Effects of Polarization and Resolution on the Performance of a SAR Automatic Target Recognition System

Leslie M. Novak, Shawn D. Halversen, Gregory J. Owirka, and Margarita Hiett

Lincoln Laboratory is investigating the use of high-resolution fully polarimetric synthetic-aperture radar (SAR) imagery to detect and classify stationary ground targets. This article summarizes a study in which data collected by the Lincoln Laboratory 33-GHz SAR were used to perform a comprehensive comparison of automatic target recognition (ATR) performance for several combinations of polarization and resolution. The Lincoln Laboratory baseline ATR algorithm suite was used, and it was optimized for each polarization/resolution case. The polarizations evaluated were HH and PWF; the resolutions evaluated were 1 ft × 1 ft and 1 m × 1 m. The data set used for this study contained approximately 74 km^2 of clutter (56 km^2 of mixed clutter and 18 km^2 of highly cultural clutter) and 136 tactical target images divided equally between tanks and howitzers.

In support of the Critical Mobile Targets (CMT) program sponsored by the Advanced Research Projects Agency (ARPA), Lincoln Laboratory has developed a complete end-to-end automatic target recognition (ATR) system that utilizes two-dimensional synthetic-aperture radar (SAR) imagery [1]. The ATR system is intended to address the surveillance and targeting aspects of the ARPA Warbreaker program; this program requires a sensor that can search large areas and also provide fine enough resolution to detect and identify the presence of missile launch systems such as a SCUD.

Currently there is much interest in examining SAR ATR performance trade-offs with regard to polarization and resolution. For example, what performance gain does sensor data at 1-ft × 1-ft resolution provide over data at 1-m × 1-m resolution? Similarly, what performance gain does fully polarimetric data provide over single-polarization data? Several factors contribute to the difficulty in answering such questions. SAR sensors operated by various organizations have different polarization and resolution capabilities. Data collections vary with respect to target types, target deployments, countermeasures, season, and geographical region. Organizations evaluating ATR performance typically test their own ATR algorithms by using a single data set collected by one SAR sensor. Because of these factors, direct and meaningful comparisons of ATR performance trade-offs with regard to polarization and resolution are difficult.

Lincoln Laboratory has a large database of clutter and target-in-clutter imagery, which can be used for comprehensive ATR performance comparisons for various combinations of polarization and resolution. The data were gathered by the Lincoln Laboratory Millimeter-Wave (33 GHz) SAR sensor [2], a high-resolution (1 ft × 1 ft) fully polarimetric Ka-band SAR with a 7-km standoff range. Because the SAR sensor data are fully polarimetric, the linear polarization components—HH (horizontal transmit, horizontal receive), HV (horizontal transmit, vertical receive), and VV (vertical transmit, vertical receive)—can be combined to form any desired polarization combination (e.g., LL, LR, RR, HH–VV, and PWF). Also, by processing less than the full bandwidth of the SAR data, we can produce imagery that emulates imagery from a lower-resolution sensor.

The Lincoln Laboratory ATR system has three main stages: detection (or prescreening), discrimination, and classification. In the detection stage, a two-parameter constant false-alarm rate (CFAR) detector is used as a prescreener to select candidate targets in a SAR image on the basis of local brightness. In the discrimination stage, a target-sized two-dimensional matched filter accurately locates the position of candidate targets and determines their orientation. Then discrimination features (including textural, size, contrast, and polarimetric features) are calculated and used to reject cultural and natural-clutter false alarms [3, 4]. In the classification stage, a two-dimensional pattern-matching algorithm provides additional false-alarm rejection and then classifies the remaining detections by target type (i.e., tank or howitzer).

This article provides a direct performance comparison of the combinations of polarization and resolution cases we tested. Also, because large amounts of clutter and target data were used, the results are statistically significant down to low false-alarm rates.

The next section gives an overview of the Lincoln Laboratory SAR sensor and the data processing techniques used in this study. The following section gives an overview of the ATR system used in this study, and describes the three stages (detection, discrimination, and classification) in detail. A subsequent section summarizes the performance studies. First, we compare performance using 1-ft × 1-ft resolution data and PWF and HH polarizations. Then we compare performance using 1-m × 1-m resolution data and PWF and HH polarizations. We also give the performance results for a dual-resolution scheme that uses 1-m × 1-m PWF data in the prescreener stage and 1-ft × 1-ft data in the discrimination and classification stages. The appendix discusses the features used in the discrimination stage.

The SAR Sensor and Data Processing Techniques

To evaluate the performance of the ATR system, we used high-resolution (1 ft × 1 ft) fully polarimetric target data and clutter data gathered by the Lincoln Laboratory SAR sensor at a 22.5° depression angle

(a) (b)

FIGURE 1. The Lincoln Laboratory synthetic-aperture radar (SAR) sensor. (a) The sensor platform is a Gulfstream G1 aircraft, shown here in flight. (b) The radome, located at the bottom of the aircraft, contains the sensor antenna.

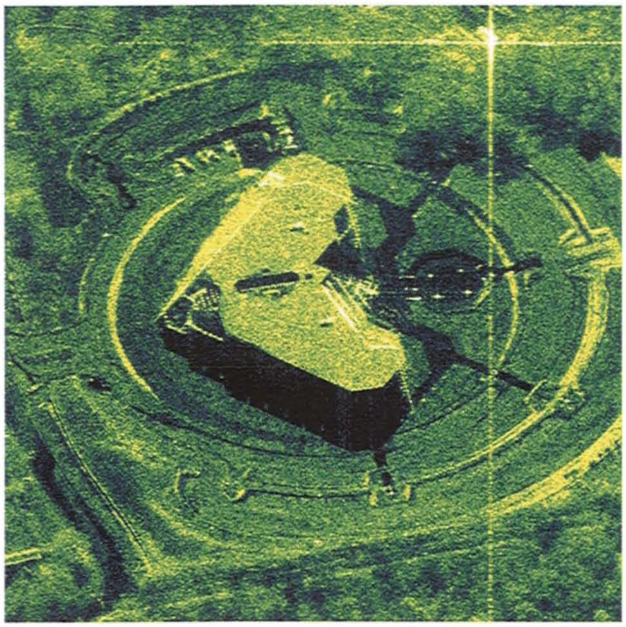

FIGURE 2. A 1-ft × 1-ft resolution PWF SAR image of an office building in Lincoln, Massachusetts. The high resolution and polarimetric processing allow for near-optical image quality. Cars are not seen because the parking lot was empty when the SAR data were obtained.

FIGURE 3. An aerial photograph of the office building in Lincoln, Massachusetts, shown in Figure 2. Comparison with the 1-ft × 1-ft resolution PWF SAR image in Figure 2 demonstrates the high quality of data gathered with the Lincoln Laboratory SAR sensor.

and a 7-km slant range. Figure 1 shows the sensor, which collects single polarization (HH, HV, and VV) data that can be combined linearly to form any desired transmit and receive polarization combination. In addition, these three channels can be processed by using an optimal polarimetric processing technique known as the polarimetric whitening filter (PWF) [5, 6]. The performance of the ATR system was evaluated by using single-polarization-channel radar imagery (HH) and PWF imagery. In addition, these data were reprocessed to a lower resolution in order to emulate a 1-m × 1-m sensor.

Figure 2 shows an example of the imagery gathered by the Lincoln Laboratory SAR sensor; this figure is a 1-ft × 1-ft resolution PWF image of an office building located in Lincoln, Massachusetts. The high-resolution data and polarimetric processing allow for near-optical image quality, which is evident when we compare the SAR image with the corresponding aerial photograph of the office building shown in Figure 3. The SAR image was obtained under clear weather conditions (and during off hours, which explains the lack of vehicles in the parking lot). However, the quality and resolution of the SAR image would not have been degraded 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 SAR sensor can be used during the day or at night. In addition, SAR sensors can image large areas from a long distance.

