



# Target-Centered Models and Information-Theoretic Segmentation for Automatic Target Recognition

MICHAEL D. DEVORE

*Electronic Systems and Signals Research Laboratory, Dept. of Electrical Engineering, Washington University, St. Louis, MO 63130*

mdd2@cis.wustl.edu

JOSEPH A. O'SULLIVAN

*Electronic Systems and Signals Research Laboratory, Dept. of Electrical Engineering, Washington University, St. Louis, MO 63130*

jao@ee.wustl.edu

*Received May 9, 2001; Revised May 2, 2002; Accepted May 10, 2002*

**Abstract.** We present an approach to automatic target recognition (ATR) from synthetic aperture radar (SAR) imagery which combines advantages of both model-based and template-based approaches. Prior observations are used to estimate the statistical properties of reflectance over regions in the training scene. These target-centered statistical models can then be used to estimate the statistical properties of sensor output for arbitrary pose. Two-sided hypothesis tests which are maximally powerful at the most likely alternative are developed in an information-theoretic framework to address target model segmentation and confuser rejection. Segmentation of target from clutter is performed in the target-centered coordinate system using all prior observations to produce a consistent segmentation over all poses. We present performance and computation complexity results as a function of segmentation threshold, confuser-rejection threshold, and operating conditions for publicly available SAR data.

**Key Words:** automatic target recognition, synthetic aperture radar, statistical hypothesis testing, performance-complexity, resource consumption rate, information-based segmentation, information-based confuser rejection, MSTAR

## 1. Introduction

Automatic target recognition (ATR) systems are generally based on some characterization of sensor output conditioned on the class, state, and pose of objects which may be in a scene. In the problem of ATR based upon synthetic aperture radar (SAR) images, several approaches to object characterization have been reported in the literature. The family of approaches generally referred to as model-based methods represent various attempts to characterize the physical structure of the target classes of interest which are then used to predict the SAR sensor output for arbitrary combinations of target class, pose, etc. Classification of an object in some observation is performed by searching over the hypothesis space for the parameter combination that yields a predicted observation closest to the actual observation. The physical structure of the targets can be represented by CAD models as in the Moving and Stationary Target Acquisition and Recognition (MSTAR) ATR system [19] with features or signatures predicted from electromagnetic modeling tools such as described by Keydel and Lee [9]. A variation on this approach is to extract information about primitive shapes comprising

a target class from sets of training images of those targets as in the work of Richards, et al. [13] and Chiang, et al. [1].

The family of approaches to ATR often referred to as template-based methods attempt to build a database of typical sensor outputs from training images of a wide range of target classes and parameterizations such as reported by Owirka, et al. [12] and Ross, et al. [14]. Such databases represent a direct model of the sensor output for each target class bypassing a model of the target structure itself. The database is typically searched for the sensor output which best resembles an actual observation, and the corresponding target class and parameters are declared as the ATR system output. In what may be viewed as a hybrid between model-based and template-based approaches, some have sought to start with a model describing the physics of SAR imaging and from this model determine a statistical model for the sensor output [2], [11]. Unknown parameters of the statistical model are estimated from training data and stored as a function of target class, pose, etc. Similar to the template-based approaches, a target in an observation is classified by searching the parameter database for the set of parameters which, when applied to the statistical sensor model, best describe the observation.

Model-based approaches have the advantage of being able to predict sensor output at arbitrary target pose, but construction of the models can be problematic given limited training data. Template-based methods often bin sample data into some partitioning of pose-space and determine average templates over each pose subset. Such partitioning has the advantage of not requiring large amounts of training data but suffers from an inability to describe the likely sensor output at an exact target pose. In this paper, we consider methods which combine the ability of model-based approaches to predict signatures of targets in previously unseen poses with the statistical sensor modeling and ability to employ training images enjoyed by template-based approaches. The idea is to estimate the statistical properties not of pixels in an image but of patches in the scene as a function of viewing angle. The result is a collection of distribution parameter functions defined in a target-centered coordinate system. Verbout, et al. [17] build three-dimensional templates which they use to predict SAR images under previously unseen squint angles but still maintain templates for each combination of depression and azimuth angles. By comparison, we develop target models in which each location on the target surface has an associated function of depression and azimuth that characterizes the region of a SAR image corresponding to the target location. These target models and associated statistical functions comprise deformable templates which can accommodate an infinite variety of target pose and radar power fluctuation.

Given an estimated model of a training scene, it is desirable to segment the model into target and clutter regions. Also, given a SAR image, it is desirable to infer whether or not the image is of a target for which a model exists. Both of these problems can be phrased as a two-sided hypothesis test  $\theta = \theta_0$  versus  $\theta \neq \theta_0$ , where  $\theta_0$  characterizes a known distribution. No uniformly most powerful test, which has maximum power for any given significance across all possible alternatives, exists for this problem. Often, researchers impose restrictions on the set of tests under consideration and select the test which is most powerful across the subset. For example, the restriction to unbiased tests, for which the power is larger than the significance level, yields the uniformly most powerful unbiased

(UMPU) test (cf. [15, p. 349]). By contrast, we employ the test which maximizes power at the most likely alternative, yielding a decision based upon the empirical relative entropy between distributions parameterized by  $\hat{\theta}$  and  $\theta_0$ . Segmentation affects the way target classification is performed because the physical dimensions of segmented target models varies with target class, and the size of a target's projection onto the slant plane varies with pose. We discuss modifications to the generalized likelihood ratio test that accommodate varying lengths of observation vectors and relate segmentation to the problem of dimensionality reduction.

In Section 2, we describe conditional probabilistic models for SAR imagery, two-sided hypothesis testing, and the use of variable-sized target images in classification and confuser rejection. Section 3 describes the use of training data to estimate target-centered models of the training scene, segmentation of the training scene to obtain target models, and prediction of image statistical properties from arbitrary pose. Section 4 contains an analysis of the results obtained by applying these methods to a variety of multi-class ATR problems using SAR data from the Moving and Stationary Target Acquisition and Recognition (MSTAR)<sup>1</sup> program. Conclusions follow in Section 5.

## 2. Model-Based ATR

In this section we briefly review the conditionally complex Gaussian model for SAR imagery and its implications for ATR. We discuss similarities between image segmentation and dimensionality reduction, and we present a classification algorithm motivated by this connection. Finally, we address the problem of deciding whether an image is of a target class not represented among the available models.

