

Noise and Degradation Reduction for Signal and Image Processing via Non-
Adaptive Convolution Filtering

Benjamin A. Bjerke

Thesis submitted to the faculty of the Virginia Polytechnic Institute and State University in
partial fulfillment of the requirements for the degree of

Master of Science

In

Mechanical Engineering

Michael J. Roan, Chair

Alexander Leonessa

Cory Papenfuss

7/22/2013

Blacksburg, VA

Keywords: Signal Processing, Image Processing, Synthetic Aperture Radar, Vibration,
Convolution

Noise and Degradation Reduction for Signal and Image Processing via Non-Adaptive Convolution Filtering

Benjamin A. Bjerke

ABSTRACT

Noise and degradation reduction is of significant importance in virtually all systems where these phenomena are present, specifically in the fields of signal and image processing. The effect of image processing on target detection is of significant interest because noise and degradations can greatly reduce the effectiveness of detection algorithms, due to the presence of high intensity noise which is often mistaken as a target. In signal processing, noise in vibration data, or any time-series data, can reduce the accuracy of measurement and can prevent the passing of useful information.

Many filters that have been developed are designed to reduce a single class of noise, such as Wiener and Frost filters. When these filters are applied to types of noise that they were not designed for, the effect of the noise reduction can be greatly reduced. The proposed Two-Stage Non-Adaptive Convolution (TSNAC) filter significantly reduces both additive and multiplicative noise in these two unique systems.

The performance of these filters is compared through several Image Quality (IQ) metrics. It will be shown that the proposed TSNAC filter reduces noise and degradations more effectively in both SAR images and synthetic vibration data than the competing filters. It will show higher IQ scores, greater computational efficiency in target detection, and significant improvement in signal restoration of simulated vibration data.

Acknowledgements

First, I would like to thank my advisor Dr. Michael Roan for his support, both financial and otherwise. Without his guidance and motivation, achieving this degree and completing this thesis would not have been possible. Dr. Roan pointed me in the right direction while still letting me explore things on my own. I could not have asked for a better advisor.

I would also like to thank Dr. Cory Papenfuss and Dr. Alexander Leonessa for their time and input as members of my committee. They provided insight and an additional point of view which was extremely helpful in my research and thesis. They reminded me how different things can look from another perspective and I thank them for that.

I would like to thank my lab mates Mike Milo and Brad Harris for making the work environment both enjoyable and productive. They kept the mood of the office light while still maintaining focus on research and the various tasks at hand.

I would like to thank Jill, Brian, and Morgan for their belief and support. They kept me grounded and reminded me what is really important in life and to not take things too seriously.

Finally, I would like to thank my parents for their financial support and guidance throughout my academic career. Without their stability and example I would not have been able to make it this far. Their belief in me throughout my life has given me the confidence to achieve things in all aspects of my life that otherwise would not have been possible. Thank you.

Table of Contents

Acknowledgements.....	iii
Table of Contents.....	iv
List of Figures.....	vi
List of Tables.....	viii
Nomenclature.....	ix
1 Introduction.....	1
1.1 Motivation.....	1
1.2 Goals.....	2
1.3 Structure.....	3
2 Noise Background.....	4
2.1 Additive Noise.....	4
2.1.1 Gaussian Noise.....	4
2.2 Multiplicative Noise.....	5
2.2.1 Speckle.....	5
3 Synthetic Aperture Radar Background.....	7
3.1 RAR vs. SAR.....	7
3.2 The Doppler Effect.....	9
3.3 Challenges of SAR.....	11
3.3.1 Ground vs. Slant Range.....	11
3.3.2 Shadows.....	11
3.3.3 Speckle.....	11
4 Vibration Data Background.....	13
4.1 Preventative Maintenance.....	13
4.2 Spectrogram Background.....	14
5 Source of Noise.....	17
5.1 Source of Speckle in SAR images.....	17
5.2 Source of Noise in Time-Series data.....	18
6 Noise and Degradation Reduction.....	19
6.1 Moving-Average Filter.....	19
6.2 Weighted Filters.....	21

6.2.1	Gaussian Filtering	21
6.3	Wiener Filtering	23
6.4	The Frost Filter	26
6.5	TSNAC Filtering.....	27
7	Quality Metric Backgrounds.....	31
7.1	Non-Referential Image Quality Background	31
7.2	Distance Regularized Level Set Evolution Background	34
7.3	Rotationally-Invariant Symbolic Histogram Method Background	38
7.3.1	Training Phase.....	38
7.3.2	Testing Phase	41
8	Results.....	42
8.1	Qualitative SAR Results	42
8.2	Quantitative SAR Results	47
8.2.1	NRIQ Results	47
8.2.2	DRLSE Results	48
8.2.3	RISH Results.....	53
8.2.4	Spectrogram Results	55
8.2.5	Computational Efficiency	59
9	Conclusions and Future Work.....	60
	References.....	63

List of Figures

Figure 1. (a) Example of a Gaussian distribution and (b) Gaussian noise added to an image.....	5
Figure 2. Example of a Gamma distribution and (b) an image multiplied by speckle noise.	6
Figure 3. Range Resolution and Swath Width of side-looking SAR.....	9
Figure 4. Azimuth coordinate of side-looking SAR.	10
Figure 5. Example SAR image from CARABAS-II Database	12
Figure 6. Example of a Ball Bearing Assembly.....	13
Figure 7. (a) Uncorrupted time-series data and (b) Spectrogram of uncorrupted time-series data.	15
Figure 8. (a) Corrupted time-series data and (b) Spectrogram of corrupted time-series data.....	15
Figure 9. Example of speckle contained in a SAR image.....	18
Figure 10. 2-D Convolution Example.....	20
Figure 11. (a) Original Image, (b) 5x5 MA filter applied, and (c) 10x10 MA filter applied.	20
Figure 12. (a) Normal/Gaussian Distribution and (b) Weighted Gaussian kernel.	22
Figure 13. (a) Unfiltered speckle region and (b) Gaussian filtered speckle region.	22
Figure 14. Target region with first stage of TSNAC filter overlaid.....	28
Figure 15. Window of second stage of TSNAC filter.....	28
Figure 16. (a) Square MA filter showing non-circular symmetry and (b) Circular MA filter showing circular symmetry.	29
Figure 17. Target region with second stage of TSNAC filter overlaid.....	30
Figure 18. (a) <i>Lena</i> #1, (b) <i>Lena</i> #2, (c) <i>Lena</i> #3, (d) <i>Lena</i> #4, (e) <i>Lena</i> #5, and (f) <i>Lena</i> #6.....	33
Figure 19. Windowing of RISH training phase.	38
Figure 20. Cumulative Density Functions of different SAR targets.....	39
Figure 21. Example of the Markov state dimensionality reduction.....	40
Figure 22. Transition Matrix probability setup.....	40
Figure 23. Transition Matrix example.	41

Figure 24. (a) Raw SAR image, (b) Wiener filtered, (c) Frost filtered, and (d) TSNAC filtered.	43
Figure 25. (a) Raw SAR image high-intensity speckle region, (b) Wiener filtered image, (c) Frost filtered image, and (d) TSNAC filtered image.	44
Figure 26. (a) Raw SAR image low-intensity speckle region, (b) Wiener filtered image, (c) Frost filtered image, and (d) TSNAC filtered image.	45
Figure 27. Raw SAR image target region, (b) Wiener filtered image, (c) Frost filtered image, and (d) TSNAC filtered image.	46
Figure 28. (a) Sample of raw SAR image, (b) Sample of Wiener filtered image, (c) Sample of Frost filtered image, and (d) Sample of TSNAC filtered image.	48
Figure 29. (a) Original SAR image, 0 iterations, (b) 10 iterations, (c) 1235 Iterations, and (d) Final Level-Set Function.	49
Figure 30. (a) Original SAR image, 0 iterations, (b) 494 iterations, (c) 910 Iterations, and (d) Final Level-Set Function (Double-Well).	51
Figure 31. DRLSE Results for (a) Gaussian filter, (b) Wiener filter, (c) Frost filter, and (d) TSNAC filter.	52
Figure 32. RISH Target Detection Results.	54
Figure 33. (a) Spectrogram of uncorrupted time-series data and (b) Spectrogram of time-series data with additive noise.	55
Figure 34. Wiener filtered spectrogram.	56
Figure 35. Frost filtered spectrogram.	56
Figure 36. TSNAC filtered spectrogram.	57
Figure 37. (a) Corrupted signal of interest and (b) Restored signal of interest via TSNAC.	58

List of Tables

Table 1. Relative and Absolute AQI results for <i>Lena</i> images.	34
Table 2. Relative and Absolute AQI Results.	47
Table 3. DRLSE Results.	52
Table 4. Spectrogram Filtering Results.	58
Table 5. Computation time of filters for one SAR image and one corrupted spectrogram.	59

Nomenclature

Variables

a	Shape parameter of Gamma Distribution
b	Scale parameter of Gamma Distribution
β_r	Frequency bandwidth of pulse
c	Speed of light
D	Diffusion rate
$\varepsilon(\phi)$	Energy function
$\varepsilon_{ext}(\phi)$	External energy
$E $	Expected value
$f(x,y)$	Uncorrupted image
f_o	Observed frequency
f_s	Emitted frequency
\hat{F}_s	Fourier transform of estimate of uncorrupted image
$G_s(u,v)$	Fourier transform of original image
H_a	Height of antenna
$H_s(u,v)$	Fourier transform of point spread function
$I(x,y)$	Original image
$I'(x,y)$	Tested image
$\bar{I}(x,y)$	Local image mean
k	Frequency variable
K_1, K_2	Normalizing constants
L_a	Length of antenna
λ	Wavelength of pulse
n	Time variable
η	Incident angle of pulse
$\eta(x,y)$	Noise
p	Potential function
$P[n,k]$	Discrete space-frequency distribution of image
R	Slant range
R_α	Rényi Entropy
$R_p(\phi)$	Distance regularization term
$S_f(u,v)$	Power spectra of uncorrupted signal
$\hat{S}_n(u,v)$	Power spectra of noise
S_{per}	Periodogram estimate of power spectra of uncorrupted signal
σ	Standard deviation
u,v	Fourier coordinates
ϕ	Level-set function
v	Velocity
X	Uncorrupted signal
Y	<i>Corrupted signal</i>

Z

Noise signal

Abbreviations

<i>AQI</i>	Anisotropic Quality Index
<i>CDF</i>	Cumulative Density Function
<i>DRLSE</i>	Distance Regularized Level Set Evolution
<i>FAR</i>	False Alarm Rate
<i>FPR</i>	False Positive Rate
<i>HVS</i>	Human Visual System
<i>IQ</i>	Image Quality
<i>MA</i>	Moving Average
<i>MSE</i>	Mean Square Error
<i>NRIQ</i>	Non-Referential Image Quality
<i>PDF</i>	Probability Density Function
<i>POI</i>	Pixel of Interest
<i>PSF</i>	Point Spread Function
<i>PSNR</i>	Peak Signal to Noise Ratio
<i>RAR</i>	Real Aperture Radar
<i>RISH</i>	Rotationally-Invariant Symbolic Histogram
<i>SAR</i>	Synthetic Aperture Radar
<i>SNR</i>	Signal to Noise Ratio
<i>STFT</i>	Short-Time Fourier Transform
<i>TPR</i>	True Positive Rate
<i>TSNAC</i>	Two-Stage Non-Adaptive Convolution

1 Introduction

Measuring physical phenomena is vital in all fields of science, particularly engineering. Accomplishing this is made difficult by the introduction of noise into a system. Noise can be classified as any (typically) unwanted signal in a system. The two main classes of noise are additive and multiplicative [1]. Additive noise consists of Gaussian noise, Uniform noise, and salt and pepper noise and, for the purposes of this thesis, the main example of multiplicative noise is speckle [2] [3]. Additive noise presents itself in mechanical systems such as in accelerometer data, while multiplicative noise is typically found in signal processing systems such as radar imagery. Both of these topics will be discussed in detail in this thesis.

