
Performance of a High-Resolution Polarimetric SAR Automatic Target Recognition System

Leslie M. Novak, Gregory J. Owirka, and Christine M. Netishen

■ Lincoln Laboratory is investigating the detection, discrimination, and classification of ground targets in high-resolution, fully polarimetric, synthetic-aperture radar (SAR) imagery. This paper summarizes our work in SAR automatic target recognition by discussing the prescreening, discrimination, and classification algorithms we have developed; data from 5 km² of clutter and 339 targets were used to study the performance of these algorithms. The prescreener required a low threshold to detect most of the targets in the data, which resulted in a high density of false alarms. The discriminator and classifier stages then reduced this false-alarm density by a factor of 100. We improved target-detection performance by using fully polarimetric imagery processed by the polarimetric whitening filter (PWF), rather than by using single-channel imagery. In addition, the PWF-processed imagery improved the probability of correct classification in a four-class (tank, armored personnel carrier, howitzer, or clutter) classifier.

THE ARPA-SPONSORED WARBREAKER program is a broad-based advanced technology program to develop new weapons technology that can locate and destroy critical mobile targets such as SCUD launch systems and other highly mobile platforms. Automatic target recognition (ATR) is an important candidate technology for this effort. To address the surveillance and targeting aspects of the Warbreaker program, Lincoln Laboratory has developed a complete, end-to-end, 2-D synthetic-aperture radar (SAR) ATR system. This system requires a sensor that can search large areas and also provide fine enough resolution to detect and identify mobile targets in a variety of landscapes and deployments.

The Lincoln Laboratory ATR system has three basic stages: detection (or *prescreening*), discrimina-

tion, and classification (see Figure 1). In the prescreening stage, a two-parameter constant-false-alarm-rate (CFAR) detector selects candidate targets in a SAR image by examining the amplitude of the radar signal in each pixel of the image. In the discrimination stage, a target-sized 2-D matched filter accurately locates the candidate targets and determines their orientation. Then texture-discrimination features (standard deviation, fractal dimension, and weighted-rank fill ratio) are used to reject natural-clutter false alarms [1]. In the classification stage, a 2-D pattern-matching algorithm rejects cultural-clutter false alarms (i.e., man-made objects that are not targets) and classifies the remaining detections by target type (tank, armored personnel carrier, or howitzer).

To evaluate the performance of the ATR system,

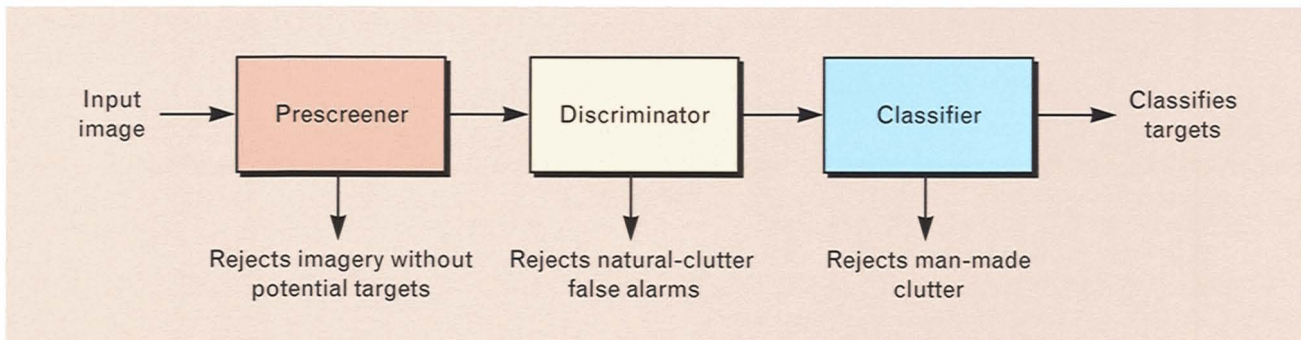


FIGURE 1. Block diagram of the SAR automatic target recognition system. The prescreener locates candidate targets based on the brightness of pixels in the input image, the discriminator rejects natural-clutter false alarms, and the classifier rejects non-target cultural clutter and classifies targets by vehicle type.

we used high-resolution (1 ft by 1 ft), fully polarimetric target data and clutter data gathered by the Lincoln Laboratory millimeter-wave SAR [2] at a depression angle of 22.5° and a slant range of 7 km. We demonstrated the robustness of the ATR system by testing it against targets both with and without radar camouflage.

Figure 2 is an example of the quality of the imagery gathered by the Lincoln Laboratory SAR. In this image of a golf course in Stockbridge, New York, the high resolution of the SAR resolves individual trees and bushes as well as small objects such as the flagpole located in the center of the putting green. This particular SAR image was obtained under clear weather conditions; the quality and resolution of the image would not have been degraded, however, by dense fog or thick cloud cover. Thus a SAR sensor has a significant advantage over optical sensors. SAR image quality is not dependent on weather conditions, and the sensor can be used at any time of day or night. In addition, SAR sensors can perform other tasks, such as searching large areas from a long distance.