### 2.1. SAR Image Model

The scene being imaged by the sensor platform represents a source of information which may be regarded as probabilistic in nature with some *a priori* probability mass function on the class  $A$  of objects in the scene,  $P_A(a)$ , and some conditional probability density on the pose  $\Psi$  that a given object class is likely to present,  $p_{\Psi|A}(\psi|a)$ . In the most general case,  $\Psi$  is a vector with continuous-valued components that capture the three rotational degrees of freedom for a target, the three translational degrees of freedom, and the states of articulated components. Information from the scene is collected via the sensing platform which yields an observation  $\mathbf{R}$  that has a probabilistic dependence  $p_{\mathbf{R}|A,\Psi}(\mathbf{r}|a, \psi)$  on the scene. The observation  $\mathbf{R}$  represents a measurement in vector form, and for the SAR imaging scenario we can take  $\mathbf{R}$  to be an  $mn$  length vector with components taken row-wise from an  $m \times n$  SAR image.

The complex Gaussian distribution for  $p_{\mathbf{R}|A,\Psi}$  is introduced for high resolution radar in [7] and is extended to SAR in [11]. According to this model, a resolution cell in the scene contains a number of point-like scatterers with complex return signals that combine coherently to produce a complex value  $R_i$  in the corresponding pixel. If the number of scatterers is large and no one dominates within a resolution cell, the real and imaginary

components of  $R_i$  follow an approximate Gaussian distribution. If the phase of the scatterers is uniformly distributed, the pixel value will have a complex mean of zero. Assuming that the pixels in a SAR image represent nonoverlapping resolution cells in the scene, the complex pixel values are independent from one another and the conditional density function for an entire SAR image is a product over all pixel  $i$  of the conditional density for each pixel. We denote by  $\sigma_i^2(a, \psi)$  the variance of pixel  $i$  when target  $a$  presents itself with pose  $\psi$ .

The power in the radar echo is a function of the distance to the target. Additionally, it is known that the power transmitted by a radar platform can fluctuate from one device to another within the same model line and for any given unit the power output can drift from a nominal value over time. This power fluctuation has the effect that the pixels in a reconstructed SAR image may be uniformly scaled by some constant  $c$ . We incorporate this constant into the model by assuming that the received SAR image  $\mathbf{R} = c\mathbf{Y}$ , where  $\mathbf{Y}$  represents the image received under some nominal power. Then  $R_i$  is a complex, conditionally Gaussian random variable with variance  $c^2\sigma_i^2(a, \psi)$  and the conditional density for  $\mathbf{R}$  becomes

$$p_{\mathbf{R}|A,\Psi}(\mathbf{r}|a, \psi) = \prod_i \frac{1}{\pi c^2 \sigma_i^2(a, \psi)} e^{-|r_i|^2 / c^2 \sigma_i^2(a, \psi)}. \quad (1)$$

This conditional density is a continuous-valued function defined over the discrete parameter  $A$  and the continuous pose parameter  $\Psi$ . In Section 3 we consider target models that allow us to approximate  $p_{\mathbf{R}|A,\Psi}$  for arbitrary  $\Psi$ .

## 2.2. Two-Sided Hypothesis Testing

It is well known that the likelihood ratio test maximizes the power over all other tests for any given significance level for the case of a simple hypothesis  $\mathcal{H}_0$  against a simple alternative  $\mathcal{H}_A$ . For an observation  $\mathbf{r}$  with independent components, this test is equivalent to choosing  $\mathcal{H}_A$  if  $\sum_{i=1}^N \ln[p_A(r_i)/p_0(r_i)] > \gamma$ . In many cases, we need to test a simple  $\mathcal{H}_0$ ,  $\mathbf{r} \sim p_0$ , against a compound alternative  $\mathcal{H}_A$ ,  $\mathbf{r} \not\sim p_0$ . This is a two-sided hypothesis testing problem for which there is no uniformly most powerful test against all alternatives. We therefore seek a test based upon the most likely alternative. The resulting test will form the basis for both confuser rejection and target model segmentation.

In the case of zero-mean complex Gaussian distributions,  $\mathcal{H}_0$  is an assertion that  $r_i \sim \mathcal{CN}(0, \sigma_i^2)$ . For any given alternative,  $r_i \sim \mathcal{CN}(0, \xi_i^2)$ , the likelihood ratio test becomes

$$\sum_{i=1}^N \ln\left(\frac{\sigma_i^2}{\xi_i^2}\right) + \frac{|r_i|^2}{\sigma_i^2} - \frac{|r_i|^2}{\xi_i^2} > \gamma. \quad (2)$$

From a single observation component  $r_i$ , the most likely alternative is  $r_i \sim \mathcal{CN}(0, |r_i|^2)$ .

Substituting  $\xi_i^2 = |r_i|^2$  into the log-likelihood ratio above yields

$$\sum_{i=1}^N \ln \left( \frac{\sigma_i^2}{|r_i|^2} \right) + \frac{|r_i|^2}{\sigma_i^2} - 1 > \gamma. \quad (3)$$

The left side of this inequality is the sum of relative entropies between the distributions  $\mathcal{CN}(0, \sigma_i^2)$  and  $\mathcal{CN}(0, |r_i|^2)$ . This relative entropy is defined in terms of the natural logarithm and so has units of nats.

### 2.3. Hypothesis Testing Algorithms

It is pointed out in [11] that in experiments with SAR data the likelihood function  $p_{\mathbf{R}|A, \Psi}$  is often very sharply peaked in the vicinity of the true target and pose. In this case, the generalized likelihood ratio test (GLRT) for target classification and maximum *a posteriori* (MAP) estimates for pose yield results nearly identical to the minimum probability of error classifier and minimum mean-square error estimator, respectively. The algorithm yielding the GLRT inferred target  $\hat{a}_{GLRT}$  and MAP estimate for pose  $\hat{\psi}_{MAP}$  can be represented in a single equation as

$$\begin{bmatrix} \hat{a}_{GLRT} \\ \hat{\psi}_{MAP} \\ \hat{c}^2 \end{bmatrix} = \underset{[a, \psi, c^2]^T}{\operatorname{argmax}} p(\mathbf{r}|a, \psi, c^2) p(\psi|a) P(a). \quad (4)$$

Given that  $\mathbf{r}$  is hypothesized to have come from a target  $a$  at pose  $\psi$ , the maximum-likelihood estimator for  $c^2$  is

$$\hat{c}^2(a, \psi) = \frac{1}{N} \sum_i \frac{|r_i|^2}{\sigma_i^2(a, \psi)}, \quad (5)$$

where  $N$  is the number of components in  $\mathbf{r}$ .