1.1 Motivation

Noise can distort, block, or remove valuable information from all types of systems. When noise is present in a system, it is difficult to determine what portion of the signal is due to the phenomena of interest and what is due to noise. This often results in inaccurate measurements and representations of these phenomena. In order to prevent this, numerous noise reduction techniques have been, and continue to be developed, which can be used on different types of systems. These methods range in complexity from simple non-adaptive smoothing algorithms to much more complex and computationally intensive, adaptive filters which use locally estimated parameters to preserve wanted signal while removing noise [3] [4]. There are three primary issues with the existing filters used throughout the field. One, for the simple non-adaptive filters in particular; significant denoising can occur, however, it is often at the expense of the signal of interest. That is, either the edges and important detail of an image is not preserved or, for vibration data, the signal of the measured phenomena is degraded. Two, for the more complex adaptive filters, the current algorithms either do not significantly remove noise or do so at great

temporal and computational expense. For several fields, including military surveillance, real-time computation is a necessity, making these methods obsolete. Three, the developed filters are only effective on either additive or multiplicative noise but typically not both. This presents issues when dealing with different types of data. The comparison of these filtering methods and the proposed new method will be described in great detail later.

In most cases, the goal of noise reduction techniques is to increase the Signal to Noise Ratio (SNR) of the image. That is, removing excess noise while preserving the original signal. In image processing, this means removing noise, such as speckle, while still preserving the edges within an image. In terms of vibration data, this means removing noise without disrupting useful signal.

1.2 Goals

The goal of this thesis is to present a novel approach to denoising different types of data, specifically Synthetic Aperture Radar (SAR) images, corrupted by multiplicative speckle, and simulated vibration data, corrupted by additive Gaussian white noise. This will be accomplished by using a single filter called the Two-Stage Non-Adaptive Convolution (TSNAC) Filter. This filter will significantly remove noise while preserving the signal of interest and it will do this in real-time. This filter will demonstrate that significant despeckling and noise reduction is possible without substantial signal loss and computational expense. The filter will also show to be effective on different types of noise. It will be compared to various state-of-the-art filters used today. The proposed filter will display efficient noise reduction in both sets of data while other filters tested do not.

The quality of these filters will be compared both objectively and subjectively, using Distance Regularized Level-Set Evolution (DRLSE) and the Rotationally-Invariant Symbolic Histogram (RISH) methods, along with several error metrics.

1.3 Structure

The structure of this thesis is as follows. Chapter 2 will be an introduction into the different types of noise. Chapter 3 provides a background on Synthetic Aperture Radar and how it works. Chapter 4 is a background on preventative maintenance and the application of filtering on synthetic time-series vibration data. Chapter 5 is an introduction to the sources of noise for each of the different types of data. Chapter 6 is a background on each of the different filters presented in this thesis. Chapter 7 is a background on each of the image quality metrics used to determine the effectiveness of each of the filters. This includes Non-Referential Image Quality metrics, Distance Regularized Level Set Evolution target detection, and Rotationally-Invariant Symbolic Histogram target detection. Chapter 8 presents the results of the filters for each of the image quality metric tests. Finally, Chapter 9 presents conclusions and future research.

2 Noise Background

For the purposes of this thesis, there are two classes of noise, additive and multiplicative. For both cases, noise is any unwanted signal in the data. An example of additive noise studied in this thesis is background noise in vibration data. Multiplicative noise is shown in the form of speckle in coherent SAR images.

2.1 Additive Noise

Additive noise can be broken down into many categories including Gaussian, salt and pepper, and Brownian [2]. This thesis focuses on Gaussian noise.

2.1.1 Gaussian Noise

Gaussian noise is normally distributed noise that is added to an existing signal. The noise and the resulting signal take the following form.

$$Z \sim N(0, \sigma^2) \quad (1)$$


where

$$Y = X + Z \sim N(X, \sigma^2) \quad (2)$$


This shows that the noise is normally distributed with a mean of 0 and a variance of σ^2 . The probability density function (PDF) of a normal distribution is shown in the following equation. [2] [5].

$$p(Z) = \frac{1}{\sqrt{2\pi}\sigma} e^{\left(-\frac{1}{2\sigma^2}(z-0)^2\right)} \quad (3)$$

This equation is plotted in Figure 1(a). When this noise is added to the actual signal, X , the resulting signal, Y , is normally distributed around X with a variance of σ^2 . The resulting signal, Y , when X is a signal of ones and $Z \sim N(0,4)$, is shown in the following figure. Note: The characteristics of the noise (mean and variance) greatly affect the resulting image.



(a)



(b)

Figure 1. (a) Example of a Gaussian distribution and (b) Gaussian noise added to an image.

2.2 Multiplicative Noise

Multiplicative noise is primarily caused by the complex scattering of signals. In SAR images this results in speckle.

2.2.1 Speckle

While Gaussian noise is normally distributed, fully-developed speckle takes the form of a gamma distribution, governed by the following equation [2] [6] [7].


$$Z \sim \text{Gamma}(a, b) \quad (4)$$

where

$$Y = X \cdot Z \quad (5)$$

This equation shows that a gamma distribution's shape is determined by a and its scale by b . The probability density function (PDF) of a Gamma distribution is shown in the following equation and plotted below [5]. The result of a signal of ones multiplied by Gamma distributed noise is also shown below.

$$p(Z) = \frac{b^a}{\Gamma(a)} Z^{a-1} e^{-bZ} \quad (6)$$



(a)



(b)

Figure 2. Example of a Gamma distribution and (b) an image multiplied by speckle noise.

Further details on how speckle noise develops in SAR imagery are described in detail in Section 3.3.3.

3 Synthetic Aperture Radar Background

Synthetic Aperture Radar (SAR) is an imaging system that produces high resolution images while being robust to weather and light conditions [8]. Because of this, SAR has a wide range of military and scientific applications [8]. In this case, resolution is defined as the minimum distance between two objects that can be seen as separate objects in the image [9] [10]. For images with poor resolution, two objects can be seen as a single object which should be avoided in virtually all applications. SAR has more or less replaced Real Aperture Radar (RAR) in these applications due to its superior resolution.

3.1 RAR vs. SAR

RAR uses a straight antenna mounted to an aircraft that is pointed parallel to the flight path. This antenna emits electromagnetic pulses pointed towards the ground and perpendicular to the flight path [9]. These pulses then reflect off the terrain and return to a receiver. The time interval of the returning pulses determines the distance from the antenna to the ground or object. For both RAR and SAR, the range, or slant, resolution (perpendicular to flight path) is determined by the incident angle, η , and the frequency bandwidth of the pulse, B_R [9] [10]. For SAR, pulse bandwidths range from 10 to 40 MHz resulting in a slant resolution of 3.7 to 15 meters [10]. The equation for ground range resolution is shown below [9].

$$\text{Ground Range Resolution} = \frac{c}{2B_R \sin \eta} \quad (7)$$

The azimuth resolution (parallel to flight path) in RAR is determined by the slant range, R , the wavelength, λ , and the length of the antenna, L_a [9] [10] [11]. This is shown in the following equation.

$$\text{Azimuth Resolution} = \frac{\lambda R}{L_a} \quad (8)$$

Another parameter of interest is the swath width which is the perpendicular distance that the transmitted pulses cover [9]. This is dependent on the wavelength, λ , slant range from the observer to the midpoint of the swath, R_m , height of the antenna, H_a , and the inclination angle, η .

$$\text{Swath Width} \approx \frac{\lambda R_m}{H_a \cos \eta} \quad (9)$$

These parameters can be visualized in the figure below. The main issue with RAR is poor azimuth resolution. For RAR, finer ground range resolution can be achieved by increasing the pulse bandwidth and the incident angle. Azimuth resolution, however, cannot be easily improved because changes to the range and wavelength are not very effective from space [10]. The length of the aperture can improve the resolution; however, in order to produce useful images, the length of the antenna would have to be extremely large which is not possible in most applications [9]. An example illustrating this is RADARSAT. RADARSAT is a Canadian satellite that has been in orbit since 1995 [9]. It orbits at 792 km above Earth's surface and operates at 5.3 GHz. With an antenna length of 15m, its azimuth resolution is 3km [9]. This means that if two objects are within 3 km, they will be seen as one object in the resulting image. With SAR, the azimuth resolution becomes 9 m [9]. This is possible because SAR takes advantage of the Doppler Effect to create a much larger, synthetic aperture.




Figure 3. Range Resolution and Swath Width of side-looking SAR.

3.2 The Doppler Effect

The Doppler Effect is an observed shift in frequency caused by the relative motion between an observer and source. When the source approaches the observer, the observed frequency is

$$f_o = \left[\frac{1 + \left(\frac{v}{c}\right)}{1 - \left(\frac{v}{c}\right)} \right]^{\frac{1}{2}} f_s. \quad (10)$$

When the source moves away from the observer, the observed frequency is

$$f_o = \left[\frac{1 - \left(\frac{v}{c}\right)}{1 + \left(\frac{v}{c}\right)} \right]^{\frac{1}{2}} f_s \quad (11)$$

where v is the velocity of the source, c is the speed of light, and f_s is the emitted frequency [9].

This shows that the observed frequency for an incoming force is higher than f_s , and the observed

frequency for a receding force is lower than f_s [9]. This allows for targets at the same range but different azimuth coordinates to be located individually even with very small separation because they have slightly different R -directional velocities relative to the observer, resulting in different frequency shifts. The azimuth coordinate with the frequency shift, f_d is then

$$x = \frac{f_d \lambda R}{2v_{rel}}. \quad (12)$$

Although using the Doppler Effect does not affect the range resolution, it does improve azimuth resolution. The azimuth resolution using the Doppler frequency shift, assuming a constant Doppler shift while the source is in the beam, takes the following form [9].

$$\text{Azimuth Resolution} = \left(\frac{\lambda R}{2v_{rel}} \right) \left(\frac{L_a v_{rel}}{R \lambda} \right) = \frac{L_a}{2} \quad (13)$$

This equation shows that azimuth resolution improves with a decreased antenna length. This is true as long as the antenna is long enough to create proper interference patterns [9]. These parameters can be seen in the figure below.




