Radar Target Identification Using Spatial Matched Filters

L.M. Novak, G.J. Owirka, and C.M. Netishen
MIT Lincoln Laboratory

Abstract

The application of spatial matched filter classifiers to the synthetic aperture radar (SAR) automatic target recognition (ATR) problem is being investigated at MIT Lincoln Laboratory. Initial studies investigating the use of several different spatial matched filter classifiers in the framework of a 2D SAR ATR system are summarized. In particular, a new application is presented of a shift-invariant, spatial frequency domain, 2D pattern-matching classifier to SAR data. Also, the performance of this classifier is compared with three other classifiers: the synthetic discriminant function, the minimum average correlation energy filter, and the quadratic distance correlation classifier.

Introduction

In support of the DARPA-sponsored Critical Mobile Target (CMT) program, MIT Lincoln Laboratory has developed a complete, end-to-end, 2D SAR automatic target recognition (ATR) system. This baseline ATR system performs three basic functions: first, a CFAR (constant false alarm rate) detector locates candidate targets in a SAR image on the basis of radar amplitude. Next, a target-size matched filter is used to accurately locate the candidate targets and determine their orientation; textural discriminants (fractal dimension, standard deviation, and ranked fill ratio) are then used to reject natural-clutter false alarms [1]. Finally, a pattern-matching classifier is used to reject cultural false alarms (man-made clutter discretely) and classify the remaining detections.

High resolution ($0.3 \text{ m} \times 0.3 \text{ m}$), fully polarimetric target and clutter data gathered by the Lincoln Laboratory Millimeter-wave Sensor have been used to evaluate the performance of the ATR system [2]. Prior to ATR processing, an optimal polarimetric processing technique known as the PWF (polarimetric whitening filter) is used to combine the HH, HV, and VV polarization components into minimum-speckle SAR imagery [3]. This processing technique has been shown to improve the performance of the detection, discrimination, and classification algorithms by reducing the clutter variance and by enhancing the target signature relative to the clutter background. The robustness of the ATR system has been demonstrated by testing it against targets with and without radar camouflage.

The ultimate goals of this system are to operate in a wide-area search mode, maintain a very low false alarm density on the order of 1 false alarm per 1000 km^2 search area, and provide a high probability of detection ($P_d \approx 0.8$). To meet the stringent false alarm requirement, it is essential for the classifier to reliably reject man-made clutter discretely. Also, to maintain a high probability of correct classification (P_{cc}), the classifier must provide good separation between classes and must be robust with respect to target variability.

We have recently implemented several spatial matched filter classifiers, as possible alternatives to the pattern-matching classifier used in the baseline ATR system. This paper describes them and presents preliminary performance results. Preprocessing methods and target variability issues are addressed. Performance results are given for a shift-invariant template matcher that we have recently developed. Algorithms are compared by presenting classifier-performance confusion matrices, which indicate the probability of correct and incorrect classification. The ability of each classifier to reject cultural false alarms (buildings, bridges, etc.) is quantified.

ATR System Review

This paper focuses on target classification algorithms. This section describes our complete SAR ATR system (detection, discrimination, and classification), in order to place the classification stage into a more general system context. A simplified block diagram of the multi-stage ATR system is shown in Figure 1; each stage of the multi-stage system is briefly discussed below.

Stage 1. In the first stage of processing, a two-parameter CFAR detector [4] is used as a prescanner; this stage of processing locates potential targets in the image on the basis of radar amplitude. Since a single target may produce multiple detections, the CFAR detections are clustered (grouped together). Then a $128 \text{ pixel} \times 128 \text{ pixel}$ region of interest (ROI) around the centroid of each cluster is passed to the next stage of the algorithm for further processing.

Figure 1: Simplified block diagram of baseline automatic target recognition (ATR) system.

Stage 2. The second stage of processing, called discrimination, takes as its input each ROI and analyzes it. The goal of discrimination processing is to reject natural-clutter false alarms while accepting real targets. This stage consists of three steps: (1) determining the position and orientation of the detected object, (2) computing simple textural features, and (3) combining the features into a discrimination statistic that measures how "target-like" the detected object is.

