

MATHEMATICAL PROPERTIES OF THE
PSEUDOMEDIAN FILTER

APPROVED:

Supervisor: _____

Dr. John A. Pearce

Dr. Jonathan W. Valvano

MATHEMATICAL PROPERTIES OF THE
PSEUDOMEDIAN FILTER

by

MARK ALLEN SCHULZE, B.A., B.S.E.E.

THESIS

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

MASTER OF SCIENCE IN ENGINEERING

THE UNIVERSITY OF TEXAS AT AUSTIN

August 1990

Acknowledgments

Many people have helped me as I have pursued this research and written this thesis. Everyone in the Biomedical Engineering department has been very helpful and I thank them all for their kind assistance over the past two years. In particular, the recommendations and guidance of Dr. Pearce have improved the quality of my work immensely. Dr. Valvano also has provided some very valuable suggestions, especially in the development of Chapter 4. My coworkers in Dr. Pearce's laboratory have aided this work in many ways, and specifically I thank Eric Chan and Xiaolu Zheng for their darkroom assistance. Thanks also to Ed Loewenstein for always seeming to give me a key insight at the right moment during my research.

I am grateful to the National Science Foundation for providing me with a generous fellowship, and to all those who helped me apply for and receive the fellowship. Most importantly, I thank my father, mother, and brother, Keith, for their constant love and encouragement.

Mark Schulze
August 1, 1990

This material is based upon work supported under a National Science Foundation Graduate Fellowship.

Any opinions, findings, conclusions or recommendations expressed in this publication are those of the author and do not necessarily reflect the views of the National Science Foundation.

ABSTRACT

MATHEMATICAL PROPERTIES OF THE PSEUDOMEDIAN FILTER

by

MARK ALLEN SCHULZE, B.A., B.S.E.E.

SUPERVISING PROFESSOR: DR. JOHN A. PEARCE

The pseudomedian filter was designed to be a computationally efficient alternative to the median filter. However, a thorough analysis of the pseudomedian filter reveals some important differences between its response and that of the median filter. Several theorems describe the set of signals that are invariant to pseudomedian filtering, and show that this set is a subset of the set of signals invariant to median filtering, with the difference between the sets consisting only of fast-fluctuating signals. The pseudomedian filter does not completely remove impulses, as does the median filter, but both filters preserve edges. The responses of these filters to edges and impulses contrasts with those of the average and midrange filters, which neither preserve edges nor remove impulses. A generalization of the filters to continuous time reveals characteristics of the filter responses to periodic signals, particularly the ability of the pseudomedian filter to block high frequency signals that the median filter cannot. The response of the median filter to high-frequency periodic signals resembles that of the average filter, whereas the response of the pseudomedian filter resembles that of the midrange filter. A square-shaped two-dimensional

definition for the pseudomedian filter preserves sharp corners and fine details better than the square-shaped two-dimensional median filter. As is true for the one-dimensional filters, the two-dimensional median filter is susceptible to high-frequency periodic noise and the two-dimensional pseudomedian filter is not. Pseudomedian- and midrange-filtered images often have a "blocky" appearance, while similar median- and average-filtered images do not. These properties of the pseudomedian filter distinguish it from the median, average, and midrange filters and show its superior performance on signals and images without highly impulsive noise and with fine details, sharp corners, or high frequency periodic noise.

Table of Contents

Chapter 1	Introduction	1
	Pseudomedian Filter Definition.....	2
	Overview.....	3
Chapter 2	Root Signal Analysis	4
	Root Signal Theory.....	4
	Convergence to Root Signals.....	12
Chapter 3	Response to Edges and Impulses	13
Chapter 4	Continuous-Time Analysis	22
	Continuous-Time Filter Analogues.....	22
	Response Curves.....	26
	Distortion Analysis.....	44
Chapter 5	The Two-Dimensional Pseudomedian Filter	52
	Examples.....	58
	Applications.....	61
Chapter 6	Conclusions	75
	Bibliography	78

List of Figures

Figure 2.1	Diagram of signal segment considered for Lemma 3.....	8
Figure 3.1	One-point impulse filtered by 5-wide pseudomedian, median, average, and midrange filters.....	14
Figure 3.2	Two-point impulse filtered by 5-wide pseudomedian, median, average, and midrange filters.....	15
Figure 3.3	Three-point impulse filtered by 7-wide pseudomedian, median, average, and midrange filters.....	16
Figure 3.4	Response of 15-wide filters to spike noise on a ramp signal....	17
Figure 3.5	Sawtooth wave filtered by 5-wide pseudomedian, median, average, and midrange filters.....	19
Figure 3.6	Response of 5-wide pseudomedian and median filters in edge jitter situations.....	20
Figure 4.1	Response of pseudomedian and median filters to triangle waves.....	26
Figure 4.2	Response of midrange and average filters to triangle waves.....	28
Figure 4.3	Response of pseudomedian and median filters to sinusoidal waves.....	29
Figure 4.4	Response of midrange and average filters to sinusoidal waves..	30
Figure 4.5	Response of pseudomedian and median filters to square waves.....	33
Figure 4.6	Response of midrange and average filters to square waves.....	34
Figure 4.7	Response curves for 75-wide median and pseudomedian filters acting on triangle waves.....	38
Figure 4.8	Response curves for 75-wide average and midrange filters acting on triangle waves.....	38
Figure 4.9	Response curves for 75-wide median and pseudomedian filters acting on sinusoidal waves.....	39

Figure 4.10	Response curves for 75-wide average and midrange filters acting on sinusoidal waves.....	39
Figure 4.11	Response curves for 75-wide median and pseudomedian filters acting on square waves.....	40
Figure 4.12	Response curves for 75-wide average and midrange filters acting on square waves.....	40
Figure 4.13	Response curves for 7-wide median and pseudomedian filters acting on triangle waves.....	41
Figure 4.14	Response curves for 7-wide average and midrange filters acting on triangle waves.....	41
Figure 4.15	Response curves for 7-wide median and pseudomedian filters acting on sinusoidal waves.....	42
Figure 4.16	Response curves for 7-wide average and midrange filters acting on sinusoidal waves.....	42
Figure 4.17	Response curves for 7-wide median and pseudomedian filters acting on square waves.....	43
Figure 4.18	Response curves for 7-wide average and midrange filters acting on square waves.....	43
Figure 4.19	Correlation curves for 75-wide median and pseudomedian filters acting on triangle waves.....	49
Figure 4.20	Correlation curves for 75-wide average and midrange filters acting on triangle waves.....	49
Figure 4.21	Correlation curves for 75-wide median and pseudomedian filters acting on sinusoidal waves.....	50
Figure 4.22	Correlation curves for 75-wide average and midrange filters acting on sinusoidal waves.....	50
Figure 4.23	Correlation curves for 75-wide median and pseudomedian filters acting on square waves.....	51
Figure 4.24	Correlation curves for 75-wide average and midrange filters acting on square waves.....	51
Figure 5.1	3x3 square two-dimensional filter window.....	53

Figure 5.2	135° corner filtered by 5x5 square median and pseudomedian filters.....	56
Figure 5.3	90° corner filtered by 5x5 square median and pseudomedian filters.....	57
Figure 5.4	45° corner filtered by 5x5 square median and pseudomedian filters.....	57
Figure 5.5	Pseudomedian- and median-filtered geometrical objects.....	67
Figure 5.6	Average- and midrange-filtered geometrical objects.....	67
Figure 5.7	Pseudomedian- and median-filtered geometrical objects corrupted with Gaussian noise.....	68
Figure 5.8	Average- and midrange-filtered geometrical objects corrupted with Gaussian noise.....	68
Figure 5.9	Segmented filtered images from Figures 5.7 and 5.8.....	69
Figure 5.10	Pseudomedian- and median-filtered geometrical objects corrupted with salt-and-pepper noise.....	69
Figure 5.11	Pseudomedian- and median-filtered image of tiger corrupted with Gaussian noise.....	70
Figure 5.12	Average- and midrange-filtered image of tiger corrupted with Gaussian noise.....	70
Figure 5.13	Pseudomedian- and median-filtered image of biker corrupted with Gaussian noise.....	71
Figure 5.14	Average- and midrange-filtered image of biker corrupted with Gaussian noise.....	71
Figure 5.15	Pseudomedian- and median-filtered image of feline corneal endothelial cells (window size 5x5).....	72
Figure 5.16	Pseudomedian- and median-filtered image of feline corneal endothelial cells (window size 9x9).....	72
Figure 5.17	Filtered thermograms of a circular blackbody with vertical linescan.....	73
Figure 5.18	Filtered thermograms of a class ring.....	73
Figure 5.19	Filtered thermograms of a class ring in pseudocolor.....	74
Figure 5.20	Filtered thermograms of the author's face in pseudocolor.....	74

Table of Nomenclature

$a(t_0)$	averaging filter output at time t_0 , continuous time
$A(f)$	response curve for average filter as a function of frequency
CORR	correlation
cy/w	cycles per window (unit of frequency)
DC	direct current (constant signal component)
f	frequency variable
GAIN	gain (of a filter)
i	arbitrary integer constant
j	arbitrary integer constant
k	arbitrary constant
m	arbitrary integer constant
max	maximum operator
min	minimum operator
$m(t_0)$	median filter output at time t_0 , continuous time
MSE	mean square error
$M(f)$	response curve for median filter as a function of frequency
n	arbitrary integer constant
N	arbitrary integer constant used to specify window size, discrete time overall window size (discrete time) is $2N+1$
p	arbitrary integer constant
$p(t_0)$	pseudomedian filter output at time t_0 , continuous time
PMED	pseudomedian operator
$P(f)$	response curve for pseudomedian filter as a function of frequency
q	arbitrary integer constant
$r(t_0)$	midrange filter output at time t_0 , continuous time
$R(f)$	response curve for midrange filter as a function of frequency
t_0	arbitrary constant (point in time under consideration), continuous time

T	period of input signal (periodic signals only)
w	window width, continuous time
W	window width, discrete time
$\{x_k\}$	input signal, discrete time
$x(t)$	input signal, continuous time
$\{y_k\}$	output signal, discrete time
$y(t)$	output signal, continuous time
	weighting factor for α -trimmed mean filters ($0 < \alpha < 1$)

Chapter 1

Introduction

Linear filtering is a well-established method for extracting signals from noisy environments. There are many tools available for analyzing linear filters, and the properties of these filters are well-understood. Nevertheless, in some situations a linear filter, even an optimal linear filter, does not perform adequately. In these cases, the special properties of certain nonlinear filters are required. Unfortunately, nonlinear filters are much more difficult to analyze than linear filters, and as a result the properties of many nonlinear filters are poorly understood.

Median filtering is one of the most common nonlinear techniques used in signal processing. Among the first to demonstrate the benefits of taking medians of sample data were Borda and Frost [31], and the median filter itself is generally ascribed to Tukey [29]. The properties of the median filter have been studied in quite some depth since these early uses [6, 7, 8, 9, 10]. The edge-preserving and impulse-removing properties of this filter are usually the most desirable features, and although the median filter is not conceptually complex, its computation can become quite cumbersome. The problem of efficient computation of the filter led Pratt, Cooper, and Kabir [1] to propose the "pseudomedian" filter, a filter with properties similar to the median filter, which can be more efficiently calculated. The pseudomedian filter often produces results that are very similar indeed to those produced by the median filter, and many of the theoretical properties of the pseudomedian are the same or nearly the same as the corresponding properties of the median filter. However, the response of the pseudomedian filter to high-frequency oscillations and to impulses is quite different from the response of the median filter. In many circumstances, the response of the pseudomedian filter is preferable.

Pseudomedian Filter Definition

The median filter is described as operating on a discrete signal. A window of width $2N+1$ sample points slides across the signal. The output of the filter is the median of the $2N+1$ values in the window, and this output is the filtered value at the sample in the center of the window.

The pseudomedian filter is also described for a discrete signal and a window width $2N+1$. However, the output of the pseudomedian operator, PMED, is the average of the maximum of the minima and the minimum of the maxima of the $N+1$ sliding subsequences of length $N+1$ in the window. This definition is illustrated by the equations below for $N=1$ and $N=2$.

$$\text{PMED}\{a,b,c\} = \frac{1}{2} \max\{(\min\{a,b\}, \min\{b,c\})\} + \frac{1}{2} \min\{(\max\{a,b\}, \max\{b,c\})\}$$

where the values in the window (width = $2N+1 = 3$) are $\{a,b,c\}$

$$\begin{aligned} \text{PMED}\{a,b,c,d,e\} = & \frac{1}{2} \max\{(\min\{a,b,c\}, \min\{b,c,d\}, \min\{c,d,e\})\} \\ & + \frac{1}{2} \min\{(\max\{a,b,c\}, \max\{b,c,d\}, \max\{c,d,e\})\} \end{aligned}$$

where the values in the window (width = $2N+1 = 5$) are $\{a,b,c,d,e\}$

Once again, the filtered output for the window is the filtered value at the sample in the center of the window.

The relationship of the definition of the pseudomedian to the median is illustrated by noting that the median can be defined as the maximum of the minima (or, equivalently, the minimum of the maxima) of all subsequences of length $N+1$ in the window. There are $\binom{2N+1}{N+1}$ such subsequences, and so the pseudomedian uses only a small subset of all possible subsequences in the window. It is worthwhile to notice that each of the subsequences used in the pseudomedian has the value at the center of the window as an element; this indicates that the pseudomedian will exhibit a more "center-weighted" response than the median.

Several observations about the pseudomedian filter can be made at this point. First, since the pseudomedian is defined as an average of two signal values, its output is not limited to values in the unfiltered signal. For example, an unfiltered signal consisting only of integer values does not necessarily result in a pseudomedian-filtered output of only integer values, as would be the case for the median-filtered output. This averaging method also indicates that the pseudomedian filter may have a "more linear" response than the median filter. Results that support this are given in later chapters. Second, the pseudomedian filter has many of the same general properties of the median filter: it exhibits a "low-pass" type of response in many cases, and preserves edges while reducing impulse noise. Finally, the pseudomedian filter can be implemented by an algorithm that is theoretically of lower order than most algorithms for the median filter. The fast median filtering algorithm developed by Huang, *et al.* [5] is a very fast implementation of the median filter, however, and a similar implementation of the pseudomedian filter would not be as efficient. Other algorithms for the median filter are either similar in speed or noticeably slower than most implementations of the pseudomedian filter.

Overview

This thesis develops a comprehensive system of results that will show how the response of the pseudomedian filter is similar to and differs from that of the median filter and some other linear and nonlinear filters, including the average and midrange filters. Among the results are a complete root signal analysis and a comparison of responses to edges and impulses. Also included are analogous continuous-time constructs for the filters, and a type of frequency response curves (developed using these constructs) for the filters acting on certain periodic signals. Distortion of these periodic signals induced by the filters is also examined using correlation measures. In addition, this thesis introduces a two-dimensional filter that is logically equivalent to the one-dimensional pseudomedian filter and develops some of its properties.