Figure 4. Azimuth coordinate of side-looking SAR.

3.3 Challenges of SAR

SAR uses the Doppler Effect to greatly increase the azimuth resolution, resulting in much clearer images. However, there are several unique characteristics of SAR imagery that negatively affect picture quality. One class of these negative characteristics is called Geometric Distortion which includes slant vs. ground range issues and shadows among others [9] [10] [12].

3.3.1 Ground vs. Slant Range

The observer emits pulses at a constant rate which are reflected by objects and returned to the observer along the line of sight. However, a constant sampling rate along the range, R , does not correspond to a constant rate along the ground, resulting in variable ground resolution throughout the swath, dependent on an objects angle from the observer [10] [12]. The difference in ground range resolution can be as large as 25% [12]. In order to create a constant resolution, resampling is required [10]. However, this results in objects appearing closer together than they actually are because the observer is viewing from an angle [10].

3.3.2 Shadows

Shadows also affect the quality of SAR images. When an emitted pulse hits a vertical structure, the pulse cannot reach the area directly behind the structure because of the viewing angle [9]. This creates an area where no returns were received, resulting black pixels near vertical objects. Therefore, the area behind the object is not represented in the SAR images.

3.3.3 Speckle

Another class of unique SAR characteristics involves the interaction between the emitted pulses and the ground. These interactions include several different types of scattering, polarization, and speckle [9].

Speckle takes the form of seemingly random white and black pixels throughout a SAR image. It is the result of the constructive and destructive interference of the coherent radar pulses [9] [10]. This occurs because there are often multiple objects within a pixel that interact with the pulse as pixels can correspond to multiple square meters in area [9] [10]. Non-coherent radar does not have this problem because it only needs to detect the total intensity of the returned pulses. Coherent radar is necessary for SAR because amplitude and phase is needed to utilize the Doppler Effect as opposed to just the amplitude. The removal of speckle has and continues to be extensively studied and there are several methods that attempt it that will be discussed later. For this thesis, the CARABAS-II SAR image database is used [13].




Figure 5. Example SAR image from CARABAS-II Database

4 Vibration Data Background

The main goal of this thesis is to develop of filter that can be applied to different types of data. The other type that will be used to verify this is simple time series vibration data. Analyzing vibration data is extremely important because nearly all mechanical systems contain some form of vibration [14] [15] [16] [17]. In some cases vibration is desired and therefore designed into the system, i.e. musical instruments. In other cases, however, vibration can be accidentally introduced into a system by a number of ways including improper fitting or defective components, poor lubrication, etc. These sources of vibration can negatively impact the performance and lifetime of these mechanisms. Because of this, much work is being done in the field of preventative maintenance [17].

4.1 Preventative Maintenance

Preventative maintenance is the repairing or replacing of components before a catastrophic failure of the mechanism occurs. Preventative maintenance can reduce the downtime of machines, repair costs, and possible injury to operators [17].

In rotating machinery, such as rolling-element bearings, preventative maintenance is used in order to detect faults in the bearings before failure occurs. These faults can occur in virtually any component of the bearing including the ball, outer raceway, inner raceway, and the cage. The figure below shows the general structure of a bearing [16] [17].




Figure 6. Example of a Ball Bearing Assembly

These faults can result in premature and catastrophic failure of machinery. The faults result in unique frequencies in the vibration data that can be [15] [16] [17]. However, the introduction of noise into the system can make detection difficult. This is why filtering vibration data is often necessary.

4.2 Spectrogram Background

TSNAC and the other filters discussed in this thesis are all based on the intensity of pixels in an image. In order to use these filters on vibration data, the time series data must be converted into an intensity image. Vibration, and all other time series data, can be represented in this manner via a spectrogram.

A spectrogram is a visual representation of the frequencies and amplitudes of those frequencies displayed over time. A spectrogram is created by taking the Short-time Fourier Transform (STFT) of time series data [18]. This is done by breaking the sample up into overlapping bins via a Hamming window [18]. Each of the bins is then transformed by the following equation in order to find the magnitude of its frequency spectra

$$X(n, \omega) = \sum_{m=-\infty}^{\infty} x[m]w[n-m]e^{-j\omega n} \quad (14)$$

where $x[]$ and $w[]$ are the signal and window respectively [19].

The transformed bins then correspond to a vertical line in the spectrogram which represents the magnitude versus frequency for a specific instant. These lines are then laid side-by-side in order to form the complete spectrogram [18] [20]. The result is a plot of frequency over time with the intensity gradient corresponding to the amplitude of the signal. An example of time series data and its corresponding spectrogram is shown below. The signal being processed here is the following simple sinusoid.

$$f(t) = \sin(2\pi t) \tag{15}$$

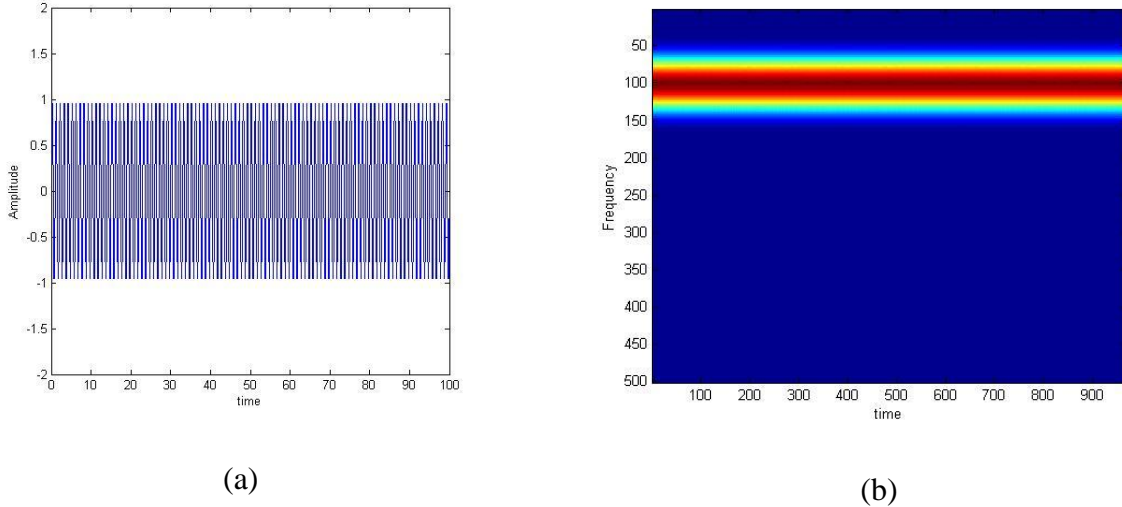


Figure 7. (a) Uncorrupted time-series data and (b) Spectrogram of uncorrupted time-series data.

The spectrogram clearly shows that the main frequency at play is 100 Hz. The color/intensity of the bands corresponds to the magnitude of that frequency. When noise is introduced into the system, such as additive random Gaussian noise, both the time series data and the spectrogram differ significantly from the uncorrupted data.




Figure 8. (a) Corrupted time-series data and (b) Spectrogram of corrupted time-series data.

The effect of the random noise is clear in both the time series data and the spectrogram. Less information about the signal of interest is available when noise is introduced into the system. In order to isolate the signal from the noise, filtering techniques must be used.

The effectiveness of the filters will be judged based on the Mean Square Error (MSE) and Peak Signal to Noise Ratio (PSNR). The MSE calculates the pixel to pixel difference between two images and takes the form

$$MSE = \frac{1}{MN} \sum_{y=1}^M \sum_{x=1}^N [I(x, y) - I'(x, y)]^2 \quad (16)$$

where M and N are the dimensions of the image, I is the original, uncorrupted image, and I' is the image being tested [21]. A low MSE means that the two images are similar and vice versa.

The PSNR is the ratio of the maximum power of a signal and the noise corrupting and is given by

$$PSNR = 20 \log_{10} \left(\frac{255}{\sqrt{MSE}} \right) \quad (17)$$

The higher quality filters will have a low MSE and, therefore, a high PSNR.

5 Source of Noise

As stated earlier, the goal of this thesis is to develop a filtering algorithm that successfully removes noise and other degradations in both image and time-series data. The corruptions in these types of data are distinct and result from very different phenomena.

5.1 Source of Speckle in SAR images

Speckle is a very common occurrence in SAR images. It is the result of the coherent nature of SAR. In RAR, non-coherent radar may be used which only records the amplitude of the returned signal. This does not result in the formation of speckle because RAR uses total intensity detectors which detect the total intensity of the incoming radiation and no phase is recorded [10]. However, SAR images must record the amplitude and the phase of the incoming pulse in order to utilize the Doppler Effect to form a synthetic aperture. As a result, speckle is introduced in the image.

Speckle presents itself as black and white pixels scattered over an image [9]. It is caused by the constructive and destructive interference of multiple scatterers which can result in either extremely strong or weak return signals [9]. These pixels, or collection of pixels, degrade the image and can prevent accurate target detection and classification. Speckle is multiplicative, meaning it is dependent on the strength of the signal and is more difficult to remove than other types of noise [22]. An example of speckle in a SAR images is shown below.




Figure 9. Example of speckle contained in a SAR image.

The removal of speckle in SAR images is of great interest in the field of image processing; particularly in its application of military surveillance. Removing speckle in SAR imagery improves image quality, target detection and recognition, as well as the computational efficiency of these algorithms.

5.2 Source of Noise in Time-Series data

The noise present in many mechanical systems, however, is additive rather than multiplicative. It can be caused by a number of different phenomena. Some of these include, improper fitting of components, insufficient lubrication resulting in friction, sensor interference, and picking up unwanted signals. Regardless of the cause of the noise, these events introduce noise into the system which must be removed via filters in order to process recorded data accurately.

6 Noise and Degradation Reduction

As stated earlier, there are numerous algorithms that significantly reduce speckle. However, while reducing the signals resulting from speckle, they also reduce the signal of interest. In terms of target detection, these algorithms will have a low false positive rate (FPR) but will also have a low true positive rate (TPR). Ideally, a filter would result in target detection with a high TPR and a low FPR.

These filters are called non-adaptive filters because they treat the entire image the same. There are other, more complex techniques that determine the amount of filtering based on the statistics of a specified area of an image and are called adaptive filters. These methods typically perform better.

This thesis will focus on three image filters. Two are adaptive filters. The third is the proposed TSNAC Filter. Each of these filters will be described in detail in the following sections.

