Multiscale Segmentation and Anomaly Enhancement of SAR Imagery

by

Charles H. Fosgate B.S., Northeastern University (1994)

Submitted to the Department of Electrical Engineering and Computer Science in partial fulfillment of the requirements for the degree of

Master of Science in Electrical Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 1996

© Massachusetts Institute of Technology 1996. All rights reserved.

Author	Department of Electrical Engineering and Computer Science May 17, 1996
Certified by	Alan S. Willsky Professor of Electrical Engineering Thesis Supervisor
Certified by Research	Hamid Krim h Scientist, Laboratory for Information and Decision Systems Thesis Supervisor
Certified by	Ronald D. Chaney Technical Staff, MIT Lincoln Laboratory Thesis Supervisor
Accepted by	Frederick R. Morganthaler Chairman, Departmental Committee on Graduate Students

Multiscale Segmentation and Anomaly Enhancement of SAR Imagery

by

Charles H. Fosgate

Submitted to the Department of Electrical Engineering and Computer Science on May 17, 1996, in partial fulfillment of the requirements for the degree of Master of Science in Electrical Engineering

Abstract

We present an efficient multiscale approach to the segmentation of natural clutter, specifically grass and forest, and to the enhancement of anomalous SAR image regions. The methods we propose exploit the coherent nature of SAR sensors. In particular, they characterize the scale-to-scale statistical differences in imagery of various terrain categories due to radar speckle. To achieve this, we employ a recently introduced class of multiscale stochastic processes that provide a powerful framework for describing random processes and fields that evolve in scale. We build models representative of each relevant category of terrain (i.e. grass and forest), and use them to direct subsequent decisions on pixel classification, segmentation, and anomaly presence. The scale-autoregressive nature of our models allows for efficient likelihood calculation over SAR image windows. We use these likelihoods as the basis for image pixel classification as well as grass-forest boundary estimation. In addition, anomaly enhancement comes with little further computation. Specifically, the residuals produced by our models in predicting fine-scale imagery from coarser-scale representations are theoretically uncorrelated. Thus, anomalous pixel regions are enhanced and pinpointed by noting regions whose residuals display a high level of correlation throughout scale. We evaluate the performance of our techniques through testing on 0.3 meter SAR data gathered with Lincoln Laboratory's millimeter-wave SAR.

Thesis Supervisor: Alan S. Willsky

Title: Professor of Electrical Engineering

Thesis Supervisor: Hamid Krim

Title: Research Scientist, Laboratory for Information and Decision Systems

Thesis Supervisor: Ronald D. Chaney

Title: Technical Staff, MIT Lincoln Laboratory

Acknowledgments

First and foremost, I would like to thank my research advisor, Alan S. Willsky, for his guidance and support over the last two years. His energy and enthusiasm is highly contagious and has been a motivating factor in my research. I have thoroughly enjoyed our interaction on both an academic (research) and personal (basketball) level along the way.

Several other individuals have contributed, in some manner, to the completion of this work and deserve mention. Thesis supervisors Hamid Krim and Ronald Chaney both offered extensive mentoring throughout the process as well as invaluable proofing of the final manuscript. Clem Karl provided additional insight and food for thought during weekly "grouplet" meetings. Also, Bill Irving not only laid the foundation for this work, but showed great patience during endless questioning. Thank you all.

My experience at MIT has been highlighted by the many people I have come to know. In particular, those SSG members that, over the past few years, have served as an extended family: Mike Daniel, Paul Fieguth, Austin Frakt, Ben Halpern, Terrence Ho, Bill Irving, Seema Jaggi, Andy Kim, Rachel Learned, Cedric Logan, Ilya Polyak, and Mike Schnieder. Good luck to all of you in future endeavors.

I would like to acknowledge all family members for their support. Curt and Ann, for listening to my gripes; Renée and Richard, for keeping me going during the tough times (and for the diet tips); & Mom and Dad, for not losing faith years ago. Thank you.

Finally, I would like to thank my wife, Michelle, for her patience during the many long weeknights and weekends over the past few years as well as her unwavering support throughout.

Contents

	Ackı	nowledg	ments	5
	List	of Figu	res	11
	List	of Tabl	es	13
1	Intr	oducti	on	15
	1.1	Radar	Speckle	16
	1.2	Multis	cale Modeling	18
	1.3	Thesis	Contributions	19
		1.3.1	Multiscale Terrain Segmentation in SAR Imagery	19
		1.3.2	Multiscale Anomaly Enhancement	21
	1.4	Thesis	Organization	22
2	Bac	kgroui	\mathbf{ad}	25
	2.1	Synthe	etic Aperture Radar	25
		2.1.1	Stripmap Mode SAR	26
	2.2	Multis	scale Modeling of SAR Imagery	31
		2.2.1	Multiscale Framework	32
		2.2.2	Multiresolution SAR Image Sequence Generation	33
		2.2.3	Multiresolution SAR Image Sequence Mapping	35
		2.2.4	Scale-Autoregressive Model Class	37
		2.2.5	State Augmentation	38
		2.2.6	Multiscale Model Identification	39

	2.3	Multis	scale Likelihood Calculation	40
		2.3.1	General Multiscale Likelihood Calculation	40
		2.3.2	Multiscale SAR Likelihood Calculation	42
3	Nat	ural C	Clutter Segmentation	47
	3.1	Mode	l Selection	48
3.2		Regio	nal Pre-Classification	48
		3.2.1	Problem Formulation	50
·		3.2.2	Single versus Double Thresholding	51
	3.3	Classi	fication Refinement	53
	3.4	Exper	rimental Results	56
		3.4.1	Model Construction	56
		3.4.2	Algorithm Specification	59
		3.4.3	Segmentation Performance	63
4	And	omaly	Enhancement	71
	4.1	Lincol	In Laboratory Baseline ATR System	72
		4.1.1	CFAR Detection	74
	4.2	Multis	scale Anomaly Enhancement	75
		4.2.1	Multiscale CFAR Statistic	76
		4.2.2	Combined Multiscale CFAR Statistic	77
		4.2.3	Implementation	79
	4.3		mance Comparison	82
5	Con	clusio	ns & Extensions	95
	5.1	Concl	usions	95
		5.1.1	Natural Clutter Segmentation	95
		5.1.2	Anomaly Enhancement	97
	5.2	Evten		08

5.2.1	Extension to Fully Polarimetric Data
5.2.2	SAR Image Compression
5.2.3	Bright Tree-line Localization
A Log–Rayl	eigh Variance 105 egical Operators 109

~ (.

List of Figures

1-1	Simple speckle example	17
1-2	Basic segmentation process	19
1-3	Basic segmentation approach	20
2-1	Stripmap SAR geometry	26
2-2	Slant range stripmap SAR geometry	29
2-3	q^{th} order multiscale tree	33
2-4	SAR quadtree mapping	36
2-5	Multiscale measurement set notation	41
2-6	SAR image state-augmentation mapping	43
3-1	Pixel blocks and windows for segmentation	50
3-2	Initial classification sequence	53
3-3	Classification refinement	55
3-4	Multiscale model training imagery	57
3-5	Residual distribution fits	59
3-6	Statistical determination of thresholds Ψ_g and Ψ_f	61
3-7	Sample mean of statistic ℓ versus percent forest	63
3-8	Segmentation results, image #1	66
3-9	Segmentation results, image #2	67
3-10	Segmentation results, image #3	68
3-11	Probability of misclassification versus swath width	69

4-1	ATR algorithm block diagram	72
4-2	CFAR detector	74
4-3	Morphological filtering of segmentation results	80
4-4	SAR imagery of corner reflectors	83
4-5	SAR imagery of two man-made structures	83
4-6	SAR imagery of an L-shaped structure	84
4-7	Anomaly enhancement of imagery in Figure 4-4	85
4-8	Anomaly enhancement of imagery in Figure 4-5	86
4-9	Anomaly enhancement of imagery in Figure 4-6	87
4-10	Anomaly enhancement of imagery in Figure 4-4 (6-scale model)	88
4-11	Anomaly enhancement comparison block diagram	90
4-12	Anomaly enhancement performance plots for imagery in Figure 4-4 .	92
4-13	Anomaly enhancement performance plots for imagery in Figure 4-5 $$.	92
4-14	Anomaly enhancement performance plots for imagery in Figure 4-6 .	93
5-1	Segmentation improvement example	103
B-1	Binary morphological erosion	110
B-2	Binary morphological opening	111
B-3	Binary morphological closing	111

List of Tables

3.1	Multiscale model coefficients	57
3.2	Forest model standard deviations	58
3.3	Assigned thresholds Ψ_g and Ψ_f	62
41	Peak and average statistic values	91

Chapter 1

Introduction

In recent years there has been a growing interest in Synthetic Aperture Radar (SAR) imaging, in applications ranging from remote sensing to surface surveillance and automatic target recognition (ATR). For applications such these, the classification of various categories of terrain for delineation (i.e., segmentation) and enhancement of anomalous image regions play critical roles in subsequent analysis for target detection and recognition. In light of typical coverage rates (capable of exceeding $1km^2/s$) of an air-borne SAR, it is crucial to devise efficient (preferably parallelizable) algorithms capable of performing these tasks while, at the same time, meeting the daunting computational demands of the resulting data collection.

In this thesis, we apply multiscale approaches to the SAR image segmentation and anomaly enhancement problems that take full advantage of the coherent nature of SAR image formation. In particular, we build on the idea initially introduced in [12, 25] of characterizing and exploiting the scale-to-scale statistical variations in SAR imagery due to radar *speckle*. Radar speckle is a phenomenon inherent in coherent sensors that gets its name from the grainy or "speckly" appearance of the log-magnitude image. It has historically been thought of as a nuisance and an obstacle to overcome as it generally corrupts the visual quality of the image. In fact methods, such as the polarimetric whitening filter (PWF) [20], have been developed that successfully

minimize the effects of radar speckle and improve overall visual quality. Since we are not interested in visual image quality, but rather the statistical properties of terrain categories that come about as a consequence of this radar speckle, we do not attempt in any manner to reduce it, but rather to take advantage of its statistical character. Since speckle plays a major role in the discussions to follow, we provide a brief introduction to radar speckle and radar scatterers.

1.1 Radar Speckle

In representing the relationship between the emitted waveform and the received echo in remote sensing, it is often useful to model the terrain as a discrete set of scatterers. With each scatterer we associate a radar reflectivity, (i.e., attenuation and complex phase). The primary goal in radar imaging is to estimate the terrain radar reflectivity from the received echos. Ignoring noise, multipath, and other propagation losses, the received echo from each scatterer will be an attenuated, time/phase-shifted version of the transmitted waveform. As an example, suppose there exists a single scatterer, β , within a region of interest with gain G_{β} and phase $e^{j\phi_{\beta}}$. The received echo from the transmitted waveform s(t) will be:

$$G_{\beta} e^{j\phi_{\beta}} s(t-2\tau),$$

where τ represents the one-way propagation time from the transmitter to the scatterer.

Now consider a region of terrain (resolution cell) that is small in size with respect to its range from the transmitting source. If we model the terrain within this cell as consisting of several scatterers, β_i , the received echo, r(t), may be approximated as,

$$r(t) \approx \left(\sum_{i} G_{\beta_i} e^{j\phi_{\beta_i}}\right) s(t-2\tau).$$

We see that for each resolution cell we are essentially measuring the complex sum of the contributions from all scatterers residing within the cell.

To relate this to the concept of radar speckle, we first note that SAR is a coherent sensing device. Hence, unlike optical sensors that sense the magnitude of the returned echo, SAR sensors acquire a measure of the complex radar reflectivity function of the illuminated terrain. As stated above, a SAR image pixel represents a measure of the coherent sum of the complex radar reflectivity of all radar scatterers within each resolution cell. As a result, both constructive and destructive interference occurs. Smooth variations in the true reflectivity log-magnitude ¹ between adjacent resolution cells may accordingly be visually perceived in the sensed image as sharp gradients due to random variations in scatterer phase. Moreover, as image resolution is varied (e.g., as we consider a sequence of images at a succession of scales), the set of scatterers contributing to each pixel changes and thus so does the resulting interference. Consequently, in looking at imagery at successive scales, we also expect to see some degree of statistical variability.

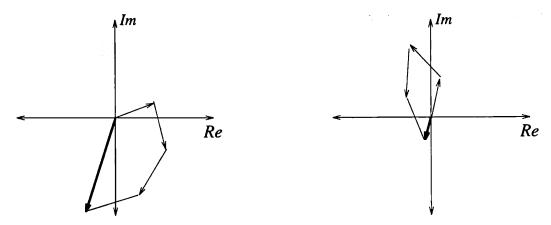


Figure 1-1: Example displaying resolution cell magnitude variation due to the effects of speckle for two distinct cells. The bold vectors in each plot represent the return due to the sum of the scatterers for each cell.

¹It is standard practice to view the log-detected image. In log-detection, each pixel, p_i in the original image is represented by $20 \times log(|p_i|)$ in the log-detected image. This reduces the dynamic range for easier visualization.

As a simple example of the effects of radar speckle we display, in Figure 1-1, a vector plot of hypothetical radar returns from two distinct resolution cells, each containing four discrete scatterers whose reflectivities display similar magnitude but random phase. The magnitude of the returns for each resolution cell differ greatly, as is seen by the bold vectors, even though the scatterers in these two cells have similar overall magnitude and likely appear similar visually.

1.2 Multiscale Modeling

To capture the resulting variability of SAR imagery across scale, we employ a recently introduced class of multiresolution stochastic models [1, 8] that provides a powerful framework for describing random fields that evolve in scale. This framework uses a pyramidal tree structure in which each tree node corresponds to a pixel at a particular image location and resolution; the offspring of the node, in turn, correspond to the pixels, in the same location, at the next finer scale, (i.e., each pixel at a given scale is subdivided, at the next scale into 4 pixels). The statistical variability of this set of imagery for any particular terrain type is then captured by a scale-recursive stochastic model. The exact model and further insight will be provided in Chapter 2.

In [12, 13] this modeling approach was exploited to enhance the performance of the so-called discrimination stage of the ATR system developed at Lincoln Laboratory [15, 22]. This was accomplished by performing a series of likelihood ratio tests on regions of interest for the purpose of classification. In this thesis we employ multiresolution models for classification in a similar manner, yet for a different purpose. Namely, we perform a series of classifications over image regions for the purpose of image segmentation. This segmentation, in turn, provides a means to enhance anomalous regions by identifying statistical outliers given background classifications.

1.3 Thesis Contributions

The work in this thesis represents an extension of the work performed by W.W. Irving in [12, 13]. Our contributions to this area include applying multiscale SAR image classification to the segmentation of SAR imagery and subsequently, to anomalous pixel detection. To accomplish this, we employ multiscale models of SAR imagery, as introduced in [12], and use these models for terrain classification and the identification of statistical outliers. The utility of each of these contributions lies primarily, yet not exclusively, in the application to ATR. We will see that each offers a potential increase in the accuracy of ATR systems such as the baseline system currently under development at Lincoln Laboratory.

1.3.1 Multiscale Terrain Segmentation in SAR Imagery

By image segmentation one refers to the identification of pixel regions that display similar characteristics under some metric. To perform this in an automated fashion, one would provide an image as input and may, for example, receive an N-valued image out. Each of the values, $1, 2, \ldots, N$, would then provide an index to a particular category. For example, to segment grass and forest, we wish to obtain a two-valued output image as in Figure 1-2. In this figure, dark and light grey represent categorization as forest and grass respectively.

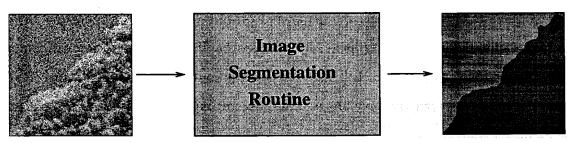


Figure 1-2: Desired output image for a binary segmentation of grass and forest.

Our contribution to this area is a methodology to segment various categories of natural clutter in SAR imagery. We specifically demonstrate the utility of this methodology for the segmentation of regions of trees and forest from open fields and grass. Such segmentation is of considerable value in assisting an ATR system. For example, if a densely forested region is identified, then detection of anomalous pixels in such a region is unnecessary since high-frequency SAR is incapable of providing imagery of targets under forest canopies. Furthermore, targets of interest often attempt to conceal themselves near tree lines, thus accurate estimation of tree lines can be of considerable value in focusing further ATR algorithm attention.

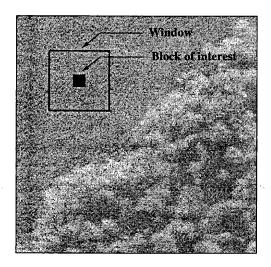


Figure 1-3: Successive classification segmentation approach.

Our segmentation approach focuses on multiscale classification of SAR imagery. By multiscale classification, we refer to classification based on multiscale models. It is assumed that these models are generated off-line (i.e., before processing begins). The classifications are accomplished by means of multiscale log-likelihood ratio tests. As was demonstrated in [12, 13], these multiscale log-likelihood calculations may be performed in an extremely efficient manner. Accordingly, this yields a highly efficient means of classification.

As mentioned above, we exploit the efficiency of existing multiscale classification techniques to achieve a segmentation of our image. The procedure is thus performed as a series of classifications of disjoint regions making up the image. Each region is classified based on a window of pixels centered on the region as is illustrated in Figure 1-3. If a terrain boundary is detected within the window region, the classification is deferred, otherwise the resulting classification holds for the block of interest. This procedure is repeated for each block within the image independently. By deferring decisions on regions suspected of containing terrain boundaries, we avoid possible misclassification due to model mismatch and only classify those regions that offer a high level of confidence (e.g., homogeneous regions of either category) in the assigned classification.

Each region receiving a classification of "defer" is subject to further scrutiny. This is accomplished in a hierarchical manner. In particular, we subdivide the window region surrounding the block of interest into quadrants and independently consider the classification of each quadrant in a manner analogous to the above procedure. Each quadrant classified again as "defer" is again subdivided and the process repeated. The procedure terminates upon determining a majority-rule classification for the entire window region or when the sub-quadrants become small enough such that statistical significance is lost. This classification is assigned to the original block of interest. A positive aspect of this approach lies in the fact that it requires little further computation following the original classification. This, and many other points central to this procedure are discussed in detail in Chapter 3.

1.3.2 Multiscale Anomaly Enhancement

By segmenting regions of natural clutter, as outlined in the previous subsection, we gain the ability to identify pixels that are anomalous with respect to the statistical properties of the now-defined background (i.e., segmentation). Such identification plays a key role in the initial stage of Lincoln Laboratory's baseline ATR system. This initial stage, detection, serves to identify those image regions that are statistically inconsistent with their background. One may imagine that man-made or cultural clutter would yield statistical properties inconsistent with those of natural clutter.

Hence, statistical outliers in imagery consisting primarily of natural clutter could possibly correspond to targets of interest.

Presently, the commonly accepted CFAR [21, 22] detection algorithm is performed to accomplish the intentions of the detection stage. This routine estimates, for each pixel, the second-order statistics of the background and creates a statistic that, if consistent with its background, would have zero-mean and unit variance. Thus outliers may be identified by merely thresholding the CFAR detected image. As mentioned above, our multiscale segmentation also provides a means of anomalous pixel identification. This is accomplished by noting the statistical properties of each terrain model. Those pixels whose statistical properties are inconsistent with those of the model corresponding to the pixel's classification are identified as being anomalous. We will see that not only do we acquire a measure of anomalous behavior throughout scale, but additionally we acquire this at very low computational cost. We present the full discussion, several variants on this procedure, and results in Chapter 4.