In order to determine the position and orientations of the detected object, a target-size rectangular template is placed on the image and of slid and rotated until the energy within the template is maximized. This operation is computationally

quick, since it is performed only on the ROIs obtained by the Stage 1 CFAR detector. Mathematically, this operation is equivalent to processing with a 2D matched filter where the orientation of the detected object is unknown.

In step 2 of the discrimination stage, three textural features are calculated:

(1) The standard deviation of the data within the target-size template. This feature measures the statistical fluctuation of the data. Targets typically exhibit significantly larger standard deviations than natural clutter.

(2) The fractal dimension of the pixels in the ROI. This feature provides information about the spatial distribution of the brightest scatterers of the detected object. It is complementary to the standard deviation feature, which depends only on the intensities of the scatterers and not on their spatial locations.

(3) The ranked fill ratio of the data within the target-size template. This feature measures the fraction of the total target energy contained in the brightest 5% of the detected object's scatterers. For targets, a significant portion of the total power is due to a small number of very bright scatterers; for natural clutter, the total power is distributed more evenly among the scatterers.

Reference [1] provides a detailed description of the three textural discrimination features used in the baseline ATR system.

Stage 3. The third and final stage of ATR processing is classification. Here a 2D pattern-matching algorithm is used to (1) reject cultural false alarms caused by man-made clutter discretes (buildings, bridges, etc.) and (2) classify the remaining detected objects. A four-class classifier (tank, APC, self-propelled gun, and clutter) is used. Those detected objects that pass the discrimination stage are matched against stored reference templates of the tank, APC, and gun targets. 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 with the highest match score.

Data Description

For this preliminary study, tactical target data of tanks, APCs, and self-propelled guns gathered in the 1989 Stockbridge, New York, data collection was used. Figure 2 shows photographs of each of the tactical targets. The target data consisted of imagery collected in spotlight mode at 1 deg azimuthal increments. In evaluating the performance of each classification algorithm, every other image (i.e. odd frame numbers) was used for training, and algorithm testing was performed using the even numbered frames. All of the targets used in these evaluations were bare (i.e. uncamoouflaged).

Tank

APC

Gun

Figure 2: Photographs of the three tactical targets used in these studies.

The clutter data consisted of 100 cultural clutter discretes selected from 50 km of Stockbridge stripmap imagery. The clutter discretes had all been detected by the Stage 1 algorithm and passed by the Stage 2 discriminator. In addition, they were selected as visually target-like. Figure 3 shows a PWF-processed fully polarimetric image of one of the clutter discretes used in these classifier studies — an above-ground swimming pool located behind a house; Figure 4 shows the corresponding optical (ground truth) photograph of this scene.

Figure 3: A SAR image of a cultural clutter discrete — an above-ground swimming pool located behind a house.

Figure 4: An optical photograph of the cultural clutter scene shown in Figure 3.

The classification algorithms were tested using fully polarimetric (HH, HV, and VV) data, optimally combined using the PWF. Some testing was also performed on single-channel (HH) data; however, the results confirmed earlier classification results [5], which showed that PWF-processed imagery generally provides better classification performance than single-channel (HH) imagery. Therefore, the HH results are not discussed further in this paper.

We also investigated several methods of image preprocessing, such as absolute amplitude, normalized amplitude, and normalized dBs. All of the algorithms evaluated in this paper showed improved performance when dB normalization was applied; this is consistent with results presented in reference [5].

Description of the Spatial Matched Filter Classifiers

In this section, we describe four classifiers that were investigated as alternatives to the pattern-matching classifier used in the baseline ATR system. The classifiers are the synthetic discriminant function (SDF), the minimum average correlation energy (MACE) filter, the quadratic distance correlation classifier (QDCC), and the shift-invariant 2D pattern-matching classifier.

The basic structure of the SDF and MACE filter is characterized in the frequency domain by the expression [6,7],

$$H = A \bar{X} \left(\bar{X}^\dagger A \bar{X} \right)^{-1} U \quad (1)$$

where H denotes the DFT of the spatial matched filter. The matrix $\bar{X}$ is composed of a set of target training vectors obtained by taking the DFTs of the target training images. The vector U represents a set of constraints imposed on the values of the correlation peaks obtained when the training vectors are run through the spatial matched filter. The matrix A represents a positive definite weighting matrix that can be selected to produce various types of spatial correlation filters.

