-wide horizontal trimmed mean ($\alpha=0.2$) + 3-wide vertical triangular	5.54%	15.16%	9.23%
5-wide horizontal mean + 3-wide vertical triangular	5.61%	16.63%	8.41%

The assumption that the isothermal contours of these images are circular is probably not realistic, since the laser beam and thermal camera are unlikely to be perfectly normal to the surface. Therefore, the contours are more accurately described by ellipses. For the images in Figures 5.3–5.10, the best-fit ellipses for most of the contours have eccentricity equal to approximately 0.4. This means that the minor axis of the ellipse is about 90% as long as the major axis. This deviation from circular is not severe, so the conclusions of the above analysis are probably not significantly altered by the circularity assumption. However, to verify the results further, a “convexity” measure will be used to quantify the smoothness of the contours. This measure is a ratio of the area inside the contour to its perimeter squared, normalized so that for a perfect circle the ratio is 1. This “convexity” measure, C , is defined in equation (5.2) below, and it can effectively distinguish between elliptical shapes with smooth and rough borders. An ellipse with eccentricity equal to 0.4 has a value for the convexity measure of approximately 0.996, very close to that of a true circle. A square has a convexity measure of about 0.785.

$$C = 4 \frac{A}{p^2} \quad (5.2)$$

This convexity measure is computed at 13 different gray levels from 70 to 190 (in increments of 10 gray levels) on each of the images in this section. The results are summarized in Table 5.3 below. Using this measure, the performance of Ryu’s technique is not as clearly superior to the other methods. The square median filter and the methods using the trimmed and standard horizontal mean filters all give higher values for the convexity on average than Ryu’s method. The

5×5 round LOCO filter also performs nearly as well as the established techniques. However, the square LOCO filter again is shown to be unsuited to this application.

Table 5.3. Convexity measure of isothermal contours for various thermal image filtering techniques.

Filter	Convexity measure (C) at gray level 70	Convexity measure (C) at gray level 190	Mean convexity measure for gray levels 70–190 (increments of 10)
(none)	0.837	0.887	0.856
Ryu technique (5-wide horizontal median + 3-wide vertical triangular)	0.924	0.955	0.944
Square LOCO 3×3	0.917	0.952	0.908
Round LOCO 5×5	0.940	0.972	0.935
Weighted round LOCO 5×5	0.928	0.972	0.933
Median 5×5	0.948	0.969	0.948
5-wide horizontal trimmed mean ($\tau=0.2$) + 3-wide vertical triangular	0.941	0.967	0.947
5-wide horizontal mean + 3-wide vertical triangular	0.948	0.967	0.948

Since it is not possible to find a “truth” image corresponding to an acquired thermal image, there is no way to tell from the above analysis how accurately the filters estimate the temperatures from the acquired images. To check the accuracy of the filters, a control image that provides a known truth for

comparison is needed. Since the irradiation and heating profiles of a laser spot are approximately Gaussian, a reasonable phantom image for this application is a Gaussian intensity profile. Figure 5.11 is such a phantom, a 48×48 image of a Gaussian profile with a standard deviation of 7 pixels. The amplitude of the profile is 200 gray levels, and the background level is at a gray level of 20, so the range of gray levels in the phantom is 20 to 220. Figure 5.12 is the phantom with normally distributed noise added with a standard deviation of 4.5. This is slightly higher than the noise level observed in the acquired images shown previously. Figures 5.13–5.19 are the results of filtering the noisy phantom with the various schemes described above, shown in pseudocolor to enhance the visibility of the isothermal contours.

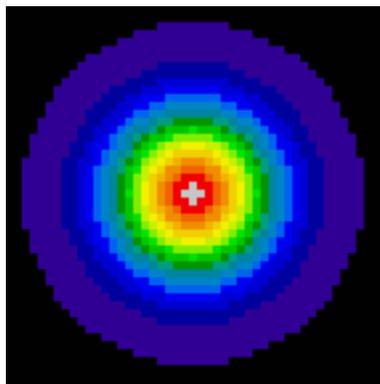


Figure 5.11. Phantom image (Gaussian profile).

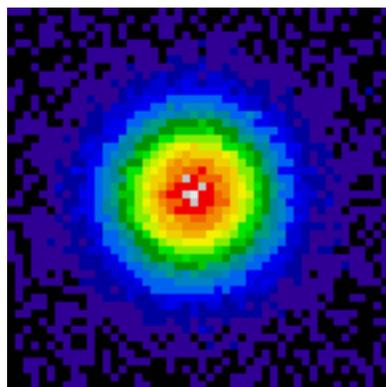


Figure 5.12. Noisy phantom image (Normal density, $\sigma = 4.5$).

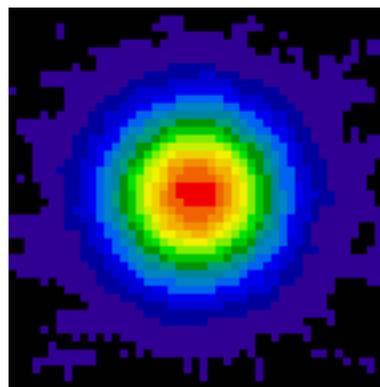


Figure 5.13. Phantom filtered by 5-wide horizontal median filter and 3-wide vertical triangular filter (Ryu [53]).

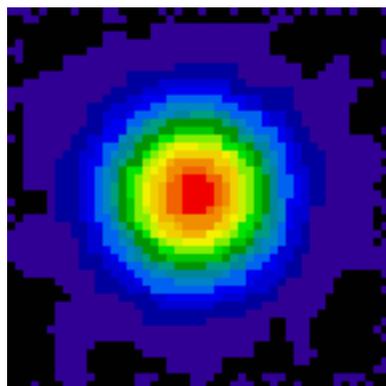


Figure 5.14. Phantom filtered by 3×3 square LOCO filter.

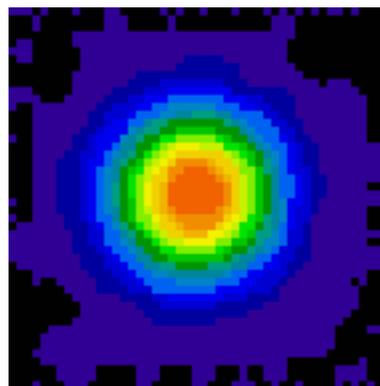


Figure 5.15. Phantom filtered by 5×5 "round" LOCO filter.

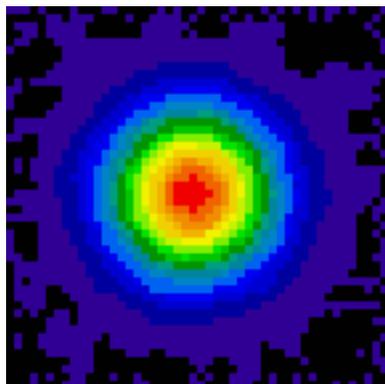


Figure 5.16. Phantom filtered by weighted 5×5 “round” LOCO filter.

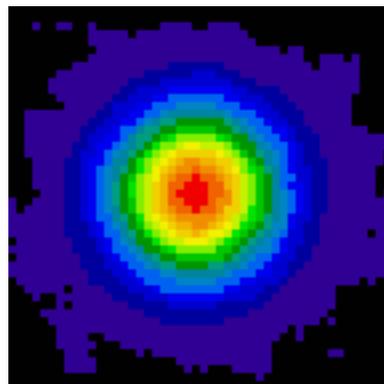


Figure 5.17. Phantom filtered by 5×5 square median filter.

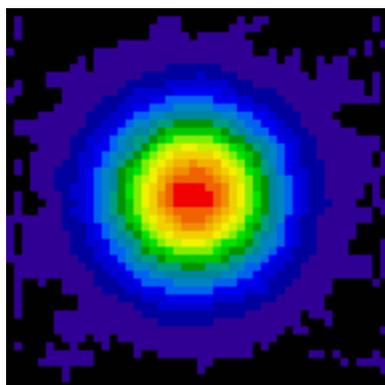


Figure 5.18. Phantom filtered by 5-wide horizontal trimmed mean filter ($\tau=0.2$) and 3-wide vertical triangular filter.

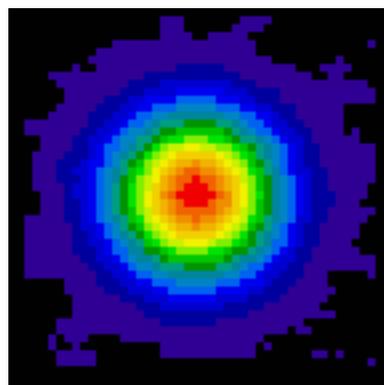


Figure 5.19. Phantom filtered by 5-wide horizontal mean filter and 3-wide vertical triangular filter.

The quality measures used to estimate filter performance for the acquired images can also be used to evaluate performance on the phantom image. Table 5.4 below summarizes the mean percentage increase in the standard deviation of the radius of the contours in the filtered phantom images for gray levels 70 to 195. The standard deviation is expressed as an increase this time because the filtered images are compared to the noiseless original image. Table 5.4 also lists the

mean convexity measure for contours at levels 70 to 190. The convexity measure, as expected, shows that the original phantom image has the most convex contours ($C = 0.97$). The contours of the original phantom image are the closest approximation to circles possible in the digital domain. The noisy phantom has an average convexity measure of about 0.90, and the various filtering algorithms raise the convexity to values mostly around 0.95. The standard deviation of the radius for the noisy phantom image is nearly 10% higher than for the noiseless phantom image, but the filters are all effective at reducing this increase to the range of 2% to 6%. The relative performance of the filters indicated by these measures is slightly different from the performance measure used for the acquired images. The trimmed and horizontal mean methods and square median filter seem to exhibit the best performance on the phantom image using these measures, followed by Ryu's algorithm and the LOCO filters.

Table 5.4. Mean percentage increase in standard deviation of “isothermal” contour radius and mean convexity measure for various filtering techniques on the thermal image phantom.

Filter	Mean percentage increase in standard deviation of radius from noiseless phantom (gray levels 70–195 in increments of 5)	Mean convexity measure for gray levels 70–190 (increments of 10)
(original phantom)	—	0.966
(noisy phantom)	9.96%	0.903
Ryu technique (5-wide horizontal median + 3-wide vertical triangular)	3.03%	0.953
Square LOCO 3 × 3	6.02%	0.946
Round LOCO 5 × 5	4.01%	0.951
Weighted round LOCO 5 × 5	3.65%	0.948
Median 5 × 5	4.50%	0.957
5-wide horizontal trimmed mean ($\alpha=0.2$) + 3-wide vertical triangular	2.33%	0.961
5-wide horizontal mean + 3-wide vertical triangular	3.48%	0.961

The main advantage of evaluating filter performance on a phantom image is that the “correct” isothermal contours are known and thus the accuracy of the filters in restoring these contours from the noisy image can be measured. The area enclosed by an isothermal contour at a given level is compared to the “ideal” area enclosed by the contour at the same level on the noiseless phantom image. “False positives” are points that are included in the filtered contour at a given

level but are not in the original contour at the same level. Similarly, “false negatives” are points that are included in the original contour at a given level but not in the filtered contour at the same level. The numbers of false positives and false negatives are given as percentages of the total area of the true contours in the noiseless phantom image. The average percentages of false positives and false negatives for the various filtered images for 25 different gray level thresholds from 70 to 190 in 5-level increments are given in Table 5.5. This table points out important limitations of some of the filtering algorithms. For example, the 5×5 median filter has a very low number of false positives, but an extremely high number of false negatives. This indicates that the median filter substantially underestimates the area of the isothermal contours. This is one reason that the standard 2-D median filter is not recommended for thermal image processing. The other measures (standard deviation of the radius and convexity) indicated that the median filter performed quite well. However, the phantom image analysis shows that although the median filter yields approximately correct shapes for the contours, it often significantly underestimates the size of the contours. A similar problem afflicts the algorithms with the trimmed and standard horizontal means; however, the LOCO filters and Ryu’s algorithm are much more balanced in their estimation of the contour areas.

Table 5.5. Mean percentage of false positive and negative identifications of points inside “isothermal” contour radius and mean convexity measure for various filtering techniques on the thermal image phantom.

Filter	False Positives (percentage of area)	False Negatives (percentage of area)
(original phantom)	0	0
(noisy phantom)	4.17%	2.29%
Ryu technique (5-wide horizontal median + 3-wide vertical triangular)	1.26%	1.76%
Square LOCO 3×3	2.64%	2.63%
Round LOCO 5×5	1.83%	2.91%
Weighted round LOCO 5×5	1.66%	2.72%
Median 5×5	0.57%	4.50%
5-wide horizontal trimmed mean ($\alpha=0.2$) + 3-wide vertical triangular	0.47%	1.93%
5-wide horizontal mean + 3-wide vertical triangular	0.33%	3.01%

Summary

Although the LOCO filter does not provide a quantitative improvement over previous filtering techniques for thermal images, the shape control it provides is a tool that is potentially useful in thermography. Thermography also illustrates the tradeoffs inherent in morphological filtering, particularly between noise reduction and preservation of amplitude features such as signal peaks. The weighted LOCO filter generally does a better job of preserving the shape of the temperature profile than the unweighted LOCO filter does, but the weighted filter

does not remove noise as well as the unweighted filter does. The importance of choosing the correct shape for the structuring element in morphological filtering is underscored by the poor performance of the square structuring element shape on the circular features in the laser spot image. The examples shown above also demonstrate that the median filter is a good estimator of the center of the distribution of the noisy image gray levels despite indications from robust estimation research that a trimmed or standard mean might be more appropriate.

5.3. MAGNETIC RESONANCE IMAGING (MRI)

5.3.1. Introduction

Magnetic resonance imaging, or MRI, is based on the phenomenon of nuclear magnetic resonance that is also used for spectroscopy in chemistry and physics. One advantage that MRI has over standard X-rays and X-ray based computed tomography techniques is that the magnetic resonance signal is able to penetrate bone and other radio-opaque tissues with very little attenuation or distortion [65]. MRI is also a fundamentally three-dimensional technique, so the data can be viewed in any desired section. MR images are noted for their excellent soft tissue contrast and low noise. However, the amount of noise in the image is related to the acquisition time for the image. Imaging of physiological events in “real time” requires rapid acquisition, so images acquired under these circumstances are noisier than other MR images.

A common image analysis task for magnetic resonance images of the brain is to separate the different types of tissues into different segments [66]. For

example, the relative amounts of white matter, gray matter, and cerebrospinal fluid may be of diagnostic interest for a variety of reasons (*e.g.*, hydrocephalus), and the presence and amount of any abnormal tissue may also be of interest. The MLV filter works well as a pre-filter for this problem, since it smoothes noise in homogeneous regions and sharpens the boundaries between regions. In addition, post-processing for noise reduction and contrast enhancement of MR images is of clinical interest because it allows the quality of images taken with faster acquisition times to be enhanced to an appropriate level.

Gerig, *et al.* [14] used a technique called anisotropic diffusion to reduce noise and improve the contrast of magnetic resonance images. The effects of this algorithm are similar to those of the MLV filter; however, anisotropic diffusion requires many iterations to yield a good result, while the MLV filter produces comparable improvement with only one pass. Both techniques require the selection of parameters that affect the output. Anisotropic diffusion requires choosing the number of iterations (since one iteration is typically not enough to provide much improvement) and a parameter that affects both the level of noise smoothing and the ability of the filter to preserve edges. For the MLV filter, the size and shape of the structuring element must be specified. The size of the structuring element directly influences the degree of noise smoothing by the filter, and the size and shape both affect the preservation of fine details in the image.

5.3.2. Anisotropic Diffusion

Definition

The use of anisotropic diffusion for image filtering is described by Perona and Malik [67, 68]. Anisotropic diffusion tends to smooth across regions with small differences between pixels while preserving large differences. Each point in the image is changed by an amount related to the divergence of the gradient of the image at that point multiplied by a diffusion function. If $f(\mathbf{x}, t)$ denotes a family of images where \mathbf{x} is the space variable and t indicates the iteration, filtering by anisotropic diffusion is given by equation (5.3) below.

$$f(\mathbf{x}, t + \Delta t) = f(\mathbf{x}, t) + \text{div} [c(\mathbf{x}, t) \nabla f(\mathbf{x}, t)] \quad (5.3)$$

The symbol ∇ denotes the gradient operator, and “div” denotes the divergence operator. The function $c(\mathbf{x}, t)$ is the diffusion function, and Δt is an integration constant that affects the stability of the iterative process. The diffusion function c is monotonically decreasing as the gradient of f increases; the most commonly used diffusion function is given in equation (5.4) below. The parameter λ in equation (5.4) is the diffusion parameter that specifies which edge slopes will be preserved and which will be smoothed.

$$c(\mathbf{x}, t) = \exp \left[- \frac{|\nabla f(\mathbf{x}, t)|^2}{\lambda} \right] \quad (5.4)$$

In the one-dimensional discrete case, the gradient of f is estimated by taking the differences of neighboring points, and the divergence term in equation (5.3) is the difference of two weighted differences, one to the left of the center point (l) and one to the right (r), as given in equations (5.5) and (5.6) below.

$$l = \frac{1}{x^2} \exp - \frac{f(\mathbf{x},t) - f(\mathbf{x} - \mathbf{x},t)}{x^2} (f(\mathbf{x},t) - f(\mathbf{x} - \mathbf{x},t)) \quad (5.5)$$

$$r = \frac{1}{x^2} \exp - \frac{f(\mathbf{x} + \mathbf{x},t) - f(\mathbf{x},t)}{x^2} (f(\mathbf{x} + \mathbf{x},t) - f(\mathbf{x},t)) \quad (5.6)$$

Using the above equations, equation (5.3) reduces to:

$$f(\mathbf{x},t + t) = f(\mathbf{x},t) + t (r - l) \quad (5.7)$$

For anisotropic diffusion to be stable in one dimension, the integration constant t must be no greater than $\frac{1}{3}$ [14].