These inference methods assume that all components of the observation vector  $\mathbf{R}$  are relevant to the recognition problem. In practice, however, the components of the observation vector that are germane may vary with the hypothesis considered. This is the case if a model dependent segmentation is employed. The set of pixels corresponding to target regions in the scene clearly varies with target class  $a$ , due to varying target sizes and shapes, and with pose  $\psi$ , due to target location in the image, etc. Model segmentation is very closely related to the subject of dimensionality reduction in which one seeks a lower dimension subset of statistics to characterize an observation vector for use in a recognition system. Dimensionality reduction can result in lower error rates in cases where the class statistics are estimated from a finite set of training data. As the clutter distribution is estimated from training samples, its inclusion can drastically hinder recognition accuracy and we seek a method of reducing the dimensionality of the estimated statistics to reduce the impact of inaccurately known clutter.

Schmid and O'Sullivan [16] consider an approach to dimensionality reduction based on hard thresholding of the information of each measurement component relative to an assumed "null-hypothesis." This null-hypothesis can be interpreted as representing the distribution of an uninformative component, and selection is performed based upon whether the empirical relative entropy between the component value and the null-hypothesis exceeds some threshold. Components of an observation vector which are used in computing likelihoods varies with target class. To accommodate this variation, likelihoods of the dimensionality reduced statistics under each class are computed relative to likelihoods under the null-hypothesis. They derive asymptotic analytical bounds on the probability of error in a two-class case and demonstrate lower error bounds than without segmentation for a two-class problem with simulated high resolution radar profiles.

DeVore, Schmid, and O'Sullivan [4] apply this method to ATR from SAR imagery and show that as few as 250 components from each of four target classes can be retained while decreasing the classification error rate by several percentage points. For individual target classes, the error rate was drastically reduced for classes with small training sets and was not increased even for those with large training sets. Their dimensionality reduction method is equivalent to the hypothesis testing problem in (3). We apply this method to target model segmentation and select a varying number of components from each class rather than some fixed-small number. Details and examples of this segmentation are provided in Section 3.2. Here we discuss appropriate modification of the GLRT to accommodate class dependent component selection.

Application of their method requires a definition of the null-hypothesis which we take to be a complex Gaussian distribution with zero mean and variance equal to that of the average background clutter,  $\xi^2$ . Let  $\mathbf{I}(a, \psi)$  be a binary mask with the value one in each pixel location hypothesized to be on target when object  $a$  is imaged with pose  $\psi$  and zero in all other pixel locations. The modified GLRT becomes

$$\begin{bmatrix} \hat{a}_{GLRT} \\ \hat{\psi}_{MAP} \\ \hat{c}^2 \end{bmatrix} = \operatorname{argmax}_{[a, \psi, c^2]^T} p(\psi|a)P(a) \prod_{i: I_i(a, \psi)=1} \frac{p(r_i|a, \psi)}{p_{R|\xi^2}(\mathbf{r}|\xi^2)}. \quad (6)$$

Assuming uniform distribution on  $a$  and  $\psi$ , substituting the complex Gaussian probability distribution in (1), and noting that maximizing a positive function is equivalent to maximizing its logarithm yields

$$\begin{aligned} \begin{bmatrix} \hat{a}_{GLRT} \\ \hat{\psi}_{MAP} \end{bmatrix} &= \operatorname{argmax}_{[a, \psi]^T} \sum_i \left[ -\ln \hat{c}^2(a, \psi) \sigma_i^2(a, \psi) - \frac{|r_i|^2}{\hat{c}^2(a, \psi) \sigma_i^2(a, \psi)} + \ln \xi^2 + \frac{|r_i|^2}{\xi^2} \right] I_i(a, \psi) \\ &= \operatorname{argmax}_{[a, \psi]^T} \left\{ -N(a, \psi) \ln \hat{c}^2(a, \psi) + \sum_i \left[ -\ln \sigma_i^2(a, \psi) - 1 + \ln \xi^2 + \frac{|r_i|^2}{\xi^2} \right] I_i(a, \psi) \right\}, \end{aligned} \quad (7)$$

where

$$\hat{c}^2(a, \psi) = \frac{1}{N(a, \psi)} \sum_i \frac{|r_i|^2}{\sigma_i^2(a, \psi)} I_i(a, \psi) \quad (8)$$

and  $N(a, \psi)$  is the number of nonzero elements of the segmentation mask  $\mathbf{I}(a, \psi)$ .

#### 2.4. Confuser Rejection

A common problem in ATR systems concerns the ability to properly note those received images that contain objects not represented in the model database, so called confuser targets, rather than misclassifying them. Such images may be formed from “actual” targets that have not been modeled and included in the database or they may be of clutter that appears target-like. One approach to this problem of confuser rejection can be found in classical statistical hypothesis testing. Denote the most likely combination of target class and pose determined from the previous algorithms as the null-hypothesis,  $\mathcal{H}_0$ , which represents an assertion that  $\mathbf{r}$  is an image of target class  $a_0$  with pose  $\psi_0$ . That is,  $\mathcal{H}_0$  represents an assertion that  $\text{var}(r_i) = \hat{c}\sigma_i^2(a_0, \psi_0)$ , for  $i = 1, 2, \dots, N$ . Denote the alternate hypothesis by  $\mathcal{H}_A$  which represents an assertion that  $\mathbf{r}$  is not an image of  $a_0$  with pose  $\psi_0$ . That is,  $\mathcal{H}_A$  represents an assertion that  $\text{var}(r_i) \neq \hat{c}\sigma_i^2(a_0, \psi_0)$ , for at least one  $i = 1, 2, \dots, N$ . As discussed in Section 2.2, we conduct a likelihood ratio test against the most likely alternative. From (3), we reject the image  $\mathbf{r}$  as a confuser if the relative entropy between these two distributions exceeds some threshold  $\gamma$  as

$$\begin{aligned} D(p(\cdot | \text{var}(r_i) = |r_i|^2) \| p(\cdot | \text{var}(r_i) = \hat{c}^2 \sigma_i^2(\hat{a}, \hat{\psi}))) \\ = \sum_i \left[ \ln \left( \frac{\hat{c}^2 \sigma_i^2(\hat{a}, \hat{\psi})}{|r_i|^2} \right) + \frac{|r_i|^2}{\hat{c}^2 \sigma_i^2(\hat{a}, \hat{\psi})} - 1 \right] > \gamma. \end{aligned} \quad (9)$$

### 3. Model Building

For our purposes, a target model consists of both a mathematical description and a collection of parameters sufficient to characterize the output of the sensor platform given some target and pose. As mentioned in the introduction, some approaches to ATR rely on CAD representations of targets which, combined with simulation of the sensor’s physical properties, allow the prediction of sensor outputs. An alternative approach of constructing target models from training data is described in this section. We first describe how the probabilistic description of sensor output from Section 2.1 can be combined with training data to estimate the radar reflection statistics over the surface of a target. Through this approach, model training becomes an issue of parameter estimation. Following this is a

discussion of how these estimates can be used to characterize the sensor output for an arbitrary pose. Finally, we discuss segmentation of the estimated parameters into target and clutter regions.

