L.M. Novak, M.C. Burl, R.D. Chaney, and G.J. O[redacted]

Optimal Processing of Polarimetric Synthetic-Aperture Radar Imagery

The Advanced Detection Technology Sensor can detect, discriminate, and classify stationary ground targets—during the day or night—even through cloud cover, fog, smoke, dust, or rain. The sensor is a coherent, fully polarimetric, 35-GHz synthetic-aperture radar (SAR) with a resolution of 1 ft × 1 ft. And, to minimize SAR speckle while preserving image resolution, it uses the polarimetric whitening filter, our recently developed method for processing fully polarimetric data into SAR imagery.

The Advanced Detection Technology Sensor (ADTS) is a fully polarimetric, 35-GHz synthetic-aperture radar. The polarimetric capability of the radar is used to enhance the quality of the imagery taken from a small aircraft; the synthetic aperture permits data to be processed to a resolution of 1 ft by 1 ft at a slant range of 7 km. The sensor was developed to provide a high-quality database of clutter and target-in-clutter imagery, and to perform stationary-target detection, discrimination, and identification.

An example of the quality of ADTS images is presented in Fig. 1. This synthetic aperture radar (SAR) image of a farmhouse in Stockbridge, N.Y., has undergone the optimal speckle-reduction processing described in this article. Because of its high resolution, the ADTS can resolve individual trees and bushes, as well as the house itself. This image was obtained under clear weather conditions. However, the quality and resolution of the ADTS SAR would not be degraded in the presence of dense fog or thick cloud cover. Thus the ADTS sensor has a significant advantage over optical sensors: the image quality is not dependent on weather conditions, and the sensor can be used either during the day or at night.

Figure 2(a) shows the ADTS and lists some of its specifications. The ADTS is an airborne, instrumentation-quality radar carried on a Gulfstream G1 aircraft (Fig. 2[b]). One reason that this aircraft was selected was its low operating costs. The radar antenna is housed in a specially designed radome mounted beneath the aircraft. The radar transmitter, dual-channel receiver, and digital recording equipment are all carried by the airplane

Since one of the principal objectives of the Advanced Detection Technology program is to evaluate the benefits of fully polarimetric radar data for stationary-target detection, polarization purity is essential. To achieve polarization purity, a corrugated horn antenna with a Fresnel lens was designed, providing very pure horizontally and vertically polarized transmit waves. The radome was designed to minimize cross coupling between the horizontal and vertical polarizations. The radar transmits horizontal and vertical polarizations on alternate pulses; dual receiver channels measure both returns simultaneously. Inertial velocity estimates compensate for aircraft motion between the horizontal and vertical transmit pulses.

An in-scene calibration array—comprised of several high-quality trihedrals and dihedrals oriented at 0°, 22.5°, and 45°—is used for polarimetrically calibrating the imagery. The polarimetric calibration scheme is described in Ref. 1.

During each flight, data are gathered and digitized in real time with a 28-channel Ampex recorder. The data are then brought to the Lincoln Laboratory ground processing facility, where SAR image formation is performed. Special-purpose, high-speed digital processing hardware is used to construct the imagery and perform the polarimetric calibration.

A SAR is a radar that synthesizes a long

Fig. 1—ADTS SAR image of farmhouse in Stockbridge, N.Y. (1-ft × 1-ft resolution). The sensor was flown at an altitude of 2 km with a look-down (depression) angle of 22.5°, giving a slant range of 7 km. PWF processing was used to produce this minimum-speckle image. The radar is located at the top of the image, looking down; therefore, the radar shadows go toward the bottom of the page.

aperture as an aircraft flies along its path. Thus a SAR can achieve cross-range resolutions that could otherwise be attained only with a long antenna. In SAR mode the ADTS has 1-ft × 1-ft resolution. Range resolution is achieved by using 600-MHz bandwidth pulses. To achieve 1-ft azimuth resolution, a synthetic aperture of approximately 150-m length is constructed by processing 1 sec of data as the plane flies.

SAR processing can produce high-resolution images, but the process is subject to a considerable amount of speckle in the images because of the coherent nature of the image process. Figure 3 presents a side-by-side comparison of a single-channel image and of the same image after polarimetric whitening filter (PWF) processing. Note that the PWF processing has reduced the image speckle significantly, making the dirt roads and other features more visible. The reduction of image speckle is essential for good target detection, discrimination, and classification. Reducing image speckle can

also improve the performance of image segmentation algorithms, as will be discussed later in this article.

Noncoherent spatial averaging of high-resolution pixel intensities can reduce image speckle. We have significantly reduced image speckle by averaging 4-pixel × 4-pixel clusters from ADTS data into effective 1-m × 1-m resolution pixels. However, the speckle reduction was obtained at the cost of degraded image resolution.

A new technique, the polarimetric whitening filter, uses a polarimetric method of speckle reduction that preserves image resolution [2, 3]. This algorithm processes the complex (HH, HV, VV) data into full-resolution pixel intensities in a way that minimizes SAR image speckle. We have recently begun testing the PWF on actual polarimetrically calibrated ADTS data. And, at 1-ft resolution, the PWF has reduced the clutter log standard deviation by 2.5 dB relative to the single-channel system.

It is possible to combine the two methods of speckle reduction, by first reducing the speckle polarimetrically while preserving the 1-ft × 1-ft resolution, and then using noncoherent spatial averaging of the 1-ft × 1-ft pixel intensities into 1-m × 1-m resolution pixels. This method yielded a clutter log standard deviation 0.5 dB lower than noncoherent averaging of HV data and 0.7 dB lower than noncoherent averaging of HH data.

(a)

(b)

System Parameters	
Frequency	33 GHz
Resolution (SAR)	1 ft × 1 ft
Beamwidth	2°
Polarization Isolation	1000/1
Sensitivity (SAR Mode) (7-km Range)	Signal to Noise = 10 for 1/1000 m^2 Radar Cross Section

Fig. 2—The Advanced Detection Technology Sensor. (a) The sensor platform is a Gulfstream G1 aircraft, shown in flight. (b) The ADTS radome, located at the bottom of the aircraft, was built by the Loral Corp.

Polarimetric Clutter Model

Speckle has long been recognized as a fundamental problem in SAR imaging systems. Speckle reduction improves the visual quality of SAR imagery, and it permits the application of sophisticated image-processing and machine-vision algorithms to the SAR-image interpretation problem. Although spatial, noncoherent averaging of high-resolution pixel intensities reduces speckle, it does so at the expense of image resolution. But with the recent availability of fully polarimetric SAR data, it is possible to reduce speckle polarimetrically while preserving image resolution.

