

April, 1996

LIDS-P 2327

Research Supported By:

ARPA contract F19628-95-C-0002

AFOSR grant F49620-93-I-0604

AFOSR grant F49620-05-1-0083

Multiscale segmentation of SAR imagery*

Fosgate, C.H.

Krim, H.

Willsky, A.S.

Irving, W.W.

Chaney, R.D.

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE APR 1996		2. REPORT TYPE		3. DATES COVERED 00-04-1996 to 00-04-1996	
4. TITLE AND SUBTITLE Multiscale segmentation of SAR imagery				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Massachusetts Institute of Technology, Laboratory for Information and Decision Systems, 77 Massachusetts Avenue, Cambridge, MA, 02139-4307				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 12	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Multiscale segmentation of SAR imagery*

C.H. Fosgate H. Krim A.S. Willsky

MIT Laboratory for Information and Decision Systems
Cambridge, Massachusetts 02139

W.W. Irving

Alphatech, Inc.
Burlington, MA 01803

R.D. Chaney

Lincoln Laboratory, Massachusetts Institute of Technology
Lexington, Massachusetts 02173-9108

ABSTRACT

In this paper, we propose an efficient multiscale approach for the segmentation of natural clutter, specifically grass and forest, in synthetic aperture radar (SAR) imagery. This method exploits the coherent nature of SAR sensors. In particular, we exploit the characteristic statistical differences in imagery of different clutter types, as a function of scale, due to radar *speckle*. 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 category of clutter of interest (i.e. grass and forest), and use these models to segment the imagery into these two clutter classes. The scale-autoregressive nature of the models allows extremely efficient calculation of the relative likelihoods of different clutter classifications for windows of SAR imagery, and we use these likelihoods as the basis for classifying image pixels and for accurately estimating forest-grass boundaries. We evaluate the performance of the technique by testing it on 0.3 meter SAR data gathered with the Lincoln Laboratory Millimeter-Wave SAR.

Keywords: Multiscale, SAR, Speckle, Segmentation

1 INTRODUCTION

In recent years, there has been a growing interest in synthetic aperture radar (SAR) imaging for applications ranging from remote sensing to surface surveillance and automatic target recognition (ATR). For applications such as these, the classification of various categories of clutter is quite important, and their delineation (i.e. segmentation) can play a key role in the subsequent analysis for target detection, recognition, and image compression. In light of typical coverage rates (exceeding $1 \text{ km}^2/\text{s}$) of an air-borne SAR, it is of great importance to devise efficient (preferably parallelizable) algorithms capable of meeting the daunting computational demands of the resulting data collection.

*This work was sponsored in part by the Advanced Research Projects Agency under Air Force Contract F19628-95-C-0002 and Air Force Office of Scientific Research Grant F49620-93-1-0604 and in part by Air Force Office of Scientific Research under Grant F49620-95-1-0083. Opinions, interpretations, conclusions, and recommendations are those of the authors and are not necessarily endorsed by the United States Air Force.

In this paper, we apply a multiscale approach to the SAR image segmentation problem that exploits the coherent nature of SAR image formation. In particular, we build on the idea of characterizing and exploiting the scale-to-scale statistical variations in SAR imagery due to radar speckle.^{6–8} A SAR image measures the coherent sum of the returns from all radar scatterers within each resolution cell (pixel). As the resolution of the imagery changes, the set of complex scatterers contributing to each resolution cell changes; similarly, the relative contribution of each scatterer to the resolution cell changes. Because this is a coherent process, there can be constructive and destructive interference among the scatterers. Consequently, there is a statistical variation in the imagery from resolution to resolution (i.e. from scale to scale). The statistics of the variation depend on the distribution of the scattering elements in the environment. More specifically, the statistics depend on the type of the clutter being imaged.

To exploit this phenomenon, we employ a recently introduced class of multiscale stochastic models.^{1,2} These models provide a powerful framework for describing random fields that evolve in scale. The 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 correspond to the pixels, in the same location, at the next finer scale. The statistical variability of the multiscale imagery is captured by a scale-recursive stochastic model for each clutter type.