The SAR image in Figure 2 was formed by using PWF processing, which optimally combines the HH, HV, and VV polarization components. PWF processing 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 such as the building are more sharply defined. As a result, PWF-processed imagery is visually much clearer than single-polarization-channel imagery. For comparison, Figure 4 shows a single-polarization-channel image (HH) of the same scene. The speckle inherent in single-polarization SAR imagery obscures much of the detail visible in the PWF SAR image in Figure 2. In addition, PWF-processed imagery improves the performance of all three stages of the ATR system

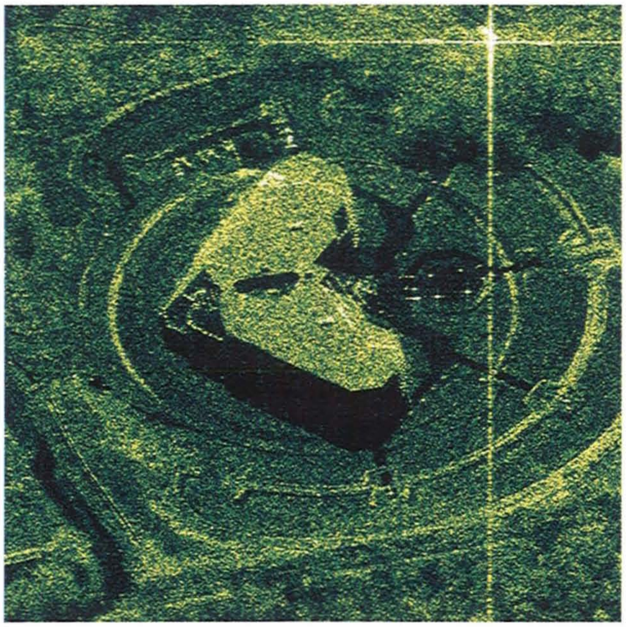

FIGURE 4. A 1-ft × 1-ft resolution HH SAR image of the office building in Lincoln, Massachusetts. Although many features are visible in this image, the speckle inherent in single-polarization SAR imagery obscures much of the detail seen in the 1-ft × 1-ft PWF SAR image.

compared to the performance achieved by using single-polarization-channel imagery. The improvement occurs because PWF processing reduces the clutter variance and thus enhances target signatures relative to the clutter background.

In addition, all the target data and clutter data were reprocessed to 1-m × 1-m resolution in order to investigate end-to-end ATR system performance for a reduced-resolution sensor. Figures 5 and 6 show examples of 1-m × 1-m PWF-processed SAR data and single-polarization HH SAR data, respectively. The 1-m × 1-m HH image in Figure 6 is clearly dominated by speckle. PWF processing reduces the speckle to an acceptable level, producing an image of good visual quality as shown in Figure 5. However, the PWF-processed image in Figure 5 does not approach the near-optical image quality of the 1-ft × 1-ft PWF image shown in Figure 2.

The Baseline ATR System

This section describes the three stages—detection, discrimination, and classification—of our baseline

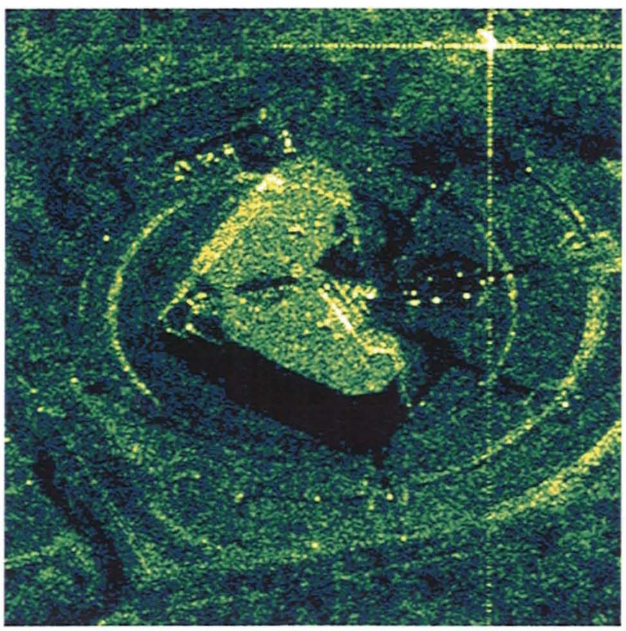

FIGURE 5. A 1-m × 1-m resolution PWF SAR image of the office building in Lincoln, Massachusetts. This image has been oversampled to match the pixel spacing of the 1-ft × 1-ft SAR images in Figures 2 and 4. Some image detail has been lost at the reduced resolution, but the speckle reduction provided by PWF processing produces a clear image.

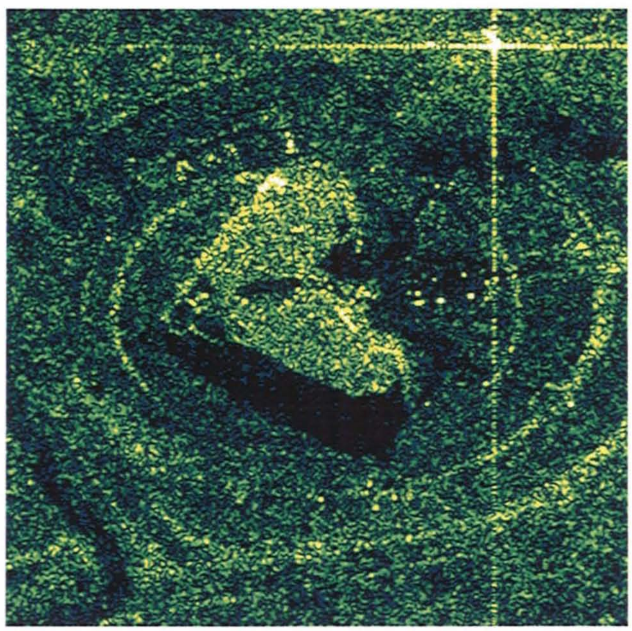

FIGURE 6. A 1-m × 1-m resolution HH SAR image of the office building in Lincoln, Massachusetts. This image has been oversampled to match the pixel spacing of the 1-ft × 1-ft SAR images in Figures 2 and 4. Even though this single-polarization image has the same resolution as Figure 5, it is evident that this image is dominated by speckle, which results in a lower-quality image.

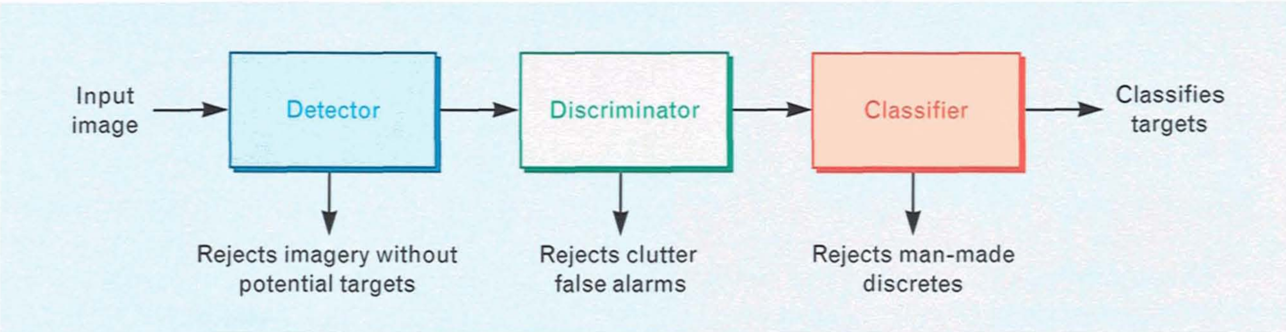

FIGURE 7. A block diagram of the three-stage SAR ATR system. The detector, or prescreener, selects candidate targets in the image; the discriminator rejects natural-clutter false alarms while accepting real targets; and the classifier rejects man-made clutter and classifies the remaining target detections by vehicle type.

SAR ATR system. Figure 7 shows a block diagram of the system. The three stages are described below.

Stage 1: Detection and Prescreening

In the first stage of processing, a two-parameter CFAR detector [7] is used as a prescreener; this stage of processing selects candidate targets in the image on the basis of local brightness (i.e., by searching for bright returns in units of dB). Figure 8 is a sketch of the two-parameter CFAR detector used by the prescreener. The detector is defined by the rule

$$\frac{X_t - \hat{\mu}_c}{\hat{\sigma}_c} > K_{\text{CFAR}}, \qquad (1)$$

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 that controls the false-alarm rate. If the detection statistic calculated in Equation 1 exceeds the constant K_{CFAR}, the test cell is declared to be a target; if not, the test cell is declared to be clutter. Since the SAR image is in dB units, the detector rule defined in Equation 1 is equivalent to the expression

$$\frac{T/C}{\hat{\sigma}_c} > K_{\text{CFAR}},$$

where $T/C = X_t - \hat{\mu}_c$ is the peak target-to-average clutter ratio (in dB); the quantity $\hat{\sigma}_c$ is referred to as the log standard deviation of the clutter.