We have, in fact, developed an optimal method of processing polarimetric radar data into pixel intensity that minimizes image speckle. This method is based on a mathematical model that characterizes fully polarimetric radar returns from clutter. By using this pola-

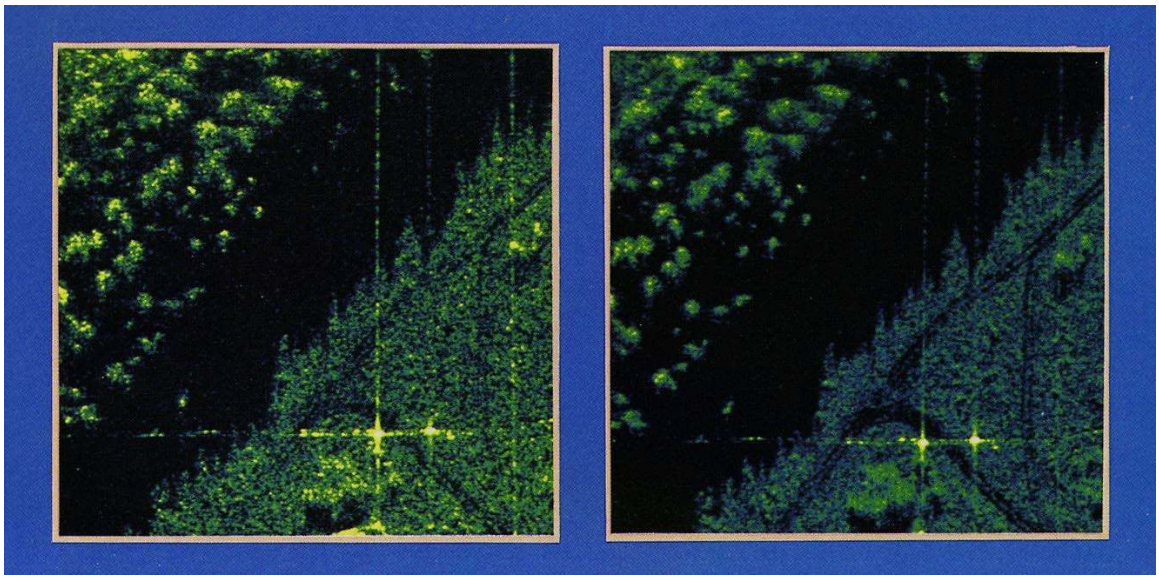

(a) (b)

Fig. 3—(a) HH image of trees and a meadow with several dirt roads. (b) The same image after PWF processing. The bright spots in the center of each image are registration reflectors. The dirt roads, which are virtually invisible in the single-channel image, can be clearly seen in the PWF image.

rimetric clutter model, we can derive an algorithm that shows how fully polarimetric data can be combined into minimum-speckle imagery.

In this article a non-Gaussian product model is used to characterize clutter. Note that with a Gaussian model, each pixel of clutter in a spatially homogeneous region of an image has the same average polarimetric power. A number of authors have stated that it is more realistic to assume that ground clutter and sea clutter, for example, are spatially nonhomogeneous. A non-Gaussian model consistent with this more realistic assumption has been proposed [4–6], and, in fact, the Gaussian model is actually a special case of the non-Gaussian model.

We assume that the radar measurement vector $\mathbf{Y}$ consists of three complex elements: HH, HV, and VV. Therefore,

$$\mathbf{Y} = \begin{pmatrix} \mathrm{HH} \\ \mathrm{HV} \\ \mathrm{VV} \end{pmatrix} = \begin{pmatrix} \mathrm{HH_I} + j\mathrm{HH_Q} \\ \mathrm{HV_I} + j\mathrm{HV_Q} \\ \mathrm{VV_I} + j\mathrm{VV_Q} \end{pmatrix}$$

where $\mathrm{HH_I}$ and $\mathrm{HH_Q}$, for example, are the in-phase and quadrature components of the complex HH measurement. $\mathbf{Y}$ is assumed to be the product of a complex Gaussian vector $\mathbf{X}$ (representing the speckle) and a spatially varying texture variable $\sqrt{g}$. That is,

$$\mathbf{Y} = \sqrt{g}\,\mathbf{X}.$$

The joint probability density function (PDF) of the complex vector $\mathbf{X}$ is given by the expression

$$\mathrm{f}(\mathbf{X}) = \frac{1}{\pi^3 |\Sigma|} \exp\left(-\mathbf{X}^\dagger \Sigma^{-1} \mathbf{X}\right)$$

where $\Sigma = \mathrm{E}\left(\mathbf{X}\mathbf{X}^\dagger\right)$ is the polarization covariance matrix. The vector $\mathbf{X}$ is zero mean $\left[\mathrm{E}\,(\mathbf{X}) = 0\right]$. The covariance matrices that we use for clutter data take the following form (in linear-polarization basis):

$$\Sigma = \sigma_{\mathrm{HH}} \cdot \begin{pmatrix} 1 & 0 & \rho\sqrt{\gamma} \\ 0 & \varepsilon & 0 \\ \rho^*\sqrt{\gamma} & 0 & \gamma \end{pmatrix} \qquad (1)$$

where

$$\sigma_{\mathrm{HH}} = \mathrm{E}\left(|\mathrm{HH}|^2\right),\ \varepsilon = \frac{\mathrm{E}\left(|\mathrm{HV}|^2\right)}{\mathrm{E}\left(|\mathrm{HH}|^2\right)},\ \gamma = \frac{\mathrm{E}\left(|\mathrm{VV}|^2\right)}{\mathrm{E}\left(|\mathrm{HH}|^2\right)},$$

$$\text{and}\quad \rho = \frac{\mathrm{E}\left(\mathrm{HH} \cdot \mathrm{VV}^*\right)}{\left[\mathrm{E}\left(|\mathrm{HH}|^2\right) \cdot \mathrm{E}\left(|\mathrm{VV}|^2\right)\right]^{\frac{1}{2}}}.$$

We model the product multiplier g as a gamma-distributed random variable. This assumption is not universal; the log-normal and Weibull models are also widely used. But if we assume that the gamma distribution is reasonable, the PDF of the product multiplier g is specified by

$$f_G(g) = \frac{1}{\bar{g}}\left(\frac{g}{\bar{g}}\right)^{(\upsilon-1)} \frac{1}{\Gamma(\upsilon)} \exp\left(\frac{-g}{\bar{g}}\right)$$

where the parameters $\bar{g}$ and υ are related to the mean and variance of the random variable g:

$$\begin{aligned} E(g) &= \bar{g}\upsilon \\ E\left(g^2\right) &= \bar{g}^2\upsilon(\upsilon+1). \end{aligned} \tag{2}$$

With the assumption that g is gamma distributed, the PDF of the resulting vector $\mathbf{Y} = \sqrt{g}\mathbf{X}$ is the modified Bessel function, or generalized K-distribution, given in Ref. 4.

$$f(\mathbf{Y}) = \frac{2}{\pi^3 \bar{g}^{\upsilon}\Gamma(\upsilon)|\mathbf{\Sigma}|} \cdot \frac{K_{3-\upsilon}\left(2\sqrt{\frac{\mathbf{Y}^\dagger\mathbf{\Sigma}^{-1}\mathbf{Y}}{\bar{g}}}\right)}{\left(\bar{g}\mathbf{Y}\,\mathbf{\Sigma}^{-1}\mathbf{Y}\right)^{(3-\upsilon)/2}}.$$

If we set $\bar{g} = \frac{1}{\upsilon}$ so that the mean of the texture variable is unity, then in the limit as $\upsilon \to \infty$, this model reduces to the Gaussian model.

Minimum-Speckle Image Processing

In this section we consider how to process the three complex measurements—HH, HV, and VV (i.e., the vector **Y**)—into pixel intensity in a way that minimizes speckle. The clutter product model is used to derive the optimal method of polarimetric speckle reduction, which can be interpreted as a PWF. Then the amount of speckle reduction that can be achieved by using the PWF is determined theoretically.

