Target Prescreening Based on a Quadratic Gamma Discriminator

JOSE C. PRINCIPE, Senior Member, IEEE
ALEX RADISAVLJEVIC
JOHN FISHER, III
University of Florida
MARGARITA HIETT
LESLIE M. NOVAK
Lincoln Laboratory

This work presents the development, analysis and validation of a new target discrimination module for synthetic aperture radar (SAR) imagery based on an extension of gamma functions to 2-D. Using the two parameter constant false-alarm rate (CFAR) stencil as a prototype, a new stencil based on 2-D gamma functions is used to estimate the intensity of the pixel under test and its surroundings. A quadratic discriminant function is created from these estimates, which is optimally adapted with least squares in a training set of representative clutter and target chips. This discriminator is called the quadratic gamma discriminator (QGD). The combination of the CFAR and the QGD was tested in realistic SAR environments and the results show a large improvement of the false alarm rate with respect to the two-parameter CFAR, both with high resolution (1 ft) fully polarimetric SAR and with one polarization, 1 m SAR data.

Manuscript received February 24, 1995; revised May 23 and September 3, 1997 and March 21, 1998.

IEEE Log No. T-AES/34/3/06007.

This research was partially supported by DARPA Grants N60921-93-C-A335 and F33615-97-1-1019.

Authors' current addresses: J. C. Principe, Dept. of Electrical and Computer Engineering, University of Florida, 451 Engineering Bldg., P.O. Box 116130, Gainesville, FL 32611-6130; A. Radisavljevic and J. Fisher, III, Computational NeuroEngineering Laboratory, EB-451, Electrical and Computer Engineering, University of Florida, Gainesville, FL 32611; M. Hiett and L. M. Novak, MIT Lincoln Laboratory, P.O. Box 73, Lexington, MA 02173.

I. INTRODUCTION

In a typical target detection/recognition task an automatic target recognition (ATR) system is required to search large areas of clutter containing relatively few targets. Because it is prohibitive to run computationally intensive target recognition algorithms in the entire image it is necessary to divide the ATR task into submodules, commonly a front-end prescreener followed by a classifier [1]. The goal of the front-end prescreener is to identify points of interest which contain potential targets. Functionally, the front-end prescreener is divided into the subtasks of rejecting imagery without potential targets (detection) and screening of false alarms (discrimination) which are mostly created by bright clutter areas (tree tops) and man-made clutter. The two-parameter CFAR (constant false alarm rate) detector has been widely utilized in synthetic aperture radar (SAR) imagery [1]. The goal of the discriminator is to eliminate as many false alarms as possible from the CFAR output.

In this work the local amplitude features used for the CFAR test are further combined to construct a quadratic discriminant function, which is known to be optimal for classification of Gaussian distributed classes. We call this discriminator the quadratic gamma discriminator (QGD) because the local intensity features are extracted with a 2-D extension of the integrands of the gamma function. The QGD utilizes the same local amplitude features to classify the CFAR outputs in the classes of targets and nontargets. Since the CFAR is already implemented in practical ATR systems, the addition of the QGD increases the choices for discriminators with very little extra computational effort.

The analyzing functions for extracting local 1st- and 2nd-order statistics are 2-D radially symmetric gamma functions, a direct extension of their 1-D counterpart [2], which have been successfully applied in various adaptive signal processing problems. The 1-D gamma filter is built around a cascade of recursive delays, whose impulse responses are the gamma bases. The recursive parameter serves as an additional degree of freedom that adapts the scale of the gamma basis to the characteristics of the signal. In addition, this previous work on gamma filters [2] makes adaptation equations readily available.

We start with a brief description of the Lincoln ATR system, followed by a description of the QGD by specifying the 2-D gamma functions. We then address how the feature set used by the prescreener is computed using these functions. Next, the QGD structure is presented and we show how the QGD generalizes the conventional two-parameter CFAR detector. A method of computing detector parameters is described which is essentially a training process based on exemplars from both classes. Finally, we

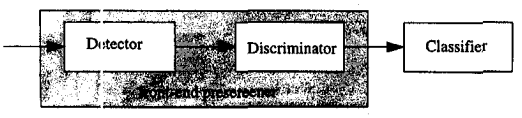

Fig. 1 Main stages of Lincoln ATR system.

show some detection results. Performance tests have been conducted by MIT/Lincoln Laboratory involving a realistic data set that better represents the actual large area detection problem. The results of these experiments in the form of receiver operating characteristic (ROC) curves are also presented here.

II. LINCOLN SAR ATR SYSTEM

A. System Overview

Lincoln Laboratory was one of the pioneers in developing an-end-to-end high resolution ATR system exploiting two-dimensional SAR imagery [1]. The SAR data was collected with a 33 GHz (Ka-band) high resolution (1 ft × 1 ft), fully polarimetric sensor developed by Loral [3]. The sensor is mounted in a Gulfstream G1 aircraft and has been utilized to collect large sets of imagery containing both tactile targets, man-made and natural clutter.

Fig. 1 shows the three main stages of the Lincoln ATR system: detection, discrimination and classification [1]. The role of detection is to find regions in the imagery that can contain potential targets, and it is built around a CFAR detector. The CFAR is triggered by cells that display an intensity significantly different from the immediate surroundings. Since metallic objects in the open produce high intensity scatters at these frequencies, this test is an effective compromise to separate metallic objects from most of the background clutter when the computational bandwidth of the end-to-end system is constrained (as it must be in practical systems) [4]. The false alarms (i.e., CFAR detections that do not correspond to the objects of interest) are further reduced by the next stage, the discriminator.