1.4 Thesis Organization

In Chapter 2, we provide preliminary discussions on several topics central to this thesis. In particular, we introduce the basics of SAR image formation, an overview of multiscale stochastic processes with an emphasis on multiscale modeling of SAR imagery, and describe the efficient calculation of likelihoods for discrimination between competing models. Chapter 3 describes the central part of this work, namely the classification of SAR imagery for subsequent segmentation and the identification of boundaries between different terrain regions. We begin the chapter by describing the generation of multiscale models for each terrain category, and proceed by detailing the procedure in a theoretical framework. We finish by addressing several key issues, including block and window size selection, and offer experimental results using actual SAR data. Chapter 4 details the application of these multiscale segmentation results

3.5

to the enhancement and subsequent identification of anomalous regions within the imagery. We begin by providing an overview of the baseline ATR system being developed at Lincoln Laboratory and of CFAR detection. The efficient multiscale anomaly enhancement procedure is defined and evaluated. Experimental results are provided, again using actual SAR data. We finish by providing conclusions regarding the utility of each contribution as well as possible paths for future extensions to this research.

.

.

×X

٠<u>٠</u>

Chapter 2

Background

We provide in this chapter some background material on topics central to this thesis. In Section 2.1 we discuss Synthetic Aperture Radar (SAR) imaging with an emphasis on stripmap mode processing. In section 2.2 we describe the modeling of multiscale stochastic processes, concentrating on multiscale modeling of SAR imagery. We conclude with an overview of multiscale likelihood calculation and describe efficient methods for their calculation.

2.1 Synthetic Aperture Radar

Synthetic Aperture Radar is a methodology for producing extremely high resolution imagery. This is accomplished by merging individual radar returns, collected at discrete intervals, to simulate a *synthetic* array. Finer resolutions are achieved by merging a larger number of contiguous returns. This is analogous to array processing concepts in which a larger number of array elements provides a finer beamwidth and consequently, higher resolution. The fundamental objective of SAR is to increase the azimuthal resolution capabilities of conventional radar by exploiting differences in range rate along the azimuthal direction.

SAR systems generally operate in one of two modes, stripmap and spotlight.

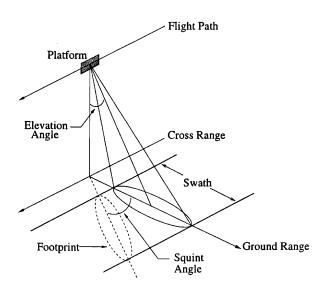


Figure 2-1: Stripmap mode SAR . The antenna beam is fixed perpendicular to the flight path and at some elevation angle with respect to the ground. The resulting footprint traces a swath parallel to the path of the platform.

Each mode serves a particular imaging purpose. Spotlight mode, for example, is characterized by a refocusing of the antenna beam such that a particular terrain region is illuminated at each look. Accordingly, this mode is generally used to gather high resolution imagery of a limited ground region. Stripmap SAR imaging, on the other hand, provides imagery over ground *swaths*, spanning several kilometers in length. In this thesis we will focus on the application of multiscale image processing techniques to stripmap SAR imagery. Towards that end, we provide an in-depth discussion of stripmap mode SAR processing.

2.1.1 Stripmap Mode SAR

Stripmap SAR Geometry

Figure 2-1 [19] shows the geometry associated with a side-looking stripmap SAR. In this figure, the antenna is focused at a *squint* angle of 0 degrees (i.e., perpendicular to the flight path), and at a fixed elevation angle with respect to the ground. Clearly,

operation at various elevation and squint angle combinations is possible, however for simplicity we will assume throughout this discussion that the antenna remains focused in this manner. The portion of the ground illuminated by the antenna's beam at any given moment is referred to as the *footprint* and the ground region illuminated by this footprint over time, the *swath*. The range to the center of the footprint is commonly referred to as the *slant range*. All future references to range in this section will imply slant range, which via simple trigonometric manipulations will yield true ground range.

Pulse Compression

The ideal waveform for pulsed radar systems is impulsive in nature. Impulsive waveforms provide both the narrow spatial extent required for high range resolution and the energy required for high signal-to-noise ratio (SNR) in the measured echo. Impulsive waveforms, however, are non-optimal for many realistic applications. Their limited spatial extent requires extremely high power transmission levels to obtain SNR reasonable for detection. Moreover, typical transmitters are often incapable of meeting such high power requirements. This limitation may be circumvented by applying a technique commonly referred to as pulse compression. Pulse compression is made possible by transmitting a high bandwidth pulse-like waveform and processed the echo so as to compress the received energy in time. The use of high bandwidth waveforms reduces the peak power requirements of the transmitter by redistributing the energy throughout the spectrum. Furthermore, the compression process allows for finer range resolution. As with simple pulsed radar systems, the temporal delay of the compressed pulse may be used to estimate true target range. SAR processing exploits pulse compression to gain high range and azimuthal resolution.

ţ,

Stripmap SAR Processing

Stripmap SAR processing [19] exploits pulse compression to obtain fine range resolution. The high bandwidth waveform typically transmitted is the linear FM chirp. This waveform is characterized by a constant rate of change in frequency and is given by,

$$s(t) = e^{j\pi\alpha t^2} \operatorname{rect}\left(\frac{t}{T_p}\right),$$
 (2.1)

where,

$$\operatorname{rect}(t) \triangleq \left\{ egin{array}{ll} 1 & |t| \leq rac{1}{2} \\ 0 & \operatorname{otherwise} \end{array} \right.$$

The constant $\alpha > 0$, determines the rate at which the frequency varies and is consequently called the *chirp rate*, while the temporal extent of the waveform is determined by T_p . This waveform is modulated on a carrier frequency, f_c , for transmission. The transmitted waveform is,

$$\Re e\left\{s(t)e^{j2\pi f_c t}\right\}. \tag{2.2}$$

The received echo is a function of pulse transmission time index number, n, integration time interval, Δ , and time. In Figure 2-2, we see this scenario where we obtain a return from the radar reflectivity function q(x, y) over the entire radar footprint at discrete spacings of distance Δ .

We model the return as,

$$r(t, n\Delta) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} w(x, y - n\Delta) G(x, y) \Re e^{\left\{s\left(t - \frac{2R(x, y - n\Delta)}{c}\right)\right\}} dx dy.$$

$$\left. e^{j\left[2\pi f_c\left(t - \frac{2R(x, y - n\Delta)}{c}\right) + \Phi(x, y)\right]}\right\} dx dy.$$
(2.3)

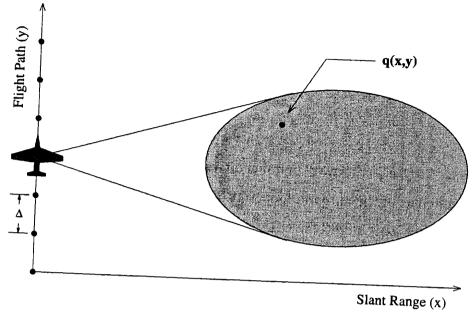


Figure 2-2: Slant range stripmap SAR geometry

In the above, the complex radar reflectivity, q(x, y), that we would like to estimate is,

$$q(x,y) = G(x,y) e^{j\Phi(x,y)},$$
 (2.4)

where G(x,y) is the two-dimensional gain function and $\Phi(x,y)$ the associated phase. The antenna spatial gain function is represented in (2.3) by $w(x,y-n\Delta)$. This function uniformly decreases away from the direction of transmission and accounts for gain roll off. The contribution of this function vanished outside the radar footprint, thus allowing us to use infinite regions of integration. The range to a scatterer at location (x,y) with the n^{th} pulse transmission is represented in (2.3) by $R(x,y-n\Delta)$. We see that the return at any time for a distinct pulse transmission corresponds to an integration over the entire radar footprint of delayed, phase-shifted, and attenuated versions of the transmitted waveform.

For simplicity, we will assume that the antenna gain function, w(x, y), is separable and may be approximated by the product, $w_{az}(y)w_{ra}(x)$. The quadrature demodu-

lated ¹ return may, as a result, be written as,

$$\overline{r}(t, n\Delta) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} w_{ra}(x) w_{az}(y - n\Delta) q(x, y) s \left(t - \frac{2R(x, y - n\Delta)}{c} \right) e^{-j2kR(x, y - n\Delta)} dx dy,$$
(2.5)

where $k = 2\pi f_c/c$ is the wave number at carrier frequency f_c . Due to the narrow azimuthal extent of the antenna footprint, the range, as a function of azimuthal position, should vary only slightly over the footprint and may be approximated by x. This slight variation over azimuth will, however, have a drastic effect on the exponential term in (2.5). Accordingly, we approximate $R(x, y - n\Delta)$ in the exponential term of (2.5) by the first two terms in its Taylor series expansion:

$$R(x, y - n\Delta) = \sqrt{x^2 + (y - n\Delta)^2} \approx x + \frac{(y - n\Delta)^2}{2x}$$

We may now, as a result, approximate the quadrature modulated echo as,

$$\overline{r}(t, n\Delta) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} w_{ra}(x) w_{az}(y - n\Delta) q(x, y) s\left(t - \frac{2x}{c}\right) e^{-j2kx} e^{-jk(y - n\Delta)^2/x} dx dy.$$
(2.6)

By incorporating a change of variables and evaluating (2.6) at time $t = 2R_0/c$, we obtain the spatial form of (2.6):

$$\tilde{r}(x, n\Delta) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} w_{ra}(\rho) w_{az}(\sigma - n\Delta) q(\rho, \sigma) s\left(\frac{2(R_0 - \rho)}{c}\right)$$

$$e^{-j2k\rho} e^{-jk(\sigma - n\Delta)^2/\rho} d\rho d\sigma. \tag{2.7}$$

Range and azimuth compression may now be obtained by exploiting the range

¹Quadrature demodulation is a process in which the received signal is down-converted and separated into its in-phase (real) and quadrature (imaginary) components. This enables representation of the underlying complex nature of a real radar return.

and azimuth dependent chirp terms in (2.7). In particular, cross-correlation with a matched filter of s(t) evaluated for scatterers at range R_0 ,

$$h_{ra} = s^*(2R_0/c),$$
 (2.8)

provides the desired range compression. Similarly, azimuthal compression is possible due to the chirp-like exponential term in (2.7). Since this term is azimuth-dependent, cross-correlation (at some range $\rho = R_0$) with the matched filter,

$$h_{az} = e^{-jk(n\Delta)^2/R_0} \operatorname{rect}\left(\frac{n\Delta}{L}\right), \tag{2.9}$$

provides azimuthal compression for a synthetic aperture of length L.

2.2 Multiscale Modeling of SAR Imagery

We use as a basis for the representation of SAR imagery, the multiscale framework introduced in [1]. This concept was introduced in [12], where multiscale models of SAR imagery were used to discriminate between regions of cultural and natural clutter. We provide here the methodology for developing multiscale stochastic models of SAR imagery. We begin with an overview of the multiscale framework as the basis for modeling and analysis of multiscale stochastic processes. This is followed by a discussion of multiscale SAR image modeling. In particular, we note that the models are constructed starting with some finest-scale SAR training image. Accordingly, we identify a technique for generating a multiresolution sequence of imagery and subsequently mapping it into the multiscale framework. We then motivate the research undertaken herein by identifying a restricted class of multiscale models for SAR imagery. We finish by providing a precise method for multiscale model parameter identification.

×

2.2.1 Multiscale Framework

The multiscale framework developed in [1, 12] enables one to systematically describe and analyze processes that evolve in scale. For those processes that lend themselves to a multiresolution description (i.e., SAR imagery), this framework provides efficient, highly parallelizable methods for modeling and analysis. Within this framework, processes are mapped to indexed nodes on trees. The trees are graphs in which each node, starting with a $root\ node$, is connected to a specified number of "child" nodes. Thus each level of the tree may represent a distinct scale/resolution in the representation of a random process, with the resolutions proceeding from coarse to fine as the tree is traversed from top to bottom (root node to terminal nodes). An example of a q^{th} order tree (q children per node) is displayed in Figure 2-3. This mapping provides a means of visualization as well as a representation of inter-nodal relationships crucial to the development of dynamical models.

Both an upward shift (coarser scale) operator, $\overline{\gamma}$, and a downward shift (finer scale) operator, α_n , are defined for these trees. Hence, if we let s denote any node on the tree, $s\overline{\gamma}$, $s\overline{\gamma}^2$, etc., respectively refer to the parent, grandparent, etc. of node s. Similarly, $s\alpha_1$, $s\alpha_2$, ..., $s\alpha_q$ refer to the q children of node s. The state values at the indexed tree nodes (initialized by the random vector x_0 at the root node) may thus be described by the forward model (coarse-to-fine recursion)

$$x(s) = A(s)x(s\bar{\gamma}) + B(s)w(s), \qquad (2.10)$$

where w(s) represents white driving noise and A(s), B(s) are matrices of appropriate size. A model for the coarse-to-fine evolution of the multiscale stochastic process is identified by defining A(s) and B(s) for all s.

One of the fundamental characteristics of these multiscale processes is that they exhibit the Markov property. As a result, conditioned on the state value at node s, the processes defined on each of the distinct subtrees extending away from node s are

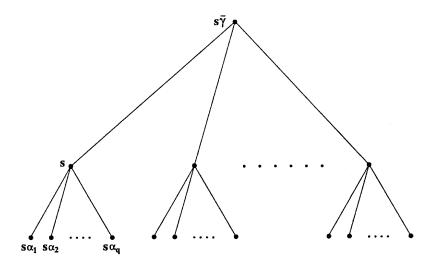


Figure 2-3: The first three levels of a q^{th} order tree are shown. The parent of node s is given by $s\overline{\gamma}$ and the children are given by $s\alpha_1$ through $s\alpha_q$.

mutually independent. This property is a direct consequence of the additive white driving noise term, w(s), and results in highly efficient signal processing algorithms for multiscale stochastic processes defined accordingly.

2.2.2 Multiresolution SAR Image Sequence Generation

The natural starting point for multiscale modeling of SAR imagery is a multiresolution sequence of SAR images which, in some sense, captures the essence of the terrain being modeled. To attain a model for regions of forest, for example, the logical choice would be a sequence of SAR images, at various resolutions, of some particular representative region of forest. Because the azimuthal resolution of SAR imagery is dependent upon the number of returns (i.e., synthetic aperture length) employed in post-processing, the consideration of SAR imagery at multiple resolutions comes naturally. With direct access to the raw radar echos, one may accordingly construct SAR imagery at several resolutions simultaneously. In this thesis, however, we merely have access to the finest-scale SAR imagery formed from the raw returns. Hence, we provide the system by which we emulate SAR imagery at several coarser resolutions from some

**

finest-scale image for the purpose of multiscale modeling and analysis.

The "coarsening" process is accomplished with the following succession of steps:

- Coherent averaging
- Log-detection
- Mean subtraction

The coherent averaging is accomplished by summing blocks of pixels in the finest-scale complex valued image, \mathcal{I}_0 . For example, to obtain the complex valued image at resolution m, \mathcal{I}_m , $2^m \times 2^m$ pixel blocks in \mathcal{I}_0 are summed. This procedure is attractive in that it may be recursively applied. Specifically, the complex valued image at resolution m is acquired from the complex valued image at resolution m-1 as,

$$\mathcal{I}_m(k,l) = \mathcal{I}_{m-1}(\lceil k/2 \rceil, \lceil l/2 \rceil), \qquad (2.11)$$

where $\lceil x \rceil$ is the smallest integer greater than or equal to x. The impetus behind this approach lies in the treatment of each pixel as a resolution cell. Each complex pixel value in this representation depicts the coherent sum of the returns from the individual scatterers within the resolution cell. Consequently, in reducing the resolution by a factor of two (i.e., doubling the resolution cell size), the individual scatterers in four adjacent cells in the finer resolution image will contribute to one cell in the coarser image. Although this is an oversimplification, it does provide for a reasonable and efficient approximation to true coarser resolution imagery. Moreover, this is essentially a low-pass filtering and decimation operation, and may thus be viewed as reducing SAR pulse bandwidth and synthetic aperture length simultaneously. We note that this operation is best suited for finest-scale imagery, \mathcal{I}_0 , consisting of $M \times N$ pixels for $M = 2^p$ and $N = 2^q$ for some p and q.

Although the first step provides a multiresolution sequence of SAR imagery, further processing is performed to ensure that the data is well behaved. Log-detection of each complex valued image, \mathcal{I}_m , achieves, among other factors, a reduction in image dynamic range. The log-detected image, \mathcal{I}'_m , is calculated as,

$$\mathcal{I}'_{m}(k,l) = 20 \log_{10} [|\mathcal{I}_{m}(k,l)|].$$
 (2.12)

Note that the imagery is converted to log-magnitude because the multiscale recursive models that we consider have empirically proven most effective using this representation. Use of the complex imagery directly is impractical because of the variability of the phase. In addition, the log-magnitude of the imagery provides statistically more well-behaved residuals than the magnitude imagery.

As a final step, we subtract out the sample mean of each image. The resulting image, I_m , is computed as,

$$I_m(k,l) = \mathcal{I}'_m(k,l) - \hat{\mu}_m,$$
 (2.13)

where,

$$\hat{\mu}_{m} = \left(\frac{4^{m}}{N \times M}\right) \sum_{i=0}^{\frac{M}{2^{m}}-1} \sum_{j=0}^{\frac{N}{2^{m}}-1} \mathcal{I}'_{m}(i,j), \tag{2.14}$$

and the finest-scale image contains $M \times N$ pixels. This process offsets any spurious DC gain variations between successive stripmap SAR passes.

2.2.3 Multiresolution SAR Image Sequence Mapping

We now consider a multiscale sequence, I_0, I_1, \ldots, I_L , of SAR images, produced via the method described above in Section 2.2.2, where I_L and I_0 correspond to the coarsest and finest resolution images respectively. We would like to create a mapping for each pixel in this sequence of images to nodes on a multiscale tree. In doing so, we produce a multiscale process that may be used as a starting point for model development.

As previously mentioned, we assume a dyadic variation in resolution between images at successive scales. Hence, for a finest-scale image I_0 with resolution $\delta \times \delta$ consisting of an $M \times N$ array of pixels (with $M, N = 2^p, 2^q$ for some p, q), each coarser resolution image I_m has $2^{-m}M \times 2^{-m}N$ pixels and resolution $2^m\delta \times 2^m\delta$. In this arrangement, each pixel in image I_m will correspond to four "child" pixels in image I_{m-1} . This indicates that a 4^{th} order tree, or quadtree, provides a natural order for the mapping. Furthermore, each node s on the quadtree can be thought of as having associated with it a 3-tuple (m, k, l), where m denotes scale and (k, l) denotes two-dimensional image pixel location. That is, each node s on the tree is associated with one of the pixels $I_m(k, l)$ corresponding to pixel (k, l) of SAR image I_m . As an example, Figure 2-4 illustrates a multiscale sequence of three SAR images, together with the quadtree mapping. Here the finest-scale SAR imagery is mapped to the finest level of the tree, and each coarse scale representation is mapped to successively higher levels. Furthermore, we use the notation I(s) to indicate the pixel mapped to node s.

