

# A Multiresolution Approach to Discrimination in SAR Imagery

WILLIAM W. IRVING  
ALPHATECH, Inc.

LESLIE M. NOVAK  
Lincoln Laboratory

ALAN S. WILLISKY, Fellow, IEEE  
M.I.T.

We develop and test a new algorithm for discriminating man-made objects from natural clutter in synthetic-aperture radar (SAR) imagery. This algorithm exploits characteristic variations in speckle pattern as image resolution is varied from course to fine. We model these variations as an autoregression in scale, and then use the autoregressive model to define a multiresolution log-likelihood ratio discriminant. We incorporate this discriminant into the existing Lincoln Laboratory SAR system for automatic target recognition (ATR), and test the augmented system by applying it to millimeter-wave SAR imagery having 0.3 m resolution and representing 56 square kilometers of terrain. At a probability of detection of 0.95, the addition of the multiresolution discriminant reduces the number of natural-clutter false alarms by a factor of six.

Manuscript received May 1, 1995; revised January 7 and August 8, 1996.

IEEE Log No. T-AES/33/4/06840.

The work of William Irving and Alan Willisky was supported by the Advanced Research Projects Agency under Grant F49620-93-1-0604, by the Air Force Office of Scientific Research under Grant F49620-95-1-0083 and by the Army Research Office under Grant DAAL03-92-G-0115. The work of Leslie Novak was supported by the Advanced Research Projects Agency under Air Force Contract F19628-95-C-0002.

Authors' addresses: W. W. Irving, ALPHATECH, Inc., Executive Place III, 50 Mall Rd., Burlington, MA 01803-4562; L. M. Novak, Massachusetts Institute of Technology, Lincoln Laboratory, 244 Wood St., Lexington, MA 02173-9108; A. S. Willisky, Dept. of Electrical Engineering, Massachusetts Institute of Technology, Cambridge, MA.

0018-9251/97/\$10.00 © 1997 IEEE

## I. INTRODUCTION

In this work, we both develop and test a new algorithm for discriminating man-made objects from natural clutter in synthetic-aperture radar (SAR) imagery. In the context of automatic target recognition (ATR), this algorithm provides a computationally fast and simple way to reject regions of exclusively natural clutter, so resources can be focused on the classification of remaining regions, containing man-made objects that are potentially targets.

The novel feature of our algorithm is its exploitation of characteristic variations in speckle pattern, for imagery of natural clutter and of man-made objects, as image resolution is varied from coarse to fine. These characteristic variations arise, because for natural clutter, there is typically a large number of equivalued scatterers in a resolution cell, while for man-made objects, there is typically only a small number of prominent scatterers [1]; this difference leads to very different statistics for the variation in speckle pattern as resolution changes.

The fact that speckle has multiresolution characteristics is also noted and exploited in [1]. However, in contrast to that work, where the different characteristics of natural clutter and man-made objects are used to analyze individual image pixels, we model and exploit these characteristics over entire blocks of imagery. To capture mathematically these scale-to-scale variations, we use a direct scale-recursive extension of the widely used method of autoregressive modeling of time series. To be more specific, let us suppose coherent processing has been applied to a complex-valued SAR image  $\mathcal{I}$ , to create a sequence of images,  $I_0, I_1, \dots, I_L$ , where each successive image has half the resolution (in both range and cross-range) of its predecessor in the sequence. Fig. 1 illustrates such a multiresolution sequence of images, mapped onto a so-called quadtree, for the case  $L = 2$ . We model the statistical coupling of the images  $I_0, I_1, \dots, I_L$  using an autoregression in scale, wherein each pixel in a given finer resolution image is related to its coarser resolution ancestor pixels by a white-noise-driven autoregressive relationship.




Fig. 1. Illustration of multiresolution sequence of 3 SAR images, together with quadtree onto which we map pixel values.

For our discrimination application, we build a pair of multiscale stochastic models: one model for SAR imagery of natural clutter and another for imagery of man-made objects. We then use these models to define a multiresolution discriminant. This discriminant is the likelihood ratio for distinguishing between natural clutter and man-made objects, given a multiresolution sequence of SAR images. Thanks to the structure of our models, the calculation of these likelihoods turns out to be a computationally simple task.

We incorporate our new discriminant into an existing SAR ATR system developed at Lincoln Laboratory [2, 3]. We then apply the augmented ATR algorithm to a data set of HH polarization, millimeter-wave SAR imagery having 0.3 m resolution and representing 56 square kilometers of terrain. As shown in Section IV, this augmented system leads to an improvement in receiver operating characteristics, compared with the standard Lincoln system. For instance, at a probability of detection of 0.95, the number of natural-clutter false alarms produced by the augmented system is one-sixth the number of false alarms produced by the original system.

The body of this work is divided into five major sections. In Section II, we describe in greater detail the Lincoln SAR ATR system. Then in Section III, we describe our class of multiscale models for SAR imagery, our procedure for identifying these models, and our procedure for calculating our multiresolution discriminant. In Section IV, we show the results of our testing of these discriminators. Finally, we summarize the main points and suggest directions for future work in SAR applications of multiresolution-based techniques.

## II. LINCOLN SAR ATR SYSTEM

The Lincoln SAR ATR system is conveniently decomposed into a sequence of three processors: a *prescreener*, a *discriminator*, and a *classifier* (see Fig. 2). The prescreener searches through imagery representing many square kilometers of terrain, and outputs a collection of so-called regions of interest (ROIs). Each ROI is a subimage extracted from the original SAR data set and centered at a possible target location; collectively, all ROIs represent only a small fraction of this original data set. The discriminator applies further processing to distinguish between two kinds of ROIs: those containing man-made objects and those containing natural clutter. All ROIs that appear to contain natural clutter are discarded. Finally, the classifier assigns each remaining ROI to a predefined target category, or to a *none-of-the-above* category if the ROI appears to contain man-made clutter; in our study, we do not concern ourselves with the third (i.e., classification) stage, and consequently, no detailed




Fig. 2. Illustration of input-output operation of SAR ATR system described in [2].

description of this stage is provided and no further mention of it is made.