The discriminator utilizes many features complementing the amplitude test implemented by the two-parameter CFAR, such as texture (including fractal dimension), size, contrast, and polarimetric features [1]. We propose here to use the QGD as an additional discriminator with the advantage of low computational cost and high performance. This stage basically rejects high radar returns from natural clutter (such as tree tops) and many man-made discretes. The remaining detects are further scrutinized in the classifier that rejects man made discretes and is able to classify the objects of interest.

This staged system has the appeal of providing good performance and an efficient implementation. In fact the CFAR is a computationally simple algorithm and the only one in the ATR signal processing chain that must run over the entire imagery. The following stages just analyze the cells that trigger the CFAR, so more computationally demanding but also more accurate algorithms can be implemented. The CFAR has to work at or very close to 100% detection probability. The number of false alarms created by the detection stage dictates the computational bandwidth of the entire system. Presently the Lincoln Laboratory system achieves 0.3 false alarms per km^2 for a probability of detection of 90% [5].

B. Polarimetric Whitening Filter

One of the advantages of radar for surveillance is the all-weather, 24 h capability of the sensor, which is unlike optical or infrared sensors. The shortcoming is the lack of resolution due to the wavelength and the intrinsic noise produced by the image formation which is called speckle. A useful improvement in SAR processing was the invention of the polarimetric whitening filter or PWF [6, 7]. A cell in fully polarimetric SAR is represented by an 8 component vector x formed by four (complex) polarized measurements (VV, HV, HH, VH), where VV stands for vertical send/vertical receive, HV horizontal send/vertical receive, etc. PWF is a linear projection of the fully polarimetric data x which produces an intensity cell y with a minimum amount of speckle. The goal is to design a weighting matrix W that minimizes the variance normalized by the mean of the radar return,

$$\frac{\mathrm{var}(y)}{E^2(y)} = \frac{\mathrm{var}(x^\dagger W x)}{E^2(x^\dagger W x)} \tag{1}$$

where † denotes the Hermitian transpose. The formulation of this optimization problem with Lagrange multipliers shows that this solution whitens the polarimetric measurements. So PWF has the dual role of reducing the data (from an 8-dimensional vector to a real value) and producing a scalar projection where the speckle is minimized.

C. CFAR Prescreening

The process of target prescreening should be a fast and robust computational algorithm based on a few local statistical measurements. Intensity of a test cell relative to the local mean and variance is one such useful prescreener that has been shown to work well for identifying potential targets in SAR images [1]. Under the assumption that the radar returns from clutter are Gaussian distributed a test that declares as targets test cells whose amplitudes stand out from their surroundings according to

$$y = \frac{x_0 - \overline{x}}{\hat{\sigma}} \quad \begin{matrix} > T_{\mathrm{CFAR}} & \text{(target)} \\ \leq T_{\mathrm{CFAR}} & \text{(clutter)} \end{matrix} \tag{2}$$

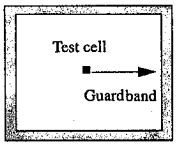

Fig. 2. CFAR detector stencil.

behaves as a CFAR detector. The assumption of Gaussian distributed clutter is not met in high resolution millimeter wave SAR, but experience has shown that (2) is an effective test for detecting targets in clutter [1]. In (2) x_0 is the intensity of the single pixel under analysis, $\bar{x}$ and σ are estimates of the mean and standard deviation of the clutter in the local neighborhood of x_0, and T_{CFAR} is the threshold. An important characteristic of this detector is the so-called *guard area* which ensures that the clutter statistics are collected at some distance away from the cell under test, thus prohibiting target pixels from contaminating the estimate of clutter mean and standard deviation (see Fig. 2).

Average pixel intensity radially away from a test cell is clearly a useful source of information. Delicate issues are the accuracy of the local statistical estimation and the assumption of Gaussian distributions. We introduce a novel target discriminator that also exploits this local intensity neighborhood level information by employing a class of analyzing functions which can be adaptively scaled (and shaped) so as to capture the regions (around the test cell) containing more discriminatory information. In contrast to the two-parameter CFAR detector, where the decision rule and the guard band radius are assigned a priori based on ad-hoc considerations (target size), we relax this constraint and allow all parameters of the discriminator to be optimized through a training process to maximize the separability of the two classes. Particularly in SAR where the target detail controls the reflectivity, physical dimensions alone should be used as mere indicators of a suitable guard band.

III. QUADRATIC GAMMA DISCRIMINATOR

A. Gamma Functions for Local Feature Extraction

A circularly symmetric 2-D gamma kernel family $g_{n,\mu}(k,l)$ given by

$$g_{n,\mu}(k,l) = C(\sqrt{k^2+l^2})^{n-1}e^{-\mu\sqrt{(k^2+l^2)}} \qquad \Omega = \{(k,l);\ -N \le k,\ l \le N\} \tag{3}$$

is used to estimate local statistics in concentric bands around the test cell (0,0). In (3) the integers k and l are the two space coordinates, N specifies the square region of support Ω and C is a normalization constant (unit volume under the surface). As in the

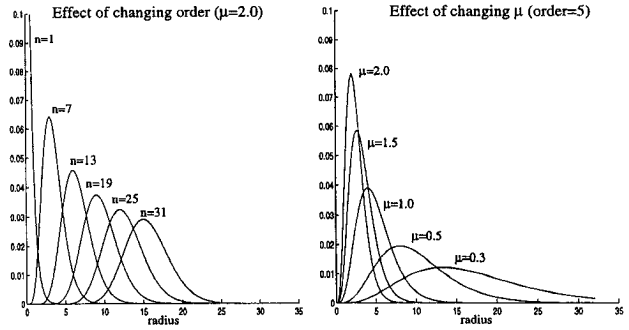

Fig. 3. Gamma function parameters.