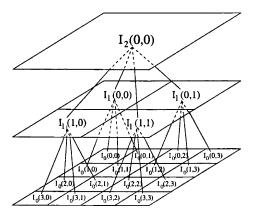


Figure 2-4: Sequence of three multiscale SAR images mapped onto a quadtree. The pixel value at scale m and position (k, l) is denoted by $I_m(k, l)$.

*

2.2.4 Scale-Autoregressive Model Class

In this thesis, as in [12, 13], we focus on a specific class of multiscale models, namely scale-autoregressive models of the form

$$I(s) = a_1(s)I(s\overline{\gamma}) + a_2(s)I(s\overline{\gamma}^2) + \ldots + a_R(s)I(s\overline{\gamma}^R) + w(s), \ a_i(s) \in \mathbb{R}, \ (2.15)$$

where w(s) represents white driving noise, and R the order of the regression. Restriction to such a class of models has several motivating factors, most significantly, the underlying physical process. As previously mentioned in Chapter 1 and illustrated in Figure 1-1, SAR imagery will inherently exhibit the phenomenon of radar speckle. Specifically, SAR imagery will display this phenomenon throughout scale due to the inclusion of different sets of scatterers in a given resolution cell for imagery at various resolutions. As a result, although one would expect to be able to predict, to some extent, a pixel value from its ancestors, the degree of success in doing so will be limited by speckle. The coefficients in (2.15) capture the prediction of the pixel at node s from its ancestors, while the white noise term w(s) depicts the unpredictability in this prediction due to speckle.

For homogeneous regions of texture, we have experimentally found that the prediction coefficients (the $a_i(s)$ in (2.15)) tend to be spatially constant at any given scale. That is, the coefficients, $a_1(s), \ldots, a_R(s)$, in (2.15) depend only on the scale of node s, denoted by m(s), and may thus be represented as $a_{1,m(s)}, \ldots, a_{R,m(s)}$. Furthermore, the probability distribution for w(s) depends only on m(s). Thus, identifying both the regression coefficients as well as the probability distribution for w(s) at each scale completely specifies the model. This provides further motivation for the scale-autoregressive model restriction. Namely, the model may be identified with a small number of coefficients per scale. Multiscale models of homogeneous regions of SAR

10

imagery are accordingly distinguished by models of the form

$$I(s) = a_{1,m(s)}I(s\overline{\gamma}) + a_{2,m(s)}I(s\overline{\gamma}^2) + \ldots + a_{R,m(s)}I(s\overline{\gamma}^R) + w(s), \quad a_{i,m(s)} \in \mathbb{R}. \quad (2.16)$$

2.2.5 State Augmentation

By implementing the scale-autoregressive model class introduced in the previous section, we clearly violate the Markov property that is so crucial to efficient processing of multiscale processes. We circumvent this dilemma by following the procedure of state augmentation used in converting autoregressive time series models to state space models. Namely, we associate to each node s an R-dimensional vector of pixel values, where R is the order of the regression in (2.16). The components of this vector correspond to the SAR image pixel associated with node s and its first R-1 ancestors. Specifically, we define

$$\mathbf{x}(s) = \begin{bmatrix} I(s) & I(s\overline{\gamma}) & \cdots & I(s\overline{\gamma}^{R-1}) \end{bmatrix}^{T}. \tag{2.17}$$

Accordingly, the recursion in (2.10) takes on the form

$$\mathbf{x}(s) = \begin{bmatrix} a_{1,m(s)} & a_{2,m(s)} & \cdots & a_{R-1,m(s)} & a_{R,m(s)} \\ 1 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & \cdots & 0 \\ \vdots & & \ddots & & \vdots \\ 0 & \cdots & 0 & 1 & 0 \end{bmatrix} \mathbf{x}(s\overline{\gamma}) + \begin{bmatrix} 1 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix} w(s). \quad (2.18)$$

It is clear from (2.18) that by associating the vector x(s) in (2.17) with each node s, the vector values at each node depend solely on the values at its parent. This confirms that the Markov property is preserved.

ì

٠ ا

2.2.6 Multiscale Model Identification

As previously mentioned, identification of multiscale models for homogeneous regions of SAR imagery may be accomplished by specifying the model coefficients $a_{1,m(s)}, \ldots, a_{R,m(s)}$ in (2.16) and the probability distribution for w(s) at each scale. This is initiated by selecting a finest-scale homogeneous region of SAR imagery that sufficiently reflects the nature of the terrain being modeled (i.e., training region). Ideally, the pixel size of this training image will be a factor of 2 in both range and cross-range, allowing for multiresolution image sequence generation following the procedure in Section 2.2.2. Accordingly, a multiresolution sequence of SAR images, I_0, I_1, \ldots, I_L , is generated from this image. The regression coefficients for each scale k are obtained by the standard least-squares minimization,

$$\mathbf{a}_{k} = \arg\min_{\mathbf{a}_{k} \in \mathbb{R}^{R}} \left\{ \sum_{\{s \mid m(s)=k\}} \left[I(s) - a_{1,k} I(s\overline{\gamma}) - \dots - a_{R,k} I(s\overline{\gamma}^{R}) \right]^{2} \right\}, \quad (2.19)$$

where

$$\mathbf{a}_k = \left[a_{1,k} \ a_{2,k} \ \dots \ a_{R,k} \right]^T,$$

and R represents the regression order. By varying R, we can achieve a tradeoff between computational complexity and model accuracy. We have found R=3 to be more than adequate for the modeling of SAR imagery of grass and forest.

Providing a statistical characterization for the prediction error residuals at each scale k, completes the model description. To accomplish this, we first evaluate the error residuals in the prediction of each scale of I_0, I_1, \ldots, I_L (generated from the training image) using the coefficients calculated in (2.19). In other words, we use the \mathbf{a}_k found in (2.19) to evaluate all $\{w(s) \mid m(s) = k\}$ in (2.16). Theoretical distributions which provide good fits to the normalized histograms of these residuals at each scale are subsequently chosen.

2.3 Multiscale Likelihood Calculation

Multiscale stochastic models offer insight into a broad array of problems. In addition, as mentioned in Subsection 2.2.1, they are quite attractive in that they afford for extremely efficient signal processing algorithms. Central to this thesis is the efficient manner in which likelihood ratios may be computed for competing multiscale models. Specifically, the Markov property associated with the multiscale modeling framework allows for likelihood calculations that are indeed extremely efficient and parallelizable. We offer here the methodology for multiscale likelihood calculation both for general multiscale stochastic models as well as our restricted scale-autoregressive model class.

2.3.1 General Multiscale Likelihood Calculation

Consider a set of multiscale measurements, Y, on a q^{th} order tree. The likelihood of attaining this set of measurements is defined by some multiscale stochastic model, or equivalently, some joint probability distribution, $p_{\mathbf{y}}(Y)$. Out aim is to outline a tractable method to calculate such likelihoods given a set of multiscale measurements.

We begin by defining some notation. For any node, s, we let Y_s represent the set of finer-scale multiscale measurements extending from and including s. Using this notation, the measurement subsets extending from the q children of s are referred to as $Y_{s\alpha_1}, Y_{s\alpha_2}, \dots, Y_{s\alpha_q}$. This syntax may be extended to the q children of $s\alpha_1$, for example, where these q subsets are referred to as $Y_{s\alpha_1\alpha_1}, Y_{s\alpha_1\alpha_2}, \dots, Y_{s\alpha_1\alpha_q}$. This set notation is displayed in Figure 2-5 for a third-order tree. Furthermore, for this discussion we will refer to the individual measurement at node s as y_s .

If we define the root node as s_0 , the likelihood of obtaining the full set of measurements associated with s_0 and all finer-scale representations, Y_{s_0} , is given by the joint distribution $p_{\mathbf{y}}(Y_{s_0})$ or, by Bayes' law,

$$p_{y_s}(y_{s_0}) p_{\mathbf{y}}(Y_{s_0\alpha_1}, Y_{s_0\alpha_2}, \cdots, Y_{s_0\alpha_q} \mid y_{s_0}).$$
 (2.20)

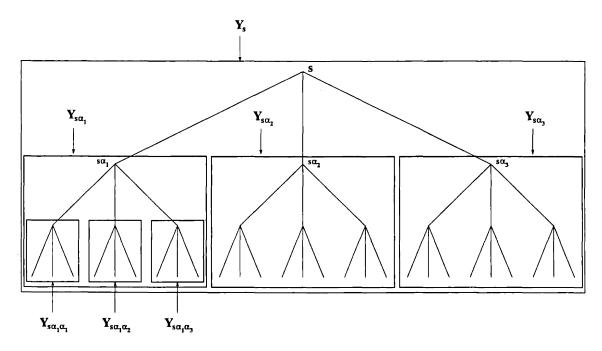


Figure 2-5: Multiscale measurement set notation

The q sets are, due to the Markov property, independent and (2.20) may accordingly be rewritten as,

$$p_{y_s}(y_{s_0}) \left[p_{\mathbf{y}}(Y_{s_0\alpha_1} \mid y_{s_0}) p_{\mathbf{y}}(Y_{s_0\alpha_2} \mid y_{s_0}) \cdots p_{\mathbf{y}}(Y_{s_0\alpha_q} \mid y_{s_0}) \right]. \tag{2.21}$$

Within the square brackets in (2.21), there are q terms of the form $p_{\mathbf{y}}(Y_{s_0\alpha_i} \mid y_{s_0})$. Reapplication of Bayes' law allows each of these terms to be rewritten as,

$$p_{y_s}\left(y_{s_0\alpha_i} \mid y_{s_0}\right) p_{\mathbf{y}}\left(Y_{s_0\alpha_i\alpha_1}, Y_{s_0\alpha_i\alpha_2}, \cdots, Y_{s_0\alpha_i\alpha_q} \mid y_{s_0}, y_{s_0\alpha_i}\right), \tag{2.22}$$

or equivalently,

$$p_{y_s}\left(y_{s_0\alpha_i} \mid y_{s_0}\right) p_{\mathbf{y}}\left(Y_{s_0\alpha_i\alpha_1}, Y_{s_0\alpha_i\alpha_2}, \cdots, Y_{s_0\alpha_i\alpha_q} \mid y_{s_0\alpha_i}\right). \tag{2.23}$$

As before, in the progression from (2.20) to (2.21), independence of the individual

sets allows each of the terms in (2.23) to be rewritten as,

$$p_{\mathbf{y}_s}\left(y_{s_0\alpha_i}\mid y_{s_0}\right)\left[p_{\mathbf{y}}\left(Y_{s_0\alpha_i\alpha_1}\mid y_{s_0\alpha_i}\right)p_{\mathbf{y}}\left(Y_{s_0\alpha_i\alpha_2}\mid y_{s_0\alpha_i}\right)\cdots p_{\mathbf{y}}\left(Y_{s_0\alpha_i\alpha_q}\mid y_{s_0\alpha_i}\right)\right]. \tag{2.24}$$

By recursively following this scheme down each branch of the multiscale tree to the terminal nodes, the final representation for the likelihood is,

$$p_{\mathbf{y}}\left(Y_{s_0}\right) = p_{y_s}\left(y_{s_0}\right) \prod_{s \in \mathcal{S}} p_{y_s}\left(y_s \mid y_{s\overline{\gamma}}\right), \qquad (2.25)$$

where S is the set of all tree nodes, excluding the root node s_0 , for which there exist measurements. If, as is commonly the case, we are interested in log-likelihoods, further simplification is achieved. The log-likelihood for the multiscale set of measurements is merely the sum of the conditional log-likelihoods at each node,

$$\log \{p_{\mathbf{y}}(Y_{s_0})\} = \log \{p_{y_s}(y_{s_0})\} + \sum_{s \in S} \log \{p_{y_s}(y_s \mid y_{s\overline{\gamma}})\}.$$
 (2.26)

2.3.2 Multiscale SAR Likelihood Calculation

Likelihood calculations are central to the classification of and discrimination between various categories of SAR imagery. For example, in [12] the calculation of likelihoods provided the basis for discrimination between SAR imagery of natural and man-made clutter. Given a region of SAR imagery, I, we can calculate the individual likelihoods of obtaining the sequence of images, I_0, I_1, \ldots, I_L , resulting from successive coarsening of I. These likelihoods may subsequently be used to obtain a sufficient statistic for SAR image terrain classification.

Consider a complex-valued region of SAR imagery, I, and a multiscale model, M_1 , for SAR imagery representing a particular terrain category. We begin by creating a multiresolution sequence of SAR images, I_0, I_1, \ldots, I_L , from I following the procedure described in Subsection 2.2.2. This sequence of images is subsequently mapped to

vector valued nodes on a quadtree as outlined in Subsection 2.2.5. As previously mentioned, this mapping preserves the Markov nature of the multiscale models and enables us to follow the efficient likelihood calculation techniques described by (2.25) and (2.26) in the preceding subsection. As an example, in Figure 2-6 we illustrate this mapping for a third order regression (R=3) and a multiscale model spanning 4 scales (P=4). In this figure, S_i represents the set of all nodes at scale i, and X_i the imagery mapped to these nodes. We illustrate, with the bold images to the right of each scale, the fact that data at each scale after the mapping corresponds to prior data at several scales (i.e., SAR images at several resolutions). Note that the dyadic tree representation of the multiscale process is used solely for simplicity in presentation and implies nothing about the process itself.

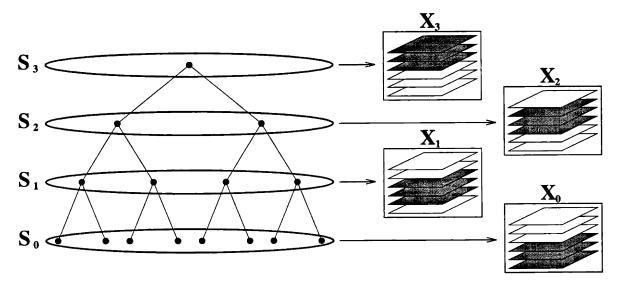


Figure 2-6: Illustration of the mapping involved via state-augmentation. The S_i represents the set of all nodes at resolution i, and X_i the SAR imagery mapped to these respective node sets.

We may now evaluate the likelihood of obtaining I_0, I_1, \ldots, I_L given the multiscale model, M_1 . We first note, using the notation in Figure 2-6, that this likelihood is statistically equivalent to the likelihood of obtaining the data X_0, X_1, \cdots, X_P . Hence,

we may directly apply (2.25) to evaluate this likelihood,

$$p_{I_0, I_1, ..., I_L \mid M_1} (I_0, I_1, ..., I_L \mid M_1) = p_{\mathbf{x}} (\mathbf{X}_{\mathbf{P}}) \prod_{s \in \{S\}} p_x (x(s) \mid x(s\overline{\gamma}), M_1), \quad (2.27)$$

where S denotes the set of nodes $\{S_0, S_1, \dots, S_{P-1}\}$. From (2.18) it is clear that the individual densities, $p_x(x(s)|x(s\overline{\gamma}), M_1)$, are *statistically* equivalent to the prediction error residual distributions,

$$w(s) = I(s) - \left[a_{1,m(s)}I(s\overline{\gamma}) + a_{2,m(s)}I(s\overline{\gamma}^{2}) + \dots + a_{R,m(s)}I(s\overline{\gamma}^{R}) \right], \quad (2.28)$$

given the particular multiscale SAR terrain model M_1 . Consequently, we may reformulate (2.27) as,

$$p_{I_0, I_1, \dots, I_L \mid \mathbf{M}_1} (I_0, I_1, \dots, I_L \mid M_1) = p_{\mathbf{x}} (\mathbf{X}_{\mathbf{P}}) \prod_{s \in \{\mathcal{S}\}} p_w (w(s) \mid M_1).$$
 (2.29)

In Chapter 3 of this thesis, we compute a statistic to determine which of two competing multiscale models provides the best match for a SAR image region. This statistic, ℓ , is defined as the log-likelihood ratio of obtaining a multiscale image sequence for each of the two models. This statistic, for models M_1 and M_2 , is given by:

$$\ell = \log \left\{ p_{I_0, I_1, \dots, I_L \mid M_1} (I_0, I_1, \dots, I_L \mid M_1) \right\} - \log \left\{ p_{I_0, I_1, \dots, I_L \mid M_2} (I_0, I_1, \dots, I_L \mid M_2) \right\}.$$
(2.30)

Using the results in (2.29), it is clear that this is equivalent to,

$$\ell = \log \left\{ p_{\mathbf{x}}\left(\mathbf{X}_{\mathbf{P}}\right) \prod_{s \in \{\mathcal{S}\}} p_{w}\left(w(s) \mid M_{1}\right) \right\} - \log \left\{ p_{\mathbf{x}}\left(\mathbf{X}_{\mathbf{P}}\right) \prod_{s \in \{\mathcal{S}\}} p_{w}\left(w(s) \mid M_{2}\right) \right\}$$

$$= \prod_{s \in \{S\}} \log \{ p_w (w(s) \mid M_1) \} - \log \{ p_w (w(s) \mid M_2) \}.$$
 (2.31)

This statistic may be calculated by merely summing the individual residual loglikelihood differences given the two models. Hence, classification may be achieved in an extremely efficient manner.

Chapter 3

Natural Clutter Segmentation

In this chapter we describe a procedure for the classification and ultimate segmentation of SAR imagery containing natural clutter. The starting point for the development of this procedure is the construction of multiscale stochastic models for homogeneous regions of SAR imagery of various terrain categories, as described in the preceding section. As an illustration of the utility of this approach we focus on segmenting and distinguishing between two broad terrain categories; namely, forest and grass. In addition to providing a two class example of this approach, such segmentation provides information that may be exploited for ATR. As discussed in Chapter 1, a forest-grass segmentation may identify vast regions of forest. Hence, one may envision an algorithm that utilizes this information to increase efficiency by subsequently ignoring these regions since high-frequency SAR is incapable of penetrating forest canopies. Moreover, targets of interest often attempt concealment by hiding near grass/forest boundaries. By identifying such regions, we may efficiently focus future processing.

3.1 Model Selection

Upon choosing the terrain categories of interest for the particular application (i.e. grass and forest as we have chosen here), we must first construct multiscale models representative of each class. This is performed by following the procedure outlined in Chapter 2 using homogeneous image regions of each terrain category. In doing so, we obtain models that capture the scale-to-scale evolution for imagery of homogeneous regions of each terrain. We intentionally stress that the models represent homogeneous regions of SAR imagery. We make this explicit since we will inevitably encounter SAR imagery containing both terrain categories, thus fitting neither model well, and we will have to devise a method to effectively deal with such situations.

We note that the segmentation approach we describe here is readily extended to additional terrain categories and finer level discrimination of each terrain category. This may be accomplished by simply creating additional terrain models and classifying via an M-ary hypothesis test.

3.2 Regional Pre-Classification

Our initial approach to the segmentation problem may be viewed as a series of independent classifications of disjoint blocks, \mathcal{B} , of pixels. Each block of pixels in the image is classified using a likelihood ratio test (LRT) whereby we exploit the efficient calculation of multiscale likelihoods as described in the previous chapter. Specifically, each block of pixels is classified based on a window of pixels, \mathcal{W} , centered on that particular block. We then slide this window from block to block, and repeat the classification as illustrated in Figure 3-1.