6.1 Moving-Average Filter

The moving-average (MA) filter is the simplest type of filter discussed in this thesis. It replaces the value of the pixel of interest (POI) with the average of a window of neighboring pixel values using convolution [23] [24] [25] [26] [27]. The window shape and size selection is discussed later. A 2-D convolution example using a 3x3 square average filter is shown below. The input matrix (left) is convolved with the 3x3 averaging window (middle) and the result for the POI is shown in the output matrix (right).




Figure 10. 2-D Convolution Example.

Moving-average filters are considered linear filters because the output image is a linear combination of the input image [26]. Moving-average filters are also classified as smoothing filters because they remove noise by smoothing/blurring the image. These filters do significantly reduce noise; however, it can also blur edges and other elements of interest if care is not taken. The degree of smoothing is dependent on the size and shape of the neighborhood. An example of square moving-average filter applied to an image with different sized windows is shown below.




Figure 11. (a) Original Image, (b) 5x5 MA filter applied, and (c) 10x10 MA filter applied.

Common alternative shapes of moving-average filters are rectangular and circular (disk/pillbox) [23]. Rectangular filters have a greater smoothing effect in the direction of the longer side. That is, if a filter has a height of 5 pixels and a length of 10 pixels, the resulting image will be smoothed more horizontally than vertically.


Circular filters have the unique feature of being circularly symmetric [23]. This means that pixels that are the same distance horizontally, vertically, and diagonally will have the same smoothing effect on the POI. This concept will be discussed in greater detail later.

6.2 Weighted Filters

The MA filters discussed thus far have all weighted the pixels within the window evenly. That is, a pixel that is a distance of 1 from the POI has the same smoothing effect as a pixel a distance of 2 away from the POI. This does not always have to be the case. The most common method for weighting pixels within the window is called Gaussian Filtering.

6.2.1 Gaussian Filtering

The Gaussian filter is the same as the moving-average filter with the exception that it does not weight each pixel in the window evenly. Instead, it weights them according to a Gaussian, or Normal distribution shown below [23] [28]. The corresponding convolution window is also shown below.




(a)

0.011	0.084	0.011
0.084	0.619	0.084
0.011	0.084	0.011


(b)

Figure 12. (a) Normal/Gaussian Distribution and (b) Weighted Gaussian kernel.

This filter does not remove noise nearly as well as an evenly weighted MA filter but, generally, it does preserve edges well. Take a SAR image for instance. If the POI is on a speckle region, the high level speckle be weighted the greatest in determining the output pixel value. That is, if the input POI is high then the output POI will also be relatively high which does not reduce speckle. An example of a Gaussian filter applied to a SAR speckle region is shown below. For this same reason, however, edges contained in images are still preserved which is ideal.



(a)



(b)

Figure 13. (a) Unfiltered speckle region and (b) Gaussian filtered speckle region.

6.3 Wiener Filtering

As stated earlier, the Wiener filter is an adaptive filter. This means that the filter changes based on the statistics calculated in the current window. The MA filter is an example of a non-adaptive filter because it remains constant throughout the entire image or dataset.

The Wiener filter is an image restoration technique that aims to recover an original, uncorrupted image that has been corrupted in some way based on a priori knowledge of the degradation [24]. This corruption can be due to either a degradation function, noise, or both [4] [24] [29] [30]. The degradation function, otherwise known as the Point Spread Function (PSF) is essentially how an imaging system inherently spreads a point of light. PSFs can be estimated by image observation, experimentation, or modeling [24].

Estimation by image observation is done by choosing a subimage within the corrupted image that has a high SNR. This sub image is then unblurred as much as possible through a variety of methods. The Fourier transform of the original subimage and the unblurred subimage are found and then divided and that quotient is the Fourier transform of the PSF. This is shown in the equation below

$$H_s(u,v) = \frac{G_s(u,v)}{\hat{F}_s(u,v)} \quad (18)$$

where $H_s(u,v)$ is the Fourier transform of the PSF of the subimage, $G_s(u,v)$ is the Fourier transform of the original subimage, and $\hat{F}_s(u,v)$ is the Fourier transform of the unblurred subimage [24]. Assuming that the degradation function is position invariant, the PSF of the subimage can be extended to the entire image.

Estimation by experimentation is done by using the same imaging system and capturing an impulse of light because a linear, spatially-invariant system is characterized entirely by its impulse response [24]. Based on the resulting image, the PSF can be approximated.

Estimation by modeling is the most common method for PSF estimation. This method can include anything from modeling atmospheric turbulence to linear motion blur. For SAR images in particular, the atmospheric turbulence model is often used and takes the following form [24].

$$H(u, v) = e^{-k(u^2+v^2)^{5/6}} \quad (19)$$

The Wiener filter assumes an image of the following form

$$g(x, y) = H[f(x, y)] + \eta(x, y) \quad (20)$$

where $g(x, y)$ is the original image, $H[]$ is the degradation function, $f(x, y)$ is the uncorrupted image, and $\eta(x, y)$ is the noise [24]. It is easy to see that the Wiener filter is designed for additive noise. Despite this, it is still widely used for denoising SAR images with multiplicative noise [31] [32] [33]. The Wiener Filter is a type of inverse filter. That is, it attempts to deconvolve the uncorrupted image and the degradation function. However, standard inverse filters focus solely on deconvolution and make no effort in removing noise.

The Wiener filter attempts to minimize the mean square error (MSE) between the input image, f and its theoretical, uncorrupted image, \hat{f} . This relationship takes the following form [24].

$$e^2 = E\{(f - \hat{f})^2\} \quad (21)$$

Assuming that either the noise or the uncorrupted image has zero mean and that the intensity of \hat{f} is a linear function of the intensity in f , the minimum of this error function in the frequency domain is given by

$$\hat{F}(u, v) = \left[\frac{1}{H(u, v) \frac{|H(u, v)|^2}{|H(u, v)|^2 + S_\eta(u, v) / S_f(u, v)}} \right] G(u, v) \quad (22)$$

where $|H(u, v)|^2 = H^*(u, v)H(u, v)$, $S_\eta(u, v) = |N(u, v)|^2$ which is the power spectrum of the noise, and $S_f(u, v) = |F(u, v)|^2$ which is the power spectrum of the uncorrupted image (Derivation on pp. 354-355 Lin) [24] [29] [30] [31] [34] [35].

Determining these power spectra require some a priori knowledge. In the case of additive Gaussian white noise, the power spectra, $S_\eta(u, v)$ is simply the variance of the noise [29] [30] [35]. If it is not Gaussian noise, however, there are other methods of determining it. One is to take a homogenous region of the image, assuming that the signal is low in this area and calculating the variance of that region.

Estimating the power spectra for the signal is more complex. The most widely used method is the periodogram estimate [29] [35]. This takes the form

$$S_f = \frac{S_{per} - S_\eta}{|H|^2} \quad (23)$$

where

$$S_{per} = \frac{1}{N^2} [Y(k, l)Y(k, l)^*] \quad (24)$$

where N is the number of pixels in the observation window and $Y(k, l)$ is the Fourier transform of the observation [29]. If estimating these power spectra is not possible or desirable, an estimation for Equation (22) is often used and takes the form

$$\hat{F}(u, v) = \left[\frac{1}{H(u, v) \left[|H(u, v)|^2 + K \right]} \right] G(u, v) \quad (25)$$

where K is a user defined constant [24].

The adaptability of the Wiener Filter comes from the estimation of the power spectra. In the non-adaptive Wiener filter, the noise and signal power spectra are calculated and assumed constant for the entire image. For the adaptive filter, however, these power spectra are changed based on the statistics of the current neighborhood [35]. The result is that in neighborhoods where the variance is low the filtering effect is high and in areas where the variance is high, i.e. target regions, the effect of filtering is minimal [29] [30] [35].

6.4 The Frost Filter

The Frost filter is actually a class of Wiener filter that has been modified to reduce multiplicative noise [36]. The Frost filter assumes an input image of the form

$$g(x, y) = H[f(x, y)\eta(x, y)] \quad (26)$$

where $H[]$ is the degradation function, $f()$ is the uncorrupted image, and $\eta[]$ is the multiplicative noise [3]. Like the Wiener filter, it is also based on minimizing the MSE between an uncorrupted image and an image corrupted by a degradation function and noise. The impulse response of the Frost filter is the following

$$m'(t) = K_1 \alpha e^{(-\alpha|t|)} \quad (27)$$

where

$$\alpha = K_2 \left(\frac{\sigma_i^2}{\bar{I}^2} \right) \quad (28)$$

where K_1 and K_2 are normalizing constants, t is the pixel being filtered, and σ and \bar{I} are the local standard deviation and local mean [3] [36] [37]. The Frost filter can also be adaptive in that the

degree of filtering is dependent on the local variance. In areas with high variance, little or no filtering is done. In areas of low variance, significant filtering is done.

6.5 TSNAC Filtering

The proposed filter in this thesis is a Two-State Non-Adaptive Convolution (TSNAC) filter. This filter is a combination of two average convolution filters. The first is a square, evenly-weighted average filter, similar to the one described earlier in this section. The main difference is in the selection of the neighborhood size. As stated earlier, a typical neighborhood size for a moving average filter is 3×3 . This is often the default selection because in most cases it provides reasonable smoothing while still maintaining most of the image detail. However, this has shown to not always be the case with more complex images such as those from SAR, especially when they are being used for target detection. When this is the case, particular attention must be paid to the size of the targets.

In target detection, maintaining the intensity of the targets while simultaneously reducing noise is of utmost importance. If too much smoothing occurs, target detection algorithms will not pick up the targets, resulting in a low TPR. If there is too little smoothing, the clusters of speckle will be classified as targets, resulting in a high FPR. Therefore, in order to keep the signal of targets strong while still significantly decreasing noise, the size of the neighborhood must be chosen based on the size of the target.

Through much experimentation, the optimal size of the window is roughly equal to the size of the average target region. This is ideal because, in general, the size of the targets relative to speckle regions is fairly large. This means that the target regions will stay relatively the same intensity while the intensity of the speckle regions, however, will be greatly reduced because of their small size relative to the window. This concept will be shown in greater detail later. This

first stage provides baseline filtering that maintains target intensity while reducing image noise. For the CARABAS-II dataset, the target regions are essentially two 5x5 square regions. Because of this, the size of the square average filter was chosen to be 5x5 pixels. This is shown in the figure below.




Figure 14. Target region with first stage of TSNAC filter overlaid.

The second stage of the filter is also an average filter; however, it has a disk or circular window rather than a square. An example of this window is shown below.