As illustrated in Figure 8, the test cell is in the center of a boundary stencil, and the cells in the boundary stencil are used to estimate the local mean and log standard deviation of the clutter. A guard area exists so that the target will not affect the clutter estimates. The stencil size varies with resolution and the type of target of interest; for tactical targets the stencil contains 160 cells for 1-m × 1-m data and 640 cells for 1-ft × 1-ft data. If the amplitude distribution of the clutter were Gaussian, then the CFAR detector would provide a constant false-alarm rate for any given value of K_{CFAR} [7]. Because the clutter distributions of

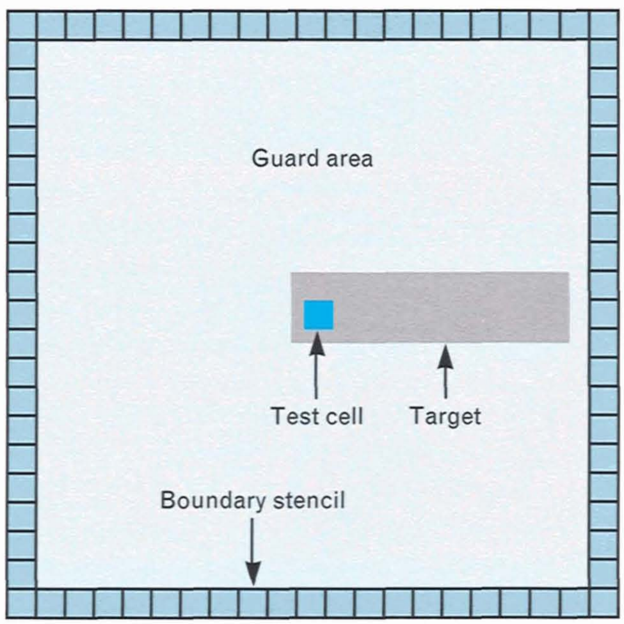

FIGURE 8. Diagram of the CFAR detection algorithm. First-order and second-order clutter statistics are estimated from the data in the boundary stencil and compared with the test cell. The guard area prevents target scatters from affecting the clutter estimates.

high-resolution data are only approximately Gaussian [8], the CFAR detector does not always yield a constant false-alarm rate; however, the CFAR detector is still a reasonable algorithm for detecting targets in clutter.

Because a single target can produce multiple CFAR detections, the detections in target-sized regions are clustered together. Then a 128-ft × 128-ft region of interest (ROI) around each cluster center is extracted and passed to the discrimination stage of the ATR system for further processing.

Stage 2: Discrimination

The discriminator analyzes each 128-ft × 128-ft ROI that it receives from the detector. The goal of discrimination processing is to reject clutter false alarms while accepting real targets, thereby reducing the computational load in the classification stage. The discrimination stage consists of three steps: (1) determining the position and orientation of a detected object, (2) computing textural, size, contrast, and polarimetric features, and (3) combining the features into a discrimination statistic that measures how target-like the detected object is.

To determine the position and orientation of the object, the discrimination algorithm places a target-sized rectangular template on the image; the template is slid in range and cross-range and rotated until the energy within the template is maximized. This object position estimate is more accurate than the estimate produced in the prescreening stage. The estimation operation is also computationally reasonable because it is performed only on the regions of interest passed by the prescreener. Because the target is assumed to be brighter than the surrounding clutter, the operation is similar to processing an image with a two-dimensional matched filter when the orientation of the target is unknown.

The second step in the discrimination stage calculates twelve discrimination features (nine for single-polarization data). These features, which are listed in Table 1, are described in detail elsewhere [3, 4]. The appendix gives a brief description of the features used in this stage.

The third and final step in the discrimination stage combines a subset of the twelve discrimination features into a single discrimination statistic. This subset varies with polarization and resolution. The choice of features used to determine the discrimination statistic is discussed in more detail in the section entitled "Performance of the ATR System." The discrimination statistic is calculated as a quadratic distance metric,

$$d(\mathbf{X}) = \frac{1}{n}(\mathbf{X} - \mathbf{M})^T \hat{\Sigma}^{-1}(\mathbf{X} - \mathbf{M}), \qquad (2)$$

where n is the number of features in the subset, $\mathbf{M}$ and $\hat{\Sigma}$ are estimates of the mean vector and covariance matrix of the features in the subset, and $\mathbf{X}$ is the vector of corresponding features measured from the detected object [9]. The quantities $\mathbf{M}$ and $\hat{\Sigma}$ are estimated by using a set of training feature vectors calculated from an independent set of non-netted target-training images (targets that are not concealed by radar-scattering camouflage nets).

The distance metric $d(\mathbf{X})$ represents a statistical measure of the distance of the detected object from the target class [10]. A theoretical analysis of the performance of this discriminator has been performed [4, 9], under the assumption that the features measured from the detected object are jointly Gaussian. Many of the features listed in Table 1 clearly do not satisfy this assumption. Thus, in this study, the

Table 1. List of Discrimination Features

Textural features	Standard deviation Fractal dimension Weighted fill
Size features	Mass Diameter Normalized rotational inertia
Contrast features	Peak CFAR Mean CFAR Percent bright CFAR
Polarimetric features (fully polarimetric data only)	Percent pure (odd or even) Percent pure even Percent bright even

thresholds were set empirically, as explained in the following section.

For an actual target, the quadratic distance metric calculated by using Equation 2 is small, and the corresponding ROI will be passed from the discrimination stage to the classification stage for further processing. Most clutter false alarms produce a large value for the distance metric $d(\mathbf{X})$, and will be rejected in the discrimination stage. Some man-made clutter discretes that are target-like (such as trucks and other vehicles) will pass the discrimination stage; therefore, the classification stage of the ATR system must have the ability to reject them.

Stage 3: Classification

A mean-squared-error (MSE) pattern-matching classifier is used to reject cultural false alarms caused by man-made clutter discretes such as trucks and small buildings, and to classify the remaining target detections by vehicle type. In these studies we implemented three-class classifiers (tank, howitzer, and clutter) for both 1-ft × 1-ft imagery and 1-m × 1-m imagery. The pattern-matching references used in these classifiers were constructed by averaging five consecutive spotlight-mode images of a target collected at a 22.5° depression angle and at 1° increments of azimuth, yielding seventy-two smoothed references for each of the targets at the given depression angle. For an operational system to be able to accommodate a range of depression angles, an expanded set of pattern-matching references would have to be constructed at various depression angles.

To construct the reference templates for the MSE classifier, we choose thresholds for the pixel values (converted to dB) of non-netted target-training images to eliminate all but the brightest and dimmest 1.5% of the data. This procedure reduces the original target image to a collection of target shadow pixels and bright target pixels. A binary image of these pixels is then generated and morphologically processed (by using a series of dilations and erosions with various binary kernels) to create a binary mask; all of the pixels in the original (dB) image that are not under the mask are used to calculate an average dB clutter level $\hat{\mu}_c$. The target-training image U_i is then normalized by $R_i = U_i - \hat{\mu}_c$.

With this normalization scheme, the average clutter level has been removed, but the energy of the reference target relative to the clutter background is retained; therefore, howitzer templates (for example) will have more total power than tank templates, making howitzers and tanks more separable. Because of the normalization, this approach is not affected by errors in absolute radar-cross-section calibration of the sensor data. After normalization, the target images are windowed with a rectangular mask that is slightly larger than the size of the target and oriented at the same angle as the target. This windowing eliminates the influence of nearby clutter in the classifier templates.

When an ROI is passed to the classifier, the mean clutter level is calculated and the ROI is normalized, as described above. The MSE is calculated for all of the classifier templates in the reference library, and the template that minimizes the error is chosen as the best match. For a given reference template, the error ε is calculated for the test image under the reference template window by

$$\varepsilon = \frac{\sum_{i=1}^{N} (R_i - T_i)^2}{N},$$

where N is the number of pixels in the reference template window (which can vary for different reference target sizes and different aspect angles), and R_i and T_i denote pixel values in the reference template and test image, respectively. The reference template is slid a small amount in range and cross-range to compensate for any small target centroid errors, and the minimum error over these translations is taken as the best match for that template. Because of the normalization of the reference images and test images, the error ε is a measure of the relative difference in total power between the two, which aids in the rejection of clutter discretes.