The image in Figure 2 was constructed from fully polarimetric SAR data that were processed with a technique known as the polarimetric whitening filter (PWF) [3]. PWF processing optimally combines the HH (horizontal transmit, horizontal receive), HV (horizontal transmit, vertical receive), and VV (vertical transmit, vertical receive) polarization components of the radar return. This polarimetric combination enhances the quality of the imagery in two ways: (1) the amount of speckle in the imagery is minimized, and (2) the edges of objects in the image (such

as the pond) are sharper. As a result, PWF-processed imagery is visually clearer than single-polarimetric-channel imagery. In addition, PWF-processed imagery improves the performance of all three stages of the ATR system (compared with the performance achieved by using single-polarimetric-channel imagery) because PWF processing reduces clutter variance and enhances target signatures relative to the clutter background.

This article begins with an overview of the three stages of the baseline ATR system. Next we describe the performance of the ATR system with both camouflaged and uncamouflaged targets, and then we compare performance using PWF data with performance using single-channel (HH) data. We also present details of the three discrimination features used in our studies, with particular emphasis on the fractal-dimension feature. Finally, we discuss future improvements to the discrimination and classification stages.

Overview of the Baseline ATR System

This section describes our three-stage baseline SAR ATR system, which is illustrated in Figure 1 by a simplified block diagram. The three stages—prescreener, discriminator, and classifier—are described below.

Stage 1: Prescreener

In the first stage of processing, a two-parameter CFAR detector [4] is used as a prescreener; this stage of processing identifies potential targets in the image on the basis of radar amplitude (i.e., by searching for bright returns). Computation time for this stage of processing is significantly reduced by operating the

detector at a reduced resolution (1 m by 1 m) rather than at the full resolution (1 ft by 1 ft).

Figure 3 is a sketch of the two-parameter CFAR detector used by the prescreener; the detector is defined by the rule

$$\begin{aligned} \frac{X_t - \hat{\mu}_c}{\hat{\sigma}_c} &> K_{\text{CFAR}} \Rightarrow \text{target}, \\ \frac{X_t - \hat{\mu}_c}{\hat{\sigma}_c} &\leq K_{\text{CFAR}} \Rightarrow \text{clutter}, \end{aligned} \quad (1)$$

where X_t is the amplitude of the test cell, $\hat{\mu}_c$ is the estimated mean of the clutter amplitude, $\hat{\sigma}_c$ is the estimated standard deviation of the clutter amplitude,

and K_{CFAR} is a constant threshold value that defines the false-alarm rate. As shown in the figure, the test cell is at the center of a defined local region, and the 80 cells in the boundary stencil are used to estimate the mean and standard deviation of the local clutter. The guard area ensures that no target cells are included in the estimation of the clutter statistics. If the detection statistic calculated in Equation 1 exceeds K_{CFAR} , the test cell is declared to be a target pixel; if not, it is declared to be a clutter pixel.

When the amplitude distribution of the clutter is Gaussian, the CFAR detector provides a constant false-alarm rate for any given K_{CFAR} [5]. Because the clutter distributions of high-resolution data are typi-




FIGURE 2. High resolution (1 ft by 1 ft) synthetic-aperture radar (SAR) image of a golf course near Stockbridge, New York. Polarimetric whitening filter (PWF) processing was used to produce this minimum-speckle image. The radar is located at the top of the image; therefore, the radar shadows go toward the bottom of the page. Notice that the SAR can resolve details as small as the flagpole in the putting green near the center of the image.

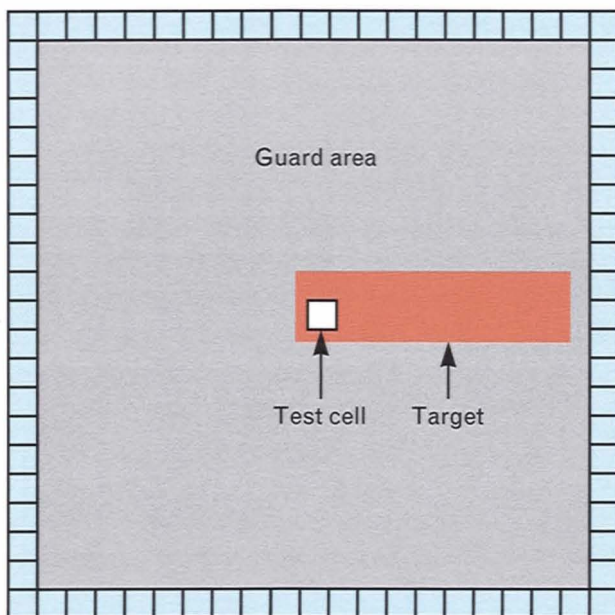


FIGURE 3. The prescreener CFAR detector. The amplitude of the test cell is compared with the mean and standard deviation of the clutter. The boundary consists of 80 cells that are used for clutter statistics estimation. Each cell in the boundary consists of 16 raw pixels that are noncoherently averaged. The guard area ensures that no target cells are included in the clutter statistics estimation.

cally not Gaussian [6], however, the detector does not always yield a constant false-alarm rate. In spite of this fact, the detector given by Equation 1 still proves to be an effective algorithm for detecting targets in clutter.