The measure of speckle we use is the ratio of the standard deviation of the image pixel intensities to the mean of the intensities (s/m):

$$\frac{s}{m} = \frac{\text{st. dev. } (y)}{\text{mean } (y)} \tag{3}$$

where the random variable y denotes pixel intensity. Given the measurements HH, HV, and VV, we wish to construct an image from the quadratic

$$y = \mathbf{Y}^\dagger\mathbf{A}\mathbf{Y} = g\mathbf{X}^\dagger\mathbf{A}\mathbf{X} \tag{4}$$

where **A** is a weighting matrix that is assumed to be Hermitian symmetric and positive definite, thus keeping y positive. To find the optimal weighting matrix $\mathbf{A}^*$ (i.e., the one that results in an image whose pixel intensities have the minimum possible s/m), we use the following results:

$$E\left(\mathbf{X}^\dagger\mathbf{A}\mathbf{X}\right) = \text{tr}\,(\mathbf{\Sigma}\cdot\mathbf{A}) = \sum_{i=1}^{3}\lambda_i \tag{5}$$

$$\text{VAR}\left(\mathbf{X}^\dagger\mathbf{A}\mathbf{X}\right) = \text{tr}\,(\mathbf{\Sigma}\cdot\mathbf{A})^2 = \sum_{i=1}^{3}\lambda_i^2 \tag{6}$$

where E is the expected value, tr is the trace, VAR is the variance, and λ_1, λ_2, and λ_3 are the eigenvalues of the matrix $\Sigma\cdot\mathbf{A}$. Combining Eqs. 2 through 6 yields

$$\left(\frac{s}{m}\right)^2 = \left[\frac{\text{VAR}\,(y)}{E^2(y)}\right] \times \left[\frac{E\left(g^2\right)}{E^2(g)} \cdot \frac{\text{VAR}\left(\mathbf{X}^\dagger\mathbf{A}\mathbf{X}\right)}{E^2\left(\mathbf{X}^\dagger\mathbf{A}\mathbf{X}\right)} + \frac{\text{VAR}(g)}{E^2(g)}\right]$$

$$= \left[\frac{\upsilon+1}{\upsilon} \cdot \frac{\sum_{i=1}^{3}\lambda_i^2}{\left(\sum_{i=1}^{3}\lambda_i\right)^2} + \frac{1}{\upsilon}\right]$$

By using the approach of Cadzow [7], we can show that matrix $\mathbf{A}^*$ must make the eigenvalues of $\mathbf{\Sigma}\cdot\mathbf{A}$ all equal

$$\lambda_1 = \lambda_2 = \lambda_3 .$$

Therefore, the optimal solution is

$$\mathbf{A}^* = \mathbf{\Sigma}^{-1}.$$

Note that any constant multiple of $\mathbf{A}^*$ is also optimal, because the s/m is invariant to scale change.

This solution is equivalent to applying a whitening filter to the polarimetric vector **Y** prior to forming the image. The vector **Y** is passed through the whitening filter $\mathbf{\Sigma}^{-1/2}$ to obtain

$$\mathbf{W} = \mathbf{\Sigma}^{-\frac{1}{2}}\mathbf{Y} = \sqrt{g}\,\mathbf{\Sigma}^{-\frac{1}{2}}\mathbf{X}.$$

The elements of **W** are independent, complex Gaussian random variables with equal expected power. The covariance of **W** is a scaled identity matrix; thus **W** is said to be white. Now the optimal solution of the speckle-reduction problem is simply the noncoherent sum of powers in the elements of **W**:

$$y = \mathbf{W}^{\dagger}\mathbf{W},$$

hence the name—polarimetric whitening filter.

Observe that, as illustrated in Fig. 4, this process is simply a change of polarimetric basis from linear polarization (HH, HV, VV) to a new basis given by

$$\left[\mathrm{HH},\ \frac{\mathrm{HV}}{\sqrt{\varepsilon}},\ \frac{\left(\mathrm{VV} - \rho^{*}\sqrt{\gamma}\mathrm{HH}\right)}{\sqrt{\gamma\left(1-|\rho|^{2}\right)}}\right].$$

In the new basis, the three polarimetric channels are uncorrelated and have equal expected power. Thus the optimal way to reduce speckle polarimetrically is to sum the powers noncoherently in these three polarimetric channels.

We have shown that the PWF processes the polarimetric measurement vector **Y** in a way that minimizes SAR image speckle. Furthermore, the PWF is the maximum-likelihood estimate (MLE) of the spatial multiplier *g*. The MLE is also an unbiased, minimum-variance estimator of *g;* i.e., it achieves the Cramer-Rao lower bound (see Appendix 1 for proofs).

Next we shall determine the amount of speckle reduction that can be achieved by using the PWF. Although the PWF solution is independent of the PDF of the spatial multiplier *g* in the product model, the resulting s/m after speckle reduction does depend upon $f_G(g)$. Thus the s/m for the PWF is

$$\left(\frac{s}{m}\right)_{\mathrm{PWF}} = \sqrt{\frac{1}{3}\left(1+\frac{4}{\upsilon}\right)} = \sqrt{\frac{1}{\upsilon}+\frac{1}{3}\left(1+\frac{1}{\upsilon}\right)} \quad (7)$$

and the s/m for a single polarimetric channel image is

$$\left(\frac{s}{m}\right)_{\mathrm{HH}} = \sqrt{1+\frac{2}{\upsilon}} = \sqrt{\frac{1}{\upsilon}+\left(1+\frac{1}{\upsilon}\right)}. \quad (8)$$

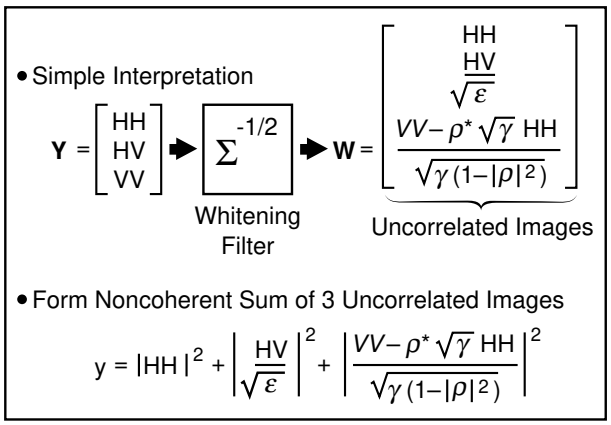

Fig. 4—Minimum-speckle image processing. Y is a complex vector containing the three linear-polarization measurements. Using the whitening filter gives a new polarization basis W. In this basis the three elements are uncorrelated and have equal expected power. The PWF image is then y, the noncoherent sum of the uncorrelated images.

The υ parameter of the gamma multiplier appears in Eqs. 7 and 8 because the s/m includes fluctuations in the texture variable *g*. For the ideal speckle-free image, in fact, fluctuations in the terrain reflectivity across the image are still present, so the s/m is given by

$$\left(\frac{s}{m}\right)_{\mathrm{ideal}} = \sqrt{\frac{1}{\upsilon}}.$$

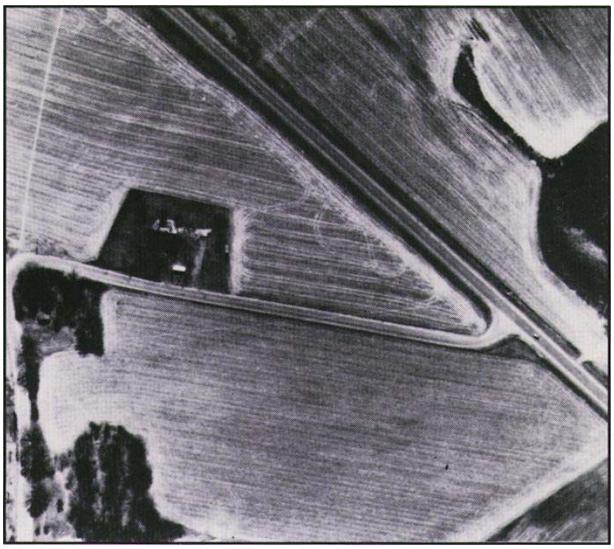

Fig. 5—Optical photograph of farm area near Stockbridge, N.Y. This aerial photograph was taken in April 1989. Some visible features include furrows in plowed fields, farmhouse (left of center), tree-lined river edge (right), tree area (bottom left), road intersection (right of center). Note that, unlike radar imagery, aerial photography of this quality can only be taken under clear conditions.