We choose to base each classification on a window of pixels exceeding the size of the block being classified due to the fact that the block will often, if not always, be quite small (on the order of 1 to 8 pixels-square). Such small block sizes provide little or no

statistical information in a multiresolution sense. One of our primary assertions then, is that if the imagery within the larger window, W, belongs exclusively to a single homogeneous terrain class, we may classify a block of pixels central to this window with a very small probability of error. In Section 3.4.3 we provide experimental results backing this claim. Moreover, we later devise a procedure to cope with those situations in which imagery of two (or more) terrain classes are present within \mathcal{W} . The size of the window surrounding \mathcal{B} thus represents a tradeoff between the local accuracy and statistical significance in the subsequent classification. In choosing a larger window size one gains a level of statistical significance and would thus expect more accurate classifications in distinguishing homogeneous regions of either terrain type. However, larger window sizes increase the probability that the window region is inhomogeneous. In Section 3.4 we examine this tradeoff by considering window sizes given by 2^N , where $N \in II$, and choose a size that yields excellent classification of homogeneous regions of terrain. An important point to note here is that, by using a window of pixels around each block for subsequent classification, we run into potential problems near the edges of each image being segmented. Rather than implement some sort of extrapolation routine to emulate the statistics of the data outside the actual image, we merely segment the portion of the image such that the window is always fully contained within the image. This procedure does neglect a small ring around the border of each image, yet poses no practical concern. One may readily envision a system to segment stripmap SAR imagery that overcomes this inconvenience. For example, as the imagery is acquired, it is separated into overlapping sub-images. Each sub-image overlaps its neighbor in such a manner that the concatenation of the segmentations of each sub-image results in the segmentation of the full set of imagery gathered (with the exception of the very outermost edges).

The size of the block of pixels, \mathcal{B} , being classified in each step is more of an algorithmic issue. This size determines a tradeoff between efficiency and resolution in the segmented image. For example, using each window to classify a single pixel

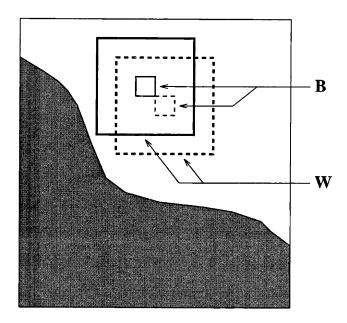


Figure 3-1: Illustration displaying two pixel blocks being classified as well as the two window regions used in the classification.

is much more computationally burdensome than using the same window to classify a 4×4 block of pixels. However, by classifying blocks of pixels simultaneously, the resolution of the resulting segmentation is reduced.

We should note that in a fully optimal method one would also postulate a spatial random field model for terrain type, thus capturing, for example, the fact that the classification of a given pixel is very likely to be the same as its neighbors [9]. The use of such models, however, increases the computational complexity of the classification algorithm considerably and will not be considered in this work.

3.2.1 Problem Formulation

To perform the LRT, we must first define our hypotheses and measurements. The hypotheses of interest to us are: 1) the window region is homogeneous grass-like terrain, and 2) the window region is homogeneous forest-like terrain. Correspondingly, we define H_g and H_f as hypotheses corresponding to situations (1) and (2) respectively. The measurements available in each case are the sequence of images that result from

repeated dyadic reductions in resolution of the region W. We define this multiresolution sequence of images as I_0, I_1, \ldots, I_L , where I_0 and I_L correspond to the finest and coarsest scale representations of W, respectively. Accordingly, the likelihood ratio is defined as,

$$\Lambda = \frac{p_{I_0, I_1, \dots, I_L \mid H_g}(I_0, I_1, \dots, I_L \mid H_g)}{p_{I_0, I_1, \dots, I_L \mid H_f}(I_0, I_1, \dots, I_L \mid H_f)}.$$
(3.1)

We may use this likelihood ratio to perform the LRT by comparing it to a threshold, η . Specifically, the LRT becomes,

$$\Lambda \underset{H_f}{\overset{H_g}{\geq}} \eta. \tag{3.2}$$

Furthermore, the log-likelihood, $\ell = \ln (\Lambda)$, provides an equivalent test,

$$\ell \underset{H_f}{\overset{H_g}{\geq}} \ln (\eta) \stackrel{\triangle}{=} \Psi, \tag{3.3}$$

which is much more efficient since multiplications in Λ become additions in ℓ . The statistic ℓ is given by,

$$\ell = \ln \left[p_{I_0, I_1, \dots, I_L \mid H_g} (I_0, I_1, \dots, I_L \mid H_g) \right] - \ln \left[p_{I_0, I_1, \dots, I_L \mid H_f} (I_0, I_1, \dots, I_L \mid H_f) \right].$$
(3.4)

3.2.2 Single versus Double Thresholding

Given our two hypotheses and the statistic, ℓ , our goal now is to perform a likelihood ratio test and classify each pixel in the block, \mathcal{B} accordingly. As stated in (3.3), the likelihood ratio test involves comparing the value of ℓ to a threshold and classifying depending upon which side of the threshold ℓ falls. As an initial approach to the problem, we set the threshold Ψ in (3.3) to zero. This corresponds to having no

prior knowledge as to the terrain makeup of region W. This LRT performs quite well, as confirmed in Section 3.4, given that W consists of homogeneous terrain from either of our modeled terrain classes. This, however, is not always the case. When the window crosses a boundary between the two terrain categories, both are present within the window simultaneously. Thus, the region is inhomogeneous and fits neither terrain model well, and the validity of the subsequent classification is questionable. Indeed, if we do nothing to counter this phenomenon, the resulting classification will be dubious in the vicinity of terrain boundaries. In fact, as we illustrate in the next section, this procedure has a tendency to extend classification as forest into grassy regions, resulting in an apparent bias in the estimation of the grass-forest boundary.

To overcome this problem we devise both a method to detect proximity to grass-forest boundaries as well as a procedure to accurately perform the subsequent classification. The former of these is accomplished through a very simple modification to the decision rule based on the test statistic ℓ . Specifically, rather than comparing ℓ to the single threshold, $\Psi = 0$, and making a forest-or-grass classification based on this comparison, we compare ℓ to two thresholds:

 $\ell > \Psi_g$ Classify as Grass

 $\Psi_g \,>\, \ell \,>\, \Psi_f$ Defer decision, as a terrain boundary may be present within ${\cal W}$

 $\ell < \Psi_f$ Classify as Forest

The thresholds in this approach are chosen such that when ℓ exceeds Ψ_g or is less than Ψ_f , we have a high level of confidence that \mathcal{W} is a homogeneous region of grass or forest, respectively. In fact, we may set these thresholds to meet specified probabilities of correct classification, as we address in Section 3.4. When the value of ℓ falls between the two thresholds, it is expected that a terrain boundary is present within \mathcal{W} . Hence, we defer the classification of \mathcal{B} for further refinement.

The resulting test structure for the initial pixel block classification is illustrated in Figure 3-2. Here, the box designated "defer" corresponds to the test statistic, ℓ , falling between the two thresholds, Ψ_g and Ψ_f , and thus the window requires additional scrutiny. As we describe in Section 3.4, examination of empirical distributions of ℓ over windows containing boundaries and with differing percentages of forest and grass allows us to set Ψ_g and Ψ_f .

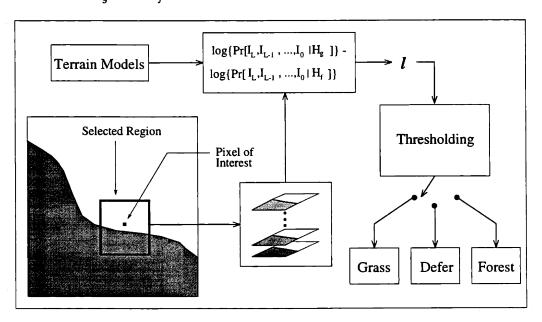


Figure 3-2: Sequence of steps involved in initial classification of pixels. 1) Creation of multiscale sequence from square region surrounding pixel block of interest. 2) Evaluation of statistic ℓ . 3) Thresholding to determine if the block should be classified as grass or forest, or should be flagged as potentially containing a grass-forest boundary.

3.3 Classification Refinement

The procedure as we have described it so far requires a computationally efficient method to deal with the deferred decisions caused by the potential presence of terrain boundaries. The structure of the multiscale likelihood calculation allows us to perform this additional task with very little computational overhead and in fact in a fashion that can be viewed as a replication of the procedure of Figure 3-2 at a hierarchy of

scales. Recall that the objective of this classification process is to classify a block of pixels centered in a window as either forest or grass. Consequently, in a region suspected of containing a single terrain boundary, we simply wish to determine on which side of the boundary this block resides. If we assume that the size of the window $\mathcal W$ is sufficiently small such that at most one boundary resides within the window and that the boundary is relatively smooth and linear, the resulting classification reduces to determining which of the two hypothesized terrain types make up the majority of the imagery within the window.

Consequently, we proceed to reclassify each deferred block by considering each of the four quadrants of W independently. For example, if we denote the four quadrants of W as W_i for i = 1, 2, 3, 4, and the four corresponding multiresolution sequences of images as $I_0^i, I_1^i, \ldots, I_L^i$ again for i = 1, 2, 3, 4, then the log-likelihood ratio for each quadrant, ℓ_i , will be given as,

$$\ell_{i} = \ln \left[p_{I_{0}^{i}, I_{1}^{i}, \dots, I_{L}^{i} \mid \mathbf{H}_{g}} (I_{0}^{i}, I_{1}^{i}, \dots, I_{L}^{i} \mid H_{g}) \right] - \ln \left[p_{I_{0}^{i}, I_{1}^{i}, \dots, I_{L}^{i} \mid \mathbf{H}_{f}} (I_{0}^{i}, I_{1}^{i}, \dots, I_{L}^{i} \mid H_{f}) \right].$$
(3.5)

Due to the structure of the multiscale models, the likelihood calculation described in Chapter 2 may be carried out in a pyramidal fashion. That is, each ℓ_i may be obtained by considering the individual likelihoods used in calculating ℓ over the sub-trees corresponding to each quadrant. The classification procedure described in Section 3.2 may then be performed on each of these quadrants independently based upon each ℓ_i with the interpretation that each classification be assigned to all pixels within the quadrant. This procedure may be repeated recursively on each smaller sub-window or sub-quadrant deferred in the previous hierarchical step in an analogous manner. A final classification is established when either a majority rule decision of the pixels in the entire windowed region, \mathcal{W} , may be determined or statistical significance is lost due to decreasing quadrant size. In the latter, the central block of pixels is classified

according to a majority rule of all pixels classified to that point.

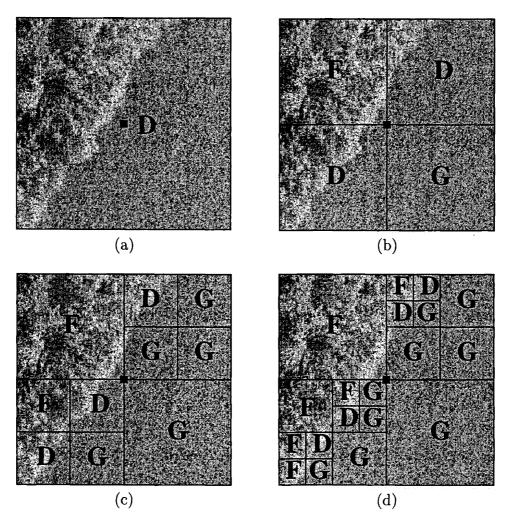


Figure 3-3: Potential boundary refinement of typical SAR window region: (a) Window region classified as boundary. b) Region divided into quadrants, with classification results marked for each quadrant. (c) & (d) Regions still classified as boundary further sub-divided and reclassified. Classifications as G, F, and D refer to grass, forest, and defer respectively.

The hierarchical succession of classifications in this "progressive refinement" procedure is pictorially summarized in Figure 3-3. In each frame, the block being classified is represented by the solid square box central to each image. Figure 3-3(a) represents the original windowed region, W. Based on the test illustrated in Figure 3-2 this win-

: a

dow has been identified as potentially containing a grass-forest boundary. Note that the correct classification for the block is grass, yet given the window size and boundary proximity, the classification has been deferred. Figures 3-3(b) through 3-3(d) display the successive subdivision and reclassification (F, G, and D represent classification) as forest, grass, and defer respectively), where at each stage only the regions deferred at the preceding stage are subject to further examination. In this example, 9/16 of the region is ultimately classified as grass, and thus the center block is correctly classified. Central to this procedure is the fact that with proper organization, each of the log-likelihood ratios needed for every sub-window of potential interest (e.g., the full window in Figure 3-3(a) and each of the sub-windows shown in the other parts of Figure 3-3) can be calculated with the same total computational effort as that needed to calculate the single likelihood for the entire window. This provides an effective and highly efficient means of classifying those blocks in the vicinity of terrain boundaries.

3.4 Experimental Results

To evaluate the performance of this approach, we have applied it to 0.3 meter resolution HH polarization SAR data gathered over Stockbridge, New York with Lincoln-Laboratory's millimeter-wave SAR [11]. In the following subsections we describe the construction of the models on which the subsequent experiments are based, discuss the details of the algorithm design (namely the setting of window sizes and decision thresholds), and illustrate the performance of our algorithm for classification and segmentation.

3.4.1 Model Construction

To apply the preceding segmentation approach, we first construct multiscale SAR models for imagery of grass and forest from homogeneous regions of imagery for each terrain. We used the two 256 pixel-square regions of grass and forest displayed

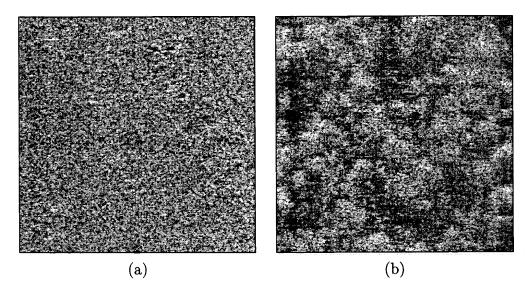


Figure 3-4: 256 pixel homogeneous regions of grass (a) and forest (b) used in multiscale model creation.

in Figure 3-4 as "training" data for this purpose. We chose to use a third-order regression for each model and to build models for the prediction of each of the three finest resolution images $(\delta \times \delta, 2\delta \times 2\delta, \text{ and } 4\delta \times 4\delta, \text{ with } \delta = 0.3m)$. As described in Chapter 2, this implies a third-order model with a four-level tree and for the prediction of each of the three finest scales there are three coefficients, a_1 , a_2 , a_3 , to be specified. Using the method described in Section 2.2.6 we determined the coefficient values given in Table 3.1. Note that the coefficients for the forest model are consistently larger, indicating higher scale to scale correlation. This inter-scale correlation is consistent with our expectations due to the scatterer distribution in the image. Grassy regions

Grass model coefficients				
Resolution	a_1	a_2	a_3	
$\delta \times \delta$	0.5263	0.0720	-0.0029	
$2\delta \times 2\delta$	0.3135	0.0313	-0.0064	
$4\delta \times 4\delta$	0.2278	0.0169	-0.0006	

Forest model coefficients				
Resolution	a_1	a_2	a_3	
$\delta \times \delta$	0.5842	0.1257	0.0669	
$2\delta \times 2\delta$	0.5005	0.1222	0.0683	
$4\delta \times 4\delta$	0.4584a	0.1292	0.0250	

(a) (b)

Table 3.1: Model coefficients for third order regression in scale. (a) Grass model coefficients. (b) Forest model coefficients.

tend to have a large number of equi-valued scatterers and, as a result, a great number of scatterers migrate in or out of each resolution cell as resolution is varied. Hence, one would expect SAR imagery of grassy regions to display less scale-to-scale correlation.

To complete the models, we must also specify the distributions for the w(s) in (2.16) for each scale and each terrain type. For imagery of grass, we have found that a Log-Rayleigh distribution,

$$p_{w(s)}(W(s)) = k \exp \left[k W(s) - \gamma - \exp \left(k W(s) - \gamma\right)\right]$$

$$k = \frac{\ln(10)}{10}$$

$$\gamma \approx 0.57721566 \text{ (Euler's constant)},$$

provides a good fit for the residuals at each scale. The residuals for imagery of forest, on the other hand, was best approximated with zero-mean Gaussian distributions. In particular, the standard deviations chosen for the Gaussian densities representing w(s) for each of the three scales predicted are given below in Table 3.2.

Resolution	σ	
$\delta imes \delta$	5.3724	
$2\delta imes 2\delta$	6.1811	
$4\delta \times 4\delta$	6.5056	

Table 3.2: Forest Model Residual Standard Deviations

In Figure 3-5, we display, an example of the accuracy of the fits of the chosen distributions. In particular, in Figure 3-5(a) and (b) we display normalized histograms for the residuals in the prediction of regions of grass (forest) using the grass (forest) model coefficients given in Table 3.1 as well as the chosen distributional fits (solid lines).

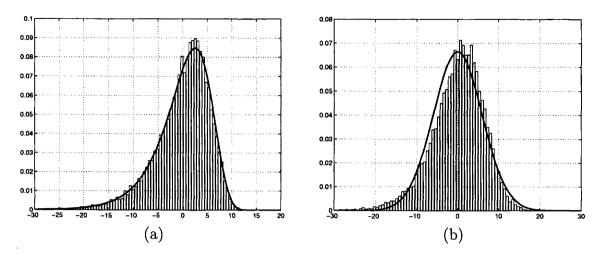


Figure 3-5: Histograms of residuals in prediction of second-finest resolution for (a) Grass model (b) Forest model. Solid line represents (a) Log-Rayleigh distribution (b) Gaussian distribution.

3.4.2 Algorithm Specification

To perform the classification and segmentation procedure described in Sections 3.2 and 3.3, we must first specify 1) The size of the block of pixels, \mathcal{B} , being classified in each step; 2) the size of the window, \mathcal{W} , around each block used to perform the classification; and 3) the thresholds, Ψ_g and Ψ_f , for each stage in the hierarchical procedure.

For the examples we display here, we somewhat arbitrarily choose to use a block size for \mathcal{B} of 4×4 pixels. This provides an order of magnitude increase in efficiency over classification of individual pixels (i.e., 1×1 blocks) with minimal resolution loss in the final segmentation. If we had used a block size of a single pixel, we would have obtained similar results, but at the cost of efficiency. As previously mentioned, the choice of the size of the window \mathcal{W} determines a balance between the statistical significance of the subsequent likelihood ratio test and the local accuracy of the results. To determine a reasonable window size, we gathered statistical data about ℓ for several square window pixel sizes and each terrain category. For each size and terrain category, we evaluated ℓ as in (3.4) in several hundred distinct situations for

which the particular window size covered homogeneous imagery of the given terrain class exclusively. This allowed us to form histograms for the statistic ℓ for each terrain category and window size, thus observing our ability to differentiate between homogeneous regions of grass and forest as window size is varied.