In this paper, we hypothesize that the scale-to-scale variation of distinct clutter types will differ in a statistically significant manner. If this is the case, very efficient algorithms associated with the multiscale models can be applied to calculate likelihoods for the classification of individual SAR image pixels and the subsequent segmentation of SAR imagery. We demonstrate the utility of the multiscale methodology for the segmentation of regions of trees and forest from open fields and grass. Such segmentation can be useful for ATR systems. For example, if a densely forested region is identified, performing target detection in such a region is unnecessary because high-frequency SAR is incapable of providing imagery of targets under the forest canopy. Moreover, targets of interest often attempt to conceal themselves near tree lines. Thus, accurate estimation of tree lines can be used to focus attention of ATR algorithms on areas of particular interest.

In the next section, we describe the multiscale framework and its application to SAR image segmentation. In Section 3, we describe the segmentation of SAR imagery into different clutter classes. In Section 4, we evaluate the performance of this approach when applied to 0.3 meter resolution SAR imagery collected by the Lincoln Laboratory Millimeter-Wave SAR.⁵

2 MULTISCALE MODELS OF SAR IMAGERY

This section describes a multiscale framework for analyzing SAR imagery. Subsection 2.1 describes the mapping of multiscale imagery onto a *quadtree* structure. Subsection 2.2 describes a class of stochastic models for describing and analyzing multiscale processes that are mapped onto quadtrees. Subsection 2.3 specifies an autoregressive model for SAR imagery. And, Subsection 2.4 describes the estimation of parameters within the multiscale SAR models.

2.1 Quadtree Interpretation of SAR Imagery

The starting point for model development is a multiscale sequence $I_L, I_{L-1}, \dots, I_0$ of SAR images, where I_L and I_0 correspond to the coarsest and finest resolution images, respectively. The resolution varies dyadically between images at successive scales. More precisely, we assume that the finest-scale image I_0 has a resolution of $\delta \times \delta$ and consists of an $N \times N$ array of pixels (with $N = 2^M$ for some M). Hence, each coarser resolution image I_m has $2^{-m}N \times 2^{-m}N$ pixels and resolution $2^m\delta \times 2^m\delta$. Each pixel, $I_m(k, l)$, is obtained by taking the coherent sum of complex fine-scale imagery over $2^m \times 2^m$ blocks, performing log-detection (computing 20 times the log-magnitude), and correcting for DC gain variations by subtracting the mean value.³ (Note that the imagery is converted to log-magnitude because multiscale recursive models, described below, have empirically proved most effective when using this representation. Use of the




Figure 1: Sequence of three multiscale SAR images mapped onto a quadtree. The pixel value at scale m and position (k, l) is denoted by $I_m(k, l)$.

complex imagery directly is impractical because of the variability of the phase in the imagery. Also, the log-magnitude of the imagery provides more well-behaved residuals, in a statistical sense, than the magnitude imagery.) Each pixel in image I_m corresponds to four “child” pixels in image, I_{m-1} . This indicates that a 4^{th} order tree, or *quadtree*, is natural 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 1 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 .

2.2 Multiscale Stochastic Models

In this subsection, we describe a general multiscale modeling framework^{1,2} and its applicability to the SAR quadtree representation. Under this framework, a multiscale process is mapped onto nodes of a q^{th} order *tree*, where q is dependent upon the manner in which the process progresses in scale. A q^{th} order tree is one in which each *node*, starting at some root node, branches off to q child nodes. (As described above, the appropriate representation for a multiscale SAR image sequence is $q = 4$, a quadtree.) Each level of the tree can be viewed as a distinct scale 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). Letting s denote any node on the tree and $s\bar{\gamma}$ its parent, the state elements at these nodes are related by the coarse-to-fine recursion

$$\mathbf{x}(s) = \mathbf{A}(s)\mathbf{x}(s\bar{\gamma}) + \mathbf{B}(s)\mathbf{w}(s). \quad (1)$$

In this recursion, $\mathbf{A}(s)$ and $\mathbf{B}(s)$ are matrices of appropriate size and the term $\mathbf{w}(s)$ represents white driving noise. The matrix $\mathbf{A}(s)$ captures the deterministic progression from node $s\bar{\gamma}$ to node s , i.e., the part of $\mathbf{x}(s)$ predictable from $\mathbf{x}(s\bar{\gamma})$, while the term $\mathbf{B}(s)\mathbf{w}(s)$ represents the unpredictable component added in the progression. An attractive feature of this framework is the efficiency it provides for signal processing algorithms. This stems from the *Markovian* property of the multiscale model class, which states that, conditioned on the value of the state at any node s , the processes defined on each of the distinct subtrees extending away from node s are mutually independent.