In discrete two-dimensional space, the gradient can be estimated using four or eight neighbors. Using four neighbors, the iteration equation becomes:

$$f(\mathbf{x},t + t) = f(\mathbf{x},t) + t (e - w + n - s) \quad (5.8)$$

The terms e , w , n , and s represent the weighted differences between the center pixel and its east, west, north, and south neighbors, respectively. These terms are formulated as in equations (5.5) and (5.6).

Three parameters must be chosen when implementing anisotropic diffusion: the diffusion parameter σ , the integration constant t , and the number of iterations. Selection of the integration constant is usually not difficult, since any value less than the maximum allowable value usually works well. However, if a small value of t is chosen, the number of iterations required to achieve an acceptable result goes up. The maximum allowable t for 1-D is $\frac{1}{3}$, and for 2-D with 4-connected neighbors as in equation (5.8) the maximum t is $\frac{1}{5}$. Since the improvement with each iteration is only incremental, usually many iterations of anisotropic diffusion are desired. The best results are achieved when the algorithm is allowed to run to convergence, which may require hundreds or even

thousands of iterations. In one dimension, these requirements are not extreme, because the computations involved in each iteration are not intensive. In two dimensions, however, even 100 iterations are very time-consuming, so the algorithm is not usually run to convergence. Typical 2-D applications only run about three iterations, which is not even close to convergence [14].

Selection of the diffusion parameter is trickier than choosing the other parameters. The value of represents the magnitude of the gradient that gives the maximum change between iterations; that is, edges of magnitude are modified the most, while constant regions and edges of magnitude much greater than are virtually unchanged. An illustration of this is shown in Figure 5.20 below.

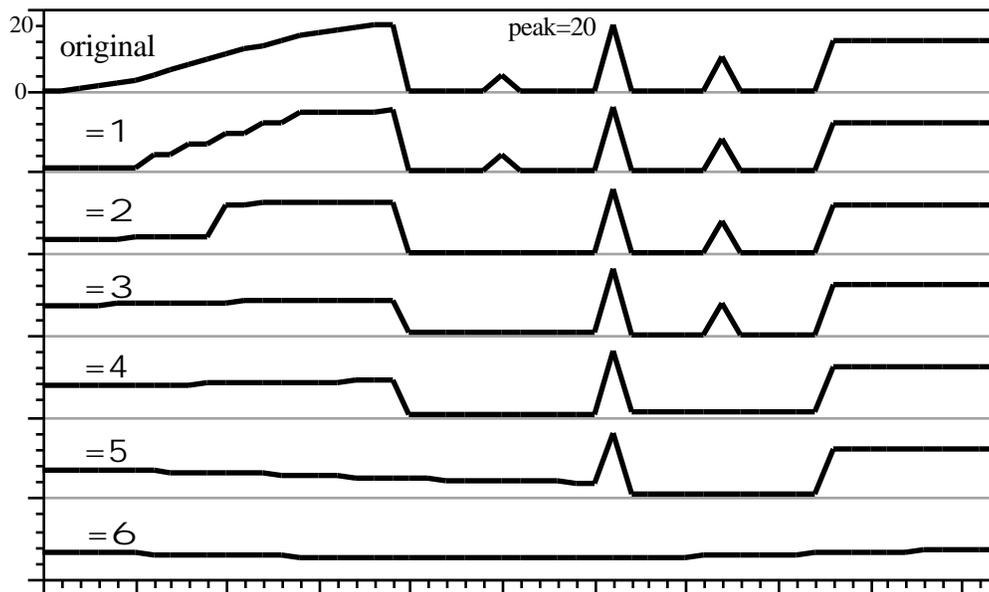


Figure 5.20. Anisotropic diffusion of a signal with ranging from 1 to 6.

The original signal at the top of Figure 5.20 has a slow ramp edge going from 0 to 20, three impulses of size 5, 20, and 10, and a step edge of height 15. The six signals below it are the results of anisotropic diffusion with diffusion parameter λ ranging from 1 to 6, iterated to convergence with integration constant $t = 0.2$. For $\lambda = 1$, the impulses and step edge are unchanged, but the ramp edge is converted to several “stair steps” of height approximately 3.2. For $\lambda = 2$, the smallest impulse (amplitude 5) is removed, and the ramp edge is converted to a step edge of height approximately 10.2. The ramp edge is smoothed out completely for $\lambda = 3$, and the impulse of amplitude 10 is removed for $\lambda = 4$. The features of amplitude 20 are not removed until $\lambda = 6$. These values all roughly agree with the heuristic that differences greater than 3λ are removed by iterated anisotropic diffusion. The number of iterations for anisotropic diffusion to converge to within 0.001 for this signal ranges from 46 iterations at $\lambda = 1$ to 550 iterations at $\lambda = 5$.

Comparison to MLV Filter

The MLV filter operates quite differently from anisotropic diffusion. Since it is based on the concepts of mathematical morphology, the size and shape of the structuring element (N) of the MLV filter determine its behavior. The MLV filter thus preserves structures of a certain size and shape, whereas anisotropic diffusion preserves structures of a particular slope or amplitude. This difference can be seen by comparing the results of anisotropic diffusion in Figure 5.20 to the results of one iteration of MLV filtering with various structuring element sizes ($|N|$) in Figure 5.21 below.

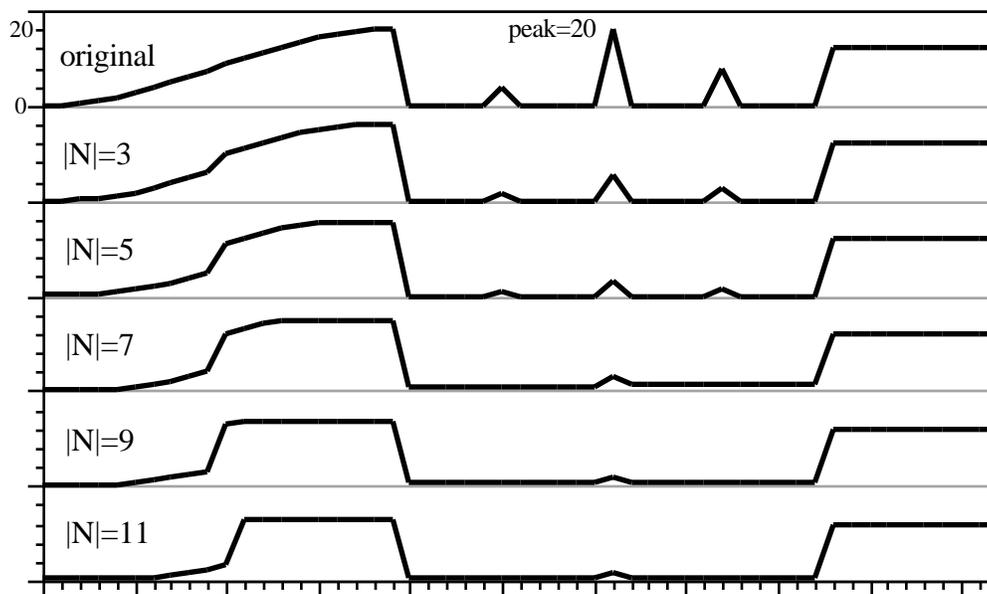


Figure 5.21. MLV filtering of a signal with $|N|$ ranging from 3 to 11.

Several differences between the MLV-filtered signals and the anisotropic diffusion results are prominent. First, all sizes of the MLV filter reduce the amplitude of the impulses, and for $|N| \geq 7$ the smaller two impulses are completely removed. Second, the ramp edge is sharpened by all sizes of the MLV filter as well, and as the size of the structuring element increases, so does the amount of sharpening. Finally, the step edge at the right end of the signal is preserved exactly by all the filter sizes shown in the figure. The amplitude and (to a certain extent) the slope of an edge are not important in MLV filtering; rather, the width of a feature or “pulse” are the significant factors. Therefore, long constant areas with small steps between them are preserved by the MLV filter. In contrast, the results of anisotropic diffusion do not depend on the size of the

feature, but rather on the amplitude or slope of the feature. This means that one iteration of the MLV filter is much more effective at reducing the effects of outlying values than even many iterations of anisotropic diffusion.

In two dimensions, similar differences are observed between anisotropic diffusion and the MLV filter. Figure 5.22 illustrates how the edge enhancing properties of anisotropic diffusion depend heavily on the value of the diffusion parameter λ , and that these properties can change drastically with only a small change in λ . A 30×30 image of a 2-D rectangular pulse of amplitude 1 and of size 14×14 is shown in Figure 5.22(a) as a 3-D mesh plot. The edges are blurred by successive passes of the simple averaging filter to give the blurred image in Figure 5.22(b). Figure 5.22(c)–(g) illustrate anisotropic diffusion iterated to convergence for various values of λ ranging from $\frac{1}{100}$ to $\frac{1}{5}$. For values of λ that are less than about $\frac{1}{10}$ the amplitude of the pulse, little sharpening of the ramp edges is observed. When λ is about $\frac{1}{6}$ the amplitude of the feature, excellent edge sharpening is seen, but the amplitude of the pulse is reduced by about 30%. However, when λ is $\frac{1}{5}$ of the amplitude or greater, the edges are smoothed across and the feature is not preserved. This rapid change from the best edge sharpening to edge removal is a troubling feature of anisotropic diffusion, because the amplitudes of features in images are not typically known in advance. Also, features of various amplitudes often need to be sharpened, and a single value of λ does not provide good contrast enhancement in this situation. Note also that the amplitude of the rectangular pulse is changed significantly for the values of λ that provide the most contrast enhancement.

The results of MLV filtering with various structuring element sizes from 3×3 to 9×9 , shown in Figure 5.22(*h*)–(*j*), show that the MLV filter enhances edges without regard to the amplitude of the feature, and does not modify the amplitude in most cases. However, the size of the feature is very important in MLV filtering; if the structuring element of the filter becomes larger than the feature, the feature is almost completely removed by the filter. This is evidenced by an MLV filter with a 17×17 structuring element, which reduces the rectangular pulse to a 6×6 area in the center of the pulse with amplitude less than 0.05; that is, virtually to zero. The more reasonably sized structuring elements shown in Figure 5.22(*h*)–(*j*) do a good job of enhance the edges of the pulse and preserving its amplitude.

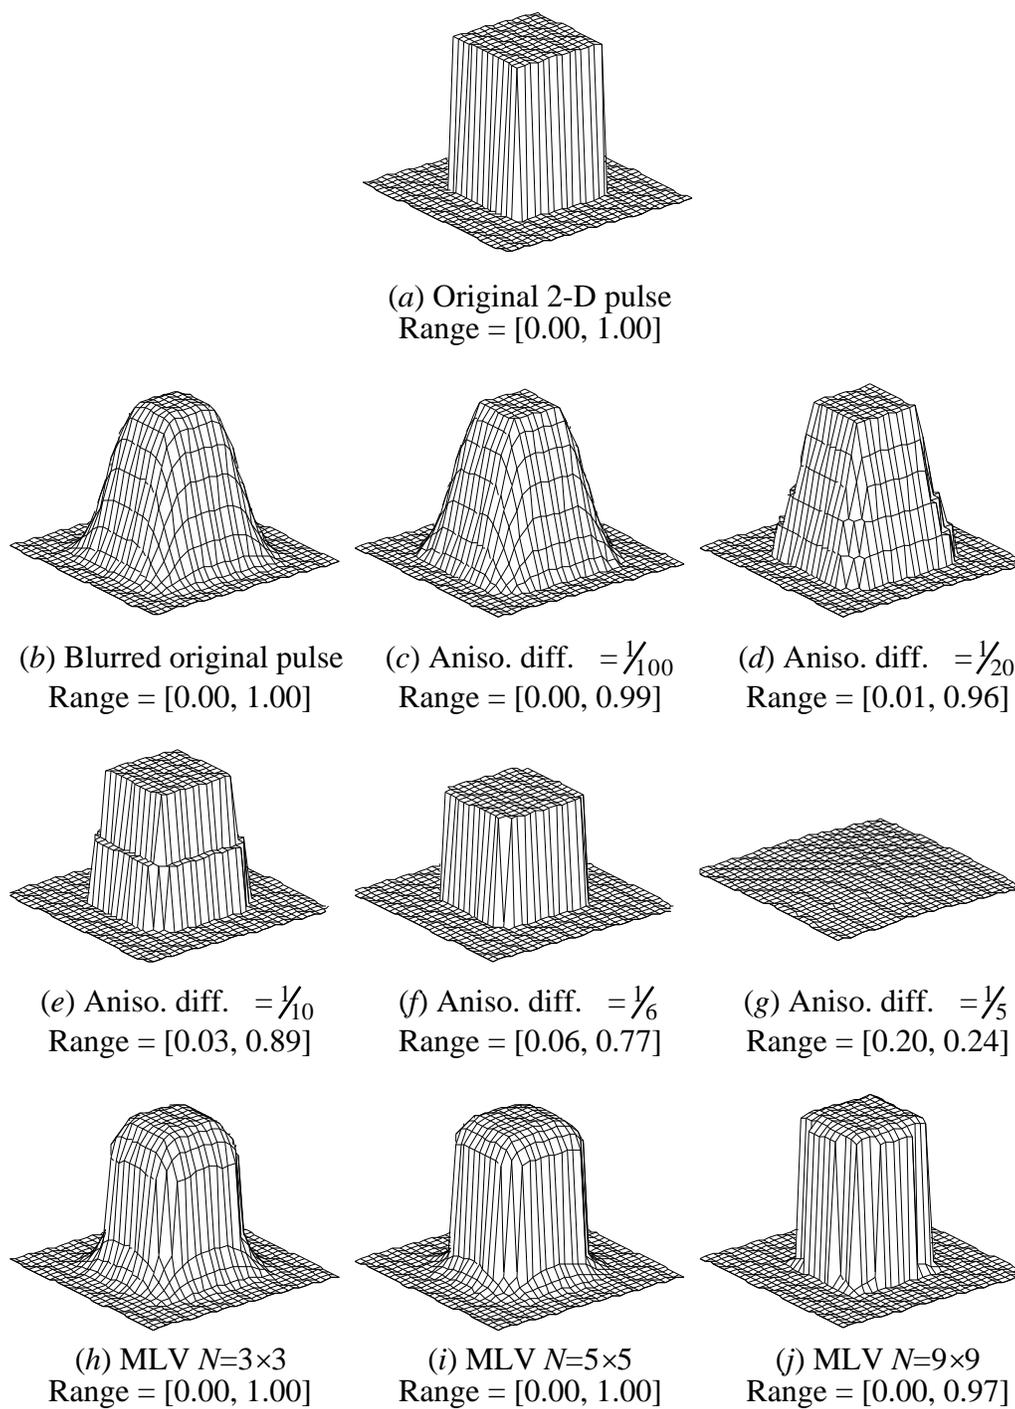


Figure 5.22. Comparison of anisotropic diffusion and MLV filtering for a two-dimensional rectangular pulse.

5.3.3. Noise Estimation

Since a “noiseless original” or “truth” image corresponding to the acquired (noisy) MRI cannot be found, standard techniques for estimating image quality and the signal-to-noise ratio cannot be used. Gerig, *et al.* [14] estimated the noise in MRIs by finding the 8×8 regions with the smallest standard deviation in the tissue and background areas of the image. The assumption behind this method is that an area with the smallest variation represents a homogeneous region of the tissue (or background) and therefore the standard deviation of this region represents the amount of noise in these regions. This method will automatically choose areas with the lowest noise as well, so the noise estimate represents a least lower bound. The same 8×8 regions are used to estimate the noise in the filtered image when comparing it to the unfiltered image.

A similar technique used to estimate noise in an image where a noiseless version is not available is to find the peak of the histogram of the local standard deviations of an image. The standard deviation of the $m \times m$ region centered at each point in the image is computed, and the histogram of these values is found. The standard deviation of the noise is estimated to be approximately equal to the histogram bin that contains the greatest number of local standard deviations. The reasoning behind this technique is that local values of the standard deviation that are due mainly to the edges in the image will be spread throughout the histogram, whereas those due mainly to the noise will be concentrated within a small range of values. One problem with this technique is that it is typically impossible to find a unique mode of a distribution of real numbers. Therefore, there must be

some criterion to establish the correct bin size () for the histogram to provide the proper tradeoff between resolution in the standard deviation and accuracy in estimating the peak of the histogram (mode). Also, the size of local regions over which the standard deviations are computed must be chosen appropriately for this technique to provide meaningful data.