The υ parameter is closely related to the log standard deviation (σ_C in dB) of the texture component of the clutter. This relation is derived in Appendix 2, and values of the υ parameter for clutter regions that have σ_C of 1 dB to 3 dB are tabulated. By using Eqs. 7 and 8 and the results given in Appendix 2, we can calculate the reduction in the standard-deviation-to-mean ratio achieved with the PWF (relative to a single-polarimetric-channel image). For clutter with a spatial log standard deviation of 1 dB, the s/m ratio of single-channel data is 1.66 times larger (i.e., 4.4 dB larger) than the s/m ratio of PWF data. For clutter with a spatial log standard deviation of 3 dB, the s/m ratio of single-channel data is 1.45 times larger (i.e., 3.2 dB larger) than that of PWF data. In the next section, these theoretical predictions will be compared with actual measurements made on ADTS data.

Speckle-Reduction Results

In the preceding section, we determined that the optimal polarimetric processing for speckle reduction is the PWF, and we derived formulas to calculate the amount of speckle reduction achievable by using the PWF. In this section we show ADTS imagery of a farm located near Stockbridge, N.Y. Then we use polarimetrically calibrated ADTS clutter data to calculate polarization covariances of trees, grass, mixed scrub, and shadows. Finally, we apply PWF processing to these clutter data and compare the actual amount of speckle reduction with the theoretical prediction.

Figure 5 shows a photograph of the Stockbridge farm that was imaged by the ADTS. The corresponding SAR image of the farm is shown in Fig. 6. The SAR image was constructed by first reducing speckle polarimetrically (using PWF processing while preserving the 1-ft × 1-ft resolution) and then using noncoherent spatial averaging of the 1-ft × 1-ft PWF pixel intensities into effective 1-m by 1-m resolution pixels.

The area of the image shown in Fig. 6 is approximately 500 m by 500 m. Clearly visible in the SAR image are several roads, plowed fields, a farmhouse surrounded by several trees, and a larger set of trees located below the farmhouse. Note that, although the optical photograph gives an excellent image under conditions of good visibility, only the SAR image would be unaffected by such phenomena as rain, cloud cover, or fog.

Higher-resolution PWF images (512 ft × 512 ft at 1-ft × 1-ft resolution) of the farmhouse area and tree area are shown in Figs. 1 and 7. In general PWF processing enhances the imagery in two ways: (1) the amount of speckle in the imagery is reduced, and (2) the edges of roads, fields, and other objects are more clear. The improvement in the sharpness of edges is due to the use of polarimetric measurements—the HV and VV polarizations fill in edges and areas of objects that do not show up in the HH polarization.

Polarization-Covariance Calculations

Four types of clutter regions were considered in these studies: shadows, grass, mixed scrub, and trees. Each region contained 2000 to 8000 pixels, representing an area of several hundred square meters. We evaluated the clutter polarization-covariance parameters σ_{HH}, ε, γ, and ρ as defined in Eq. 1. The results are given in Table 1.

Since the standard-deviation-to-mean ratio is invariant with respect to scale, and since the normalized polarization covariance parameters (ε, γ, ρ) of trees, scrub, and grass were found to be very similar (Table 1), a single clutter polarization covariance could be used to whiten the entire image without sacrificing speckle-reduction performance significantly. The images in this article were all created with a single PWF, which used only the grass polarization-covariance parameters.

The polarization-covariance parameters of targets are quite different from those of trees, grass, and other types of clutter. The effect of PWF processing on target detection was theoretically analyzed in Ref. 8, where it was shown that the performance of a PWF detector was comparable to that of an optimal polarimetric detector.

Fig. 6—SAR image of farm area (1-m × 1-m resolution) corresponding to aerial photograph shown in Fig. 5. This image was formed by first applying PWF processing to the 1-ft coherent data, then spoiling (4 × 4 noncoherent averaging) to an effective 1-m resolution.

PWF Speckle Reduction

To validate the speckle-reduction formulas, we selected four clutter regions—shadow, grass, mixed scrub, and trees. For each region we calculated the standard-deviation-to-mean ratios for the HH, HV, and VV data, and for the PWF data. The results are given in Table 2, and they show lower—and better—numbers for the PWF approach.

Notice that the s/m ratios for trees are larger than for mixed scrub, and the s/m ratios for mixed scrub are larger than for grass. This effect occurs because the s/m ratio depends not only on speckle, but also on terrain roughness.

The single-polarimetric-channel s/m ratio given in Table 2 can be used to compute the rms standard-deviation-to-mean ratio for each clutter region. Equation 8 was used to estimate the approximate υ for each region, and Equation 7 was used to predict the s/m ratio of the PWF data. In Table 3, the theoretical predictions are compared with the actual measured values, and the agreement between the theory and the measurements turns out to be very good—within 5% in all cases.

Although the s/m figures given in Table 2 clearly show that the PWF reduces SAR speckle, the more important question is whether the clutter log standard deviation has a correspond-

ing decrease, because the log standard deviation directly affects target-detection performance. Shown in Table 4 are the log standard deviations computed from the 1-ft-resolution HH, HV, and VV data, and the 1-ft PWF data.

The PWF reduces the log standard deviation by approximately 2.0 to 2.7 dB over single-channel 1-ft data.

PWF versus Adaptive PWF

In Ref. 2 an adaptive PWF was proposed and analyzed. This algorithm adaptively estimates polarization covariances of various regions of clutter over the image and uses these same estimated covariances to minimize the speckle within each clutter region. Since the polarization covariances of grass, trees, and shadow regions were found to be very similar (Table 1), we investigated the improvement in speckle reduction achievable through the use of adaptive PWF.

We compared the log standard deviation of

Fig. 7—PWF image of trees (1-ft × 1-ft resolution). This image, taken from an area near the farmhouse, shows the detail obtained from PWF processing of ADTS data. Note the bright leading edges of the trees and the radar shadows directly behind the trees.

Table 1. Polarization Covariance Parameters of ADTS Data

	σ_{HH}	ε	γ	$\rho\sqrt{\gamma}$
Trees	0.256	0.16	0.89	0.61
Mixed	0.098	0.19	1.08	0.60
Grass	0.086	0.19	1.03	0.53
Shadow	0.006	0.43	1.18	0.49

the clutter regions after they were whitened with the covariance of grass versus the result when they were whitened with the proper covariance (e,g., trees whitened with tree covariances). As the data in Table 5 indicate, the use of the proper covariances made virtually no change in the log standard deviations. Thus, because the covariances of the different regions are so similar, the extra computational time required for adaptive PWF is not warranted.

Polarimetric Averaging versus Spatial Averaging

As mentioned earlier in this article, speckle can be reduced by noncoherent spatial averaging, or spoiling, of the high-resolution data. However, spatial averaging degrades image resolution. The log standard deviations of 1-ft and 1-m PWF data are clearly superior to the 1-ft and 1-m single-polarimetric-channel data, as is shown by the results presented in Table 6 for the PWF and HH channels.