Chapter 2

Root Signal Analysis

This chapter develops some formal properties of the pseudomedian filter. Gallagher and Wise [2] and Tyan [3] first demonstrated the root signal set (that is, the set of signals unchanged by filtering) for the median filter. The root signal analysis that follows demonstrates the close relationship between the median and pseudomedian filters while pointing out important and useful differences.

Root Signal Theory

As previously stated, the pseudomedian filter operates on a discrete window in a sliding-window fashion, where the output at a particular point is the pseudomedian of the values in a $2N+1$ wide window centered at the point. That is,

$$\{y_k\} = \text{PMED} \{ x_{k-N}, \dots, x_k, \dots, x_{k+N} \}$$

The filter is assumed to be operating on a signal that extends to infinity in both directions, such that the filter window is always filled with signal values. However, the analysis presented below also holds true for extended signals. These signals are finite in length and have N constant points equal to the first value in the signal appended to the beginning of the signal, and N constant points equal to the last value appended to the end. (See also [2].)

The signal characteristics defined below create a precise vocabulary for the theorems presented in this chapter. The terms apply to a signal to be operated on by a filter of window width $2N+1$.

1. A *constant neighborhood* is an area of $N+1$ or more consecutive points that have the same value.
2. An *edge* is a monotonic sequence of points between constant neighborhoods such that the edge and constant neighborhoods combined are monotonic.

3. An *impulse* is an area of N or fewer points that is surrounded on either side by identically-valued constant neighborhoods.
4. An *oscillation* is any area that is not included in a constant neighborhood, edge, or impulse.
5. A *root signal* of a filter is a signal that is unchanged by that filter.

The first result below associates two of the above definitions with a single property.

Lemma 1: A signal of arbitrary length consists only of constant neighborhoods and edges if and only if each subsequence of length $N+2$ in the signal is monotonic.

Proof: Suppose a signal consists only of constant neighborhoods and edges. Since edges are by definition monotonic and are separated by constant neighborhoods, a nonmonotonic subsequence of the signal must contain a constant neighborhood. But since a constant neighborhood has at least $N+1$ consecutive equally valued points, any length $N+2$ subsequence including a constant neighborhood has at most only one point not equal to the constant neighborhood, and thus any such subsequence must be monotonic. Therefore every length $N+2$ subsequence of a signal consisting solely of constant neighborhoods and edges must be monotonic.

Now suppose that every length $N+2$ subsequence of a signal is monotonic. For a change in trend to occur (that is, from nondecreasing to nonincreasing or vice versa), there must be a subsequence of at least $N+1$ equally valued points between the sections of increase and decrease. Thus the areas of equal value between trend changes are constant neighborhoods, and the monotonic regions between these areas are edges. Therefore all signals where every length $N+2$ subsequence is monotonic consist only of constant neighborhoods and edges.

QED

Two simple observations concerning the pseudomedian filter may now be made. The value in the center of the filter window is important because

this value is in every subsequence taken while computing the pseudomedian. Since this is so, no subsequence can have its minimum be greater than this value, and none can have its maximum be less than this value. Therefore the maximum of the minima is restricted to values greater than or equal to the value in the center of the window, and likewise the minimum of the maxima must be less than or equal to the value in the center of the window.

Observation 1: The minimum of the maxima of the length $N+1$ subsequences taken for a length $2N+1$ pseudomedian is greater than or equal to the value of the point at the center of the window.

Observation 2: The maximum of the minima of the length $N+1$ subsequences taken for a length $2N+1$ pseudomedian is less than or equal to the value of the point at the center of the window.

The pseudomedian is thus the average of two values, one of which is greater than or equal to the value in the center of the filter window, and the other of which is less than or equal to the value in the center of the window.

The following two properties are important features of the pseudomedian filter and are shared with the median filter. They follow directly from the above observations.

Property 1: The pseudomedian of a monotonic sequence of length $2N+1$ is the value in the center of the sequence. That is, if $x_{-N} \dots x_0 \dots x_N$, then

$$\text{PMED} \{x_{-N}, \dots, x_0, \dots, x_N\} = x_0.$$

Proof: The maximum of the first subsequence in the window is x_0 , which is the value in the center of the window. Since at least one subsequence has x_0 as its maximum, by Observation 1 the minimum of the maxima is x_0 .

The minimum of the last subsequence in the window is x_0 , which is the value in the center of the window. Since at least one subsequence has x_0 as its minimum, by Observation 2 the maximum of the minima is x_0 .

Therefore, $\text{PMED} \{x_{-N}, \dots, x_0, \dots, x_N\} = x_0$.

QED

Property 2: The pseudomedian of a sequence of length $2N+1$ that contains a constant neighborhood is the value of the constant neighborhood.

Proof: Since a constant neighborhood is an area of at least $N+1$ consecutive equally valued points, then at least one of the length $N+1$ subsequences taken while computing the pseudomedian has its maximum and its minimum equal to the value of the constant neighborhood. By Observation 1, the minimum of the maxima is the value of the center point, and by Observation 2, the maximum of the minima is the value of the center point. The center point must be within the constant neighborhood, because the constant neighborhood is more than half the length of the window. Therefore, the pseudomedian of a length $2N+1$ sequence containing a constant neighborhood is the value of the constant neighborhood.

QED

The following lemma establishes a sufficient condition for a signal to be a root of the pseudomedian filter.

Lemma 2: A signal consisting only of constant neighborhoods and edges is a root of a pseudomedian filter of window width $2N+1$.

Proof: For a signal to be a root, the pseudomedian of each window must be the value in the center of the window. For a given signal, there are two possible cases for each $2N+1$ length window.

Case I: The points in the window are monotonic.

By Property 1, the pseudomedian is the value in the center of window.

Case II: The points in the window are nonmonotonic.

Since edges are monotonic and are separated by constant neighborhoods, a nonmonotonic window must contain a constant neighborhood. By Property 2, the pseudomedian of such a window is the value of the constant neighborhood. The

center point of the window must be within the constant neighborhood, and therefore the pseudomedian of the window is equal to the value in the center of the window.

Thus the pseudomedian of each $2N+1$ window is the value in the center of the window, and therefore a signal consisting only of constant neighborhoods and edges is a root signal.

QED

The next lemma is a more restrictive version of Lemma 6.A.1 proved in Tyan [3] for the median filter. The additional restrictions enforced by the pseudomedian filter here are the basis for proving that the root signals from Lemma 2 are in fact *necessary* and sufficient conditions, whereas there exists an additional class of infinite-length root signals for the median filter.

Lemma 3: Let $n < m$ and let $x_n < x_i < x_m$ for all $n < i < m$. If $\text{PMED}\{x_{n-p}, \dots, x_{n+p}\} = x_n$ and if $\text{PMED}\{x_{m-q}, \dots, x_{m+q}\} = x_m$, where $n-p < m-q$ and $n+p < m+q$, then $x_j \leq x_n$ for all $n-p < j < n$ and $x_j \leq x_m$ for all $m < j < m+q$.

Recast in less symbolic language for $p=q=N$, Lemma 3 states that if a signal is strictly increasing from point n to point m , and $\text{PMED}\{x_{n-N}, \dots, x_{n+N}\} = x_n$ and $\text{PMED}\{x_{m-N}, \dots, x_{m+N}\} = x_m$, then the signal is monotonically nondecreasing from point $n-N$ to point $m+N$.

Proof of Lemma 3: See Figure 2.1 below for a diagram of the signal segment under consideration.




Figure 2.1. Diagram of signal segment considered for Lemma 3.

Consider first the window from $n-p$ to $n+p$. The proof will be given for the case where the maximum of the minima and minimum of the maxima of the length $p+1$ subsequences must be x_n . (The value of x_n may be achieved by averaging two unequal values when computing the pseudomedian, but signals that achieve this result do not fit the hypothesis of the theorem. The proof for this case is omitted.) Since x_n is the center point of the window, if there is at least one length $p+1$ subsequence with x_n as its maximum, by Observation 1 the minimum of the maxima will be x_n . For this to be true, the first subsequence in the window must have x_n as its maximum, since all other subsequences include $x_{n+1} > x_n$. Thus $x_j \leq x_n$ for all $n-p \leq j < n$.

Now consider the window from $m-q$ to $m+q$. Again, the minimum of the maxima and the maximum of the minima must both be x_m . Since x_m is the center point of the window, if there is at least one $q+1$ subsequence with x_m as its minimum, by Observation 2 the maximum of the minima will be x_m . All subsequences in the window except the rightmost one include $x_{m-1} < x_m$, so the rightmost subsequence must have x_m as its minimum. Therefore, $x_j \geq x_m$ for all $m < j \leq m+q$.

QED

The last two lemmas needed to demonstrate the necessary conditions for root signals also derive from Tyan's work on median filters [3]. Lemma 4 is a restatement of Tyan's Theorem 6.2 (with virtually identical proof) and Lemma 5 is related to Tyan's Theorem 6.3.

Lemma 4: If there is at least one monotonic subsequence of length $N+1$ in a root signal of a pseudomedian filter of length $2N+1$, then every length $N+2$ subsequence of the root signal is monotonic.

Proof: Follows from Lemma 3. Proof for median filter given in Tyan [3] applies, with Lemma 3 substituted for the corresponding median filter lemma.

Lemma 5: Every root signal of a length $2N+1$ pseudomedian filter has at least one subsequence of length $N+1$ that is monotonic.

Proof: Consider a root signal $\{x_k\}$ of a length $2N+1$ pseudomedian filter. Assume that this root does not have any monotonic subsequences of length $N+1$. Consider a particular transition to a higher value at, say, $k=0$, so $x_0 < x_1$. Then Lemma 3 requires that $x_i \leq x_0$ for each i , where $-N \leq i < 0$. Thus $x_{-N} \leq \dots \leq x_1 \leq x_0$ is a monotonic subsequence of the length $N+1$, which contradicts the assumption that the root signal $\{x_k\}$ does not have any monotonic subsequences of length $N+1$. Therefore, every root signal has at least one subsequence of length $N+1$ that is monotonic.

QED

The five lemmas proved above combine to show a necessary and sufficient condition for root signals (finite/extended or infinite) for a length $2N+1$ pseudomedian filter.

Theorem 1: A necessary and sufficient condition for a signal to be invariant under pseudomedian filtering is that the signal consist only of constant neighborhoods and edges.

Proof: Any signal, and thus any potential root signal of a length $2N+1$ pseudomedian filter, can be ascribed to one of the following cases:

Case I: At least one length $N+1$ subsequence of the signal is monotonic;
or

Case II: The signal has no monotonic subsequences of length $N+1$.

By Lemma 4, all root signals belonging to Case I are everywhere monotonic of length $N+2$; that is, every length $N+2$ subsequence of a root signal satisfying Case I is monotonic. By Lemma 1, all such signals consist only of constant neighborhoods and edges.

By Lemma 5, there are no root signals that satisfy Case II.

All (root and non-root) signals must satisfy one of the above cases, but all root signals of the pseudomedian filter must satisfy Case I, and furthermore must only consist of constant neighborhoods and edges. Therefore,

all root signals of the pseudomedian filter consist solely of constant neighborhoods and edges. By Lemma 2, all signals consisting of only constant neighborhoods and edges are root signals of the pseudomedian filter.

QED

Theorem 1 is a powerful result that holds both for finite-length signals that are extended as described earlier and for "infinite-length" signals. It is the same result as the necessary and sufficient condition for extended finite-length root signals of the median filter shown in Gallagher and Wise [2]. This leads immediately to the results below:

Corollary 1: An extended finite-length root signal of a median filter of length $2N+1$ is a root signal of a pseudomedian filter of length $2N+1$.

Corollary 2: A root signal of a pseudomedian filter of length $2N+1$ is a root signal of a median filter of length $2N+1$.

Recall that an extended finite-length signal has N constant points equal to the first value in the signal appended to the beginning of the signal, and N constant points equal to the last value appended to the end. The root signal set of the pseudomedian filter is thus identical to that of the median filter for finite-length signals. However, if the set of roots under consideration is extended to include signals that extend to infinity in either direction, another type of root signal of the median filter emerges. These "Type II" or "fast-fluctuating" roots are necessarily bi-valued; that is, they take on only two values, and do not contain any constant neighborhoods [3]. There is also a related set of signals that may be termed "oscillatory" roots of the median filter; these signals are not invariant to each pass of a median filter, but the original signal will recur after two or more passes of the filter. The pseudomedian filter does not have any fast-fluctuating or oscillatory roots.

In practice, all signals are finite in length, but Type II and oscillatory median filter roots are still important. Even a small section of such a root within a finite signal creates unattenuated oscillations in the output of the median filter; only at the endpoints of the section does the filter cause any change. The pseudomedian filter, however, usually reduces such bi-valued fast-fluctuating

sections to a constant DC level in one pass by averaging the two values in the section throughout.

Convergence to Root Signals

Another topic of interest in root signal analysis is convergence of a repeatedly filtered signal to a root signal. For the median filter, Gallagher and Wise [2] have shown that a signal of length L will become a root signal after at most $(L-2)/2$ successive passes of a median filter. No such upper limit on the number of passes required to yield a root signal can be derived for the pseudomedian filter. Indeed, some signals (for example, any impulse) will never in theory be reduced to a root signal by successive passes of a pseudomedian filter. However, each successive pass results in a signal that is "closer" to a root signal than the previous pass, and as the number of passes becomes very large, the resulting signal approaches a root signal. These observations indicate that repeated pseudomedian filtering will result in either:

- i. a root signal after a finite number of passes; or
- ii. a sequence of signals that converges to a root signal.

Note that the resulting root signal in either case often is not identical to the corresponding median filter root signal.

The root signal analysis in this chapter demonstrates some provocative similarities and differences between the pseudomedian and median filters. Although the root signal sets of the two filters are identical when restricted to finite-length signals, the root signals to which the filters converge for a given input signal are not the same. The reasons for these similarities and differences are explained further in the succeeding chapters.

Chapter 3

Response to Edges and Impulses

Root signal analysis points out some of the most useful similarities and differences between the pseudomedian and median filters. However, the root signal set of a filter is not a complete description of the behavior of a filter. For example, the average filter, which has its output defined as the average of all values in its sliding window, has a root signal set (for infinite-length signals) consisting only of straight lines (constant values and constant-slope lines). The midrange filter, with output equal to the average of the maximum value and minimum value in the sliding window, has the same root signal set as the average filter. But in many cases, the output of the midrange filter resembles the output of the pseudomedian more closely than that of either the median or average filter. Clearly, additional analysis is needed to improve understanding of the pseudomedian filter.