One way to show the validity of this technique is to demonstrate that it can correctly estimate the standard deviation of noise with known properties that is added to an MR image, assuming that the magnitude of the added noise is much greater than that of the inherent noise in the image. Figure 5.23 below shows an original MR image of a human head in coronal section. The size of the image is 256×256 pixels, and gray values in the image range from 34 to 225. In Figure 5.24, Gaussian noise with mean of $\mu = 0$ and standard deviation of $\sigma = 10.0$ is added to the original image, resulting in gray values from 0 to 243. Computing the standard deviation of the noisy image over 9×9 local windows yields a histogram with a clear peak of 1,043 points at the bin centered at a value of $\bar{x} = 9.95$ when the bin size is $\Delta x = 0.1$. This means that the estimation technique yields a standard deviation estimate in the range $9.90 \leq \sigma \leq 10.0$, very close to the true value of $\sigma = 10.0$. Using smaller local windows yields less distinct peaks near the correct value; however, when computing the variance of just the added noise, a window size of at least 9×9 is necessary to achieve consistently accurate results. Larger window sizes also show a peak near the expected value of $\sigma = 10.0$, but larger regions also include many more signal features and are thus less likely to satisfy the general assumption that the standard deviation is due to the noise and not the signal.



Figure 5.23. Scaled original MR image of human head (coronal section).

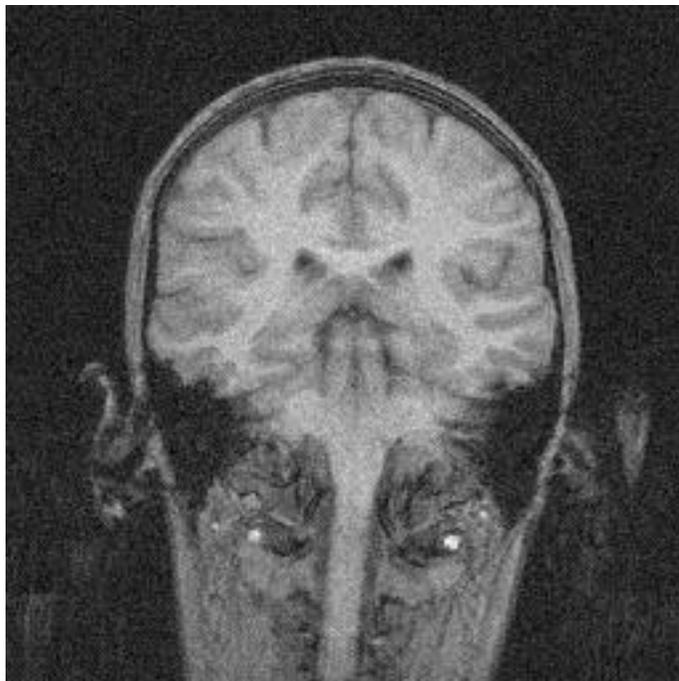


Figure 5.24. MR image with additive Gaussian noise ($\mu=0$, $\sigma=10.0$).

A criterion is needed to ensure that the histogram bins are sufficiently wide not to yield a spurious peak. If the bins are too narrow, the histogram will be noisy and false peaks will likely occur. If the bins are too wide, the estimated value of the standard deviation may include some variation due to the signal. One simple way to provide a criterion for the bin size is to require that the peak histogram bin contain at least a certain percentage of the overall number of points. Experiments on the images in this section indicate that a peak bin containing at least 1.5% of the total number of image pixels (that is, about 1,000 points for a 256×256 image) is required to prevent spurious histogram peaks from appearing. Best performance of this noise estimation method still depends somewhat on heuristics for the choice of histogram bin size to get a histogram with a clearly defined peak, but the above criterion at least gives some objective measure to use to assist in the choice.

5.3.4. Results

The noise estimation technique described above is used in this section to compare the performance of anisotropic diffusion to that of MLV filtering on MR images. The original MR image used for this section is shown in Figure 5.25 below. It is a full grayscale (0 to 255) version of the image shown in Figure 5.23. The peak of the histogram of the 9×9 local standard deviations is in the bin centered at $\bar{\sigma} = 1.88$ with a bin size $\Delta\sigma = 0.08$. There are 1,011 points in this bin. This means that the standard deviation of the noise of the original image is estimated to lie in the range from $\sigma = 1.84$ to $\sigma = 1.92$.

The MLV filter with a square structuring element reduces the noise in this MR image by an amount that agrees with the statistical properties derived in Section 3.3.5. The filtered images for structuring elements of size 3×3 , 5×5 , and 9×9 are shown in Figures 5.26–5.28 below. The results of the noise estimation process for the MLV filtered images are given in Table 5.6 below for these filter sizes and for 2×2 and 4×4 filters as well. The 3×3 filter reduces the noise by an estimated 61%, and the 5×5 filter by about 74%. The 9×9 filtered image gives a noise estimate of 0, indicating that the noise is nearly completely removed by the MLV filter with a large structuring element. However, large structuring elements also introduce significant distortions to the image, as is seen in Figure 5.28. The filtered images are much sharper than the original image, and boundaries between different tissue types are much more distinct after MLV filtering.

Gerig, *et al.* [14] used three iterations of 2-D anisotropic diffusion to filter single-slice MR images. They suggested that the value of the diffusion parameter be chosen based on the estimated standard deviation of the noise in the image, σ , with $1.5 \sigma < \epsilon < 2.0 \sigma$, and showed results for three iterations of the diffusion process. For the original MR image in Figure 5.25, the noise estimate is $\sigma = 1.88$, so the suggested range of ϵ is about 2.8 to 3.8. Using the fastest allowable rate of change ($\epsilon = 0.2$), the results of three iterations of anisotropic diffusion on the original MR image are shown for $\epsilon = 3.0$ in Figure 5.29 and for $\epsilon = 20$ in Figure 5.30. Noise estimates for the images after anisotropic diffusion are given in Table 5.6 for $\epsilon = 2.0, 3.0, 4.0$, and 20. Larger values of ϵ for the diffusion process yield lower noise estimates, and the $\epsilon = 4.0$ and $\epsilon = 20$ values

are distinct improvements over the smaller values of λ . These data contradict the suggestion of Gerig, *et al.* for selecting the value of λ for a small number of iterations.

To improve the performance of anisotropic diffusion, the algorithm may be run for more iterations. The effect of running anisotropic diffusion for 15 iterations on the original MR image is shown for $\lambda = 3.0$ in Figure 5.31 and for $\lambda = 20$ in Figure 5.32. The noise estimates for these images and for 15 iterations at $\lambda = 2.0$ and $\lambda = 3.0$ are also listed in Table 5.6. Note that the most common value of the local 9×9 variance for the images with $\lambda = 3.0$ is exactly 0. This means that these images have a very low noise level. Secondary peaks in the histograms of these images are found at $\sigma = 0.495$, which indicates that many of the local standard deviations have values near 0.5 because of image features and not from noise. Therefore, when the noise level in this image drops below 0.5, it is difficult to estimate it accurately using the method of Section 5.3.3. Note also that the filtered image with $\lambda = 20$ (Figure 5.32) does not exhibit contrast enhancement between the white and gray matter. This is because the high value of λ has smoothed across edges of intermediate size (up to approximately a difference of 60 gray levels). The image with $\lambda = 3$ (Figure 5.31) has most of the noise removed; however, contrast enhancement is quite subtle compared to the MLV filter. As a pre-filter for a segmentation application, the MLV filter does a much better job of removing spurious features in the image, whereas anisotropic diffusion removes the noise and subtly improves contrast to yield a faithful representation of the original image.



Figure 5.25. Original MR image of human head (coronal section).



Figure 5.26. MLV filtered MR image ($N = 3 \times 3$).

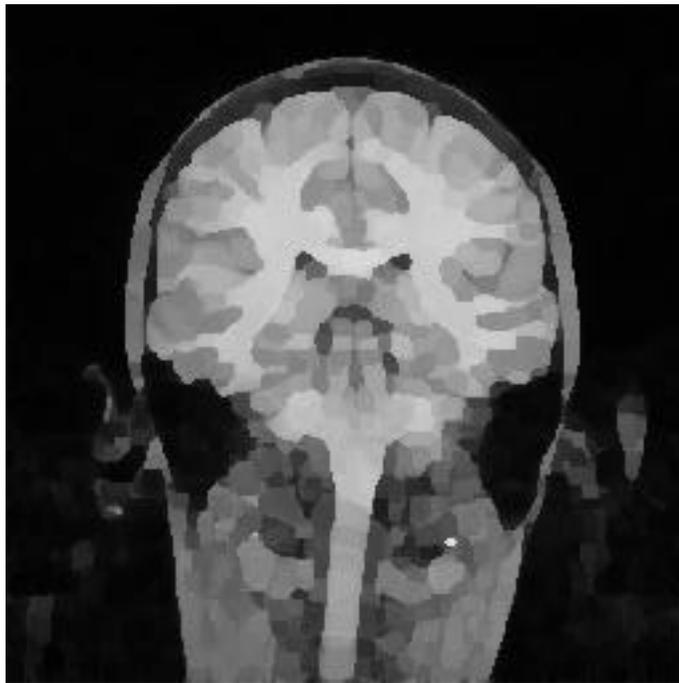


Figure 5.27. MLV filtered MR image ($N = 5 \times 5$).



Figure 5.28. MLV filtered MR image ($N = 9 \times 9$).



Figure 5.29. Anisotropic diffusion of MR image (3 iterations with $\sigma = 3.0$).

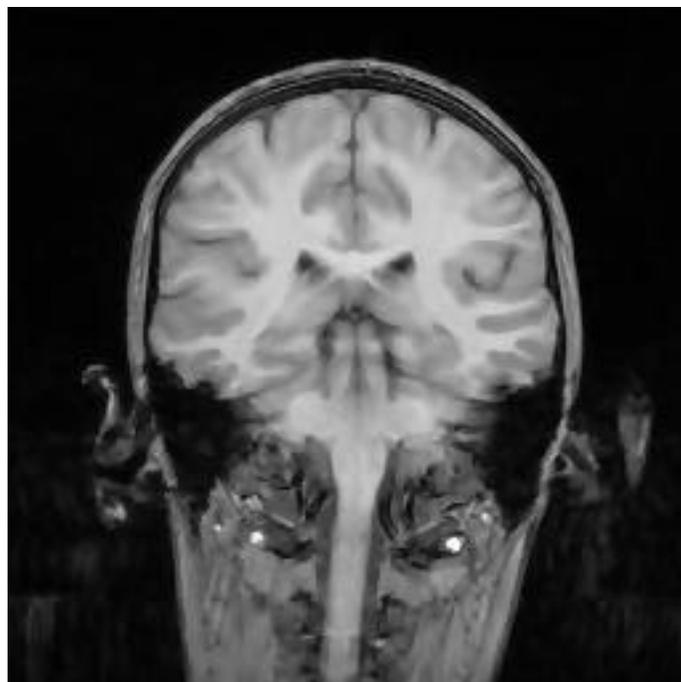


Figure 5.30. Anisotropic diffusion of MR image (3 iterations with $\sigma = 20.0$).



Figure 5.31. Anisotropic diffusion of MR image (15 iterations with $\lambda = 3.0$).

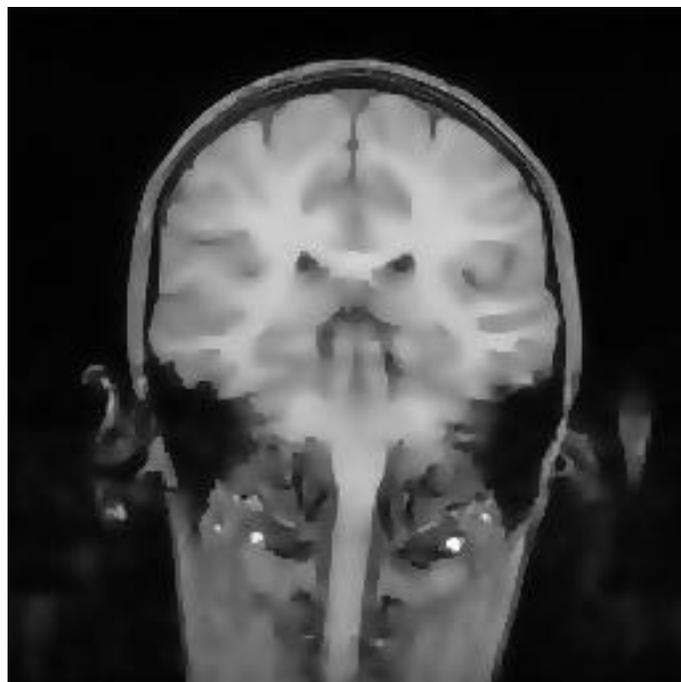


Figure 5.32. Anisotropic diffusion of MR image (15 iterations with $\lambda = 20$).

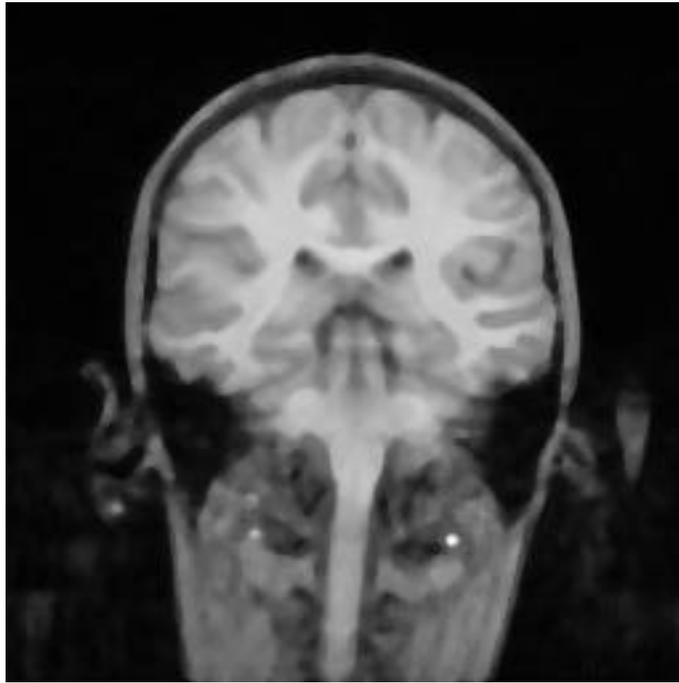


Figure 5.33. Median filtered MR image (5×5 square filter window).

Table 5.6. Noise estimation in filtered MR images.

Filter	Size or Parameter	Number of Iterations	Bin center —	Bin size	Counts in Bin
(None)	—	—	1.88	0.08	1011
MLV	$N = 2 \times 2$	1	1.015	0.07	1154
MLV	$N = 3 \times 3$	1	0.725	0.05	1130
MLV	$N = 4 \times 4$	1	0.49	0.02	1196
MLV	$N = 5 \times 5$	1	0.495	0.01	1143
MLV	$N = 9 \times 9$	1	0.00	(any)	3220
Aniso. diff.	= 2.0	3	1.45	0.10	1066
Aniso. diff.	= 3.0	3	1.15	0.10	1057
Aniso. diff.	= 4.0	3	0.57	0.06	1048
Aniso. diff.	= 20.0	3	0.495	0.03	1136
Aniso. diff.	= 2.0	15	0.65	0.10	1176
Aniso. diff.	= 3.0	15	0.00	(any)	1364
Aniso. diff.	= 4.0	15	0.00	(any)	1639
Aniso. diff.	= 20.0	15	0.00	(any)	1795
Median	$W = 5 \times 5$	1	0.04875	0.025	1120

Another way to quantify the enhancement of MR images by the MLV filter and anisotropic diffusion is to use an image quality measure. Unfortunately, the most widely used image quality measures require a “noiseless” original image. This is true for the mean squared error (MSE), mean absolute error (MAE), and related measures, which compute a distance from the “noiseless” image to the

filtered image to estimate the fidelity of the filtering process. In medical imaging, a “noiseless” original image is almost never available, so other techniques must be used to estimate image quality.

One such image quality measure is the Lorenz Information Measure, or LIM [63, 69, 70]. Briefly, the LIM measures the degree of nonuniformity of the histogram of an image by ranking the probabilities of each gray level, integrating over all gray levels, and normalizing. Images with nearly uniform probabilities for all gray levels have a high LIM, whereas images that mostly consist of only a few different gray levels have a low LIM. The extreme values are $LIM = \frac{1}{2}$ for uniformly distributed noise (or, *e.g.*, a uniform ramp) and $LIM = 0$ for a constant image. Generally, images with lower values for the LIM are considered to be of better quality; however, as with any image quality measure, the LIM must be used with care. Very low values for the LIM may be achieved by destroying the signal content along with the noise. This weakness is a common problem when estimating quality or noise levels without having a noiseless signal or image with which to compare filtered results; the noise estimation technique of Section 5.3.3 is also likely to yield “good” results when the signal is removed along with the noise.

The LIM for the original MR image (Figure 5.25) is 0.2263. For the filtered images, the LIM values are listed in Table 5.7. For all MLV filter sizes shown and for all anisotropic diffusion parameters shown, the LIM is lower for the filtered image, indicating that the noise is reduced by filtering. However, the LIM for all MLV filter sizes bigger than 2×2 is smaller (better) than the LIM for any of the anisotropically diffused images, including those with 15 iterations of

the diffusion. The higher LIMs for anisotropic diffusion may be partially due to its property of preserving even very small features of large amplitude, whereas the MLV filter smoothes over features smaller than its structuring element, regardless of amplitude. The noise reduction estimates in Table 5.7 are computed directly from the standard deviation estimates in Table 5.6 and are shown for comparison.

Table 5.7. Image quality estimation in filtered MR images.