For grass regions, PWF data at 1-ft resolution were measured to have a log standard deviation of 3.0 dB. At the same resolution, the HH data were measured to have a log standard deviation of 5.7 dB and the HV data had a log standard deviation of 5.6 dB. Thus polarimetric averaging improved the results over single-channel data by 2.7 dB. Noncoherently spoiling the single-channel data to 1-m resolution (i.e., 4×4 averaging) reduced σ_C by approximately 3.7 dB, 1 dB better than polarimetric averaging—but image resolution was sacrificed for the improvement. Spoiling the PWF data to 1-m resolution yielded the clutter log standard deviation of 1.3 dB, an 0.7-dB improvement over the 1-m HH data, and an 0.5-dB improvement over the 1-m HV data. Similar results were obtained for tree clutter and for shadows.

Table 3. *s/m* Theory vs. Measurement

	Predicted	*Measured*
Trees	1.13	1.10
Mixed	0.90	0.94
Grass	0.70	0.67
Shadow	0.58	0.60

Polarimetric Segmentation

Recent improvements in the resolution and quality of polarimetric SAR imagery now permit the use of advanced image-processing and machine-vision techniques in the analysis of radar imagery. Until recently, a major impediment to the use of such techniques had been the

Table 2. *s/m* of ADTS Data

	HH	*HV*	*VV*	*PWF*
Trees	1.59	1.69	1.39	1.10
Mixed	1.43	1.27	1.38	0.94
Grass	1.12	1.06	1.16	0.67
Shadow	0.99	0.99	1.02	0.60

Table 4. Log Standard Deviations (dB) of 1-ft-Resolution Data

	HH	*HV*	*VV*	*PWF*
Trees	6.67	6.14	6.35	4.13
Mixed	6.28	6.00	6.36	3.90
Grass	5.67	5.59	5.67	2.97
Shadow	5.20	5.24	5.46	2.81

Table 5. Log Standard Deviations (dB) of PWF vs. Adaptive PWF

	PWF	*APWF*
Trees	4.13	4.10
Mixed	3.90	3.89
Grass	2.97	2.97
Shadow	2.18	2.17

presence of speckle in the imagery. Progress in polarimetric SAR sensor technology, coupled with the optimal processing described in this paper, have made it possible to reduce image speckle significantly—while preserving very high resolution. Hence the development of machine-vision techniques may have significant potential for SAR image analysis.

In this section we consider the application of advanced image-processing and machine-vision techniques to clutter segmentation—the problem of partitioning a SAR image into regions of homogeneous terrain types (grass regions, tree regions, roads). Partitioning an image in this way is desirable because such classical radar detection algorithms as constant false-alarm rate (CFAR) assume a homogeneous clutter background when setting the detection threshold. In nonhomogeneous regions, the assumption of homogeneity may cause the CFAR threshold to be set too high, thus reducing detection performance along, for example, tree lines. By first segmenting the clutter, we can set the CFAR detection threshold more appropriately, which will improve detection performance.

Earlier attempts to partition clutter regions in polarimetric SAR imagery have focused on the use of Bayes classifiers on a pixel-by-pixel basis. J.A. Kong et al. [9], showed that an optimal Bayes classifier provides unacceptable error (37.4%) in discriminating between grass and trees. S.R. De Graaf [10] showed that excellent performance can be obtained by the optimal classifier for some clutter types, but that exceptionally poor performance is obtained for others. H.A. Yueh et al. [11] developed an optimal Bayes classifier for normalized polarimetric data, i.e., data that don't include absolute radar cross-section information. Although this classifier is less sensitive to changes in the operating characteristics of the radar than an optimal Bayes classifier that uses the absolute radar cross section, its performance is poorer. Hence we can conclude that a Bayes classifier applied to individual pixels typically provides unacceptable results.

Improved segmentation results have been obtained by coupling a spatial processing method and a pixel-by-pixel PWF classifier. In this scheme, each pixel is classified by using the PWF output as the discrimination statistic. The initial segmentation is then improved by using a voting filter, in which isolated pixels with a different classification from surrounding pixels are relabeled to match the classification of their neighbors. The output of the PWF is a reasonable choice for a discrimination statistic because it minimizes the spread (standard deviation) of the clutter PDF, as illustrated in Figs. 8 and 9. Figure 8 shows the PDFs of the HH channel for selected areas of shadow, tree, and grass clutter. Figure 9 shows the PDFs of the

Table 6. Log Standard Deviations of 1-ft- and 1-m-Resolution Data

	1-ft HH	*1-ft PWF*	*1-m HH*	*1-m PWF*
Trees	6.7	4.1	3.5	2.8
Grass	5.7	3.0	2.0	1.3
Shadow	5.2	2.8	1.7	1.0

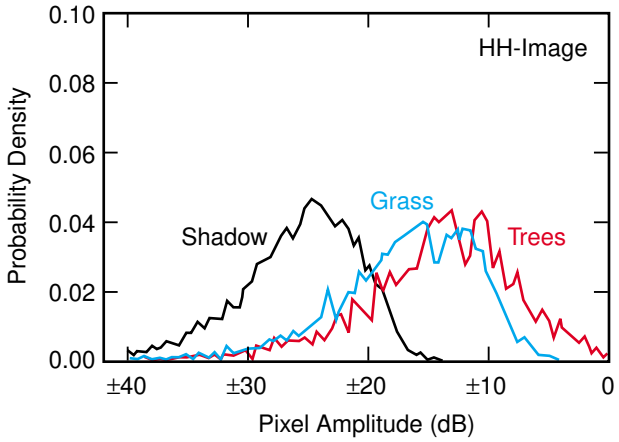

Fig. 8—Probability density functions (PDF) of HH data (1-ft × 1-ft resolution). The PDFs of the pixel amplitudes for three clutter types (shadow, grass, and trees) show that the single-polarimetric-channel data has considerable overlap in the distributions, making it hard, for example, to separate grass and trees on the basis of amplitude.

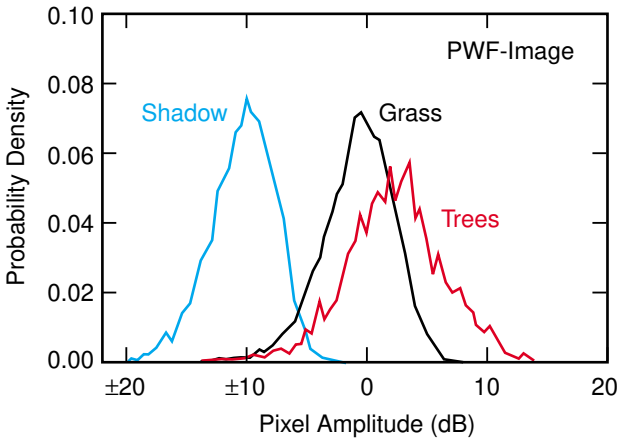

Fig. 9—PDFs of PWF data (1-ft × 1-ft resolution). Here the PDFs of the pixel amplitudes for the same clutter data as in Fig. 8 have much less overlap in their distributions (smaller standard deviation) than in the single-channel counterpart.

PWF for these same areas of clutter. These clutter types are much more separable in the PWF imagery than in the HH imagery, clearly showing the advantage of the polarimetric method. The histogram in Fig. 9 can be used to select thresholds that minimize the misclassification error. A sample PWF image (at 1-ft × 1-ft resolution) is shown in Fig. 10; the segmentation of this image into three clutter types (grass, trees, and roads and shadows) is shown in Fig. 11.

An alternate approach to segmentation that uses minimum speckle PWF imagery is currently under development. Preliminary results of this edge-based segmentation algorithm applied to Fig. 6 (the farm scene at 1-m × 1-m resolution) are shown in Fig. 12. This algorithm uses a Marr-Hildreth edge detector [12] to identify boundaries between relatively bright regions and dim regions in the image. The boundaries are grouped to form regions with a threshold relaxation technique that incorporates a smoothness constraint. The algorithm produces a binary image in which relatively dim areas (shadows, rivers, and roads) are represented as black, and relatively bright areas (grass, fields, and trees) are represented as white. Comparison of the original image (Fig. 6) with the segmentation output (Fig. 12) indicates that the coarse features of the image have been accurately captured by the segmentation process.