1-D case (Fig. 3), the shape of the functions $g_{n,\mu}(k,l)$ is controlled by the kernel order n and the parameter μ [12]. The 2-D functions in (3) are obtained by rotating these 1-D curves around the radius zero axis. For $n = 1$ the 2-D gamma function is a radially symmetric decaying exponential, and with increasing $n = 2,3,\ldots$ the functions form rings around the origin with increasing radius. For each $n > 1$ the parameter μ controls the effective radius where the peak occurs as well as the width of each concentric ring, since the peak of the ring occurs at a radius given by n/μ.

The gamma kernels provide an interesting compromise between global and local approximators: the kernel is a family of local approximators with different shapes, spanning an ever increasing region of support when n is increased. Suppose the goal is to approximate a two-dimensional function $x(k,l)$ in the region Ω as

$$\hat{x}(k,l) = \sum_{n=1}^{M} w_n g_{n,\mu}(k,l) \tag{4}$$

where $\hat{x}(k,l)$ is the approximation to $x(k,l)$ and $M < N$. M is the number of basis functions and normally is defined a priori. A gamma kernel of order $M < N$ with the appropriate μ parameter can expand the area about the fiducial point where the function $x(k,l)$ is well approximated by $\hat{x}(k,l)$. In fact, as shown in Fig. 3(b), μ controls the region of support of the kernel for a given order, so we call μ the scale of the approximation. In an adaptive signal processing framework, the μ parameter can be automatically adjusted to provide the best performance for a given approximation metric even after the decision of the order has been made. More details about the properties of the gamma kernels can be found in [2] which treats the 1-D case, but the properties extend to the 2-D case using (3).

Thinking about the stencil as effecting an image projection into a set of local bases is very appealing because it brings approximation theory to stencil design. The 2-D gamma kernels can construct a predefined radially symmetric stencil, i.e., a mask over the input image space. Fig. 4 shows the amplitude mask constructed with the combination of $g_{1,\mu}$ and $g_{15,\mu}$ (for simplicity, the dependence on k,l is dropped

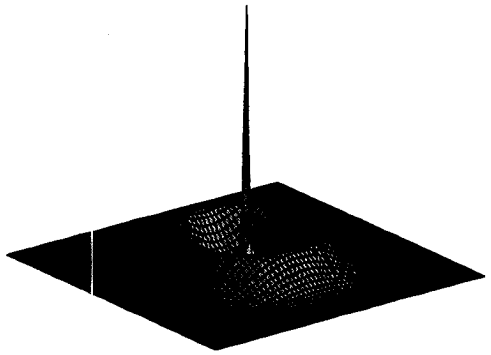

Fig. 4. Stencil for QGD which combines two gamma kernels ($n = 1, n = 15$) for $\mu_1 = \mu_{15} = 0.6$.

from the notation). This stencil resembles the CFAR stencil. In this application $g_{1,\mu}$ implements the mask to estimate the intensity of the cell under test, and $g_{15,\mu}$ the mask to estimate the intensity in the surrounding neighborhood.

In our experiments, only two gamma kernels of orders 1 and 15 were used to preserve the analogy with the two-parameter CFAR detector. For applications where sudden change of amplitude around a pixel location represents characteristics of the targets, a combination of kernels $g_{1,\mu}$ and $g_{15,\mu}$ will provide the required mask to implement the CFAR test. An added advantage of the gamma stencil is that the scale parameter μ controls the extent of the kernels (i.e., the size of the region used to estimate the local image intensity) and may be adapted for best performance depending upon the target and the clutter characteristics. This is useful since there are indications that some regions around the test cell contain more discriminatory information than others, as the need for the guard band in CFAR detector suggests. Here we will be independently adapting the scale for each gamma kernel, i.e., we will have two scales determined by μ_1 and μ_{15}.

B. Local Features and the Decision Rule

The role of the QGD is to further separate targets from nontargets in the cells that trigger the CFAR. However, unlike other discriminators used in SAR, the QGD uses the same features as the CFAR but constructs a quadratic discriminant function of the local image intensity of the cell under test and its surroundings estimated by the gamma stencil. An efficient way to implement a quadratic discriminant is to compute first the second order expansion of the original features and subsequently combine them linearly [8].

Features are extracted from the local neighborhood of a cell under test (the fiducial point) by centering gamma kernels on that cell and taking the inner product with the kernels over Ω of the image amplitude x and the local image power x^2. Here, we consider x to be a scalar reflectivity measure derived from fully polarimetric SAR data obtained using the PWF [4]. We now introduce a more compact notation and we specify the local image features as projections of amplitude and power onto gamma kernels,

$$\begin{aligned} y_{n,\mu} &\equiv \sum_{i,j\in\Omega} g_{n,\mu}(i,j)x_{\mathrm{PWF}}(i,j) \\ y^2_{n,\mu} &\equiv \sum_{i,j\in\Omega} g_{n,\mu}(i,j)x^2_{\mathrm{PWF}}(i,j) \end{aligned} \tag{5}$$

where $n = 1, 15$ for a total of 4 features. Note that (5) can be interpreted as 2-D finite impulse response (FIR) filtering operations of the original image and its power with the gamma stencil $g_{n,\mu}(k,l)$ in locations specified by the CFAR detector. The output $y_{n,\mu}$ can be viewed as an estimate of the 1st-order moment of the local image amplitude, while $y^2_{n,\mu}$ can be interpreted as the 2nd-order moment estimation which are sufficient to compute the local variance.