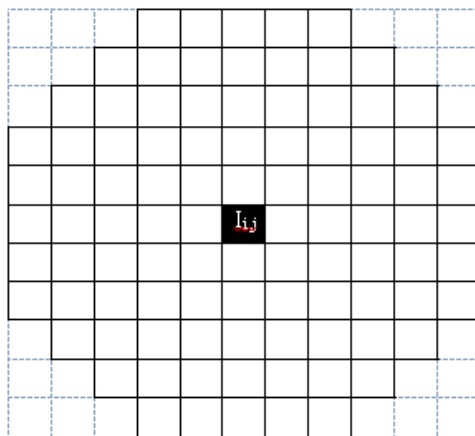


Figure 15. Window of second stage of TSNAC filter.

This disk filter still evenly weights each pixel in the window. The difference between it and the first stage square average filter is that it is circularly symmetric [23]. This means that all pixels the same distance from the POI will all be equally involved in determining the new value of the POI. That is, pixels the same distance from the POI horizontally, vertically, and diagonally will all affect the new value. This differs from the square average filter where not all equidistant pixels affect the POI. This difference is displayed in the figure below. Notice that for a square filter, a pixel a distance, D , away from the POI diagonally is contained in the window while pixels a distance, D , from the POI vertically and horizontally are not. For the circular window, all pixels a distance, D , away from the POI will be included in the window provided D is less than the radius of window.




Figure 16. (a) Square MA filter showing non-circular symmetry and (b) Circular MA filter showing circular symmetry.

Similar to the first stage, the size of the second stage filter is also dependent on the size of the targets. In this case, however, the window contains the entire target region, not just one of the 5x5 regions. This filter ensures that the entire target region, being both 5x5 regions, has relatively the same intensity. Because it has a circular window and is therefore circularly symmetric, virtually all pixels in the target region affect the intensity of the POI when it is in the center of the target. Also, this larger window greatly increases the smoothing of the filter, added on to the smoothing done by the first stage. An example of the window overlaid onto a target region is shown below.




Figure 17. Target region with second stage of TSNAC filter overlaid.

The TSNAC filter provides significant smoothing while maintaining target detail. It performs as well or better in a variety of tests than more complicated and computationally expensive adaptive filters. Details on the results of these tests will be shown in Section 8.2.

7 Quality Metric Backgrounds

As stated earlier, the quality of the filters discussed in this thesis are based on three metrics: The Non-Referential Image Quality via the Anisotropy Quality Index and the target detection capabilities via the Distance Regularized Level Set Evolution (DRLSE) and the Rotationally-Invariant Symbolic Histogram (RISH) methods.

7.1 Non-Referential Image Quality Background

The quality of many different types of filters and images, including SAR images, is based on a comparison between an uncorrupted image and a modified or filtered image. Examples of these metrics are the Mean Square Error (MSE) and the Peak Signal to Noise Ratio (PSNR). However, uncorrupted images are not always available in practice such as in the CARRABAS-II database. This presented the need for a Non-Referential (NR) Image Quality (IQ) measure that does not require an uncorrupted image.

The Anisotropy Quality Index (AQI) method measures the anisotropy of an image and has been shown to provide similar results to the Human Visual System (HVS) method which models how humans perceive images and their quality [38]. Because of this, images can be ranked based on their anisotropy from best to worst quality.

Anisotropy is essentially the property of directional-dependence. Through experimentation, it has been shown that degradation (blurring and noise) depreciates an images directional dependence and therefore its anisotropy [38]. So images with less degradation will have a higher anisotropy than those with more degradation.

The process of calculating the anisotropy of an image is fairly complex. It begins by calculating the local-directional entropy which is accomplished via the Rényi entropy [38]. The

entropy of an image is essentially the amount of information it contains or the amount of information per symbol [38] [39]. The Rényi entropy in space-frequency domain has the form

$$R_\alpha = \frac{1}{1-\alpha} \log_2 \left(\sum_n \sum_k P^\alpha[n, k] \right) \quad (29)$$

where $P[n, k]$ is discrete space-frequency distribution of the image [38]. This can be shown on a point wise basis where $\alpha = 3$ by

$$\check{R}_3 = \frac{1}{2} \log_2 \left(\sum_k \check{P}^3[n, k] \right). \quad (30)$$

Gabarda used a Pseudo-Wigner Distribution (PWD) to model $\check{P}[n, k]$ which takes the form

$$W_z[n, k] = 2 \sum_{m=-N/2}^{N/2-1} z[n+m] z^*[n-m] e^{-2i(2\pi m/N)k} \quad (31)$$

where n and k are the time and frequency variables, respectively, m is the shifting parameter, and $z[n]$ is a 1-D sequence of image data [38]. This window can be shifted to cover the entire image to develop the full PWD.

The full image PWD must then be normalized to create of probability distribution from the existing frequency distribution in order to preserve the property

$$\sum_n \sum_k \check{P}[n, k] = 1 \quad (32)$$

There are several options for normalization but Quantum Normalization is used in Gabarda. This normalization is shown in the following equations.

$$\tilde{P}[n,k] = \frac{Q[n,k]}{\sum_k Q[n,k]} \quad (33)$$

where

$$Q[n,k] = P[n,k]P^*[n,k] \quad (34)$$

Equation (33) can then be plugged back into Equation (30) to get the probability distribution of the entropy. The standard deviation of this distribution is the AQI which measures image quality. The AQI can also be calculated relative to the highest quality image in the set. This provides a normalized method for image quality comparison. Any example of the relative and absolute AQI comparisons for six images is shown below using the standard *Lena* test image [24]. In the example, it is clear that the higher quality images have a higher relative and absolute AQI compared to their lower quality counterparts.

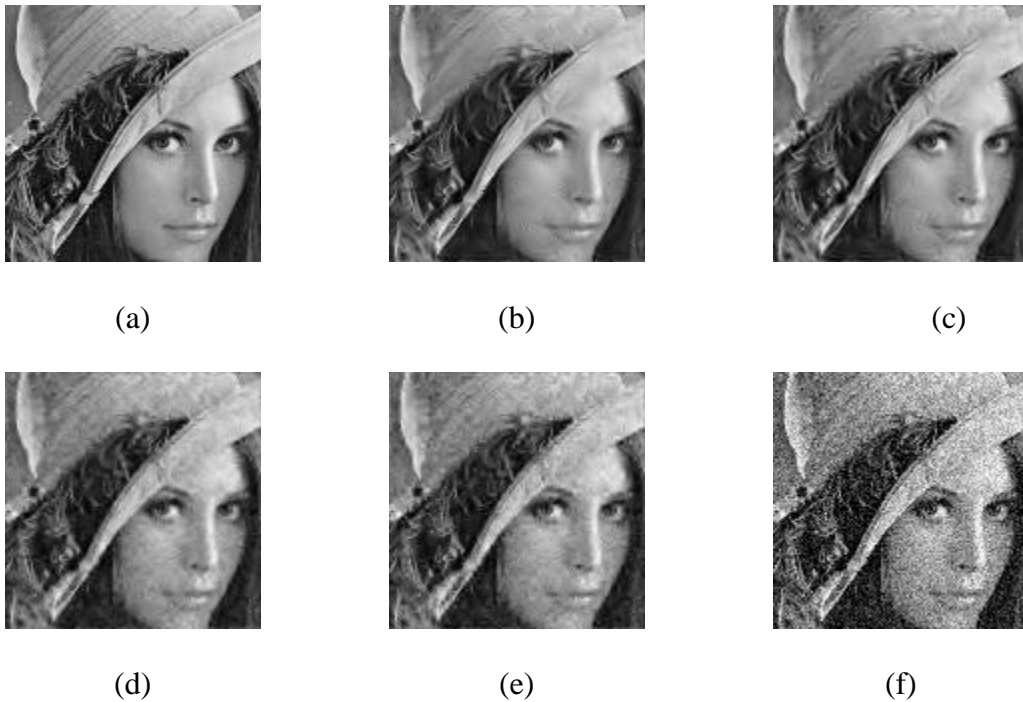


Figure 18. (a) *Lena* #1, (b) *Lena* #2, (c) *Lena* #3, (d) *Lena* #4, (e) *Lena* #5, and (f) *Lena* #6.

Table 1. Relative and Absolute AQI results for *Lena* images.

	AQI relative	AQI absolute
1	1	0.042737
2	0.82	0.035076
3	0.80	0.034009
4	0.71	0.030238
5	0.70	0.030022
6	0.51	0.02183

7.2 Distance Regularized Level Set Evolution Background

Level set methods are commonly used in image processing and segmentation. These techniques can be very useful in detecting the edges of objects both accurately and quickly in simple images, i.e., images with little to no noise. The performance of level set methods do, however, decrease with more complex images, i.e., images with a lower SNR. When level-set methods are applied to these images, irregularities in the level set function evolution can occur which can affect the accuracy of edge/target detection. In the past, a process called reinitialization has been used which replaces the existing function with a signed distance function [40]. Reinitialization has been shown to be effective to an extent, however, it is often unclear when and how reinitialization should be performed as well as its effect on accuracy and stability. To remove the need for reinitialization, a technique called Distance Regularization was developed [40].

Distance Regularized Level Set Evolution (DRLSE) differs from conventional level-set method in that it does not require reinitialization. In order to preserve the integrity of the level-set function, a distance regularization term is included in an energy functional which governs the evolution of the level set function. The energy function is defined as

$$\varepsilon(\phi) = \mu R_p(\phi) + \varepsilon_{ext}(\phi) \quad (35)$$

where $\mu > 1$ is a constant, $R_p(\phi)$ is the distance regularization term, $\varepsilon_{ext}(\phi)$ is an external energy term, and ϕ is the level set function. The distance regularization term is defined as

$$R_p(\phi) \equiv \int_{\Omega} p(|\nabla \phi|) dx \quad (36)$$

where p is a potential function.

The potential function, p serves two purposes. One is to smooth the level set function, ϕ . This is necessary because the level set function must be stable, particularly near the zero level to ensure accuracy. Another purpose of the potential function is to make certain that the signed distance property $|\nabla \phi| = 1$ is satisfied. Li uses both single-well and double-well potential functions, i.e. potential functions with one and two minimums respectively [40]. In Li, a single-well potential function with a unique minimum value of 1 is used and is shown below which satisfies the signed distance property.

$$p = p_1(s) \equiv \frac{1}{2}(s-1)^2 \quad (37)$$

However, a single-well potential function can cause problems that will be described later in this section. Instead, a potential function that maintains the signed distance property only near the zero level set is used. That is, $|\nabla \phi| = 0$ away from the zero level and $|\nabla \phi| = 1$ near the zero level [40]. In order to accomplish this, a potential function with a minimum at 0 and 1 is chosen. This is called a double-well potential. Both single and double-well potentials have their uses. For example, single-well potentials are very useful for region-based models while double-well potentials are good for both edge and region-based models.

The derivation of the level set evolution equation for both the single and double-well potential is based on the solution of the gradient flow equation

$$\frac{\partial \phi}{\partial t} = - \frac{\partial F}{\partial \phi} \quad (38)$$

where $F(\phi)$ is an energy functional [40]. $\frac{\partial F}{\partial \phi}$ is the Gateaux derivative of the function $F(\phi)$.