For the application of segmenting different types of clutter in SAR imagery, a multiscale model can be constructed for each clutter class. To specify each model, it is necessary to determine the appropriate coefficients in the matrix, $\mathbf{A}(s)$, and the statistical properties of the driving noise, $\mathbf{w}(s)$. Once the models have been specified, a likelihood ratio test can be derived to segment the imagery into the clutter classes.

Consider the problem of segmenting forested regions from grass regions in SAR imagery. For each pixel in the image, we choose between two hypotheses: the pixel is part of a grass region (H_g) or a forest region (H_f). The log likelihood ratio test for classifying each pixel based on multiscale imagery is given by

$$\ell = \log [p_{I_L, I_{L-1}, \dots, I_0 | H_g}(I_L, I_{L-1}, \dots, I_0 | H_g)] - \log [p_{I_L, I_{L-1}, \dots, I_0 | H_f}(I_L, I_{L-1}, \dots, I_0 | H_f)] . \quad (2)$$

By invoking the state space interpretation of multiscale imagery and exploiting the Markovian property of the multiscale models, the log likelihood ratio test for the two competing hypotheses can also be written,

$$\ell = \sum_s \log [p_{\mathbf{x}(s) | \mathbf{x}(s\bar{\gamma}), H_g}(\mathbf{X}(s) | \mathbf{X}(s\bar{\gamma}), H_g)] - \sum_s \log [p_{\mathbf{x}(s) | \mathbf{x}(s\bar{\gamma}), H_f}(\mathbf{X}(s) | \mathbf{X}(s\bar{\gamma}), H_f)] . \quad (3)$$

Here, $p_{\mathbf{x}(s) | \mathbf{x}(s\bar{\gamma}), H_g}$ and $p_{\mathbf{x}(s) | \mathbf{x}(s\bar{\gamma}), H_f}$ are the conditional distributions for $\mathbf{x}(s)$ given $\mathbf{x}(s\bar{\gamma})$ for the two hypothesized models. In the next section, we will show that this likelihood test can be computed efficiently in terms of the distributions for $\mathbf{w}(s)$ under the two hypotheses.

2.3 Scale-Autoregressive SAR Model

In this paper, we focus on a specific class of multiscale models, namely scale-autoregressive models^{6,7} of the form

$$I(s) = a_1(s)I(s\bar{\gamma}) + a_2(s)I(s\bar{\gamma}^2) + \dots + a_P(s)I(s\bar{\gamma}^P) + w(s), \quad a_i(s) \in \mathbb{R} \quad (4)$$

where $w(s)$ is white driving noise. For homogeneous regions of texture, the prediction coefficients (the $a_i(s)$ in (4)) are constant with respect to image location for any given scale. That is, the coefficients, $a_1(s), \dots, a_P(s)$, depend only on the scale of node s (denoted by $m(s)$), and thus will be denoted by $a_{1,m(s)}, \dots, a_{P,m(s)}$. Furthermore, the probability distribution for $w(s)$ depends only on $m(s)$. Thus, specifying both the scale-regression coefficients and the probability distribution for $w(s)$ at each scale completely specifies the model.

Following the procedure of state augmentation used in converting autoregressive time series models to state space models, we associate to each node s a P -dimensional vector of pixel values, where P is the order of the regression in (4). The components of this vector correspond to the SAR image pixel associated with node s and its first $P - 1$ ancestors. Specifically, we define

$$\mathbf{x}(s) = [I(s) \quad I(s\bar{\gamma}) \quad \dots \quad I(s\bar{\gamma}^{P-1})]^T . \quad (5)$$

The recursion in (1) takes on the form

$$\mathbf{x}(s) = \begin{bmatrix} a_{1,m(s)} & a_{2,m(s)} & \dots & a_{P-1,m(s)} & a_{P,m(s)} \\ 1 & 0 & 0 & \dots & 0 \\ 0 & 1 & 0 & \dots & 0 \\ \vdots & & \ddots & & \vdots \\ 0 & \dots & 0 & 1 & 0 \end{bmatrix} \mathbf{x}(s\bar{\gamma}) + \begin{bmatrix} 1 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix} w(s). \quad (6)$$