### 3.1. Model Parameter Estimation

Given a probabilistic sensor description and a collection of training data for a target representing a wide variety of target pose, we seek estimates of the contribution to sensor output statistics from all locations on the target surface. For a given target  $a$ , the region  $z$  of the target surface that corresponds to pixel  $i$  is dependent upon the pose  $\psi$  and we can represent this dependence as a function  $z_i(\psi)$ . The region  $z$  corresponding to one pixel at a given pose will contribute to the variance of different pixels when viewed at some other pose. The model of a target  $a$  will consist of the model for SAR imagery in (1) and estimates of the variance function  $\zeta_z^2(a, \psi)$  for the resolution cells  $z_i(\psi)$  on the target. SAR imaging is inherently insensitive to translations along the radar line of sight, and translations perpendicular to the line of sight can be well modeled by simple shifts of the resulting image so that  $\zeta_z^2$  need not incorporate variation due to translation.

The variance  $\zeta^2$  is a function of the region of the target surface being imaged and its pose relative to the radar platform. Because we are considering storage and computation on digital computer systems, we must restrict the class of variance functions to those with finite representations. In practice, these functions will be estimated from a finite collection of training data so representations that strike a reasonable balance between approximation error in representation and estimation error in the function parameters is desirable. Several methods of regularization for function estimation have been reported. For an overview of many of these methods in an image formation context, see O'Sullivan, et al. [10]. Here, we consider a restriction of  $\zeta^2$  to functions that are piecewise constant over both the scene and pose. The piecewise constant regions in the scene will be the  $z_i(\psi_0)$  corresponding to the pixels  $i$  in an image of the target at some nominal pose  $\psi_0$ . The piecewise constant regions in pose will follow from a uniform discretization of the components of  $\psi$ .

Suppose we collect a set of  $M$  power-normalized training images  $\{\mathbf{T}'_1, \mathbf{T}'_2, \dots, \mathbf{T}'_M\}$  from some fixed depression angle with position vectors  $\{\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_M\}$  in the slant plane and with azimuth angles  $\{\phi_1, \phi_2, \dots, \phi_M\}$ . Neglecting articulated components, the transformation  $U(-\phi_j, -\mathbf{s}_j) \in \mathbf{SE}(2)$  maps points  $(r, c)$  in image  $\mathbf{T}'_j$  into points  $(x, y)$  in the slant plane. The image  $\mathbf{T}_j = U(-\phi_j, -\mathbf{s}_j) [\mathbf{T}'_j]$  is a representation of the radar return from the scene registered with nominal pose  $\psi_0 = [\phi_0, \mathbf{s}_0]^T = [0, \mathbf{0}]^T$  and the pixels in this image correspond to the regions in the scene over which the variance function will be approximated as piecewise constant. The collection of all such transformed data,  $T = \{\mathbf{T}_1, \mathbf{T}_2, \dots, \mathbf{T}_M\}$ , constitutes a registered set of training data from which the variance function  $\zeta^2$  can be estimated.

For a fixed target position  $\mathbf{s} = \mathbf{0}$ , the variance  $\zeta^2$  is a function of slant plane location and azimuth angle. Let the  $2\pi$  radians of target azimuth be discretized into  $N_w$  uniform intervals with the  $k$ th interval centered at  $\theta_k = 2\pi/N_w$ . An estimate of the variance function  $\zeta_i^2(a, \theta_k)$  can be determined from the registered training images within some training

interval of width  $d$  as

$$\hat{\zeta}_i^2(a, \theta_k) = \frac{1}{|\mathcal{I}_k|} \sum_{\mathbf{r} \in \mathcal{I}_k} |r_i|^2, \quad (10)$$

where  $\mathcal{I}_k = \{\mathbf{T}_j \in \mathcal{T} : \phi_j \in [\theta_k - d/2, \theta_k + d/2]\}$ . It is possible to define  $d > 2\pi/N_w$ , so that training images formed from azimuth angles near the edges of the training intervals contribute to both neighboring intervals. The left panel of Figure 1 shows a SAR magnitude image in a logarithmic intensity scale of target class  $a = \text{T-72 tank}$  with azimuth angle  $\phi = 135^\circ$ . The middle panel shows a registered image formed by a nearest-neighbor rotation of the data. The right panel shows an image of the variance estimate  $\hat{\zeta}^2(\text{T-72}, 135^\circ)$  formed from training data with  $N_w = 64$  windows and with training interval width  $d = 11.25^\circ$ .

### 3.2. Target Model Segmentation

The estimation method from the previous section yields estimated variances for regions in the scene as a function of orientation relative to the radar platform. Although the corresponding pixels vary with orientation, a region in the scene that corresponds to the surface of a target remains on the target regardless of the angle from which it is viewed. Segmentation of target regions from clutter in this target-centered reference frame is a problem of target model segmentation. This is in contrast with the segmentation of SAR imagery as in [6] and [18].

Segmentation of the model can be phrased as a hypothesis testing problem which pits a hypothesis  $\mathcal{H}_0$ , an assertion that a given region is clutter, against the alternate  $\mathcal{H}_A$ , an assertion that the region is not clutter. The clutter hypothesis is defined by a zero-mean complex Gaussian distribution with variance  $\xi^2$ . As in Section 2.2, we conduct a test against the most likely alternative, which is that the pixel comes from a complex Gaussian




Figure 1. SAR magnitude image on a logarithmic scale of a target T-72 from the MSTAR dataset with azimuth angle  $135^\circ$ , the same image registered in pose, and estimated variances determined from registered images with azimuth near  $135^\circ$ .




Figure 2. Average relative entropy over all orientations between the distributions parameterized by  $\zeta_i^2$  and  $\xi^2 = 0.0028$  for the T-72 tank. The left panel shows an image of the relative entropy values on a logarithmic intensity scale. The right panel shows the histogram of average relative entropy values. The sharp peak at the value 2 indicates that approximately 800 scene regions have average relative entropies greater than or equal to 2.

distribution with conditional variance  $\hat{\zeta}_i^2(a, \theta)$ ). For each target  $a$ , we reject the clutter hypothesis if the average over all orientations of the relative entropy between the distribution in (1) parameterized by  $\hat{\zeta}^2$  and that parameterized by the clutter level  $\xi^2$  surpasses some threshold  $\eta$ . The scene region  $i$  is taken to be on the target if

$$E_{\theta} [D(p(\cdot|\sigma^2 = \hat{\zeta}_i^2(a, \theta)) \| p(\cdot|\sigma^2 = \xi^2))] = \frac{1}{N_w} \sum_k \left( \frac{\hat{\zeta}_i^2(a, \theta_k)}{\xi^2} + \ln \frac{\xi^2}{\hat{\zeta}_i^2(a, \theta_k)} - 1 \right) > \eta. \quad (11)$$

We denote the target model segmentation by a binary mask  $\mathbf{S}(a)$  which has  $S_i(a) = 1$  if region  $i$  satisfies the above inequality and zero otherwise.