### A. Description of Prescreening Algorithm

The prescreener uses a two-parameter CFAR (constant false-alarm rate) algorithm. This algorithm is a pixel-by-pixel adaptive image processing scheme that takes advantage of the fact that in SAR imagery, targets typically appear brighter than nontargets. At each pixel, this algorithm performs the following test:

$$\text{Declare } \begin{cases} \text{target present} \\ \text{target absent} \end{cases} \quad \text{if } \frac{x - \hat{\mu}_c}{\hat{\sigma}_c} \begin{cases} > \\ \leq \end{cases} K$$

where  $x$  is the intensity of the test pixel under consideration,  $\hat{\mu}_c$  and  $\hat{\sigma}_c$  are estimates of the local clutter mean and standard deviation, respectively, and  $K$  is a threshold parameter. To avoid contaminating  $\hat{\mu}_c$  and  $\hat{\sigma}_c$  with pixels from a spatially extended target, these estimates are computed using an annular window at a prespecified radius from the test pixel. Under certain idealized conditions [4], this detection procedure will indeed yield a CFAR, which is a function of the threshold value  $K$ . However, even under more realistic conditions where the CFAR property may not prevail, the algorithm is often employed, because it provides a fast, simple, and reasonable detection algorithm.

In many instances, the CFAR algorithm will yield multiple detections on a single target. These detections are clustered together to form an ROI, with the centroid of the cluster aligned with the center of the ROI. Each ROI is passed on to the discriminator for further processing.

### B. Description of Discrimination Algorithm

The main idea underlying the second stage discriminator is that an ROI has a number of characteristics, or so-called *features*, whose statistical distribution will significantly depend on whether the ROI represents a man-made object or natural clutter. The measured values of these features are combined to form a single, scalar-valued measure of the so-called distance of the ROI from the class of targets of interest. All ROIs having a large distance are labeled *nontarget* and the remaining ROIs are labeled *target*.

1) *Description of Features:* There are nine features available to the standard Lincoln

discriminator. These features span the categories *texture*, *size* and *contrast*, and are measured using finest resolution imagery only.

The values of the textural features (i.e., *standard deviation*, *fractal dimension* and *ranked fill-ratio*) are calculated using the pixel values in a special target-sized region within the ROI. This region is chosen in the following way. A target is hypothesized to exist in the ROI; then, the position and orientation of this hypothesized target are estimated by applying simple matched-filtering with a rectangular-shaped target-sized template. We denote the resulting target-sized region by  $\mathcal{T}$ . The *standard-deviation* feature is calculated as simply the sample standard deviation of the log-detected pixel values in the region  $\mathcal{T}$ . The *fractal-dimension* feature provides an estimate of the Hausdorff dimension of the spatial distribution of the brightest 50 pixel values in the region  $\mathcal{T}$  [3]; the details of this feature's calculation are too involved to include here. The *rank-fill-ratio* feature is the percentage of power that is contained in the brightest five percent of the pixels in the region  $\mathcal{T}$ .

The values of the size features (i.e., *mass*, *diameter*, and *rotational inertia*) are calculated using the pixel values in the region of the ROI containing the ROI's principal object. In a rough sense, this principal object is defined to be the bright blob near the center of the ROI, and is found by applying morphological processing [3]; we denote the resulting principal-object region by  $\mathcal{P}$ . The *mass* feature is calculated as simply the number of pixels in the region  $\mathcal{P}$ . The *diameter* feature is equal to the length of the diagonal of the smallest rectangle (either horizontally oriented or vertically oriented) that encloses the region  $\mathcal{P}$ . The *rotational-inertia* feature is calculated as the second mechanical moment of the region  $\mathcal{P}$  around its center of mass, normalized by the inertia of a square having equal mass.

The values of the contrast features (i.e., *peak CFAR*, *mean CFAR*, and *percent bright CFAR*) are calculated by first applying the CFAR algorithm (see Section IIA) to a log-detected version of the ROI. In this way, we obtain a CFAR image, in which the pixel value at location  $(k, l)$  is equal to the CFAR statistic at that location. The values of the contrast feature are calculated using the pixel values in the region  $\mathcal{P}$  of the CFAR image. The *peak-CFAR* feature is simply the maximum value of the CFAR image within the region  $\mathcal{P}$ . The *mean-CFAR* feature is the sample mean of the CFAR image within the region  $\mathcal{P}$ . The *percent-bright-CFAR* feature is the percentage of pixels in the region  $\mathcal{P}$  of the CFAR image that exceed a certain CFAR value.

## 2) Processing Features into Discrimination

**Decision:** A so-called one-class classification scheme [3, 5] is used to process the features into a discrimination decision. To describe this scheme,

we assemble the measured values of our (scalar) features into a vector  $Z$ . We assume that the conditional probability density function (pdf)  $P(Z | \text{target})$  is known, and this pdf alone is used to make the discrimination decision, using the rule

$$\text{Declare } \begin{cases} \text{target present} \\ \text{target absent} \end{cases} \quad \text{if } P(Z | \text{target}) \begin{cases} > \\ \leq \end{cases} T' \quad (1)$$

where  $T'$  is a threshold parameter.

To implement the decision rule in (1), we must of course specify the conditional pdf  $P(Z | \text{target})$ . In this regard, an empirical analysis in [3] demonstrated that for many choices for features, the conditional distribution of  $Z$  is approximately Gaussian. We invoke the Gaussian assumption, thereby allowing us to express the decision rule (1) in the more explicit form

$$\text{Declare } \begin{cases} \text{target present} \\ \text{target absent} \end{cases} \quad \text{if } (Z - M_t)^T \Sigma_t^{-1} (Z - M_t) \begin{cases} \leq \\ > \end{cases} T \quad (2)$$

where  $T$  is another threshold parameter, and where  $M_t$  and  $\Sigma_t$  are the mean and covariance, respectively, of the Gaussian distribution. To implement this rule, estimates of  $M_t$  and  $\Sigma_t$  are first calculated off-line, using training imagery.

In the numerical experiments Section IV, we actually use an optimized version of the decision rule in (2), in which a simple modification is incorporated. This modification is motivated by the empirical observation that the diameter size feature works best when it is used in isolation. In the first stage of the modified rule, the diameter feature is evaluated; only ROIs having a diameter within a prespecified range are passed to the second stage, while the others are assigned to the nontarget class. In the second stage, the remaining ROIs are processed using the quadratic discriminator described in (2).

## III. MULTIREOLUTION MODEL FOR SAR IMAGERY

In this section, we develop our multiscale stochastic models for SAR imagery, and we define our new multiresolution-based discriminant. Our model for natural clutter, hereinafter referred to as our *natural-clutter model*, is specifically designed to describe imagery of grass, while our model for man-made objects, hereinafter referred to as our *man-made model*, is specifically designed to describe imagery of tactical targets. A natural question is whether these simple modeling choices lead to a discriminant that is robust to variations within each of the two large classes. We see in Section IV that the answer is yes.