Each detected object that passes the discrimination stage is compared to each of the stored references. If none of the matches is below some maximum allowable score, the detected object is classified as clutter; otherwise, the detected object is assigned to the class (tank or howitzer) with the best match score. For an

MSE classifier, the lower the score the better the match; i.e., a perfect match has an MSE score of 0.0.

Performance of the ATR System

This section summarizes the performance of the detection, discrimination, and classification stages of the SAR ATR system. A total of 56 km^2 of high-resolution (1 ft × 1 ft) clutter data were processed by the ATR algorithms, along with 136 netted target images. The targets were M48 tanks and M55 howitzers, which have similar dimensions. These target data and clutter data were collected at Stockbridge, New York, and are referred to in this article as the Stockbridge data set. The clutter data contained a significant number of man-made discretes such as buildings, roads, cars, and electrical lines. The target-training data were taken from a set of non-netted target images. The target test data were collected under realistic deployment conditions in which the targets (with the same orientations and articulations) were covered with radar camouflage netting.

We show performance curves at all three ATR stages (detection, discrimination, and classification) for each combination of polarization and resolution that were studied. First the detection, or prescreening, performance curve was generated. Then an operating point on that curve was selected and the detected objects (targets and false alarms) corresponding to that operating point were passed to the discrimination stage to obtain the discrimination performance curve. Finally, an operating point on the discrimination curve was selected and the detected objects corresponding to that operating point were passed to the classification stage to obtain the classification performance curve. To compare results in a consistent manner among combinations of polarization and resolution, we set the operating point for the prescreener at probability of detection $P_D = 1.0$, for the discriminator at $P_D = 1.0$, and for the classifier at $P_D = 0.9$. These fixed probability-of-detection levels allow us to compare ATR performance for the combinations of polarization and resolution by examining only the false-alarm rates at the output of each stage.

In the discrimination stage, the location and orientation of each detected object was determined. The discrimination features (textural, size, contrast, and polarimetric features) were calculated and the best subset of features (for each polarization/resolution case) was used by the discriminator to reject both natural-clutter false alarms and man-made clutter discretes. For a given polarization/resolution case, all possible combinations of three or more features were tested against the test targets and the clutter false alarms passed to the discriminator by the detection stage. The best combination of features was used in the discriminator for that case. Thus, the comparisons presented in this article are for the best possible discrimination performance for each case.

To test the robustness of these discrimination feature sets against a different data set, we ran the ATR algorithm suite (using the same discrimination feature sets and threshold settings) on an additional 18 km^2 of highly cultural clutter collected near Ayer, Massachusetts. The results obtained were comparable to the results with the original Stockbridge data set (see the next section).

In the classification stage, the two-dimensional MSE pattern-matching classifier was applied to those detected objects which had passed the discrimination stage. In the classification stage, the MSE threshold was set to produce a probability of detection $P_D = 0.9$. Therefore, of the 136 netted test targets input to the MSE classifier, a total of 122 test targets were detected (i.e., classified as either a tank or howitzer). The two-dimensional MSE pattern-matching classifier was also applied to the clutter false alarms that passed the discrimination stage; these clutter false alarms were either rejected by the classifier (i.e., declared "clutter") or were classified into tank or howitzer categories. Classification confusion matrices were tabulated that summarize the ability of the MSE classifier to reject clutter false alarms and assign the remaining detected objects into tank and howitzer classes.

Performance Using 1-ft × 1-ft Resolution Data

Figure 9 presents the performance curves of all three stages of the ATR system for the PWF-processed data. Note that the discrimination stage reduces the number of false alarms by approximately a factor of five, compared to the detection stage, and the MSE classifier shows a further reduction in false alarms by more than an order of magnitude. The false alarms at each stage were categorized as either man-made discretes or natural-clutter false alarms. Table 2 gives the breakdown for each stage with respect to the type of false alarm.

Also included in Table 2 is the best subset of features used in the discrimination stage. For PWF and HH polarizations, the best subset of features was found by examining all possible combinations of three or more features and determining which subset produced the best discrimination performance using the netted target data set and the 56 km^2 of Stockbridge clutter false alarms. As the data in Table 2 show, the set of 3457 clutter false alarms from the detection stage is nearly equally divided between natural-clutter false alarms (bright tree lines, berms, isolated trees) and man-made discretes (buildings, nontarget vehicles, roads, electrical lines). The best subset of discrimination features was able to reject—nearly equally—both the man-made discretes and the natural-clutter false alarms. The size features (diameter and normalized rotational inertia) are effective at rejecting objects much larger or much smaller than the targets of interest. The textural features (fractal dimension and standard deviation) are effective at rejecting natural-clutter false alarms. The polarimetric feature (percent pure even) takes advantage of the even-bounce nature of targets such as tanks and howitzers, and the contrast feature (mean CFAR) is effective in rejecting false alarms that do not have a sufficient mean local brightness.

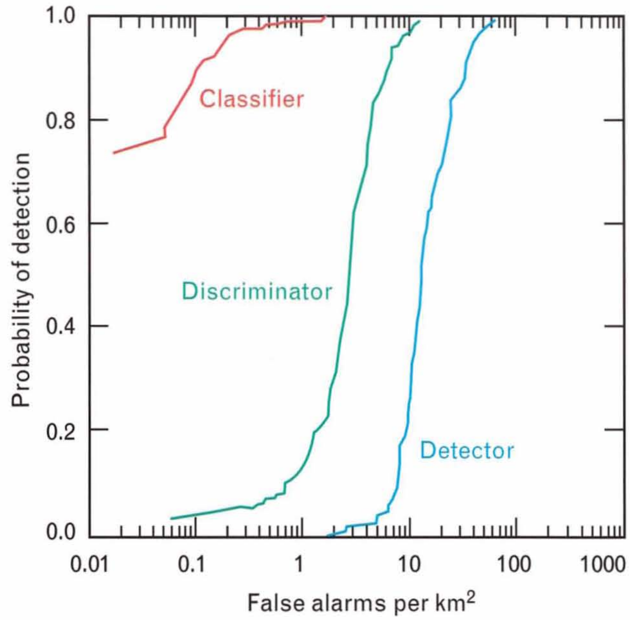

FIGURE 9. Performance curves for the three stages of the ATR system using 1-ft × 1-ft resolution PWF data. All algorithm training was performed by using non-netted tactical target imagery; testing was then performed on netted targets. To produce false-alarm statistics, 56 km^2 of clutter data collected in Stockbridge, New York, were used.

The results presented in Table 2 also indicate the clutter rejection capabilities of the MSE classifier. In the classification stage, with the probability of detection $P_D = 0.9$, nearly 99% of the clutter false alarms presented to the classifier were rejected as being nontargets. Finally, the false-alarm density achieved at $P_D = 0.9$ is calculated from the data in the table to be

Table 2. Clutter False-Alarm Statistics and Categories for PWF Data at 1-ft x 1-ft Resolution*

Stage	*Probability of detection*	*Total number of false alarms*	*Number of man-made-discrete false alarms*	*Number of natural-clutter false alarms*
Detection	1.0	3457	1760	1697
Discrimination**	1.0	709	444	265
Classification	0.9	6	4	2

* Target data: 136 tactical targets (netted)
Clutter data: 56 km^2 Stockbridge clutter

** Best subset of discrimination features:
1. Fractal dimension
2. Standard deviation
3. Mean CFAR
4. Percent pure even
5. Diameter
6. Normalized rotational inertia

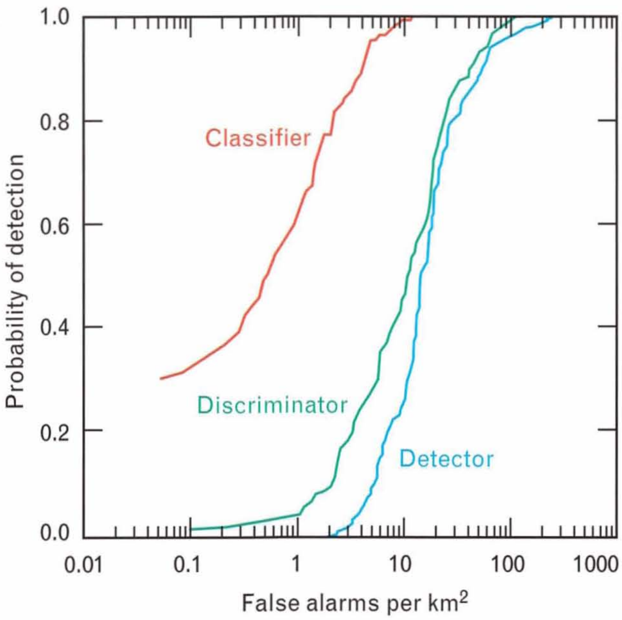

FIGURE 10. Performance curves for the three stages of the ATR system using 1-ft × 1-ft resolution HH data. The discrimination stage does not provide as large a false-alarm reduction as when using 1-ft × 1-ft resolution PWF data.

approximately 0.1 false alarms per km^2 (6 false alarms in 56 km^2 total clutter area).