Synthetic Discriminant Function (SDF)

If the weighting matrix is set equal to the unity matrix ($A = I$), the spatial matched filter defined in equation (1) reduces to

$$H = \bar{X} \left(\bar{X}^\dagger \bar{X} \right)^{-1} U \quad (2)$$

which is the conventional projection SDF [7]. Figure 5 presents a simplified block diagram showing how the SDF classifier was implemented: three separate SDF filters were constructed using the tank, APC, and gun training data. Also shown in the figure is a typical output obtained when a tank training image was run through the tank SDF filter. The output correlation peak is 1.0 (as expected).

Figure 5: Synthetic discriminant function (SDF). Simplified block diagram shows the separate SDF correlators for the three target classes. Graph shows a typical correlator output.

Minimum Average Correlation Energy (MACE) Filter

The MACE filter is obtained by letting $A = D^{-1}$ in equation (1), where

$$D = \sum_{i=1}^N D_i \quad (3)$$

and N is the number of training images.

Each D_i matrix is constructed by placing the power spectrum of the i th training vector along the diagonal:

$$D_i = \begin{bmatrix} |\bar{X}_i(1)|^2 & & & \\ & |\bar{X}_i(2)|^2 & & \\ & & \ddots & \\ & & & |\bar{X}_i(p)|^2 \end{bmatrix} \quad (4)$$

where p is the dimension of the training vectors.

Thus, the MACE filter is defined in the frequency domain by the expression

$$H = D^{-1} \bar{X} \left(\bar{X}^\dagger D^{-1} \bar{X} \right)^{-1} U \quad (5)$$

It can be shown [6] that the effect of D^{-1} (the inverse of the average target power spectrum) in equation (5) is equivalent to applying a whitening filter to the target data. Note that since the D matrix is diagonal, its inversion is trivial.

Figure 6 presents a block diagram showing how the MACE classifier was implemented: three separate MACE filters were constructed using the tank, APC, and gun training data. Also shown in the figure is a typical output obtained by running a tank training image through the tank MACE filter. As with the SDF (Figure 5), the correlation peak is 1.0. Here, however, the correlation value is very small elsewhere. The MACE filter is designed to minimize the average correlation energy away from the peak.

Figure 6: Minimum average correlation energy (MACE) filter. Simplified block diagram shows the separate MACE correlators. Graph shows a typical correlator output.

Quadratic Distance Correlation Classifier (QDCC)

In the QDCC, the DFT of the spatial matched filter is expressed by

$$H = S^{-1}(m_1 - m_2) \quad (6)$$

where m_1 and m_2 are the means of the DFTs of the training images for class 1 and class 2, respectively. S is a diagonal similarity matrix, which is defined by

$$S = \begin{bmatrix} \frac{1}{N} \sum_{i=1}^N (X_i - M_1)^\dagger (X_i - M_1) \\ + \frac{1}{N} \sum_{i=1}^N (Y_i - M_2)^\dagger (Y_i - M_2) \end{bmatrix} \quad (7)$$

where M_1 and M_2 are matrices with the elements of m_1 and m_2 placed on the main diagonal, and X_i and Y_i are the i th training vectors from class 1 and class 2 (also in diagonal matrix form), respectively.

In this classifier, quadratic distance measures are calculated as follows:

$$d_z(1) = |Z^* H - M_1^* H|^2 \quad (8)$$

$$d_z(2) = |Z^* H - M_2^* H|^2 \quad (9)$$

where Z is a matrix constructed by placing the elements of the DFT test vector z on the main diagonal. The quadratic distances $d_z(1)$ and $d_z(2)$ measure the similarity of the test input z to class 1 and class 2, respectively. Note that the development of equations (6)-(9) is only valid for a two-class problem. For an n -class problem the QDCC requires $n(n-1)$ filters.