Filter	Size or Parameter	Number of Iterations	Noise Reduction Estimate	LIM
(None)	—	—	0	0.2263
MLV	$N = 2 \times 2$	1	46.0%	0.2189
MLV	$N = 3 \times 3$	1	61.4%	0.2093
MLV	$N = 4 \times 4$	1	73.9%	0.2022
MLV	$N = 5 \times 5$	1	73.7%	0.1956
MLV	$N = 9 \times 9$	1	100%	0.1710
Aniso. diff.	= 2.0	3	22.9%	0.2233
Aniso. diff.	= 3.0	3	38.8%	0.2209
Aniso. diff.	= 4.0	3	69.7%	0.2195
Aniso. diff.	= 20.0	3	73.7%	0.2151
Aniso. diff.	= 2.0	15	65.4%	0.2186
Aniso. diff.	= 3.0	15	100%	0.2158
Aniso. diff.	= 4.0	15	100%	0.2143
Aniso. diff.	= 20.0	15	100%	0.2102
Median	$W = 5 \times 5$	1	74.1%	0.2190

Another nonlinear filter that is often used to suppress noise while preserving edges is the median filter. However, the median filter does not improve the contrast of images because it does not sharpen edges or preserve corners, unlike the MLV filter. The results of 5×5 square median filtering of the original image are shown in Figure 5.33. The edges in the median-filtered image are not nearly as distinct as in the MLV-filtered images, although both filters remove noise quite well. The 5×5 median filter reduces the estimate of the noise standard deviation to about 0.49, roughly the same as for the 4×4 and 5×5 MLV filters, but does not provide any contrast enhancement.

5.3.5. Evaluating Filter Performance Using a Phantom Image

Since a noiseless original MR image is not available for the examples given above, there is no way to measure the accuracy of the tissue boundaries after filtering. One way to determine if the filtering algorithms introduce significant spatial distortions is to use a phantom image that simulates the structure of the real images. Since a “truth” image is known under these circumstances, the amount of distortion introduced by filtering can be measured. It is significantly more difficult to form a meaningful model of a magnetic resonance image of a human head than to form a model for the thermographic problem in the previous section. However, a simple phantom image can be designed to indicate whether the filters introduce significant distortion to tissue boundaries.

A well-known phantom image for structures of the human head is the Shepp and Logan phantom [50, 71], which is intended to simulate the head as

visualized by X-ray computed tomography (CT). The Shepp and Logan head phantom is specified by 10 ellipses of given sizes and relative amplitudes on a coordinate lattice ranging from -1 to 1 in both the x and y directions. This phantom is not appropriate for MRI, however, because the relative amplitudes and sizes of the features are not typical for MR images.

Using the real MR image of Figure 5.25 as a guide, I have developed a head phantom for MRI that models the major features of MR images with ellipses. This phantom image is diagrammed in Figure 5.34. There are five ellipses in this phantom, which correspond to the skull and various brain structures (gray matter, white matter, and ventricles). The approximate shapes and relative amplitudes of the ellipses are also modeled from the real image of Figure 5.25. Table 5.8 below lists the coordinates, major and minor axis lengths, rotation angles, and relative gray levels of the five ellipses in the phantom. Note that the coordinate lattice is assumed to range from -1 to 1 in the x and y directions and that a rotation angle of 0 for an ellipse means that its major axis is vertical. Absolute gray levels in the phantom range from 0 to 2 . In Figure 5.35, the phantom is shown on a 256×256 pixel grid with the gray levels scaled by a factor of 100 ; that is, the skull has a gray level of 200 in the figure.

Table 5.8. Component ellipses of MRI head phantom.

Name	Center	Major axis	Minor axis	Rotation	Gray level
Skull	(0, 0)	0.667	0.667	0	+2
Gray Matter	(0, 0)	0.637	0.605	90°	-1
White Matter	(0, 0)	0.488	0.450	0	+0.5
Ventricle 1	(-0.14, 0.08)	0.066	0.059	90°	-1
Ventricle 2	(0.14, 0.08)	0.066	0.059	0	-1

To simulate the effects of noise on the image and on the filtering algorithms, Gaussian noise with zero mean and standard deviation of 10 gray levels was added to the original phantom. The resulting noisy image is shown in Figure 5.36. This noise level is more than five times as great as the noise estimate for the real image in Figure 5.25. This higher noise level is needed to demonstrate clear differences among the filtering schemes. The noisy phantom image after filtering is shown for several selected filters in Figure 5.37–5.41. The same effects that were seen for these filters in the real images above are again seen in the phantom images below.

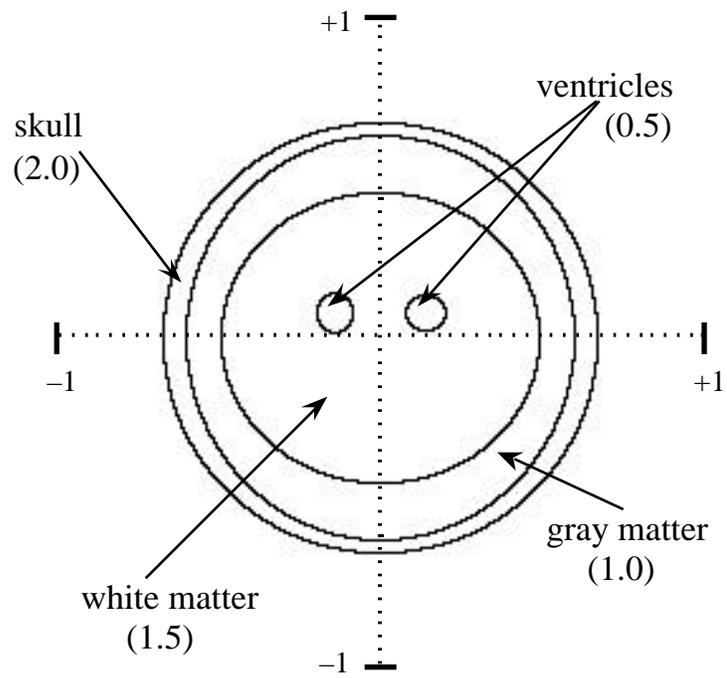


Figure 5.34. Diagram of MRI head phantom.

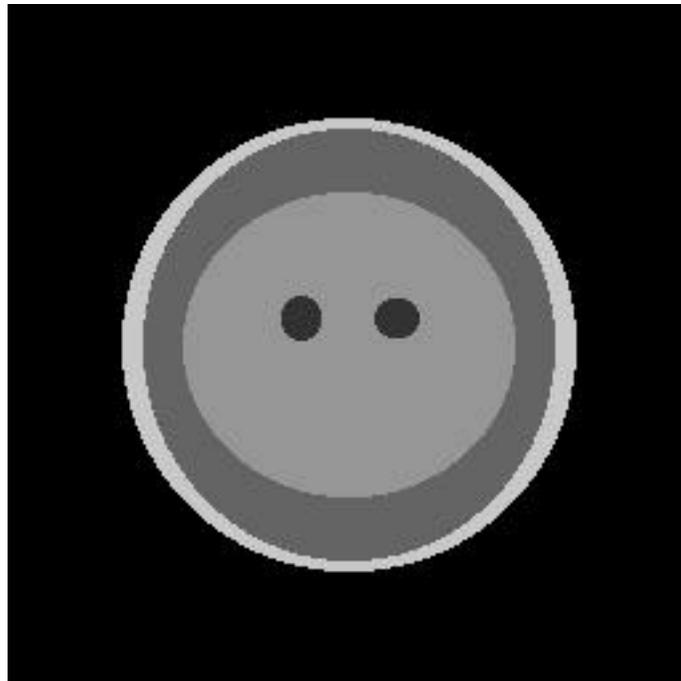


Figure 5.35. Image (256×256) of MRI head phantom.

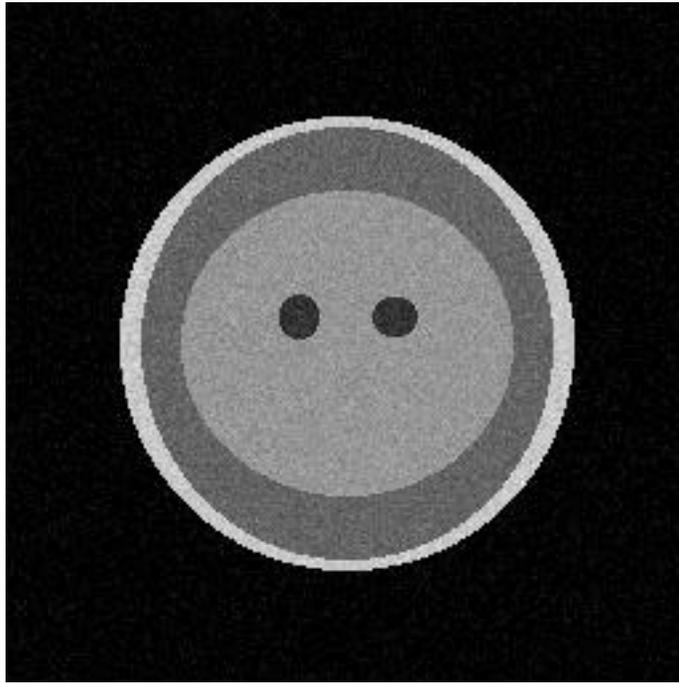


Figure 5.36. Head phantom with additive Gaussian noise ($\sigma = 10$).

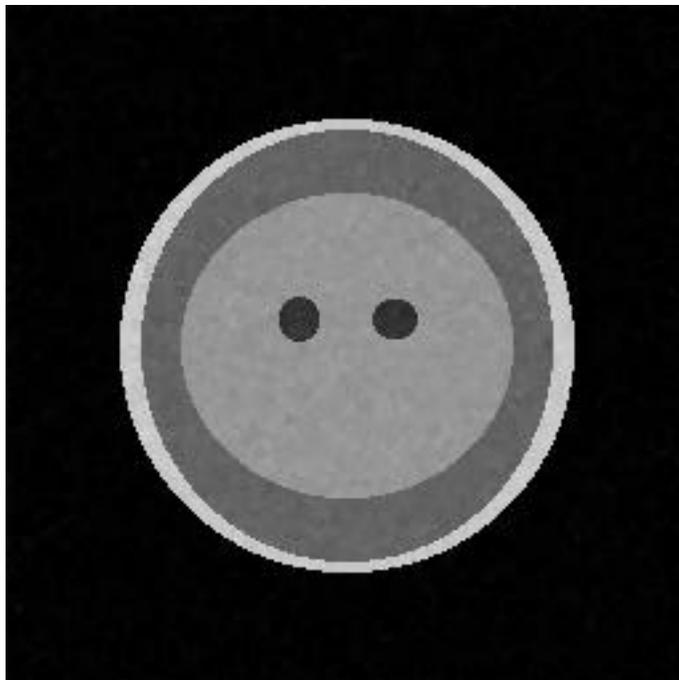


Figure 5.37. MLV filtered head phantom ($N = 3 \times 3$).

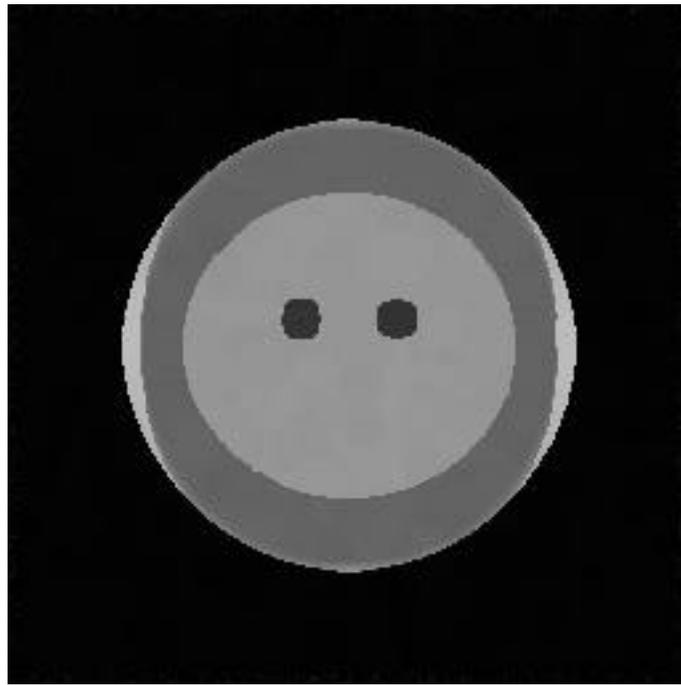


Figure 5.38. MLV filtered head phantom ($N = 9 \times 9$).

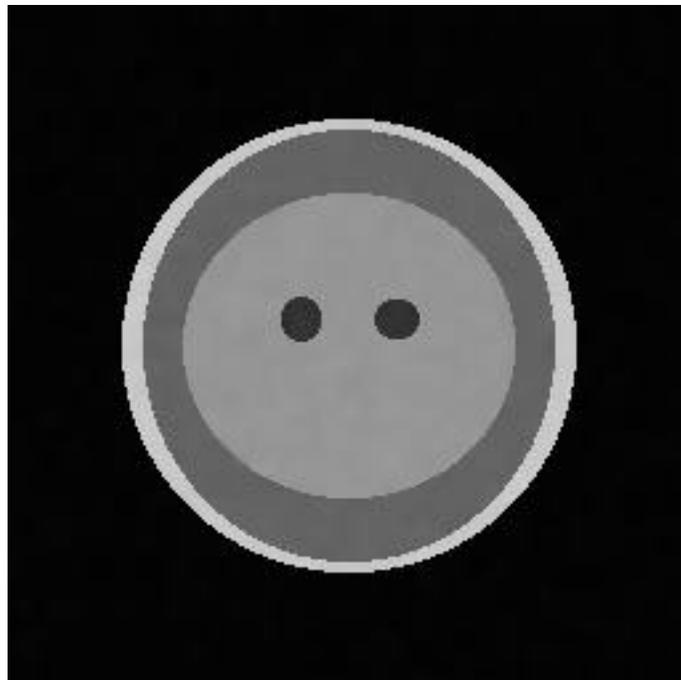


Figure 5.39. Anisotropic diffusion of head phantom (15 iterations with $\sigma = 20$).

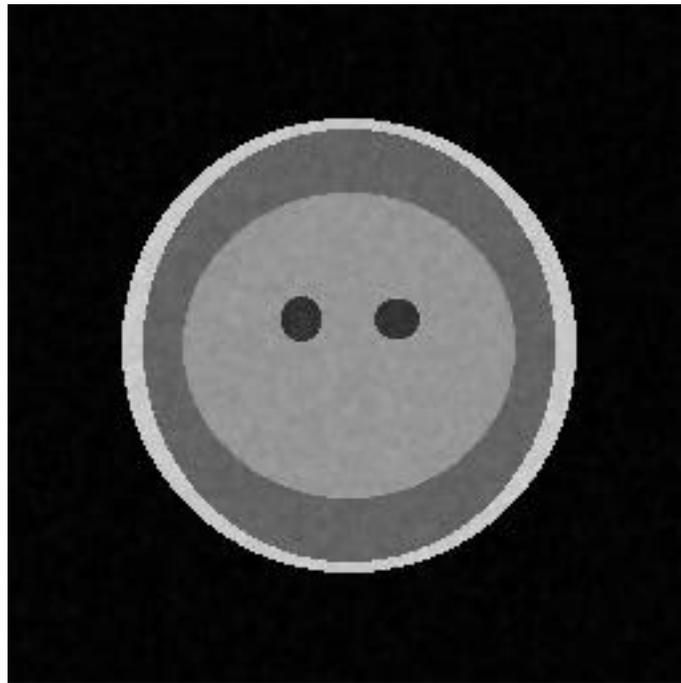


Figure 5.40. Anisotropic diffusion of head phantom (3 iterations with $\sigma = 30$).

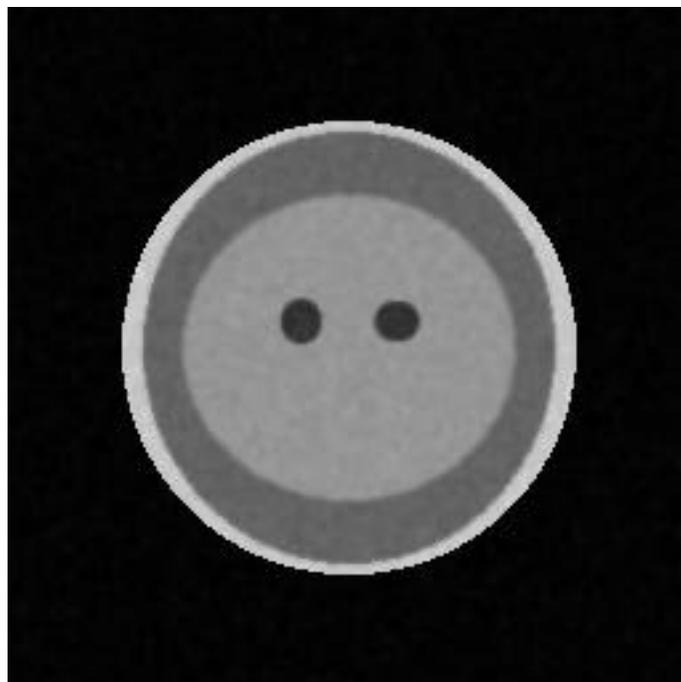


Figure 5.41. Anisotropic diffusion of head phantom (3 iterations with $\sigma = 100$).

