Trained Detection of Buried Mines in SAR Images via the Deflection-Optimal Criterion

Russell B. Cosgrove, Peyman Milanfar, *Senior Member*, and Joel Kositsky

***Abstract*—In this paper, we apply a deflection-optimal linear-quadratic detector to the detection of buried mines in images formed by a forward-looking ground-penetrating synthetic aperture radar. The detector is a linear-quadratic form that maximizes the output SNR (deflection), and its parameters are estimated from a set of training data. We show that this detector is useful when the signal to be detected is expected to be stochastic, with an unknown distribution, and when only a small set of training data is available to estimate its statistics. The detector structure can be understood in terms of the singular value decomposition; the statistical variations of the target signature are modeled using a compact set of orthogonal "eigenmodes" (or principal components) of the training dataset. Because only the largest eigenvalues and associated eigenvectors contribute, statistical variations that are underrepresented in the training data do not significantly corrupt the detector performance. The resulting detection algorithm is tested on data that are not in the training set, which has been collected at government test sites, and the algorithm performance is reported.**

***Index Terms*—Automatic target recognition, buried mines, deflection, detection, synthetic aperture radar (SAR), training, principal components.**

I. Introduction

THE DETECTION of buried mines is a vexing problem of broad military and humanitarian concern. For instance, any ground-based mobile armed force is concerned with the detection and clearing of minefields in its path. Long after conflicts are settled, the very same minefields pose a grave hazard to civilian populations who reoccupy the area of conflict. Indeed, many innocent civilians are killed or injured each year by land mines that go otherwise undetected.

In this paper, we present our recent efforts in developing a framework for the detection of buried mines. The techniques developed in this paper can be considered a general class of automatic target detection/recognition algorithms and, as such, are broadly applicable to many sensor modalities. However, the present work concentrates on the case where observations are collected using a dual-polarization forward-looking ground-penetrating synthetic aperture radar (FLGP SAR) sensor [1]. The sensor in question was developed over the past four years under the generous support of the U.S. Army's Fort Belvoir Communications–Electronics Research Development and Engineering Center Night Vision and Electronic Sensors Directorate Counter-Mine Program. The promising performance of the proposed methods is demonstrated on actual field data collected during two separate field experiments.

Since the early days of SAR [2], many algorithms have been developed for postprocessing the resulting (complex) SAR imagery, to produce classification of pixels into regions with specific properties of interest. One class of algorithms seeks to adaptively model the background "clutter" and detect targets as deviations from the model predictions (e.g., [3]–[8]). Another class of algorithms uses prior knowledge to construct a target signature and compares the image under test to the signature (e.g., [9]–[14]). The latter comparison may be accomplished in a variety of ways, ranging from optimized neural networks (e.g., [13] and [15]) to statistically optimal likelihood functions (e.g., [14] and [16]). Of interest in our case is whether a mine is buried underground in the region being imaged. Because the resolution of the FLGP SAR is sufficient to resolve multiple-scattering centers on mines, we will take the approach of constructing a spatial target signature using prior knowledge. We will derive a statistically optimal test statistic using a set of training data to characterize the target statistics.

In general, construction of a target signature can proceed either through physics-based modeling or through use of training data. Physics-based modeling can provide useful intuition on target behavior [17] and can be more practical than constructing a large field experiment to collect training data. However, physics-based modeling is sensitive to the specific parameters of the simulations and makes strong, sometimes unjustified, assumptions about how the targets of interest truly behave in a real imaging environment. Hence, quantitative application of a physics-based model requires a calibration procedure based in training data. Calibration to actual sensor data is, in itself, a challenging and interesting problem, which we wish to avoid here.

The approach we take is motivated by the fact that we can, in many instances, have access to training data from the sensor in question. In the particular case of the FLGP SAR, several field experiments were carefully planned and carried out so as to collect a statistically significant number of realizations of images of several types of mines buried in a variety of soil types and depths [1]. We construct a "learned" statistical model of the buried target from data that are collected under conditions that are extremely difficult, if not impossible, to model theoretically. Our learned model of the target signature is constructed by considering a low-dimensional approximation of a target-plus-clutter covariance matrix computed from a limited set of training data.

The detector we apply is optimal in the sense that it maximizes the generalized SNR at its output, otherwise known as

Manuscript received December 31, 2002; revised June 29, 2004. This work was supported by the U.S. Army Communications–Electronics Research Development and Engineering Center Night Vision and Electronic Sensors Directorate Counter-Mine Program.