Following our goal of implementing a quadratic discriminant by feature expansion of all the first- and second-order terms we expand the original four features into an 8 value feature vector

$$\begin{aligned} f_\mu(x) = [&y_{1,\mu} \quad y_{15,\mu} \quad y^2_{1,\mu} \quad y^2_{15,\mu} \quad (y_{1,\mu})^2 \\ &(y_{15,\mu})^2 \quad (y_{1,\mu})(y_{15,\mu}) \quad 1]^T. \end{aligned} \tag{6}$$

This feature vector enables a direct computation of the local neighborhood variance as needed in the two-parameter CFAR (2) [1]. As explained in the following section, these terms enable a one-to-one comparison of the gamma discriminator and the two-parameter CFAR detector and are also used to study the benefits of using adaptive feature extraction and adaptive weights.

At this stage of our research we are using the quadratic decision function which is specified by a linear weighting of the quadratic features of (6),

$$\begin{aligned} y &= w^T \cdot f_\mu(x) \\ &= w_1(y_{1,\mu}) + w_2(y_{15,\mu}) + w_3(y^2_{1,\mu}) + w_4(y^2_{15,\mu}) + \cdots \\ &\quad + w_5(y_{1,\mu})^2 + w_6(y_{15,\mu})^2 + w_7(y_{1,\mu})(y_{15,\mu}) + w_8 \\ &> T \quad \text{(target)} \\ &\leq T \quad \text{(clutter)} \end{aligned} \tag{7}$$

where the weights $(w_1, w_2, \ldots, w_8)$ are obtained during a training procedure. The discrimination between the two input classes (targets and clutter) is done using a single threshold T. In order to validate the quality of the discriminant function proposed in (7), T will be varied in the tests to create an ROC curve. Following the validation in realistic data sets, T is selected such that the required probability of detection is achieved. The above equations are summarized in the discriminator structure (Fig. 5) which consists

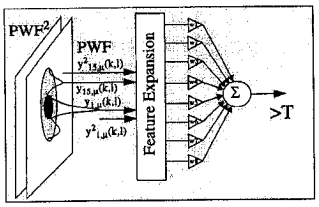

Fig. 5. QGD structure. PWF and PWF^2 represent inputimage plane and its square.

of the feature extraction block followed by feature expansion and the decision block.

We call this classifier the QGD, since the discriminant function is a quadratic function of the image amplitude features extracted by the gamma stencil.

C. Gamma Discriminator Generalizes the CFAR Detector

To draw a parallel between the gamma discriminator decision function (7) and the two parameter CFAR detector (2), we first observe that $\bar{x}$ and σ^2 are estimates of the local mean and variance around the test cell (indexed by 0). The standard deviation can be computed from the first and the second moments as

$$\hat{\sigma} = \sqrt{x^2 - \bar{x}^2}. \tag{8}$$

Next we rewrite the condition for target detection by cross multiplying and squaring both sides of the two parameter CFAR equation (2)

$$x_0^2 - 2x_0\bar{x} + \bar{x}^2 - T_{\text{CFAR}}^2 x^2 + T_{\text{CFAR}}^2 \bar{x}^2 > 0. \tag{9}$$

Equation (9) can be interpreted as a restricted quadratic discriminant function with a fixed set of weights for a given T_{CFAR}. We call (9) restricted because only some of the quadratic terms of the intensities are present here. The quantities that appear in this equation correspond closely to some of those used in gamma discriminator (7). Specifically, $y_{1,\mu}$ has the same role as x_0, and $y_{15,\mu}$ corresponds to the local mean $\bar{x}$. Similarly, for the second moment we have correspondence between $y_{15,\mu}^2$ and $\bar{x}^2$. Hence the two-parameter CFAR corresponds to a QGD with the following weights:

$$[0 \quad 0 \quad 0 \quad -T_{\text{CFAR}}^2 \quad 1 \quad 1+T_{\text{CFAR}}^2 \quad -2 \quad 0]$$

Hence, the QGD formulation preserves the similarity to the two-parameter CFAR detector (the types of features used) but has generalized it with respect to 1) the shape of the kernels used for the mean and variance estimation, 2) the number of features, 3) the selection of the weights of the decision function which are not chosen a priori but are adapted with the goal of optimally separating targets from nontargets.

Unlike the CFAR, the QGD parameters need to be optimized through a training procedure. For that reason we expect it to be able to achieve better discrimination performance, but the choice of the training data set becomes crucial. The linear and the bias terms which are used to complete the quadratic feature set bring an additional flexibility to the gamma discriminator in terms of sensitivity to absolute intensity measures. With the CFAR detector a target is detected when the intensity of a pixel exceeds its surroundings by some threshold relative to the local mean and standard deviation. A pixel with low intensity may therefore be classified as a detection if the surrounding clutter is smooth and has low amplitude [11]. Although this is generally a valid assumption, it is also reasonable to assume that the absolute radar return from a target could carry discriminatory information, especially when the SAR images are calibrated. By adding the terms $y_{1,\mu}$, $y_{15,\mu}$ (corresponding to x_0 and $\bar{x}$) and the bias term which is not present in the CFAR detector, we allow the discriminator to become sensitive to absolute intensity measures as well. We have noticed significant improvement in discriminant results when these terms were added, but it is not clear yet how robust they are with respect to radar calibration errors.