One of the most useful characteristics of the median filter is its ability to completely remove impulses in one pass [2]. (A precise definition for the term "impulse" for the median and pseudomedian filter is given in Chapter 2.) In applications where impulse noise is a major component, such as thermographic imaging, the median filter and modifications of it are used widely. But, despite the fact that the root signal set of the pseudomedian filter is a subset of that of the median filter, the pseudomedian filter does not completely remove impulses.

Consider an impulse of one point, for instance $0, \dots, 0, 1, 0, \dots, 0$. Whenever the pseudomedian filter window of length $2N+1$ is not centered on the impulse, there are $N+1$ consecutive zeroes in the window, and by Property 2, the output is zero. But when the window is centered on the impulse, Property 2 does not apply. Observation 1 requires that the minimum of the maxima be one, and it is clear that the maximum of the minima is zero. The pseudomedian is the average of these two values, $1/2$. The impulse thus

remains in the filtered signal, but with its amplitude attenuated by one-half. See Figure 3.1.




Figure 3.1. One-point impulse filtered by 5-wide pseudomedian, median, average, and midrange filters.

For more generic impulses, the pseudomedian filter generally passes the impulse at one-half amplitude. An illustration is shown in Figure 3.2. However, the pseudomedian filter may change the shape as well as amplitude of impulses that are at least three points wide. An example of this is given in Figure 3.3. Although the specific shape of the impulse is slightly modified by the pseudomedian filter, the general shape is passed at one-half amplitude.




Figure 3.2. Two-point impulse filtered by 5-wide pseudomedian, median, average, and midrange filters.




Figure 3.3. Three-point impulse filtered by 7-wide pseudomedian, median, average, and midrange filters.

Figures 3.1-3.3 also illustrate the response of the average and midrange filters to impulses. Note that the response of these filters differs radically from the median and pseudomedian because the output at every point within a half window width in both directions from the impulse is influenced by the impulse. The pseudomedian and median filters do not change any values in the constant neighborhoods surrounding an impulse; in general, the average and midrange filters do. One notable exception is for zero-average impulses (that is, where the sum of the deviations from the constant neighborhood is zero). In this case, the average and midrange filters only give deviations from the constant neighborhood at points roughly one-half window width from the impulse. This can be seen from Figure 3.2 by noting that a similar zero-average impulse would cause the flat areas in the filtered signals to be all the same value. The median and pseudomedian filters behave the same for zero-power impulses

as they do for other impulse types. Also, the responses of the average and midrange filters depend strongly upon the filter window size, whereas the responses of the median and pseudomedian filters do not vary as window size changes, so long as the window is more than twice as wide as the impulse.

Another interesting type of impulse response to consider is a spike on a constant-slope ramp signal. Although such a spike is not an "impulse" as defined in Chapter 2, the response of the filters in this situation is important. Figure 3.4 gives an example.

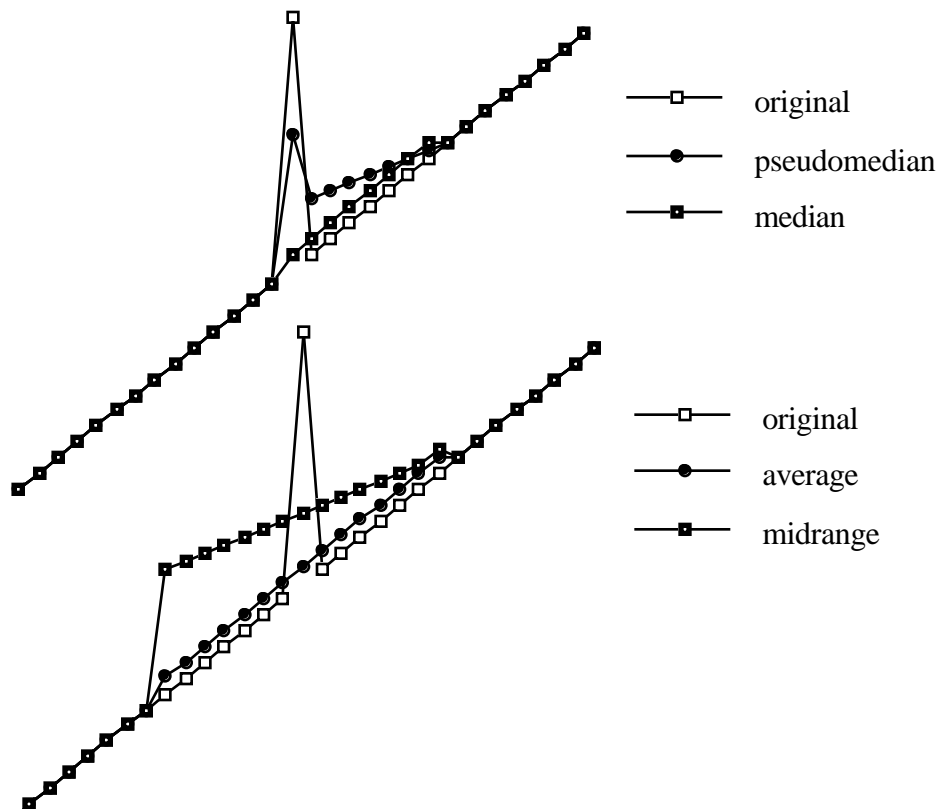


Figure 3.4. Response of 15-wide filters to spike noise on a ramp signal.

The above example indicates some significant differences between the filters. The median filter only exhibits a small bias for values on one side of the spike, while removing the spike completely. The pseudomedian filter attenuates the spike by about one-half, and has a bias on one side of the spike that is initially quite large, then slowly decreases. The midrange filter exhibits a bias not unlike the pseudomedian filter, that is initially large and decreases. However, the bias is on both sides of the spike, and is much larger at the start. The average filter, as expected, exhibits only a very slight bias on both sides of the spike.

The response of the median filter to edges is one of its most celebrated aspects. The root signal analysis of the median filter proves that it preserves edges exactly in signals free from noise or other edge interference. The pseudomedian filter is likewise an edge-preserving filter, and thus resembles the median filter in a way that many other related filters do not. Figure 3.5 demonstrates the response of the median, pseudomedian, average, and midrange filters to a sawtooth wave. The average and midrange filters introduce severe distortions in this example. Note that the pseudomedian filter output has lower mean square error and maximum amplitude error than the median filter output, since every point in the pseudomedian-filtered signal is the same as in the original or closer to the original than the corresponding median-filtered point. However, the median and pseudomedian filter do an equally good job of preserving edge location.




Figure 3.5. Sawtooth wave filtered by 5-wide pseudomedian, median, average, and midrange filters.

The alpha-trimmed mean filters, as developed by Bednar and Watt [13] and by Restrepo and Bovik [14], yield results that lie between the median and average filter response, or between average and midrange filter response. The α -inner mean filter [14] results in outputs that are normalized linear combinations of the median and average filters, where α is a weighting factor such that $\alpha=1$ corresponds to the average filter and $\alpha=1/(2N+1)$ corresponds to the median filter. The α -outer mean filter [14] gives normalized linear combinations of the average and midrange filters, where α weights the combination so that $\alpha=1$ is the average filter and $\alpha=0$ is the midrange filter.

The edge response of the median filter to signals with noise is of concern in many applications. Impulse-like noise near edges can cause the median filter output to have the edges offset in either direction by up to several pixels. This phenomenon is called "edge jitter" [6]. In applications which

require precise feature placement, edge jitter can be a major problem. While the pseudomedian filter does not eliminate the signal defects that lead to edge jitter, its edge placement in these situations is often more accurate.

The differences between the response of the median and pseudomedian filters to edge jitter situations can be illustrated by a simple example (see Figure 3.6). Edge jitter is often caused by an impulse that occurs within one-half window width of an edge. This impulse causes the values in the window to be skewed in the direction of the impulse when placed in order. Thus the median filter will indicate that the feature extends further toward this impulse than it actually does because of the skewed values. The pseudomedian filter is influenced more by the value in the center of the window than the median filter is; this is clear from the comparison of the filters' responses to impulses. For a pseudomedian-filtered signal, an impulse near an edge usually causes only a small "foothill" or "valley" near the edge, and the edge will appear in the correct location, although often reduced in amplitude.

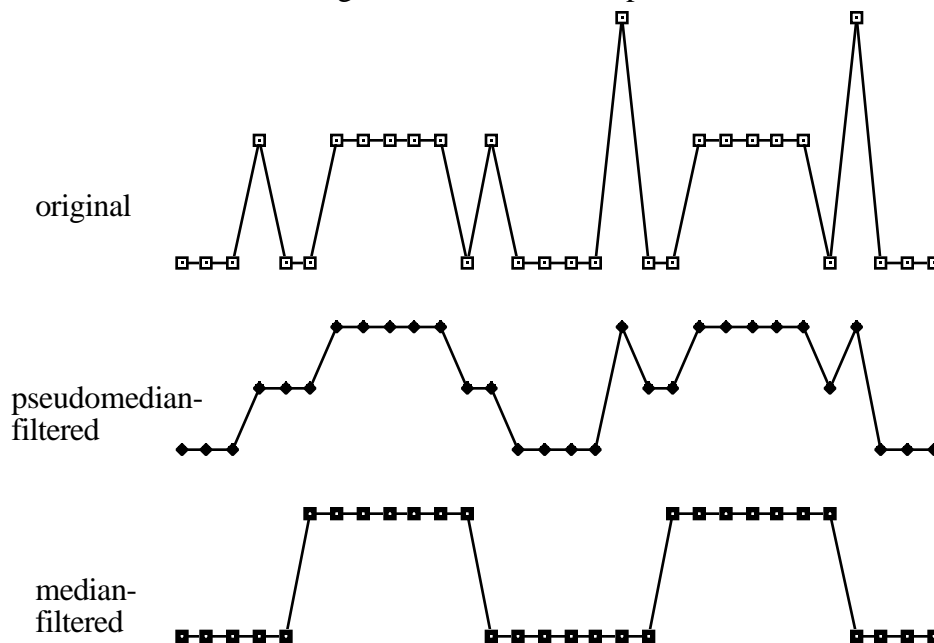


Figure 3.6. Response of 5-wide pseudomedian and median filters in edge jitter situations.

Figure 3.6 also demonstrates that the pseudomedian filter output for signals consisting mostly of impulses and edges bears more resemblance to the original signal than the median filter output does. More precisely, the mean square error and maximum absolute error between the filtered and original signal are both smaller for the pseudomedian filter than for the median filter in most impulse and edge situations like Figure 3.6. (Note that lower errors in this situation are not necessarily desirable if the impulses are considered noise.) These properties of the pseudomedian filter may be advantages in some applications, but are likely disadvantages in others.

While the average and midrange filters do not appear to have much in common with the median and pseudomedian filters in their response to impulses and edges, the similarities between these filters will be clarified by examining their response to periodic signals. This is the subject of Chapter 4.

Chapter 4

Continuous-Time Analysis

This chapter introduces another method useful for understanding the differences between the pseudomedian and median filters. Although there are a number of complications which limit the applications of this analysis, the basic ideas that it presents are very helpful. In general, this analysis develops continuous-time constructs analogous to the filters that are defined in discrete time. The peak responses of these constructs to periodic signals are easily obtained, and from these a type of "response curve" for the filters in continuous time can be derived. Also, a construct based upon statistical correlation is used to approximate of the amount of distortion introduced by the filters.

Continuous-Time Filter Analogues

A continuous-time analysis of discrete filters clearly requires many assumptions that greatly influence the validity of the results. Despite the tenuous connection between this development and the true operation of the filters, much insight can be gained from these results. Before beginning the analysis, several of the most important assumptions must be clarified.

First, the response curves derived below do not show the response of the filters to any combination of frequencies; rather, they provide only the response to a certain periodic signal at a single frequency. The response curves are derived by noting the response of the filter to a single-frequency periodic signal at a given frequency, and deducing how this response varies as the frequency of the signal is changed. The superposition principle that may be applied for linear filters does not apply to the nonlinear filters studied here.

Second, the term "response" is used in a relaxed sense here, as it means only the amplitude and phase of the output wave relative to the input wave without regard to any distortion of the wave caused by the filter. The

values presented as "amplitude response" indicate only the ratio of the peak values of the filtered signal to the peak values of the unfiltered signal. In some cases, especially for "amplitude responses" near zero and for the median filter, the shape of the signal is quite distorted. However, in this study the filtered signals always have a base period identical to the unfiltered signals. Acceptance of this distortion dilutes the impact of the results slightly, but a brief study of the distortion characteristics at the end of this chapter helps clarify this problem. Also, the "phase response" indicates only the most basic relationship of the filter output to the original signal. In this study, the results for "phase response" were always either 0 or 180° , so I have simplified the response curves by indicating regions where the output signal is 180° out of phase with negative values for "amplitude response." For the signals presented, this negative "amplitude response" is the practical effect of this phase error.

Third, each filter has a different response curve for every different type of periodic signal. Because the filters are nonlinear, the response of a filter to one type of periodic signal is different from its response to another type of periodic signal. However, the results given in this chapter indicate certain similarities that reveal important characteristics of each filter.

Finally, the actual response of the filters to a signal sometimes deviates significantly from what the continuous-time response curves predict. Ultimately, this is due to the assumptions inherent in generalizing the discrete-time filters to continuous-time and in interpreting the continuous-time results in discrete time. The more immediate reasons for these differences from the predicted response are sampling effects and errors introduced by the discrete windows (which often are only a few samples wide). Although these problems are important, the results from the continuous-time analysis provide valuable insights to the behavior of the median and pseudomedian filter.

A generalization of the median filter to continuous time begins the response analysis. The median of a digital sequence is the value in the sequence where one-half of the remaining values are greater than or equal to the value and one-half are less than or equal to the value. To generalize this definition, first

consider a continuous-time definition of the mean of a function over a finite window:

$$a(t_0) = \frac{1}{w} \int_{t_0-w/2}^{t_0+w/2} f(t) dt$$

where $a(t_0)$ = mean of $f(t)$ over $t_0-w/2 \leq t \leq t_0+w/2$

[$a(t_0)$ = average-filtered value at $t=t_0$ for window length w]

The above expression is used as the continuous-time definition of the average filter. An alternate definition of the average that will help to generalize the median is given below.

$$\int_{t_0-w/2}^{t_0+w/2} [f(t) - a(t_0)] dt = 0$$

The above equation requires only a simple modification to express the balance of the *number* of values that is achieved by the median, instead of the balance of the *area* that defines the mean. The signum function, $\text{sgn}(x)$, is applied to the expression under the integral. This function returns a value of +1 for arguments greater than zero, 0 for arguments equal to zero, and -1 for arguments less than zero. The definition is then:

$$\int_{t_0-w/2}^{t_0+w/2} \text{sgn} \{ f(t) - m(t_0) \} dt = 0$$

where $m(t_0)$ = median of $f(t)$ over $t_0-w/2 \leq t \leq t_0+w/2$

[$m(t_0)$ = median-filtered value at $t=t_0$ for window length w]

This definition is only valid, however, for functions $f(t)$ that only take on the median value (m) a finite number of times over the the window $t_0-w/2 \leq t \leq t_0+w/2$. When a function takes on the median value an infinite number of times over the window, there sometimes exists no value for m such that the equality holds. But, at least one value m_0 will exist where the evaluated integral is greater than zero for $m < m_0$ and the evaluated integral is less than zero for $m > m_0$. The median of $f(t)$ over the window is then the value of m_0 where

the absolute value of the evaluated integral is minimized. This is shown in the expression below.