Since the features in the head phantom are about the same size as the features in the real images, the MLV filter should have roughly the same effect on the phantom as on the real images for the same structuring element sizes. However, selection of the anisotropic diffusion parameter (λ) is related to the noise level in the images. Since the noisy phantom image has a higher noise level than the real MR images, the diffusion parameter must be increased in the phantom examples to achieve noise smoothing results similar to those for the real images. Values of λ ranging from about 1.5 to about 10 were used for the real images, which correspond to values of λ for the noisy phantom image of 15 to 100.

To measure the performance of the filters for removing noise without distorting tissue boundaries, the number of false positives and false negatives will be found for a simple segmentation algorithm acting on the filtered results. The gray levels in the original phantom are 0 for the background, 200 for the skull, 100 for the gray matter, 150 for the white matter, and 50 for the ventricles. The simplest segmentation algorithm for this image is simple thresholding, where the values 0 to 25 are assumed to be background; 26 to 75, ventricles; 76 to 125, gray matter; 126 to 175, white matter; and 176 to 255, skull. “False positives” are points in the filtered segmentation that are not in the corresponding segmentation of the original phantom. “False negatives” are points that are not in the filtered segmentation that are in the corresponding original segmentation. The number of false positives and negatives for each classification is expressed as a percentage of the true area of that class. For the head phantom, the tissue classes have the following areas: background (B), 42555 pixels; skull (S), 3128 pixels; white

matter (W), 10873 pixels; gray matter (G), 8566 pixels; and ventricles (V), 414 pixels. The results for a variety of filtering techniques and parameters are summarized in Table 5.9 below. (“A.D.” in the table stands for anisotropic diffusion.)

Table 5.9. Percentage of points misclassified by various filters acting on phantom image.

iterations \times filter (size/parameter)	False Negatives (%)					False Positives (%)				
	B	S	G	W	V	B	S	G	W	V
(noisy phantom)	0.49	0.74	1.30	1.13	1.45	0.01	1.92	0.77	0.65	65.9
1 \times MLV (3 \times 3)	0.02	0	0	0.04	0	0	0	0.05	0	1.69
1 \times MLV (5 \times 5)	0.03	4.25	0.05	0.07	0	0.01	0	1.52	0.09	2.66
1 \times MLV (9 \times 9)	0.05	74.0	0.08	0.29	7.49	1.31	0	17.0	3.45	5.31
3 \times A.D. (λ =15)	0.12	0.42	0.44	0.45	0.97	0.01	0.67	0.35	0.30	16.9
3 \times A.D. (λ =20)	0.02	0.13	0.11	0.07	0	0	0.06	0.07	0.08	3.38
3 \times A.D. (λ =30)	0	0	0	0.01	0	0	0	0.01	0	0
3 \times A.D. (λ =100)	0	7.70	2.25	0.47	2.25	0	0	1.10	3.99	0
15 \times A.D. (λ =15)	0.02	0.13	0.11	0.06	0	0	0.10	0.04	0.09	2.42
15 \times A.D. (λ =20)	0	0	0.01	0	0	0	0	0	0.01	0
1 \times Median (5 \times 5)	0	0	0.18	0.01	9.42	0	0	0.47	0.14	0

The above results indicate that the MLV filter and anisotropic diffusion do not introduce significant distortions of the tissue boundaries in MR images when their parameters are chosen reasonably. The most significant distortions noted in Table 5.9 are for: 1) the 9 \times 9 MLV filter, which has an inappropriately large structuring element that smoothes over most of the “skull” pixels and distorts the ventricles; 2) anisotropic diffusion with λ = 100, which is too large for this

application and has problems segmenting the skull and ventricles; and 3) the median filter, which also distorts the ventricles. All the other filtering techniques listed in the table show improvement in most of the tissue classes. Anisotropic diffusion with $\lambda = 30$ run for 3 iterations and with $\lambda = 20$ for 15 iterations yield the best overall results; in each case, only 1 pixel is misclassified out of 65 536 in the entire image. However, the MLV filter with a 3×3 structuring element performs nearly as well in only one iteration, yielding a total of 11 misclassified points. The 2×2 MLV filter actually performs slightly better on the phantom, leading to only 7 misclassified points, but its result is not as smooth as the result of 3×3 MLV filtering. Clearly, the MLV filter and anisotropic diffusion (with appropriate parameter choices) do not cause significant distortions in the MR phantom image, which indicates that their use on real MR images will not cause significant changes in tissue volume estimation.

5.3.6. Conclusions

This section has illustrated the noise smoothing and contrast enhancing properties of the MLV filter and demonstrated how these properties are based on the size and shape of image features. This behavior is expected because the MLV filter structure is closely related to mathematical morphology. The behavior of the MLV filter contrasts with the edge enhancing properties of anisotropic diffusion, which sharpens edges based on their amplitude and slope. The MLV filter generally provides more smoothing and contrast enhancement in one iteration, whereas anisotropic diffusion offers more subtle improvements and requires many iterations to yield a significantly enhanced result. Since the MLV

filter greatly accentuates the boundaries between different tissue types in MR images, it is useful as a pre-filter for automatic segmentation by computer. Anisotropic diffusion is more useful for improving the appearance of MR images than as a part of a segmentation algorithm, since it does not tend to cluster pixels together based on location as much as the MLV filter does.

5.4. ULTRASOUND

Another biomedical imaging modality that often requires image enhancement is ultrasound imaging. Ultrasound images are very noisy, due to echo dropout and speckle [72], and objects within an image can be difficult to distinguish because signal is only present at boundaries where acoustic properties change. Therefore, an object is theoretically only visible in an ultrasound image as an outline. Taking measurements of these objects (for example, measuring bone length or ventricular volume) with a computer requires extensive pre-processing.

Most of the measurements taken on ultrasound images do not require any knowledge of the echo amplitude (gray level) at a point; ultrasound image processing usually involves segmentation of a region of interest or detection of a boundary. For this reason, the bias induced by morphological filtering is not a problem; on the contrary, it is actually helpful for enhancing ultrasound images. Not surprisingly, then, several studies of the use of morphological filters in ultrasound images have already been published. Klingler, *et al.* [56] used morphological operators to detect the left ventricular endocardial border in short-

axis echocardiograms. They used opening over a large hexagonal structuring element to find a background to subtract from the original image, then applied closing with a smaller element to the background-subtracted result to improve connectivity in the remaining features. After thresholding, they discovered that the endocardial contour was still incomplete, and so they applied dilation and a thinning algorithm to achieve the final result. Their contours were in good agreement with those drawn by human observers. This application shows the potential of morphological operators for feature detection in ultrasound images.

Another application of morphological filters in ultrasound imaging was described by Thomas, Peters, and Jeanty [57]. They used a procedure similar to [56] to segment the fetal femur and measure its length to estimate gestational age. Opening was again used to provide a background level, which was subtracted off. Then they performed contrast enhancement, followed by thresholding, closing, and opening (with a different structuring element, so as not to form a CO). The basic scheme in both of these applications is to remove the background using a large opening, then using smaller closings and dilations in combination with thresholding. The opening is biased low, which ensures that the background that is subtracted off is not over-estimated. Closing and dilation, which are biased high, are then used to improve the connectivity of the contours that result from thresholding. Since the morphological operators are stack filters (see Chapter 2), they commute with thresholding; for example, closing followed by thresholding gives the same result as thresholding at the same level followed by closing. The morphological operators are a natural choice for improving the connectivity of thresholded ultrasound images, since dilation and closing reconnect

discontinuities caused by dropout, and the opening removes speckle and other impulse-type defects.

The linear combinations of morphological operators (*e.g.*, the pseudomedian and LOCO filters) are not useful in ultrasound imaging, however, since they only partially correct for dropout and speckle-type defects in images. Value-and-criterion filters (such as the MLV filter) do not possess these properties either, and are not particularly good for processing ultrasound images. In this case, the standard morphological operators are clearly the best choice of the nonlinear filtering schemes presented here. The median filter is of some utility in ultrasound imaging for removing noise (especially speckle) and smoothing out image contours, but the LOCO, pseudomedian, and MLV filters are not generally useful for feature extraction and segmentation problems in ultrasound imaging.

The effect of various filters is illustrated by using them as pre-filters for the segmentation of a 2-D echocardiogram to determine the region of interest for a 3-D reconstruction of the left ventricle of the heart. Sweeping an ultrasound transducer through the angle subtended by the heart produces a sequence of 2D sections that may be used for a 3D reconstruction of the heart. Since reconstruction is a computationally complex problem, a fast method of determining the region to be reconstructed would greatly speed the process by eliminating points from consideration that do not need to be used in the reconstruction. Since there are small details in many ultrasound images that must be preserved, closing is a logical filtering choice because it helps remove echo dropout by connecting bright regions. Completely isolated bright points are also

passed by closing, however, and it helps to use some method of outlier removal, such as median filtering, to reduce speckle in the "closed" image.

Filtering results for this type of ultrasound image are shown in Figures 5.42–5.53. The original image, a short-axis parasternal section echocardiogram, is Figure 5.42. The mitral valve is clearly visible in this image as a nearly horizontal thin feature in the center of the image. Background subtraction is not performed for this image because DC baseline shift and sidelobe energy (common ultrasound image defects) are not prominent in the image. The result of closing with a 3×3 square structuring element is shown in Figure 5.43. Note that this operation preserves the valve while reducing echo dropout, making the images features more contiguous. Figure 5.44 shows the closed result after median filtering with a 3×3 square window. The median filter removes the spurious bright points scattered about the image and smoothes the blocky contours caused by closing the image with a square structuring element. Figures 5.45–5.47 correspond to Figures 5.42–5.44 after thresholding at the 20% level (that is, the top 80% of the grayscale is shown). Figure 5.47 is a smooth, connected region that would serve well as a region of interest for a 3-D reconstruction. Since morphological and ranked-order filters commute with thresholding, Figures 5.46 and 5.47 can also be found by thresholding the original image and then applying the closing and median filters.

For comparison, results of LOCO and MLV filtering are shown in Figures 5.48 and 5.49. The 2-D structuring element shape for the LOCO filter that has the best chance of preserving the mitral valve in the image is the 3×3 plus-shaped structuring element. However, even this small, line-based shape does not

preserve the valve well enough to ensure that it is part of the thresholded image, Figure 5.50. Note that LOCO filtering does not commute with thresholding, since it is the average of two operations that yield different binary contours. The OC operation removes the valve almost completely, leaving the LOCO result with values too low to be part of the thresholded image. Also, since the closing operation introduces a positive bias, the threshold must be adjusted downward for an unbiased filter such as the LOCO, so the threshold for Figure 5.50 is at the 11% level, 9% lower than for the closing-based algorithms given in the earlier examples. This level gives approximately the same area in the segmented region for the LOCO filter as for the closing-based algorithms. The contours of the LOCO filtered image are not very smooth, and the shape of the ventricle seems distorted.

The MLV filter with a 3×3 square structuring element does a better job than the LOCO filter of preserving the valve in the image, but still does not yield very pleasing contours. The threshold for the MLV filtered image is at the 11% level also (as for the LOCO filter) and the resulting region of interest is shown in Figure 5.51. The mitral valve is nearly complete in this image, but not nearly as well defined as for the closing algorithms. The shape of the ventricle is also somewhat distorted by the MLV filter. Admittedly, the MLV filter is not designed for ultrasound image processing, so the mediocre results are not unexpected. The value-and-criterion filter structure could probably be used to design a filter specifically for ultrasound image processing; however, the resulting filter is likely to bear a close resemblance to the morphological closing. The

properties of the closing are nearly ideal for feature extraction and segmentation in ultrasound applications.

Another technique that has been investigated for application in ultrasound imaging is anisotropic diffusion [58]. However, anisotropic diffusion does not smooth the contours in the image and does not reconnect broken features. It does remove low-level speckle quite well, though, and accentuates differences between regions that are approximately homogeneous. Anisotropic diffusion is mostly suitable as a pre-filter that removes some noise and speckle in ultrasound applications, since it does not improve connectivity. Lamberti and Sgallari [58] use anisotropic diffusion this way, as a pre-filter for the morphological operators. An illustration of the effects of anisotropic diffusion on an ultrasound image is shown in Figure 5.52–5.53. Figure 5.52 is the original image filtered by 15 iterations of anisotropic diffusion with a diffusion parameter of $\epsilon = 15$ and $t = 0.2$. The filtered image thresholded at 11% (the same as for the LOCO and MLV examples previously) is shown in Figure 5.53. The resulting contour is fairly smooth, but not as contiguous as when morphological closing is used.



Figure 5.42. Original echocardiogram.

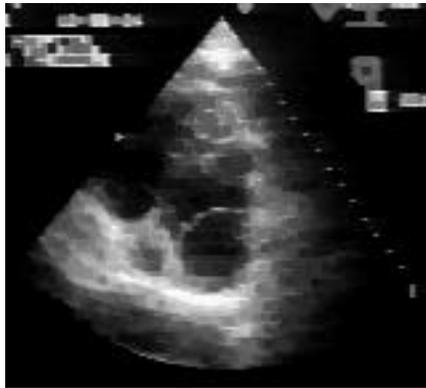


Figure 5.43. Closing (3×3) of original image.

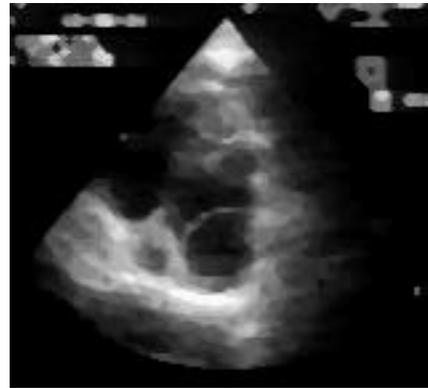


Figure 5.44. Median filtering (3×3) of closed image.



Figure 5.45. Threshold of original echocardiogram.



Figure 5.46. Threshold of Fig. 5.43.



Figure 5.47. Threshold of Fig. 5.44

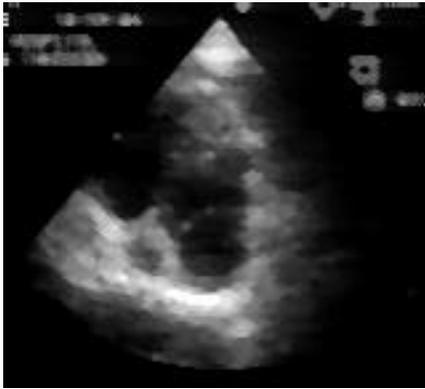


Figure 5.48. LOCO filtering (plus-shaped 3×3) of original image.

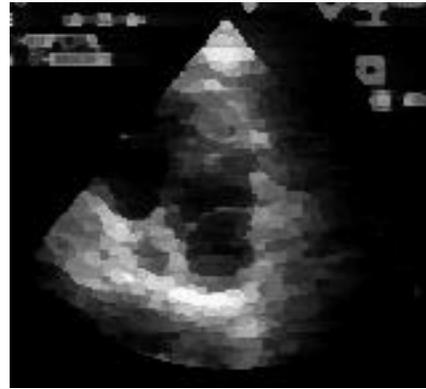


Figure 5.49. MLV filtering (3×3) of original image.



Figure 5.50. Threshold of Fig. 5.48.



Figure 5.51. Threshold of Fig. 5.49



Figure 5.52. Anisotropic diffusion (15 iterations, $\sigma=15$) of original image.



Figure 5.53. Threshold of Fig. 5.52.

5.5. CONCLUSIONS

This chapter has demonstrated the utility of nonlinear filters for the enhancement of medical images, including thermography, MRI, and ultrasound. Morphology-based filters allow processing by shape and size, which can be particularly useful for feature extraction and segmentation problems in medical image analysis. For MRI, the desired processing usually includes contrast enhancement and noise reduction. These enhancements simplify the segmentation process on MR images. The morphology-based MLV filter, which is one of the new filters in the value-and-criterion structure introduced in Chapter 3, does an excellent job of sharpening edges and reducing noise in MR images. The MLV filter yields excellent results in only one pass, as opposed to anisotropic diffusion, which requires many iterations. The MLV filter is also computationally more efficient than similar edge-preserving smoothing filters.

In thermographic imaging, it is important that the gray levels in the image be estimated very accurately. Filters that bias the image, such as the standard morphological operators, cannot be used in this application. To have the shape control of morphology when processing thermal images, one must use the linear combinations of morphological operators (such as the LOCO filter) described in Chapter 3. These filters reduce noise and yield accurate gray level estimates, but the structuring element shape and size affect the results. Although the LOCO filter does not yield a quantitative improvement over some of the previously used techniques for thermography, the shape control of the filter makes a qualitative difference that is useful in some filtering situations.

In ultrasound imaging, the standard morphological operations work quite well for removing speckle and reconnecting broken features. Since most processing tasks in ultrasound are related to segmentation or feature extraction, filter bias is not a concern in this application. The new filters defined in Chapter 3 are of little use in ultrasound, then, because the standard morphological operations fit the application so well. Therefore, this chapter illustrates applications for which the new morphology-based nonlinear filters are very helpful, and points out applications where the new filters are not an improvement over existing techniques. These examples, together with the filter analyses in Chapter 4, clarify the particular situations for which the new filters are and are not useful.

6. Conclusions

This work extends the theory of nonlinear image processing by introducing new filter structures and analysis methods, and demonstrates the utility of these new techniques on a variety of medical images. The filters are based on generalizations of mathematical morphology, which is itself a relatively recent development in image processing. One of these generalizations is the class of linear combinations of morphological operators. This filter class includes the previously defined midrange and pseudomedian filters and leads to the definition of the LOCO filter. The deterministic properties of the linear combinations are similar to those of the constituent morphological filters; however, the linear combinations are statistically unbiased, unlike the conventional morphological operators. This is important in applications like thermography, where the LOCO filter allows shape-based filtering as in standard morphology without introducing a statistical bias.