Only those test cells whose amplitudes stand out from the surrounding cells are declared to be targets. The higher we set the threshold value of K_{CFAR} , the more a test cell must stand out from its background for the cell to be declared a target. Because a single target can produce multiple CFAR detections, the detected pixels are clustered (grouped together) by the detector if they are within a target-sized neighborhood. Then a 120-ft-by-120-ft region of interest of full resolution (1 ft by 1 ft) data around each cluster centroid is extracted and passed to the discrimination stage of the algorithm for further processing.

Stage 2: Discriminator

The discrimination stage takes as its input the 120-ft-by-120-ft regions of interest passed to it by the prescreener, and it analyzes each region of interest at

full resolution (1 ft by 1 ft). The goal of discrimination processing is to reject the regions containing natural-clutter false alarms while accepting the regions containing real targets. This stage consists of three steps: (1) determining the position and orientation of a detected object, (2) computing simple textural features, and (3) combining the features into a discrimination statistic that measures how targetlike the detected object is.

In the first step of the discrimination stage the algorithm determines the position and orientation of the target by placing a target-sized rectangular template on the image. The algorithm then slides and rotates the template until the energy within the template is maximized. The position estimate produced in the discrimination stage is more accurate than the position estimate produced in the prescreening stage. This operation is computationally feasible because it is performed only on the selected high-resolution regions of interest passed by the prescreener, and not on the entire image. Mathematically, this operation is equivalent to processing the data in the region of interest with a 2-D matched filter for the case when the orientation of the target is unknown.

In the second step of the discrimination stage the algorithm calculates three textural features: (1) the standard deviation of the data within the target-sized template, (2) the fractal dimension of the pixels in the region of interest, and (3) the weighted-rank fill ratio of the data within the template. The standard deviation of the data within the template is a statistical measure of the fluctuation of the pixel intensities; targets typically exhibit significantly larger standard deviations than natural clutter. The fractal dimension of the pixels in the region of interest provides information about the spatial distribution of the brightest scatterers of the detected object. It complements the standard-deviation feature, which depends only on the intensities of the scatterers and not on their spatial locations. The weighted-rank fill ratio of the data within the template measures the fraction of the total power contained in the brightest 5% of the detected object's scatterers. For targets, a significant portion of the total power comes from a small number of very bright scatterers; for natural clutter, the total power is distributed more evenly among the scatterers.

In the third step of the discrimination stage the algorithm combines the three textural features into a single discrimination statistic; this discrimination statistic is calculated as a quadratic distance measurement (see the accompanying article entitled "Discriminating Targets from Clutter" by Daniel E. Kreithen et al.). Most natural-clutter false alarms have a large quadratic distance and are rejected at this stage. Most man-made clutter discretely (such as buildings and bridges) pass the discrimination stage; therefore, the next stage—the classifier stage—must have the ability to reject them.

Stage 3: Classifier

A 2-D pattern-matching classifier rejects cultural false alarms caused by man-made clutter discretely and then classifies target detections by vehicle type. In our studies we implemented a four-class classifier (tank, armored personnel carrier, howitzer, and clutter) using high-resolution (1 ft by 1 ft) PWF imagery. Detected objects that pass the discrimination stage are matched against stored references of the tank, armored personnel carrier, and howitzer. If none of the matches exceeds a minimum required score, the detected object is classified as clutter; otherwise, the detected object is assigned to the class (tank, armored personnel carrier, or howitzer) with the highest match score.

The pattern-matching references used in the classifier were constructed by averaging five consecutive

spotlight-mode images of a target collected at 1° increments of azimuth, yielding 72 smoothed images of each of the targets. Figure 4 shows typical pattern-matching references for the three targets at a particular aspect angle.

Performance of the Baseline ATR System

This section describes the performance of the prescreening, discrimination, and classification stages of the baseline SAR ATR system. Clutter data from 5 km^2 of ground area were processed through the ATR-system algorithms, along with data for 162 camouflaged targets and 177 uncamouflaged targets. The camouflaged target data used in this study represent a difficult scenario in which the test targets were realistically deployed and covered with radar camouflage. The training data used to design the ATR system were taken from the uncamouflaged targets. The clutter data contained a moderate number of man-made clutter discretely.

The CFAR detection threshold in the prescreener was set relatively low to obtain a high initial probability of detection (P_D) for the target data. At the output of the prescreener, $P_D = 1.00$ was obtained for the uncamouflaged targets, while $P_D = 0.82$ was obtained for the camouflaged targets. At this CFAR threshold, a false-alarm density of approximately 30 false alarms per km^2 (FA/ km^2) was obtained. The initial detection processing was carried out at reduced resolution (1 m by 1 m).