The development of image-analysis and machine-vision algorithms for SAR imagery is in its infancy. But improvements in SAR sensor technology and in polarimetric processing capabilities—such as the polarimetric whitening filter—have opened a wide range of possibilities for automatic interpretation of SAR imagery.

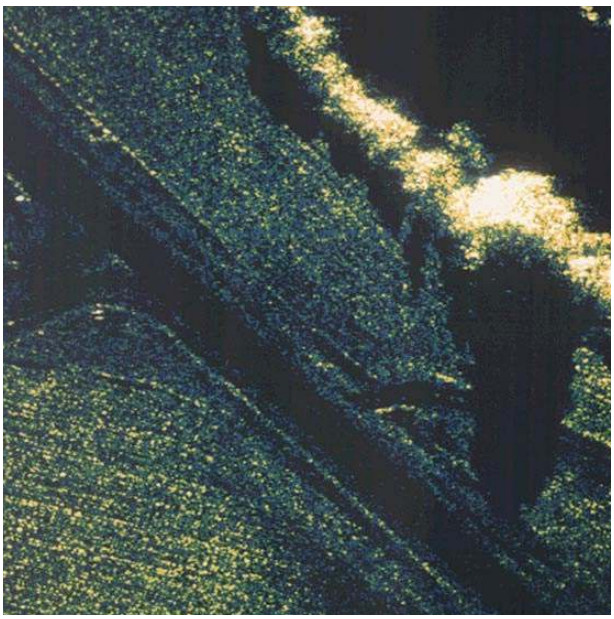

Fig. 10—PWF image of road intersection (1-ft × 1-ft resolution). Note that this image has several transition boundaries: river-tree, tree-shadow, shadow-meadow, and meadow-road.

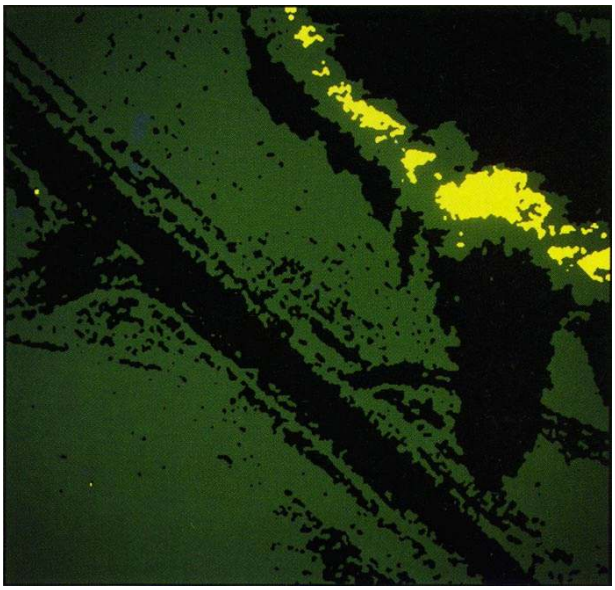

Fig. 11—Segmented version of PWF road intersection image (shown in Fig. 10) into three clutter types—grass, trees, and roads and shadows.

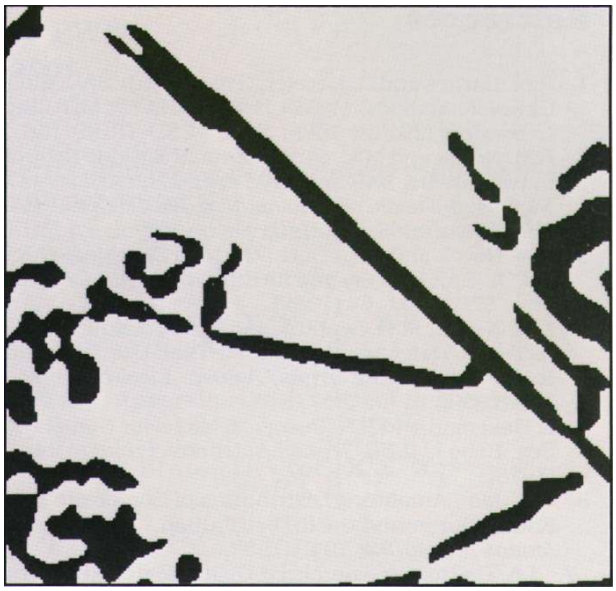

Fig. 12—Segmented version of PWF farm image (shown in Fig. 6) into two clutter types: bright and dim.

Right now, clutter segmentation can improve the performance of classical target-detection and target-classification algorithms. In the future, machine-vision algorithms applied to SAR imagery should lead to more sophisticated and robust approaches to target detection and classification.

Conclusions

We have investigated polarimetric speckle reduction by using the polarimetric whitening filter; the polarimetric whitening filter reduced synthetic-aperture-radar image speckle without degrading the spatial resolution of the image. Results obtained with ADTS SAR data show that the PWF reduces speckle—and it significantly reduces clutter log standard deviations. At 1-ft resolution, the log standard deviation of clutter is reduced by approximately 2.5 dB relative to a single-channel system. This improvement in log standard deviation provides improved detection performance. The PWF technique also has potential application to the problem of clutter segmentation.

Another potential application of minimum-speckle imagery is in the discrimination of man-made objects from natural tree clutter by using texture features [13]. We are planning future studies of PWF-processed SAR data to evaluate the improvement in texture discrimination by using PWF data.

Acknowledgments

The ADTS and special-purpose ground processing equipment were built by the Loral Corp. of Phoenix, Ariz. The Advanced Detection Technology program at MIT Lincoln Laboratory is sponsored by DARPA under contract No. AF19628-90-C-0002.

References

1. R.M. Barnes and D.J. Blejer, "Polarimetric SAR Antenna Characterization," *Project Report STD-15,* MIT Lincoln Laboratory (28 July 1989), DTIC #ESD-TR-89-166.
2. L.M. Novak and M.C. Burl, "Optimal Speckle Reduction in Polarimetric SAR Imagery," *Proc. 22nd Asilomar Conf. Signals, Systems, and Computers, Pacific Grove, CA, 31 Oct.–2 Nov. 1988,* p. 781.
3. L.M. Novak and M.C. Burl, "Optimal Speckle Reduction in POL-SAR Imagery and Its Effect on Target Detection," *Proc. SPIE* **1101,** 84 (1989).
4. L.M. Novak, M.B. Sechtin, and M.J. Cardullo, "Studies of Target Detection Algorithms That Use Polarimetric Radar Data," *IEEE Trans. Aerosp. Electron. Syst.* **25,** 150 (1989).
5. E. Jakeman and P.N. Pusey, "A Model for Non-Rayleigh Sea Echo," *IEEE Trans. Antennas Propag.* **24,** 806 (1976).
6. J.K. Jao, "Amplitude Distribution of Composite Terrain Radar Clutter and the K-Distribution," *IEEE Trans. Antennas Propag.* **32,** 1049 (1984).
7. J.A. Cadzow, "Generalized Digital Matched Filtering," *Proc. 12th Southeastern Symp. on System Theory, Virginia Beach, VA, 19–20 May 1980,* p. 307.
8. R.D. Chaney, M.C. Burl, and L.M. Novak, "On the Performance of Polarimetric Target Detection Algorithms," *Proc. 1990 IEEE International Radar Conf., Arlington, VA, 7–10 May 1990,* p. 520.
9. J.A. Kong, A.A. Swartz, H.A. Yueh, L.M. Novak, and R.T. Shin, "Identification of Terrain Cover Using the Optimum Polarimetric Classifier," *J. Electromag. Waves Appl.* **2,** 171 (1988).
10. S.R. De Graaf, "Statistical Classification of Polarimetric SAR Images," *1989 International Symp. Digest—Antennas and Propagation, San Jose, CA, 26–30 June 1989,* p. 1349.
11. H.A. Yueh, A.A. Swartz, J.A. Kong, R.T. Shin, and L.M. Novak, "Bayes Classification of Terrain Cover Using Normalized Polarimetric Data," *J. Geophys. Res. B* **93,** 15,261 (1988).
12. D. Marr and E. Hildreth, "Theory of Edge Detection," *Proc. R. Soc. London Ser. B* **207,** 187 (1980).
13. M.C. Burl, G.J. Owirka, and L.M. Novak, "Texture Discrimination in Synthetic Aperture Radar Imagery," *Proc. 23rd Asilomar Conf. Signals, Systems, and Computers, Pacific Grove, CA, 30 Oct.–1 Nov. 1989,* p. 399.