The left panel of Figure 2 shows the average relative entropy on a logarithmic intensity scale for the T-72 tank assuming a background clutter level of  $\xi^2 = 0.0028$  (determined experimentally from the MSTAR dataset) and with variance estimates described in the example of the previous section. In some SAR images used in estimation, the tank turret was slightly skewed resulting in the appearance of two barrels at a slight angle. This is an example of the model inaccuracies that can arise if the data are not registered in articulation as well as pose. The gray inner ring surrounding the pixels on target correspond to the regions of the scene masked by the target shadow at varying orientations and the grey-speckled outer ring corresponds to background clutter. The right panel shows the histogram of the average relative entropy values. The isolated peak at the value two indicates that approximately 800 resolution cells have an average relative entropy of two or greater.

Segmentation is a function of the threshold  $\eta$  and larger values result in smaller numbers of regions being included in the segmentation. Figure 3 shows segmented target models of the T-72 tank on a logarithmic intensity scale for several decreasing values of  $\eta$  from 125




Figure 3. Model of target T72 trained over 64 windows of width  $11.25^\circ$  at various segmentations resulting from different threshold levels. The segmentation threshold level and the number of resolution cells with average relative entropies surpassing the threshold are given beneath each image.

to zero nats assuming a background clutter level of  $\xi^2 = 0.0028$ . Regions of the scene that have no corresponding SAR image pixels for some azimuth angles are explicitly chosen to represent clutter regions resulting in the circular shape in the segmentation for  $\eta = 0$ .

### 3.3. Sensor Output Characterization

The segmented target models derived in the previous sections represent the variability in the radar return for regions on the target as a function of azimuth angle in a coordinate system centered at the target. The output of a SAR radar platform is in a sensor-centric coordinate system, and to form a statistical characterization of the resulting data we need a transformation of the target model. SAR imaging is inherently insensitive to translations in the cross range direction and translations in the range direction affect only the power in the received data which is accommodated by the estimate  $\hat{\rho}^2$  in (8). However, the radar return from a scene is sensitive to the orientation of the target relative to the radar platform. Thus, to estimate the variances of pixel values in an image of target  $a$  with pose consisting of location  $\mathbf{s}$  and azimuth angle  $\theta$ , we find the  $\theta_k = 2\pi k/N_w$  closest to  $\theta$  and transform the target model into the sensor coordinate system. That is, the image variance estimate is  $\hat{\sigma}^2(a, \psi) = U(a, \psi)[\hat{\zeta}^2(a, \theta_k)]$  and  $\mathbf{I}(a, \psi) = U(a, \psi)[S(a)]$  is a binary pixel mask in the image. Note that while  $\hat{\zeta}^2$  is a piecewise constant function, the transformation  $U$  is a function of continuous  $\psi$  so  $\hat{\sigma}^2$  is also a function of continuous-valued pose.




Figure 4. Left panel shows a segmented target model for a T-72 tank illuminated at  $135^\circ$  azimuth. Right panel shows the corresponding sensor output model.

Figure 4 shows a target model on a logarithmic scale, segmented with  $\eta = 1$  for a T-72 tank when illuminated from  $135^\circ$  azimuth and the corresponding model for sensor output. The sensor output model is a characterization of the SAR image in the left panel of Figure 1.

### 3.4. Performance Sensitivity

We view the tradeoff between performance and complexity as the fundamental criterion for comparing alternate ATR algorithms. Each implementation delivers some level of accuracy (i.e. correct classification rate or average pose estimation error) and entails some complexity (i.e. storage or processing requirements). ATR is a difficult problem with many interrelated parameters governing the final system. In evaluating a general approach to ATR, we seek to decouple the effects of these parameters in order to understand their relative importance. DeVore, et al. [3] consider ATR performance as a function of resolution in pose space. In this section, we consider issues related to variation in segmentation.

The segmentation threshold  $\eta$  represents a free parameter in the design of an ATR system as described in the previous sections. Figure 3 indicates that for the T-72 data, a threshold value of 1 nat nicely separates target from clutter pixels. In the examples of the next section we consider the ATR accuracy as a function of  $\eta$ . This evaluation is used to determine whether this same threshold works well across a wide range of target classes and sheds light on the sensitivity of ATR to the segmentation threshold. We are also interested the encoding of estimated sensor modules  $\hat{\sigma}^2(a, \psi)$  for rapid communication from the model database to the ATR processor and successive computation of the object of the argmax function in (7). The ability to refine initial classifications of SAR image content is potentially useful as operational demands on the ATR system change in response to the

complexity or level of threat in the environment. Also, it offers the potential for restricting the parameter space of the search based on initial approximate calculations of the most likely target classes.

The number of nonzero pixels in the segmentation mask increases with decreasing  $\eta$ . By communicating from the database the pixel locations and values with average relative entropy greater than some  $\eta_1$  for all possible hypotheses  $(a, \psi)$  and computing the log likelihood ratios in (7), an initial classification and pose estimation can be made. If  $\eta_1$  is large so that relatively few pixels are included on average, the classification can be made quickly but with increased error rate. The initial classification can be refined by communicating only the locations and values of the pixels with average relative entropy between  $\eta_2$  and  $\eta_1$  for all possible hypotheses. If the subexpressions in the object of the argmax function are retained for each hypothesis and if an average of  $N_2$  pixels are added by decreasing the entropy threshold, the estimate  $\hat{c}^2(a, \psi)$  in (8) can be refined with  $N_2$  divide and accumulate operations and a single division. The initial classification can be refined by computing the logarithm of the updated  $\hat{c}^2$ , a multiplication by the new  $N(a, \psi)$ , and a logarithm and divide accumulate operation for each of the  $N_2$  additional pixels. This process can be continued until the all components of the sensor output models have been communicated from the database and a final classification has been reached. The additional computational cost in supporting the successive refinement from  $\eta_1$  to  $\eta_2$  rather than determining a classification for  $\eta_2$  directly is one logarithm, multiply, divide, and subtract operation for each hypothesis. This is the same additional cost in going from  $\eta_n$  to  $\eta_{n+1} < \eta_n$ . The additional communication cost for supporting successive-refinement depends on the representation of pixel locations and values used.