FIGURE 4. Typical pattern-matching reference templates for (a) a tank, (b) an armored personnel carrier, and (c) a howitzer. These pattern-matching references are used to classify detected objects by vehicle type. They are constructed by averaging five consecutive spotlight-mode images of a target collected at 1° increments of azimuth.

Table 1. Overview of ATR System Performance.

	FA/km^2	P_D Uncamouflaged Targets*	P_D Camouflaged Targets
After prescreening	30	1.00	0.82
After discrimination	3.0	1.00	0.75
After classification	0.3	1.00	0.70

* The uncamouflaged target test data was used for algorithm training.

Each detected region of interest (containing a potential target) was passed to the discrimination stage for further processing at full resolution (1 ft by 1 ft). The discrimination stage determined the location and orientation of each detected object, and then calculated the textural features (standard deviation, fractal dimension, and weighted-rank fill ratio) that were used to reject natural-clutter discretes. Discrimination processing reduced the false-alarm density by a factor of 10, to 3 FA/km^2 . No uncamouflaged targets were rejected by the textural-feature tests; thus the initial P_D of 1.00 was maintained. A few of the camouflaged targets were rejected by the textural-feature tests, which resulted in $P_D = 0.75$ for these targets.

In the classification stage of processing, the 2-D pattern-matcher was applied to those detections which had passed the discrimination stage. Classification processing reduced the false-alarm rate by another factor of 10, to approximately 0.3 FA/km^2 . No uncamouflaged targets were rejected by the pattern matcher (resulting in $P_D = 1.00$ for these targets), but some camouflaged targets were incorrectly classified as clutter (resulting in $P_D = 0.70$ for these targets).

Table 1 summarizes the performance of all three stages of the ATR system. The uncamouflaged target data were used for training the discrimination and classification stages. The thresholds of the algorithms were set so that perfect performance was achieved with the uncamouflaged data. Once the thresholds had been set in this way, the clutter and camouflaged targets were processed.

Figure 5 illustrates how clutter rejection was imple-

mented by the pattern-matching classifier. As the figure shows, most of the clutter discretes had correlation scores below the threshold value of 0.7, and thus were rejected (i.e., classified as clutter). Detected objects with correlation scores equal to or greater than 0.7 were considered to be sufficiently targetlike, and were classified according to their highest correlation scores (as tank, armored personnel carrier, or howitzer). Figure 5 also indicates that only a small fraction of the camouflaged targets were declared to be clutter because of low correlation scores.

The second function of the classifier is to assign

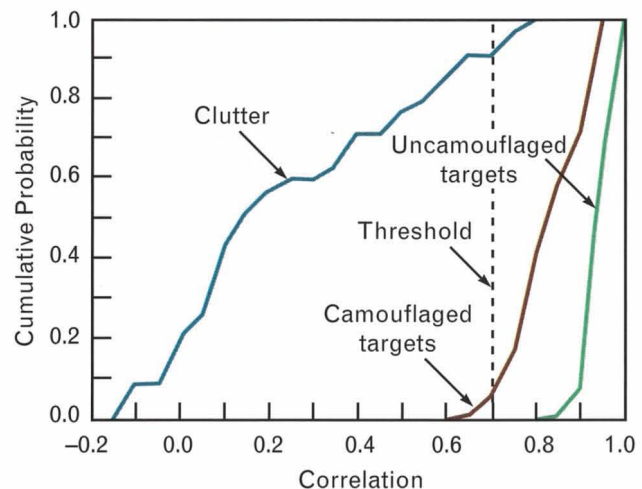


FIGURE 5. Overview of classifier performance. The clutter and target data were processed separately. Most of the clutter had correlation scores below the threshold value of 0.7, and was rejected. Detected objects above the threshold were classified according to their highest correlation scores.

Table 2. Correlation Pattern-Matching Classifier Performance

(training data: uncamouflaged targets;
test data: camouflaged targets and clutter discretized)

	Percent Classified as			
	Tank	APC	Howitzer	Clutter
Tank	89%	11%	0%	0%
APC	0%	96%	0%	4%
Howitzer	0%	13%	71%	16%
Clutter	0%	14%	0%	86%

objects accepted as targets to target classes (tank, armored personnel carrier, howitzer). Table 2 shows the classification performance of the baseline classifier as a confusion matrix that tabulates the correct and incorrect classifications. Recall that the classifier used templates constructed from uncamouflaged targets; the classification results shown in Table 2 are for clutter discretized and camouflaged test targets that passed the detection and discrimination stages. At the output of the classification stage, 70% of the camou-


flagged targets were classified as targets, and 85% of those targets were correctly classified by vehicle type.

The four-class, 2-D pattern-matching algorithm used in this study was implemented with normalized dB references, which provided the best overall performance among five different reference schemes that were tested.