The other new class of filters introduced here is the value-and-criterion filter structure. This structure is based on morphological opening and closing, but allows the use of both linear and nonlinear operators within the window structure. One useful filter designed with this structure is the Mean of Least Variance (MLV) filter, which takes the mean of the window with the smallest variance within a set of windows defined by the morphological structure. This filter reduces noise and enhances edges in images. Its statistical and deterministic properties resemble those of the averaging filter more than the morphological

filters, with several important differences. The MLV filter is valuable in applications like MRI, where noise smoothing and contrast enhancement are the primary goals.

This dissertation also advances the analysis of nonlinear filters by introducing new methods for understanding their behavior. A method of finding the response of nonlinear filters to continuous time periodic signals of various frequencies was described. This technique bears some resemblance to Fourier analysis, but its results are much more limited because of the nonlinear nature of the filters. The breakdown point is another new technique that measures the robustness of filters to outlying signal values. The breakdown point can be used to help design filters with almost any desired level of outlier rejection ability. Corner response analysis is introduced to help quantify the ability of filters to preserve or remove two-dimensional features. Filters that have similar edge responses (such as the median and morphological filters) may have drastically different responses to corners of various angles, and this is illustrated intuitively and quantitatively by polar plots of the corner response of the filters.

To show how the new filter design and analysis techniques introduced here are used in real-world imaging applications, examples from various biomedical imaging fields (thermography, magnetic resonance, and ultrasound) are given. The unbiased linear combinations of morphological operators are good candidates for shape-based filtering of thermograms, since temperature information often must be extracted from the gray levels of these images. By choosing a structuring element shape that matches the shape of the thermal

features in the image, the LOCO filter performs well for removing noise and reconstructing smooth thermal contours.

In magnetic resonance and ultrasound imaging, preserving accurate gray levels is not as important as in thermal imaging. Processing goals for these images are more likely to include reducing noise while preserving or enhancing boundaries between tissue regions. In magnetic resonance imaging, contrast enhancement may improve the visual quality of images acquired very quickly, or may be used as a part of algorithms to segment the image into various tissue types automatically. The MLV filter is ideally suited to this task, since it provides noise smoothing and edge enhancement by averaging within homogeneous regions and away from edges. Furthermore, it yields excellent results in a single iteration, unlike other proposed methods such as anisotropic diffusion. In ultrasound imaging, however, the speckle and dropout defects are not amenable to filtering by the LOCO or MLV filters. The standard morphological operators (such as opening and closing) provide better enhancement in ultrasound imaging, thus illustrating a case where the behavior of the conventional filters is preferred.

This dissertation details important new developments in the theory, analysis, and application of nonlinear filtering. Filtering structures are devised to overcome shortcomings in the responses of morphological filters, and new analysis techniques assist in understanding the behavior of both the standard and new nonlinear filters. Appropriate filters for a given application are chosen by using information about the imaging modality and the results of the various analysis methods. Improved options now exist for nonlinear image processing because of the introduction of these new filter design and analysis tools.

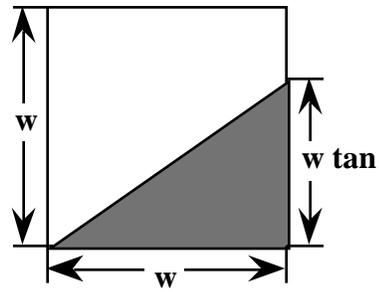
Appendix: Geometric Derivations of Corner Responses

This appendix outlines the derivation of the fractional corner preservation in continuous space for various 2-D filters. The results of these computations are summarized in Section 4.4.2 (Chapter 4). The geometric analysis is shown here for the derivations of the response of morphological opening with square and circular structuring elements, the square midrange filter, and the plus-shaped median filter. The response values at specific angles for the median and averaging filter are also derived in this appendix.

A.1. TOTAL AREA OF CORNER

Recall that the fractional corner preservation $r(\theta, \phi)$ is defined as the corner area preserved by the filter divided by the total area of the corner. The first step in finding the corner response of a filter is then to determine the total area of the region of interest for corners of all angles θ . For corners aligned with the x -axis ($\theta = 0$), the area is easily computed. There are four separate cases corresponding to different ranges of angles. Each of these cases is illustrated in Figure A.1, with the expression for the total area A given below each diagram. The final expression for the total area of a corner is given in equation (A.1).

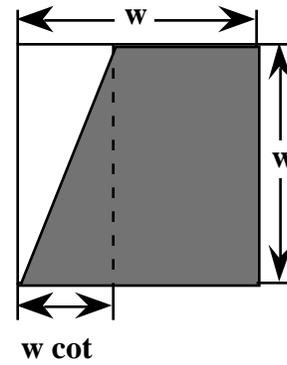
$$A = \begin{cases} w^2 \left(\frac{1}{2} \tan^2 \theta \right), & 0 \leq \theta < \frac{\pi}{4} \\ w^2 \left(1 - \frac{1}{2} \cot^2 \theta \right), & \frac{\pi}{4} < \theta < \frac{3\pi}{4} \\ w^2 \left(2 + \frac{1}{2} \tan^2 \theta \right), & \frac{3\pi}{4} < \theta < \pi \end{cases} \quad (\text{A.1})$$



$$0 < \theta < \frac{\pi}{4}$$

$$A = \frac{1}{2} w \cdot w \tan \theta$$

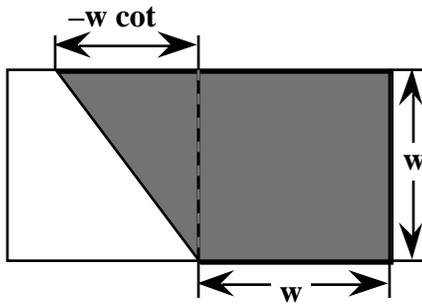
$$= w^2 \left(\frac{1}{2} \tan \theta \right)$$



$$\frac{\pi}{4} < \theta < \frac{\pi}{2}$$

$$A = \frac{1}{2} w \cdot w \cot \theta + w (w - w \cot \theta)$$

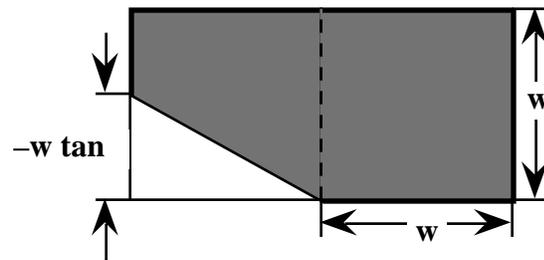
$$= w^2 \left(1 - \frac{1}{2} \cot \theta \right)$$



$$\frac{\pi}{2} < \theta < \frac{3\pi}{4}$$

$$A = w^2 + \frac{1}{2} w (-w \cot \theta)$$

$$= w^2 \left(1 - \frac{1}{2} \cot \theta \right)$$



$$\frac{3\pi}{4} < \theta < \pi$$

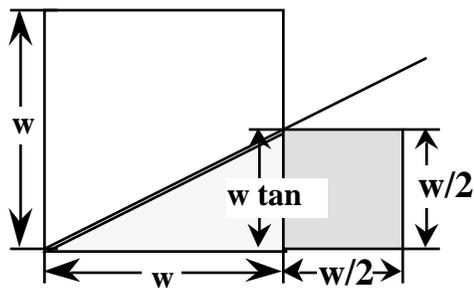
$$A = 2w^2 - \frac{1}{2} w (-w \tan \theta)$$

$$= w^2 \left(2 + \frac{1}{2} \tan \theta \right)$$

Figure A.1. Derivation of total area of corners with angles in the range $0 < \theta < \pi$.

A.2. CORNER RESPONSE OF MORPHOLOGICAL OPENING

The fractional corner preservation $r(\theta, \mathcal{Q})$ of morphological opening with a square structuring element, given in equation (4.14), is found using geometry. Recall that for the compound morphological operators (opening, closing, OC, and CO) the overall filter window is found by compounding the structuring element. Thus, a square overall window of size $w \times w$ corresponds to a square structuring element of size $\frac{w}{2} \times \frac{w}{2}$. Opening preserves all parts of a corner where the structuring element fits completely within the corner, as illustrated in Figure A.2 below. The first diagram shows the case where the structuring element does not fit anywhere inside the corner within the $w \times w$ region of interest. This is true for angles from 0 to the arctangent of $\frac{1}{2}$, or roughly 27° , as shown in the calculation beneath the diagram. (In the diagrams, the smaller square corresponds to the structuring element and the larger square to the region of interest of the corner.) The corner is completely removed by opening in this case, so the preserved area is 0 and therefore $r = 0$ for $0 < \theta < \tan^{-1}\left(\frac{1}{2}\right)$. The second and third cases are for angles in the range $\tan^{-1}\left(\frac{1}{2}\right) < \theta < \frac{\pi}{2}$. In these cases, the preserved area is the total area minus the area of a triangle of height $\frac{w}{2}$ at the tip of the corner. The expression for the total area of the corner changes at 45° , so there are two different expressions for the fractional preservation for angles in this range. For angles greater than 90° , opening preserves the corner exactly and therefore $r = 1$.

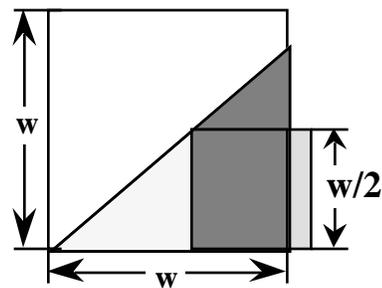


$$w \tan < \frac{w}{2}$$

$$0 < \tan^{-1}\left(\frac{1}{2}\right)$$

$$P = 0$$

$$r = 0$$

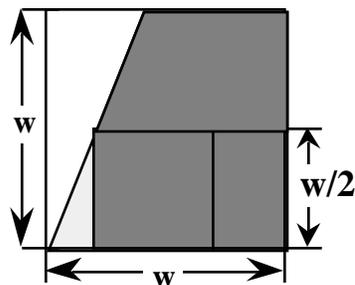


$$\tan^{-1}\left(\frac{1}{2}\right) < \frac{\pi}{4}$$

$$P = A - \frac{1}{2} \frac{w}{2} \left(\frac{w}{2} \cot \right)$$

$$r = \frac{A - w^2 \frac{1}{8} \cot}{A} = 1 - \frac{\frac{1}{8} \cot}{\frac{1}{2} \tan}$$

$$= 1 - \frac{1}{4} \cot^2$$



$$\frac{\pi}{4} < \frac{\pi}{2}$$

$$P = A - \frac{1}{2} \frac{w}{2} \left(\frac{w}{2} \cot \right)$$

$$= w^2 \left(1 - \frac{1}{2} \cot - \frac{1}{8} \cot \right)$$

$$= w^2 \left(1 - \frac{5}{8} \cot \right)$$

$$r = \frac{1 - \frac{5}{8} \cot}{1 - \frac{1}{2} \cot}$$

Figure A.2. Derivation of fractional corner preservation for morphological opening with a square structuring element.

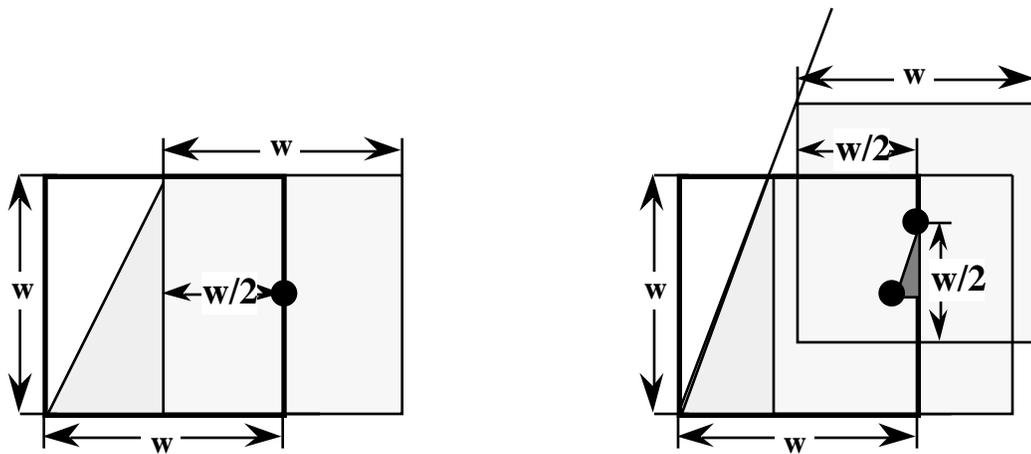
A.3. CORNER RESPONSE OF EROSION AND THE MIDRANGE FILTER

The derivation of the corner response of the midrange filter, equation (4.17), is similar to the derivation shown above, with two important differences. First, since the midrange filter is the average of erosion and dilation, its corner response is the average of the corner response of erosion and dilation. Second, since erosion and dilation are not compound morphological operators, the structuring element is the same size as the overall filter window. For a square structuring element, both the overall window and the structuring element are of size $w \times w$. Dilation has a corner response equal to unity for all angles, so the only geometric calculations that must be made are for erosion. Erosion only gives output where a structuring element centered at a point is completely within the area of the corner. Therefore, the corner response of erosion is zero for a much wider range of acute angles than for opening. The first diagram in Figure A.3 illustrates the smallest angle for which erosion has a nonzero response in the region of interest, $\theta = \tan^{-1}(2)$. For angles smaller than this, the corner response is zero. The second diagram illustrates the case where the preserved area is a small triangle, angles in the range $\tan^{-1}(2) < \theta < \tan^{-1}(3)$. The upper bound here is where the triangle reaches the upper right corner of the region of interest, at $\tan \theta = \frac{3w/2}{w/2}$. Figure A.4 shows the derivation of the corner response of erosion for angles greater than $\tan^{-1}(3)$, up to 135° . For angles greater than 135° , the derivations are similar to the cases for smaller angles in Figure A.3, and are not shown here. The complete expression for the corner response of erosion is given in equation (A.2). The fractional corner preservation of the midrange filter,

equation (4.17), is found by averaging the erosion corner response with one:

$$r_{\text{midr}} = \frac{1}{2}(1 + r_{\text{erode}}).$$

$$r_{\text{erode}}(\beta, 0) = \begin{cases} 0, & |\beta| \leq \tan^{-1}(2) \\ \frac{\tan^2|\beta| - 4\tan|\beta| + 4}{8\tan|\beta| - 4}, & \tan^{-1}(2) < |\beta| \leq \tan^{-1}(3) \\ \frac{2\tan|\beta| - 5}{8\tan|\beta| - 4}, & \tan^{-1}(3) < |\beta| \leq \frac{\pi}{2} \\ \frac{1}{4}, & \frac{\pi}{2} < |\beta| \leq \frac{3\pi}{4} \\ \frac{2 - \cot|\beta|}{16 + 4\tan|\beta|}, & \frac{3\pi}{4} < |\beta| \leq \tan^{-1}\left(-\frac{1}{3}\right) \\ \frac{8 + 9\tan|\beta|}{16 + 4\tan|\beta|}, & \tan^{-1}\left(-\frac{1}{3}\right) < |\beta| \end{cases} \quad (\text{A.2})$$



$$0 \quad \tan^{-1} \left(\frac{w}{w/2} \right)$$

$$0 \quad \tan^{-1}(2)$$

$$P = 0$$

$$r = 0$$

$$\tan^{-1}(2) < \tan^{-1}(3)$$

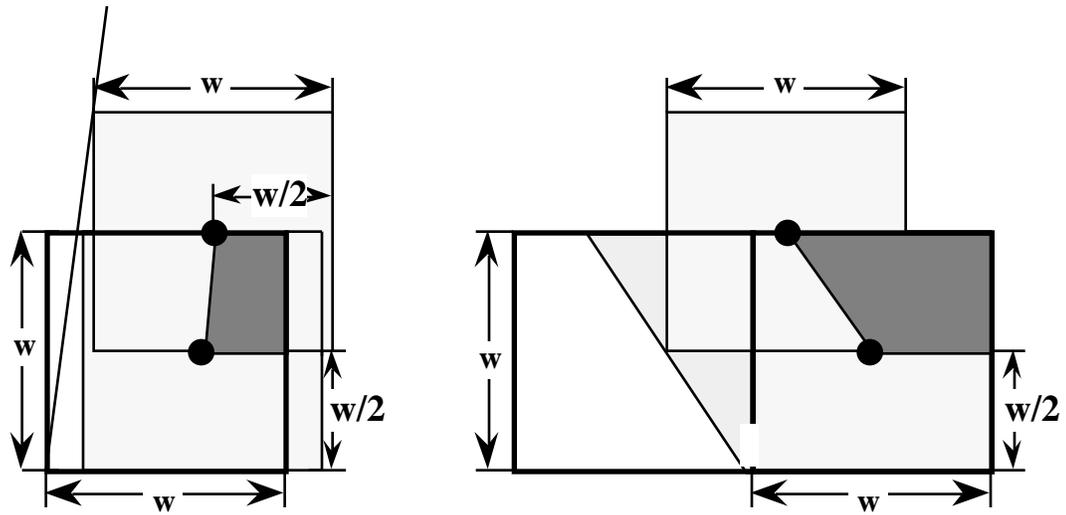
$$P = \frac{1}{2} \left(\frac{w}{2} - w \cot \theta \right) \left(\frac{w}{2} - w \cot \theta \right) \tan \theta$$

$$= w^2 \frac{1}{8} (\tan^2 \theta - 4 + 4 \cot \theta)$$

$$r = \frac{w^2 \frac{1}{8} (\tan^2 \theta - 4 + 4 \cot \theta)}{w^2 \left(1 - \frac{1}{2} \cot \theta \right)}$$

$$= \frac{\tan^2 \theta - 4 \tan \theta + 4}{8 \tan \theta - 4}$$

Figure A.3. Derivation of fractional corner preservation for morphological erosion with a square structuring element (small angles).



$$\tan^{-1}(3) < \frac{\pi}{2}$$

$$P = \frac{1}{2} \left(\frac{w}{2} - \frac{w}{2} \cot \theta \right) + \frac{w}{2} \left(\frac{w}{2} - 3 \frac{w}{2} \cot \theta \right)$$

$$= w^2 \left(\frac{1}{4} - \frac{5}{8} \cot \theta \right)$$

$$r = \frac{w^2 \left(\frac{1}{4} - \frac{5}{8} \cot \theta \right)}{w^2 \left(1 - \frac{1}{2} \cot \theta \right)}$$

$$= \frac{2 \tan \theta - 5}{8 \tan \theta - 4}$$

$$\frac{\pi}{2} < \frac{3}{4}$$

$$P = \frac{w}{2} \left(\frac{w}{2} + \frac{1}{2} \frac{w}{2} \cot \theta \right) + \frac{1}{2} \frac{w}{2} \left(\frac{w}{2} \cot \theta \right)$$

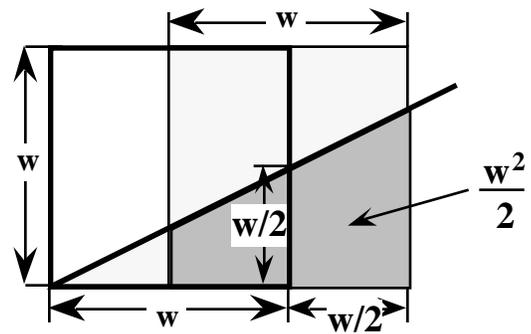
$$= w^2 \frac{1}{4} \left(1 + \frac{1}{2} \cot \theta \right)$$

$$r = \frac{1}{4}$$

Figure A.4. Derivation of fractional corner preservation for morphological erosion with a square structuring element (larger angles).