### A. Generation of Multiscale Image Sequences

We begin with complex-valued SAR imagery, formed to the highest resolution available. Each pixel value in this imagery represents a measurement of both the amplitude and phase of the radar reflectivity of the scatterers within a resolution cell. In all of our work, we use HH polarization imagery gathered with the Lincoln Laboratory millimeter-wave SAR [6]. Although this choice will affect the specifics of the models we build, our general procedures should be more broadly applicable.

From this full resolution imagery, we assume that ROIs have been extracted. For each (complex-valued) ROI  $\mathcal{I}$ , we create a multiresolution sequence of images,  $I_0, I_1, \dots, I_L$ . This sequence is created directly from  $\mathcal{I}$ , with no dependence on the rest of the SAR data set from which  $\mathcal{I}$  was extracted. We assume for simplicity that the image  $\mathcal{I}$  has the same resolution in both range and cross-range, and we denote this resolution by  $\delta$  (for the Lincoln Laboratory millimeter-wave SAR,  $\delta = 0.3$  m). We denote by  $\mathcal{I}(k, l)$  the measured reflectivity at range/cross-range position  $(k, l)$ . Finally, for convenience only, we assume that this image array is square, consisting of  $N \times N$  pixels, where  $N = 2^M$  for some integer  $M$ .

1) *Finest Scale Image*: The finest scale image is denoted by  $I_0$ ; it has resolution  $\delta \times \delta$  and consists of  $N \times N$  pixels. We create  $I_0$  from an intermediate, log-detected image  $I'_0$ , where  $I'_0(k, l) = 20 \log_{10} |\mathcal{I}(k, l)|$ . To eliminate dependence on the absolute calibration of the radar sensor (which is susceptible to spurious fluctuations), we normalize  $I'_0$ , thereby yielding  $I_0$ , with  $I_0(k, l) = I'_0(k, l) - C_0$ . Here,  $C_0$  is set equal to the sample mean of the pixel values in  $I'_0$ . Because  $I_0$  has no dependence on absolute calibration, our decisionmaking exploits only the *relative* variation of image intensity, with respect to a mean level of zero.

2) *Coarse-Scale Images*: We denote the coarser scale images by  $I_1, I_2, \dots, I_L$ , respectively. We form these images from the original complex, fine-scale data  $\mathcal{I}$  by sequentially applying three processing steps: a) lowpass filtering, b) decimation, and 3) log-detection with normalization. For image  $I_m$ , this processing ultimately yields a  $2^{-m}N \times 2^{-m}N$  square image having resolution  $2^m\delta \times 2^m\delta$ .

To describe the first processing step, we denote the inverse discrete Fourier transform of the 2-D complex data  $\mathcal{I}$  by  $\tilde{\mathcal{I}}$ , where

$$\tilde{\mathcal{I}}(p, q) = \frac{1}{N^2} \sum_{k=0}^{N-1} \sum_{l=0}^{N-1} \mathcal{I}(k, l) \times \exp\left(j \frac{2\pi k}{N} p\right) \exp\left(j \frac{2\pi l}{N} q\right).$$

In terms of  $\tilde{\mathcal{I}}$ , we define  $I'_m$  via

$$I'_m(k, l) = \sum_{p=0}^{N-1} \sum_{q=0}^{N-1} \tilde{\mathcal{I}}(p, q) H_m(p) H_m(q) \times \exp\left(-j \frac{2\pi k}{N} p\right) \exp\left(-j \frac{2\pi l}{N} q\right)$$

where  $H_m(p)H_m(q)$  represents a separable 2-D Hamming window with  $H_m(p)$  defined to be

$$H_m(p) = \begin{cases} 0.54 + 0.46 \cos \frac{2\pi p}{2^{-m}N} & 0 \leq p < 2^{m-1}N \text{ \& } \\ & N - 2^{m-1}N \leq p < N \\ 0 & 2^{m-1}N \leq p < N - 2^{m-1}N \end{cases}$$

In the second processing step, we decimate  $I'_m$  by a factor of  $2^m$  in both range and cross-range. The result is denoted by  $I''_m$  and is related to  $I'_m$  by

$$I''_m(k, l) = I'_m(2^m k, 2^m l), \quad 0 \leq k, \quad l < 2^{-m}N.$$

Finally, in the third step, we apply log-detection and normalization. The image  $I_m(k, l)$  is thus related to  $I''_m$  via

$$I_m(k, l) = 20 \log_{10} |I''_m(k, l)| - \frac{1}{(2^{-m}N)^2} \sum_{p=0}^{(2^{-m}N-1)} \sum_{q=0}^{(2^{-m}N-1)} 20 \log_{10} |I''_m(p, q)|.$$

3) *Mapping the Multiscale SAR Image Sequence onto a Quadtree*: The multiresolution image sequence  $I_0, I_1, \dots, I_L$  is matched quite naturally to the structure of a quadtree, and we consequently use the quadtree for all our SAR image modeling. In Fig. 1, we illustrate our convention for the correspondence between pixel values and tree nodes. To formalize this convention, we associate each node  $s$  on the quadtree with a 3-tuple  $(m, k, l)$ , where  $m$  denotes scale and  $(k, l)$  denotes 2-D location; correspondingly, we denote by  $I(s)$  the image pixel residing at node  $s$ , namely  $I_m(k, l)$ . For example, in the context of Fig. 1,  $I(0)$  corresponds to  $I_2(0, 0)$ .

We define  $m(s)$  to be the scale at which the node  $s$  resides. We also define an upward (i.e., fine-to-coarse) shift operator  $\bar{\gamma}$ , such that  $s\bar{\gamma}$  is the parent of node  $s$  and  $s\bar{\gamma}^k$  is the  $k$ th order ancestor of node  $s$ . For instance, in the context of Fig. 1, when  $s$  corresponds to the 3-tuple  $(0, 2, 3)$ , then  $s\bar{\gamma}$  corresponds to  $(1, 1, 1)$  and  $s\bar{\gamma}^2$  corresponds to  $(2, 0, 0)$ .

While our example in Fig. 1 illustrates a special case in which we have formed a complete sequence of images, down to a single-pixel image at the coarsest resolution possible (i.e.,  $L = M$ ), we allow more generally the possibility of truncating the image sequence at some image having more than a single

pixel (i.e.,  $L < M$ ). When this possibility occurs, the nodes at the coarser scales of the tree have no corresponding measurements. This additional flexibility is useful, because beyond a certain coarseness of resolution, SAR imagery conveys very little meaningful information.