D. Training From Exemplars

Given a set of training images $\{x_0, x_1, x_2, x_{P-1}\}$ centered around points of a known class (target or clutter) we compute the corresponding feature vectors $\{f_\mu(x_0), f_\mu(x_1), f_\mu(x_2), \ldots, f_\mu(x_{P-1})\}$ and arrange them as rows of the matrix $F_\mu(X)$. Since a large discriminator output is desired for test cells that belong to targets and a low output for clutter, we construct a vector $d = [d_0, d_1, d_2, \ldots, d_{P-1}]$ such that $d_i = 1$ for targets and 0 otherwise. Effectively this vector of desired response is a constraint for the optimization. Next, we compute the linear weight vector $\mathbf{w}$ by solving an overdetermined (assuming $P > 8$) system of linear equations

$$d = F_\mu(X)w \tag{10}$$

where the optimal weights $\mathbf{w}_{\mathbf{o},\mu}$ are found in the least squares sense as [9]

$$\begin{aligned} w_{o,\mu} &= \text{argmin}(\|d - F_\mu w\|_2) \\ w_{o,\mu} &= (F_\mu^T F_\mu)^{-1} F_\mu^T d \end{aligned} \tag{11}$$

where $(\cdot)^T$ and $(\cdot)^{-1}$ represent matrix transpose and matrix inversion operations, respectively. Notice that this is a parametric least squares problem, since the solution depends on the parameter μ. For the case of the gamma discriminator the optimization is nonlinear in μ (μ appears on the exponent of the kernel), which is normally much more delicate than the computation of the optimal weights $\mathbf{w}_{\mathbf{o},\mu}$ due to the problem of local minima [2]. Since our primary objective was to evaluate the full potential of this novel feature set and

to study the effect of the adaptive scale parameter on performance, we circumvented the problem of local minima by employing an exhaustive search to find the optimal μ. We computed 34 sets of optimal weights covering logarithmically the operating range of the parameter μ ($0 < \mu < 4$). The optimal value is then simply picked as the minimum of the scanned curve representing the mean square error or another measure of performance (such as false alarms).

E. Implementation Issues

We designed a test procedure that emulates a real world situation where the discriminator output is computed for each pixel in the image. Computing the feature set at each pixel of the image amounts to correlating each of the two kernels $y_{1,\mu}$ and $y_{15,\mu}$ with both the original image and the image squared. Each kernel assumes a role of an FIR filter with rectangular support (size 65×65). The four base features at point (k,l) are then obtained using translated gamma kernels as

$$f_{n,p}(k,l) = \sum_{-32 \le i,j \le 32} g_{n,\mu}(k+i, l+j) x^p(i,j) \quad (12)$$

where n stands for the kernel order and takes the values $n = 1, 15$, and $p = 1, 2$ indicates either 1st or 2nd moment. Correlations are computed in the frequency domain using fast Fourier transforms (FFTs) to obtain better computational efficiency. For processing large images and to avoid memory problems we divide the image into overlapping radix 2 windows which are individually processed and combined using an overlap and save method [10].

The gamma discriminator needs a training set, both for targets and clutter. Adaptive systems can fine tune their parameters for optimal performance, but the training data sets have to contain realistic coverage of the variety of target and clutter conditions the discriminator will find in practice. This coverage of all possible conditions may be difficult, primarily if the adaptive system has a large number of free parameters. However, the QGD has only 8 parameters, which means that the input space is projected into a small size feature space, where much of the variability of the input data is reduced through projection. Hence, in our opinion, an adaptive system with a small number of parameters (such as the gamma discriminator) can be effectively trained with manageable training set sizes. A rule of thumb we use is that the training set should have a size of at least 10 times the number of free parameters and cover a diverse set of operating conditions.

The selection of training exemplars for the target class is rather straightforward. Image chips (65×65) of the different types of targets (with the targets centered) are extracted from the data base. The selection for the clutter class is a little more involved because the training set should have exemplars that are both representative of the class and help discriminate between targets and clutter. For instance the training information of low reflectivity clutter is not as important as high reflectivity clutter such as tree tops. An understanding of the gamma discriminator is therefore necessary to establish the best training set. The results presented below show that this is achievable under realistic conditions of targets in clutter.

IV. TEST RESULTS

A. Importance of Scale in Discrimination

In our preliminary experiments, the gamma discriminator was optimized on a training set extracted from the MIT Lincoln Laboratory mission 90 pass 5, 33 GHz SAR data set collected in Stockbridge, New York. This fully polarimetric high-resolution ($1 \text{ ft} \times 1 \text{ ft}$) data was preprocessed using an optimal whitening technique known as the PWF which combines the HH, HV, and VV polarization components into minimum-speckle imagery [4].