$m(t_0)$ = the value of x that minimizes the following expression:

$$\int_{t_0-w/2}^{t_0+w/2} \text{sgn} \{ f(t) - x \} dt$$

Note for certain functions $f(t)$ and certain values of t_0 and w , it is possible to obtain more than one value for x that minimizes the above expression. These situations are akin to having an even-sized window in the discrete case, so that half of the values in the window are above where the median should be, and half are below. In the discrete case, no value in the window can be singled out as the median, and likewise in the continuous-time expression the median is indeterminate. One convention used in these situations is to average the two values on either side of the supposed "median" and consider that value as the median. This convention works in both the continuous and discrete time definitions.

It is considerably easier to derive a continuous-time expression for the pseudomedian filter. By extending the discrete definition to arbitrary length, and then converting the index variables and maximum and minimum operators to continuous-time, the following expression results:

$$p(t_0) = 1/2 \cdot \max (\{ \min (f([t_0+\beta-w/2, t_0+\beta])), 0 \} \beta \ w/2) \\ + 1/2 \cdot \min (\{ \max (f([t_0+\beta-w/2, t_0+\beta])), 0 \} \beta \ w/2)$$

where $p(t_0)$ = pseudomedian-filtered value at $t=t_0$ for window width w

β = index that takes on all values from 0 to $w/2$

The final filter considered in this continuous-time study is the midrange filter, which has as its output the average of the maximum and minimum value within the filter window. The expression for this filter in continuous time is:

$$r(t_0) = 1/2 \cdot \max \{ f([t_0-w/2, t_0+w/2]) \} + 1/2 \cdot \min \{ f([t_0-w/2, t_0+w/2]) \}$$

where $r(t_0)$ = midrange-filtered value at $t=t_0$ for window width w

Response Curves

Using these continuous-time representations for the median, pseudomedian, average, and midrange filters, one may derive the output of these filters for various input signals. Of particular interest as input are three common periodic signals: the sine wave, the triangle wave, and the square wave. The simplicity and symmetry of these waves allow development of analytic solutions for the "response" of the filters to these signals. The most convenient way to express the frequency of the waves is in cycles per window length, abbreviated cy/w in this thesis.

The triangle wave is one of the easiest periodic functions for which to determine the response of the median and pseudomedian filters. The amplitude of the filter output at peaks of the wave for waves of various frequencies is obtained, and from this a value for the filter response at that frequency is assigned. The response curves of the pseudomedian and median filters to triangle waves are shown in Figure 4.1.

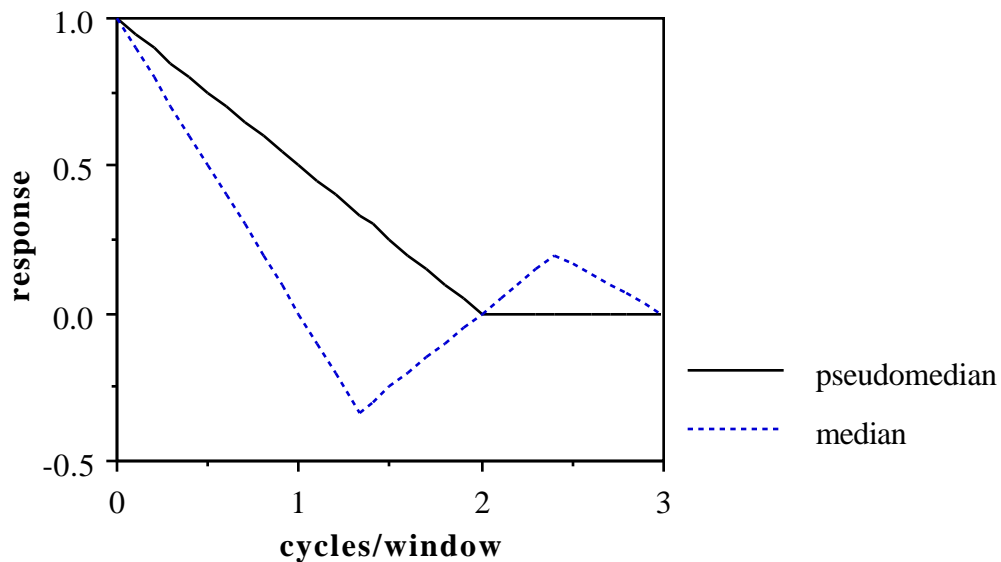


Figure 4.1. Response of pseudomedian and median filters to triangle waves.

The response of the pseudomedian filter to triangle waves can be written as follows:

$$P(f) = \begin{cases} 1 - f/2, & 0 \leq f \leq 2 \\ 0, & f > 2 \end{cases}$$

where f = frequency of triangle wave in cy/w

$P(f)$ = ratio of peaks of filtered to unfiltered signal at frequency f

The response of the median filter to triangle waves is quite different from that of the pseudomedian filter. An expression for this response curve is:

$$M(f) = \begin{cases} 1 - f, & 0 \leq f \leq 1 \frac{1}{3} \\ f/2 - 1, & 1 \frac{1}{3} < f \leq 2 \frac{2}{5} \\ 1 - f/3, & 2 \frac{2}{5} < f \leq 3 \frac{3}{7} \\ \dots & \text{etc.} \end{cases}$$

Or, more formally,

$$M(f) = \begin{cases} 1 - \frac{f}{2n+1}, & 2n + \frac{2n}{4n+1} \leq f \leq 2n+1 + \frac{2n+1}{4n+3} \\ \frac{f}{2n+2} - 1, & 2n+1 + \frac{2n+1}{4n+3} < f \leq 2n+2 + \frac{2n+2}{4n+5} \end{cases}$$

where $n = 0, 1, 2, 3, \dots$

f = frequency of triangle wave in cy/w

$M(f)$ = ratio of peaks of filtered to unfiltered signal at f

An instructive comparison for these response curves is the similar response curves that can be developed for the average and midrange filters. These curves are shown in Figure 4.2.




Figure 4.2. Response of midrange and average filters to triangle waves.

A comparison of Figure 4.2 with Figure 4.1 indicates that, for triangle waves, the responses of the median and average filters are quite similar. Also, the responses of the pseudomedian and midrange filters are related very closely; in fact, the output of a pseudomedian filter is identical (in the continuous-time analysis) to that of a midrange filter with a window one-half as wide as the pseudomedian filter window width. The response of the midrange filter is:

$$R(f) = \begin{cases} 1 - f, & 0 \leq f \leq 1 \\ 0, & f > 1 \end{cases}$$

where f = frequency of triangle wave in cy/w

$R(f)$ = ratio of peaks of filtered to unfiltered signal at f

The response of the average filter to triangle waves is as follows:

$$A(f) = (-1)^n [-f + (2n + 1) - n(n+1)/f], \quad n \leq f < n+1$$

where $n = 0, 1, 2, 3, \dots$

f = frequency of triangle wave in cy/w

$A(f)$ = ratio of peaks of filtered to unfiltered signal at f

Since sinusoidal waves are quite similar in many respects to triangle waves, response curves for the filters to sinusoidal waves may be expected to resemble closely those for triangle waves. Figure 4.3 shows the response curves of the pseudomedian and median filters to sinusoidal waves.




Figure 4.3. Response of pseudomedian and median filters to sinusoidal waves.

The analytic expressions for these response curves are given below.

For the pseudomedian filter:

$$P(f) = \begin{cases} 1/2 \cdot (1 + \cos f/2), & 0 \leq f \leq 2 \\ 0, & f > 2 \end{cases}$$

where f = frequency of sinusoidal wave in cy/w

$P(f)$ = ratio of peaks of filtered to unfiltered signal at frequency f

For the median filter:

$$M(f) = \begin{cases} \cos(f/2), & 0 \leq f < 1\ 1/3 \\ -\cos(f/4), & 1\ 1/3 < f < 2\ 2/5 \\ \cos(f/6), & 2\ 2/5 < f < 3\ 3/7 \\ \dots & \text{etc.} \end{cases}$$

Or, more formally,

$$M(f) = \begin{cases} \cos \frac{f}{4n+2}, & 2n + \frac{2n}{4n+1} < f < 2n + 1 + \frac{2n+1}{4n+3} \\ -\cos \frac{f}{4n+4}, & 2n + 1 + \frac{2n+1}{4n+3} < f < 2n + 2 + \frac{2n+2}{4n+5} \end{cases}$$

where $n = 0, 1, 2, 3, \dots$

f = frequency of sinusoidal wave in cy/w

$M(f)$ = ratio of peaks of filtered to unfiltered signal at frequency f

Although the analytic expressions given above for the responses to sinusoidal waves appear quite different from those given for triangle waves, Figures 4.1 and 4.3 show that these response curves are remarkably similar, as might be expected. To complete this comparison, response curves of the average and midrange filters are shown in Figure 4.4 and analytic expressions are listed below.

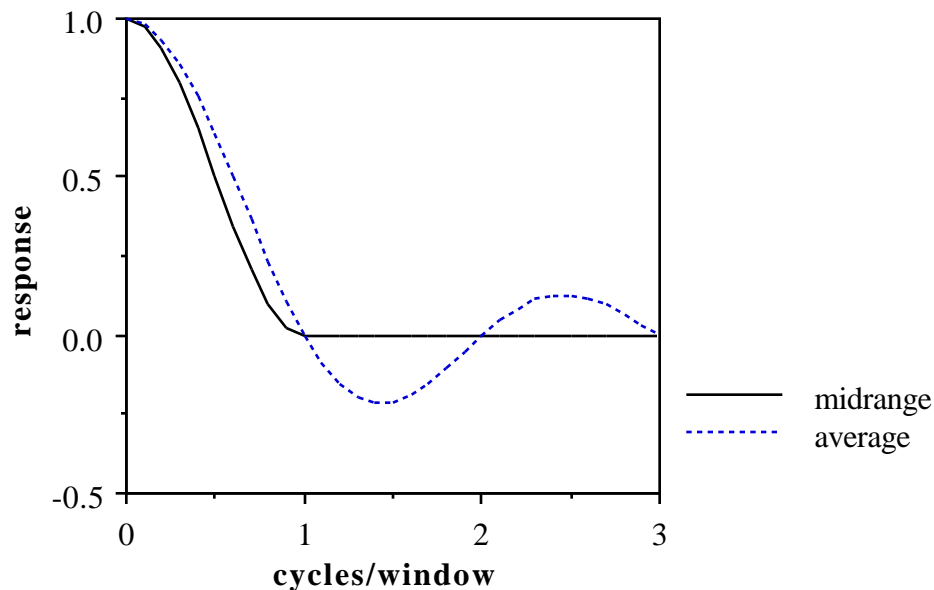


Figure 4.4. Response of midrange and average filters to sinusoidal waves.

For the midrange filter:

$$R(f) = \begin{cases} 1/2 \cdot (1 + \cos f), & 0 \leq f \leq 1 \\ 0, & f > 1 \end{cases}$$

where f = frequency of sinusoidal wave in cy/w

$R(f)$ = ratio of peaks of filtered to unfiltered signal at frequency f

For the average filter:

$$A(f) = \text{sinc } f, \quad f \geq 0$$

where f = frequency of sinusoidal wave in cy/w

$A(f)$ = ratio of peaks of filtered to unfiltered signal at frequency f

The results given above for the response of the median and pseudomedian filters to triangular and sinusoidal waves show some very interesting and significant differences between the two filters. The pseudomedian filter has a uniformly decreasing response from unity to zero as the frequency increases from zero to two cycles per window. At frequencies above two cycles per window, the pseudomedian filter will (in this continuous-time theory) always give a zero output for triangular and sinusoidal waves. In fact, for any vertically symmetrical periodic signal, the pseudomedian filter will have a constant output equal to the point of symmetry for frequencies at or above two cycles per window. The midrange filter behaves similarly, except that it has a constant output for frequencies above one cycle per window. This is due to its relationship to the pseudomedian filter described earlier for the triangle wave responses where midrange-filtered output is identical to pseudomedian-filtered output of twice the window width. This relationship holds for periodic signals that are both vertically and horizontally symmetric, as are triangular and sinusoidal waves.

The median filter, however, exhibits a much different response to these periodic signals. Its amplitude response to triangular and sinusoidal waves decreases much more rapidly than the response of the pseudomedian, dropping to zero at one cycle per window. For input signal frequencies between one and two cycles per window, the median filter has a negative amplitude response, indicating the output signal is 180° out of phase from the

original signal with a magnitude equal to the absolute value of the indicated response. This amplitude increases up to $1 \frac{1}{3}$ cycles per window, then decreases to zero at two cycles per window, above which the response once again has zero phase. Note that the response of the median filter to these periodic signals never goes to zero for all frequencies above a certain value; rather, it oscillates with decreasing amplitude. In theory, the response of the median filter is precisely zero only at frequencies of exactly an integral number of cycles per window.

The response of the median filter to triangular and sinusoidal waves is not unlike the response of the average filter. The close correspondence between the median and average filters in this case is understandable, since both filters are basically low-pass and these examples do not exploit the impulse-suppressing or edge-preserving features of the median filter. Nevertheless, it is instructive to note that the pseudomedian filter has a much different response to these signals than either the median or the average filter has.

Another periodic signal for which it is easy to analyze the effects of the filters is the square wave. Since there are only two levels (or values) in this signal, the median filter output may have only two output levels, while the pseudomedian and midrange filters may have only three. The average filter, however, is not limited like the others. Equations for the filter responses to square waves are given below, and Figures 4.5 and 4.6 illustrate these curves.