## B. Identifying the Multiscale Dynamics

To characterize jointly the scale-to-scale statistical coupling of a sequence of SAR images,  $I_0, I_1, \dots, I_L$  spanning multiple resolutions, we use a class of statistical models representing a direct scale-recursive quadtree extension of a well-known modeling technique for time-series. Specifically, we relate the SAR pixel value residing at node  $s$  (i.e.,  $I(s)$ ) to its ancestors by a linear autoregression in scale:

$$I(s) = a_{1,m(s)}I(s\bar{\gamma}) + a_{2,m(s)}I(s\bar{\gamma}^2) + \dots + a_{R,m(s)}I(s\bar{\gamma}^R) + w(s). \quad (3)$$

In this relation,  $R$  is the order of the regression, while  $a_{1,m(s)}, a_{2,m(s)}, \dots, a_{R,m(s)}$  are the real, scalar-valued regression coefficients. These regression coefficients are allowed to be scale-varying, but are restricted to be shift-invariant for any fixed scale. We will have occasion to refer collectively to the whole set of coefficients for a given scale, and for this purpose, we define the vector  $\mathbf{a}_k$  as

$$\mathbf{a}_k = (a_{1,k} \quad a_{2,k} \quad \dots \quad a_{R,k})^T.$$

The term  $w(s)$  in (3) represents the residual error in the prediction of  $I(s)$ . We assume that  $w(s)$  and  $w(\sigma)$  are statistically independent for  $s \neq \sigma$ . As we will see, this independence assumption is what allows us to develop an extremely simple procedure for likelihood calculation for the entire piece of multiresolution imagery; thus, the validation of this whiteness assumption is critical. The probability distribution of  $w(s)$  is allowed to be non-Gaussian, and furthermore the distribution is allowed to vary with scale. However, for any fixed scale, the distribution is assumed to be spatially invariant. We denote the standard deviation of the residuals at scale  $m$  by  $\sigma_m$ .

To accommodate the fact that we only have a finite number of SAR images, namely for the finest  $L + 1$  resolutions, we must statistically characterize the initial condition of the recursion in (3). This initial condition comprises the values of the pixels in the  $R$  coarsest scale images  $I_L, I_{L-1}, \dots, I_{L-R+1}$ . Since we are more interested in the scale-to-scale variation of the speckle pattern than in the initial condition of this pattern, we choose to impose no prior knowledge about the initializing values. We simply observe these images, and use the observed values to initialize the recursion in (3).

### 1) Identification of the Regression Coefficients:

For each model we build, we identify a regression




Fig. 3. (a) Image of grass, used to build our natural-clutter model. (b) Image of target-like man-made object, used to build our man-made model.

order  $R$  and corresponding regression coefficients, using complex-valued, training images generated with the Lincoln Laboratory millimeter-wave SAR. For our natural-clutter model, we use a single image, which is displayed in Fig. 3(a). This image represents a homogeneous region of grass at  $\delta \times \delta$  resolution,<sup>1</sup> and consists of  $256 \times 256$  pixels. Fig. 3(b) depicts an example of a scene containing a man-made object. In contrast to the case of natural clutter the object in Fig. 3(b) is spatially localized and nonstationary, and thus it makes no sense to build our model based on a large region of “homogeneous targets” analogous to Fig. 3(a). Thus, to build our target model, we use a training set of 64 SAR images (having  $32 \times 32$  pixels each) of howitzers, each imaged at a different aspect angle.

We convert a given training image  $\mathcal{I}$  into a multiresolution sequence of images  $I_0, I_1, \dots, I_L$ , as described in detail in Section IIIA. We then systematically consider a sequence of possible regression orders,  $R = 1, 2, 3$ . For each proposed order, we apply the autoregression to the training data and solve for the regression coefficients that minimize the sum of the squares of the residuals,

$$\sum_{\{s:m(s)=k\}} [I(s) - a_{1,k}I(s\bar{\gamma}) - \dots - a_{R,k}I(s\bar{\gamma}^R)]^2. \quad (4)$$

This least-squares approach is the most widely used method for parameter estimation, and is eminently reasonable here, especially given our absence of prior knowledge about the statistics of the residuals  $w(s)$  [6, p. 461 of Appendix II]. As we will see, the outcome of this procedure suggests both a natural model order and appropriate regression coefficients for each of the models we build.

Table I summarizes the outcome of applying our estimation procedure. The first column of the table lists the resolution of the image pixels to be predicted (i.e., the resolution of  $I(s)$  in the autoregression in (3)). The second column lists the

<sup>1</sup>As defined in Section IIIA,  $\delta = 0.3$  m.

TABLE I

Image resolution	Regression order	Natural-clutter model			Man-made model			
		Regression coefficients		Residual std. dev.	Regression coefficients		Residual std. dev.	
$\delta \times \delta$	1	0.28			5.4	0.70		7.2
	2	0.31	-0.011		5.4	0.67	0.10	7.0
	3	0.28	0.008	-0.01	5.4	0.69	0.12	0.008
$2\delta \times 2\delta$	1	0.30			5.3	0.87		7.6
	2	0.32	0.02		5.3	0.84	-0.16	7.5
	3	0.28	0.02	0.02	5.3	0.82	-0.11	0.009
$4\delta \times 4\delta$	1	0.25			5.5	0.58		8.5
	2	0.25	0.008		5.5	0.58	0.002	8.5
	3	0.25	-0.008	0.007	5.5	0.57	-0.009	0.01

Note: This table summarizes the outcome of our least-squares procedure for determining both model order  $R$  and regression coefficients  $\mathbf{a}_m$ .

proposed order of the regression. Finally, the last two sets of columns list the regression coefficients  $\mathbf{a}_m$  and the residual standard deviation  $\sigma_m$  for the natural-clutter and man-made models, respectively. For example, according to the table, the appropriate second-order regression for prediction of  $2\delta \times 2\delta$  resolution pixels  $I(s)$  in imagery of man-made objects is  $I(s) = 0.84I(s\bar{\gamma}) - 0.16I(s\bar{\gamma}^2) + w(s)$ , with the residual  $w(s)$  having standard deviation of  $\sigma_m = 7.5$ .

With regard to the natural-clutter model, the table suggests that there is no practical benefit to using a model order greater than one; the higher order regression coefficients have negligible magnitude and the standard deviation of the prediction error is not noticeably reduced by an increased model order. For these reasons, we use a first-order autoregression for the natural-clutter model. On the other hand, for the man-made model, a second-order regression appears to be preferable. In particular, the second-order regression coefficient is not negligible (at least for the  $\delta \times \delta$  and  $2\delta \times 2\delta$  images), and the standard deviation of the prediction error is reduced by increasing the model order from one to two. Because a third-order regression fails to continue this trend of increased benefit, we use a second-order autoregression for the man-made model.