We display the results for window sizes of 128, 64, and 32 square pixels in Figures 3-6(A-i) through (A-i). In addition, we have also displayed to the right of each set of histograms, in Figures 3-6(B-i) through (B-iii), the corresponding Gaussian fits for each histogram. This allows us to approximate the probability of misclassification for each window size. We see virtually no overlap between the histograms for each each terrain category in Figure 3-6(A-i) when using the larger 128 pixel window. Increasing overlap (i.e., probability of misclassification) occurs as the window size decreases as is illustrated in Figures 3-6(A-ii) and (A-iii). For a 32×32 window size, we still observe distinguishing information between the two classes, yet with a much higher probability of misclassification. For the purposes of the test results here we have chosen to use an original window size of 128×128 pixels since this size provides minimal probability of misclassification. Again, we could have chosen a larger window size, but chose not to so as to preserve the local accuracy of the results. The results above also provided a measure of statistical significance versus window pixel size. Accordingly, we decided to consider subdivisions, i.e. quadrants, in the classification refinement procedure, described in Section 3.3, down to a minimal size of 32×32 pixels.

If we were only considering classification of homogeneous regions and thus needed only to set the single threshold, Ψ , for each size window, then Figure 3-6 would provide all of the information needed to do this. However, as we indicated previously, we also wish to do a better job of classifying pixels near grass-forest boundaries, and consequently a better job of boundary estimation. To meet this end, we must identify two thresholds for each window pixel size; the threshold above which we classify as grass (Ψ_g) and the threshold below which we classify as forest (Ψ_f) . These

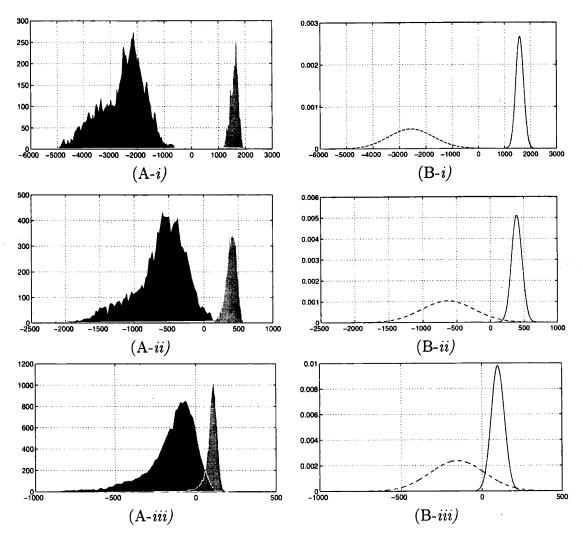


Figure 3-6: Statistical results used in determination of threshold values for each region size. Frame (A-i) displays histograms of values of ℓ from 128 pixel square homogeneous regions of forest (dark) and grass (light). Frame (B-i) shows Gaussian estimates of ℓ for each terrain category (solid line for grass and dashed for forest). Frames (A-ii) & (B-ii) and (A-iii) & (C-iii) display similar results for 64 and 32 pixel square regions respectively.

two thresholds designate the values of ℓ above and below which we observe strong evidence that the window covers homogeneous imagery of grass and forest respectively.

To observe the effects of boundary presence within the window upon the distribution of the resulting test statistic, ℓ , we gathered statistics of ℓ for the three window sizes specified above when the window region contained varying percentages of forest

and grass. In Figure 3-7(a) we display our results for the 128 pixel-square window by plotting the mean of ℓ as well as its 2σ interval versus varying percentages forest in the window (the rest of the window in each case corresponds to imagery of grass). It may be seen that the presence of even a modest amount of "contamination" by forest pixels within the grass region can alter the behavior of the test statistic significantly. For instance, in Figure 3-7(a), we see that when the window region contains approximately 50 per-cent imagery of forest, thresholding on zero will almost exclusively provide a classification as forest. This introduces an effective "bias" in the resulting decision process if only a forest-grass decision is made using the single threshold that works best for homogeneous regions of imagery ($\Psi = 0$). As a consequence, it is essential that forest-grass decisions at the level of the full 128×128 pixel region be made only when overriding evidence for one of these two hypotheses exists, i.e., when the resulting value of ℓ provides extremely strong evidence. Similarly, the decisions at the subsequent 64×64 and 32×32 pixel regions have the same characteristics. Based on these observations we are led to choosing the pairs of thresholds in Table 3.3. Here, a value above the "grass" threshold, (Ψ_g) leads to a classification as grass; a value below the "forest" threshold (Ψ_f) to classification as forest; and a value between the two, to a deferred classification. If the latter decision is made at either the 128×128 or 64×64 pixel levels, we continue down hierarchically to the next level of classification. If a deferred decision is made at the 32×32 region, we stop and give no classification to that region, as described in Section 3.3.

Region Size	Grass Threshold	Forest Threshold
128	1000	-1600
64	500	-800
32	50	0

Table 3.3: Threshold values for various window sizes. These values serve as thresholds for the likelihood ratio test for classification as grass, forest, or defer at each hierarchical level of the algorithm.

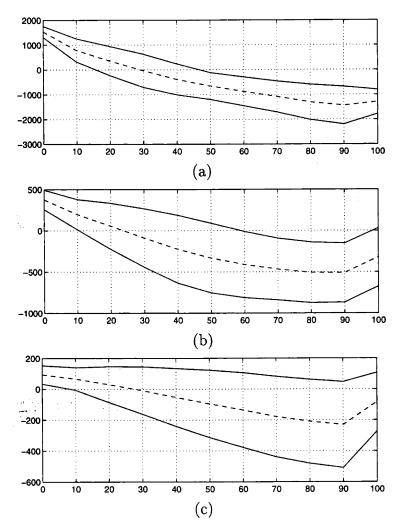


Figure 3-7: Sample mean value of the test statistic ℓ versus percentage forest in window region. Frame (a) displays plots of the test statistic mean (dashed line) as well as plots of its plus and minus 2σ points (solid lines) for a 128 pixel window. Frames (b) and (c) displays similar data for 64 and 32 pixel window sizes respectively.

3.4.3 Segmentation Performance

Given the block, window, and threshold parameters as determined in the preceding subsection, our algorithm is now completely specified, and we can apply it to SAR imagery to investigate its performance in terms of accuracy of classification and segmentation. The former of these, namely the accuracy of classification of forest and

grass pixels embedded in homogeneous regions of the same terrain type, is completely characterized by two numbers, namely the probabilities of making each type of classification error. Note that if we knew a priori that a region is homogeneous, then, based on Figure 3-6(a) we would already know that the probability of misclassification of either type is extremely small. In fact for the Gaussian probability distributions shown in Figure 3-6(a-ii), with a single threshold set at a value of 910, the probabilities of misclassifying forest as grass and grass as forest are minute and approximately equal to 1.5×10^{-5} . However, since we wish to use our algorithm in contexts in which there may be boundaries, a fairer test involves using the complete hierarchical decision procedure. To calculate the misclassification probabilities for this algorithm we have performed tests on twenty 512 pixel-square segments of SAR imagery, ten consisting of homogeneous imagery of forest and ten of grass. The resulting probabilities of misclassification were approximately 0.005 (for misclassifying grass as forest) and 0.011 (for misclassifying forest as grass). The 2σ confidence intervals for these probabilities are approximately 0.01 for misclassifying grass as forest and 0.02 for misclassifying forest as grass.

Next, we evaluate the performance of the algorithm in terms of segmentation, i.e., in terms of the accuracy in estimating the boundary between forest and grass. Figures 3-8 through 3-10 display the results of applying our algorithm to three 512 pixel-square images, each containing one or more grass-forest boundaries. The white line in each frame represents a hand-picked estimate of the boundary, and dark and light regions represent terrain classification as forest and grass respectively. In each figure we display the segmentation results for the SAR image pictured in frame (a). Frames (b) through (d) display in order: the results of classification if we were to use only a single threshold at the full 128 × 128 window level (i.e., without allowing the possibility of deferred decisions), classification results using the two thresholds as listed in Table 3.3 (mid-tone value identifies those pixels for which decision at the full 128 × 128 window level was deferred and finer subdivisions of the window were

examined), and the final results using the full hierarchical approach. Comparing parts (b) and (d) of these figures, we see that, as expected, restriction to using only the 128×128 window likelihood test leads to a bias in the estimated boundary, i.e., the estimated grass-forest boundary is pushed toward the grassy region. However, when the full hierarchical system is used, the final classification presents a more accurate identification of these boundaries.

The results in Figure 3-10(d) are not quite as promising as those in the previous two figures. This may be partially attributed to the fact that one of the primary assumptions that we make for subsequent classification is violated. Namely, we assume in Section 3.3 that the terrain boundary within the window region is relatively linear. We immediately see that the boundaries in this figure are not linear, even over the smaller 128 pixel-square blocks. This phenomenon is also present in Figure 3-9(d). Here, we see that the segmentation is degraded in the vicinity of the bend in the terrain boundary. This effect may be reduced by introducing a weighting scheme in the refinement majority rule procedure that assigns higher values to classifications closer to the central block being classified.

It is also worth noting here that the orientation of the terrain boundary with respect to the direction of illumination by the SAR sensor has a significant effect on the resulting imagery. A boundary that is observed more or less broadside frequently displays a very bright line of returns, while those facing away from the sensor often display extended shadows. Both of these are evident in the SAR images in Figures 3-9(a) and 3-10(a), and lead to larger errors in estimating the grass-forest boundaries than when the boundary has "average" illumination as in Figure 3-10(a). There are several possible manners in which to enhance performance in these cases, including the detection of bright tree lines as a line of anomalies. We illustrate this, as a possible extension to our work, in Chapter 5.

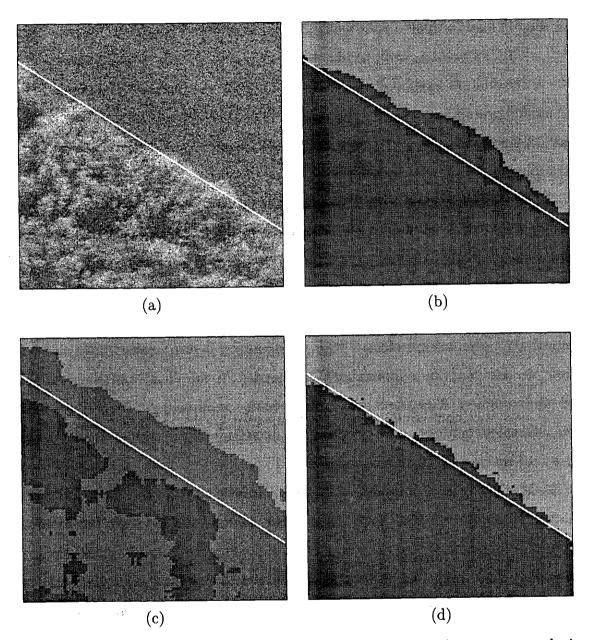


Figure 3-8: Segmentation results for a 512 pixel square image of 0.3 meter resolution stripmap SAR data. The white line in each frame represents the hand-picked estimate of the grass-forest boundary. (a) Region of original images over which segmentation was performed. (b) Initial segmentation results without refinement. (c) Preliminary classifications using two thresholds (dark = forest, light = grass, medium = deferred decision) (d) Final segmentation after refinement.

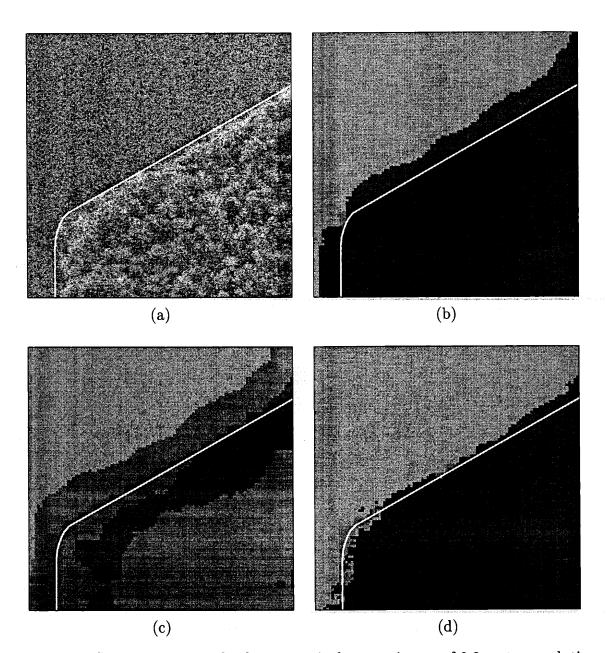


Figure 3-9: Segmentation results for a 512 pixel square image of 0.3 meter resolution stripmap SAR data. The white line in each frame represents the hand-picked estimate of the grass-forest boundary. (a) Region of original images over which segmentation was performed. (b) Initial segmentation results without refinement. (c) Preliminary classifications using two thresholds (dark = forest, light = grass, medium = deferred decision) (d) Final segmentation after refinement.

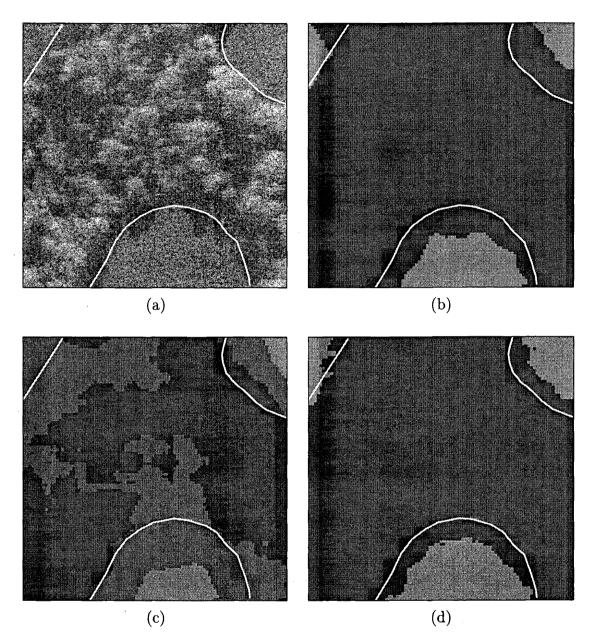


Figure 3-10: Segmentation results for a 512 pixel square image of 0.3 meter resolution stripmap SAR data. The white line in each frame represents the hand-picked estimate of the grass-forest boundary. (a) Region of original images over which segmentation was performed. (b) Initial segmentation results without refinement. (c) Preliminary classifications using two thresholds (dark = forest, light = grass, medium = deferred decision) (d) Final segmentation after refinement.

Finally, we add a quantitative performance measure that provides a feel for how accurately we are able to estimate terrain boundaries. Specifically, we define a parameter d and compute empirically the probability of misclassification of pixels that are at least d pixels away from the true boundary. Thus for d=0 we consider the misclassification probability for all pixels in the image. As d increases we ignore a swath of the image, of width 2d pixels, with d pixels on either side of the boundary. Thus, we expect the misclassification probability to decrease with increasing d, and approach the homogeneous misclassification probability evaluated above. Hence, a faster drop in misclassification probability with respect to d corresponds to a smaller effective error in estimating the boundary. In Figure 3-11 we display the results for our algorithm based on regions of Lincoln Laboratory SAR imagery containing grass-forest boundaries with average tree line illumination (circled points) and regions containing boundaries with all levels of illumination (starred points). From these tests we see that for boundaries with average illumination we are able to determine the boundary

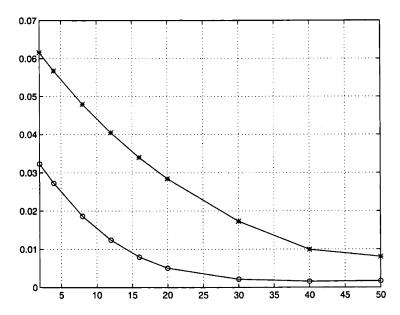


Figure 3-11: Probability of misclassification versus swath width (in pixels) on either side of grass/forest boundary for forest/grass boundaries with average tree line illumination (circled points) and boundaries with all levels of illumination (starred points).

Ø

within approximately 7 pixels (2m) with a probability of misclassification of 0.02. While if all illuminations are considered we can isolate the boundary to within approximately 27 pixels (8m) with the same probability of misclassification. Further improvement in boundary localization for either bright tree lines or ones casting shadows should be possible using the anomaly enhancement results presented in the next chapter. This concept is touched upon in Chapter 5.

Chapter 4

Anomaly Enhancement

The methods described in the previous chapter address two of the principal objectives of this thesis, namely, terrain classification and terrain boundary estimation. In this chapter we discuss a method aimed at the third of our principal objectives described in Chapter 1: the use of terrain classifications to enhance anomalies, i.e., to cause anomalies due to cultural clutter and targets to stand out from their background. Enhancement of cultural clutter image regions provides the potential for improvement in the detection of potential target regions (i.e., man-made objects) through simple thresholding. That is, higher probabilities of target detection are possible by thresholding the enhanced imagery.

A commonly accepted method to enhance anomalous pixel regions is the CFAR procedure [15, 22, 21] described in Section 4.1.1. In this procedure, a local sample mean and standard deviation are estimated from an annular region around each pixel to characterize the second-order statistics of the background. Anomalous pixels are thus identified as those that are outliers with respect to the estimated background distribution. We may pursue a similar scheme using the segmentation provided by the procedure outlined in Chapter 3. In particular, we use the prediction error residuals in conjunction with the segmentation results to identify anomalous pixels.

In this chapter we first present a brief description of the baseline ATR system being

developed at Lincoln Laboratory, and point out the utility of anomalous pixel identification. We then describe the CFAR detection routine generally used to accomplish this goal. We finish by describing our multiresolution anomalous pixel enhancement procedure and illustrate its potential with several qualitative results.

4.1 Lincoln Laboratory Baseline ATR System

The goal of any ATR system is to detect and classify objects of interest from the received sensor data with no human intervention. The object of interest, or target, varies amongst applications. For example, in military applications it may be a tactical target such as a particular tank or military vehicle. Whereas in visual navigation, the target may be some landmark representing a particular global position. The utility of an ATR system is to reduce the amount of human interaction required to identify these objects.

An ATR system [15, 22, 21] is currently under development at Lincoln Laboratory for the identification of mobile targets within SAR imagery. This system consists of three fundamental parts. The objective of each section is to reject imagery that does not contain any targets, while passing imagery that does to the next stage for further processing. Hence, each stage performs a finer scrutinization of the input imagery. In an ideal system, all targets are detected and classified and all non-targets are rejected. A high-level block diagram of this system is displayed below in Figure 4-1.

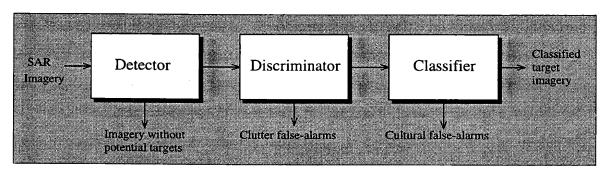


Figure 4-1: Block diagram of the Lincoln Laboratory SAR ATR algorithm

The input to this ATR system is SAR imagery formed from the raw radar returns. This imagery is sequentially passed through the three fundamental processing steps displayed above: detection, discrimination, and classification. Each step rejects portions of the input imagery that don't contain targets of interest according to specific criteria. The remaining imagery is passed to the next stage for subsequent processing. This process reduces the computational burden at each remaining stage by decreasing the overall throughput.