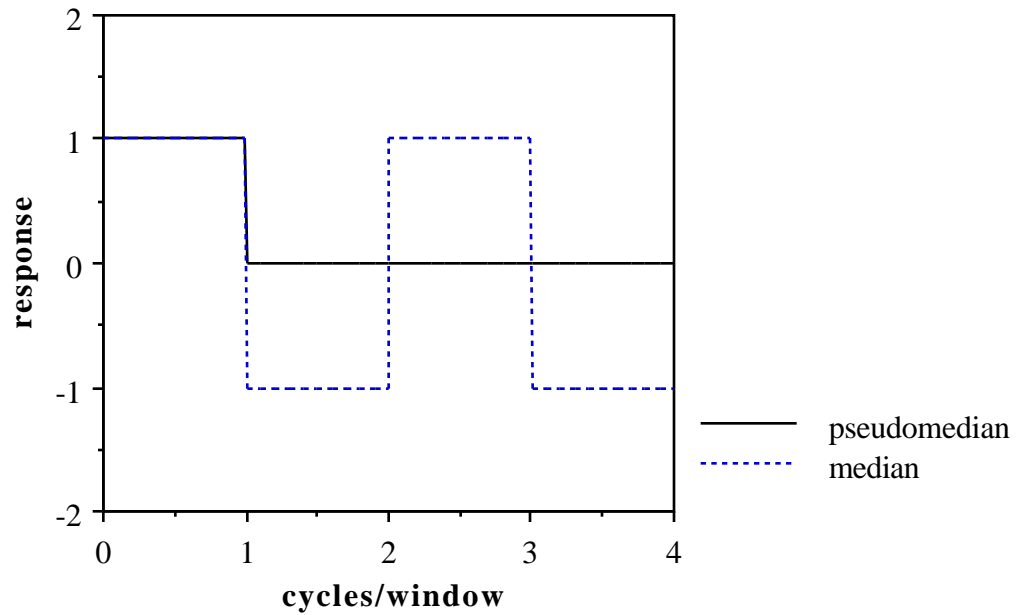


Figure 4.5. Response of pseudomedian and median filters to square waves.

For the pseudomedian filter:

$$P(f) = \begin{cases} 1, & 0 \leq f < 1 \\ 0, & f \geq 1 \end{cases}$$

where f = frequency of square wave in cy/w

$P(f)$ = ratio of peaks of filtered to unfiltered signal at frequency f

For the median filter:

$$M(f) = \begin{cases} 1, & 2n \leq f < 2n+1 \\ -1, & 2n+1 \leq f < 2n+2 \end{cases}$$

where $n = 0, 1, 2, 3, \dots$

f = frequency of square wave in cy/w

$M(f)$ = ratio of peaks of filtered to unfiltered signal at f




Figure 4.6. Response of midrange and average filters to square waves.

For the midrange filter:

$$R(f) = \begin{cases} 1, & 0 \leq f \leq 1/2 \\ 0, & f > 1/2 \end{cases}$$

where f = frequency of square wave in cy/w

$R(f)$ = ratio of peaks of filtered to unfiltered signal at frequency f

For the average filter:

$$A(f) = \begin{cases} 1 - 2n/f, & 2n \leq f < 2n+1/2 \\ (2n+1)/f - 1, & 2n+1/2 \leq f < 2n+3/2 \end{cases}$$

where $n = 0, 1, 2, 3, \dots$

f = frequency of square wave in cy/w

$A(f)$ = ratio of peaks of filtered to unfiltered signal at frequency f

Once again, the pseudomedian and midrange filter responses are related by a factor of 2 in frequency. However, responses of the median and average filters are not as similar for square waves as they are for sinusoidal or

triangle waves. Both these filters show phase inversion as they did for other periodic signals, but the median filter does not attenuate the signal at all. This is because the square waves are bi-valued, and therefore are infinite-length roots or oscillatory roots of the median filter. The root signal analysis in Chapter 2 showed that these median filter roots are not roots of the pseudomedian filter; the square wave response indicates that the pseudomedian filter output for a fast-fluctuating bi-valued signal is a constant (DC) value.

Pseudoresolution is a filtering problem related to the 180° phase error introduced by the median and average filters. For instance, if a signal has five peaks at a frequency that has a negative response from a certain filter, the filter output will show only four peaks positioned in the valleys between the original peaks. A negative response and pseudoresolution are essentially the same effect; both occur when the window is centered on an area that is, for instance, a valley between two peaks that form the majority of the window. The pseudomedian filter does not exhibit negative response or pseudoresolution since it is more center-weighted and thus is not "fooled" by such situations.

The above results clarify the behavior of the filters in a generalized theoretical sense, but do not necessarily indicate how the filters will behave on actual signals in discrete time. Although one might expect certain features of the response to hold true for the digital filters (for example, the phase inversion of the median filter at frequencies between 1 and 2 cy/w), it is not clear that the actual responses will mirror those shown in the curves, especially for high frequencies. However, for very wide filter windows, the correspondence between theory and practice should be quite good. The results shown in Figures 4.7-4.12, given for window widths of 75 and for frequencies of 0 to 4 cy/w in increments of 0.05 cy/w, are nearly identical to the continuous-time results. The value for the response was computed by averaging the ratios of the output signal value to input signal value at both positive and negative peaks of a long signal. Only at the higher frequencies do the effects of sampling and quantization cause some of the curves to differ from those derived in continuous time. The most noticeable anomalies occur for the median filter at 1.5, 2.5, and 3.75 cy/w. These frequencies are related by integer factors to the window

width of 75, and correspond to signal periods of exactly 50, 30, and 20 pixels per cycle respectively. The integer number of pixels per cycle at these frequencies causes any sampling anomalies to be repeated regularly throughout the signal and thus the sampling effects are larger at these frequencies. This effect is observed for other median filter window widths at frequencies that create an integer number of pixels per cycle.

The response curves for triangle waves (Figures 4.7 and 4.8) show values that seem too high for very low frequencies (0.05 to 0.1 cy/w). This is likely caused by quantization effects, and the experimental response curve values obtained for other frequencies are very close to the theoretical values. The median filter response to sinusoidal waves (Figure 4.9) shows the sampling anomalies described earlier, and the average filter response to square waves (Figure 4.12) is noisy at high frequencies. Otherwise, the empirically derived curves match the theoretical curves very closely.

In practice, filter windows are usually not very wide. For the median filter, widths from 3 to 11 are most common. For these widths, the responses of the filters vary significantly from the response curves derived above, especially as the frequency of the signal approaches the aliasing limit. For a given filter width W , input signal frequencies above $2/W$ cy/w (the Nyquist frequency) are aliased, and thus the response of the filter for higher frequencies is the same as that at the corresponding aliased frequency. But even for a 5-wide median filter, the phase inversion between 1 and 2 cy/w can still be observed. As the input signal frequency approaches the Nyquist frequency, the response curve deviates further and further from the theoretical values. This is shown in Figures 4.13-4.18, which are the response curves computed for 7-wide filters for frequencies from 0 to 3.5 cy/w in increments of 0.05 cy/w. Although the curves vary significantly from the theoretical values, the basic characteristics of the various filters at various frequencies are intact. For instance, the median and average filters exhibit phase inversion in the same frequency ranges, and the pseudomedian and midrange filters still generally block frequencies above 2 cy/w and 1 cy/w respectively. The results are quite noisy for all filters, although those for the average filter are significantly less

noisy and those for the median filter are the most noisy. At the Nyquist frequency (3.5 cy/w), the recurring root property of the median filter dominates its response and causes the median filter response to all periodic waves at this frequency to be -1. Still, it is clear that the continuous-time results illustrate the filters' basic behavior quite well, even for small window sizes.




Figure 4.7. Response curves for 75-wide median and pseudomedian filters acting on triangle waves.

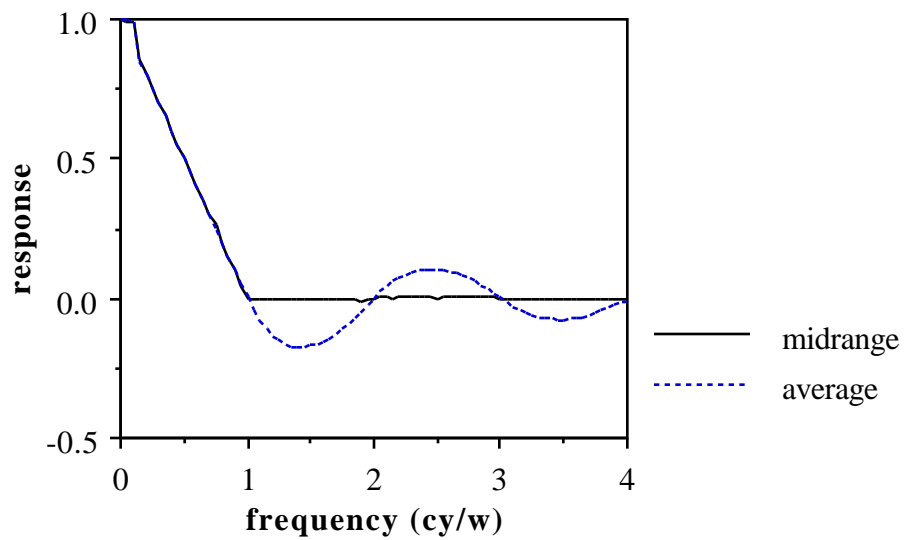


Figure 4.8. Response curves for 75-wide average and midrange filters acting on triangle waves.




Figure 4.9. Response curves for 75-wide median and pseudomedian filters acting on sinusoidal waves.




Figure 4.10. Response curves for 75-wide average and midrange filters acting on sinusoidal waves.

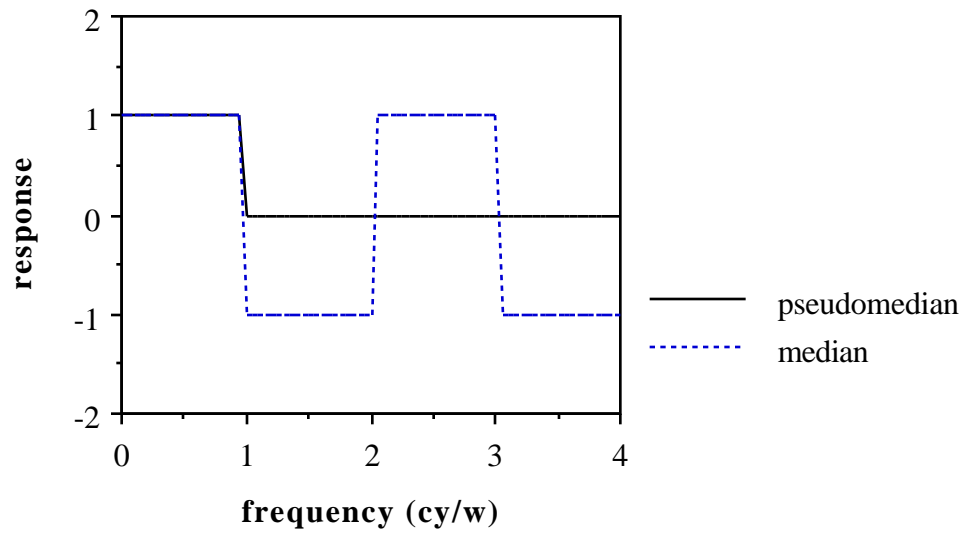


Figure 4.11. Response curves for 75-wide median and pseudomedian filters acting on square waves.




Figure 4.12. Response curves for 75-wide average and midrange filters acting on square waves.




Figure 4.13. Response curves for 7-wide median and pseudomedian filters acting on triangle waves.




Figure 4.14. Response curves for 7-wide average and midrange filters acting on triangle waves.




Figure 4.15. Response curves for 7-wide median and pseudomedian filters acting on sinusoidal waves.




Figure 4.16. Response curves for 7-wide average and midrange filters acting on sinusoidal waves.




Figure 4.17. Response curves for 7-wide median and pseudomedian filters acting on square waves.




Figure 4.18. Response curves for 7-wide average and midrange filters acting on square waves.

Distortion Analysis

Although the continuous-time response curves provide much insight into the behavior of the pseudomedian filter as compared to the other filters, these curves only represent a crude comparison of the amplitudes of the input and output signals and yield no information about the extent to which the filter modifies the shape of the waveform. A separate criterion will have to be used to quantify the distortion of these filters. The statistical correlation between the input and output signals is a good candidate to measure the changes between the two signals. However, since correlation is very sensitive to amplitude and phase differences between signals, a correlation taken directly between the input and output signals yields much the same information as the response curves, since these curves give amplitude and phase. But if the output signal is adjusted to compensate for amplitude and phase changes introduced by the filter, the correlation will measure the change in the shape of the waveform between input and output.

The statistical correlation between two amplitude-adjusted digital signals may take on any value from -1 to +1 inclusive, where +1 indicates exact correspondence between the two signals, -1 indicates 180° out-of-phase correspondence, and 0 indicates no correlation at all, as between a signal and its DC component. The correlation (CORR) is related to the mean square error (MSE) between the two signals by the equation $\text{CORR} = 1 - \text{normalized MSE}$. Values near one indicate high correlation and values near zero indicate low correlation. For phase-corrected signals, the correlation is strictly nonnegative, since negative correlations indicate a phase shift. Although in general the DC component of the signals should be removed before computing the correlation, this is not necessary in this case because the input signals have zero DC component and these filters do not induce any significant DC level distortion for the given inputs.

The correlation of the input and output signals can, in theory, be derived in continuous time, much like the above response curves. Consider a continuous-time input signal $x(t)$ and the corresponding filter output $y(t)$.

Define a gain of the filter (for a particular input signal with period T, where k is an arbitrary constant):

$$\text{GAIN} = \frac{\left(\int_k^{k+T} [y(t)]^2 dt \right)^{1/2}}{\left(\int_k^{k+T} [x(t)]^2 dt \right)^{1/2}}$$

The GAIN figure is negated for frequencies where the filtered signal is 180° out of phase with the input signal to adjust for the effect of the phase shift. This adjustment is valid because of the vertical symmetry of the signals considered. The correlation figure is then:

$$\text{CORR} = 1 - \frac{\left(\int_k^{k+T} \left| x(t) - \frac{1}{\text{GAIN}} y(t) \right|^2 dt \right)^{1/2}}{\left(\int_k^{k+T} [x(t)]^2 dt \right)^{1/2}}$$

Unfortunately, using the continuous-time definitions above to derive expressions for the correlation at various frequencies for the filters is extremely difficult, since equations must be developed for the entire output signal at all frequencies for each signal. Developing the response curves required only the peak value for each output signal. Even if continuous-time analyses were performed to obtain the correlation, the expressions for the correlation would most likely be quite cumbersome. Therefore, in this chapter the correlation results will be developed using digital implementations of the filters and signals. The definitions for the GAIN and CORR are modified as shown below. The sampled signals are $y(n)$ and $x(n)$ and the input signal has a period of N samples.

$$\text{GAIN} = \frac{\left(\begin{array}{c} k+N \\ \vdots \\ [y(n)]^2 \\ \vdots \\ n=k \end{array} \right)^{1/2}}{\left(\begin{array}{c} k+N \\ \vdots \\ [x(n)]^2 \\ \vdots \\ n=k \end{array} \right)^{1/2}}$$

$$\text{CORR} = 1 - \frac{\left(\begin{array}{c} k+N \\ \vdots \\ x(n) - \frac{1}{\text{GAIN}} y(n) \\ \vdots \\ n=k \end{array} \right)^{1/2}}{\left(\begin{array}{c} k+N \\ \vdots \\ [x(n)]^2 \\ \vdots \\ n=k \end{array} \right)^{1/2}}$$

The results for the correlation of the filtered signals with the original input signals are given in Figures 4.19-4.24 for each of the four filters acting on triangle, sine, and square waves. To ensure that the results approach the theoretical behavior of the filters, a very large window width of 75 was chosen, and input frequencies varied from 0 to 4 cy/w in 0.05 cy/w increments. As for the response curves, the correlation may vary significantly from these results for narrower windows; however, the distortion characteristics of the filters are easier to generalize when discretization effects are minimized by choosing a very wide window.