Thus, for a model of the form (4) or equivalently (6), ℓ in (3) can be calculated using

$$p_{\mathbf{x}(s) | \mathbf{x}(s\bar{\gamma})}(\mathbf{X}(s) | \mathbf{X}(s\bar{\gamma})) = p_{w(s)}(W(s)), \quad (7)$$

where

$$W(s) = I(s) - [a_{1,m(s)}I(s\bar{\gamma}) + \dots + a_{P,m(s)}I(s\bar{\gamma}^P)] , \quad (8)$$

and $p_{w(s)}$ is the probability distribution for $w(s)$. By substituting (7) into (3), the test statistic can be written

$$\ell = \sum_s \log [p_{w(s) | H_g}(w(s) | H_g)] - \sum_s \log [p_{w(s) | H_f}(w(s) | H_f)] . \quad (9)$$

Thus, the likelihood test can be computed from the residuals, $w(s)$.

2.4 Identification of Multiscale Models of SAR Imagery

In order to use the multiscale methodology just described, we need to identify the model parameters for each clutter class of interest, namely the model order P , the model coefficients $a_{1,m(s)}, \dots, a_{P,m(s)}$ for each scale, and the probability distribution for $w(s)$ at each scale. To accomplish this, we choose a homogeneous training region of SAR imagery representative of each clutter class being modeled. This region is subsequently processed to produce a sequence of images $I_L, I_{L-1}, \dots, I_0$. The regression coefficients for each scale k are obtained by a standard least-squares minimization,

$$\mathbf{a}_k = \arg \min_{\mathbf{a}_k \in \mathbb{R}^P} \left\{ \sum_{\{s \mid m(s)=k\}} [I(s) - a_{1,k}I(s\bar{\gamma}) - \dots - a_{P,k}I(s\bar{\gamma}^P)]^2 \right\}, \quad (10)$$

where

$$\mathbf{a}_k = [a_{1,k} \ a_{2,k} \ \dots \ a_{P,k}]^T.$$

By varying the regression order P , one can achieve a tradeoff between computational complexity and model accuracy. We have found that $P = 3$ is adequate for our purposes.

To obtain a statistical characterization, the prediction error residuals (the $w(s)$ in (4)) of the model at scale k , we evaluate the residuals in predicting scale k of the test homogeneous region. In particular, we use the $\mathbf{a}_k$ found in (10) to evaluate all $\{\mathbf{w}(s) \mid m(s) = k\}$ in (8). We then choose a theoretical distribution that provides a good fit to the normalized histogram of these residuals. In Subsection 4.1, we define the theoretical distributions chosen for the models for grass and forest and illustrate their accuracy for statistically representing the prediction error residuals.

3 NATURAL CLUTTER SEGMENTATION

In this section, we describe a procedure for the segmentation of SAR imagery consisting primarily of natural clutter. The starting point for this development is the construction of multiscale SAR models for different clutter types as described in the preceding section. Although we illustrate the approach by focusing on distinguishing forested regions from grass, the methods are readily extended to additional clutter types.

3.1 Basic Structure of the Segmentation Procedure

We employ the multiscale models constructed for forest and grass in the preceding section to classify individual pixels and subsequently segment regions of clutter. In a fully optimal method, one would also postulate a spatial random field model for each clutter category, capturing, for example, the fact that the classification of a given pixel is very likely to be the same as its neighbors.⁴ The use of such a model, however, can increase the computational complexity of the classification algorithm considerably. Consequently, we take a simpler approach that exploits the pyramidal structure of the likelihood calculation for the multiscale models. Specifically, we classify each individual pixel based on a test window of specified size surrounding that pixel.

The size of the window used in the pixel-by-pixel classification must be chosen carefully. A larger window provides a more accurate classification of homogeneous regions. Using larger windows, however, increases the likelihood that the window contains a clutter boundary. Thus, keeping the window size as small as possible is also desirable. As demonstrated in the next section, by examining the empirical distribution of ℓ in (2) over windows of various size for homogeneous regions of grass and forest, we can determine the tradeoff between classification accuracy and window size. This, in turn, allows us to choose a window size that yields adequate performance in classifying homogeneous regions of clutter.