The ultimate objective of the test was to quantify the reduction in false alarms achievable with the QGD. However, we would also like to understand how estimating the local statistics at different scales affects the performance of the gamma discriminator. Therefore, this first set of results should not be considered testing of the QGD, but simply a quantification of some of its characteristics. Recall that the scale can be controlled by the selection of the parameter μ. The performance of the 1-D gamma filter is dependent upon μ [2] so we would like to also quantify its effect for image processing. Due to the unavailability of actual tactical targets in mission 90 data we emulated the testing framework by choosing vehicles (cars) in the SAR imagery to represent our "target" class. We identified and labeled 16 stationary vehicles located in frame 7 of the mission 90 data. Due to the low number of vehicles available we used all of them in training, creating 10 training exemplars from each vehicle by slightly displacing the center of the chip. Most of the vehicles were taken from a parking lot and they represent a difficult detection problem due to their close proximity. The clutter class was also taken from mission 90 SAR data. Open fields and tree tops were primarily used to represent the clutter class; a total of 302 clutter chips were extracted.

Equation (11) was used to compute the weights of the discriminant function at each value of μ_{15} as described previously. The curves in Fig. 6 demonstrate how the selection of μ_{15} affects the discriminator performance in terms of two different criteria, false alarms (left curve) and mean square error (MSE)

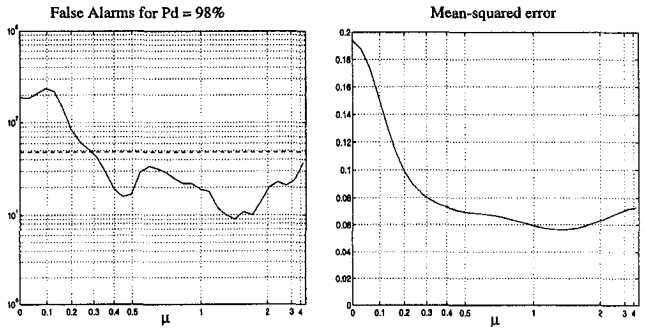

Fig. 6. Performance curves as function of μ_{15}. Dashed line represents CFAR performance on same data.

between the desired response (1 for targets and zero for clutter) and the gamma discriminator output.

A distinguishing aspect of these plots is the nonconvex shape of both curves. The number of false alarms rises to 235 for $\mu_{15} = 0.09$ and has a minimum of 9 for $\mu_{15} = 1.38$, increasing again to 35 for $\mu_{15} = 4.0$. This behavior is not monotonic. It could be due to the reduced number of training exemplars, but this is exactly the same behavior we encountered in the 1-D case for system identification [2] where plenty of data was available. Therefore we conclude that in this family of kernels, the performance is a nonconvex function of the scale. The scale at which the target and clutter statistics are estimated affect the performance of the gamma discriminator. Hence, a detection algorithm that has the ability to choose the region of support (scale) where features are estimated seems very important for improved performance.

During training with an on-line algorithm, it is difficult to adapt directly the parameter μ with the false alarm rate, but techniques have been developed to train μ with the information of the output MSE between the desired response and the output of the discriminator [2]. This curve is shown in Fig. 6 (right plot). For this case ($Pd = 98\%$) we can see that the minimum of the MSE corresponds to the minimum of the false alarm rate, i.e., adapting μ with the MSE will yield the minimum false alarm rate. However, we found that this is not always the case when the probability of detection is set at other values (notably $Pd = 100\%$). Another aspect to mention is the broad minimum which means that there are a range of scales (values of μ) for which the system performance is adequate. Alternatively, this implies that a range of target sizes can be accommodated with the same stencil parameters without affecting performance appreciably.

The two-parameter CFAR detector was implemented according to [1], and tested on the same segment of SAR data. Fig. 6 shows the number of false alarms of the two parameter CFAR (dashed line) for $Pd = 98\%$. The gamma discriminator produces a smaller number of false alarms than the two-parameter CFAR for a large range of μ values ($0.3 < \mu < 4$).

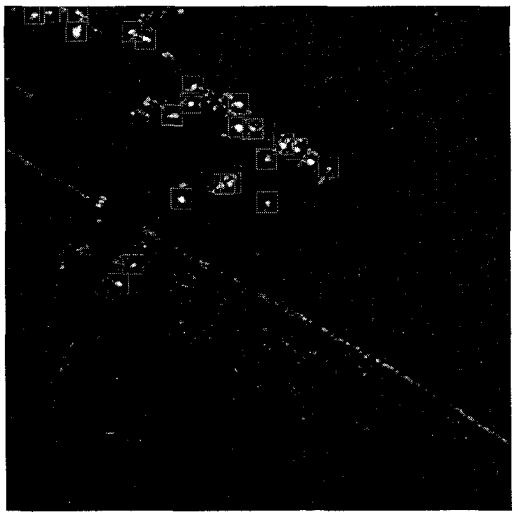

Fig. 7. Detection of vehicles in parking lot. Each square represents a detection.

A sample detection result for a section of frame 7 of the Lincoln Laboratory SAR data set is shown in Fig. 7. In this image the discriminator output is intensity coded, with brighter values indicating likely targets. Here, the threshold was set so that all vehicles are detected ($Pd = 100\%$) and at the same time the number of false alarms remains rather low. In this experiment, a simple clustering algorithm was used to merge multiple detections that were close to each other, which yields a more representative count of false alarms.

B. Large Scale Tests with Military Targets

The tests just described were conducted at the Computational NeuroEngineering Laboratory (CNEL) with the exclusive goal of providing the necessary knowledge of how the QGD handled the imagery, how it should be adapted and to test our algorithmic implementation. Under a contract from DARPA, the QGD algorithm was shared with Lincoln Laboratory for a realistic testing. The results are presented below.