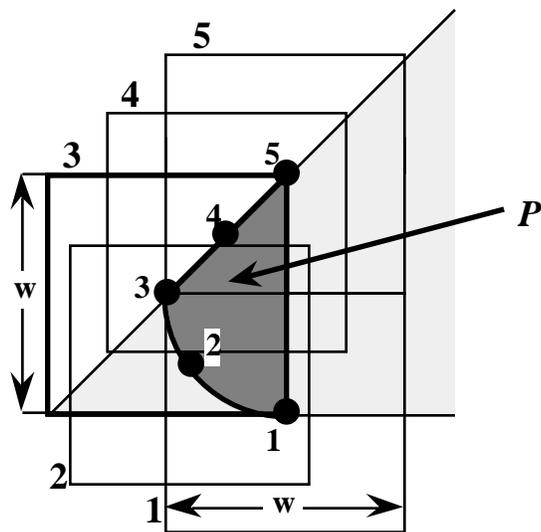
A.4. CORNER RESPONSE OF THE SQUARE MEDIAN FILTER

For the median filter with a square window, equation (4.18) gives the corner response at selected angles. The range of angles where the response is zero is the same as for morphological opening, $0 < \theta < \tan^{-1}\left(\frac{1}{2}\right)$. This is illustrated in Figure A.5. The lightly shaded square represents a filter window that has half of its area (the portion that is more darkly shaded) within the corner. The output of the median filter at the center of this window is one. At $\theta = \tan^{-1}\left(\frac{1}{2}\right)$, the points where the median filter gives an output of 1 begin to reach the region of interest, which is the square with the darker outline. For $\theta = \frac{\pi}{4}$, the preserved area is shown in the second diagram of Figure A.5. The five numbered points denote the centers of the respectively numbered squares, each of which represents a median filter window that has exactly one-half of its area in the corner (lightly shaded area). The preserved area P (darkly shaded area) lies to the right of these points and is limited to the region of interest, the darkly outlined square numbered 3. The area P is calculated using parametric equations, as shown below the diagram.



$$w \tan \theta < \frac{w}{2} \quad P = 0$$

$$0 < \theta < \tan^{-1}\left(\frac{1}{2}\right) \quad r = 0$$



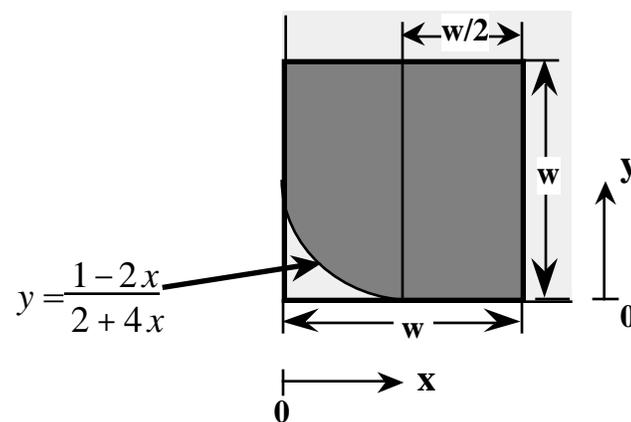
$$P = w^2 \int_{x=0}^{1/2} \int_{y=0}^{-x+\sqrt{2x}} 1 \, dy \, dx + \int_{x=1/2}^1 \int_{y=0}^{1-x} 1 \, dy \, dx$$

$$= w^2 \frac{1}{3}$$

$$r = \frac{2}{3}$$

Figure A.5. Derivation of fractional corner preservation for the square median filter (small angles).

The shape of the preserved area for a median-filtered right angle corner ($\theta = \frac{\pi}{2}$) is found in a manner similar to that for the 45° corner above. The equation for the line limiting the preserved area is shown in the diagram in Figure A.6. The calculations for the preserved area and corner preservation are shown below the diagram. The median filter preserves a binary straight edge ($\theta = \pi$) exactly, so the fractional preservation of the square median filter for a 180° angle is 1.



$$\theta = \frac{\pi}{2}$$

$$P = \frac{w^2}{2} + w^2 \int_0^{\frac{1}{2}} \left(1 - \frac{1-2x}{2+4x}\right) dx = w^2 \left[\frac{x}{2} + \int_0^{\frac{1}{2}} \frac{6x+1}{4x+2} dx \right]$$

$$= w^2 \left(\frac{5}{4} - \frac{1}{2} \ln 4 + \frac{1}{2} \ln 2 \right)$$

$$r = \frac{5}{4} - \frac{1}{2} \ln 4 + \frac{1}{2} \ln 2$$

Figure A.6. Derivation of fractional corner preservation for the square median filter for a right angle corner ($\theta = \frac{\pi}{2}$).

A.5. CORNER RESPONSE OF THE AVERAGING FILTER

For the averaging filter, the calculations for the corner response become more difficult because the output of the filter is not binary. For small angles, the response can still be determined because the corner covers a simple trapezoidal region of all the square filter windows. The limit of this case is $\tan^{-1}\left(\frac{2}{3}\right)$, as shown in Figure A.7. An expression for the output of the averaging filter over the region of interest is needed to determine the percentage of the corner intensity preserved. To simplify these calculations without any loss of generality, let the size of the filter window be 1×1 ($w = 1$). First, a function describing the corner is found:

$$f(x, y) = \begin{cases} 1, & 0 \leq y \leq x \tan \theta \\ 0, & \text{otherwise} \end{cases} \quad (\text{A.3})$$

Then, an expression for the output of the averaging filter $a(x, y)$ is found by integrating $f(x, y)$.

$$a(x, y) = \int_{x-\frac{1}{2}}^{x+\frac{1}{2}} \int_{y-\frac{1}{2}}^{y+\frac{1}{2}} f(x, y) dy dx \quad (\text{A.4})$$

For $0 \leq \theta < \tan^{-1}\left(\frac{2}{3}\right)$, the limits of the corner are more restrictive than the limits on y , yielding the following expression:

$$a(x, y) = \begin{cases} \int_{x-\frac{1}{2}}^{x+\frac{1}{2}} \int_{x \tan \theta}^{y+\frac{1}{2}} 1 dy dx, & 0 \leq x \leq \frac{1}{2} \\ \int_{x-\frac{1}{2}}^0 \int_{x \tan \theta}^0 1 dy dx, & \frac{1}{2} \leq x \leq 1 \\ \int_{x-\frac{1}{2}}^0 0 dy dx, & \text{otherwise} \end{cases} \quad (\text{A.5})$$

Equation (A.5) integrates to yield the following result for the output of the averaging filter when $0 < \tan^{-1}\left(\frac{2}{3}\right)$.

$$a(x, y) = \begin{cases} \frac{\left(x + \frac{1}{2}\right)^2}{2} \tan^{-1}\left(\frac{2}{3}\right), & 0 \leq x \leq \frac{1}{2} \\ \frac{\left(x + \frac{1}{2}\right)^2}{2} - \frac{\left(x - \frac{1}{2}\right)^2}{2} \tan^{-1}\left(\frac{2}{3}\right), & \frac{1}{2} \leq x \leq 1 \end{cases} \quad (\text{A.6})$$

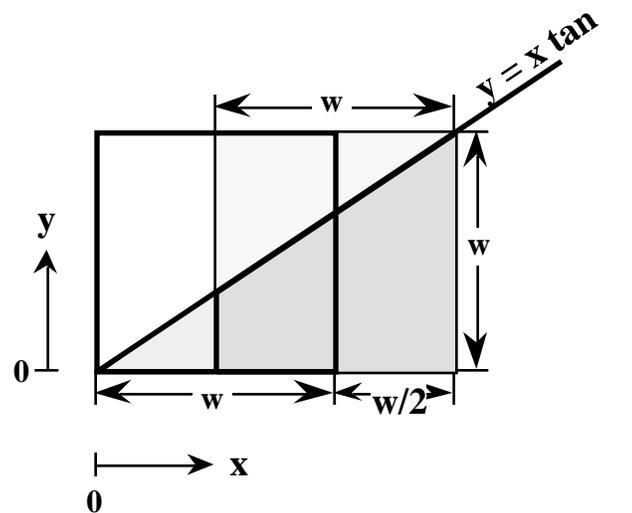
This equation simplifies to:

$$a(x, y) = \begin{cases} \frac{\left(x + \frac{1}{2}\right)^2}{2} \tan^{-1}\left(\frac{2}{3}\right), & 0 \leq x \leq \frac{1}{2} \\ x \tan^{-1}\left(\frac{2}{3}\right), & \frac{1}{2} \leq x \leq 1 \end{cases} \quad (\text{A.7})$$

Using equation (A.7) for the output of the averaging filter, the preserved intensity can be found by integrating over the region of interest; that is, the area of the original binary corner: 0 to 1 in x and 0 to $x \tan^{-1}\left(\frac{2}{3}\right)$ in y .

$$\begin{aligned} P &= \int_0^1 \int_0^{x \tan^{-1}\left(\frac{2}{3}\right)} a(x, y) dy dx \\ &= \int_0^{\frac{1}{2}} \int_0^{x \tan^{-1}\left(\frac{2}{3}\right)} \frac{\left(x + \frac{1}{2}\right)^2}{2} \tan^{-1}\left(\frac{2}{3}\right) dy dx + \int_{\frac{1}{2}}^1 \int_0^{x \tan^{-1}\left(\frac{2}{3}\right)} x \tan^{-1}\left(\frac{2}{3}\right) dy dx \end{aligned} \quad (\text{A.8})$$

Recall that the area of the original corner is given in equation (A.1), and that in this example $w = 1$. The final expressions for the preserved intensity and fractional corner preservation for the averaging filter with $0 < \tan^{-1}\left(\frac{2}{3}\right)$ are given in Figure A.7 below.



$$0 \quad \tan^{-1}\left(\frac{2}{3}\right) \quad w = 1$$

$$\begin{aligned}
 P &= \int_0^{\frac{1}{2}} \frac{\tan^2}{2} \left(x^3 + x^2 + \frac{1}{4}x \right) dx + \int_{\frac{1}{2}}^1 x^2 \tan^2 dx \\
 &= \frac{17}{384} \tan^2 + \frac{7}{24} \tan^2 \\
 &= \frac{43}{128} \tan^2
 \end{aligned}$$

$$\begin{aligned}
 r &= \frac{P}{A} = \frac{\frac{43}{128} \tan^2}{\frac{1}{2} \tan} \\
 &= \frac{43}{64} \tan
 \end{aligned}$$

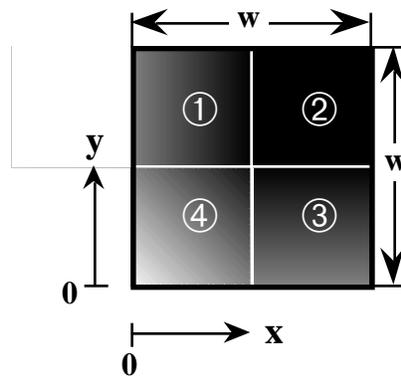
Figure A.7. Derivation of fractional corner preservation for the square averaging filter (small angles).

For $\theta = 0$ (a straight edge), the intensity preservation of the averaging filter is easily found. The top half of the area is preserved with intensity 1, and the bottom half has a linearly decreasing intensity from 1 (at the top) to $\frac{1}{2}$ (at the bottom). Therefore, the bottom half of the area is preserved at an average level of $\frac{3}{4}$, and the top half is preserved at level 1. The overall fractional preservation is then $r_{\text{avg}}(\theta, 0) = \frac{1}{2} \cdot 1 + \frac{1}{2} \cdot \frac{3}{4} = \frac{7}{8}$.

Figure A.8 illustrates the derivation of the fractional corner preservation of the averaging filter for $\theta = \frac{\pi}{2}$, a right angle. The area ② in the figure is preserved at level 1, while areas ① and ③ are preserved at an average level of $\frac{3}{4}$, because they have linearly decreasing intensity from 1 to $\frac{1}{2}$ across their entire width. Area ④, however, requires a more complex integration. The expression for the output of the averaging filter in this area is given in equation (A.9) below.

$$\begin{aligned}
 a(x, y) &= \int_0^{y+\frac{1}{2}} \int_0^{x+\frac{1}{2}} 1 \, dx \, dy \\
 &= \left(x + \frac{1}{2}\right) \left(y + \frac{1}{2}\right) \\
 &\text{for } 0 \leq x \leq \frac{1}{2} \text{ and } 0 \leq y \leq \frac{1}{2}
 \end{aligned} \tag{A.9}$$

Integrating equation (A.9) over area ④ gives the intensity preservation for that region. The final expression for the fractional corner preservation of the averaging filter for a right angle is shown in Figure A.8. This result, along with the previous results for small angles and for a straight edge, leads to the expression for the averaging filter corner response given in equation (4.19) in Chapter 4.



$$= \frac{1}{2}$$

$$w = 1$$

$$P = ① + ② + ③ + ④$$

$$= \frac{1}{4} \frac{3}{4} + \frac{1}{4} \left(1 + \frac{1}{4} \frac{3}{4} + \int_0^{\frac{1}{2}} \int_0^{\frac{1}{2}} (x + \frac{1}{2})(y + \frac{1}{2}) dx dy \right)$$

$$= \frac{1}{4} + \frac{3}{16} + \frac{3}{16} + \frac{9}{64}$$

$$= \frac{49}{64}$$

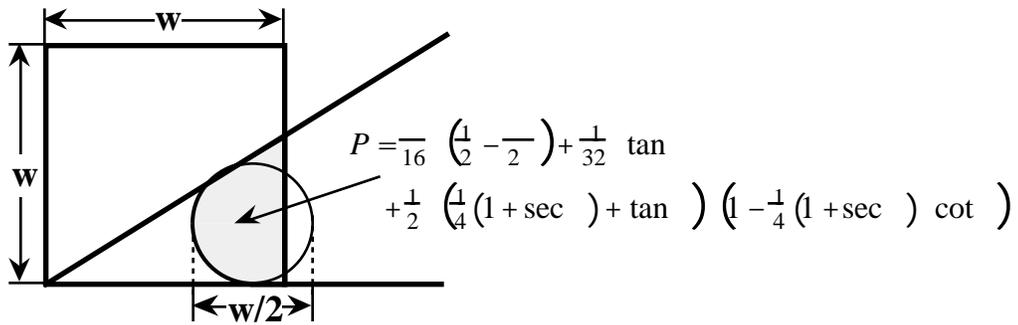
$$r = \frac{P}{A} = \frac{49/64}{1} = \frac{49}{64}$$

Figure A.8. Derivation of fractional corner preservation for the square averaging filter ($= \frac{1}{2}$).