Figure 10 presents the performance of all three stages of the ATR system for the single-polarization 1-ft × 1-ft resolution HH data. Note that the discrimination stage does not significantly reduce the false-alarm density. This result occurs partially because of the large amount of speckle contained in single-polarization data, which prevents accurate estimates of the discrimination features. Additionally, the use of polarimetric discrimination features such as percent pure even are not possible with single-polarization data. The false alarms at each stage are categorized as either man-made discretes or natural-clutter false alarms. Table 3 gives the breakdown of each stage into types of false alarms. As the data in Table 3 show, the number of false alarms coming out of the prescreener was significantly greater than that achieved by using PWF data. Most notably, the number of natural-clutter false alarms increased by nearly a factor of four (6473 natural-clutter false alarms) resulting in a total of 8739 false alarms for HH data. By using the best subset of discrimination features for HH data the discriminator rejected slightly more than half of the false alarms. Nearly all the false alarms rejected were from the natural-clutter category; only a small fraction of the false alarms from man-made discretes were rejected.

Table 3 also indicates that the MSE classifier is able to reject most of the clutter discretes even though a single-polarization channel is being used. At the final detection probability of $P_D = 0.9$, nearly 94% of the false alarms presented to the classifier were rejected as being non-targets. However, the performance is significantly worse than that achieved by using PWF imagery. The false-alarm density achieved at $P_D = 0.9$ is calculated from the data in the table to be approximately four false alarms per km^2 (229 false alarms in 56 km^2 total clutter area).

The second function of the classifier is to assign objects accepted as targets to target classes (i.e., tank or howitzer). Table 4 shows the classification perfor-

Table 3. Clutter False-Alarm Statistics and Categories for HH Data at 1-ft x 1-ft Resolution*

Stage	*Probability of detection*	*Total number of false alarms*	*Number of man-made-discrete false alarms*	*Number of natural-clutter false alarms*
Detection	1.0	8739	2266	6473
Discrimination**	1.0	3758	1928	1830
Classification	0.9	229	83	146

* Target data: 136 tactical targets (netted)
Clutter data: 56 km^2 Stockbridge clutter

** Best subset of discrimination features:
1. Standard deviation
2. Weighted fill
3. Mean CFAR
4. Mass

Table 4. Performance of MSE Classifier (Detection Probability = 0.9)*

PWF data 1-ft × 1-ft resolution		*Classified as tank*	*Classified as howitzer*	*Classified as clutter*
	Tank	62	0	6
	Howitzer	0	60	8
	Clutter	6	0	703
HH data 1-ft × 1-ft resolution		*Classified as tank*	*Classified as howitzer*	*Classified as clutter*
	Tank	60	0	8
	Howitzer	0	62	6
	Clutter	210	19	3529

* Training data: non-netted targets; test data: 136 netted targets

mance of the MSE classifier for both PWF and HH data, in the form of confusion matrices that tabulate the correct and incorrect classifications. Recall that the MSE classifier used templates constructed from non-netted targets; the classification results shown in Table 4 are for clutter discretes and netted test targets that passed the detection and discrimination stages. At the output of the classification stage, 90% of the netted targets were classified as targets, and 100% of these targets were correctly classified by vehicle type.

For comparison purposes, Figure 11 presents the end-to-end performance of the ATR system for both HH and PWF 1-ft × 1-ft resolution data. Note that the false-alarm density with PWF data is over an order of magnitude lower than with HH data.

To investigate whether the discrimination feature sets and classifier templates developed by using the Stockbridge data were scene dependent, we processed another 18 km^2 of clutter data through the ATR system with the same feature sets, templates, and thresholds used for the Stockbridge data set. This new data set was collected in Ayer, Massachusetts, and consists almost exclusively of man-made objects, which yield target-like false alarms because of the sharp edges of buildings, vehicles, railways, and roads, as shown in Figure 12. This data set is considered to be very difficult because of the large amount of man-made clutter. When the Ayer PWF data were processed through the ATR system, the prescreener produced 4133 false

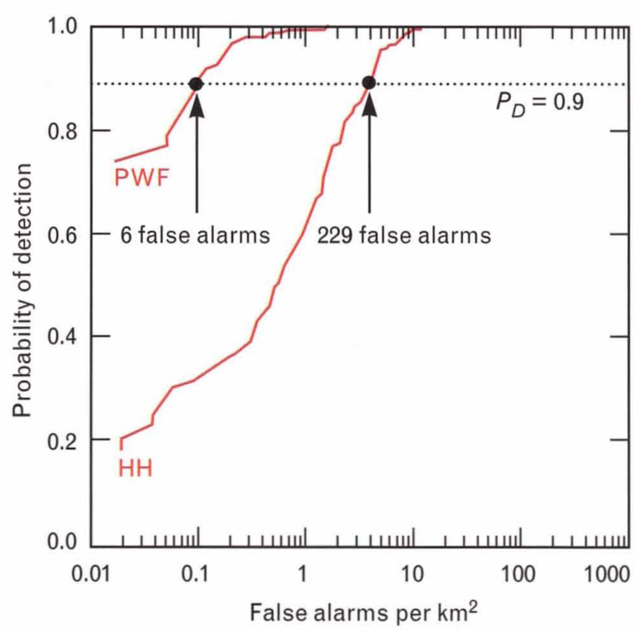

FIGURE 11. Comparison of end-to-end performance of the ATR system using 1-ft × 1-ft resolution PWF data and 1-ft × 1-ft resolution HH data from the Stockbridge data set. Overall, end-to-end performance of the ATR system is degraded when using HH data compared to using PWF data.

(a)

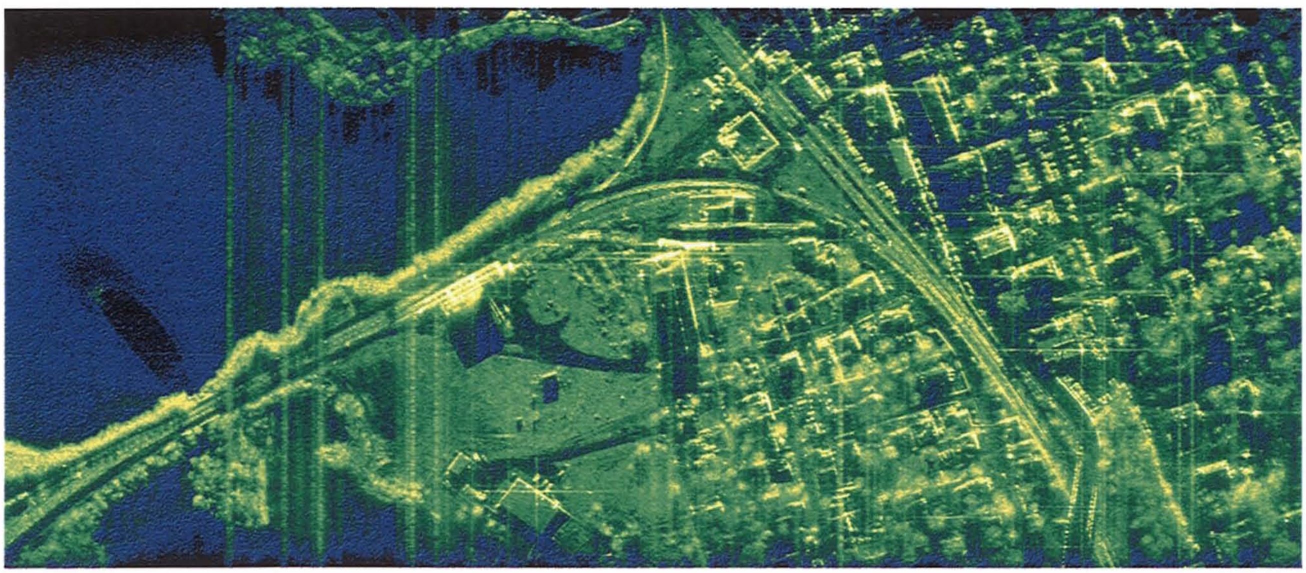

(b)

FIGURE 12. (a) Optical image and (b) SAR image of man-made cultural clutter in Ayer, Massachusetts. The houses and roads are clearly seen in the SAR image. The frozen pond appears as a black patch of very low radar returns because of the forward scattering of the radar signal. This image, which consists almost exclusively of man-made objects, is representative of the additional 18 km^2 of cultural-clutter data used for the verification of discrimination feature selection.

alarms, the discriminator rejected 3114 of these as clutter false alarms, and the classifier (P_D = 0.9) rejected 1016 of the remaining ones, resulting in only three false alarms at the output of the classification stage of the ATR system from the 18 km^2 clutter area. Thus, although the prescreener detects many target-like man-made objects, the discriminator is fairly robust, and it rejects a large percentage of the clutter false alarms that pass the prescreener. Finally, the robustness of the classifier allows the overall end-to-end performance of the ATR system to remain relatively stable.