2) *Validation of Residual Whiteness:* As we have previously noted, a critical assumption in our modeling framework is that the residuals  $w(s)$  in our models are statistically independent, both in space (for a fixed scale) and in scale. In [1], our assumption of whiteness in scale is justified, using a

theoretical model. We here validate our assumption of whiteness in space, by examining the sample correlation of the residuals that result when the appropriate autoregression is applied as a predictor to a multiresolution sequence of images.

An example of the data used in this validation for natural clutter is provided in Fig. 4. The left column of this figure displays a multiresolution sequence of three images of the region of grass we used for training. Proceeding downward, the images have resolution  $4\delta \times 4\delta$ ,  $2\delta \times 2\delta$  and  $\delta \times \delta$ , respectively. The right column of the figure displays images of prediction residuals; in keeping with Table I, the top image represents the residuals formed by the difference  $I(s) - 0.3I(s\bar{\gamma})$ , where  $I(s)$  is a pixel in the image having  $2\delta \times 2\delta$  resolution; the bottom image represents the residuals formed by the difference  $I(s) - 0.28I(s\bar{\gamma})$ , where  $I(s)$  is a pixel in the image having  $\delta \times \delta$  resolution.

At least visually, Fig. 4 suggests that the residuals are approximately uncorrelated. This is further confirmed by Fig. 5, which displays the sample correlation function of the residual image from the lower-right corner of Fig. 4. One can readily discern the impulse-like shape of this correlation function, which renders it in agreement with our model assumption. Although not shown, the same impulse-like shape is exhibited by the sample correlation functions of coarser scale residuals.

We now consider our man-made model. Proceeding in a manner parallel to our validation procedure for the natural-clutter model, we display in




Fig. 4. Images used to validate our natural-clutter model assumption that prediction residuals are white.




Fig. 6. Images used to validate our man-made model assumption that prediction residuals are white.




Fig. 5. Sample correlation function for residuals in prediction of  $\delta \times \delta$  resolution image of region of grass, using only  $2\delta \times 2\delta$  resolution image.

Fig. 6 a collection of five images, three of which are SAR images and two of which are residual images. We immediately see that the residuals do not have the same completely uncorrelated appearance we observed in Fig. 4. In this sense, the man-made model does not capture as completely the scale-to-scale statistical coupling of the multiresolution sequence of images. Nevertheless, as we see in Section IV, if we ignore this remaining correlation and apply the resulting likelihood calculation methods based on

the assumption of white residuals, we obtain a useful discrimination feature.

Finally, for both of our models, we can measure the correlation of residuals across different scales. We have empirically found that the peak correlation between residuals at different scales is roughly 0.2, which is quite modest. In general, we conclude that the correlation of measured residuals behaves in a manner consistent with our model assumption of uncorrelatedness, particularly in the case of our natural-clutter model.

### 3) Identification of the Residual Distributions:

To characterize the probability distributions of the prediction residuals  $w(s)$  for each model, we first calculate the sample cumulative distribution function (cdf) of the residuals associated with our training data (see Figs. 4 and 6). Then, we find a matching cdf that has a compact analytical form.

In Figs. 7(a) and (b), we plot both empirical cdfs and our analytical fits to them; the first figure displays the entire cdfs, while the second focuses exclusively on the upper tails. We first consider the cdfs shown for the residuals associated with our man-made model. The sample cdf (i.e., the dashed line) summarizes the aggregate statistics of the residuals in the prediction of 136 finest scale images of tactical targets. Each of these predictions is based on an autoregression

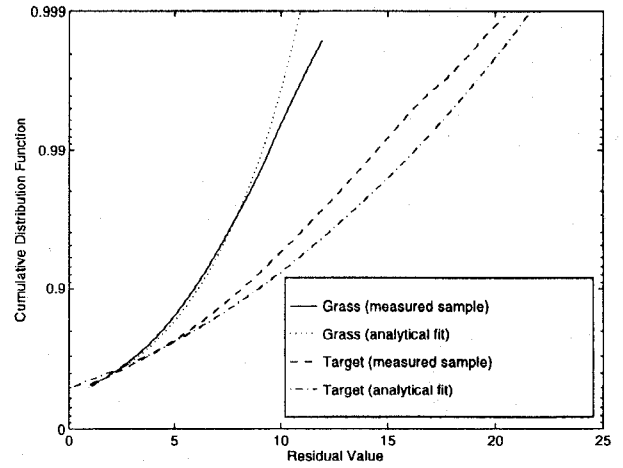



Fig. 7. (a) Sample cdfs of prediction residuals. (b) Upper tail of sample cdfs of prediction residuals.

using the coefficients listed in Table I, applied to the coarser scale images of the respective target. Our corresponding analytical fit (i.e., the dash-dot line) is with a Gaussian distribution, having mean zero and variance  $\sigma_{m(s)}^2$ , where this variance is chosen to match the sample variance associated with the training data. We note that the match between measurement and analytical fit is reasonably good. Although not shown in the figure, the same reasonably good match is also obtained for coarser scale residuals, where we continue to use a Gaussian fit.

Turning to our natural clutter model, the measured sample cdf (i.e., the solid curve in Figs. 7(a) and (b)) is based upon the residuals in the prediction of  $256 \times 256$  finest scale image of a homogeneous region of grass, using the  $128 \times 128$  second-finest scale image. The corresponding analytical fit (i.e., the dotted line) is with a zero-mean log-Rayleigh distribution,

$$P(w(s)) = k \exp[kw(s) - \gamma - \exp(kw(s) - \gamma)]$$

$$k = \frac{\ln 10}{10}, \quad (5)$$

$$\gamma \approx 0.57721566 \quad (\text{Euler's constant}).$$

4) *State-Space Representation of Models:* To facilitate the development of our multiresolution-based discriminant, we recast our final model choices in state-space form. The state  $x(s)$  is defined to be an  $R$ -dimensional vector, containing the pixel value  $I(s)$  together with the pixel values residing at the  $R - 1$  ancestors of node  $s$ . For our natural-clutter model,  $R = 1$  and for our man-made model,  $R = 2$ , and thus we have

$$x(s) = \begin{cases} I(s) & (\text{natural-clutter model}), \\ \begin{pmatrix} I(s) \\ I(s\bar{\gamma}) \end{pmatrix} & (\text{man-made model}). \end{cases} \quad (6)$$