ATR Performance Using Single-Channel Data versus Fully Polarimetric Data

We compared the performance of the ATR system using single-channel (HH) data with the performance of the system using fully polarimetric PWF data. Figure 6(a) shows an HH-polarization SAR image of a scene processed to reduced resolution (1 m by 1 m). In this image two regions of trees are separated by a narrow strip of coarse scrub. Also visible in the image, although somewhat faint, are four power-line towers located in the scrub. Figure 6(b) shows the corresponding PWF-processed image of the scene. The power-line towers have greater intensity in the PWF image than in the HH image because the PWF image includes contributions from HH, HV, and VV polarizations.

Table 3 compares ATR system performance using HH versus PWF data. The comparison was performed by using the same target and clutter data used to



(a)



(b)

FIGURE 6. Comparison of (a) HH and (b) PWF imagery. The power-line towers are more clearly visible in the PWF image because PWF processing combines data from all three polarization channels (HH, HV, and VV).

Table 3. Comparison of ATR System Performance Using HH Imagery versus PWF Imagery.

	FA/km^2	P_D HH Data	P_D PWF Data
After prescreening	30	0.65	0.82
After discrimination	3.0	0.57	0.75
After classification	0.3	0.24	0.70

generate the results in Table 1. At the output of each stage (prescreening, discrimination, and classification) the false-alarm densities were set equal for HH and PWF clutter imagery. This normalization permits us to compare the HH and PWF detection performance at each stage.

The detection performance was better with PWF data than with HH data. At the output of the detection stage $P_D = 0.82$ for PWF data and $P_D = 0.65$ for HH data. At the output of the discrimination stage $P_D = 0.75$ for PWF data and $P_D = 0.57$ for HH data. At the output of the classification stage $P_D = 0.70$ for PWF data and $P_D = 0.24$ for HH data; the P_D at the end of the classification stage represents the overall end-to-end performance of the ATR system.

Details of the Baseline Discrimination Features

This section presents details of the three baseline discrimination features: standard deviation, fractal dimension, and weighted-rank fill ratio. The equations for calculating each feature are also discussed. Because the concept of the fractal-dimension feature is fairly involved, this feature is discussed at greater length than the other two features.

Standard-Deviation Feature

The standard-deviation feature is a measure of the fluctuation in intensity, or radar cross section, in an image. The log standard deviation for a particular region is defined as the standard deviation of the radar returns (in dB) from the region. If the radar intensity in power from range r and azimuth a is denoted by $P(r, a)$, then the log standard deviation σ

can be estimated as follows:

$$\hat{\sigma} = \sqrt{\frac{S_2 - S_1^2 / N}{N - 1}},$$

where

$$S_1 = \sum_{r,a \in \text{Region}} 10 \log_{10} P(r, a) \quad (2)$$

and

$$S_2 = \sum_{r,a \in \text{Region}} [10 \log_{10} P(r, a)]^2 \quad (3)$$

and N is the number of points in the region.

Fractal-Dimension Feature

The fractal-dimension feature provides a measure of the spatial distribution of the brightest scatterers in a region. In the following paragraphs we present the formal definition of fractal dimension, along with several simple examples to illustrate the definition. We also show how to calculate the fractal dimension of detected objects in a SAR image. By using high-resolution SAR imagery gathered at Stockbridge, New York, we demonstrate how the spatial texture differences measured by the fractal-dimension feature can be used to discriminate between natural and cultural objects.

The fractal dimension of a set S in a two-dimensional space can be defined as follows:

$$\dim(S) = \lim_{\epsilon \rightarrow 0} \frac{\log M_\epsilon}{\log \left(\frac{1}{\epsilon} \right)}, \quad (4)$$

where M_ϵ = the minimum number of ϵ -by- ϵ boxes needed to cover S . (By covering S , we mean finding a set of square boxes B_i such that $\cup B_i \supseteq S$.) For small values of ϵ , the definition in Equation 4 is equivalent to writing

$$M_\epsilon \approx K \epsilon^{-\dim(S)}, \quad (5)$$

where K is a constant. This equation expresses one of the important ideas behind fractal analysis: fractal

dimension measures how certain properties of a set change with the scale of observation ε . In the following paragraphs, three specific examples clarify this idea.

Example 1. Let S be a single point. A point can be covered by one box regardless of the box size ε ; hence

$$\dim(\text{point}) = \lim_{\varepsilon \rightarrow 0} \frac{\log M_{\varepsilon}}{\log \left(\frac{1}{\varepsilon}\right)} = \lim_{\varepsilon \rightarrow 0} \frac{\log 1}{\log \left(\frac{1}{\varepsilon}\right)} = 0.$$

We could use Equation 5 to derive this same result by noting that the number of square boxes needed to cover S is independent of the box size; thus $\dim(\text{point})$ equals zero. Figure 7 summarizes this example. In addition, as long as ε is below a certain critical value, a finite set of isolated points can be covered by a fixed number of boxes (independent of ε). Therefore, a finite set of isolated points also has a dimension of zero.