By combining the data collected in Ayer with the Stockbridge data set, we expanded the amount of data processed to 74 km^2. The addition of this cultural-clutter data set has the effect of degrading the performance of the prescreener and discriminator somewhat, but degrading overall end-to-end performance only slightly. Figure 13 shows the end-to-end performance of the ATR system for HH data and PWF data. Comparing Figure 13 with Figure 11 clearly shows that end-to-end performance is degraded only slightly for both PWF and HH data by the addition of the Ayer cultural-clutter data set.

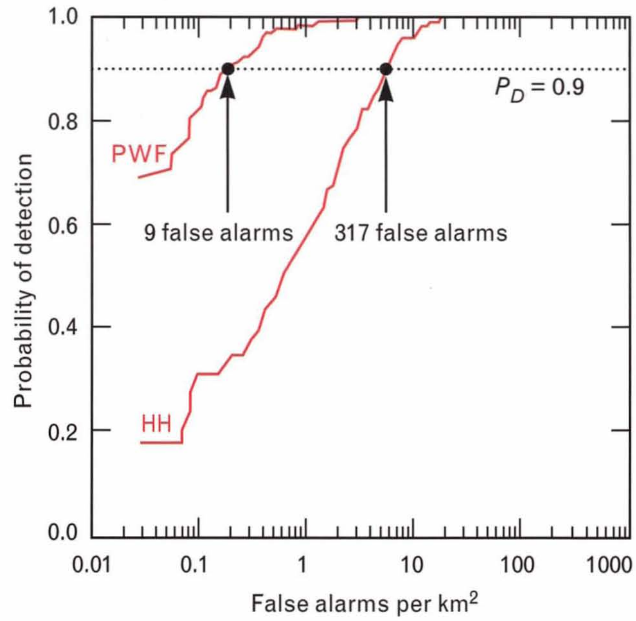

FIGURE 13. Comparison of end-to-end performance of the ATR system using 1-ft × 1-ft resolution PWF data and 1-ft × 1-ft HH data. In this case, the Stockbridge data set was combined with cultural-clutter data from Ayer, Massachusetts, to enlarge the overall test data set to 74 km^2. The addition of the cultural-clutter data set somewhat degrades the performance of the prescreener and discriminator, but end-to-end performance of the ATR system is degraded only slightly, as seen by comparing this figure with Figure 11.

Performance Using 1-m × 1-m Resolution Data

This section summarizes the performance of the ATR system using reduced-resolution SAR data. In order to emulate a reduced-resolution sensor, we reprocessed all the target data and clutter data used in the 1-ft × 1-ft resolution studies to 1-m × 1-m resolution. In addition, we examined the possibility of using lower-resolution data only for the initial detection stage; i.e., 1-m × 1-m resolution data were used for detection followed by 1-ft × 1-ft data for discrimination and classification. This combination was performed only with PWF data. Each of the three stages of the ATR system was appropriately modified to work with 1-m × 1-m resolution data. The following modifications to the algorithms were performed:

Detector. The CFAR detector used 1-m × 1-m resolution data, which provided a significant computational savings compared to the use of 1-ft × 1-ft resolution data. This computational savings occurred primarily because the amount of SAR data processed by the detector was reduced by a factor of sixteen. Additional computational savings occurred because the size of the boundary stencil was reduced by a factor of four as a result of the reduction in resolution of the SAR data.

Discriminator. The discrimination stage that processed the high-resolution 1-ft × 1-ft ROIs was used, without modification, to process 1-m × 1-m ROIs. This was possible because the reduced-resolution data were oversampled to match the 1-ft × 1-ft resolution ROIs. The oversampling was not performed in order to affect performance, but to make a reduced-resolution image have the same number of pixels as the corresponding full-resolution image. Thus the same algorithms for determining the position and orientation of the detected object and combining features into the discrimination statistic were applied. The calculation of the textural, size, contrast, and polarimetric features was also performed by using the oversampled data, and the selected set of features applied to the reduced-resolution data was optimized to maximize the false-alarm rejection performance of the discriminator.

Classifier. The MSE classifier used to process the high-resolution 1-ft × 1-ft ROIs was also used, without modification, to process the 1-m × 1-m ROIs. Again, this was done because the reduced-resolution data were oversampled to match the pixel spacing of the 1-ft × 1-ft resolution data. The MSE classifier used new pattern-matching references; these were constructed by averaging five consecutive spotlight-mode images (1-m × 1-m resolution) of a target, collected at 1° increments of azimuth, yielding seventy-two references for each of the targets.

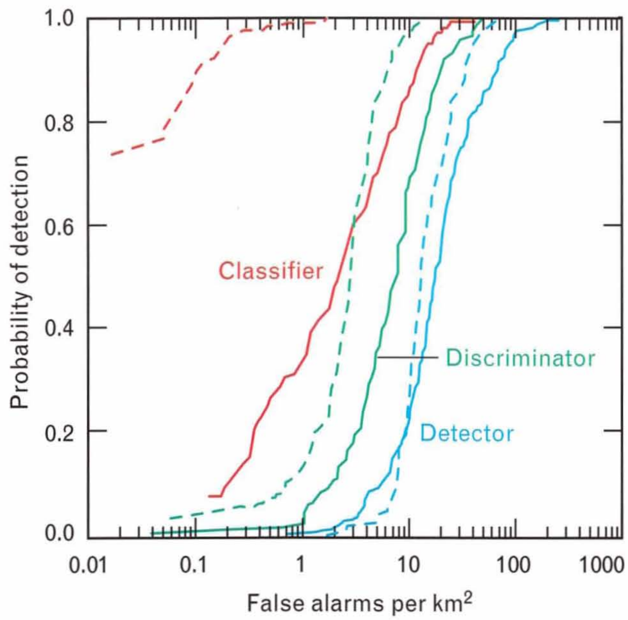

FIGURE 14. Comparison of performance curves for the ATR system using 1-m × 1-m resolution PWF data (solid lines) versus 1-ft × 1-ft resolution PWF data (dashed lines).

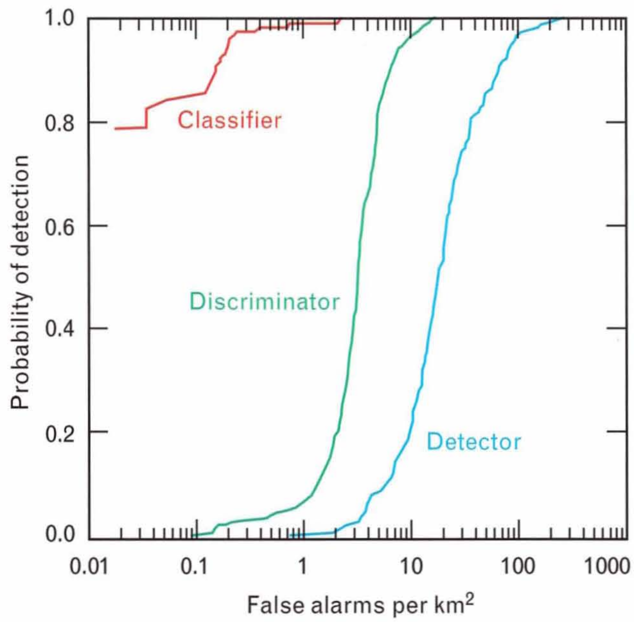

FIGURE 15. Performance curves for the three stages of the ATR system using 1-m × 1-m resolution PWF data for the detection stage, followed by 1-ft × 1-ft resolution data for the discrimination and classification stages. The overall system performance (i.e., the output of the classification stage) is almost identical to the performance of the ATR system using 1-ft × 1-ft PWF data exclusively.