Taking the partial derivative of Equation (35) with respect to ϕ yields

$$\frac{\partial \varepsilon}{\partial \phi} = \mu \frac{\partial R_p}{\partial \phi} + \frac{\partial \varepsilon_{ext}}{\partial \phi}. \quad (39)$$

Plugging this into the gradient flow Equation (38) yields

$$\frac{\partial \phi}{\partial t} = -\mu \frac{\partial R_p}{\partial \phi} - \frac{\partial \varepsilon_{ext}}{\partial \phi}. \quad (40)$$

The Gateaux derivative of $R_p(\phi)$ in Equation (36) can be found to be

$$\frac{\partial R_p}{\partial \phi} = -div(d_p(|\nabla \phi|)\nabla \phi) \quad (41)$$

where

$$d_p(s) \equiv \frac{p'(s)}{s} \quad (42)$$

Combining Equation (40) and Equation (41) gives

$$\frac{\partial \phi}{\partial t} = -\mu div(d_p(|\nabla \phi|)\nabla \phi) - \frac{\partial \varepsilon_{ext}}{\partial \phi} \quad (43)$$

which is the level set evolution equation. This equation governs the motion of the level set function as it approaches the final zero level. If the external energy functional is ignored, the level set evolution equation can be rewritten as a standard diffusion equation

$$\frac{\partial \phi}{\partial t} = div(D\nabla \phi) \quad (44)$$

where the diffusion rate D is

$$D = \mu d_p (|\nabla \phi|). \quad (45)$$

Based on this diffusion equation, it can be seen that if $|\nabla \phi| > 1$, then D is positive and the diffusion is forward. If $|\nabla \phi| < 1$, then D is negative and the diffusion is backward. This is called forward-and-backward diffusion (FAB). This forces the signed distance property to be satisfied. However, the potential function in Equation (37) results in a diffusion rate that approaches negative infinity as $|\nabla \phi|$ approaches 0. This can result in instability of the level set function which affects accuracy. In order to avoid this issue, a double-well potential can be used.

For double-well potential, the potential function

$$p_2(s) = \begin{cases} \frac{1}{(2\pi)^2} (1 - \cos(2\pi s)), & s \leq 1 \\ \frac{1}{2} (s-1)^2, & s \geq 1 \end{cases} \quad (46)$$

is used because it has two minimum points at $s = 0$ and $s = 1$ [40]. Using a similar process as before, it can be shown that if $|\nabla \phi| > 1$ then the diffusion rate is positive and the diffusion is forward. If $\frac{1}{2} < |\nabla \phi| < 1$ then the diffusion rate is negative and the diffusion is backward. If $|\nabla \phi| < \frac{1}{2}$ then the rate is positive and the diffusion is forward. This potential function results in a bounded diffusion rate which does not have the same negative effects as the unbounded rate for the single-well potential. Double-well potential functions are typically used for objects that change over time. Typical examples used are inflation and deflation of objects. Using a double-well potential function for objects that are time-invariant can result in problems that will be shown later. For time-invariant objects, single-well potentials are typically used instead.

7.3 Rotationally-Invariant Symbolic Histogram Method Background

Rotationally-Invariant Symbolic Histogram (RISH) anomaly detection is a feature extraction method made specifically for SAR images. It attempts to classify regions in the image as either a target or a non-target based on Hidden Markov Models (HMM) and Transition Matrix (TM) comparisons. The RISH algorithm employs a supervised learning method that requires training and test phases.

7.3.1 Training Phase

The training phase begins by creating windows around known target locations. The window size is based on the size of the target. In the case of the CARRABAS-II images, this is a 15x15 square-pixel window. This was chosen because it is large enough to contain the entire target region but not so large as to include clutter in the window. An example of the windowing is shown in the figure below.




Figure 19. Windowing of RISH training phase.

Next, the Cumulative Density Function (CDF) is derived based on the pixel intensities included in the window for each target. These CDFs are then used to partition the intensity values into bins. This is done by taking the second derivative of the function and creating the boundaries of the bins at the inflection points. The inflection points can clearly be seen below.




Figure 20. Cumulative Density Functions of different SAR targets.

As seen in the figure, there are typically two inflection points in the CDF, resulting in 3 bins. This means that all pixels with an intensity of less than ~60 are assigned one symbol. A different symbol is used for pixel intensities between ~60 and ~200 and greater than ~200. Two bin partitions were also used where the only partition is at the lower inflection point. These partitions result in each pixel value being represented by one of either 2 or 3 possible symbols.

For each 15x15 window, 3x3 sub windows are formed and the number of symbols in each window defines a Markov state. In the example shown below, State 1 would be 5 A's, 2 B's, and 2 C's. It is important to note that only the number of each symbol in the window defines the Markov state, not the arrangement. Because of this, the dimensionality is greatly reduced.

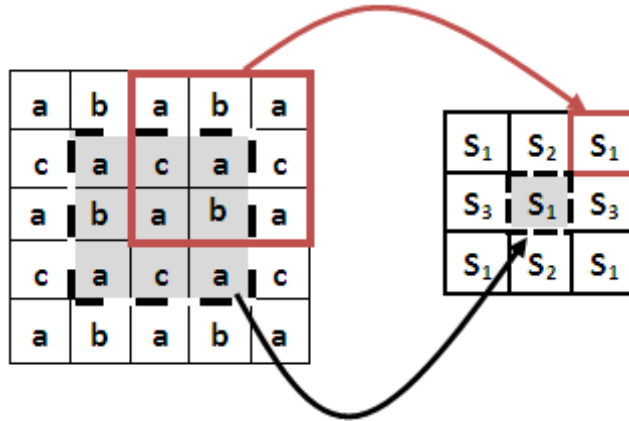


Figure 21. Example of the Markov state dimensionality reduction.

The Transition Matrices are then formed based on the probabilities of one state given the previous state. Transition matrices show the probability of going to one state given the previous state. They are size $n \times n$ where n is the number of possible states. An example of the TM setup is shown below.

	S_1	S_2	...	S_n
S_1	$S_1 S_1$	$S_2 S_1$...	$S_n S_1$
S_2	$S_1 S_2$	$S_2 S_2$...	$S_n S_2$
...
S_n	$S_1 S_n$	$S_2 S_n$...	$S_n S_n$

Figure 22. Transition Matrix probability setup.

To calculate the values in the TM, it is necessary to look at the arrangement of Markov states such as the one shown in Figure 21. For example, take S_1 in the top left corner. It is surrounded by S_2 to the right, S_1 diagonally, and S_3 below. In the empty TM, a 1 is added to the $S_2|S_1$ cell, the $S_1|S_1$ cell, and the $S_3|S_1$ cell. The result of this process is shown below. The next

step would be to complete the same process for S_2 located in position (1,2) and so on for each of the state transformations in Figure 21.

	S_1	S_2	S_3
S_1	1	1	1
S_2			
S_3			

Figure 23. Transition Matrix example.

The probabilities are obtained by normalizing each row so that the sum of the values in each row is 1, creating a row stochastic matrix. This means that given state S_1 , the subsequent state must be either S_1 , S_2 , or S_3 . This entire process is repeated for a speckle region.

7.3.2 Testing Phase

For the testing phase, the same process is completed except for, instead of choosing a region; RISH scans the entire image and forms TMs for each 15x15 subimage. Once these TMs are calculated, they are compared to the target and speckle TMs formed in the training phase. This comparison is done using the L_1 -Norm which is the sum of the absolute value of the differences between each entry in the two TMs. Based on these distances, the testing subimage is classified as either a target or a speckle region.


A comparison of RISH results for unfiltered SAR images and TSNAC filtered SAR images will be shown in later sections.

8 Results


The results for each of the tests described in the previous chapter are discussed in the following sections. The filters will first be applied to SAR images from the CARRABAS-II dataset. The results will be compared qualitatively by examination of low-intensity speckle regions, high-intensity speckle regions, and target regions. Quantitatively, the filters will be compared based on NRIQ Tests, and the target detection capabilities using DRLSE and RISH, as well as computational efficiency. The filters will then be applied to the spectrogram of synthetic vibration data and will be compared via the MSE and PSNR.

8.1 Qualitative SAR Results


The full-scene despeckled images are shown below. Qualitatively, it is difficult to decipher any differences between the images based on the complete scene.




(a)



(b)



(c)



(d)

Figure 24. (a) Raw SAR image, (b) Wiener filtered, (c) Frost filtered, and (d) TSNAC filtered.

However, when these images are zoomed in on various speckle and target regions, the differences become very apparent. The first set of images is of a high intensity speckle region.




Figure 25. (a) Raw SAR image high-intensity speckle region, (b) Wiener filtered image, (c) Frost filtered image, and (d) TSNAC filtered image.

Subjectively, it is clear that the TSNAC filter performs the most despeckling in the case of high intensity speckle. The Frost filter did smooth the noise to an extent, however, not nearly to the degree that TSNAC did. The Wiener filter performed almost insignificant despeckling to the speckle region and surrounding area. Like the Frost filter, it is designed to reduce its smoothing effect in areas of high variance.

In areas of low intensity speckle, similar results were found.




Figure 26. (a) Raw SAR image low-intensity speckle region, (b) Wiener filtered image, (c) Frost filtered image, and (d) TSNAC filtered image.

Again it is clear, subjectively, that TSNAC significantly outperforms the Frost and Wiener filters. The Wiener filter performed slightly better in the low intensity speckle region than the high intensity region due to the lower variance in the area surrounding the speckle.

The filtered results for the target region are more difficult to compare. The Wiener filter resulted in no filtering near the target region; completely preserving the signal. The Frost filter did significantly filter near the target region; however, the signal is still relatively strong. The TSNAC filter did smooth the target region the most and resulted in a slightly weaker signal.




Figure 27. Raw SAR image target region, (b) Wiener filtered image, (c) Frost filtered image, and (d) TSNAC filtered image.

The degree of smoothing of the target region is not the main issue. The main issue is difference in the intensity of the resulting target regions, relative to the intensity of the filtered speckle regions. The best filter will have the greatest difference between the intensities of the resulting target and speckle regions. The Wiener filter, for instance, preserved the signal in the target region the best. However, it also preserved the most signal due to noise in the speckle regions. Because of this, it is clear that the Frost and TSNAC filters are superior.

The Frost filter did preserve the signal of the target regions better than TSNAC. However, it also did not smooth the speckle regions nearly as well. Because of this, choosing a filter based on subjective analysis alone is difficult.

8.2 Quantitative SAR Results

The effectiveness of the filters will be judged based on AQI scores and target detection performance via DRLSE and RISH. The computational efficiency of each algorithm will also be noted.


8.2.1 NRIQ Results

The AQI was calculated for each of the filtered images. The relative and absolute results are shown in the table below. As stated earlier, the AQI is equal to the standard deviation of the Rényi entropy. A large standard deviation equates to the image having greater directional dependence and therefore greater anisotropy and image quality.