Whenever a clutter boundary is present within a test window, however, the validity of the center pixel classification is questionable. This effect, if not taken into account, can result in a classification bias near boundaries. To address this problem, we devise a method to detect the proximity of grass-forest boundaries as well as a procedure to refine the subsequent classification. Proximity to boundaries is detected via a simple modification of the decision made based on the test statistic ℓ . Specifically, rather than comparing ℓ to a single threshold to decide on a grass-or-forest classification, we compare ℓ to the two thresholds a and b :

$$\begin{aligned} \ell &> a && \text{Classify as Grass,} \\ a &> \ell &> b && \text{Defer decision (possible boundary presence),} \\ \ell &< b && \text{Classify as Forest.} \end{aligned}$$

The resulting test structure is illustrated in Figure 2, in which the box designated “defer” corresponds to the test statistic falling between the two thresholds. In such cases, a refinement procedure, described below, is used to classify the pixel. As we describe in Section 4, the choice of thresholds a and b are determined by examining empirical distributions of ℓ for windows containing boundaries with varying fractions of forest and grass.




Figure 2: Sequence of steps involved in initial pixel classification. i) Creation of multiscale sequence from the window region. ii) Evaluation of decision statistic ℓ . iii) Thresholding to determine center pixel classification.

3.2 Deferred Pixel Classification

For pixels where the classification decision has been deferred, it is necessary to determine the appropriate classification. The structure of the multiscale likelihood calculations allows us to perform this additional task as a replication of the procedure of Figure 2 at a hierarchy of scales. Recall that the objective of this process is to classify the center pixel of a window as either forest or grass. Consequently, in a region that is likely to contain a boundary, it is necessary to determine on which side of the boundary the center pixel lies. Under the assumption that the window contains only a single, smooth boundary, the classification can be made by determining which of the two hypothesized clutter types occupies the majority of the window. Consequently, each deferred pixel is reclassified by considering separately the four quadrants of the initial window region surrounding the pixel. The classification procedure described in Subsection 3.1 is then performed on each sub-window independently. This procedure is repeated recursively on each smaller sub-window classified as “defer” until either a majority rule decision of the entire windowed region is determined or statistical significance is lost due to decreasing sub-window size.

This “progressive refinement” around boundary regions is depicted in Figure 3. In each frame, the pixel of interest is exaggerated by the solid square box in the center of the image. Figure 3(a) represents the original windowed region about the center pixel. Based on the test illustrated in Figure 2 this window has been identified as potentially containing a grass–forest boundary. Note that the correct center pixel classification for this region is grass, but due to the boundary proximity, the classification has been deferred. Figures 3(b) through 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. Note that in Figure 3(d), 9/16 of the region has been classified as grass, hence the center pixel is ultimately classified as grass.




Figure 3: Boundary pixel refinement of typical SAR window region: (a) Window region deferred due to boundary presence. b) Region divided into quadrants, with ternary classification results marked for each subregion. (c) & (d) Regions still classified as boundary further subdivided and reclassified. Classifications G, F, and D refer to grass, forest, and defer respectively. Note that 9/16 of the window region is classified as grass; thus, the pixel is classified as grass.

Due to the pyramidal structure of the multiscale models, the likelihood calculation can also be carried out in a pyramidal fashion. By judiciously organizing the calculations, the likelihood ratios for all of the sub-windows within a particular window can be calculated with the same total computational effort needed to calculate likelihood for the entire window.

4 EXPERIMENTAL RESULTS

We have applied the multiscale segmentation algorithm to 0.3 meter resolution HH polarization SAR imagery. This imagery was gathered near Stockbridge, New York with the Lincoln Laboratory Millimeter-Wave SAR.⁵ In the next subsection, we describe the construction of the models on which the subsequent experiments are based. In Subsection 4.2, we discuss the details of the algorithm design, namely the choice of window sizes and decision thresholds. In Subsection 4.3, we evaluate the performance of the algorithm for clutter segmentation.