To be consistent with (3), the scale-recursive dynamics for  $x(s)$  are thus defined to be

$$x(s) = \begin{cases} a_{1,m(s)}x(s\bar{\gamma}) + w(s) & (\text{natural-clutter model}), \\ \begin{pmatrix} a_{1,m(s)} & a_{2,m(s)} \\ 1 & 0 \end{pmatrix} x(s\bar{\gamma}) + \begin{pmatrix} 1 \\ 0 \end{pmatrix} w(s) & (\text{man-made model}). \end{cases} \quad (7)$$

In this recursion,  $w(s)$  is scalar valued and white (i.e.,  $w(s)$  is independent of  $w(\sigma)$ , for  $s \neq \sigma$ ). For the natural-clutter model, the distribution of  $w(s)$  is log-Rayleigh, given by (5), while for the man-made model, the distribution is Gaussian. The values of the model-dependent regression coefficients  $a_{i,m(s)}$  are given in Table I.

### C. Calculation of Multiresolution Discriminant

Our multiscale stochastic models implicitly specify the following two pdfs:

$$P(I_0, I_1, \dots, I_{L-2} | I_{L-1}, I_L, H_i) \quad (i = 1, 2) \quad (8)$$

where  $H_0(H_1)$  denotes the hypothesis that the ROI represents natural clutter (a man-made object). As we describe below, these conditional pdfs can be calculated efficiently. Thus, with their ready availability, we are led naturally to define our multiresolution discriminant  $\ell$  to be the logarithm of the likelihood ratio,

$$\ell \equiv \log \frac{P(I_0, I_1, \dots, I_{L-2} | H_1, I_{L-1}, I_L)}{P(I_0, I_1, \dots, I_{L-2} | H_0, I_{L-1}, I_L)}. \quad (9)$$

To describe the calculation of  $\ell$ , we first introduce some convenient notation. We let  $\mathbf{a}_{k,H_0}(\mathbf{a}_{k,H_1})$  denote the  $k$ th scale regression coefficients for the



natural-clutter model (man-made model). We let  $w_{H_i}(s)$  denote the residual in the autoregressive prediction of the pixel value  $I(s)$ , using the model underlying  $H_i$ ; in keeping with the conventions established in Section IIIB4,  $w_{H_i}(s)$  is given by

$$w_{H_i}(s) = I(s) - \mathbf{a}_{m(s), H_i}^T \mathbf{x}(s\bar{\gamma}). \quad (10)$$

Finally, we define  $P(w(s) | H_0)(P(w(s) | H_1)$  to be a log-Rayleigh distribution (5) (Gaussian distribution).

The multiresolution discriminant can now be expressed in the following way:

$$\ell = \left\{ \sum_{k=0}^{L-2} \sum_{\{s; m(s)=k\}} \log[P(w_{H_1}(s) | H_1)] \right\} - \left\{ \sum_{k=0}^{L-2} \sum_{\{s; m(s)=k\}} \log[P(w_{H_0}(s) | H_0)] \right\}. \quad (11)$$

This formula is justified in detail in Appendix A. Each summand here represents a penalty associated with a single residual. This penalty provides a quantitative measure of the mismatch between actual data and the predictive fit of our models to these data.

#### IV. PERFORMANCE OF THE DISCRIMINATION ALGORITHMS

In this section, we determine the detection performance improvement that results by incorporating our new multiresolution discriminant into the standard Lincoln Laboratory discrimination algorithm. In particular, we compare the performance of two discrimination algorithms: the first is the one described in Section IIB, and the second represents extension of the first, in which the feature set used to make the discrimination decision is augmented to include the multiresolution discriminant. The structure of the decision rule in both cases is identical; the only difference between the two algorithms is that our new multiresolution feature is available to the second.

##### A. SAR Imagery Used in Study

For our study, we have used actual imagery gathered with the Lincoln Laboratory millimeter-wave SAR. All of this imagery has 0.3 m resolution (in both range and cross-range) and has HH polarization. There are two components to this data set. The first is a training data set, used to build our two multiscale stochastic models and to estimate the parameters  $M_t$  and  $\Sigma_t$  associated with the conditional pdf  $P(Z | \text{target})$ . The second component is a testing data set, used to test the discrimination algorithms.

The training data set itself has two components. The first is a SAR image of a large, homogeneous

TABLE II

<i>ROI Generation step</i>	<i>Tactical targets</i>	<i>Total FAs</i>	<i>Man-made FAs</i>	<i>Natural FAs</i>
Prescreener	136	8739	2266	6473
Size filter	136	2071	849	1222

*Note:* This table summarizes the outcome of our two steps for generating a collection of ROIs that can be used to evaluate discrimination performance.

region of grass. This is used to build our natural-clutter multiscale model. The second is a collection of 136 SAR images, each representing an uncamouflaged tactical target. This collection is used both to build our man-made multiscale model, and to estimate the parameters  $M_t$  and  $\Sigma_t$  associated with  $P(Z | \text{target})$ .

The testing data set contains imagery representing 56 square kilometers of Stockbridge, New York. This imagery includes 136 tactical targets (i.e., 68 tanks and 68 howitzers) that are realistically deployed, *all* having a single layer of radar scattering netting; this netting reduces the radar cross section of the targets by roughly 5–6 dB. The testing imagery also includes a large number of man-made clutter objects, including powerline towers, a farmhouse, a golf course clubhouse, and a junkyard (complete with buildings, a crane and old military jeeps). Finally, there are natural clutter regions of trees, grass, and shrubs. In the sequel, we refer to this testing data set as the Stockbridge imagery.

##### B. Generation of ROIs

To generate a collection of ROIs for testing purposes, we proceed in two steps, using the Stockbridge imagery. In the first step, we apply the Lincoln Laboratory prescreening algorithm (see Section IIA) to the Stockbridge imagery. We adjust the sensitivity of this algorithm so that none of the 136 tactical targets are discarded. At this sensitivity level, the prescreening algorithm yields 136 ROIs representing tactical targets and an additional 8739 ROIs representing false alarms, both natural and man-made. Each ROI consists of  $128 \times 128$  pixels, corresponding to a region of approximately 38 square-meters. In the second step, we apply the size filter described in Section IIB to the ROIs generated by the prescreening algorithm. We recall that this size filter is used by both discriminators of interest. The filter has the effect of eliminating all ROIs whose principal object's diameter is not within the range of diameters we expect a tactical target to have. Again, we adjust the sensitivity of this filter so that none of the 136 ROIs representing tactical targets are discarded. Table II summarizes the outcome of the

TABLE III

$P_D$	Natural FAs		Man-made FAs		Total FAs	
	$S_{standard}$	$S_{mr-augmented}$	$S_{standard}$	$S_{mr-augmented}$	$S_{standard}$	$S_{mr-augmented}$
0.8	52	5	563	256	615	261
0.9	117	8	676	350	793	358
0.95	191	34	753	429	944	463
1.0	300	241	786	667	1086	908

Note: This table summarizes the discrimination performance at four particular operating points of the ROC curves in Fig. 8.

two steps we have just described, by categorizing the ROIs that are generated.