Figure 7 presents a simplified block diagram showing how the QDCC classifier was implemented: to construct a three-class classifier, three separate two-class classifiers (tank vs. APC, tank vs. gun, and APC vs. gun) were used; a voting scheme was used to obtain the final tank/APC/gun result. A complete derivation of the QDCC is given in reference [8].

Figure 7: Simplified block diagram of quadratic distance correlation classifier (QDCC).

Shift-Invariant 2D Pattern-Matching Classifier

The baseline ATR system described earlier employs a 2D pattern-matching classifier that uses $0.3 \text{ m} \times 0.3 \text{ m}$ resolution PWF processed data that has been dB-normalized. The reference templates used by this classifier are generated by extracting and normalizing a rectangular “window” of data around each target training image. Since the aspect angles of the target training images vary, so does the orientation angle of the rectangular window. Also, since the sizes of targets (e.g. tank, APC, and gun) used to train the classifier differ, so do the sizes of the windows for the various targets. The stored templates are correlated against the test image; since a target may not be perfectly centered in a test image, the windowed reference templates are slid, and match scores are computed over a range of X-Y translations. The test image is assigned to the class with the best correlation score, provided the score is greater than 0.7. If the best score is 0.7 or less, the test image is called clutter. This approach is computationally inefficient.

To achieve shift invariance, we took the following approach. Reference templates are generated using a rectangular window that is fixed regardless of target size and orientation. These templates are then normalized and their DFTs taken; the DFTs are stored in a filter bank. Input test images are processed in a similar manner. A test image is normalized, and its DFT taken. Then the correlation scores are calculated by

$$\rho_i = \text{MAX} \left\{ \text{DFT}^{-1} \left[R_i \times T^* \right] \right\} \quad (10)$$

where T is the DFT of the dB-normalized test image and R_i is the i th reference template. The maximum value of the correlation surface is the correlation score. The shift-invariant 2D pattern-matching classifier is more computationally efficient than the baseline 2D pattern matcher.

Classification Performance Results

This section presents the performance results obtained using the classification algorithms described in the previous section. All of the results were obtained using $0.3 \text{ m} \times 0.3 \text{ m}$ resolution, PWF-processed imagery that had been dB-normalized.

Table 1. Performance of SDF classifier

Percent classified as (%)				
	Tank	APC	Gun	Clutter
Tank	100			
APC		100		
Gun			100	
Clutter	6	21	1	71

Table 1 is a confusion matrix that summarizes the performance results obtained using the SDF classifier. The probability of correctly classifying the test targets was 100%, and the percentage disretes rejected was 71%.

Table 2. Performance of MACE Classifier

Percent classified as (%)				
	Tank	APC	Gun	Clutter
Tank	94			6
APC		53		47
Gun			76	24
Clutter	21	3	13	63

Table 2 summarizes the performance results obtained using the MACE classifier. The average probability of correct target classification was 74%, and the percentage of cultural clutter disretes rejected was 63%.

We found that the MACE classifier was very sensitive to small aspect angle variations of the test targets. Figures 8(a) and (b) compare plots of correlation scores vs. target aspect angle for the SDF and MACE filters, respectively. For both filters, when the training data (odd numbered frames) were used as test inputs, the correlation scores were 1.0 (as expected). When the true test data (even numbered frames) were input to the SDF classifier, the correlation scores were close to 1.0 (see Figure 8(a)). However, when the true test data were input to the MACE classifier, the correlation scores were poor (see Figure 8(b)). Both SDF and MACE did a poor job of rejecting the tough cultural clutter disretes used in this study.

Figure 8: Correlation scores vs. target aspect for the SDF and MACE classifiers. For both classifiers, tank data were used for both training and testing.