Figure 14 presents the performance curves for all three stages of the ATR system for 1-m × 1-m resolution PWF data; the corresponding performance curves for 1-ft × 1-ft PWF data are included for comparison. Note that the detection stage performs similarly for both the 1-ft × 1-ft data and the 1-m × 1-m data. The discrimination stage then reduces the number of false alarms by approximately a factor of three for the 1-m × 1-m data; this indicates a degradation of performance, compared to that obtained by using 1-ft × 1-ft data. Finally, the classification stage yields only a minor reduction in the false-alarm density for the 1-m × 1-m data, in marked contrast to the order-of-magnitude reduction obtained by using 1-ft × 1-ft resolution data.

In addition to the results presented in Figure 14 for PWF data, we examined the ATR performance using 1-m × 1-m HH data. These results, which are summarized in Table 5, show an order-of-magnitude increase in false alarms for the end-to-end ATR system.

Because the performance of the detection stage was comparable for 1-m × 1-m PWF data and 1-ft × 1-ft PWF data, as shown in Figure 14, we evaluated the performance of an ATR system using 1-m × 1-m data for the detection stage, followed by 1-ft × 1-ft data in the discrimination and classification stages. One reason for using lower-resolution data for the CFAR prescreener is to save on computational resources. In particular, by reducing the resolution from 1 ft × 1 ft to 1 m × 1 m, we reduce the amount of data processed in the prescreener stage by a factor of sixteen.

Figure 15 shows the performance of all three stages of the dual-resolution ATR system. Note that the discrimination and classification stages yield a substantial reduction in the false-alarm density, and that the end-to-end performance is essentially identical to that obtained by using 1-ft × 1-ft resolution PWF data exclusively.

Table 5 compares the number of false alarms that passed through each stage with various combinations of polarization and resolution. Table 5 also includes the best set of discrimination features for each combination of polarization and resolution. As expected, each stage performs worse with 1-m × 1-m data than with 1-ft × 1-ft data. However, the use of 1-m × 1-m data for the prescreening stage followed by the use of 1-ft × 1-ft data for the discrimination and classification stages does not degrade the overall end-to-end ATR performance.

Table 6 show the performance of the 1-m × 1-m resolution MSE classifier in the form of a confusion matrix for both PWF and HH data. At the output of the classification stage, 90% of the netted targets were classified as targets. However, the number of those correctly classified into tank and howitzer categories was low, indicating that it is difficult to separate these two target types with 1-m × 1-m resolution data.

Summary

This article examines the performance of an end-to-end ATR system designed for the detection and classification of tactical targets (tanks and howitzers) in high-resolution (1 ft × 1 ft) SAR data. The study was quite extensive, using a large data set (approximately 74 km^2 of clutter), a significant number of targets (136) divided equally between tanks and howitzers, and two polarizations (HH and PWF). In addition, the data were reprocessed to examine the performance that could be achieved with a lower-resolution sensor (1 m × 1 m).

Overall, the results indicate we achieve better performance with PWF imagery than with a single-polarization channel. Also, ATR performance improved as the image resolution increased. Thus 1-ft × 1-ft PWF imagery showed the best ATR performance, with a final false-alarm count of six (in 56 km^2 of Stockbridge data) at a detection probability of 0.9. The corresponding 1-ft × 1-ft single-polarization-channel performance using HH data had a final false-alarm count of 229 (in 56 km^2 of Stockbridge data) at the same detection probability of 0.9. For both PWF and HH data, the probability of classifying the targets correctly by target type was 100%.

Table 5. False-Alarm Statistics Comparison Using Stockbridge Data (56 km^2)

	PWF 1 ft × 1 ft	*PWF 1 m × 1 m detector 1 ft × 1 ft discriminator and classifier*	*PWF 1 m × 1 m*	*HH 1 ft × 1 ft*	*HH 1 m × 1 m*
Detector output (P_D = 1.0)	3457	13,176	13,176	8739	25,873
Discriminator output (P_D = 1.0)	709	949	4401	3758	18,562
Classifier output (P_D = 0.9)	6	6	1115	229	12,610
Best set of discrimination features					
Standard deviation	✔	✔	✔	✔	
Weighted fill			✔	✔	✔
Fractal dimension	✔	✔	✔		✔
Mass				✔	✔
Diameter	✔	✔			
Normalized rotational inertia	✔	✔			
Peak CFAR					✔
Mean CFAR	✔	✔	✔	✔	✔
Percent bright even			✔		
Percent pure even	✔	✔	✔		

Table 6. Performance of MSE Classifier (Detection Probability = 0.9)*

PWF data 1-m × 1-m resolution		*Classified as tank*	*Classified as howitzer*	*Classified as clutter*
	Tank	63	0	5
	Howitzer	50	9	9
	Clutter	1051	64	3286
HH data 1-m × 1-m resolution		*Classified as tank*	*Classified as howitzer*	*Classified as clutter*
	Tank	59	3	6
	Howitzer	30	30	8
	Clutter	11,508	1102	5952

* Training data: non-netted targets; test data: 136 netted targets

For the 1-m × 1-m data, the end-to-end performance was considerably degraded. With 56 km^2 of clutter at a detection probability of 0.9, the PWF imagery produced 1115 false alarms, whereas HH imagery produced 12,610 false alarms. In addition, the percentage of detected targets that were correctly classified by target type was unacceptably low.

Figure 16 presents a comparison of the end-to-end results for PWF and HH data, for both 1-ft × 1-ft resolution and 1-m × 1-m resolution. The performance curves indicate that the performance using 1-m × 1-m HH data is poor, the performance using 1-m × 1-m PWF data and 1-ft × 1-ft HH data is roughly comparable, and the performance using 1-ft × 1-ft PWF data is superior.

Also, we have shown that we can use lower-resolution data (1-m × 1-m PWF imagery) for the initial detection stage, followed by 1-ft × 1-ft PWF data in the discrimination and classification stages, without any degradation in the end-to-end ATR performance. This type of algorithm may be suited to situations in which computation resources are at a premium.

Additional work in the field of ATR is currently being pursued. Other polarization combinations are being examined (e.g., VV, HV, and HH – VV), along with new features to exploit further the performance of the discrimination stage. Concepts for improving the performance results of the MSE classifier are under investigation. Some of these concepts may ultimately provide us with a more robust classifier. For

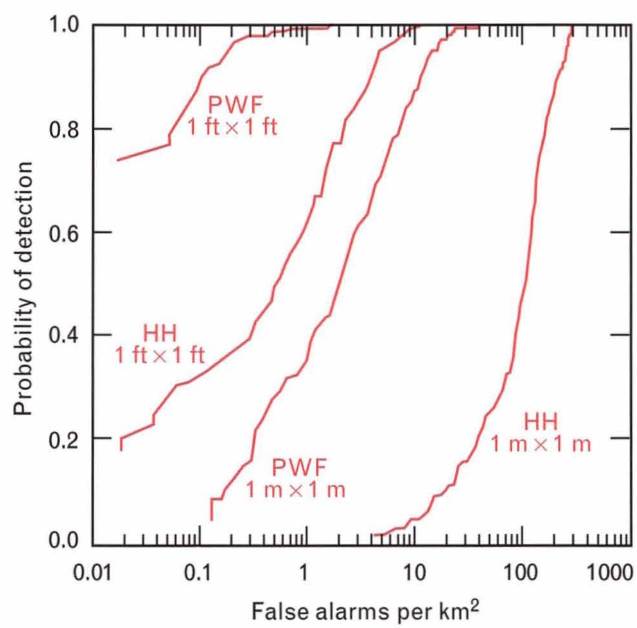

FIGURE 16. Comparison of end-to-end performance of the ATR system using both 1-ft × 1-ft resolution data and 1-m × 1-m resolution data, in HH polarization and PWF polarization.

example, one of these concepts involves the use of three-dimensional target information in the classifier templates [11].

Acknowledgments

This work was sponsored by the Advanced Research Projects Agency.