The first stage of this system, detection, identifies all potential targets within the input imagery. This is accomplished by applying the CFAR detection algorithm. This algorithm identifies pixels whose second-order statistics are inconsistent with their background. A chip (i.e., $n \times m$ pixel region) is extracted from the input imagery centered on each of these selected pixels or pixel clusters. The chips are subsequently passed along to the discrimination stage for further scrutiny. This second stage, as the name implies, strives to discriminate between chips containing cultural (i.e., manmade) clutter and those consisting solely of natural clutter. Each chip classified as natural clutter is rejected in this stage and the rest (i.e., possible targets) are passed to the final stage, classification. Here, each chip is categorized, using some pattern matching measure, as one of a pre-specified set of targets or as none-of-the-above. Each chip not fitting into the pre-specified set (i.e., none-of-the-above) is rejected, and all others are classified accordingly.

The potential utility of an anomaly enhancement algorithm lies in the detection stage. Such an algorithm would perform the same function as the CFAR detection algorithm described above. Specifically, it would enhance those pixels that are statistically inconsistent with their background such that simple thresholding can be used to identify outliers (i.e. potential objects of interest). As background, we provide a brief discussion of the CFAR detection algorithm.

4.1.1 CFAR Detection

The CFAR detector calculates sample first and second order statistics of an annular region surrounding each finest-scale pixel of interest, and uses them to provide a statistic that enhances the detection of outliers with simple thresholding.

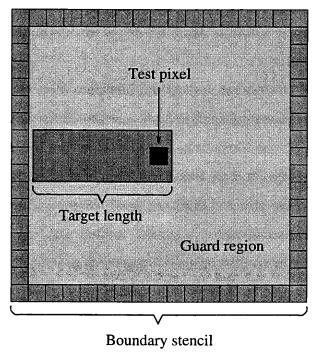


Figure 4-2: Diagram of the structure used in standard CFAR detection.

The geometric pixel-structure of the CFAR detector is displayed in Figure 4-2. The distinguishing features of the structure are: 1) the test pixel at the center; 2) the empty guard region surrounding the test pixel; and 3) the boundary stencil at the outermost edge of the guard region. The pixels in the stencil region are used to estimate the statistics of the background for the pixel of interest in the center. Therefore, the size of the guard region is generally chosen to be twice the largest target dimension in pixels. This ensures that when the center pixel lies on a target, the stencil will never overlap any portion of the same target and interfere with the estimation of background statistics.

The CFAR statistic, $\chi(k,l)$, for pixel, I(k,l) in the finest-scale log-detected (deci-

bel) image, is defined as:

$$\chi(k,l) = \frac{I(k,l) - \hat{\mu}(k,l)}{\hat{\sigma}(k,l)}. \tag{4.1}$$

where $\mu(\hat{k},l)$ and $\sigma(\hat{k},l)$ represent the background sample mean and standard deviation with the stencil centered on pixel I(k,l) respectively. Each is calculated from the pixels masked by the boundary stencil as follows:

$$\hat{\mu}(k,l) = \frac{1}{n} \sum_{(k,l) \in \mathcal{S}} I(k,l) \tag{4.2}$$

$$\hat{\mu}(k,l) = \frac{1}{n} \sum_{(k,l)\in\mathcal{S}} I(k,l)$$

$$\hat{\sigma}(k,l) = \sqrt{\frac{1}{n-1} \sum_{(k,l)\in\mathcal{S}} [I(k,l) - \hat{\mu}(k,l)]^2}$$
(4.2)

where S is the set of pixels masked by the stencil region and n is the cardinality of this set. The statistic is thus calculated by shifting this stencil over all input imagery and evaluating the statistic as in (4.1).

For pixels that are statistically consistent with their background, the CFAR statistic will have zero-mean and unit-variance. Hence, outliers (i.e., anomalies) may be detected by noting the pixels, I(k,l), for which the CFAR statistic, $\chi(k,l)$, exceeds some threshold.

4.2 Multiscale Anomaly Enhancement

As mentioned in Chapter 1, the identification of terrain boundaries and forested regions permits one to efficiently direct future ATR algorithm processing. The utility of the segmentation procedure outlined in Chapter 3 is therefore as a pre-filter for the detection stage of the three stage ATR algorithm illustrated in Figure 4-1. The segmentation identifies forested regions where targets are undetectable with high frequency SAR as well as grass-forest boundaries where subsequent processing should focus. In addition, this segmentation provides a computationally attractive method to detect potential targets of interest. Specifically, by identifying those pixels whose prediction error residuals are statistically inconsistent with the model corresponding to the pixels classification. Furthermore, since man-made objects typically exhibit very bright pixels, one would expect to observe high scale-to-scale correlation in regions where targets are present. Given the prediction error residuals calculated in the segmentation procedure of Chapter 3, as well as the concepts above, it is possible to exploit the scale-to-scale correlation to enhance and detect targets.

4.2.1 Multiscale CFAR Statistic

The ideas above suggest a method for identifying potentially anomalous pixels using what can be thought of as a multiresolution CFAR statistic. This statistic may be viewed as a synthesis of our multiscale modeling approach with the method described in [25]. Specifically, suppose a region has been identified as consisting of grass-like terrain by either the method described in the preceding chapters or some other means (such as from prior maps of the region of surveillance). Now consider a particular pixel in the finest-scale SAR imagery of that region. As described in Chapter 3, we can associate with that pixel a set of multiscale SAR imagery over a window centered at that pixel. If we let s_c denote the center node at the bottom level of the tree (the pixel under investigation), then using the parameters associated with the grass model, the residual $w(s_c)$ in (2.16) represents the error in predicting the central pixel at the finest resolution based on the SAR imagery at R coarser resolutions. Also, from our multiscale model we have a variance p_c and mean μ_c associated with this residual and can thus compute a statistic, $\zeta(s_c)$, that is analogous to the CFAR statistic,

$$\zeta(s_c) = \frac{w(s_c) - \mu_c}{\sqrt{p_c}}. \tag{4.4}$$

As with the CFAR statistic, this statistic may exploit the characteristics of the background imagery. In particular, when the pixel classification is that obtained using

the segmentation procedure in Chapter 3, we note that it is based on a window of imagery surrounding the pixel of interest. As a result, the classification characterizes the background. Accordingly, we subtract the background mean associated with the corresponding terrain model from the statistical value (each model assumes zero-mean residuals), and normalize by the corresponding terrain model standard deviation.

The multiscale segmentation algorithm computes the prediction error residuals required to determine the multiscale CFAR statistic defined in (4.4). The variance, p_c , may be evaluated from the theoretical distribution specified by the multiscale terrain models. Furthermore, because the residuals are computed at a set of scales, the scale-to-scale behavior of the CFAR statistic can be characterized and used to improve detection performance. Specifically, for a node s at the mth level from the bottom of the tree, we can also compute w(s) as in (2.16) and compare it to the corresponding variance p_m from the multiscale model. Specifically, generalizing (4.4), we can compute

$$\zeta(s) = \frac{w(s)}{\sqrt{p_m}},\tag{4.5}$$

for s on the mth level from the bottom of the tree. Note that in (4.5) we disregard the mean since for all scales of each model we assume it to be zero. By considering the set of values of $\zeta(s)$ as we traverse the path from the central pixel at the finest level back through its ancestors on the tree, we obtain a set of random variables that can be used to detect the presence of statistically significant differences between behavior around the pixel in question and that expected if the pixel were a clutter pixel and not an anomalous scatterer.

4.2.2 Combined Multiscale CFAR Statistic

By considering the set of random variables acquired by evaluating (4.5) at multiple scales, $\{\zeta(s), \zeta(s\overline{\gamma}), \dots, \zeta(s\overline{\gamma}^{P-2})\}$, we allow for several methods of obtaining statistics

to enhance the identification of anomalies. For example, one possible statistic is the sum of squared values of these random variables,

$$c_1(s_c) = \zeta^2(s_c) + \zeta^2(s_c\overline{\gamma}) + \zeta^2(s_c\overline{\gamma}^2) + \ldots + \zeta^2(s_c\overline{\gamma}^{P-2}), \tag{4.6}$$

where P represents the number of scales in the state-augmented multiscale model.

As with the standard CFAR statistic, $c_1(s_c)$ will enhance any behavior that deviates significantly from the behavior we expect for non-anomalous pixels. However, in contrast to the standard CFAR statistic, we can consider alternatives, since we have a set of statistics rather than only one. In particular, we can also consider looking for specific scale-to-scale signatures in the residuals $w(s_c)$. Indeed, as the results in [12, 25] indicate, the presence of a few dominant scatterers in man-made objects will in fact result in scale-to-scale variations in SAR imagery that have a more deterministic behavior across scale and thus do not obey a simple white noise driven autoregressive model. If we can predict this scale-to-scale behavior, we can seek the telltale signatures in the resulting residuals. A simple example of this is the statistic

$$c_2(s_c) = \left[\zeta(s_c) + \zeta(s_c\overline{\gamma}) + \zeta(s_c\overline{\gamma}^2) + \ldots + \zeta(s_c\overline{\gamma}^{P-2})\right]^2 \tag{4.7}$$

Comparing (4.6) and (4.7) the difference is that in (4.7) we sum before squaring. This corresponds to looking for a consistent over- or under-prediction of pixels at successive resolutions, as one might expect for a single strong scatterer. In such a case the statistic $c_2(s_c)$ would give a significantly larger value than $c_1(s_c)$. In addition, $c_2(s_c)$ would attenuate the effects of other types of scale-to-scale behavior. For example, if the signs of the successive values of ζ change throughout scale, then $c_2(s_c)$ will be much smaller than $c_1(s_c)$. Thus the statistic in (4.7) is far more selective in what it enhances, which will be of value if we have accurate models for what type of scatterer distributions we expect in objects of interest.

The two statistics, $c_1(s_c)$ and $c_2(s_c)$, both provide means of identifying and en-

hancing pixels that exhibit anomalous behavior. They both involve a squaring, in some sense, of the scale-to-scale statistics ($\zeta(s_c)$) though. This will enhance pixels that are outliers in the positive and negative tails of the model distribution. Due to the nature of pixels corresponding to man-made imagery, we are only concerned with the positive tail. In other words, we are only interested in those values of $c_2(s_c)$, for example, that display a consistent over-prediction. Hence, we may consider a third statistic,

$$c_3(s_c) = \left[\zeta(s_c) + \zeta(s_c \overline{\gamma}) + \zeta(s_c \overline{\gamma}^2) + \ldots + \zeta(s_c \overline{\gamma}^{P-2}) \right]. \tag{4.8}$$

This statistic will enhance only those pixels that display a consistent over-prediction and not those displaying a consistent under-prediction. As a result, we may simply threshold this statistic, as with the standard CFAR statistic, to identify potential targets of interest.

4.2.3 Implementation

The statistic in (4.8) provides a possible method of enhancing target pixels within SAR imagery. There are a few issues that still remain regarding the implementation of such a method. For example, do we compare each residual to the model distribution of the corresponding pixel's classification (grass or forest), or do we merely seek anomalous behavior over grass regions alone? It is necessary to answer such questions and, in turn, detail a feasible procedure to produce anomaly enhanced imagery.

One feasible method for anomaly enhancement may be outlined as follows:

- Multiscale segmentation.
- Morphological closing on the binary segmented image.
- Evaluation of $c_3(s_c)$ over all grass regions.

The multiscale segmentation achieves two fundamental goals; identification of forest regions and production of prediction error residuals. By identifying forest regions, we note regions over which target detection is futile. We will see that there is a minor problem in this method; target pixels are often classified as forest. We address this issue below. By generating and storing the residuals during segmentation, we improve the attractiveness of our multiscale approach. Little further computation is necessary to produce the statistic in (4.8).

A catch-22 situation arises in the method we describe above. In particular, we would like to identify grass-like regions over which to search for potential targets. Unfortunately, targets, in our two-category classification scheme, tend to be classified as forest, and would thus not be considered in our search. Moreover, it seems that in many cases, whenever a target pixel is within the classification window (as in Figure 3-1), the center pixel tend to receive a classification as forest. This may be attributed to the fact that a higher degree of scale-to-scale correlation is expected in forest regions than in grass regions due to the presence of bright scatterers. As a result, target regions are often central to block-like classifications as forest (with the block size proportional to the window size).

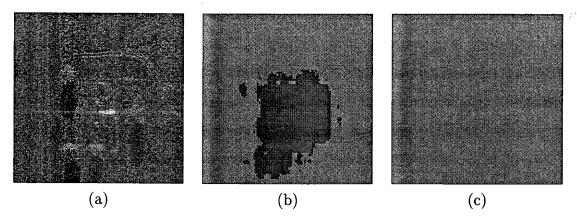


Figure 4-3: (a) Original image containing a house surrounded by several trees. (b) Results of multiscale segmentation (light and dark grey represent classifications as grass and forest respectively). (c) Classification after morphological closing operation (all pixels classified as grass).

One manner in which to counteract this phenomenon is with the use of morphological operations (see Appendix B). For example, the closing operation will reclassify all forest regions of the binary (segmented) image that have a width smaller than that of the structuring element, while preserving the general form of all other boundaries. For example, the use of a square structuring element with size proportional to the window size used in the original segmentation, will likely produce a segmentation useful for our purposes. We illustrate this point with the simple example displayed in Figure 4-3. Figure 4-3(a) shows a SAR image containing a house surrounded by several trees. As expected, the multiscale segmentation of this image in Figure 4-3(b) exhibits a sizeable area in the proximity of the dwelling classified as forest. Application of a binary morphological closing operation with a square 151 × 151 pixel stencil (to account for the segmentation window size as well as the size of the physical structure) provides the desired classification of all grass, as shown in Figure 4-3.

Small tree clusters may be reclassified as grass in the approach above and thus require investigation as potential targets. In fact this phenomenon is apparent by noting the discrepancy between the actual data and final segmentation in Figures 4-3(a) and (c) above. This, however, is of little concern since the utility of the multiscale segmentation is in identifying *extensive* forest regions for pre-filtering rather than individual trees or clusters of trees.

The statistic $c_3(s_c)$ may now be evaluated exclusively over those regions receiving a grass classification after the morphological closing of the segmented image. Since we are only interested in grass-like regions for target detection, we will use only the variance of a zero-mean Log-Rayleigh random variable (see Appendix A) in our calculations for p_m . Accordingly, the statistic $c_3(s_c)$ can be simplified as,

$$c_3(s_c) = \frac{\left[w(s) + w(s\overline{\gamma}) + \dots + w(s\overline{\gamma}^{P-2})\right]}{\sqrt{p_{LR}}}, \tag{4.9}$$

where p_{LR} denotes the variance of a Log-Rayleigh random variable. The enhanced

image that this method produces can be thresholded to determine potential detections.

4.3 Performance Comparison

We describe here a set of methods to measure the performance of the anomaly enhancement technique. We will focus on performance in the ATR environment, and, accordingly, compare it to the of the standard CFAR enhancement algorithm used in practice. As test data, we will use three HH polarized, 1 foot resolution SAR images, each containing different types of cultural clutter. We will illustrate the potential of the multiscale anomaly enhancement approach by presenting side-by-side comparisons with CFAR enhanced imagery. We further demonstrate the potential of our technique by providing qualitative measures of pixel enhancement over target regions alone. While neither of these methods yields the definitive measure of performance that a Receiver Operating Characteristic (ROC) curve would, they do demonstrate the promise that a multiscale-based anomaly enhancement technique affords.

The three SAR images used as test data are displayed in Figures 4-4 through 4-6. Each 512 pixel square image contains at least one man-made object. The highly reflective portion of each object is outlined by the black boxes in each figure. Since it is these highly reflective regions that typically set man-made objects apart from natural clutter, we will focus on the performance of each algorithm exclusively over these outlined image pixels.

For each image displayed, we evaluated both the CFAR statistic and the multiscale statistic, $c_3(s_c)$. The residuals (the w(s), $w(s\overline{\gamma})$, etc.) in (4.9) were calculated using a third-order scale-autoregressive grass model spanning four scales (P=4). The CFAR statistic was calculated using differing guard region sizes for each figure depending upon the size of the man-made objects within. Specifically, we calculated the CFAR statistic with a 121 pixel-square stencil for the imagery in Figures 4-4 and 4-5, and a

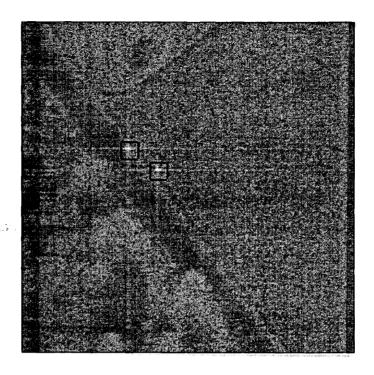


Figure 4-4: SAR imagery displaying two 8 foot corner reflectors.

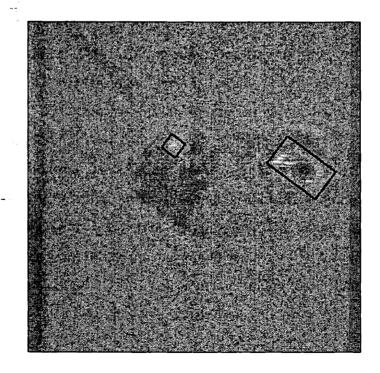


Figure 4-5: SAR imagery displaying two man-made structures.

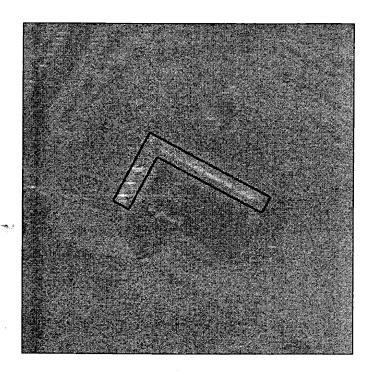


Figure 4-6: SAR imagery displaying an L-shaped structure.

201 pixel-square stencil for the imagery in Figure 4-6.

The enhanced images for each technique are displayed in Figures 4-7 through 4-9. In each figure, frame (a) and (b) show enhanced images using the CFAR enhancement method and the multiscale enhancement method respectively. In each of the three figures, it appears that the multiscale enhancement technique increases the visibility of cultural clutter regions over the CFAR enhancement method. We will see below that for these examples this is, in fact, the case. Moreover, we see that we may improve the enhancement by increasing the number of scales that the model spans. Specifically, in Figure 4-10 we show enhancement of the imagery in Figure 4-4 using each technique (multiscale enhancement in (a) and CFAR enhancement in (b)), but with the multiscale grass model increased to span six scales (P = 6). We see a clear improvement in the visibility of the corner reflectors in this enhanced image over the CFAR technique.

The blocky appearance in each of the multiscale enhanced images is an artifact of

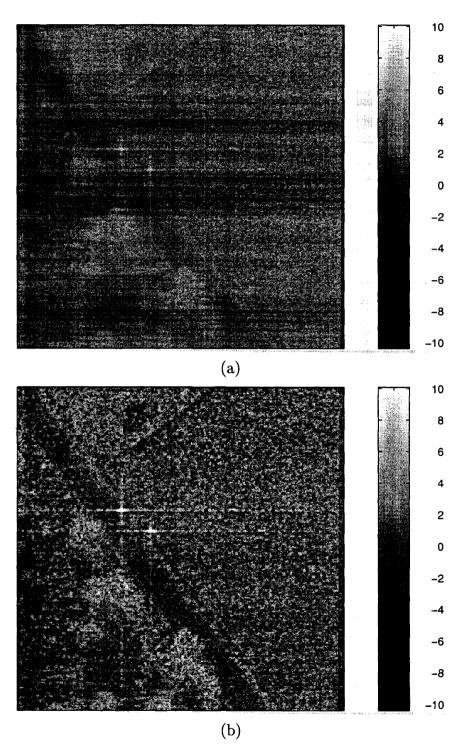


Figure 4-7: Anomaly enhanced images for imagery displayed in Figure 4-4. (a) CFAR enhancement; (b) Multiscale enhancement.