### C. Standard Lincoln Laboratory Discriminator Versus New Discriminator

We now subject the collection of remaining ROIs to a quadratic discriminator. In keeping with our objective to evaluate the effectiveness of the multiresolution discriminator, we consider two versions of the quadratic discriminator. The sole difference between the two is that each has a distinct set of features available for use in the decision process. The first set  $S_{standard}$  contains the features described in Section IIB; these are the ones traditionally available to the Lincoln discriminator. The second set  $S_{mr-augmented}$  is an augmented version of the first set, in which the multiresolution discriminator is added.

For each version of the discriminator, we evaluate the effectiveness of every possible combination of features, where these combinations are simply subsets of the available features. In this way, for each version of the discriminator, we search for the feature combination that results in the smallest number of nontarget ROIs being classified as targets, subject to the constraint that all 136 of the ROIs actually containing tactical targets are correctly classified.

The resulting optimal features, corresponding to  $S_{standard}$  are standard deviation, fractal dimension, peak CFAR, and percent bright CFAR. The optimal features corresponding to  $S_{mr-augmented}$  are peak CFAR, mean CFAR, new multiresolution discriminant. The corresponding discrimination results are summarized by the receiver operating characteristic (ROC) curves shown in Fig. 8 and in Table III.

These results merit a couple of comments. First, we note that the performance of our new discriminator, using the optimal features in the set  $S_{mr-augmented}$ , represents an improvement over the standard Lincoln discriminator. For example, at the operating point  $P_D = 0.95$ , the new discriminator reduces the number of natural-clutter false alarms by almost a factor of six. Second, we note that the optimal feature combination corresponding to the set




Fig. 8. ROC curves summarizing performance of discrimination algorithms when applied to HH-polarization SAR imagery representing 56 square kilometers of Stockbridge, New York.

$S_{mr-augmented}$  is in a certain sense consistent with the optimal feature combination corresponding to the set  $S_{standard}$ . To clarify this comment, we recall that the features *standard deviation* and *fractal dimension* are part of the optimal combination corresponding to  $S_{standard}$ . But our new multiresolution discriminant essentially captures both of these characteristics: standard deviation information is directly captured in the structure of the log-likelihood ratio, as expressed in (11), and fractal characteristics are fundamental to the structure of our multiresolution models and processes, as discussed in [8]. This observation is reinforced by the features in the optimal combination corresponding to  $S_{mr-augmented}$ , where use of the new multiresolution discriminant essentially supersedes joint use of *standard deviation* and *fractal dimension*.

## V. CONCLUSION

We have developed a new algorithm for discriminating man-made objects from natural clutter in SAR imagery. This algorithm uses an autoregression in scale to capture and exploit the

characteristically distinct variations in speckle pattern for imagery of man-made objects and of natural clutter, as image resolution is varied from coarse to fine. We tested the algorithm by applying it to a dataset of 0.3 m resolution, HH polarization SAR imagery, representing 56 square kilometers of terrain. At a probability of detection of 0.95, the number of natural-clutter false alarms generated by the new discriminator was one-sixth the number generated by the standard Lincoln discriminator.

Our results suggest a number of interesting possibilities to pursue in future work. For one, there are connections between the multiresolution approach taken here and the one taken in [1]; the two works could be combined and extended. For example, theoretical models are developed in [1] that could possibly be used to either validate or refine the multiresolution models we developed here. Furthermore, certain optimal resolutions are identified and exploited in [1]; perhaps we could use these optimal scales in lieu of the dyadic progression of scales used here.

A second possible extension of the work here would be to develop more sophisticated multiscale models. Recall, for example, the residuals in Fig. 6 resulting from application of our man-made model to a multiresolution image sequence. These residuals demonstrate that the man-made model is not completely capturing the scale-to-scale statistical coupling of the images; there is a need for better accounting of the dominant scatterers in the images. Perhaps this need could be fulfilled by developing a whole collection of man-made models, with each individual model specialized to a particular target configuration. Given this collection of models, we could then envision carrying out likelihood-based target recognition.

Another possible extension would be to exploit the multiresolution models to carry out image compression, in a manner analogous to the use of linear predictive coding for time-series (e.g., speech) compression. Finally, there is the possibility of developing multiscale models for remote sensing applications, such as classification of terrain cover, for which we could develop a number of natural-clutter models, including one for trees, another for grass, and so forth.

#### APPENDIX A. DERIVATION OF EXPRESSION FOR MULTIREOLUTION DISCRIMINANT

In Section IIIC, we described a straightforward procedure for calculating our multiresolution discriminant. Here we establish the validity of the decomposition (11) of that discriminant. We continue to use the notation that was established in Section IIIC and additionally we define  $X_k$  to be a vector containing all the state vectors  $x(s)$  at scale  $k$ .

The key to our development here is to exploit the *Markov* property of our multiscale stochastic processes. To describe this property, we first note that any given node on a quadtree can be viewed as a boundary between five subsets of nodes, where four of these subsets correspond to paths leading towards offspring and one corresponds to a path leading towards the parent.<sup>2</sup> With this boundary notion in mind, the Markov property can be stated as follows: conditioned on the value of the state at any node, the values of the states in the corresponding five subsets of nodes extending away from  $s$  are independent. Hence, in light of (6),

$$P(I_k | X_{k+1}, H_i) = \prod_{\{s; m(s)=k\}} P(I(s) | x(s\bar{\gamma}), H_i). \quad (12)$$

To establish (11), we begin by examining in greater detail the numerator and denominator terms of (9); by elementary probability,

$$\begin{aligned} P(I_0, I_1, \dots, I_{L-2} | I_{L-1}, I_L, H_i) \\ = P(I_{L-2} | I_{L-1}, I_L, H_i) \prod_{k=3}^L P(I_{L-k} | I_{L-k+1}, \dots, I_L, H_i). \end{aligned} \quad (13)$$