The data for the first test case consisted of high resolution (1 ft $\times$ 1 ft) fully polarimetric SAR imagery preprocessed using PWF. The QGD was trained on two target types using spotlight data and also man-made discretes from stripmap clutter data. A total of 135 target images (chips) were chosen for training; these were 5 deg apart in aspect angle (i.e., 5, 10, 15 deg, etc.). The clutter data used for training consisted of 100 typical man-made discretes. Evaluation of this test case was performed using spotlight target and stripmap clutter data. As in the training stage, spotlight data of two targets that were 5 deg apart in aspect angle (i.e., 3, 8, 13, 18 deg, etc.) were used for testing. The test clutter data consisted of 4,727 stripmap clutter chips extracted from a total of 56 square Km in area. Thus, the test data set for this

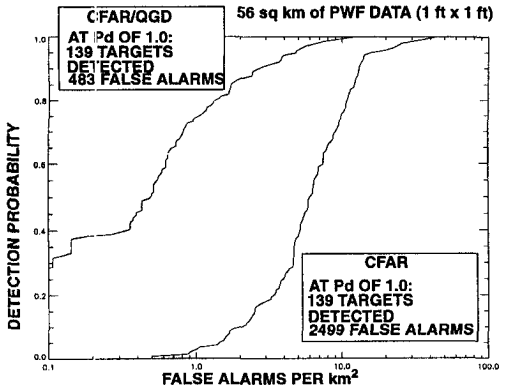

Fig. 8. Discrimination performance of Gamma discriminator versus CFAR.

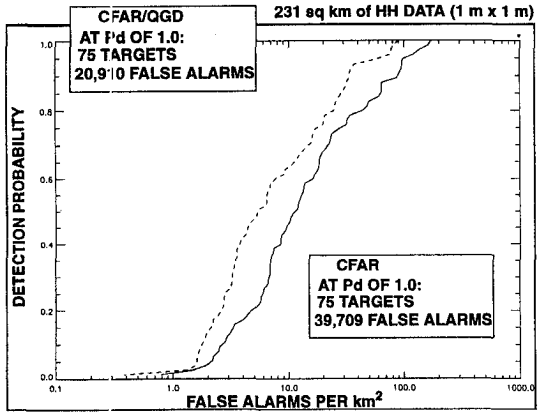

Fig. 9. Comparison of two parameter CFAR and QGD for 1 polarization, 1 m data.

experiment was composed of 139 target chips and 4,727 clutter chips.

The QGD was evaluated by running the data through the CFAR algorithm first and only the chips that triggered the CFAR were evaluated by the QGD. Then, the ROC curves were obtained by computing the cumulative number of false alarms out of each algorithm. At a $Pd = 100\%$ the CFAR algorithm detected 139 targets and had 2,499 false alarms, whereas the output of the CFAR/QGD while also detecting 139 targets, reduced the above-mentioned false alarm number to 483 (see Fig. 8).

The second test case used single channel (HH) stripmap imagery with a resolution of 1 m × 1 m. The training set for the QGD algorithm consisted of 52 target chips and 150 clutter chips that represented two types of targets and man-made clutter. The evaluation of the algorithm was performed using 75 target chips and 44,599 clutter chips, corresponding to detects of the two parameter CFAR algorithm when analyzing a 231 square Km of area. Fig. 9 shows the results in the form of an ROC curve. At $Pd = 100\%$, the two parameter CFAR algorithm had 39,709 false alarms, while the combination CFAR/QGD algorithm also detected all 75 targets and had only 19,037 false alarms.

These two experiments constitute a partial, although extensive, test of the performance of the gamma discriminator. They indicate that the algorithm is capable of generalizing the information utilized in the training set. Further testing of the QGD with other types of clutter and other selection of tactical targets is necessary to better characterize its performance. Since the original features of the QGD are the same as the CFAR, the better discriminant characteristics of the gamma discriminator are either due to a better estimation for the mean and standard deviation for targets and clutter, or to the larger feature space. We are presently investigating this topic. In tests conducted at the CNEL, the QGD applied directly to the imagery produced disappointing results, with many false detections in apparently low intensity smooth areas. This means that the QGD has higher sensitivity than the CFAR but is less specific, so the combination of the two as proposed here should be utilized.

V. SUMMARY AND CONCLUSIONS

We have presented a new discriminator module for SAR based on a 2-D extension of the gamma functions. The appeal of this discriminator is that it uses the same features extracted for the CFAR (intensity and power of the cell under test and its surroundings). The implementation that we chose uses only a subset of the gamma functions (g_1 and g_{15}). It was derived by analogy with the two parameter CFAR detector to enable a straightforward comparison with this widely used algorithm. In the gamma discriminator, the estimates of the local statistics are obtained by convolution with the g_1 (cell under test) and g_{15} kernels (local neighborhood) which have a free parameter that controls the scale and that can be adapted during training with the output error. Preliminary testing of the gamma discriminator in the Computational NeuroEngineering Laboratory shows the importance of the scale in the performance. The Lincoln Laboratory tests show that the combination CFAR/QGD can improve the false alarm rate of the two-parameter CFAR detector without affecting the probability of detection. In the 1 polarization 1 m data, the combination CFAR/QGD outperformed any of the alternate front-end prescreeners, but in general we recommend that the QGD should be utilized as an extra discriminator to reduce the number of false alarms. Novak reports [12] that the output of the QGD is weakly correlated with the other discriminators, which means that the QGD enhances the present set of discriminators for SAR. We are very encouraged with these results, and plan to investigate further the performance of the gamma discriminator and introduce several improvements.