Example 2. Let S be a line segment. For simplicity, we assume the line is 1 unit long. A single 1-unit-by-1-unit box can cover the line. If we reduce the box size to 1/2 unit by 1/2 unit, then two boxes are

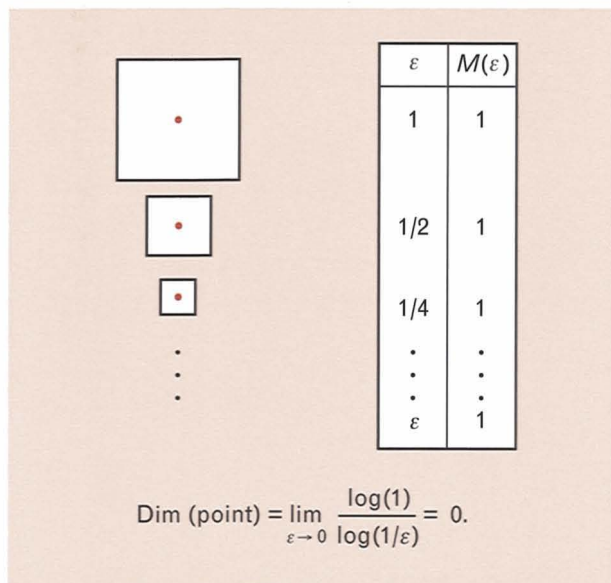


FIGURE 7. Fractal-dimension calculation for a point. As the size of the square box that covers the point decreases, the number of boxes required to cover the point remains the same (i.e., one). As a result, the fractal dimension of a point is zero.

needed to cover the line. If we reduce the box size again, to 1/4 unit by 1/4 unit, four boxes are needed to cover the line. Each time the box size is halved, the number of boxes needed to cover a line doubles; thus

$$\dim(\text{line}) = \lim_{\substack{\varepsilon \rightarrow 0 \\ n \rightarrow \infty}} \frac{\log 2^n}{\log \left(\frac{1}{2}\right)^n} = 1.$$

Figure 8 summarizes this example. It can also be shown that a finite set of isolated line segments has a fractal dimension of one.

Example 3. Let S be a square area. Again, for simplicity, we assume the square is 1 unit by 1 unit in size. A single 1-unit-by-1-unit box can cover the square. If we reduce the box size to 1/2 unit by 1/2 unit, four boxes are required. If we reduce the box size again, to 1/4 unit by 1/4 unit, 16 boxes are required. As the box size is halved, the number of boxes needed to cover the square area quadruples; thus

$$\dim(\text{square}) = \lim_{\substack{\varepsilon \rightarrow 0 \\ n \rightarrow \infty}} \frac{\log 2^{2n}}{\log \left(\frac{1}{2}\right)^n} = 2.$$

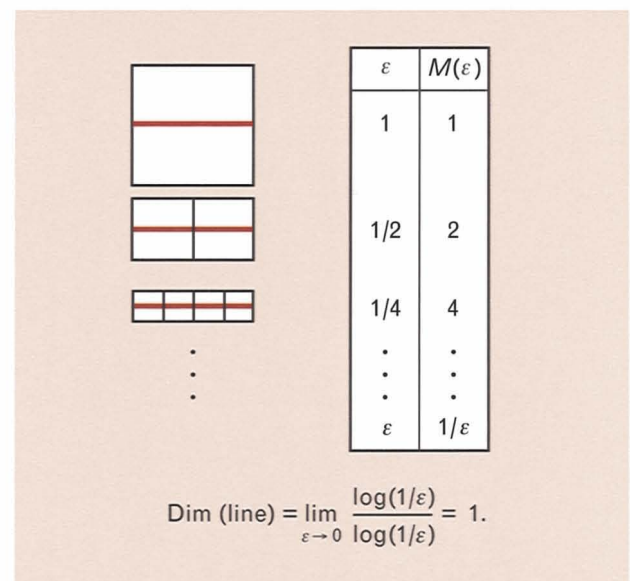


FIGURE 8. Fractal-dimension calculation for a line segment. As the size of each square box that covers the line segment is halved, the number of boxes required to cover the line segment doubles. As a result, the fractal dimension of a line segment is one.

Figure 9 summarizes this example. We used a square area in this example for convenience. Any area that can be synthesized from a finite number of square areas, however, will have a fractal dimension of two.

From these simple examples, we see that fractal dimension clearly has the potential to discriminate between certain types of objects in 2-D space. The question is, how can this feature be applied to SAR data?

The first step in applying the fractal-dimension concept to a radar image is to select an appropriately sized region of interest, and then convert the pixel values in the region of interest to binary (i.e., each pixel value equals 0 or 1). One method of performing this conversion is *amplitude thresholding*, in which all pixel values exceeding a specified threshold are converted to 1, and the remaining pixel values are converted to 0. Another method is to select the N brightest pixels in the region of interest and convert their values to 1, while converting the rest of the pixel values to 0; this second method is the approach we used (because it worked better).