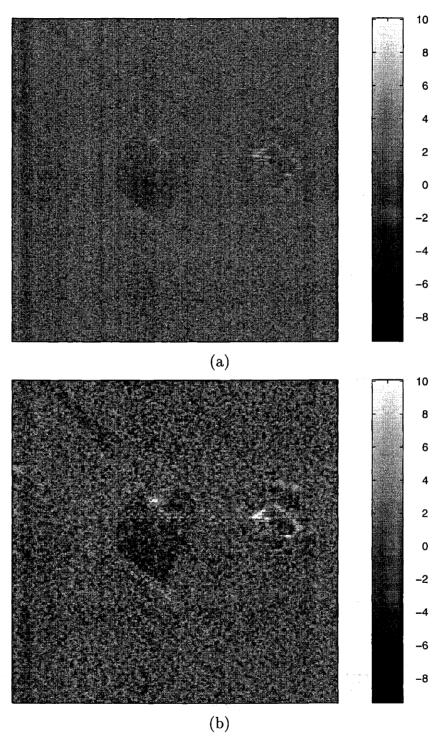


Figure 4-8: Anomaly enhanced images for imagery displayed in Figure 4-5. (a) CFAR enhancement; (b) Multiscale enhancement.

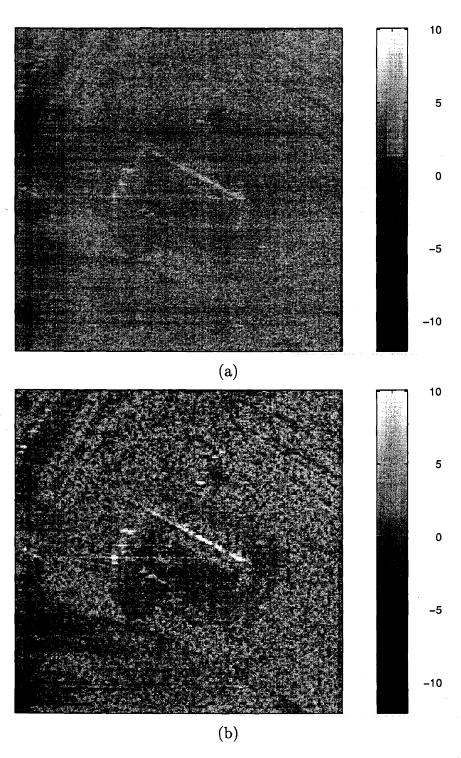


Figure 4-9: Anomaly enhanced images for imagery displayed in Figure 4-6. (a) CFAR enhancement; (b) Multiscale enhancement.

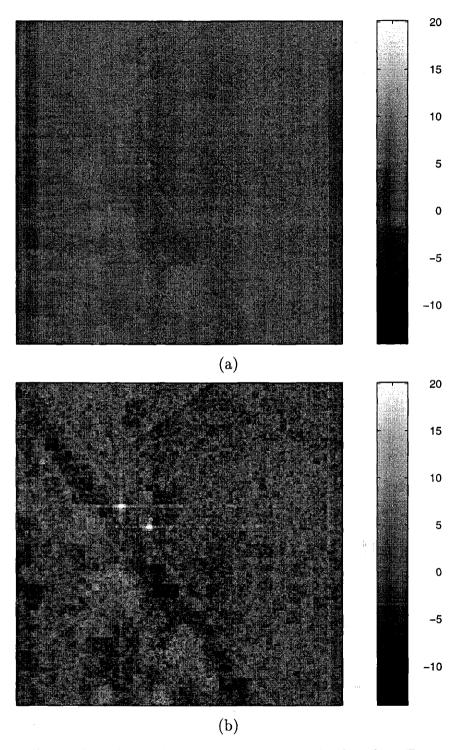


Figure 4-10: Anomaly enhanced images for imagery displayed in Figure 4-4 using a multiscale model spanning 6-resolutions. (a) CFAR enhancement; (b) Multiscale enhancement.

the multiscale structure of the statistic $c_3(s_c)$. This is apparent by considering, for example, two adjacent pixels at nodes s_1 and s_2 in the finest-scale image for which $s_1\overline{\gamma} = s_2\overline{\gamma}$. We see, by evaluating (4.9) at each node, that the enhanced pixels at these nodes may be written as,

$$c_3(s_1) = \frac{\left[v(s\overline{\gamma}) + w(s_1)\right]}{\sqrt{p_{LR}}} \tag{4.10}$$

$$c_3(s_2) = \frac{[v(s\overline{\gamma}) + w(s_2)]}{\sqrt{p_{LR}}}$$
 (4.11)

where, $v(s\overline{\gamma})$ is given as,

$$v(s\overline{\gamma}) = w(s\overline{\gamma}) + w(s\overline{\gamma}^2) + \dots + w(s\overline{\gamma}^{P-2}).$$
 (4.12)

Clearly, $c_3(s_1)$ and $c_3(s_2)$ will be highly correlated. Furthermore, similar arguments may be made for finest-scale pixels, s_3 and s_4 , for which $s_3\overline{\gamma}^2 = s_4\overline{\gamma}^2$, and so forth. Although the blocky nature of this statistic does degrade visual quality, we will find it to be of minimal consequence in the identification of objects of interest.

To perform a qualitative comparison of the two statistics, we must account for enhancement of natural clutter by each statistic due to differences in variance. For image regions that are statistically consistent with their background, the CFAR statistic will have zero mean and unit variance. The statistic, $c_3(s_c)$, on the other hand, will have zero mean but a variance of P-1 (provided that the residuals are white). We could normalize this statistic by P-1 to ensure a fair comparison, yet if the residuals are not truly white throughout scale we will not accomplish our goal. Hence, we choose an empirical method. We characterize the performance over grass regions because most forested regions will be rejected by the segmentation algorithm. Thus, we estimate the mean and standard deviation of each statistic for SAR imagery of grass. We then normalize each enhanced image such that each statistic will have zero-mean and unit variance over regions of natural clutter (grass). As a result, we may directly

compare each normalized image to determine which provides better enhancement of anomalous regions. This procedure is graphically outlined in the block diagram of Figure 4-11.

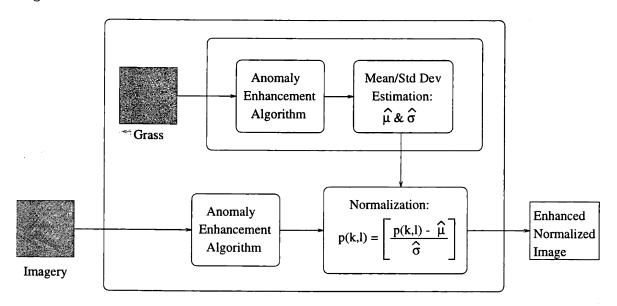


Figure 4-11: Block diagram illustrating our technique for qualitatively comparing the performance of our multiscale anomaly enhancement procedure with the standard CFAR detector.

As an initial measure, we compare the peak and average enhanced pixel values over the areas of interest (outlined in black) for each of the SAR images in Figures 4-4 through 4-6. The peak value provides a measure of probability of detection (P_D) since each image will ultimately be thresholded to detect anomalous regions, while the average value gives a measure of overall target enhancement. We evaluated three statistics: CFAR, and $c_3(s_c)$ for multiscale models spanning four and six scales. The results listed in Table 4.1 indicate the potential of a multiscale enhancement technique in increasing visibility of anomalous (i.e., man-made) regions. We see that for each SAR image, the peak value over the target is greater in the multiscale enhanced image than in the CFAR enhanced image. Furthermore, an increase in the number of model scales (4 versus 6) provides better anomalous pixel visibility.

We evaluated a second performance measure that further illustrates the ability of

	Normalized Statistic					
	CFAR		$c_3(s_c), (P=4)$		$c_3(s_c), (P=6)$	
Image	Peak	Average	Peak	Average	Peak	Average
Figure 4-4	7.40	0.91	9.61	1.13	11.32	2.13
Figure 4-5	7.38	0.65	8.50	0.94	9.97	1.62
Figure 4-6	6.62	1.03	8.19	1.32	9.82	1.97

Table 4.1: Peak and average values of each normalized statistic over target regions (outlined in black) for imagery in Figures 4-4 through 4-6.

our multiscale anomaly enhancement technique to increase the visibility of cultural clutter. This measure evaluates the number of target pixels that exceed a given threshold, for increasing thresholds. Clearly, this measure provides a performance curve that will begin at a value corresponding to the total number of target pixels and monotonically decrease to zero. The rate at which the curve decreases provides a measure of performance (i.e., a "higher" curve is indicative of increased ability for detection). We evaluated this measure for the three normalized statistics listed in Table 4.1 over each of our test images. For each of the three images, we again used the pixels outlined in black to represent target regions. The results for each of the test images in Figures 4-4 through 4-6 are displayed, respectively, in the three plots in Figures 4-12 through 4-14.

It is apparent that the multiscale technique provides better enhancement of the target regions than the CFAR algorithm. For example, for the L-shaped structure in Figure 4-6 the plotted results in Figure 4-14 show that using the multiscale enhancement technique with a four-scale model, we may set a threshold value of 8 and still detect the target. With the threshold set at this value, the CFAR algorithm will not detect this target. By allowing for higher threshold values, the multiscale enhancement technique will subsequently reduce the clutter false alarm rate and improve detection performance. Furthermore, it is clear from Figure 4-14 that the performance of the multiscale enhancement technique improves when using the larger model size (six scales).

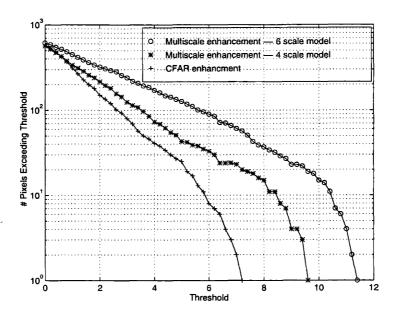


Figure 4-12: Anomaly enhancement performance plots for imagery in Figure 4-4. We display number of pixels exceeding threshold values as threshold varies for normalized enhancement statistics.

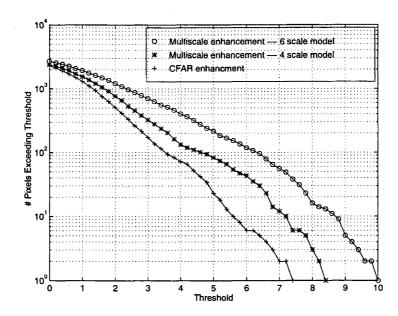


Figure 4-13: Anomaly enhancement performance plots for imagery in Figure 4-5. We display number of pixels exceeding threshold values as threshold varies for normalized enhancement statistics.

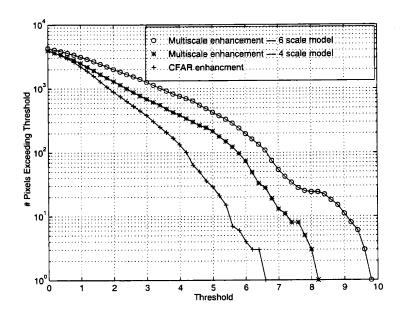


Figure 4-14: Anomaly enhancement performance plots for imagery in Figure 4-6. We display number of pixels exceeding threshold values as threshold varies for normalized enhancement statistics.

In this chapter, we have shown the potential of a multiscale technique for anomalous pixel enhancement and target detection. We compared performance with the CFAR algorithm by measuring the number of target pixels that exceeded a threshold as the threshold increased. This comparison displays the potential of a multiscale technique for anomalous pixel enhancement, yet is not entirely definitive of increased performance in an ATR system. As previously mentioned, a ROC curve is necessary to achieve such a definitive measure of performance. A ROC curve could be generated by applying the multiscale enhancement technique to several square kilometers of SAR data. The SAR data would have to contain many instances (a known quantity) of the target of interest. One could then plot probability of detection versus false alarm rate (in false alarms per kilometer squared) as a function of threshold. Such a curve could be directly compared to a ROC curve for the CFAR algorithm for a direct and definitive performance comparison.

Chapter 5

Conclusions & Extensions

Here we provide our conclusions regarding the thesis contributions and offer some natural extensions for further work.

5.1 Conclusions

We have presented here a methodology for the segmentation of background clutter and delineation of terrain boundaries as well as identification of anomalous (i.e., manmade) imagery. These methods are based on multiscale stochastic modeling of SAR imagery, and use such models to identify differences in the scale-to-scale variability and predictability of imagery of different types of terrain. We will take a look back here at each thesis contribution and offer conclusions and closing remarks.

5.1.1 Natural Clutter Segmentation

We have assembled a methodology for the segmentation of clutter in SAR imagery using multiscale SAR image models. This methodology exploits the differences in scale-to-scale variability and predictability of images of distinct clutter types due to radar speckle. Within the multiscale framework, log-likelihood test calculations are

quite efficient. Accordingly, we base the image segmentation problem on successive classifications of disjoint image regions, with the heart of each distinct classification being the aforementioned log-likelihood ratio test. Further efficiency results from the hierarchical nature of multiscale likelihood calculations. In particular, the multiscale structure yields efficient implementation of a hierarchical refinement strategy for improving the segmentation accuracy near clutter boundaries.

In Chapter 3 we outlined the segmentation procedure in detail and provided experimental results using imagery gathered with Lincoln Laboratory's millimeter-wave SAR. We observed two primary quantitative measures of performance:

- Probability of misclassification in homogeneous regions of grass or forest.
- Probability of misclassification outside swaths centered on all terrain boundaries for various swath widths.

The first measure provided a means of evaluating classification performance in homogeneous regions of natural clutter (i.e., away from terrain boundaries). We found the probabilities of misclassification to be 0.005 and 0.011 for improperly classifying grass as forest and forest as grass respectively. The second experiment furnished a performance measure on terrain boundary estimation. We concluded that, for terrain boundaries with an "average" level of illumination, we could estimate these boundaries within 7 pixels ($\approx 2m$) with 0.02 probability of misclassification. Furthermore, it was noted that boundaries displaying a higher level of illumination could likely be estimated through other means; namely, the anomaly enhancement technique outlined in Chapter 4.

We also provided, as a qualitative measure of performance, visual segmentation results for several 512×512 pixel images, each containing regions of grass and forest. For each example image we provided two primary segmentation results: 1) Simple thresholding on the statistic, ℓ , and 2) Implementation of deferred decisions via the

hierarchical refinement procedure. We observed promising improvements in boundary estimation upon implementing the refinement procedure for all images tested.

We believe that these results have demonstrated the promise of multiscale segmentation of SAR imagery. However, we do foresee several areas for improvement. In particular, in order to develop a fully automatic system, there is a need to develop a method that adapts to the nature of the SAR imagery. For example, adaptation to various squint or depression angles, types of vegetation, and presence of bright tree lines (as suggested above) would be natural extensions. Moreover, a technique that utilized fully polarimetric data would enable the incorporation of that information if available, and hence, the possibility of improved performance. This idea will be expanded upon in Section 5.2.1.

5.1.2 Anomaly Enhancement

A multiscale procedure to enhance the visibility of anomalous pixels was presented in Chapter 4. The technique is generally thought of as a by-product of the multiscale segmentation routine outlined in Chapter 3. The segmentation provides a classification for each image pixel. Correspondingly, each classification has associated with it a multiscale model. By noting those pixels whose prediction error residuals (throughout scale) are consistent outliers with respect to the associated model residual distributions, we identify anomalous pixels.

Since this technique performs a function similar to that of the standard CFAR detection technique used in the detection stage of Lincoln Laboratory's ATR system, all performance measures used the CFAR routine as a measuring stick. We evaluated two quantitative performance measures on several SAR images containing anomalous (e.g., man-made) imagery:

- Peak and average value in enhanced image over the target region.
- Number of enhanced pixels over target regions that exceeds threshold values as

threshold values vary.

For each image, we observed consistently higher peak and average values over the target area for the multiscale enhancement technique. This, however, does not translate directly to improved probability of detection. If the background clutter is also enhanced, performance may not improve or may even suffer degradation. Hence, we perform the second measure. Specifically, we measured the number of pixel values over the target regions that exceeded a threshold value for various thresholds. In order to ensure a similar probability of false alarm for grass regions, we normalized each enhanced image such that for grass-like terrain, each statistic would have zero mean and unit variance. We noted that for each man-made object tested, more pixels exceeded the threshold over the target region in the multiscale enhancement technique than in the CFAR technique for all thresholds.

Two key performance issues were also addressed. First, as a result of segmentation, the prediction error residuals are available. Hence, anomaly enhancement comes with minimal further computation. The CFAR technique, on the other hand, requires the calculation of a distinct sample mean and variance for each image pixel considered. Second, since the CFAR technique utilizes a stencil region around each pixel to gather the background statistics, a large cluster of targets could pose a potential detection problem. In contrast, this issue is of no concern for the multiscale enhancement technique considered in this thesis.

5.2 Extensions

The promise of the results obtained in this research both opens doors for future work and poses questions on techniques for improvement. As mentioned above, a natural progression for improvement in the segmentation and anomaly enhancement techniques lies in the use of fully polarimetric imagery. The use of fully polarimetric imagery ensures the fusion of all available data in the subsequent decision processes.

Moreover, our model-based segmentation leads directly to efficient techniques for SAR image compression by exploiting the scale-autoregressive nature of the terrain models. Furthermore, the anomaly enhancement techniques introduced in Chapter 4 may be fused with segmentation results in the vicinity of highly illuminated tree-lines to obtain more accurate results. We will be briefly introduce each concept here.

5.2.1 Extension to Fully Polarimetric Data

In the thesis, we formulate a methodology for the segmentation of natural clutter in SAR imagery. This formulation has assumed that we use single-polarization SAR imagery. There are, however, three polarimetric modes gathered by many SAR systems. They are listed below:

- Horizontal transmit horizontal receive (HH).
- Horizontal transmit vertical receive (HV).
- Vertical transmit vertical receive (VV).

For example, HH SAR imagery (as is used for experimentation throughout the thesis) refers to imagery gathered transmitting a horizontally polarized wavefront and measuring the horizontally polarized return. By gathering these three linear combinations, one may produce any polarization combination desired. Furthermore, additional information for classification, segmentation, and target identification may possibly be extracted by utilizing all three modes. For example, one may imagine that each mode of SAR data might contain unique information regarding distinct terrain categories. Moreover, distinguishing target information (vertical and horizontal structural makeup) may be contained in combinations of HH, HV, and VV data. We provide here the beginnings of a formulation for our image segmentation procedure that utilizes fully polarimetric SAR imagery.