#### 4. Example

In this section we present examples of the previously described methods for actual SAR data from the MSTAR dataset. First, we present experimental results for a ten-class recognition problem with forced decisions and no confuser targets. We examine classification error rates and average pose estimation error as a function of the segmentation threshold and relate them to the average number of bits transmitted between the model database and ATR processor and to the average number of floating-point operations to compute the object of the argmax function in (7). We then divide the dataset into a four-class problem with four confuser targets and examine receiver operating characteristics, probability of correct classification versus confuser rejection rate, as a function of segmentation threshold level. Finally, we consider two different four-class problems with confuser targets which represent extended operating conditions (EOC) [8], [14], in which the test images are of targets with some configuration or version variation from the training images. In each example, the target models  $\hat{\zeta}^2(a, \theta_k)$  were estimated from the training data with  $N_w = 64^\circ$  and  $d = 11.25^\circ$ . From these, the sensor output models  $\hat{\sigma}^2(a, \phi)$ , where  $\phi = [\theta, \mathbf{s}]^T$ , were generated with 256 equally spaced azimuth angles  $\theta$  and 25 target locations  $\mathbf{s}$ , corresponding to a target location unknown within 5 pixels in both range and cross-range directions.

Table 1. MSTAR dataset used in ten-class examples.

Target	Train			Test		
	Vehicles	Images	Depression	Vehicles	Images	Depression
2S1	b01	299	17°	b01	274	15°
BMP-2	9563, 9566, c21	697	17°	9563, 9566, c21	587	15°
BRDM-2	E-71	298	17°	E-71	263	15°
BTR-60	k10yt7532	256	17°	k10yt7532	195	15°
BTR-70	c71	233	17°	c71	196	15°
D7	92v13015	299	17°	92v13015	274	15°
T62	A51	299	17°	A51	273	15°
T-72	132, 812, s7	691	17°	132, 812, s7	582	15°
ZIL131	E12	299	17°	E12	274	15°
ZSU 23 4	d08	299	17°	d08	274	15°

#### 4.1. Ten Class Recognition

In this example, SAR images of ten target classes from the MSTAR dataset were partitioned into two nonoverlapping sets, one for estimation of model parameters through the methods of Section 3 and the other for evaluation through the methods of Section 2.3. Details about the dataset used for each are given in Table 1. Note that images formed with a 17° depression angle were used for estimation and those from a 15° depression angle were used for testing. Also, the target classes T-72 and BMP-2 incorporate images from three different vehicles of the same class.

Performance results for this example are shown in Figure 5 as a function of decreasing segmentation threshold,  $\eta$ , logarithmically spaced from a value of 125 nats to 1 nat. The left




Figure 5. Error rates as a function of decreasing segmentation threshold. Left panel shows the average squared Hilbert-Schmidt distance between actual and estimated target azimuth and the equivalent error in degrees. Right panel shows the percentage of misclassified test images.

graph shows the average error between estimated target azimuth,  $\hat{\theta}$ , and actual target azimuth,  $\theta$ , given that the true target class is known. The left vertical axis gives the error in terms of the squared Hilbert-Schmidt distance,  $d_{HS}^2$ , defined by Grenander, et al. [5] as the squared Hilbert-Schmidt norm of the difference between the rotation matrices,  $O(\hat{\theta}), O(\theta) \in \mathbf{SO}(\mathbf{N})$ , that describe the orientation of the target. That is, the value of  $d_{HS}^2(O(\hat{\theta}), O(\theta))$  is given by

$$d_{HS}^2(O_1, O_2) = \text{tr}[(O_1 - O_2)^T(O_1 - O_2)], \tag{12}$$

where  $\text{tr}[\cdot]$  denotes the matrix trace operation and  $(\cdot)^T$  denotes matrix transpose. The right vertical axis gives the equivalent error in degrees determined as  $\cos^{-1}(1 - d_{HS}^2/4)$ . The graph shows that the average orientation estimation error drops from over  $75^\circ$  with large segmentation thresholds down to  $9.36^\circ$  for  $\eta = 1$ . The right graph shows the percentage of test images incorrectly classified as a function of  $\eta$ . Error rates drop from 88.1% down to 2.26% over the same range of segmentation threshold.

The amount of computational resources consumed in recognizing an object in a SAR image is also related to the segmentation threshold as indicated in Figure 6. The left panel of the figure shows the average number of pixels satisfying the threshold inequality (11) for varying values of  $\eta$ . The right vertical axis of the same graph shows the average number of bits that must be transmitted between the model database and ATR processor for each  $\hat{\sigma}^2(a, \phi)$ . This value assumes that each pixel value is represented by 14 bits denoting the pixel location (sufficient for up to a  $128 \times 128$  image) and 64 bits for the floating-point representation of the variance value. For this example, there were 10 target classes and 256 azimuth angles under consideration so the total number of bits transmitted to recognize a single target image is 2560 times this value. Since variation in target location  $\mathbf{s}$  can be handled by simple image shifts in the ATR processor, no additional




Figure 6. Consumed resources as a function of decreasing segmentation threshold. Left panel shows the average number of pixels in a segmented target model and an upper bound on the number of bits to communicate the models. Right panel shows the average number of floating-point operations to compute the object of the argmax function for each combination of target class and pose assuming  $\eta$  is reduced in four steps from 125 to 1.

communication is required to incorporate the 25 target locations considered in this example.

The right panel of Figure 6 shows the average number of operations required to compute the object of the argmax function in (7), assuming that  $\eta$  is initially set to 125 nats and is successively lowered to 25, 5, and finally 1 nat, with a classification decision being reached at each level. Assuming one operation for each divide and accumulate operation and one for the division, estimation of  $c^2$  will take  $N + 1$  operations on average, where  $N$  is the average number of pixels on target from the graph in the left panel. The remaining terms in the object of the argmax function can be computed in  $N + 13$  operations assuming one operation for each divide and accumulate operation involving  $\xi^2$ , one addition for the (precomputed) term involving  $\ln \sigma^2$ , 10 operations for logarithm computation, one for multiplication by  $N$ , and a final addition to join the terms. This is a total of  $2N + 14$  operations. An additional cost of 13 operations is incurred for each successive refinement of  $\eta$  as discussed in Section 3.4, assuming 10 operations for recomputing the logarithm and one each for the redundant multiplication, division, and addition operations. For this example, the number of operations indicated in this graph must be carried out for each of the 10 target classes, 256 azimuth angles, and 25 target locations, totaling 64,000 times, in the classification of a single target image.