After converting the radar image to a binary image,




FIGURE 9. Fractal-dimension calculation for a square area. Each time the size of each square box that covers the square area is halved, the number of boxes required to cover the square quadruples. As a result, the fractal dimension of a square is two.

we let the pixels with the value 1 constitute the set S in Equation 4. A problem arises, however, when we attempt to apply the definition of Equation 4 directly to our binary image. According to the definition of fractal dimension, we need to take a limit as the box size ϵ goes to zero. The smallest meaningful value of the box size ϵ , however, is the size of one pixel. Therefore, we must develop an approximation to the formula of Equation 4.

From Equation 5 we observe that

$$\log M_\epsilon \approx -\text{dim} \log \epsilon + \log K$$

for small ϵ . Because the relation between $\log M_\epsilon$ and $\log \epsilon$ is linear for small ϵ , with the slope equal to the negative of the dimension, the fractal dimension can be approximated by using only the box counts for $\epsilon = 1$ and $\epsilon = 2$ in the following way:

$$\text{dim} = -\frac{\log M_1 - \log M_2}{\log 1 - \log 2} = \frac{\log M_1 - \log M_2}{\log 2}, \quad (6)$$

where M_1 is the number of 1-pixel-by-1-pixel boxes needed to cover the image and M_2 is the number of 2-pixel-by-2-pixel boxes needed to cover the image. Figure 10 summarizes the fractal dimensions of simple objects as they are observed in SAR imagery.

The following paragraphs provide two examples of calculating the fractal dimension of regions of interest in radar imagery. The examples use data extracted from the SAR image shown in Figure 11. The figure shows a Stockbridge, New York, clutter scene that includes trees, a street with houses on both sides, a swimming pool, and a meadow. The examples demonstrate the fractal-dimension calculation for a typical tree (natural clutter) and the rooftop of a house (cultural clutter).

Figure 12 illustrates the fractal-dimension calculation for a binary image of a tree; the binary image was formed by selecting the 50 brightest pixels from a 120-ft-by-120-ft region of interest in the image of Figure 11. The number of 1-pixel-by-1-pixel boxes needed to cover this image is identical to the number of pixels with the value 1 (i.e., M_1 equals 50). The




FIGURE 10. Fractal dimensions of simple objects in SAR imagery. (a) Points have a fractal dimension of zero, (b) lines have a fractal dimension of one, and (c) squares and L-shaped objects have a fractal dimension of two. In the text we show how to calculate the fractal dimensions of these objects by using the approximation derived in Equation 6.

minimum number of 2-pixel-by-2-pixel boxes needed to cover the image is 41; Figure 12 shows this minimal covering. By applying Equation 6, we find that the fractal dimension of the tree is 0.29. This relatively low value reflects the fact that the binary image of the tree consists primarily of isolated pixels.

Figure 13 illustrates the fractal-dimension calculation for a binary image of a house rooftop (this image was formed in the same way as the image of the tree in Figure 12). Notice that the pixels in this image are clustered into lines and areas. The number of 1-pixel-by-1-pixel boxes needed to cover the image is 50, but the minimum number of 2-pixel-by-2-pixel boxes needed to cover the image is only 21. By using Equation 6, we find that the fractal dimension of the house rooftop is 1.25. This relatively high value is caused by the clustering of the pixels. The different fractal-dimension values for the tree and the rooftop illustrate that this feature can be used to discriminate between natural clutter and cultural clutter.

Weighted-Rank Fill Ratio Feature

The third textural feature, the weighted-rank fill ratio, measures the percentage of the total energy con-

tained in the brightest scatterers of a detected object. Using the notation of Equations 2 and 3, we define the weighted-rank fill ratio η as follows:

$$\eta = \frac{\sum_{k \text{ brightest pixels}} P(r, a)}{\sum_{\text{all pixels}} P(r, a)},$$

where k is selected to correspond approximately to the brightest 5% of the detected object's pixels. For man-made objects a significant portion of the total energy comes from a small number of bright scatterers; for natural clutter the total energy is distributed more evenly among the pixels.

Future ATR System Improvements

The baseline ATR system currently uses only three features in the discrimination stage (standard deviation, fractal dimension, and weighted-rank fill ratio); we have found that these features reliably reject natural-clutter false alarms. Other discrimination features could be added that would also reject some cultural-clutter false alarms. For example, a size feature, such as length and width of the detected object, could




FIGURE 11. High resolution (1 ft by 1 ft) PWF-processed SAR image of mixed natural and cultural clutter in Stockbridge, New York. A tree and a house rooftop from this image are used in the fractal-dimension calculations illustrated in Figures 12 and 13.

reject objects too large or too small to be targets. Numerous discrimination features have been investigated, and an improved set of five features has been found to provide better discrimination performance than the baseline feature set used in this article. A detailed description of all the discrimination features investigated so far is given in the accompanying article entitled "Discriminating Targets from Clutter" by Daniel E. Kreithen et al.