Our multiscale models for single polarization SAR imagery were introduced in Chapter 2 and had the scale-autoregressive form:

$$I(s) = a_1 I(s\overline{\gamma}) + a_2 I(s\overline{\gamma}^2) + \dots + a_P I(s\overline{\gamma}^P) + w(s)$$

where P represented the order of the regression. To extend this to fully polarimetric imagery we must first define a vector, analogous to I(s) in our earlier formulation, to represent the values in all images at any multiscale node. Accordingly, we define,

$$\vec{I}(s) = [I_{HH}(s) \ I_{HV}(s) \ I_{VV}(s)]^T$$

where $I_{HH}(s)$, $I_{HV}(s)$, and $I_{VV}(s)$ represent the image values at multiscale node s in the HH, HV, and VV polarized images respectively. Our scale-recursive model now takes the form,

$$\vec{I}(s) = A_{1,m(s)}\vec{I}(s\overline{\gamma}) + A_{2,m(s)}\vec{I}(s\overline{\gamma}^2) + \dots + A_{P,m(s)}\vec{I}(s\overline{\gamma}^P) + \vec{w}(s),$$

where the $A_{i,m(s)}$ are 3×3 coefficient matrices (in the prediction of scale m(s))and $\vec{w}(s)$ represents the 3×1 vector of white driving noise. Therefore, a model is fully described by defining the coefficient matrices as well as identifying the joint probability distribution of the driving noise, $\vec{w}(s)$.

In a manner analogous to that in Chapter 2, we may translate this into a form more amenable to multiscale stochastic processes by state augmentation. If we define the state vector x(s) as,

$$x(s) = \begin{bmatrix} \vec{I}(s)^T & \vec{I}(s\overline{\gamma})^T & \cdots & \vec{I}(s\overline{\gamma}^{P-1})^T \end{bmatrix}^T,$$

the resulting multiscale dynamical equation is,

$$x(s) \ = \ egin{bmatrix} A_{1,m(s)} & A_{2,m(s)} & \cdots & A_{P-1,m(s)} & A_{P,m(s)} \ I & 0 & 0 & \cdots & 0 \ 0 & I & 0 & \cdots & 0 \ dots & \ddots & & dots \ 0 & \cdots & 0 & I & 0 \end{bmatrix} x(s\overline{\gamma}) + egin{bmatrix} I \ 0 \ 0 \ dots \ dots \ 0 \end{bmatrix} ec{w}(s).$$

At this point the problem is set up in such a manner that multiscale models may be constructed following techniques similar to those detailed in Chapter 2. Moreover, the successive regional classifications that are the heart of the segmentation procedure outlined in Chapter 3 may be performed in an identical manner. Namely, given two competing fully polarimetric terrain models, we may evaluate (again via the prediction error residuals) the log-likelihood difference of obtaining a set of fully polarimetric multiscale data given each model. This statistic may then be thresholded for classification purposes. The remainder of the procedure will remain analogous.

5.2.2 SAR Image Compression

Just as with their time series counterparts, the scale-autoregressive models we have described offer the potential for developing efficient compression algorithms. For example, given multiscale terrain models together with pixel classifications, we need only transmit the model descriptions (either by index, if one of a competing set of fixed models are used or the parameters of the model if the model is estimated on-line), the initial state (coarsest-scale image representation(s)), and the prediction error residuals, w(s). Since, as assumed in the model construction, the residuals have been approximately whitened and have a reduced dynamic range, direct scalar quantization of them should be feasible using a small number of bits. Furthermore, the image that must be encoded with a larger number of bits will correspond to the

coarsest-scale representation of the image and will thus be much smaller than the original. These two facts promise an efficient method of SAR image compression.

5.2.3 Bright Tree-line Localization

In Figure 3-11 of Subsection 3.4.3 we displayed a measure of boundary localization. Specifically, we plotted the probability of misclassification as a function of the distance from the true terrain boundary for which classifications were ignored. In this figure, we displayed the results for two boundary classes: those with average as well as high levels of illumination. We observed that boundaries with a high level of illumination exhibited extremely bright pixels along much of the boundary. As a result, the estimated boundary tended to display a bias into the grass region and performance suffered. The presence of such a distinguishable feature (bright region of scatterers), suggests that other methods, in conjunction with the multiscale segmentation, may yield improvements in boundary localization for such regions. The anomaly enhancement technique of Chapter 4 may prove useful in such cases by providing an enhanced image that may be, for example, thresholded to identify these regions. By enhancing such anomalous regions prior to thresholding, one obtains the ability to better localize the boundary.

As an example, in Figure 5-1(a) we display a SAR image with a highly distinguishable tree-line. As a result, the multiscale segmentation results for this image, given in Figure 5-1(b), exhibit a bias into the grass region. This is indicated by the hand-picked estimate (black line) in each figure. In Figures 5-1(c) and (d) we display the anomaly enhanced image and the result of processing this enhanced image with thresholding and subsequent morphological opening. We see that such processing provides an improvement in boundary estimation. We note, however, that to bring a boundary estimation improvement technique such as this to fruition, one must first possess a means of identifying those boundaries exhibiting a high level of illumination. This simple example illuminates two significant points:

5.2. Extensions 103

• Estimation of bright tree-lines may be improved with methods other than our multiscale segmentation.

Multiscale anomaly enhancement may provide a feasible means for this improvement.

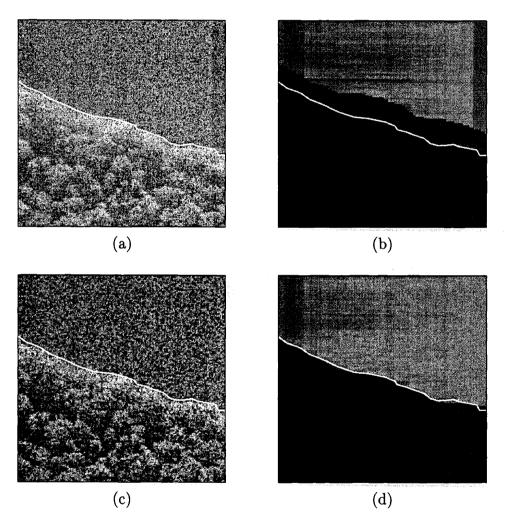


Figure 5-1: Example illustrating the potential for improvement of boundary estimation. (a) SAR image containing bright tree-line; (b) Multiscale segmentation results; (c) Multiscale anomaly enhanced image; (d) Results after performing thresholding and morphological opening on enhanced image. The white line in each frame represents a hand-picked estimate of the boundary in frame (a).

Appendix A

Log-Rayleigh Variance

To perform the anomaly enhancement calculations of (4.9) in Chapter 5, we must obtain the variance for our residual distributions in regions of grass. Our multiscale model for grass defines the residuals at all scales to be zero mean Log-Rayleigh random variables. Thus, we need only evaluate the variance of a Log-Rayleigh random variable, characterized by the probability density function,

$$p_x(X) = k c \exp\{kX - c \exp\{kX\}\},$$
 (A.1)

where,

$$k = \frac{\ln(10)}{10}$$

$$c = \exp\{-\gamma\},$$

$$\gamma \approx 0.57721566 \text{ (Euler's constant)}.$$

We begin by evaluating the moment generating function,

$$E\left[\exp\{sX\}\right] = \int_{-\infty}^{+\infty} k \ c \ \exp\{sX\} \exp\{kX - c \ \exp\{kX\}\} dX$$

$$= \int_{-\infty}^{+\infty} k \ c \ \exp\{(s+k)X - c \ \exp\{kX\}\} \ dX. \tag{A.2}$$

By choosing the change of variables,

$$u = \exp\{kX\}$$
$$du = ku dX,$$

(A.2) becomes,

$$E\left[\exp\{sX\}\right] = \int_0^{+\infty} c \, u^{\left(\frac{s}{k}\right)} \exp\left\{-uc\right\} du$$
$$= c^{-\left(\frac{s}{k}\right)} \, \Gamma\left(\frac{s}{k} + 1\right), \tag{A.3}$$

where $\Gamma(x)$ represents the Gamma function and is defined as,

$$\Gamma(x) = \int_0^{+\infty} \exp\{-t\} t^{x-1} dt. \tag{A.4}$$

As a check, we ensure that the mean of this Log-Rayleigh random variable is zero.

$$E[X] = \frac{d}{ds} \left(E\left[\exp\left\{sX\right\}\right] \right) \Big|_{s=0}$$

$$= \frac{1}{k} c^{-\left(\frac{s}{k}\right)} \left[\Gamma'\left(\frac{s}{k}+1\right) - \ln(c)\Gamma\left(\frac{s}{k}+1\right) \right] \Big|_{s=0}$$

$$= \frac{1}{k} c^{-\left(\frac{s}{k}\right)} \Gamma\left(\frac{s}{k}+1\right) \left[\Psi\left(\frac{s}{k}+1\right) - \ln(c) \right] \Big|_{s=0}$$

$$= \frac{1}{k} \Gamma(1) \left[\Psi(1) - \ln(c) \right]$$

$$= \frac{1}{k} \left[-\gamma - (-\gamma) \right]$$

$$= 0,$$
(A.5)

where $\Psi(x)$ is the Psi function defined as,

$$\Psi(x) = \frac{d}{dx} \ln \{\Gamma(x)\}\$$
$$= \frac{\Gamma'x}{\Gamma x},$$

and

$$\Gamma'(x) = \int_0^{+\infty} \exp\{-t\} \ln(t) t^{x-1} dt.$$
 (A.7)

In (A.5) we apply the basic identities,

$$\Psi(1) = -\gamma$$

$$\Gamma(1) = 1.$$

The second moment (i.e., variance since X has zero mean) is calculated as the second derivative, with respect to s, of (A.3) evaluated at s = 0. To evaluate this, we exploit the results obtained in (A.5).

$$E\left[X^{2}\right] = \frac{d^{2}}{ds^{2}} \left(E\left[\exp\left\{sX\right\}\right]\right)\Big|_{s=0}$$

$$= \frac{d}{ds} \left(\frac{1}{k} c^{-\left(\frac{s}{k}\right)} \Gamma\left(\frac{s}{k}+1\right) \left[\Psi\left(\frac{s}{k}+1\right) - \ln(c)\right]\right)\Big|_{s=0}$$

$$= \frac{1}{k^{2}} c^{-\left(\frac{s}{k}\right)} \Gamma\left(\frac{s}{k}+1\right) \left\{\left[\Psi\left(\frac{s}{k}+1\right) - \ln(c)\right]^{2} + \Psi'\left(\frac{s}{k}+1\right)\right\}\Big|_{s=0} (A.8)$$

$$= \frac{1}{k^{2}} \Gamma(1) \left\{\left[\Psi(1) - \ln(c)\right]^{2} + \Psi'(1)\right\}$$

$$= \frac{1}{k^{2}} \Psi'(1) \qquad (A.10)$$

$$= \frac{\pi^{2}}{6k^{2}}$$

In (A.8) – (A.10) we use the function $\Psi'(x)$, defined as,

$$\Psi'(x) = \frac{d}{dx}\Psi(x).$$

Furthermore, in (A.10) we make use of the identity,

$$\Psi'(1) = \frac{\pi^2}{6}.$$

Evaluating (A.11), the variance of a zero mean Log-Rayleigh random variable is,

$$Var[X] \approx 31.02538$$

Appendix B

Morphological Operators

We begin by considering the sets, A and B, each defined in a two-dimensional Euclidean space. The sets, A and B, represent the binary image to be processed and the *structuring element* to perform the processing, respectively. Different choices of B, coupled with the various morphological operations provide for many possible outcomes. These operations may, in turn, be used to perform useful image processing operations on the binary image, A.

The fundamental binary morphological operation, erosion, is defined in settheoretic notation as,

$$A \ominus B = \{x : B + x \subset A\}, \tag{B.1}$$

where \ominus is the erosion operator. The erosion operation is illustrated graphically in Figure B-1. This operation may be viewed graphically as those vectors, x, by which B may be translated such that B is contained entirely within A. The set-theoretic complement of erosion is dilation, defined as,

$$A \oplus B = \left[A^C \ominus (-B) \right]^C, \tag{B.2}$$

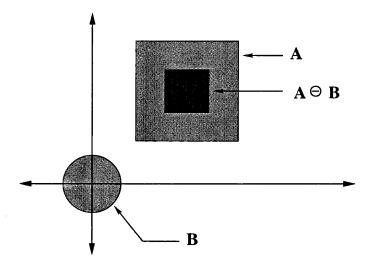


Figure B-1: Binary morphological erosion, $A \ominus B$, of set A by structuring element B.

where \oplus is the dilation operator and A^C denotes the complementary set of A. Whereas erosion is viewed as a shrinking of sets, dilation, as the name implies, is associated with the growth of sets.

Binary morphological *opening* and *closing* are defined using these two basic operations. Specifically, opening and closing are defined as,

$$A \circ B = (A \ominus B) \oplus B \tag{B.3}$$

$$A \bullet B = [A \oplus (-B)] \ominus (-B), \tag{B.4}$$

where o and • represent the opening and closing operators respectively. Again, as with erosion and dilation, opening and closing are dual operations. That is,

$$A \circ B = \left[A^C \bullet B \right]^C \tag{B.5}$$

$$A \bullet B = \left[A^C \circ B \right]^C. \tag{B.6}$$

Opening and closing perform a nonlinear filtering on the set A. Opening filters out all portions of the set that are smaller in size than the structuring element. Closing, on the other hand, filters out all portions of the set A^C that are smaller than

the structuring element. Each operation is useful in different situations. Opening may be used to reduce *pepper* noise while closing assists in the reduction of *salt* noise. Application of the two operations to sets exhibiting each category of noise are displayed in Figures B-2 and B-3. It is apparent that in these two cases the object embedded in noise is recovered.

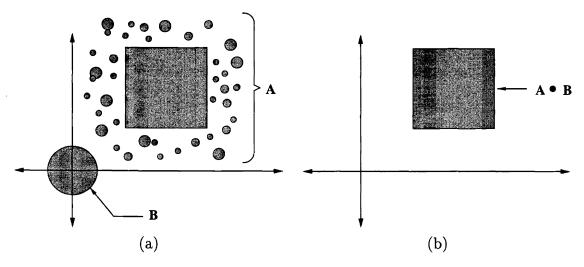


Figure B-2: (a) Binary set A displaying pepper noise and structuring element B. (b) Results after morphological opening of A with B.

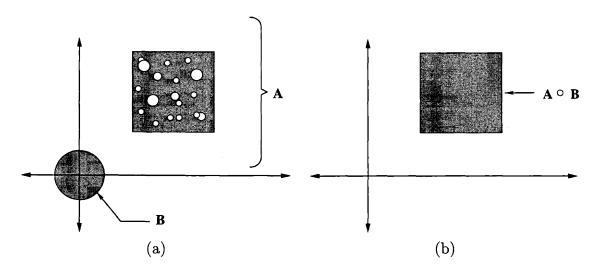


Figure B-3: (a) Binary set A displaying salt noise and structuring element B. (b) Results after morphological closing of A with B.

Bibliography

- [1] M. Basseville et al. "Modeling and Estimation of Multiresolution Stochastic Processes". *IEEE Transactions on Information Theory*, 38(2):766–784, 1992.
- [2] Walter G. Carrara, Ron S. Goodman, and Ronald M. Majewski. Spotlight Synthetic Aperture Radar. Artech House, 1995.
- [3] K.C. Chou, A.S. Willsky, A. Beneviste, and M. Basseville. "Recursive and Iterative Estimation Algorithms for Multiresolution Stochastic Processes". In *Proceedings of the IEEE Conf. on Decision and Control*, 1989.
- [4] John C. Curlander and Robert N. McDonough. "Synthetic Aperture Radar, Systems and Signal Processing". Wiley, 1991.
- [5] Edward R. Dougherty. An Introduction to Morphological Image Processing. SPIE Optical Engineering Press, 1992.
- [6] Edward R. Dougherty. Mathematical Morphology in Image Processing. Marcel Dekker, Inc., 1993.
- [7] Edward R. Dougherty and Jaakko Astola. An Introduction to Nonlinear Image Processing. SPIE Optical Engineering Press, 1994.
- [8] P.W. Fieguth, W.C. Karl, A.S. Willsky, and C. Wunsch. "Multiresolution Optimal Interpolation and Statistical Analysis of TOPEX/POSEIDON Satellite Altimetry". IEEE Trans. on Geoscience and Remote Sensing, 33(2), 1995.

- [9] S. Geman and D. Geman. "Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images". *IEEE Trans. PAMI*, 6(6), 1984.
- [10] I.S. Gradshteyn and I.M. Ryzhik. Table of Integrals Series and Products. Academic Press, 1965.
- [11] John C. Henry. "The Lincoln Laboratory 35 GHz Airborne Polarimetric SAR Imaging System". In *IEEE National Telesystems Conference*, page 353, Atlanta, GA, March 26-27 1991.
- [12] William W. Irving. "Multiresolution approach to Discriminating Targets From Clutter in SAR Imagery". In *SPIE Symposium*, Orlando, FL, April 17-21 1995.
- [13] William W. Irving. Multiscale Stochastic Realization and Model Identification with Applications to Large-Scale Estimation Problems. PhD thesis, MIT, September 1995.
- [14] Kenneth Chien ko Chou. "A Stochastic Modeling Approach to Multiscale Signal Processing". PhD thesis, MIT, May 1991.
- [15] D.E. Kreithen, S.D. Halverson, and G.J. Owirka. "Discriminating Targets from Clutter". Lincoln Laboratory Journal – Special Issue on Automatic Target Recognition, 6(1), 1993.
- [16] Alberto Leon-Garcia. Probability and Random Processes for Electrical Engineering. Addison-Wesley, 1989.
- [17] Mark R. Luettgen. "Image Processing with Multiscale Stochastic Models". PhD thesis, MIT, May 1993.
- [18] Y.L. Luke. The Special Functions and Their Approximations, volume 1. Academic Press, 1969.

- [19] D.C. Munson and R.L. Visentin. A signal processing view of strip-mapping synthetic aperture radar. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 37(12), December 1989.
- [20] L.M. Novak, M.C. Burl, R.D. Chaney, and G.J. Owirka. "Optimal Processing of Polarimetric Synthetic-Aperture Radar Imagery". Lincoln Laboratory Journal, 3(2), 1990.
- [21] L.M. Novak, S.D. Halverson, G.J. Owirka, and M. Hiett. "Effects of Polarization and Resolution on the Performance of a SAR Automatic Target Recognition System". *Lincoln Laboratory Journal*, 8(1), 1995.
- [22] L.M. Novak, G.J. Owirka, and C.M. Netishen. "Performance of a High-Resolution Polarimetric SAR Automatic Target Recognition System". Lincoln Laboratory Journal - Special Issue on Automatic Target Recognition, 6(1), 1993.
- [23] L.E. Pierce, F.T. Ulaby, K. Sarabandi, and M.C. Dobson. "Knowledge-Based Classification of Polarimetric SAR Images". *IEEE Transactions on Geoscience* and Remote Sensing, 32:1081–1086, 1994.
- [24] Merrill I. Skolnik. Introduction To Radar Systems. McGraw-Hill, 1980.
- [25] N.S. Subotic, L.M. Collins, J.D. Gorman, and B.J. Thelen. "A Multiresolution Approach to Target Detection in Synthetic Aperture Radar Data". In Asilomar Conf., Monterrey CA, 1994.
- [26] Harry L. Van Trees. Detection, Estimation, and Modulation Theory. Number Part I. John Wiley & Sons, 1968.