To simplify the right side of (13), we note that by construction (see Section IIIB4) of our state-space models, the conditioning information contained in the following two sets

$$\begin{aligned} \{I_{k+1}, I_{k+2}, \dots, I_L, H_i\} \quad \text{and} \\ \{X_{k+1}, X_{k+2}, \dots, X_{L-1}, H_i\} \end{aligned}$$

is equivalent with respect to the statistics of  $I_k$ . By combining this fact with the Markov property of the multiscale model class, we find that each element of the product on the right side of (13) can be factored as

$$\begin{aligned} P(I_k | I_{k+1}, \dots, I_L, H_i) &= P(I_k | X_{k+1}, \dots, X_{L-1}, H_i) \\ &= P(I_k | X_{k+1}, H_i), \end{aligned} \quad (14)$$

and thus, combining (12), (13), and (14),

$$\begin{aligned} P(I_0, I_1, \dots, I_{L-2} | I_{L-1}, I_L, H_i) \\ = \prod_{k=3}^L \prod_{\{s; m(s)=L-k\}} P(I(s) | x(s\bar{\gamma}), H_i). \end{aligned} \quad (15)$$

Finally, combining (9), (15), and the fact that


$$P(I(s) | x(s\bar{\gamma}), H_i) = P(w_{H_i}(s) | H_i)$$

we conclude that (11) is valid.

<sup>2</sup>The root node is an exception, having only  $q$  offspring and no parent. Also, the finest scale nodes are exceptions, each having a single parent, but no offspring.

## REFERENCES

- [1] Subotic, N. S., Collins, L. M., Gorman, J. D., and Thelen, B. J. (1994)  
A multiresolution approach to target detection in synthetic aperture radar data.  
Presented at the 1994 28th Asilomar Conference, Monterey, CA, Oct. 1994.
- [2] 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* (Special Issue on Automatic Target Recognition), **6**, 1 (Spring 1993), 11–24.
- [3] Kreithen, D. E., Halversen, S. D., and Owirka, G. J. (1993)  
Discriminating targets from clutter.  
*Lincoln Laboratory Journal* (Special Issue on Automatic Target Recognition), **6**, 1 (Spring 1993), 25–52.
- [4] Goldstein, G. B. (1973)  
False alarm regulation in log normal and Weibull clutter.  
*IEEE Transactions on Aerospace and Electronic Systems*, **AES-16** (Jan. 1973), 84–92.
- [5] Fukunaga, K., Hayes, R. R., and Novak, L. M. (1987)  
The acquisition probability for a minimum distance one-class classifier.  
*IEEE Transactions on Aerospace and Electronic Systems*, **23**, 4 (July 1987), 493–499.
- [6] Henry, J. C. (1991)  
The Lincoln Laboratory 35 GHz airborne polarimetric SAR imaging system.  
In *Proceedings of IEEE National Telesystems Conference*, Atlanta, GA, Mar. 26–27, 1991, p. 353.
- [7] Ljung, L. (1987)  
*System Identification: Theory for the User*.  
Englewood Cliffs, NJ: Prentice-Hall, 1987.
- [8] Basseville, M., Benveniste, A., Chou, K., Golden, S., Nikoukhah, R., and Willsky, A. (1992)  
Modeling and estimation of multiresolution stochastic processes.  
*IEEE Transactions on Information Theory*, **38** (Mar. 1992), 766–784.



**William Irving** received the S.B. in 1987, the S.M. in 1991, and the Ph.D. in 1995, in electrical engineering and computer science, all from M.I.T., Cambridge, MA. His Ph.D. research at M.I.T. was in the area of multiresolution statistical signal and image processing, where he developed a realization theory for a recently developed class of multiscale stochastic processes.

Prior to joining Alphatech in September, 1995, Dr. Irving was a member of the technical staff at M.I.T., Lincoln Laboratory. From 1987 to 1989, he worked full time at the Laboratory. Then, from 1989 to 1995, he participated in the Lincoln Laboratory as Staff Associate, attending graduate school at M.I.T. full time during the academic year and working at Lincoln during the summers. Currently, Dr. Irving is developing multiresolution algorithms for SAR ATR under DARPA sponsorship, and is developing a theory for SAR ATR performance as part of Wright Laboratory's MSTAR program. Also, during the spring of 1996, he taught an advanced graduate course on recursive estimation theory in the Department of Electrical Engineering and Computer Science at M.I.T.



**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.

**Alan S. Willsky** (S'70—M'73—SM'82—F'86) received both the S.B. degree and the Ph.D. degree from the Massachusetts Institute of Technology (M.I.T.), Cambridge, MA, in 1969 and 1973, respectively.

He joined the M.I.T. faculty in 1973 and his present position is Professor of Electrical Engineering. From 1974 to 1981 Dr. Willsky served as Assistant Director of the M.I.T. Laboratory for Information and Decision Systems. He is also a founder and member of the board of directors of Alphatech, Inc. Dr. Willsky has held visiting positions at Imperial College, London, L'Université de Paris-Sud, and the Institut de Recherche en Informatique et Systèmes Aléatoires in Rennes, France. His present research interests are in problems involving multidimensional and multiresolution estimation and imaging, discrete-event systems, and the asymptotic analysis of control and estimation systems.

In 1975 Dr. Willsky received the Donald P. Eckman Award from the American Automatic Control Council. He was program chairman for the 17th IEEE Conference on Decision and Control, has been an Associate Editor of several journals including the *IEEE Transactions on Automatic Control*, has served as a member of the Board of Governors and Vice President for Technical Affairs of the IEEE Control Systems Society, was program chairman for the 1981 Bilateral Seminar on Control Systems held in the People's Republic of China, and was special guest editor of the 1992 special issue of the *IEEE Transactions on Information Theory* on wavelet transforms and multiresolution signal analysis. Also in 1988 he was made a Distinguished Member of the IEEE Control Systems Society. In addition Dr. Willsky has given several plenary lectures at major scientific meetings including the 20th IEEE Conference on Decision and Control, the 1991 IEEE International Conference on Systems Engineering, the SIAM Conference on Applied Linear Algebra, 1991, and the 1992 Inaugural Workshop for the National Centre for Robust and Adaptive Systems, Canberra, Australia. He is the author of the research monograph *Digital Signal Processing and Control and Estimation Theory* and is co-author of the undergraduate text *Signals and Systems*. He was awarded the 1979 Alfred Noble Prize by the ASCE and the 1980 Browder J. Thompson Memorial Prize Award by the IEEE for a paper excerpted from his monograph.