Improvements in the baseline pattern-matching classifier will be necessary before an operational, end-

to-end ATR system can be implemented. Target signatures vary with aspect angle, depression angle, and SAR squint angle; constructing templates incorporating these variables would produce a large and computationally unwieldy set of templates. In an operational scenario, the target aspect angle is not known. Therefore, each detected object would have to be correlated with a large subset of the pattern-matching templates, which would increase computation time significantly. The accompanying article entitled "Improving a Template-Based Classifier in a SAR Automatic Target Rec-




FIGURE 12. Fractal-dimension calculation for the binary image of a tree. The 50 brightest pixels (indicated by the small black boxes) are relatively isolated, and 41 two-pixel-by-two-pixel boxes are needed to cover them, which results in a low fractal dimension of 0.29.




FIGURE 13. Fractal-dimension calculation for the binary image of a house rooftop. The 50 brightest pixels (indicated by the small black boxes) are more tightly clustered than they are for the tree in Figure 12, and only 21 two-pixel-by-two-pixel boxes are needed to cover them, which results in a higher fractal dimension of 1.25.

ognition System by Using 3-D Target Information” by Shawn M. Verbout et al. presents an approach to this classification task based on the generation of 2-D templates from 3-D models of targets.

Because the pattern-matching approach to classification requires a large number of templates, we are investigating an alternative approach to classification—


using spatial matched filters. The initial results of these studies indicate the possibility of a significant reduction in computation time and storage requirements with no reduction in performance [7].

Acknowledgments

The authors wish to acknowledge the significant contributions made by Michael C. Burl, particularly to the development of the fractal-dimension feature. The authors also wish to thank Steven M. Auerbach for his help in technical editing and Robin Fedorchuk for her outstanding typing efforts. This work was sponsored by the Advanced Research Projects Agency.

REFERENCES

1. M.C. Burl, G.J. Owirka, and L.M. Novak, “Texture Discrimination in Synthetic Aperture Radar,” *23rd Asilomar Conf. on Signals, Systems, and Computers, Pacific Grove, CA, 30 Oct.–1 Nov. 1989*, p. 399.
2. J.C. Henry, “The Lincoln Laboratory 35 GHz Airborne Polarimetric SAR Imaging System,” *IEEE Natl. Telesystems Conf., Atlanta, GA, 26–27 Mar. 1991*, p. 353.
3. L.M. Novak, M.C. Burl, R. Chaney, and G.J. Owirka, “Optimal Processing of Polarimetric Synthetic Aperture Radar Imagery,” *Linc. Lab. J.* 3, 273 (1990).
4. L.M. Novak, M.C. Burl, and W.W. Irving, “Optimal Polarimetric Processing for Enhanced Target Detection,” *IEEE Trans. Aerosp. Electron. Syst.* 29, 234 (1993).
5. G.B. Goldstein, “False-Alarm Regulation in Log Normal and Weibull Clutter,” *IEEE Trans. Aerosp. Electron. Syst.* 8, 84 (1973).
6. W.W. Irving, G.J. Owirka, and L.M. Novak, “A New Model for High-Resolution Polarimetric SAR Clutter Data,” *SPIE* 1630, 208 (1992).
7. L.M. Novak, G.J. Owirka, and C.M. Netishen, “Radar Target Identification Using Spatial Matched Filters,” *3rd ATR Systems and Technology Conf.*, June 1993.



LESLIE M. NOVAK

is a senior staff member in the Surveillance Systems group. He received a B.S.E.E. degree from Fairleigh Dickinson University in 1961, an M.S.E.E. degree from the University of Southern California in 1963, and a Ph.D. degree in electrical engineering from the University of California, Los Angeles, in 1971. Since 1977 Les has been a member of the technical staff at Lincoln Laboratory, where he has studied the detection, discrimination, and classification of radar targets. He has contributed chapters on stochastic observer theory to the series *Advances in Control Theory*, edited by C.T. Leondes (Academic Press, New York), volumes 9 and 12.



GREGORY J. OWIRKA

is an assistant staff member in the Surveillance Systems group. He received a B.S. degree (cum laude) in applied mathematics from Southeastern Massachusetts University, and he is currently working on an M.S. degree in electrical engineering at Northeastern University. Greg's current research interests are in automatic target recognition. He has been at Lincoln Laboratory since 1987.



CHRISTINE M. NETISHEN

is an assistant staff member in the Surveillance Systems group. Her research speciality is in the detection, discrimination, and classification of stationary ground vehicles in SAR imagery. She received a B.S. degree in mathematics (cum laude) from Providence College; she also spent a year at Cambridge University in England studying mathematics and physics. Christine has been at Lincoln Laboratory since 1991.