Table 2. Relative and Absolute AQI Results.

	AQI (relative)	AQI (absolute)
Unfiltered	0.30439	0.0018328
Wiener	0.62895	0.0032522
Frost	0.9262	0.0047893
TSNAC	1	0.0051709


The results of the AQI tests reinforce the claims already made. The Wiener filter did remove some speckle and even out the grey-level but it did not remove speckle as significantly as the other filters, resulting in the relatively low AQI. The Frost filter did remove speckle, however, not to the same degree as TSNAC. TSNAC greatly reduced speckle and, therefore, the degradation of the image which resulted in a high AQI. The unfiltered image predictably had the lowest AQI. The resulting despeckling and gray-levels can be seen in the figure below.




(a)



(b)



(c)



(d)

Figure 28. (a) Sample of raw SAR image, (b) Sample of Wiener filtered image, (c) Sample of Frost filtered image, and (d) Sample of TSNAC filtered image.

8.2.2 DRLSE Results

DRLSE was applied to a section of a SAR image containing 25 targets in a 5x5 grid using TSNAC and the other filters discussed in this thesis. The effectiveness of the edge detection was based on the number of targets correctly detected versus the number of falsely detected targets, i.e., TPR vs. FPR. The computational efficiency of each test is also discussed. Single and double-well potential functions were tested and compared as well.

First, DRLSE using the original Gaussian filter was applied to the original SAR image using the single-well potential function. The initial level set function contour is a rectangular region containing all of the targets (b). As the evolution begins, the region contracts and forms around each of the exterior targets until it eventually breaks off. The evolution continues until each of the targets is surrounded by its own region (c). The final level set function is shown in (d). Using the same computer as earlier, the CPU time using the single-well potential function was 72.975 seconds.






Figure 29. (a) Original SAR image, 0 iterations, (b) 10 iterations, (c) 1235 Iterations, and (d) Final Level-Set Function.


DRLSE applied to SAR images using the double-well potential function did not perform as well. Again, the initial level set function is a rectangular region that encompasses all of the targets (a). The evolution begins with the contraction of the region until it forms around the individual exterior targets (b). The evolution continues until the interior targets are detected, however, the dimmer exterior targets are no longer detected (c). The final level set function is shown in (d). This also shows the effects of the double-well potential, showing that $|\nabla\phi| = 0$ away from the zero level (bottom flat region) and, though it is difficult to see, $|\nabla\phi| = 1$ near the zero level. The signed distance property can also be seen in the SDB. Using the same parameters as the single-well potential except fewer iterations (910 vs. 1235), 14 of the 25 targets were detected with zero false positives in 80.00 seconds. This decrease in detection is due to the nature of these particular level set evolution equations. Because the images do not evolve over time, a single-well potential function is used.




(a)



(b)



(c)



(d)

Figure 30. (a) Original SAR image, 0 iterations, (b) 494 iterations, (c) 910 Iterations, and (d) Final Level-Set Function (Double-Well).

Because of the original double-well potential results, the various filtering techniques were compared only using the single-well potential.

The results are outlined in the table below. *True Positive Rate (TPR)* is defined as the percentage of targets correctly identified. *Errors* is defined as the number of false positives per image.




Figure 31. DRLSE Results for (a) Gaussian filter, (b) Wiener filter, (c) Frost filter, and (d) TSNAC filter.

Table 3. DRLSE Results.

	Computation Time	Iterations	TPR	Errors
Gaussian	72.975	1235	100	0
Wiener	111.749	2510	100	7
Frost	74.089	1610	100	6
TSNAC	41.348	794	100	0

The results show TSNAC as the most effective and efficient filter tested. TSNAC significantly reduced the number of iterations necessary, resulting in a much lower computation time while still maintaining perfect detection. Again, the computation time is extremely significant because the goal in target detection is always accuracy in real time. The Wiener and Frost filters performed similarly in terms of detection; however, the Frost filter was significantly faster.

The main reason that the Wiener and Frost filters struggled is because they simply cannot despeckled to the degree that TSNAC and the Gaussian filter can. They are designed to optimize the MSE from a theoretically uncorrupted image, not for target detection or DRLSE.

8.2.3 RISH Results

As stated earlier, the RISH target detection method consists of two phases, training and testing. To test the effectiveness of the filters, each filter was applied to the 24 SAR images in the CARRABAS-II database. These 24 images were then used in the training phase to determine the Transition Matrices for the target and speckle regions. This was also done using unfiltered images which used both an alphabet size of two and three symbols. For the testing phase, a single, unfiltered SAR image was used. It is interesting to note that the training TMs created using the unfiltered images, Wiener filtered images, and Frost filtered images were almost identical. This is intuitive because these filters do not significantly smooth areas of the image with high variance, which target and speckle regions are. Because of this, only the results using the TSNAC filtered training TMs and the unfiltered training TMs are shown.

The target detection performance results using the unfiltered and TSNAC filtered training TMs are displayed as ROC curves. ROC curves show the percent of targets correctly identified

versus the number of false alarms per square kilometer or Type II error. The results are shown below.

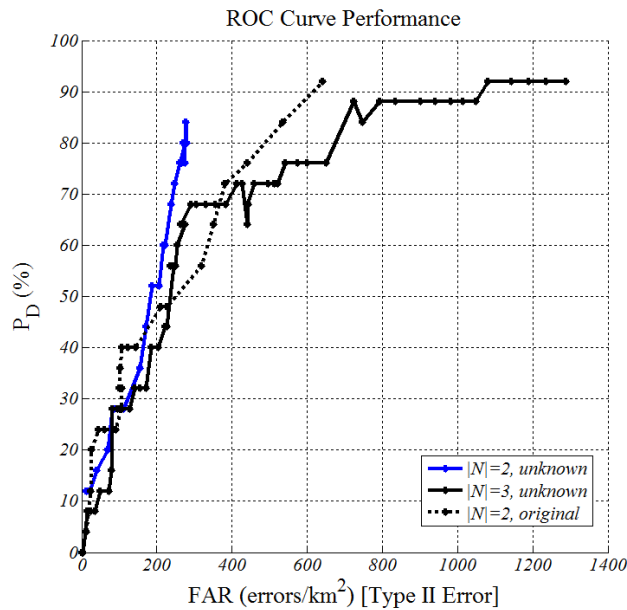


Figure 32. RISH Target Detection Results.

The ROC curves show that the TSNAC filtered TMs result in 85% detection with a reduction in the False Alarm Rate (FAR) of 55% compared to the $|N|=2$ alphabet size and 64% compared to the $|N|=3$ alphabet size. This is obviously a significant reduction in FAR.

The ROC curves also show a limitation of TSNAC filter. In order to achieve the degree of despeckling required to significantly decrease the FAR rate compared to the unfiltered images, and subsequently the Wiener and Frost filtered images, the probability of detection must be sacrificed. The probability of detection reaches between 91 and 92% for the unfiltered training images whereas the TSNAC filtered training images only result in a maximum of 85% detection. Depending on the application, this tradeoff may or may not be acceptable, however, in most cases; a decrease in the percent of correct detection of 6-7% would be negligible compared to a FAR reduction of 55-64%.

8.2.4 Spectrogram Results

As stated earlier, the effectiveness of the filters when applied to the spectrograms of simulated vibration data are judged by the MSE and PSNR. Both the MSE and the PSNR for each filter will be calculated in reference to the spectrogram of the original, uncorrupted signal. The filters will be applied to the spectrogram of the signal degraded by additive Gaussian noise and the MSE and PSNR are calculated based on those outputs.

The uncorrupted spectrogram and the noisy, unfiltered spectrogram were converted to grayscale before processing because the Wiener and Frost filters used here require a grayscale input image.

The spectrogram of the uncorrupted signal of $\sin(2\pi t)$ and the spectrogram of the signal degraded by additive Gaussian noise are shown below.




Figure 33. (a) Spectrogram of uncorrupted time-series data and (b) Spectrogram of time-series data with additive noise.

The result of the Wiener filtered spectrogram did significantly remove noise at frequencies lower than that of the signal of interest. However, because little to no filtering is

done in areas of high intensity gradients, the filtering at frequencies near that of the signal of interest is minimal. The result is shown below.




Figure 34. Wiener filtered spectrogram.

The Frost filtered spectrogram yielded little difference from the noisy input spectrogram. This is most likely due to the nature of the noise. The noise used to in the data was additive while Frost filters are designed specifically for multiplicative noise. The result is shown below.




Figure 35. Frost filtered spectrogram.

The TSNAC filter clearly performed the best compared to the other filters. Like the Wiener filtered spectrogram, it significantly reduced the effect of noise at frequencies lower than the signal of interest. However, it also reduced noise at frequencies near the signal of interest. The result is shown below.




Figure 36. TSNAC filtered spectrogram.

The window size chosen for this data was 10 pixels wide by 100 pixels long. This was chosen primarily to mimic the shape of the spectrogram of the signal of interest, similar to what was done using the SAR dataset. This results in preserving the intensity of the signal of interest and removing irregularities in the signal caused by the additive and destructive noise. This window is also large enough to significantly reduce noise at lower frequencies. These properties are shown in the figure below.




Figure 37. (a) Corrupted signal of interest and (b) Restored signal of interest via TSNAC.

The resulting MSE and PSNR values for each filter are shown in the table below.

Table 4. Spectrogram Filtering Results.

	MSE	PSNR
Noisy	372.4933	22.4196
Wiener	333.3741	22.9015
Frost	382.7006	22.3022
TSNAC	122.7850	27.2393

It is clear that the quantitative results reinforce the qualitative results. TSNAC clearly outperforms the other filters because of its ability to reduce noise while simultaneously reconstructing the original signal of interest. As with the SAR images, this high performance is due to the careful selection of the size and shape of the window. It is worth noting that before the spectrograms were converted to grayscale, the MSE for TSNAC was approximately 50% lower than shown in Table 4. This is because the hue and saturation is lost during the conversion and only the luminance of the image remains.

The other filters did not perform as well. The Wiener filter did significantly reduce noise at the lower frequencies but because it does little filtering in areas of high variance, the area around the signal of interest was barely changed. As shown by the MSE and PSNR, it did improve the image considerably.

The Frost filter actually resulted in a more dissimilar spectrogram than the input noisy spectrogram. This is most likely due to the nature of the noise. Again, the Frost filter is essentially a Wiener filter designed for multiplicative noise and the noise in this case was additive. This is why the Wiener filter, even though they are very similar, greatly outperformed the Frost filter.

8.2.5 Computational Efficiency

For target detection in particular, computational efficiency is extremely important. For national defense applications, preprocessing methods that are computationally expensive and time consuming have become obsolete. The computational times for each algorithm to filter one SAR image and the corrupted spectrogram are shown in the table below.

Table 5. Computation time of filters for one SAR image and one corrupted spectrogram.