One aspect that requires further testing is the generalization properties of the QGD to changes in

vehicle size or clutter type. An important question is to find out which conditions (particularly clutter types) must be included in the training set for unconstrained use of the discriminator, and how the performance is affected by these extended clutter mixtures. Likewise, the size and type of the tactical targets should also cover a larger selection and the effect on performance quantified. The performance curve of Fig. 6 shows that the gamma kernels should accommodate vehicles of different sizes without drastic performance degradation, but further tests of robustness are necessary. One has to remember that it is possible to use more than two gamma kernels to compensate any eventual drop in performance. If the radar resolution changes, the kernels have to be readapted.

REFERENCES

[1] Novak, L. M., Owirka, G. J., and Netishen, C. M. (1993)
Performance of a high-resolution polarimetric SAR automatic target recognition system.
Lincoln Laboratory Journal, **6**, 1 (1993), 11–23.

[2] Principe, J. C., de Vries, B., and de Oliveira, P. G. (1993)
The gamma filter—A new class of adaptive IIR filters with restricted feedback.
IEEE Transactions on Signal Processing, **41**, 2 (1993), 649–656.

[3] Henry, J. (1991)
The Lincoln Laboratory 33 GHz airborne polarimetric SAR imaging system.
In *Proceedings of the IEEE National TeleSystems Conference*, Atlanta, 1991, 353.

[4] Novak L. M., Burl, M. C., Chaney, R. D., and Owirka, G. J. (1990)
Optimal processing of polarimetric synthetic aperture radar imagery.
Lincoln Laboratory Journal, **3**, 2 (Summer 1990).

[5] Novak, L., Halversen, S., Owirka, G., and Hiett, M. (1995)
Effects of polarization and resolution on the performance of a SAR automatic target recognition system.
Lincoln Laboratory Journal, **8**, 1 (1995), 49–65.

[6] Novak, L., and Burl, M. (1990)
Optimal speckle reduction in polarimetric SAR imagery.
IEEE Transactions on Aerospace and Electronic Systems, **26**, 2 (1990), 293–305.

[7] Novak, L. M., Burl, M. C., and Irving, W. W. (1993)
Optimal polarimetric processing for enhanced target detection.
IEEE Transactions on Aerospace and Electronic Systems, **29**, 1 (1993), 234–244.

[8] Tou, J. T., and Gonzales, R. C. (1974)
Pattern Recognition Principles.
Reading, MA: Addison-Wesley, 1974.

[9] Dudgeon, D., and Mersereau, R. (1984)
Multidimensional Digital Signal Processing.
Englewood Cliffs, NJ: Prentice-Hall, 1984.

[10] Golub, G. H., and Van Loan, C. F. (1989)
Matrix Computations.
Baltimore, MD: John Hopkins, 1989.

[11] Novak, L., and Netishen, C. (1992)
Polarimetric synthetic aperture radar imaging.
International Journal of Imaging Systems and Technology, **4** (1992), 306–318.

[12] Novak, L. (1997)
Recent advances on ATR using SAR imagery.
Seminar at the University of Florida, Feb. 1997.

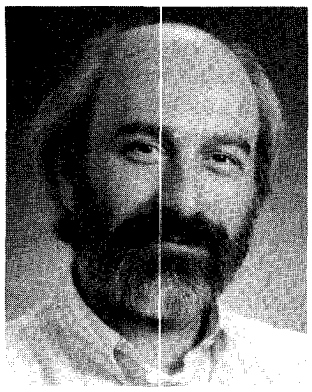

Jose C. Principe (M'83—SM'90) is Professor of Electrical and Computer Engineering at the University of Florida, Gainesville, where he teaches signal processing and artificial neural networks. He is the Founder and Director of the University of Florida Computational NeuroEngineering Laboratory (CNEL). His primary area of interest is processing of nonstationary signals with adaptive neural models.

Dr. Principe is a member of the advisory board of the University of Florida Brain Institute.

Margarita Hiett 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 Company, where she was involved in implementation of large scale real-time hardware and software systems which performed detection and discrimination of IR targets. From 1992 to 1993 she worked at Textron company where she continued her work analyzing real-time algorithms used in detection of moving IR targets. She has been at Lincoln Laboratory since 1993, where she has analyzed detection algorithms for stationary radar targets.

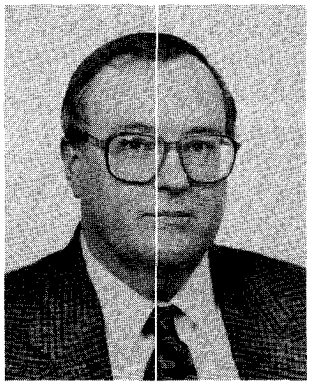

Leslie M. Novak received a B.S.E.E. degree from Fairleigh Dickinson University, Rutherford, NJ, in 1961, an M.S.E.E. degree from the University of Southern California in 1963, and a Ph.D. degree in electrical engineering from the University of California, Los Angeles, in 1971.

Since 1977 Dr. Novak has been a member of the technical staff at Lincoln Laboratory, where he has studied the detection, discrimination, and classification of radar targets and is a senior staff member in the Surveillance Systems group.

He has contributed chapters on stochastic observer theory to the series *Advances in Control* (Academic Press, New York), Vols. 9 and 12.

John Fisher, III Picture and biography not available.