Figures 4.19 and 4.20 show the correlation figures for filtered triangle waves. Once again, similarities between the median and average filters and between the pseudomedian and midrange filters are evident. As expected, the correlation curves are noisy at much the same frequencies as the response curves. The pseudomedian filter passes triangle waves with a very high correlation up to 1 cy/w, after which the correlation declines approximately linearly to zero at 2 cy/w. The median filter shows maximum correlation at the

frequencies where its response curve is at a local maximum or minimum; note also that these correlation maxima are very high (above 0.8 even for some frequencies between 3 and 4 cy/w). Contrast this with the correlation curve for the average filter: its correlation is highest very near the zero points at integer cy/w values, and at a local minimum at the points of maximum response. All nonzero correlation values for the average filter are very high, though. The midrange filter again exhibits a curve very similar to the pseudomedian filter, with the expected 2-to-1 frequency relationship.

The correlation curves for filtered sinusoidal waves are given in Figures 4.21 and 4.22. These results are very similar to those for the triangular waves. Again, the median filter correlation curve is noisiest at high frequencies, and is shaped much like the triangular wave median filter correlation curve. The pseudomedian filter results also resemble the corresponding triangular wave curve. The average filter, however, shows nearly no distortion ($CORR \approx 1$) of sinusoidal waves for all frequencies except integer cy/w values. The smoothing caused by the average filter matches well with the original shape of the sinusoidal waves. The midrange filter does not show a correlation peak at 0.5 cy/w as it does for triangular waves; instead, its correlation decreases at an increasing rate from 0 to 1 cy/w. This deviates from the behavior of the pseudomedian filter and is perhaps due to the difference in edge response of the two filters.

For square waves, the filter correlation curves are given in Figures 4.23 and 4.24. The median filter correlation is again very noisy, but for non-integer frequencies (cy/w), the correlation is generally quite high. The pseudomedian filter has a correlation of exactly one for frequencies less than 1 cy/w and exactly zero for frequencies above 1 cy/w. This is expected from the root signal analysis and response curves. The response curves also indicate that perhaps the correlation for the median filter should always be 1 or 0 as well, since square waves are always passed by the filter. (They are either roots or recurring roots.) But since the phase of non-integer cy/w frequency signals relative to the sampling affects the output of each peak, the average of these peaks over many signal periodic (that is, over a long signal) causes the

correlation to drop below 1 for some frequencies above 1 cy/w. The correlation curve for the average-filtered square waves exhibits some of the same characteristics as that for triangular waves, but the peaks near integer cy/w frequencies become lower as the frequency increases. This can be predicted from the response curve and edge response of the average filter. The midrange filter once again does not mimic the pseudomedian filter at half the frequency; instead, its correlation for square waves drops quickly from one at 0 cy/w to zero at 0.5 cy/w. This is because the midrange filter creates severe edge distortion.

The response curves and correlation curves developed above for the pseudomedian, median, midrange, and average filters demonstrate some important similarities and differences between the filters. Although the root signal sets of the pseudomedian and median filters are closely related, the filters exhibit quite different responses to periodic signals. The pseudomedian filter blocks periodic signals with frequencies higher than 2 cy/w and does not exhibit any phase inversion, while the median filter does not completely block high frequency signals and causes phase inversion at certain frequencies. The midrange filter blocks high frequency signals like the pseudomedian filter, but the vast difference in edge response between the two filters creates very important differences in their responses to many types of signals. The characteristics of the filters that are revealed in the above analysis are helpful in selecting the appropriate filter for many types of one-dimensional signals. Now that the one-dimensional filter properties have been developed in reasonable detail, Chapter 5 introduces a generalization of the pseudomedian filter to two dimensions.




Figure 4.19. Correlation curves for 75-wide median and pseudomedian filters acting on triangle waves.




Figure 4.20. Correlation curves for 75-wide average and midrange filters acting on triangle waves.




Figure 4.21. Correlation curves for 75-wide median and pseudomedian filters acting on sinusoidal waves.




Figure 4.22. Correlation curves for 75-wide average and midrange filters acting on sinusoidal waves.




Figure 4.23. Correlation curves for 75-wide median and pseudomedian filters acting on square waves.




Figure 4.24. Correlation curves for 75-wide average and midrange filters acting on square waves.

Chapter 5

The Two-Dimensional Pseudomedian Filter

Developing a two-dimensional filter definition analogous to a given one-dimensional definition is usually not a difficult task. For example, the output of an average filter in two dimensions is merely the mean of all values in the two-dimensional window under consideration. Window shape and size have a great impact on the behavior of a two-dimensional filter, and the definitions of many types of filters apply easily to many different shapes. Like the average filter, the midrange filter readily adapts to any size or shape two-dimensional window, and the median filter output for any window with an odd number of values in it is the median of the values in the window. The pseudomedian filter, unlike the other three filters under consideration, is defined in one dimension according to a specific ordering of values within the window. That is, if the order of values in a given one-dimensional window are changed, the pseudomedian of the window is in general different than it was before the change, whereas the mean, median, and midrange of the window remain the same. This difference is important when attempting to generalize the definition of the pseudomedian to two dimensions, because the locations of values within the window affect the output. The key to developing a two-dimensional definition that is analogous to the one-dimensional definition is developing a two-dimensional window with subwindows that have meaning similar to the subwindows in one dimension.

In their introduction to the pseudomedian filter, Pratt, *et al.* [1] defined a two-dimensional pseudomedian filter for square, plus-shaped, and circular windows. The definition for the plus-shaped filter merely combines subwindows from a vertical line with those from a horizontal line. This formula correlates well with the one-dimensional definition, since the center point of the window appears in all subwindows and points nearer the center

appear in more subwindows than those more distant from the center. The definition given for the circular window seems incomplete, but is apparently based on three-up, three-across plus-shaped subwindows which are translated vertically and horizontally to yield "circle-shaped" subwindows and windows of any appropriate size. Although the definition seems analogous to the one-dimensional filter, the resulting windows are unfortunately rather oddly shaped. The 7x7 window shown in their paper is not even symmetric [1].

Despite the validity of the plus-shaped and circular pseudomedian filter definitions, these filters do not correspond to the most prevalent two-dimensional window shape, the square. Plus-shaped median and average filters are not uncommon, but square-shaped median and average filters are used much more often and are perhaps better understood. The square-shaped pseudomedian filter definition given by Pratt, *et al.* is not analogous to the one-dimensional filter. They propose stringing the values in the window by rows into a one-dimensional signal, and then performing a one-dimensional pseudomedian on the signal. If instead the data are put by columns into a one-dimensional signal, the result of the pseudomedian is different. This immediately indicates that the definition is not symmetric and therefore not analogous to the one-dimensional filter.

Another indication of the problem with this square-shaped pseudomedian filter definition can be demonstrated for a 3x3 window with data given by the letters *a* through *i* as shown in Figure 5.1.

<i>a</i>	<i>b</i>	<i>c</i>
<i>d</i>	<i>e</i>	<i>f</i>
<i>g</i>	<i>h</i>	<i>i</i>

Figure 5.1. 3x3 square two-dimensional filter window.

One of the most important aspects of the pseudomedian filter definition is that values near the center appear in more subwindows than those further away from

the center; in addition, points that are equidistant from the center (in any direction) are in the same number of subwindows. If the definition of Pratt, *et al.* is used, the points in the window are put in order (a,b,c,d,e,f,g,h,i) and so the subwindows are (a,b,c,d,e) ; (b,c,d,e,f) ; (c,d,e,f,g) ; (d,e,f,g,h) ; and (e,f,g,h,i) . Points b and d are each one point away from the center point e , but b is in only two of the five subwindows while d appears in four. Even c , which is one point diagonally away from e and thus further away than either b or d , is in three subwindows. The weighting of the window is strangely skewed.

There is, however, a logical choice of subwindows that matches all important aspects of the one-dimensional pseudomedian filter. There are 4 distinct subwindows of size 2×2 in a 3×3 window: (a,b,d,e) ; (b,c,e,f) ; (d,e,g,h) ; and (e,f,h,i) . These subwindows are a logical set upon which to perform the pseudomedian, as shown below.

$$\text{PMED} = 0.5 \cdot \max\{\min(a,b,d,e), \min(b,c,e,f), \min(d,e,g,h), \min(e,f,h,i)\} \\ + 0.5 \cdot \min\{\max(a,b,d,e), \max(b,c,e,f), \max(d,e,g,h), \max(e,f,h,i)\}$$

Not only are there fewer subwindows of fewer elements in this scheme, but all the characteristics of the definition cited above are fulfilled: the center point e is in all four subwindows, points one pixel distant vertically or horizontally (b , d , f , and h) appear in two subwindows each, and points that are diagonally one pixel distant (a , c , g , and i) are in one subwindow each. Larger square windows require larger subwindows, and the progression is predictable: 5×5 windows have nine 3×3 subwindows, 7×7 windows have 25 subwindows of size 5×5 , etc. In general:

For a square two-dimensional signal segment of size $(2N+1) \times (2N+1)$, the pseudomedian is the average of the maximum of the minima of A and the minimum of the maxima of A , where A denotes the set of the $(N+1)^2$ square subwindows of size $(N+1) \times (N+1)$ inside the $(2N+1) \times (2N+1)$ window.

The correspondence between the above definition and a similarly-worded one-dimensional definition (below) is obvious.

For a one-dimensional signal segment of length $2N+1$, the pseudomedian is the average of the maximum of the minima of B and the minimum of the maxima of B , where B denotes the set of the $N+1$ subwindows of size $N+1$ inside the length $2N+1$ window.

Rigorous theoretical analysis of the properties of two-dimensional filters is much more difficult than for one-dimensional filters. Indeed, the two-dimensional median filter is not very well understood, so a comparison of the theoretical properties of the pseudomedian filter with those of the median filter in two dimensions is not on a foundation as solid as the comparison in one dimension. However, there is at least one characteristic of the two-dimensional pseudomedian filter that stands in clear contrast to the median filter: the behavior of the filter at sharp corners.

The pseudomedian filter has a "corner-preserving" characteristic that resembles the idea of constant neighborhood preservation from the one-dimensional root signal analysis. Property 2 (Chapter 2) states that the one-dimensional pseudomedian filter output for any window containing a constant neighborhood ($N+1$ points of equal value in a row) is the value of the constant neighborhood. This property also holds for the one-dimensional median filter. For the pseudomedian filter, the proof of Property 2 can be generalized for the given two-dimensional definition, yielding the following property:

Property 3: The pseudomedian of a $(2N+1) \times (2N+1)$ window that contains an $(N+1) \times (N+1)$ area of constant value (a "constant area") is the value of the constant area.

This property does not hold for the median filter, because the proof of Property 2 for the median filter rests on the idea that more than half the values in the window are in the constant neighborhood. This is clearly not true for the two-dimensional case in Property 3; in fact, just over one quarter of the values in the window must be constant when a constant area is in the window.

While both the median and pseudomedian filters will pass a straight edge in two dimensions unchanged, the median filter will "round off" corners and often give a filtered image a smooth appearance at edges. The pseudomedian filter, however, passes sharp corners meeting the constant area criterion of Property 3, and so the edges of a pseudomedian-filtered image often exhibit sharper corners and edges and a more "blocky" appearance than median-filtered images. Examples of the effects of the median and pseudomedian filters on corners of various angles are shown in Figures 5.2, 5.3, and 5.4.

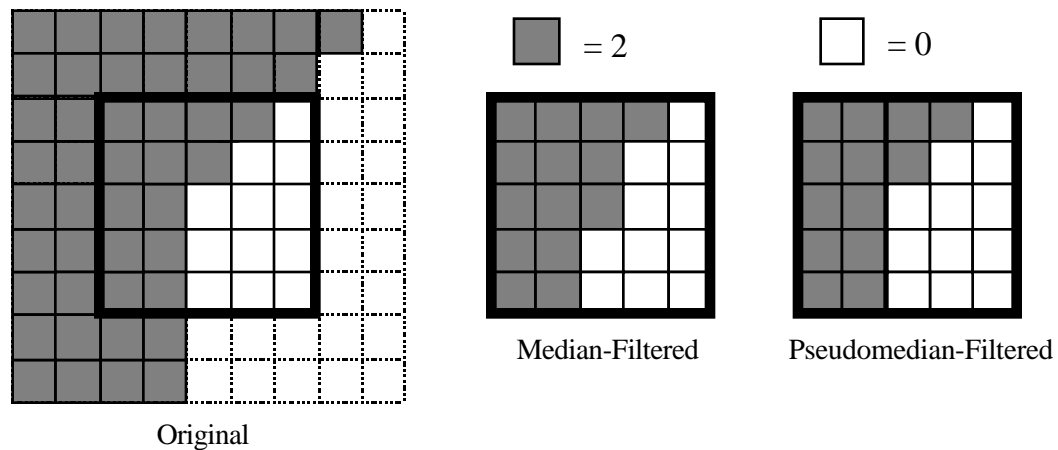


Figure 5.2. 135° corner filtered by 5x5 square median and pseudomedian filters.




Figure 5.3. 90° corner filtered by 5x5 square median and pseudomedian filters.




Figure 5.4. 45° corner filtered by 5x5 square median and pseudomedian filters.

Figure 5.4 shows that the pseudomedian filter does not change corners sharper than 90° (that is, corners which do not create constant areas) as much as the median filter does. Instead of creating a less acute angle in the corner, the pseudomedian filter fills in the sharp peak at the corner with an averaged value ("1" in the figure). This characteristic is somewhat similar to the response of the one-dimensional pseudomedian filter to impulses. For the

pseudomedian filter, the response to a one-dimensional impulse is derived from Observations 1 and 2 in Chapter 2, and these observations generalize easily to two dimensions, yielding the response to sharp corners described above. Depending upon the application, this feature of the filter may or may not be desirable. Clearly, the pseudomedian filter does not distort sharp corners as badly as does the median filter.