Appendix 1: The PWF as an Estimator

In this appendix, we show that for a given polarimetric measurement vector **Y,** the PWF is the maximum-likelihood estimate (MLE) of the clutter texture parameter g. We then show that the PWF is an unbiased, minimum-variance estimator (i.e., it achieves the Cramer-Rao lower bound) of the texture parameter g. Finally, we derive the maximum *a posteriori* (MAP) and conditional mean (Bayes) estimates of g and show that these estimates are related to the PWF. A text by H.L. Van Trees gives an excellent description of the mathematics used in this appendix [1].

Maximum-Likelihood Estimate

In MLE, the parameter g is treated as if it were deterministic (nonrandom) but unknown. We seek the value of parameter g that makes the observed vector $\mathbf{Y}$ most likely. If we denote the MLE of g for a given $\mathbf{Y}$ by g_{MLE}, then g_{MLE} is implicitly defined by

$$\frac{\partial}{\partial g}\,\mathrm{p}(\mathbf{Y}|g)\big|_{g=g_{\mathrm{MLE}}} = 0$$

where $\mathrm{p}(\mathbf{Y}|g)$ is the conditional probability density function (PDF) of the vector $\mathbf{Y}$ given g. This conditional PDF is easy to evaluate since, given g, the vector $\mathbf{Y}$ is complex Gaussian with mean $\mathbf{0}$ and covariance . Thus

$$\mathrm{p}(\mathbf{Y}|g) = \frac{1}{\pi^p|\mathbf{\Sigma}|g^p}\ \exp\left(-\mathbf{Y}^\dagger\mathbf{\Sigma}^{-1}\mathbf{Y}/g\right)$$

where p = the number of polarizations ($p = 3$ in our case). To obtain the MLE, we need to find the value of g for which $\mathrm{p}(\mathbf{Y}|g)$ is maximum. Equivalently, we can find the value of g for which $\log \mathrm{p}(\mathbf{Y}|g)$ is maximum.

$$\log\left[\mathrm{p}(\mathbf{Y}|g)\right] = -\log\left(\pi^p|\mathbf{\Sigma}|\right) - p\log g - \mathbf{Y}^\dagger\mathbf{\Sigma}^{-1}\mathbf{Y}/g.$$

Differentiating with respect to g yields

$$\frac{\partial}{\partial g}\log\left[\mathrm{p}(\mathbf{Y}|g)\right] = -\frac{p}{g} + \frac{\mathbf{Y}^\dagger\mathbf{\Sigma}^{-1}\mathbf{Y}}{g^2} = 0$$

$$g_{\mathrm{MLE}} = \frac{\mathbf{Y}^\dagger\mathbf{\Sigma}^{-1}\mathbf{Y}}{p}, \qquad \text{(A1)}$$

which is precisely the PWF solution. Next, we investigate the estimate g_{MLE} and determine (1) that the estimate is unbiased, and (2) that the estimate achieves the Cramer-Rao lower bound on variance of the error (thus g_{MLE} is an efficient estimate of the texture parameter g).

First we demonstrate that g_{MLE} is unbiased. From Eq. A1 we have

$$g_{\mathrm{MLE}} = \frac{1}{p}\mathbf{Y}^\dagger\mathbf{\Sigma}^{-1}\mathbf{Y} = \frac{1}{p}g\mathbf{X}^\dagger\mathbf{\Sigma}^{-1}\mathbf{X}.$$

Taking the expectation with respect to $\mathbf{X}$ gives

$$\mathrm{E}(g_{\mathrm{MLE}}) = \frac{1}{p}g\mathrm{E}\left(\mathbf{X}^\dagger\mathbf{\Sigma}^{-1}\mathbf{X}\right).$$

But Eq. 5 on p. 277 gives

$$\begin{aligned}\mathrm{E}\left(\mathbf{X}^\dagger\mathbf{\Sigma}^{-1}\mathbf{X}\right) &= \mathrm{tr}\left[\mathbf{\Sigma}^{-1}\mathrm{E}\left(\mathbf{X}\mathbf{X}^\dagger\right)\right] \\ &= \mathrm{tr}\left(\mathbf{\Sigma}^{-1}\mathbf{\Sigma}\right) = p.\end{aligned}$$

Thus we have verified that the MLE estimate, and g_{MLE}, is unbiased. That is,

$$\mathrm{E}(g_{\mathrm{MLE}}) = g.$$

Next we verify that g_{MLE} satisfies the Cramer-Rao lower bound, which states that the variance of any unbiased estimate, $\hat{g}$, must satisfy the inequality

$$\mathrm{VAR}\left(\hat{g}\right) \geq \frac{-1}{\mathrm{E}\left\{\dfrac{\partial^2}{\partial g^2}\log\left[\mathrm{p}(\mathbf{Y}|g)\right]\right\}}.$$

Omitting the details of the derivation, one may easily show that

$$\frac{\partial^2}{\partial g^2}\log\left[\mathrm{p}(\mathbf{Y}|g)\right] = \frac{p}{g^2} - \frac{2\mathbf{Y}^\dagger\mathbf{\Sigma}^{-1}\mathbf{Y}}{g^3}.$$

Evaluating the expectation of the above, again using Eq. 5, yields the result

$$\mathrm{E}\left\{\frac{\partial^2}{\partial g^2}\log\left[\mathrm{p}(\mathbf{Y}|g)\right]\right\} = \frac{-p}{g^2}.$$

Thus the variance of any unbiased estimator of the texture parameter g must satisfy the Cramer-Rao bound

$$\mathrm{VAR}\left(\hat{g}\right) \geq \frac{g^2}{p}.$$

Finally, we can verify that the unbiased estimate g_{MLE} achieves the lower bound.

$$\mathrm{VAR}\left(g_{\mathrm{MLE}}\right) = \mathrm{VAR}\left(\frac{1}{p}\, g\, \mathbf{X}^{\dagger}\boldsymbol{\Sigma}^{-1}\mathbf{X}\right) = \frac{g^2}{p^2}\,\mathrm{VAR}\left(\mathbf{X}^{\dagger}\boldsymbol{\Sigma}^{-1}\mathbf{X}\right).$$