The four graphs in these two figures demonstrate that large segmentation thresholds result in high recognition errors but can be very easy to compute. Clearly, recognition error rates near 80% are not going to be acceptable in practice but the error rates achievable with thresholds less than 5 nats may be, particularly if the decision must be made extremely quickly or with limited hardware capability. Also, it may be useful for a recognition system to deliver an initial classification quickly but suffering a relatively high error rate with the understanding that a more accurate decision is pending more exhaustive computation. Finally, it may be possible to exploit the recognition results with large segmentation threshold values by eliminating from consideration highly unlikely combinations of target class and pose while periodically lowering  $\eta$ .

#### 4.2. *Four Class with Confuser Rejection*

For the experiments in this section, the MSTAR dataset was partitioned into three nonoverlapping sets consisting of training, testing, and confuser images. Table 2 details the four target classes used for estimating target models. These are a subset of the ten-class training set from the previous section. Table 3 shows the testing and confuser data. The testing data consists of images of the same vehicles used in training but taken from a different depression angle. The confuser data consists of images of four vehicles not used in training. In this example, the ability to correctly classify the testing data while rejecting the confuser data is investigated.

Figure 7 shows the percentage of correctly classified testing data and percentage of confuser targets that were not rejected, for each of four segmentation threshold levels,  $\eta$ . The curves are parameterized by the confuser rejection threshold  $\gamma$  from (9). A test image is said to be correctly classified if the recognition algorithm (7) yields the correct target class  $a$  and the image is not rejected as a confuser. The rightmost point of each curve

Table 2. MSTAR dataset used for training target models in the confuser rejection examples.

Class	Vehicle	Serial No.	Depression	Images
BMP-2	#1	9563	17°	233
	#2	9566		231
	#3	c21		233
BRDM-2	#1	E-71	17°	298
BTR-70	#1	c71	17°	233
T-72	#1	132	17°	232
	#2	812		231
	#3	s7		228

shows the best correct recognition rate achievable for each segmentation threshold if confuser rejection was not of concern. With  $\eta = 1$ , a correct recognition rate of 87.7% and 6.4% confuser rejection rate are achievable with  $\gamma = 1.1$  nat. With  $\eta = 5$ , a correct recognition rate of 83.3% and 7.9% confuser rejection rate are achievable with  $\gamma = 1.1$  nat.

#### 4.3. Extended Operation Conditions

For the experiments in this section, the MSTAR dataset was partitioned into four nonoverlapping sets to investigate the ability to correctly recognize variants of modeled target classes while rejecting images of confuser targets. The data used for training are detailed in Table 4. Four target classes are represented, each from a single vehicle and incorporating two slightly different depression angles. Two sets of test data were constructed from the data in such a way that no vehicle used for training has images in

Table 3. MSTAR dataset used for evaluating confuser rejection capabilities. Target classes listed above the double line are represented in the target model database. Those below are confuser targets for this experiment.

Class	Vehicle	Serial No.	Depression	Images
BMP-2	#1	9563	15°	195
	#2	9566		196
	#3	c21		196
BRDM-2	#1	E-71	15°	263
BTR-70	#1	c71	15°	196
T-72	#1	132	15°	196
	#2	812		195
	#3	s7		191
2S1	#1	b01	15°	274
BTR-60	#1	k10yt7532	15°	195
D7	#1	92v13015	15°	274
T-62	#1	A51	15°	273




Figure 7. Correct classification rate as a function of confuser acceptance rate for each of four segmentation thresholds.

either test set and these are detailed in Table 5. The configuration variant test set consists of images of vehicles that represent minor differences in vehicle configuration from those in the training set. The version variant set consists of images of vehicles that exhibit more significant version differences from the training set. The same confuser images from the previous section were used in these experiments.

Figure 8 shows the correct classification rate and rate of confuser acceptance for both experiments. The left panel shows the results for test images that differ in configuration from the training data and the right panel shows the results for more significant version variations. The ability to simultaneously classify test images and reject confusers is

Table 4. MSTAR images used for training target models in tests under extended operation conditions.

Model	Vehicle	Serial No.	Depression	Images
BMP-2	#3	c21	15° and 17°	429
BTR-70	#1	c71	15° and 17°	429
BRDM-2	#1	E-71	15° and 17°	561
T-72	#1	132	15° and 17°	428

Table 5. MSTAR images used in the configuration variation and version variation EOC testing.

Configuration Variant Testing Set				Version Variant Testing Set			
Target	Serial No.	Depression	Images	Target	Serial No.	Depression	Images
T-72 #2	812	15° and 17°	426	T-72 #3	s7	15° and 17°	419
T-72 #4	A04	15° and 17°	573	T-72 #8	A32	15° and 17°	572
T-72 #5	A05	15° and 17°	573	T-72 #9	A62	15° and 17°	573
T-72 #6	A07	15° and 17°	573	T-72 #10	A63	15° and 17°	573
T-72 #7	A10	15° and 17°	567	T-72 #11	A64	15° and 17°	573
BMP-2 #1	9563	15° and 17°	428				
BMP-2 #2	9566	15° and 17°	427				

hindered by test images that vary from the training data as compared to the experiments of the previous section. In both the configuration and version variant tests, the target segmentation threshold  $\eta = 5$  yielded better performance than  $\eta = 1$  over most of the graphs. In both configuration and version variant experiments, the segmentation threshold  $\eta = 125$  results in more than 20% of the confuser images being closer to training data, in the sense of (9), than any of the test images.

5. Conclusions

We have presented an approach to ATR from SAR imagery in which statistical models are estimated in a target-rather than sensor-centered reference frame. These models can be used to approximate the statistics of sensor output data at arbitrary pose relative to the scene. The algorithms account for rotations and for translations both perpendicular to and parallel to the radar line of sight. A method of addressing two-sided hypothesis tests which

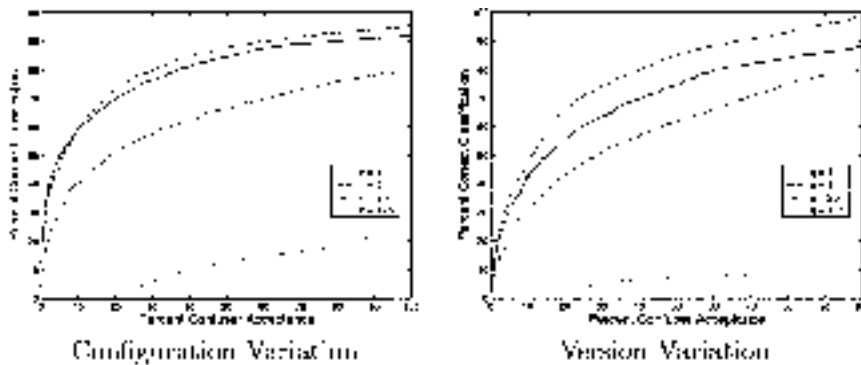


Figure 8. Correct classification rate for configuration variants and version variants as a function of confuser acceptance rate for each of four segmentation thresholds.

is maximally powerful at the most likely alternative is described and applied to both target model segmentation and confuser rejection. This hypothesis testing method yields a decision based upon the empirical relative entropy between the most likely distribution given the observed data and the distribution of the null-hypothesis.