The corner-preserving characteristic of the pseudomedian filter provides perhaps the best explanation for an observed "blockiness" of most pseudomedian-filtered images. Features of noisy images that are pseudomedian-filtered often have a "stair-step" appearance at their edges. Median-filtered images have a much smoother appearance at edges, but as a result the edges are sometimes not as steep as those in pseudomedian-filtered images. One-dimensional pseudomedian-filtered signals also exhibit a blocky appearance, much like the two-dimensional images. This effect cannot be explained by the corner-preserving characteristic, as that property only applies in two dimensions. Noisy edges or ramps, particularly those with additive high-frequency periodic noise, show a "stair-step" appearance when pseudomedian filtered in one dimension also. Although there is no obvious explanation for the blocky appearance of one-dimensional pseudomedian-filtered signals, the "stair-step" effect on some slopes is at least partially due to the high-frequency stop characteristic of the pseudomedian filter described in the root signal analysis and in Chapter 4. Examples illustrating all these effects are given below in the applications section of this chapter.

Examples

This section gives results of two-dimensional pseudomedian, median, average, and midrange filtering on several types of images. The examples begin with binary images of geometrical objects and progress through images with simulated additive noise to noisy images from real applications. This progression helps to relate the filters' behavior in situations from the artificial to the practical. Some obvious theoretical differences between the

filters often become less clear in real applications, while the effects of other differences are occasionally amplified.

An image consisting of four geometrical objects (a square, a circle, a cross, and a triangle) is the basis for the examples given in Figures 5.5 through 5.9. The original image, shown in Figures 5.5*a* and 5.6*a*, consists of only two values, one for the background (a value of 0) and one for the objects (a value of 2). On the right of Figure 5.5 are the pseudomedian- and median-filtered results for this image for a window size of 7×7 . Figure 5.6 shows 7×7 average and midrange filter outputs for this image. The pseudomedian-filtered result (Figure 5.5*b*) is identical to the original except for small regions at the points of the triangle which are reduced to one-half amplitude (a value of 1 for this example). This output is not surprising in light of the corner-preserving property (Property 3) described above. The sharp points of the triangle can be thought of as a type of two-dimensional "impulse." The median-filtered image (Figure 5.5*c*) shows more distortion of the corners of the objects in the filter, rounding them slightly; however, the edges in the median-filtered image remain very sharp. The average and midrange filter results (Figure 5.6*b,c*) show edge blurring similar to that described for these filters in Chapter 3. The average filter, however, rounds sharp corners of the objects, while the midrange filter preserves them.

Figures 5.7 and 5.8 give the filters' response to the geometrical object image when it is corrupted by a Gaussian noise process. The image that is filtered is shown in Figures 5.7*a* and 5.8*a* and is the same image as in the previous example with additive zero-mean Gaussian noise (with standard deviation, σ , equal to 0.8) to yield a signal-to-noise ratio of 8 decibels (dB). The resulting image is rescaled to fit the grayscale (hence the change in overall intensities of objects and backgrounds from the noiseless images). The results of 7×7 pseudomedian and median filtering are shown in Figure 5.7*b* and *c*, and the results of 7×7 average and midrange filtering in Figure 5.8*b* and *c*. Changes in the objects are mostly the same as those introduced by the filters in the noiseless images of Figures 5.5 and 5.6, and all the filters seem to reduce the noise quite effectively. However, a certain "blockiness" is apparent in the

pseudomedian- and midrange-filtered images, while the median- and average-filtered images generally have a smoother appearance. The pseudomedian filter causes less distortion of objects than the other filters because of its corner-preserving property.

How well each of the filters performs on the binary geometrical objects in the examples above is indicated by thresholding the filtered images to determine how many points are misclassified by the filters. Figure 5.9 shows the results of thresholding each of the filtered images in Figures 5.7 and 5.8 at a nominal value of 1 (halfway between the noiseless nominal background level of 0 and the noiseless nominal object level of 2). Points that are in the startup area at the edges of each image (for a 7×7 filter, the three lines at each edge) are removed, as these points are unfiltered. There are 3926 points in the unfiltered image that are misclassified by the thresholding operation. The pseudomedian-filtered image (Figure 5.9*a*) gives only 65 misclassified points, while in this example the median- and average-filtered images (Figures 5.9*b* and *c*) both result in 207 misclassified points (although the specific points they miss are at different locations, and in general the number of points they miss are different). The midrange filter is not really suited to this segmentation problem, since its edge response creates large areas with values close to the threshold. As one might expect, the segmented midrange-filtered image (Figure 5.9*d*) has a large number of misclassified points, 1055. Clearly, the pseudomedian filter is the most effective of the four filters in reducing noise and preserving the objects in this example.

One type of noise against which the median filter is very effective is impulse noise, such as salt-and-pepper noise. An example of this is shown in Figure 5.10*a*. This image has one-tenth of its points changed to the opposite value; that is, the selected points on the objects (value 2) are changed to the background value (0) and vice versa. The median and pseudomedian filters are applied to this image, with markedly different results. The 7×7 median filter (Figure 5.10*b*) removes nearly all of the salt-and-pepper noise, and the results are very similar for both large- and small-sized median filters. The pseudomedian filter, however, shows a much different response. The 3×3

pseudomedian filter (Figure 5.10*c*) reduces the impulsive noise, but "gray" spots remain in many locations. The objects are mostly intact, but gray rectangular areas appear in some locations, especially near the object edges. The 7x7 pseudomedian filter (Figure 5.10*d*) exhibits even more aberrant behavior; the gray areas are larger and appear throughout the image, and the objects are still visible but not very distinct. This behavior of the pseudomedian filter is related to its response to impulses, and the blocky effect is caused by nearby impulses that cause each subwindow of the overall filter window to include a impulse. The median filter, which is a strictly nonlinear operator, removes nonlinear impulse noise better than the pseudomedian filter, which is a linear combination of two nonlinear operators. But the pseudomedian filter seems to work better for linear noise processes (such as Gaussian noise).

Applications

The primary reason for choosing to use the pseudomedian filter over other linear or nonlinear filters in a given situation is its superior detail-preserving characteristics. Assuming the noise is not extremely impulsive (as is salt-and-pepper noise), the noise reduction by the pseudomedian filter is often adequate, although it cannot match the noise reduction of the median filter except in low-noise situations. Figures 5.11 through 5.14 demonstrate the tradeoffs between the various filters for two grayscale images corrupted by different levels of additive Gaussian noise. The first example is given in Figures 5.11 and 5.12. The original, noiseless image is in Figures 5.11*a* and 5.12*a*, and the image to be filtered (Figures 5.11*b* and 5.12*b*) has had simulated white noise added to it to yield a signal-to-noise ratio of 30 dB. This is a low-noise situation, so a relatively small window size of 5x5 is appropriate. The pseudomedian filter (Figure 5.11*c*) performs well because it reduces the noise without distorting the image features very much. The median and average filters (Figures 5.11*d* and 5.12*c*) exhibit much more blurring and distortion of the image, and of course the midrange filter (Figure 5.12*d*) destroys all sharp edges in the picture. The second picture (Figures 5.13 and 5.14, original in *a*, noisy in *b*) has a signal-to-noise ratio of 20 dB. A larger window size of 7x7 is

chosen because it is more effective at reducing noise than smaller window sizes. The median and average filters again show good noise reduction at the expense of image distortion. The pseudomedian filter (Figure 5.13*c*) does not distort the image features as badly as the other filters, but the noise reduction is not as good, and the image has a distinct "blockiness" to it. Thin lines in the original, such as the brake cables and the rim of the sunglasses, are severely distorted or completely removed in the median-filtered image (Figure 5.13*d*), but the pseudomedian-filtered image still shows them quite clearly. Also, details in the hair are much clearer in the pseudomedian-filtered image. The average-filtered image (Figure 5.14*c*) appears somewhat similar to the median-filtered image, although slightly more blurred. The midrange-filtered image (Figure 5.14*d*) has a blocky appearance similar to that of the pseudomedian-filtered image, but with larger blocks corresponding to the 2-to-1 window size relationship between the filters described earlier. The midrange filter, however, does not preserve edges or other details nearly as well as any of the other filters.

The above examples illustrate the differences between the filters when they are used on images with artificially added noise. In practice, however, images to be processed are already corrupted by noise which is often unknown or not well-described. Figures 5.15 through 5.18 are examples of images which were acquired with noise and other image defects that are to be reduced or removed by filtering. The first example, in Figures 5.15*a* and 5.16*a*, is a portion of a photomicrograph of feline corneal endothelial cells. The ultimate goal for this image is segmentation of the picture into individual cells. This image has many defects, including a left-to-right background level gradient, low signal strength, prominent noise, and miscellaneous point and line (impulsive) defects. Figures 5.15*b* and 5.16*b* show this image after the background level has been flattened by subtracting the best-fit fifth-order polynomial for each horizontal line from that line. The grayscale of this image has also been adjusted by limiting the original scale and adjusting the limited picture to fill the entire scale again. The noise and impulse defects are much more obvious in this image, and running vertically along the right side of the picture there seems to be some two-dimensional periodic impulse-like noise.

The pseudomedian filter is an obvious candidate for filtering this image, because although the median filter would remove the nonperiodic impulse noise, the periodic noise may be a two-dimensional fast-fluctuating root of the filter [3, 16]. Although this chapter has not explicitly examined the root structures of the two-dimensional filters, many analogies between the one-dimensional and two-dimensional behavior of the median and pseudomedian filters hold. The root signal comparison in Chapter 2 leads to the conclusion that fast-fluctuating roots of the pseudomedian filter do not exist. It is much more difficult to establish a definition of a fast-fluctuating signal in two-dimensions, but there are no periodic roots of the pseudomedian filter consisting of areas smaller than a "constant area" of the filter. This property makes the pseudomedian filter a good choice to remove the periodic impulse noise in the corneal cell images.

The difference between 5x5 pseudomedian- and median-filtered versions of the corneal cell image is shown in Figure 5.15*c* and *d*. Traces of the periodic noise are still visible in the median-filtered image, but the impulse-like defects elsewhere in the picture (a vertical line, several spots, and a plus-shaped cursor) are removed completely. By contrast, the pseudomedian filter eliminates the periodic noise but does not fully remove the impulsive defects. The edges of the cells are smoother in the median-filtered image, and blockier in the pseudomedian-filtered image. This difference is even more apparent in the images filtered by 9x9 pseudomedian and median filters shown in Figures 5.16*c* and *d*. The median-filtered image does not show any evidence of the periodic noise at this window width, but otherwise the differences between the two filtered images are similar to those in Figure 5.15. The wider windows seem to give results that are more suitable for image segmentation. The pseudomedian-filtered images have sharper edges than the median-filtered images, although the cell shapes are more squared off because of the blockiness of the pseudomedian filter. The sharp edges of the pseudomedian-filtered image probably would allow more tolerance in the choice of threshold for segmentation, but a segmented median-filtered image would exhibit smoother, more rounded cell edges.

Figure 5.17*a* is an unfiltered thermographic image of a circular blackbody and a graph of a vertical linescan of the image at the leftmost edge of the blackbody. Thermographic images taken with a flying-spot scanner are characterized by a particular set of defects, including a horizontal streak artifact and prominent noise that is somewhat impulsive [27]. The streak artifact appears as the oscillations on the plateau in the linescan graph. Since the noise process in thermographic images is one-dimensional, horizontal sliding filters are often used on these images; the impulsive nature of the noise gives the median filter many advantages for thermographic images. Ryu [27] demonstrated that a filtering method for thermographic images which gives good visual results consists of 5-wide horizontal non-recursive median filtering followed by 3-wide vertical triangular-weighted non-recursive average filtering. The vertical triangular filter is meant to reduce the effect of the horizontal streak artifact by blurring the image slightly in the vertical direction. Although Ryu's method does not create much distortion in the image, often better noise and artifact removal is desired, and often two-dimensional median filters are used on thermal images. Two-dimensional median filtering usually blurs the image more than other methods, and some window widths exhibit fast-fluctuating or oscillatory roots at edges affected by the horizontal streak artifact.

The two-dimensional pseudomedian filter seems a likely candidate for use on thermographic images, since it reduces noise while limiting distortion of fine details and overall blurring. In addition, there are no fast-fluctuating or recurring root problems for the pseudomedian filter. These properties suggest that it may be well-suited to filtering thermographic images. Figure 5.17*b* shows the 5x5 pseudomedian-filtered image and its corresponding vertical linescan; Figure 5.17*c* is a 5x5 median-filtered image and linescan. An image filtered by Ryu's 5-wide horizontal median and 3-wide vertical triangular filtering method is in Figure 5.17*d* along with its vertical linescan. The pseudomedian filter does indeed eliminate the oscillations on the edge of the object; however, the oscillations are replaced by "stairsteps" along the ramp up the edge. The overall effect in the image is not unlike the streaks in the original, although in actuality the oscillations are gone and only a blocky edge remains.

The 5x5 median filter exhibits phase inversion for the oscillations along the flat part of the object while smoothing the ramp up the edge quite well. Ryu's method gives a relatively flat plateau with fairly smoothly ramped edges, but the original oscillations are still slightly visible all along the linescan. Although the pseudomedian filter removes the oscillations while preserving the edge steepness and height, its blockiness or "stairstep" effect limits its use to situations where these types of edges are acceptable.

Figures 5.18 through 5.20 further demonstrate the various types of filtering for thermographic images. A high-contrast thermographic image of an object with fine details is shown in Figure 5.18*a* in grayscale, and in Figure 5.19*a* in pseudocolor. Figure 5.20*a* is a low-contrast pseudocolor thermographic image of the author's face. In each figure, the *b* image is filtered by a 5x5 pseudomedian filter, the *c* image by a 5x5 median filter, and the *d* image by Ryu's horizontal median/vertical triangular filtering method. Most of the noticeable differences between the filtered images have been described before: blocky edges for the pseudomedian-filtered image, rounded corners and slightly blurred details for the median-filtered image, and vestiges of the streak artifact in the image filtered by Ryu's method. Among the effects of the pseudomedian filter that are discernable in the images are the "stairstep" effect along diagonal edges in all the images and less clipping of the sharp peak at the top of the double chevron in Figure 5.18*b*. Once again, the corner-preserving property of the pseudomedian filter may be a significant advantage when filtering certain types of images.

The examples given in this chapter illustrate the differences among the major types of filters examined in this thesis. The average and median filters are well-established image processing tools, and situations where they are used to best advantage are quite common and well-known. The pseudomedian filter is another option for image processing that behaves much like the median filter, but possesses some unique properties that may make it superior in some instances, for example, when fine details need to be preserved. This chapter has introduced the corner-preserving property and the related "stairstep" effect for the two-dimensional pseudomedian filter, and briefly described two-

dimensional fast-fluctuating or recurring root properties analogous to those for the one-dimensional median and pseudomedian filters. This extension of the one-dimensional descriptions of the pseudomedian filter and its properties to two dimensions creates a new image processing tool which has fairly predictable behavior and is useful for images with certain characteristics, such as objects with sharp corners.

Figure 5.5. (a) Original image of geometrical objects. (b) 7x7 pseudomedian-filtered image. (c) 7x7 median-filtered image.

Figure 5.6. (a) Original image of geometrical objects. (b) 7x7 average-filtered image. (c) 7x7 midrange-filtered image.