4.1 Model Construction

In order to apply these methods, it is necessary to construct multiscale models for SAR imagery of grass and forest from homogeneous regions of clutter. We have chosen to use a third-order regression for each model and to build models for the prediction of each of the three finest resolution images ($\delta \times \delta$, $2\delta \times 2\delta$, and $4\delta \times 4\delta$, with $\delta = 0.3$ m). As described in Subsection 2.4, this implies a third-order model with a four-level tree. For the prediction of each of the three finest resolution images, three coefficients (a_1 , a_2 , and a_3) must be specified. Using the method described in Subsection 2.4, we determined the coefficient values given in Table 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 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.

Grass			
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

(a)

Forest			
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
$4\delta \times 4\delta$	0.4584	0.1292	0.0250

(b)

Table 1: Model coefficients for third order regression in scale. (a) Grass model coefficients. (b) Forest model coefficients.

To complete the models, it is necessary to also specify the distributions for $w(s)$ in (4) for each scale and each clutter type. For imagery of grass, a log-Rayleigh distribution,

$$\begin{aligned}
 p_{w(s)}(W(s)) &= k \exp[k W(s) - \gamma - \exp(k W(s) - \gamma)], \\
 k &= \frac{\ln(10)}{10}, \\
 \gamma &\approx 0.57721566 \text{ (Euler's constant),}
 \end{aligned}$$

provided a good fit. For imagery of trees, a zero-mean Gaussian distribution provided a good fit; in particular, the standard deviations chosen for the Gaussian density of $w(s)$ for each of the three scales was:

$$\delta \times \delta : 5.3724, \quad 2\delta \times 2\delta : 6.1811, \quad 4\delta \times 4\delta : 6.6056$$

Examples showing histograms and the accuracy of the resulting fits are given in Figure 4.




Figure 4: 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.

4.2 Algorithm Specification

It is necessary to specify both the size of the window used to compute the likelihood test and the thresholds used at each stage in the hierarchical procedure. As we have discussed, the choice of window size involves a tradeoff between the statistical significance of the test and the possibility of clutter boundaries intersecting the window. On row (i) of Figure 5 we have displayed histograms of the value of the test statistic ℓ computed over homogeneous regions of both clutter types using windows of different sizes (namely 128×128 , 64×64 , 32×32). In addition, we have also displayed on row (ii) the corresponding figures using Gaussian fits to the histograms of ℓ . There is virtually no overlap between

the histograms for the larger 128×128 pixel window, with increasing overlap as the window size decreases. With a 32×32 window there is distinguishing information between the two classes, but an increased overlap in the histograms compared to a 128×128 window. For the test results here, we chose a test window with size 128×128 . We also chose to consider subdivisions down to a size of 32×32 for subsequent classification of “deferred” pixels.




Figure 5: Statistical results used in determination of threshold values for each window size . Frame (A-i) displays histograms of values of ℓ from 128×128 pixel homogeneous regions of forest (dark) and grass (light). Frame (A-ii) shows Gaussian estimates of ℓ for each clutter category (solid line for grass and dashed for forest). Frames (B-i) & (B-ii) and (C-i) & (C-ii) display similar results for 64×64 and 32×32 pixel regions, respectively.

In practice, the segmentation algorithm will be confronted with nonhomogeneous imagery. When the test window contains a boundary, the likelihood test is corrupted. For example, we have found that even modest amounts of “contamination” by a forested region in a predominantly grassy region can significantly alter the test statistic ℓ . This effectively introduces a bias in the decision process. Hence, it is essential that forest-grass decisions at each hierarchical level are made only if there is overriding evidence for one of the two hypotheses. Because of this effect, pairs of threshold values are used to make a ternary classification decision, as described in Subsection 3.1: If the likelihood statistic, ℓ , is above the “grass” threshold, the pixel is classified as grass. If ℓ is below the “forest” threshold, the pixel is classified as forest. If the value falls between the two thresholds, the classification is deferred. The values of the thresholds used by the algorithm are listed in Table 2. Whenever the classification is deferred, the classification is determined the hierarchical refinement procedure defined in Subsection 3.2.

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

Table 2: Threshold values for various window sizes. These values serve as thresholds in the ternary hypothesis test for classification as grass, forest, or defer at each hierarchical level of the algorithm.