	SAR Computation Time (seconds)	Spectrogram Computation Time (seconds)
Wiener	0.643618	0.046211
Frost	304.816809	87.008568
TSNAC	0.098455	0.043659

Clearly, TSNAC is the most computationally efficient filter with the Wiener filter slightly behind. The filter time for the Frost filter is so long that it is unusable in many real-world applications, regardless of its filtering performance.

9 Conclusions and Future Work

This thesis introduced the novel TSNAC filter. It was shown that, rather than choosing an arbitrary default window size for each stage, effective filtering can result from choosing the window sizes based on characteristics of the input images. Though much more simplistic and computationally efficient, TSNAC outperformed the widely accepted Wiener and Frost filters in a number of different tests. The goal of this thesis was to develop a filtering algorithm that could be used on different types of data, with different types of degradations, and TSNAC accomplished that, successfully filtering out multiplicative speckle in SAR imagery and additive Gaussian noise in simulated vibration data.

The tests used in this thesis have shown that TSNAC is a very effective filter. It outperformed all of the tested filters in AQI score, and DRLSE and RISH target detection. For the synthetic vibration data, it also reduced noise while maintaining relevant signal better than the other filters. While doing this, the TSNAC filter was also much more computationally efficient.

However, these tests also presented some of the limitations of TSNAC. The first is that in order for TSNAC to reduce noise to the degree that it does, a sacrifice must be made in preserving the signal of interest. This was shown particularly in the RISH target detection results. Up to 85% true detection, TSNAC greatly outperformed the unfiltered, Wiener filtered, and Frost filtered data. However, TSNAC filtering was not able to achieve greater than 85% true detection whereas Wiener and Frost filtering could, albeit with a much greater False Alarm Rate. Another limitation is that, for SAR images in particular, the effectiveness of TSNAC filtering is dependent on the size of the target regions relative to the size of the speckle regions. The reason TSNAC can significantly remove speckle while still maintaining much of the signal of interest is

due to the fact that the target regions were significantly larger than most speckle regions. If the targets were closer to the size of the speckle regions, the window size would have to be changed. This would result in either a reduction in smoothing of speckle regions or the signal of interest would not be maintained as well, depending on the change.

While TSNAC filtering does significantly reduce both additive and multiplicative noise while maintaining the signal of interest in different types of systems, there is still room for improvement. Future work includes looking into variable selection of the filtering window size based on the local variance. Wiener and Frost filters determine the filtering effect based on the local variance but in a different way. In regions of low variance, the Wiener and Frost filters are essentially very strong low-pass filters, that is, the cutoff frequency is low. In regions of high variance, a weak low-pass filter is used to preserve the signal. In other words, they essentially weight the filtering effect based on the local variance. For adaptive TSNAC filtering, the size of the window of the filter could be determined based on the local variance. For example, if the local variance is low, then the filter would have a larger window in order to increase the smoothing effect. When the local variance is high, the window could be made smaller to reduce the smoothing effect and to preserve the signal.

For spectrogram analysis, an area of interest is the inverse spectrogram. This would allow the user to see the effect that the filter has on the time series data itself, rather than the spectrogram of the time series data. This has been attempted before but is very difficult because of the overlapping that is necessary to create the spectrogram. The process would involve taking the inverse Fourier transform of each of the columns that form the spectrogram and essentially undoing the overlapping to retrieve the filtered time series data. This would allow a direct

comparison between the input and output of the filter as well as make comparisons easier between the TSNAC filter and other filters.

Overall there are several areas of interest that could be extensively studied that could improve TSNAC filtering, namely its adaptivity and in the evaluation of its effect on the input signal. If more time were available, these topics would be both interesting and challenging additions to this research.

References

- [1] L. Xiao-Qin, "Stochastic resonance in a bistable system with coloured correlation between additive and multiplicative noise," *Chinese Physics*, vol. 13, no. 8, pp. 1201-1209, 2004.
- [2] S. Dangeti, "Denoising Techniques - A Comparison (Master's Thesis)," Louisiana State University, 2003.
- [3] V. S. Frost, J. A. Stiles, K. Shanmugan and J. C. Holtzman, "A Model for Radar Images and its Application to Adaptive Digital Filtering of Multiplicative Noise," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 4, no. 2, pp. 157-166, 1982.
- [4] N. Wiener, *Extrapolation, Interpolation, and Smoothing of Stationary Time Series, with Engineering Applications*, Cambridge: Technology Press of Massachusetts Institute of Technology, 1950.
- [5] J. Gill, *Bayesian Methods: A Social and Behavioral Sciences Approach*, Boca Raton: CRC Press, 2009.
- [6] L. Gagnon and A. Jouan, "Speckle Filtering of SAR Images - A Comparative Study Between Complex-Wavelet-Based and Standard Filters," in *Wavelet Applications in Signal and Image Processing V*, San Diego, 1997.
- [7] S. Intajag and S. Chitwong, "Speckle NOise Estimation with Generalized Gamma Distribution," in *SICE-ICASE International Joint Conference*, Bexco, 2006.
- [8] "Synthetic Aperture Radar," Lockheed Martin, 2013. [Online]. Available: <http://www.lockheedmartin.com/us/products/sar.html>.
- [9] F. Paul, "Theory of Synthetic Aperture Radar," Atlantis Scientific, 1997. [Online]. Available: http://www.geo.uzh.ch/~fpaul/sar_theory.html.
- [10] S. McCandless and C. Jackson, "Chapter 1. Principles of Synthetic Aperture Radar," National Oceanic and Atmospheric Administration, [Online]. Available: http://www.sarusersmanual.com/ManualPDF/NOAASARManual_CH01_pg001-024.pdf.
- [11] M. Lopez, "What is Synthetic Aperture Radar?," Sandia National Laboratories, 2005. [Online]. Available: <http://www.sandia.gov/RADAR/whatis.html>.
- [12] L.-L. Fu and B. Holt, "Seasat Views Oceans and Sea Ice with Synthetic-Aperture Radar," NASA, Pasadena, 1982.
- [13] "VHF Change Detection Problem Set," The Sensor Data Management System, 2013. [Online]. Available: https://www.sdms.afrl.af.mil/index.php?collection=vhf_change_detection.
- [14] D. Inman, *Engineering Vibration*, Upper Saddle River: Prentice Hall, 2008.
- [15] "Calculations for Bearing Defect Frequencies," Update International, 2005. [Online]. Available: <http://www.update-intl.com/VibrationBook8g.htm>.
- [16] D. Felten, "Understanding Bearing Vibration Frequencies," EASA, 2003. [Online]. Available: <http://electromotores.com/PDF/InfoT%C3%A9cnica/EASA/Understanding%20Bearing%20Vibration%20Frequencies.pdf>.
- [17] S. Lacey, "An Overview of Bearing Vibration Analysis," *Maintenance and Asset*

- Management*, vol. 23, no. 6, 2008.
- [18] "Spectrogram," MathWorks, 2007. [Online]. Available: <http://www.mathworks.com/help/signal/ref/spectrogram.html>.
 - [19] R. Gutierrez-Osuna, "Short-time Fourier Analysis and Synthesis," Texas Agriculture and Mining University, [Online]. Available: <http://research.cs.tamu.edu/prism/lectures/sp/16.pdf>.
 - [20] J. Smith, "Spectrograms," in *Mathematics of the Discrete Fourier Transform (DFT) with Audio Applications, Second Edition*, 2007.
 - [21] S. Kumar, "An Introduction to Image Compression," DebugMode, 2001. [Online]. Available: <http://www.debugmode.com/imagecmp/>.
 - [22] H. Aresenault and G. April, "Information Content of Images Degraded by Speckle Noise," in *International Conference on Speckle*, San Diego, 1985.
 - [23] S. W. Smith, "Linear Image Processing," in *The Scientist and Engineer's Guide to Digital Signal Processing*, 1998.
 - [24] R. C. Gonzalez and R. E. Woods, *Digital Image Processing*, Upper Saddle River: Prentice Hall, 2008.
 - [25] R. Fisher, S. Perkins, A. Walker and E. Wolfart, "Mean Filter," University of Edinburgh, 2003. [Online]. Available: <http://homepages.inf.ed.ac.uk/rbf/HIPR2/mean.htm>.
 - [26] "Filters," Biomathematics and Statistics Scotland, [Online]. Available: <http://www.bioss.ac.uk/people/chris/ch3.pdf>.
 - [27] "Convolution and Filtering," Aalborg Univeristy, [Online]. Available: <http://www.slideserve.com/kara/image-processing-3-convolution-and-filtering>.
 - [28] "fspecial," MathWorks, 2006. [Online]. Available: <http://www.mathworks.com/help/images/ref/fspecial.html>.
 - [29] "Wiener Filtering," Rice University, 1999. [Online]. Available: <http://www.owlnet.rice.edu/~elec539/Projects99/BACH/proj2/wiener.html>.
 - [30] T. Veldhuizen, "The Wiener Filter," University of Edinburgh, 1998. [Online]. Available: http://homepages.inf.ed.ac.uk/rbf/CVonline/LOCAL_COPIES/VELDHUIZEN/node15.html.
 - [31] V. Pascazio and G. Schirinzi, "Reduction of Multiplicative Noise in SAR Imaging by Wiener Filter," *IEEE*, 1996.
 - [32] h. P. Zhang and P. Zhou, "Feature Enhancement of SAR Imagery using Wiener Filtering in Multi-Wavelet Domain," *Multispectral Image Processing*, 2007.
 - [33] T. Jin, Z. Zhou and Q. Song, "Subaperture Wiener Filter Construction for Time-varying RFI Suppression in UWB SAR," *IEEE*, 2007.
 - [34] V. I. Ponomarev and A. B. Pogrebniak, "Adaptive Wiener Filter Implementation for Image Processing," in *MMET*, 1996.
 - [35] J. S. Lim, *Two-Dimensional Signal and Image Processing*, Upper Saddle River: Prentice Hall, 1990.
 - [36] V. Lukin, V. Melnik, V. Chemerovsky, J. Astola and K. Saarinen, "Hard-switching Adaptive Filter for Speckle Image Processing," Kharkov Aviation Institute, Kharkov.
 - [37] Z. Shi and K. B. Fung, "A Comparison of Digital Speckle Filters," Canadian Crown,

Ottawa, 1994.

- [38] S. Gabarda and G. Cristóbal, "Blind Image Quality Assessment through Anisotropy," *Journal of the Optical Society of America*, vol. 24, no. 12, 2007.
- [39] "Image Entropy," Cornell University, [Online]. Available: <http://www.astro.cornell.edu/research/projects/compression/entropy.html>.
- [40] C. Li, C. Xu, C. Gui and M. D. Fox, "Distance Regularized Level Set Evolution and Its Application to Image Segmentation," *IEEE Transactions on Image Processing*, vol. 19, no. 12, 2010.