These algorithms were demonstrated on actual SAR data in a variety of experiments designed to indicate the relative sensitivity to variation in segmentation threshold, confuser rejection threshold, and operating conditions. The experiments indicate that a segmentation threshold of 1 nat for the relative entropy between the most likely distribution and the clutter distribution yields consistently low error rates at a complexity of under 800 pixels per target and 1600 floating-point operations per hypothesized combinations of target class and orientation. This threshold yields a classification error rate of under 3% for a ten class problem under standard operating conditions and a Pcc-Pfa curve exhibiting a sharp knee in a four-class confuser rejection problem.

### *Acknowledgments*

This work was supported in part by the US Army Research Office grant DAAH04-95-1-0494, by the Office of Naval Research grant N00014-98-1-06-06, and by the Boeing Foundation.

### **Note**

1. The MSTAR dataset is publicly available and can be requested through the Sensor Data Management System web site at the URL <http://www.mbvlab.wpafb.af.mil/public/sdms/>.

### **References**

1. H.-C. Chiang, R.L. Moses, and L.C. Potter, "Model-Based Classification of Radar Images," *IEEE Transactions on Information Theory*, vol. 46, no. 5, Aug. 2000, pp. 1842–1854.
2. M.D. DeVore, A.D. Lanterman, and J.A. O'Sullivan, "ATR Performance of a Rician Model for SAR Images," in F.A. Sadjadi (Ed.), *Automatic Target Recognition X, Proc. of SPIE*, vol. 4050, 2000.
3. M.D. DeVore, J.A. O'Sullivan, S. Anand, and N.A. Schmid, "Probabilistic Approach to Model Extraction from Training Data," in E.G. Zelnio (Ed.), *Algorithms for Synthetic Aperture Radar Imagery VIII, Proc. of SPIE*, 2001.
4. M.D. DeVore, N.A. Schmid, and J.A. O'Sullivan, "Analytical and Experimental Performance-Complexity Tradeoffs in ATR," in *Proceedings of the Thirty-Fourth Asilomar Conference on Signals, Systems, and Computers*, Oct. 2000.
5. U. Grenander, M.I. Miller, and A. Srivastava, "Hilbert-Schmidt Lower Bounds for Estimators on Matrix Lie Groups for ATR," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 20, no. 8, Aug. 1998, pp. 790–802.
6. W.W. Irving, L.M. Novak, and A.S. Willsky, "Multiresolution Approach to Discriminating Targets from Clutter in SAR Imagery," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 33, no. 4, Oct. 1997, pp. 1157–1169.

7. S.P. Jacobs, Automatic Target Recognition Using High Resolution Radar Range-Profiles. PhD thesis, Washington University, St. Louis, Missouri, May 1997.
8. E.R. Keydel and S.W. Lee, "Signature Prediction for Model-Based Automatic Target Recognition," in E.G. Zelnio and R.J. Douglass (Eds.), *Algorithms for Synthetic Aperture Radar Imagery III, Proc. of SPIE*, vol. 2757, 1996, pp. 306–317.
9. E.R. Keydel, S.W. Lee, and J.T. Moore, "MSTAR Extended Operating Conditions: A Tutorial," in E.G. Zelnio and R.J. Douglass (Eds.), *Algorithms for Synthetic Aperture Radar Imagery III, Proc. of SPIE*, vol. 2757, 1996, pp. 228–242.
10. J.A. O'Sullivan, R.E. Blahut, and D.L. Snyder, "Information-Theoretic Image Formation," *IEEE Transactions on Information Theory*, vol. 44, no. 6, Oct. 1998, pp. 2094–2123.
11. J.A. O'Sullivan, M.D. DeVore, V. Kedia, and M.I. Miller, "Automatic Target Recognition Performance for SAR Imagery Using a Conditionally Gaussian Model," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 37, no. 1, Jan. 2001, pp. 91–108.
12. G.J. Owirka, S.M. Verbout, and L.M. Novak, "Template-Based SAR ATR Performance Using Different Image Enhancement Techniques," in *Algorithms for Synthetic Aperture Radar Imagery VI, Proc. of SPIE*, vol. 3721, 1999, pp. 302–319.
13. J.A. Richards, J.W. Fisher III, and A.S. Willsky, "An Expectation-Maximization Approach to Target Model Generation from Multiple SAR Images," in E.G. Zelnio (Ed.), *Algorithms for Synthetic Aperture Radar Imagery VII, Proc. of SPIE*, vol. 4053, 2000, pp. 652–664.
14. T.D. Ross, S.W. Worrell, J.C. Mossing, and M.L. Bryant, "Standard SAR ATR Evaluation Experiments Using the MSTAR Public Release Data Set," in E.G. Zelnio (Ed.), *Algorithms for Synthetic Aperture Radar Imagery V, Proc. of SPIE*, vol. 3370, 1998, pp. 566–573.
15. G.G. Roussas, *A Course in Mathematical Statistics*, Academic Press, 2nd ed., 1997.
16. N.A. Schmid and J.A. O'Sullivan, "Thresholding Method for Reduction of Dimensionality," *IEEE Transactions on Information Theory*, vol. 47, no. 7, Nov. 2001, pp. 2903–2920.
17. 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," *Lincoln Laboratory Journal*, vol. 6, no. 1, Spring 1993, pp. 53–76.
18. R.A. Weisenseel, W.C. Karl, D.A. Castañón, G.J. Power, and P. Douville, "Markov Random Field Segmentation Methods for SAR Target Chips," in E.G. Zelnio (Ed.), *Algorithms for Synthetic Aperture Radar Imagery VI, Proc. of SPIE*, vol. 3721, 1999, pp. 462–473.
19. J. Wissinger, R.B. Washburn, N.S. Friedland, A. Nowicki, D.R. Morgan, C. Chong, and R. Fung, "Search Algorithms for Model Based SAR ATR," in E.G. Zelnio and R.J. Douglass (Eds.), *Algorithms for Synthetic Aperture Radar Imagery III, Proc. of SPIE*, vol. 2757, 1996, pp. 279–293.