Table 3. Performance of QDCC Classifier

Percent classified as (%)				
	Tank	APC	Gun	Clutter
Tank	100			
APC		100		
Gun			100	
Clutter	1	5		94

Table 3 summarizes the performance of the QDCC classifier. The probability of correct classification the test targets was 100%, and the percentage of cultural clutter discretes rejected was 94%. As discussed previously, for our three-class (tank, APC, gun) problem, the QDCC was implemented using three two-class classifiers, and a voting scheme was used to obtain the final results. The QDCC calculates ratios of distances to the various classes; the final tank/APC/gun decision is based on these ratios. Figure 9 shows a 3D plot of the three distance ratios:

$$r_{\text{APC}/\text{tank}} = \frac{d_z(\text{APC})}{d_z(\text{tank})} \quad (11)$$

$$r_{\text{gun}/\text{tank}} = \frac{d_z(\text{gun})}{d_z(\text{tank})} \quad (12)$$

$$r_{\text{APC}/\text{gun}} = \frac{d_z(\text{APC})}{d_z(\text{gun})} \quad (13)$$

where for example, $d_z(\text{tank})$ is the distance of test vector z to the tank class. The QDCC was found (1) to be relatively insensitive to target aspect angle variation, (2) to provide good separation between classes for the test targets, and (3) to reject a high percentage of the cultural clutter discretes.

Figure 9: Distance ratios obtained by the quadratic correlation classifier.

Table 4. Performance of shift-invariant 2D pattern matching classifier

Percent classified as (%)				
	Tank	APC	Gun	Clutter
Tank	100			
APC		100		
Gun			100	
Clutter		2	4	94

Table 4 summarizes the performance of the shift-invariant 2D pattern-matching classifier. The probability of correct classification for the test targets was 100%, and the percentage of cultural clutter discretes rejected was 94%. Figure 10 shows the best match score from each filter bank vs. target aspect angle; the test inputs were all tanks, and the highest match scores were all from the tank reference class. Similar results for the APC and gun test inputs are shown in Figures 11 and 12. These results illustrate the very good separation between classes produced by the shift-invariant 2D pattern-matching classifier.

Figure 10: Best match scores from the three filters vs. target aspect angle. Input test target was a tank.

Figure 11: Best match scores from the three filters vs. target aspect angle. Input test target was an APC.

Figure 12: Best match scores from the three filters vs. target aspect angle. Input test target was a self-propelled gun.

Summary

This paper compared the performance of several spatial matched filter classifiers in the framework of a SAR ATR system. The QDCC and shift-invariant 2D pattern-matcher provided the best overall performance in terms of probability of correct target classification and rejection of cultural clutter discretely. Much work remains to be done before we can declare which classifier is the best SAR ATR classifier. We need to evaluate the QDCC and the shift-invariant 2D pattern matcher using a much larger database of targets and clutter false alarms gathered under a variety of conditions including snow clutter. This preliminary study used 35 deg of target aspect angle to design and test each of the filters. A complete classifier would probably use twelve 30 deg filters to cover the complete 360 deg of target aspect angle. Also, we need to test other classifier approaches, such as one based on eigenvalue/eigen-image concepts developed by Turk and Pentland [9].

References

- [1] M. C. Burl, et al, "Texture Discrimination in Synthetic Aperture Radar", 23rd Asilomar Conf. on Signals, Systems, and Computers, Pacific Grove, California, November (1989).
- [2] J. C. Henry, "The Lincoln Laboratory 35 GHz Airborne Polarimetric SAR Imaging Radar System", IEEE National Telesystems Conf., Atlanta, Georgia, March (1991).
- [3] L. M. Novak, et al, "Optimal Processing of Polarimetric Synthetic Aperture Radar Imagery", Lincoln Laboratory J. 3(2), (Summer 1990).
- [4] G. B. Goldstein, "False Alarm Regulation in Log-Normal and Weibull Clutter", IEEE Trans. Aerospace Electron. Syst. January (1973).
- [5] L. M. Novak, et al, "Performance of a High Resolution SAR Automatic Target Recognition System", Lincoln Laboratory J. 6(1), (Spring 1993).
- [6] A. Mahalanobis, et al., "Minimum Average Correlation Energy Filters", Appl. Optics 26, 3633-3640 (1986).
- [7] D. Casasent et al, "Correlation Synthetic Discriminant Functions", Appl. Optics 25, 2343-2350 (1986).
- [8] A. Mahalanobis et al. "A Quadratic Distance Classifier for Multi-Class SAR ATR Using Correlation Filters", SPIE Conf. on SAR, Los Angeles, California, January (1993).
- [9] M. Turk et al, "Eigenfaces for Recognition", J. Cognitive Neuroscience 3 (1), (1991).