Evaluating the above, using Eq. 6 on p. 277, yields the result

$$\mathrm{VAR}\left(g_{\mathrm{MLE}}\right) = \frac{g^2}{p}.$$

Maximum *a Posteriori* Estimate

The maximum *a posteriori* (MAP) estimate of g (also known as the conditional mode estimate) is defined as

$$\frac{\partial}{\partial g}\left[\mathrm{p}(g|\mathbf{Y})\right]\Big|_{g=g_{\mathrm{MAP}}} = 0.$$

Although the MAP estimate appears to be very similar to the MLE estimate, the MAP estimate treats g as a random variable with a known PDF. This can easily be seen from the Bayes theorem.

$$\mathrm{p}(g|\mathbf{Y}) = \frac{\mathrm{p}(\mathbf{Y}|g)\ \cdot\ \mathrm{f_G}(g)}{\mathrm{p}(\mathbf{Y})}.$$

The MAP estimate is obtained by differentiating this expression with respect to g and setting the result to 0. (Note that we can neglect p(**Y**).) Proceeding in this way, we find

$$\frac{\partial}{\partial g}\left[\mathrm{p}(g|\mathbf{Y})\right] = \frac{\partial}{\partial g}\left[g^{(\upsilon-1-p)}\ \exp\left(-g/\bar{g} - \mathbf{Y}^{\dagger}\Sigma^{-1}\mathbf{Y}/g\right)\right]$$

$$= g^{(\upsilon-1-p)}\ \exp(\cdot)\times\left(-\frac{1}{\bar{g}} + \frac{\mathbf{Y}^{\dagger}\Sigma^{-1}\mathbf{Y}}{g^2}\right) + \exp(\cdot)\times(\upsilon-1-p)\ g^{\upsilon-2-p}$$

$$= g^{(\upsilon-2-p)}\exp(\cdot)\times\left[-\frac{g}{\bar{g}} + \mathbf{Y}^{\dagger}\boldsymbol{\Sigma}^{-1}Y/g\ +(\upsilon-1-p)\right].$$

Therefore,

$$g^2 - \bar{g}(\upsilon-1-p)g - \bar{g}\mathbf{Y}^{\dagger}\boldsymbol{\Sigma}^{-1}\mathbf{Y} = 0.$$

Because this is a quadratic equation in g, the solution is given by

$$g_{\mathrm{MAP}} = \frac{\bar{g}(\upsilon-1-p)}{2}\cdot\left[1 \pm \sqrt{1 + \frac{4\,\mathbf{Y}^{\dagger}\Sigma^{-1}\mathbf{Y}}{\bar{g}(\upsilon-1-p)^2}}\right].$$

Observe that g_{MAP} is a nonlinear function of the PWF.

Conditional Mean Estimate

Finally, we present the conditional mean estimate (also known as the Bayes estimate).

$$g_{\mathrm{Bayes}} = \int_0^{\infty} g\,\mathrm{p}(g|\mathbf{Y})\,\mathrm{d}g.$$

The derivation is complex; therefore, only the final result is presented.

$$g_{\mathrm{Bayes}} = \left(\bar{g}\mathbf{Y}^{\dagger}\boldsymbol{\Sigma}^{-1}\mathbf{Y}\right)^{\frac{1}{2}}\cdot\frac{\mathrm{K}_{p-\upsilon-1}\left(2\sqrt{\mathbf{Y}^{\dagger}\boldsymbol{\Sigma}^{-1}\mathbf{Y}}/\bar{g}\right)}{\mathrm{K}_{p-\upsilon}\left(2\sqrt{\mathbf{Y}^{\dagger}\boldsymbol{\Sigma}^{-1}\mathbf{Y}}/\bar{g}\right)}.$$

Again, the estimate is found to be a nonlinear function of the PWF.

Reference

1. H.L. Van Trees, *Detection, Estimation, and Modulation Theory, Part I* (Wiley, New York, 1968), chap. 2.

Appendix 2: Relating σ_c and υ

The log standard deviation of the texture variable g (denoted σ_c) is defined to be

$$\sigma_c = \sqrt{\mathrm{VAR}\left(10 \log_{10} g\right)}. \quad \text{(dB)}.$$

The relation between σ_c and the shape parameter υ of the gamma PDF can be described as

$$\mathrm{E}(\log g) = \frac{1}{\Gamma(\upsilon)\bar{g}^{\upsilon}} \int_0^{\infty} (\log g)\, g^{\upsilon-1} \exp\left(-\frac{g}{\bar{g}}\right) \mathrm{d}g.$$

From Ref. 1, Eq. 4.352.1, we have

$$\mathrm{E}\left(\log g\right) = \Psi(\upsilon) + \log \bar{g}$$

$$\mathrm{E}\left[(\log g)^2\right] = \frac{1}{\Gamma(\upsilon)\bar{g}^{\upsilon}} \int_0^{\infty} (\log g)^2\, g^{\upsilon-1} \exp\left(-\frac{g}{\bar{g}}\right) \mathrm{d}g.$$

Ref. 1, Eq. 4.358.2, also gives

$$\mathrm{E}\left[(\log g)^2\right] = \left[\Psi(\upsilon) + \log \bar{g}\right]^2 + \zeta(2, \upsilon).$$

Thus

$$\mathrm{VAR}\left(\log g\right) = \zeta(2, \upsilon)$$

and

$$\mathrm{VAR}\left(10 \log_{10} g\right) = \mathrm{VAR}\left(\frac{10 \log g}{\log 10}\right) = (4.34)^2\, \zeta(2, \upsilon).$$

Table A lists the value of the gamma parameter as a function of the clutter standard deviation for typical clutter standard deviations.

Table A. Gamma Parameter (υ) vs. σ_c (dB)

σ_c	υ
1.0	19.3
1.5	8.9
2.0	5.2
2.5	3.5
3.0	2.6

Reference

1. I.S. Gradshteyn and I.M. Ryzhik, *Table of Integrals, Series, and Products* (Academic Press, New York, 1980), pp. 576, 578.

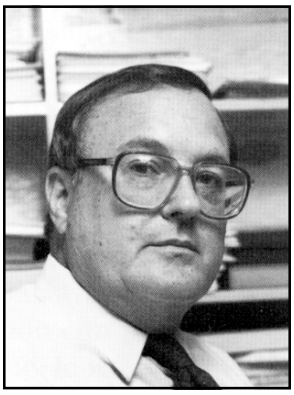

LESLIE M. NOVAK is a senior staff member in the Battlefield Surveillance Group. He received a B.S.E.E. degree from Fairleigh Dickinson University in 1961, an M.S.E.E. degree from the University of Southern California in 1963, and a Ph.D. degree from the University of California, Los Angeles, in 1971. Since 1977 Les has been a member of the technical staff at Lincoln Laboratory, where he has studied detection, discrimination, and classification of radar targets. He has contributed chapters on stochastic observer theory to the series *Advances in Control Theory,* edited by C.T. Leondes (Academic Press, New York), volumes 9 and 12.

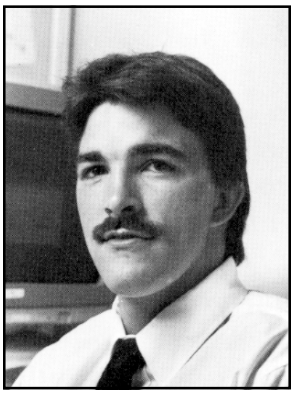

MICHAEL C. BURL is an associate staff member in the Battlefield Surveillance Group. The focus of his research is on the detection, discrimination, and classification of stationary battlefield vehicles. Michael joined Lincoln Laboratory three years ago. He received a B.S. degree in applied mathematics and electrical engineering from the California Institute of Technology, where he was named 1987 Caltech Athlete of the Year. Michael is currently taking classes at MIT as a special graduate student.