Figure 5.7. (a) Original image of geometrical objects corrupted by Gaussian (white) noise. Signal-to-noise ratio = 8 dB. (b) 7x7 pseudomedian-filtered image. (c) 7x7 median-filtered image.

Figure 5.8. (a) Original image of geometrical objects corrupted by Gaussian (white) noise. Signal-to-noise ratio = 8 dB. (b) 7x7 average-filtered image. (c) 7x7 midrange-filtered image.

Figure 5.9. (a) Segmented 7x7 pseudomedian-filtered image from Figure 5.7b, threshold at one-half of full grayscale. (b) Segmented 7x7 median-filtered image from Figure 5.7c. (c) Segmented 7x7 average-filtered image from Figure 5.8b. (d) Segmented 7x7 midrange-filtered image from Figure 5.8c.

Figure 5.10. (a) Original image of geometrical objects corrupted by salt-and-pepper noise. One-tenth of points changed. (b) 7x7 median-filtered image. (c) 3x3 pseudomedian-filtered image. (d) 7x7 pseudomedian-filtered image.

Figure 5.11. (a) Original image of tiger. (b) Image corrupted with Gaussian (white) noise. Signal-to-noise ratio = 30 dB. (c) 5x5 pseudomedian-filtered image. (d) 5x5 median-filtered image.

Figure 5.12. (a) Original image of tiger. (b) Image corrupted with Gaussian (white) noise. Signal-to-noise ratio = 30 dB. (c) 5x5 average-filtered image. (d) 5x5 midrange-filtered image.

Figure 5.13. (a) Original image of biker. (b) Image corrupted with Gaussian (white) noise. Signal-to-noise ratio = 20 dB. (c) 7x7 pseudomedian-filtered image. (d) 7x7 median-filtered image.

Figure 5.14. (a) Original image of biker. (b) Image corrupted with Gaussian (white) noise. Signal-to-noise ratio = 20 dB. (c) 7x7 average-filtered image. (d) 7x7 midrange-filtered image.

Figure 5.15. (a) Original image of feline corneal endothelial cells. (b) Modified original with background gradient removed and contrast enhanced to full grayscale. (c) 5x5 pseudomedian-filtered image. (d) 5x5 median-filtered image.

Figure 5.16. (a) Original image of feline corneal endothelial cells. (b) Modified original with background gradient removed and contrast enhanced to full grayscale. (c) 9x9 pseudomedian-filtered image. (d) 9x9 median-filtered image.

Figure 5.17. (a) Unfiltered thermogram of a circular blackbody with vertical linescan at leftmost edge of object. (b) 5x5 pseudomedian-filtered image with linescan. (c) 5x5 median-filtered image with linescan. (d) Image filtered by 5-wide horizontal median and 3-wide vertical triangular filters with linescan.

Figure 5.18. (a) Unfiltered thermogram of a class ring. (b) 5x5 pseudomedian-filtered image. (c) 5x5 median-filtered image. (d) Image filtered by 5-wide horizontal median and 3-wide vertical triangular filters.

Figure 5.19. (a) Unfiltered thermogram of a class ring in pseudocolor. (b) 5x5 pseudomedian-filtered image. (c) 5x5 median-filtered image. (d) Image filtered by 5-wide horizontal median and 3-wide vertical triangular filters.

Figure 5.20. (a) Unfiltered thermogram of the author's face in pseudocolor. (b) 5x5 pseudomedian-filtered image. (c) 5x5 median-filtered image. (d) Image filtered by 5-wide horizontal median and 3-wide vertical triangular filters.

Chapter 6

Conclusions

This thesis has analyzed the pseudomedian filter several different ways in order to enhance the understanding of the filter's operation and how it compares to several other types of filters. Although the pseudomedian filter was originally developed to mimic the median filter, its behavior differs in important ways. Root signal analysis, a rigorously theoretical comparison of the pseudomedian and median filters, reveals an exciting difference between them: the lack of fast-fluctuating or oscillatory roots for the pseudomedian filter. This result presages the findings of the continuous-time analysis concerning the behavior of median and pseudomedian filters on high-frequency periodic signals.

Although the root signal sets indicate that the pseudomedian filter may be a good substitute for the median filter, this analysis alone is not a complete representation of their behavior. A comparison of the response of the pseudomedian filter to edges and impulses with the responses of the median, average, and midrange filters, shows significant differences in the filter outputs for these common signal features. Excellent edge response and impulse suppression are prized characteristics of the median filter. While the edge response of the pseudomedian filter is just as sharp as that of the median filter, the pseudomedian filter does not completely remove impulses. However, the pseudomedian filter also does not usually spread the effects of an impulse across neighboring sections of the signal as much as the average or midrange filters do.

To further assist in understanding the differences among the various filters in this study, a method of generalizing the response of the filters to periodic signals of various frequencies is developed. Although this method relies on the use of continuous-time constructs to describe inherently digital

operators, the results illustrate some important properties of the filters. A similarly-developed distortion measure helps describe the extent to which the filters change the shape of the signals at various frequencies. This analysis brings out some similarities between the pseudomedian and midrange filters, which had shown no resemblance in the previous examinations. Both those filters do not exhibit any phase inversion to periodic signals of any frequency, whereas both the median and average filters have phase inversion and pseudoresolution problems for some frequency ranges of periodic signals.

The square two-dimensional pseudomedian filter which corresponds logically to the one-dimensional definition displays many properties similar to those derived in one dimension. In addition, the two-dimensional filter has a corner-preserving property that clearly distinguishes it from the two-dimensional median filter. Almost all noticeable differences among the filtered images in this thesis can be explained easily by the various properties and response comparisons developed for the four filters. Thus, the methods used in this thesis to compare and contrast the pseudomedian, median, average, and midrange filters can be used to explain as well as predict their behavior in many different situations.

Although the results of this research answer many questions about the pseudomedian filter, there are some areas that invite further investigation. The "blockiness" exhibited by the two-dimensional pseudomedian filter is obviously related to the shape of the filter window (and subwindows), which in this study is square. Circle-, diamond-, plus-, and cross-shaped subwindows may all be suitable to developing valid definitions for two-dimensional pseudomedians. These shapes would have some similarities to the square pseudomedian, but certain properties (such as the corner-preserving property) would surely be modified and the filtering results would be different. Another important area not considered in detail is the implementation of the pseudomedian filter and how it compares to the implementation of other filters. The pseudomedian is the most conceptually complex of the four filters studied, but it was initially developed as a more efficient substitute for the median filter. The theoretical order of the algorithm for computing the pseudomedian is indeed

lower than that of most algorithms for computing the median, but the fast median algorithm of Huang, *et al.* [5] appears to be faster than any pseudomedian implementation. Since a thorough study of the filter algorithms has not been performed, no specific conclusions can be drawn about the relative speeds of the filters. However, the differences in computation time for the various filters in this study may be small, and if so the best reason for using one instead of another is the difference in the responses of the filters. This study clearly shows that the pseudomedian filter response is significantly different from those of the other filters for most signals and images.

The methods of analysis used in this thesis were developed to help draw precise comparisons between nonlinear and linear filters. The particular methods used were designed to emphasize the most useful aspects of the filters, especially the pseudomedian and median filters. The properties of the filters revealed by these analyses are theoretically interesting, but more importantly, are practically useful in predicting and understanding the effects of the filters on signals and images. By knowing the characteristics of each filter, one may choose the most appropriate filter for each type of signal or image.

Bibliography

1. Pratt, William K., Ted J. Cooper, and Ihtisham Kabir. "Pseudomedian filter." *Proc SPIE*, v. 534 (1985), pp. 34-43.
2. Gallagher, Jr., Neal C. and Gary L. Wise. "A theoretical analysis of the properties of median filters." *IEEE Trans Acoust Speech Signal Process*, v. ASSP-29 n. 12 (1981), pp. 1136-1141.
3. Tyan, S. G. "Median filtering: deterministic properties." in T. S. Huang, ed., *Two-Dimensional Digital Signal Processing II*, Berlin: Springer-Verlag, 1981, pp. 197-217.
4. Justusson, B. I. "Median filtering: statistical properties." in T. S. Huang, ed., *Two-Dimensional Digital Signal Processing II*, Berlin: Springer-Verlag, 1981, pp. 161-196.
5. Huang, T. S., G. J. Yang, and G. Y. Yang. "A fast two-dimensional median filtering algorithm." *IEEE Trans Acoust Speech Signal Process*, v. ASSP-27 n. 1 (1979), pp. 13-18.
6. Bovik, Alan Conrad. "Effect of median filtering on edge estimation and detection." *IEEE Trans Pattern Anal Mach Intell*, v. PAMI-9 n. 2 (1987), pp. 181-194.
7. Bovik, Alan Conrad. "Streaking in median filtered images." *IEEE Trans Acoust Speech Signal Process*, v. ASSP-35 n. 4 (1987), pp. 493-503.
8. Nodes, Thomas A. and Neal C. Gallagher, Jr. "The output distribution of median type filters." *IEEE Trans Commun*, v. COM-32 n. 5 (1984), pp. 532-541.

9. Nodes, Thomas A. and Neal C. Gallagher, Jr. "Median filters: some modifications and their properties." *IEEE Trans Acoust Speech Signal Process*, v. ASSP-30 n. 10 (1982), pp. 739-746.
10. Nodes, Thomas A. and Neal C. Gallagher, Jr. "Some results on the median filtering of signals and additive white noise." *Proc 19th Ann Allerton Conf on Commun Control Computing* (1981), pp. 99-108.
11. Arce, Gonzalo R. and Michael P. McLoughlin. "Theoretical analysis of the max/median filter." *IEEE Trans Acoust Speech Signal Process*, v. ASSP-35 n. 1 (1987), pp. 60-69.
12. Arce, Gonzalo R. and Russell E. Foster. "Detail-preserving ranked-order based filters for image processing." *IEEE Trans Acoust Speech Signal Process*, v. ASSP-37 n.1 (1989), pp. 83-98.
13. Bednar, J. Bee and Terry L. Watt. "Alpha-trimmed means and their relationship to median filters." *IEEE Trans Acoust Speech Signal Process*, v. ASSP-32 n.1 (1984), pp. 145-153.
14. Restrepo, Alfredo and Alan Conrad Bovik. "Adaptive trimmed mean filters for image restoration." *IEEE Trans Acoust Speech Signal Process*, v. ASSP-36 n.8 (1988), pp. 1326-1337.
15. Peterson, Steven R., Yong Hoon Lee, and Saleem A. Kassam. "Some statistical properties of alpha-trimmed mean and standard type M filters." *IEEE Trans Acoust Speech Signal Process*, v. ASSP-36 n.5 (1988), pp. 707-713.

16. Astola, Jaakko, Pekka Heinonen, and Yrjö Neuvo. "On root structures of median and median-type filters." *IEEE Trans Acoust Speech Signal Process*, v. ASSP-35 n. 8 (1987), pp. 1199-1201.
17. Astola, Jaakko, Pekka Heinonen, and Yrjö Neuvo. "Linear median hybrid filters." *IEEE Trans Circuits and Syst*, v. CAS-36 n. 11 (1989), pp. 1430-1438.
18. Neejärvi, Jukka, Pekka Heinonen, and Yrjö Neuvo. "Sine wave responses of median type filters." *Proc 1988 IEEE Int Symp Circuits and Syst*, v. 2 (1988), pp. 1503-1506.
19. Nieminen, Ari, Pekka Heinonen, and Yrjö Neuvo. "Median-type filters with adaptive substructures." *IEEE Trans Circuits and Syst*, v. CAS-34 n. 7 (1987), pp. 1275-1290.
20. Heinonen, Pekka and Yrjö Neuvo. "FIR-median hybrid filters with predictive FIR substructures." *IEEE Trans Acoust Speech Signal Process*, v. ASSP-36 n.6 (1988), pp. 892-899.
21. Coyle, Edward J. and Jean-Hsang Lin. "Stack filters and the mean absolute error criterion." *IEEE Trans Acoust Speech Signal Process*, v. ASSP-36 n.8 (1988), pp. 1244-1254.
22. Ko, S.-J., Y. H. Lee, and A. T. Fam. "Selective median filters." *Proc 1988 IEEE Int Symp Circuits and Syst*, v. 2 (1988), pp. 1495-1498.
23. Kuwahara, M., K. Hachimura, S. Eiho, and M. Kinoshita. "Processing of RI-angiocardigraphic image" in K. Preston, Jr. and M. Onoe, eds., *Digital Processing of Biomedical Images*. New York: Plenum, 1976, pp. 187-202.

24. Lee, Yong-Hwan and Sawasd Tantaratana. "Decision-based order statistic filters." *IEEE Trans Acoust Speech Signal Process*, v. ASSP-38 n.3 (1990), pp. 406-420.
25. Davis, L. S. and Azriel Rosenfeld. "Noise cleaning by iterated local averaging." *IEEE Trans Syst Man Cybernetics*, v. SMC-8 n. 9 (1978), pp. 705-710.
26. Lev, Amos, Steven W. Zucker, and Azriel Rosenfeld. "Iterative enhancement of noisy images." *IEEE Trans Syst Man Cybernetics*, v. SMC-7 n. 6 (1977), pp. 435-442.
27. Ryu, Zee Man. "The effect of nonlinear filtering on the resolution of calibrated thermographic images." Ph.D. dissertation, University of Texas at Austin (1986).
28. Marmolin, Hans. "Subjective MSE measures." *IEEE Trans Syst Man Cybernetics*, v. SMC-16 n. 3 (1986), pp. 486-489.
29. Tukey, J. W. *Exploratory Data Analysis*. Reading, MA: Addison-Wesley, 1971.
30. Rosenfeld, Azriel and Avinash C. Kak. *Digital Picture Processing*. Orlando: Academic, 1982.
31. Borda, Robert P. and James D. Frost, Jr. "Error reduction in small sample averaging through the use of the median rather than the mean." *Electroenceph Clin Neurophysiol*, v. 25 (1968), pp. 391-392.

Vita

Mark Allen Schulze was born in Fort Worth, Texas, on April 24, 1966, the son of Arthur Edward Schulze and Sharon Havemann Schulze. After graduating from Sharpstown High School in Houston, Texas in 1984, he enrolled in Rice University, Houston, Texas. He received the Bachelor of Arts degree with a major in managerial studies and the Bachelor of Science in Electrical Engineering degree, both summa cum laude, from Rice in 1988. He was employed by TRW, Redondo Beach, California, in the summer of 1987 and by National Instruments, Austin, Texas, in 1988 and 1989. The National Science Foundation awarded him a Graduate Fellowship beginning in 1989, and he used the first year of the award to continue the studies he had begun at the Graduate School of the University of Texas at Austin in September, 1988. Mr. Schulze is a member of Phi Beta Kappa, Tau Beta Pi, and Eta Kappa Nu.

This thesis was typed by Mark A. Schulze