4.3 Segmentation Performance

Figure 6 displays the results of applying the algorithm to two images (i and ii). The white line in each frame represents a hand-picked estimate of the boundary, and dark and light regions represent clutter classification as forest and grass respectively. Each row shows the segmentation results for the SAR image pictured in frame (A). Frames

(B), (C), and (D) display (in order): the results of classification using only a binary test on the full 128×128 window (i.e., without allowing the possibility of deferred decisions), the result of the ternary classification at the full window size (the mid-tone value identifies those pixels for which decision at the full 128×128 window level was deferred), and the final results using the full hierarchical approach. Comparing parts (B) and (D) in each row, we see first that, as expected, restriction to using only the 128×128 window likelihood test leads to a bias in the estimated boundary: the tree line is pushed into the grassy region. However, when the full hierarchical system is used, the final segmentation yields more accurate boundary estimation.




Figure 6: Results for two images of 0.3 meter resolution stripmap SAR data. The white line in each frame represents a manual estimate of the grass-forest boundary. (A) Original images. (B) Binary segmentation results (without allowing deferred decisions). (C) Preliminary classifications using ternary classification scheme (dark grey = forest, light grey = grass, medium grey = deferred decision) (D) Final results after refinement procedure.

4.4 Computational Requirements

The multiscale segmentation algorithm is extremely efficient compared to most segmentation approaches. The multiscale algorithm requires computation of order N , where N is the number of pixels in the image to be segmented. Furthermore, the algorithm is parallelizable because distinct regions of clutter can be segmented independently. These desirable properties are a direct result of the pyramidal structure of the multiscale algorithms. In particular, the likelihood test for a particular pixel is not dependent on the outcome of its neighbors. Thus, iterative cycles to propagate the effects of neighboring pixels are unnecessary. This property makes the multiscale segmentation algorithm much more practical than most competing segmentation algorithms.

In contrast, consider simulated annealing algorithms that assume Markov random field models.⁴ Under these models, the outcome of the likelihood test for an individual pixel is highly dependent on the outcome for each of its neighbors. Consequently, such algorithms require many iterative cycles to converge. In fact, such algorithms require order Ne^N computations to converge theoretically. (This results from the fact that order e^N iterations are required

for convergence and order N computations are required for each iteration.) In practice, such algorithms require many iterative cycles to converge approximately. Consequently, simulated annealing algorithms are impractical to implement.

5 FUTURE WORK

The multiscale methodology described in this paper is applicable to other problems related to SAR imagery. For example, the approach leads directly to efficient SAR image compression algorithms. Given a set of precomputed multiscale models, the compression algorithm would determine the model that most efficiently represents each block of imagery; an index would be transmitted for each block to specify the chosen model. The residuals, $w(s)$, could be transmitted more efficiently than the original imagery because the residuals are uncorrelated by design.

In addition, the methodology is applicable to military target detection. When a target is present in the test window, the clutter models do not hold; the residuals and the likelihoods are likely to exhibit anomalous behavior. If a test for such anomalies is constructed, it could be used as a mechanism for target detection. Furthermore, it is possible to construct equivalent multiscale models for targets; this could lead to direct likelihood tests for target detection.

6 CONCLUSION

In this paper, we have described a methodology for the segmentation of clutter using multiscale models of SAR imagery. This methodology exploits the differences in scale-to-scale variability and predictability of images of distinct clutter types. Within this methodology, likelihood test calculations are very efficient. The efficiency results from hierarchical nature of the multiscale algorithms. This structure also yields efficient implementation of a hierarchical refinement strategy for improving the accuracy of segmentation near clutter boundaries.

7 REFERENCES

- [1] M. Basseville et al. "Modeling and Estimation of Multiresolution Stochastic Processes". *IEEE Transactions on Information Theory*, 38(2):766-784, 1992.
- [2] 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.
- [3] Charles H. Fosgate. "Multiscale Segmentation and Anomaly Enhancement of SAR Imagery". Master's thesis, MIT, May 1996. In preparation.
- [4] S. Geman and D. Geman. "Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images". *IEEE Trans. PAMI*, 6(6), 1984.
- [5] 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.
- [6] William W. Irving. "Multiresolution approach to Discriminating Targets From Clutter in SAR Imagery". In *SPIE Symposium*, Orlando, FL, April 17-21 1995.
- [7] William W. Irving. *Multiscale Stochastic Realization and Model Identification with Applications to Large-Scale Estimation Problems*. PhD thesis, MIT, September 1995.
- [8] 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.