REFERENCES

1. L.M. Novak, G.J. Owirka, and C.M. Netishen, "Performance of a High-Resolution Polarimetric SAR Automatic Target Recognition System," *Linc. Lab. J.* **6**, 11 (1993).
2. J.C. Henry, "The Lincoln Laboratory 33 GHz Airborne Polarimetric SAR Imaging System," *IEEE National Telesystems Conf., Atlanta, GA, 26–27 Mar. 1991*, p. 353.
3. M.C. Burl, G.J. Owirka, and L.M. Novak, "Texture Discrimination in Synthetic Aperture Radar Imagery," *23rd Asilomar Conf. on Signals, Systems, and Computers, Pacific Grove, CA, 30 Oct.–1 Nov. 1989*, p. 399.
4. D.E. Kreithen, S.D. Halversen, and G.J. Owirka, "Discriminating Targets from Clutter," *Linc. Lab. J.* **6**, 25 (1993).
5. 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).
6. L.M. Novak, M.C. Burl, and W.W. Irving, "Optimal Polarimetric Processing for Enhanced Target Detection," *IEEE Trans. AES* **29**, 234 (1993).
7. G.B. Goldstein, "False-Alarm Regulation in Log-Normal and Weibull Clutter," *IEEE Trans. AES* **19**, 84 (1973).
8. W.W. Irving, G.J. Owirka, and L.M. Novak, "A New Clutter Model for High Resolution SAR Data," *SPIE* **1630**, 208 (1992).
9. K. Fukunaga, R.R. Hayes, and L.M. Novak, "The Acquisition Probability for a Minimum Distance One-Class Classifier," *IEEE Trans. AES* **23**, 493 (1987).
10. R.O. Duda and P.E. Hart, *Pattern Classification and Scene Analysis*, John Wiley, New York (1973), pp. 24–32.
11. S.M. Verbout, W.W. Irving, and A.S. Hanes, "Improving a Template-Based Classifier in a SAR Automatic Target Recognition System by Using 3-D Target Information," *Linc. Lab. J.* **6**, 53 (1993).

APPENDIX: DETAILS OF THE DISCRIMINATION FEATURES

THIS APPENDIX PRESENTS DETAILS of the discrimination features used in the studies of the automatic target recognition (ATR) system described in this article. The types of features used in the discrimination stage of the ATR system include textural, size, contrast, and polarimetric features.

Lincoln Discrimination Features

Standard-Deviation Feature. The standard-deviation feature is a measure of the fluctuation in intensity (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.

Fractal-Dimension Feature. The fractal-dimension feature provides a measure of the spatial dimensionality of the detected object [1]. This feature estimates the Hausdorff dimension of the spatial distribution of the N brightest scatterers in the region of interest. For example, a straight line has a Hausdorff dimension of one and a solid rectangle has a Hausdorff dimension of two. Various other space-filling objects with holes have Hausdorff dimensions that fall between one and two. An isolated point (or a disjoint set of isolated points) has a Hausdorff dimension of zero.

Weighted-Fill Feature. The third textural feature, the weighted-fill ratio, measures the percentage of the total energy contained in the brightest five percent of the scatterers of a detected object. For man-made objects a significant portion of the total energy is reflected by a small number of bright scatterers; for natural clutter the total energy is distributed more evenly among the pixels.

ERIM Discrimination Features

Nine of the discrimination features used in these studies were developed at the Environmental Research Institute of Michigan (ERIM). These features were provided to Lincoln Laboratory under the Strategic Target Algorithm Research (STAR) contract. Instead of using a target-sized rectangular box as a preliminary step in feature calculation, the discrimination stage computes the ERIM features from the pixels contained in a target-shaped blob obtained by performing morphological operations. These operations serve both as a method of grouping spatially related detections from the prescreener and as a method of estimating the size, shape, and orientation of the detected object. There are three categories of ERIM discrimination features: size-related features, contrast-based features, and polarimetric features. Each of these three categories contains three features. We describe each feature in the following paragraphs.

ERIM Size Features. The three size-related features (mass, diameter, and normalized rotational inertia) utilize only the binary image created by the morphological operations. The mass feature is computed by simply counting the number of pixels in the morphological blob. The diameter is calculated as the length of the diagonal of the smallest rectangle (either horizontally oriented or vertically oriented) that encloses the blob. The normalized rotational inertia is the second mechanical moment of the blob around its center of mass, normalized by the inertia of an equal mass square.

ERIM Contrast Features. The contrast-based features (peak CFAR, mean CFAR, and percent bright CFAR) are determined by the CFAR algorithm. The CFAR statistic (given in Equation 1 in the main text) is computed for each pixel to create a CFAR image. The peak CFAR feature is the maximum value in the CFAR image contained within the target-shaped blob. This quantity is usually identical to the basic CFAR detection statistic used in the prescreener algorithm. The mean CFAR feature is the average of the CFAR image taken over the target-shaped blob. The percent bright CFAR feature is the percentage of pixels within the target-shaped blob that exceed a CFAR threshold empirically determined from the target-training data.

ERIM Polarimetric Features. The polarimetric discrimination features require calibrated, fully polarimetric data. These features are based on a transformation of the linear polarization basis in which the Lincoln Laboratory MMW SAR sensor gathers data to an even-bounce, odd-bounce basis described by the equations

$$E_{\text{odd}} = \frac{|\text{HH} + \text{VV}|^2}{2}$$

and

$$E_{\text{even}} = \frac{|\text{HH} - \text{VV}|^2}{2} + 2|\text{HV}|^2 ,$$

where HH (horizontal transmit, horizontal receive), HV (horizontal transmit, vertical receive), and VV (vertical transmit, vertical receive) are the three linear components of the reflected energy. The odd-bounce channel (E_{odd}) given by the first equation corresponds to the radar return from a flat plate or a trihedral; the even-bounce channel (E_{even}) corresponds to the radar return from a dihedral. The motivation for this polarimetric basis resides in the fact that few dihedral structures exist in natural clutter, whereas dihedral structures are plentiful on most man-made targets; therefore, natural clutter tends to return less even-bounce energy than man-made objects.

The polarimetric features are calculated from the even-bounce and the odd-bounce images. The percent-pure feature is the fraction of pixels within the target-shaped blob for which at least a certain fraction of the scattered energy falls in either the even-bounce or the odd-bounce channel. Percent pure even is the fraction of "pure" pixels within the target-shaped blob for which at least a certain fraction of the scattered energy falls in the even-bounce channel. The percent-bright-even feature is the fraction of pixels within the target-shaped blob that exceed a certain value in the CFAR image described above, and which are mainly even-bounce scatterers.

REFERENCES

1. M.C. Burl, G.J. Owirka, and L.M. Novak, "Texture Discrimination in Synthetic Aperture Radar Imagery," *23rd Asilomar Conf. on Signals, Systems, and Computers, Pacific Grove, CA, 30 Oct.–1 Nov. 1989,* p. 399.

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 he 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.

SHAWN D. HALVERSEN
is an associate staff member in the Surveillance Systems group. His research speciality is in the detection and discrimination of stationary targets. He received a B.A. degree in mathematics and a B.S. degree in physics from the University of South Florida, and an M.A. degree in mathematics from the University of Wisconsin. He has been at Lincoln Laboratory since 1990.

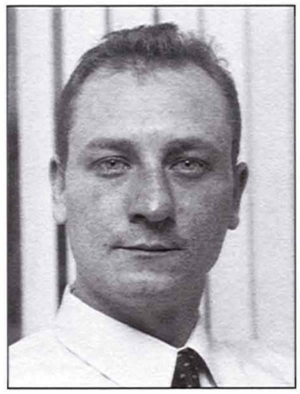

GREGORY J. OWIRKA
is an assistant staff member in the Surveillance Systems group. He received a B.S. degree (cum laude) in applied mathematics from Southeastern Massachusetts University, and he is currently working on an M.S. degree in electrical engineering at Northeastern University. Greg's current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

MARGARITA HIETT
is an assistant staff member in the Surveillance Systems group. She received a B.S. degree in electrical engineering from the University of Massachusetts at Amherst. From 1986 to 1992 she held a technical staff position at Raytheon, where she was involved in the implementation of large-scale real-time hardware and software systems for detection and discrimination of IR targets. From 1992 to 1993 she worked at Textron, where she continued her work analyzing real-time algorithms for detection of moving IR targets. She has been at Lincoln Laboratory since 1993, and her current research work is in the analysis of detection algorithms for stationary radar targets.