A.6. CORNER RESPONSE OF THE PLUS-SHAPED MEDIAN FILTER

For the plus-shaped median filter, the techniques used to derive the fractional preservation are similar to those used for the square median filter and the morphological filters. Figure A.9 illustrates the area preserved by the plus-shaped median filter for acute corners. The “arms” of the filter window are infinitesimally thin, so the corner is preserved only where one-half of the total length of the arms fits within the original corner. The filtered corners have straight-line edges, unlike the results of square median filtering. The calculation of the preserved area is shown in Figure A.9 for $0 < \theta < \tan^{-1}\left(\frac{1}{2}\right)$. The derivations for other angles are similar. For $\theta \geq \tan^{-1}\left(\frac{1}{2}\right)$, more than half of the filter window is in the corner at all locations in the corner, so the corner is preserved exactly ($r = 1$). The complete expression for the fractional corner preservation of the plus-shaped median filter is given in equation (4.20).

The expression for the fractional preservation of morphological opening with a circular structuring element is given in equation (4.21). The derivation of this expression is quite involved, but follows exactly the same lines as the derivation for opening with a square structuring element. The circular structuring element has diameter $\frac{w}{2}$ because opening is a compound operator. An example of the calculation for small angles is shown in Figure A.10.



$$P = \frac{1}{16} \left(\frac{1}{2} - \frac{1}{2} \right) + \frac{1}{32} \tan \theta + \frac{1}{2} \left(\frac{1}{4} (1 + \sec \theta) + \tan \theta \right) \left(1 - \frac{1}{4} (1 + \sec \theta) \cot \theta \right)$$

$$\frac{4}{(4 \sin \theta = 1 + \cos \theta)}$$

$$P = \frac{1}{32} \left(- + \tan \theta \right) + \frac{1}{2} \left(\frac{1}{4} + \frac{1}{4} \sec \theta + \tan \theta \right) \left(1 - \frac{1}{4} \cot \theta - \frac{1}{4} \csc \theta \right)$$

$$r = 1 - 2 \cot \theta \left[\frac{1}{16} \frac{\sin \theta}{1 - \cos \theta} - \frac{1}{32} \left(- \right) \right]$$

Figure A.10. Derivation of fractional corner preservation for circle-shaped morphological opening (small angles).

Bibliography

1. I. Pitas and A. N. Venetsanopoulos, *Nonlinear Digital Filters: Principles and Applications*. Boston: Kluwer Academic, 1990.
2. A. V. Oppenheim and R. W. Schaffer, *Discrete-Time Signal Processing*. Englewood Cliffs, New Jersey: Prentice Hall, 1989.
3. J. W. Tukey, *Exploratory Data Analysis*. Reading, Massachusetts: Addison-Wesley, 1971.
4. R. P. Borda and J. D. Frost, Jr., "Error reduction in small sample averaging through the use of the median rather than the mean," *Electroencephalography and Clinical Neurophysiology*, vol. 25, pp. 391-392, 1968.
5. P. Maragos and R. W. Schaffer, "Morphological filters—Part I: Their set-theoretic analysis and relations to linear shift-invariant filters," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 35, no. 8, pp. 1153-1169, 1987.
6. P. Maragos and R. W. Schaffer, "Morphological filters—Part II: Their relations to median, order-statistic, and stack filters," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 35, no. 8, pp. 1170-1184, 1987.
7. B. I. Justusson, "Median filtering: Statistical properties," in *Two-Dimensional Digital Signal Processing II: Transforms and Median Filtering*, T. S. Huang, Editor. Berlin: Springer-Verlag, 1981. pp. 161-196.
8. S. G. Tyan, "Median filtering: Deterministic properties," in *Two-Dimensional Digital Signal Processing II: Transforms and Median Filtering*, T. S. Huang, Editor. Berlin: Springer-Verlag, 1981. pp. 197-217.
9. N. C. Gallagher, Jr. and G. L. Wise, "A theoretical analysis of the properties of the median filter," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 29, no. 6, pp. 1136-1141, 1981.
10. G. Matheron, *Random Sets and Integral Geometry*. New York: Wiley, 1974.
11. J. Serra, *Image Analysis and Mathematical Morphology*, Vol. 1. London: Academic, 1982.

12. J. Serra, Editor. *Image Analysis and Mathematical Morphology*, Vol. 2: Theoretical Advances. London: Academic, 1988.
13. R. L. Stevenson and G. R. Arce, "Morphological filters: Statistics and further syntactic properties," *IEEE Trans. Circuits Syst.*, vol. 34, no. 11, pp. 1292-1305, 1987.
14. G. Gerig, O. Kübler, R. Kikinis, and F. A. Jolesz, "Nonlinear anisotropic filtering of MRI data," *IEEE Trans. Med. Imag.*, vol. 11, no. 2, pp. 221-232, 1992.
15. S. R. Sternberg, "Biomedical image processing," *IEEE Computer*, vol. 16, no. 1, pp. 22-34, 1983.
16. S. R. Sternberg, "Grayscale morphology," *Comp. Vision, Graphics, Image Process.*, vol. 35, pp. 333-355, 1986.
17. P. D. Wendt, E. J. Coyle, and N. C. Gallagher, "Stack filters," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 34, no. 4, pp. 898-911, 1986.
18. J. P. Fitch, E. J. Coyle, and N. C. Gallagher, "Median filtering by threshold decomposition," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 32, pp. 1183-1188, 1984.
19. J. P. Fitch, E. J. Coyle, and N. C. Gallagher, "Threshold decomposition of multidimensional ranked order operations," *IEEE Trans. Circuits Syst.*, vol. 32, pp. 445-450, 1985.
20. E. J. Coyle and J.-H. Lin, "Stack filters and the mean absolute error criterion," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 36, no. 8, pp. 1244-1254, 1988.
21. H. A. David, *Order Statistics*, 1st ed. New York: Wiley, 1970.
22. R. V. Hogg and A. T. Craig, *Introduction to Mathematical Statistics*. 4th ed. New York: Macmillan, 1978.
23. L. Koskinen, J. Astola, and Y. Neuvo, "Morphological filtering of noisy images," in *Visual Communications and Image Processing '90*, M. Kunt, Editor, *Proc. SPIE*, vol. 1360, pp. 155-165, 1990.
24. J. Astola, L. Koskinen, and Y. Neuvo, "Statistical properties of discrete morphological filters," in *Mathematical Morphology in Image Processing*, E. R. Dougherty, Editor. New York: Marcel Dekker, 1993. pp. 93-120.
25. H. A. David, *Order Statistics*, 2nd ed. New York: Wiley, 1980.

26. A. C. Bovik, T. S. Huang, and D. C. Munson, Jr., "A generalization of median filtering using linear combinations of order statistics," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 31, no. 6, pp. 1342-1350, 1983.
27. A. Restrepo and A. C. Bovik, "Adaptive trimmed mean filters for image restoration," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 36, no. 8, pp. 1326-1337, 1988.
28. W. K. Pratt, T. J. Cooper, and I. Kabir, "Pseudomedian filter," in *Architectures and Algorithms for Digital Image Processing II*, F. J. Corbett, Editor, *Proc. SPIE*, vol. 534, pp. 34-43, 1985.
29. W. K. Pratt, *Digital Image Processing*, 2nd ed. New York: Wiley, 1991.
30. M. A. Schulze and J. A. Pearce, "Continuous time analysis of the response of the pseudomedian and related filters to periodic signals," in *Nonlinear Image Processing III*, E. R. Dougherty, J. Astola, and C. G. Boncelet, Jr., Editors, *Proc. SPIE*, vol. 1658, pp. 177-188, 1992.
31. M. A. Schulze, *Mathematical Properties of the Pseudomedian Filter*. M.S. Thesis, University of Texas at Austin, 1990.
32. M. A. Schulze and J. A. Pearce, "Some properties of the two-dimensional pseudomedian filter," in *Nonlinear Image Processing II*, E. R. Dougherty, G. R. Arce, and C. G. Boncelet, Jr., Editors, *Proc. SPIE*, vol. 1451, pp. 48-57, 1991.
33. J. Song and E. J. Delp, "A study of the generalized morphological filter," *Circuits, Syst., Signal Process.*, vol. 11, no. 1, pp. 227-252, 1992.
34. M. A. Schulze and J. A. Pearce, "Linear combinations of morphological operators: The midrange, pseudomedian, and LOCO filters," in *Proc. 1993 IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. V, pp. 57-60, 1993.
35. E. J. Gumbel, *Statistics of Extremes*. New York: Columbia University Press, 1958.
36. J. Y. Hsiao and A. A. Sawchuk, "Supervised textured image segmentation using feature smoothing and probabilistic relaxation techniques," *IEEE Trans. Patt. Anal. Mach. Intell.*, vol. 11, no. 12, pp. 1279-1292, 1989.
37. M. Kuwahara, K. Hachimura, S. Eiho, and M. Kinoshita, "Processing of RI-angiographic images," in *Digital Processing of Biomedical*

- Images*, K. Preston, Jr. and M. Onoe, Editors. New York: Plenum, 1976. pp. 187-202.
38. F. Tomita and S. Tsuji, "Extraction of multiple regions by smoothing in selected neighborhoods," *IEEE Trans. Syst., Man, Cyber.*, vol. 7, pp. 107-109, 1977.
 39. M. Nagao and T. Matsuyama, "Edge preserving smoothing," *Comp. Graphics Image Process.*, vol. 9, pp. 394-407, 1979.
 40. J.-S. Lee, "Digital image enhancement and noise filtering by use of local statistics," *IEEE Trans. Patt. Anal. Mach. Intell.*, vol. 2, pp. 165-168, 1980.
 41. J.-S. Lee, "Refined filtering of image noise using local statistics," *Comp. Graphics Image Process.*, vol. 15, pp. 380-389, 1981.
 42. J. Neejärvi, P. Heinonen, and Y. Neuvo, "Sine wave responses of median type filters," in *Proc. 1988 IEEE International Symposium on Circuits and Systems*, pp. 1503-1506, 1988.
 43. J. Neejärvi and Y. Neuvo, "Sinusoidal and pulse responses of morphological filters," in *Proc. 1990 IEEE International Symposium on Circuits and Systems*, vol. 3, pp. 2136-2139, 1990.
 44. J. P. Fitch, E. J. Coyle, and N. C. Gallagher, "The analog median filter," *IEEE Trans. Circuits Syst.*, vol. 33, no. 1, pp. 94-102, 1986.
 45. F. R. Hampel, "The breakdown points of the mean combined with some rejection rules," *Technometrics*, vol. 27, no. 2, pp. 95-107, 1985.
 46. D. F. Andrews, P. J. Bickel, F. R. Hampel, P. J. Huber, W. H. Rogers, and J. W. Tukey, *Robust Estimates of Location: Survey and Advances*. Princeton: Princeton University Press, 1972.
 47. H. G. Longbotham and N. Barsalou, "The LMS, an adaptive optimal order statistic filter," in *Nonlinear Image Processing*, E. J. Delp, Editor, *Proc. SPIE*, vol. 1247, pp. 78-88, 1990.
 48. J. B. Bednar and T. L. Watt, "Alpha-trimmed means and their relationship to median filters," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 32, no. 1, pp. 145-153, 1984.
 49. C. G. Boncelet, Jr., R. Hardie, R. Hakami, and G. R. Arce, "LUM filters for smoothing and sharpening," in *Nonlinear Image Processing II*, E. R. Dougherty, G. R. Arce, and C. G. Boncelet, Jr., Editors, *Proc. SPIE*, vol. 1451, pp. 70-74, 1991.

50. A. Rosenfeld and A. C. Kak, *Digital Picture Processing*, 2nd ed. Vol. 1. Orlando: Academic, 1982.
51. A. C. Bovik, T. S. Huang, and D. C. Munson, "The effect of median filtering on edge estimation and detection," *IEEE Trans. Patt. Anal. Mach. Intell.*, vol. 9, pp. 181-194, 1987.
52. H. Longbotham and D. Eberly, "The WMMR filters: A class of robust edge enhancers," *IEEE Trans. Signal Process.*, vol. 41, no. 4, pp. 1680-1685, 1993.
53. Z. M. Ryu, *The Effect of Nonlinear Filtering on the Resolution of Calibrated Thermal Images*. Ph.D. Dissertation, University of Texas at Austin, 1986.
54. C.-H. H. Chu and E. J. Delp, "Impulsive noise suppression and background normalization of electrocardiogram signals using morphological operators," *IEEE Trans. Biomed. Eng.*, vol. 36, no. 2, pp. 262-273, 1989.
55. P. E. Trahanias, "An approach to QRS complex detection using mathematical morphology," *IEEE Trans. Biomed. Eng.*, vol. 40, no. 2, pp. 201-205, 1993.
56. J. W. Klingler, Jr., C. L. Vaughan, T. D. Fraker, Jr., and L. T. Andrews, "Segmentation of echocardiographic images using mathematical morphology," *IEEE Trans. Biomed. Eng.*, vol. 35, no. 11, pp. 925-934, 1988.
57. J. G. Thomas, I. Peters R. A., and P. Jeanty, "Automatic segmentation of ultrasound images using morphological operators," *IEEE Trans. Med. Imag.*, vol. 10, no. 2, pp. 180-186, 1991.
58. C. Lamberti and F. Sgallari, "Anisotropic diffusion and morphological approaches for echocardiography image processing," in *Proc. Fifth European Signal Processing Conference*, L. Torres, E. Masgrau, and M. A. Lagunas, Editors. Amsterdam: Elsevier Science, 1990. pp. 939-942.
59. B. D. Thackray and A. C. Nelson, "Semi-automatic segmentation of vascular network images using a rotating structuring element (ROSE) with mathematical morphology and dual feature thresholding," *IEEE Trans. Med. Imag.*, vol. 12, no. 3, pp. 385-392, 1993.
60. J. Samarabandu, R. Acharya, E. Hausmann, and K. Allen, "Analysis of bone X-rays using morphological fractals," *IEEE Trans. Med. Imag.*, vol. 12, no. 3, pp. 466-470, 1993.

61. D. L. Wilson and C. Bertram, "Morphological enhancement of coronary angiograms," in *Proc. Computers in Cardiology*, J. Meyer, Editor. Los Alamitos, California: IEEE Computer Society Press, 1990. pp. 313-316.
62. L. Vincent and B. Masters, "Morphological image processing and network analysis of cornea endothelial cell images," in *Image Algebra and Morphological Image Processing III*, P. D. Gader, E. R. Dougherty, and J. C. Serra, Editors, *Proc. SPIE*, vol. 1769, pp. 212-225, 1992.
63. T. J. McMurray, *Thermal Damage in Tissues: Birefringence Image Analysis*. Ph.D. Dissertation, University of Texas at Austin, 1994.
64. R. V. Hogg, "Adaptive robust procedures: A partial review and some suggestions for future applications and theory," *J. Amer. Statist. Assoc.*, vol. 69, no. 348, pp. 909-923, 1974.
65. K. K. Shung, M. B. Smith, and B. M. W. Tsui, *Principles of Medical Imaging*. San Diego: Academic, 1992.
66. M. E. Brummer, R. M. Mersereau, R. L. Eisner, and R. R. J. Lewine, "Automatic detection of brain contours in MRI data sets," *IEEE Trans. Med. Imag.*, vol. 12, no. 2, pp. 153-166, 1993.
67. P. Perona and J. Malik, "Scale space and edge detection using anisotropic diffusion," in *Proc. IEEE Computer Society Workshop on Computer Vision*, K. Price, Editor. Los Alamitos, California: IEEE Computer Society Press, 1987. pp. 16-22.
68. P. Perona and J. Malik, "Scale-space and edge detection using anisotropic diffusion," *IEEE Trans. Patt. Anal. Mach. Intell.*, vol. 12, no. 7, pp. 629-639, 1990.
69. S. K. Chang, *Principles of Pictorial Information Systems Design*. Englewood Cliffs, New Jersey: Prentice Hall, 1989.
70. T. McMurray and J. A. Pearce, "Theoretical and experimental comparison of the Lorenz information measure, entropy, and the mean absolute error," in *Proc. 1994 IEEE Southwest Symposium on Image Analysis and Interpretation*, A. Jain, Editor, pp. 24-29, 1994.
71. L. A. Shepp and B. F. Logan, "The Fourier reconstruction of a head section," *IEEE Trans. Nucl. Sci.*, vol. 21, no. 1, pp. 21-43, 1974.
72. R. F. Wagner, S. W. Smith, J. M. Sandrik, and H. Lopez, "Statistics of speckle in ultrasound B-scans," *IEEE Trans. Sonics Ultrasonics*, vol. 30, no. 3, pp. 156-163, 1983.

Vita

Mark Allen Schulze was born in Fort Worth, Texas on April 24, 1966, the son of Sharon Havemann Schulze and Arthur Edward Schulze. After graduating as valedictorian of Sharpstown High School in Houston, Texas in 1984, he enrolled in Rice University, Houston, Texas. He received the Bachelor of Arts degree in managerial studies and the Bachelor of Science degree in electrical engineering, both *summa cum laude*, from Rice University in May 1988. While at Rice he was elected to the Phi Beta Kappa, Tau Beta Pi, and Eta Kappa Nu honor societies. He was employed at TRW Electronic Systems Group, Redondo Beach, California, in the summer of 1987 and at National Instruments, Austin, Texas, in 1988 and 1989. In August of 1990, he received the Master of Science degree in electrical engineering from the University of Texas at Austin. In 1990-91 he attended the University of Southern California, Los Angeles, California, and in September 1991 he returned the Graduate School of the University of Texas at Austin. He was the recipient of a National Science Foundation Graduate Fellowship from 1989 to 1992 and an AT&T Bell Laboratories Ph.D. Scholarship from 1992 to 1994. He was employed at AT&T Bell Laboratories, Murray Hill, New Jersey, in the summer of 1993.

This dissertation was typed by the author.