R. B. Cosgrove and J. Kositsky are with SRI International, Menlo Park, CA 94025 USA (e-mail: russell.cosgrove@sri.com; joel.kositsky@sri.com).

P. Milanfar is with the Department of Electrical Engineering, University of California, Santa Cruz, CA 95064 USA (e-mail: milanfar@ee.ucsc.edu).

Digital Object Identifier 10.1109/TGRS.2004.834591

the deflection [22]. In the present scenario, the deflection can be thought of a measure of the separation of the probability density functions (pdfs) associated with the presence or absence of targets (mines) in a SAR image. The optimal test statistic from the point of view of probability of detection versus probability of false alarm is determined by the Neyman–Pearson (NP) criterion to be the likelihood ratio $p(x;H_1)/p(x;H_0)$ [18], where $p(x;H_1)$ and $p(x;H_0)$ are the pdfs associated with the presence and absence of targets, respectively. We find below that application of the NP criterion is impractical in the present scenario, since in general no pdf is available under H_1, where instead only a limited set of training data may be available. We show that the deflection-optimal detector is an appropriate substitute, as it produces an effective detector structure that can be constructed with limited knowledge of the pdfs and with even a few training samples. While the deflection-optimal detector has been studied extensively before, its use in the presence of training data, we believe, is novel and of great practical importance. This extended use of the deflection-based framework and its application to the detection of targets in SAR images form the core contribution of this paper.

The study of detectors that maximize the deflection actually predates the likelihood ratio theory. Baker [19] derived the deflection-optimal detector for Gaussian statistics. Chevalier and Picinbono [20] derived the general deflection-optimal linear-quadratic detector for complex data. An application to cellular mobile communications is given by Shikh-Bahaei [21]. The deflection-optimal linear-quadratic detector is known to be useful when the statistics are not Gaussian [21]–[23], since its parameters can be computed using only the moments through fourth order. A general discussion on the deflection as a performance criterion can be found in [23].

The FLGP SAR has been designed and built at SRI International. Field tests of the radar have been conducted at various government test sites, consisting of dirt and gravel lanes with mines buried at depths ranging from near-zero to 15 cm. This paper will focus on type TM-62M metal antitank mines, of which there were between 20–44 per lane.

The organization of the paper is as follows. In Section II, the deflection-optimal detector is presented. In Section III, we describe the characteristics of the FLGP SAR, and of imaged mines, and describe how to define a feature vector for input to the detector. In Section IV, we apply the results of Sections II and present the detection performance on a complementary test dataset, which does not include the training data. Concluding remarks are given in Section V.

II. Statistical Detection Theory

Statistical detection theory provides a framework for deciding among a set of mutually exclusive hypotheses regarding a set of data. We will consider a decision between: 1) the H_0 hypothesis that the data consist of random noise $\mathbf{e}$ (clutter) and 2) the H_1 hypothesis that the data consists of noise $\mathbf{e}$ plus an unknown signal $\mathbf{s}$. In the classical setting (see Section II-A below), pdfs for the data under both hypotheses are assumed to be known (at least to within some parameters). In our proposed framework, however, we do not impose this structure on the data under the H_1 hypothesis. Instead, the (first and second order) statistics of the data under the H_1 hypothesis will be characterized by the training data $\{\tilde{\mathbf{x}}_1, \tilde{\mathbf{x}}_2, \ldots, \tilde{\mathbf{x}}_N\}$, and we will use only this statistical characterization to develop an effective detector. To simplify the exposition, we assume that $\mathbf{e}$ has mean zero and covariance $\mathbf{I}$ (the identity). In practice, the clutter mean $\boldsymbol{\mu}_0$ and covariance $\mathbf{C}_0$ can be estimated from either training data or from the image under test,[1] and a prewhitening step must be applied to arrive at the white noise assumption we make here.

The hypotheses can thus be summarized as

$$H_1 : \mathbf{x} = \mathbf{s} + \mathbf{e} \quad H_0 : \mathbf{x} = \mathbf{e}$$

where we assume that the clutter $\mathbf{e}$ is a stochastic variable distributed as a complex Gaussian with zero mean and covariance $\mathbf{I}$ (i.e., $\mathbf{e} \sim \mathcal{CN}(\mathbf{0}, \mathbf{I})$). As for the statistics of the data $\mathbf{s}+\mathbf{e}$ under the H_1 hypothesis, we consider the signal from the mine to be a deterministic signature measured through a random medium (due to being buried), and we do not attempt to characterize it specifically by way of a model. Instead, we use the training set $\{\tilde{\mathbf{x}}_1, \tilde{\mathbf{x}}_2, \ldots, \tilde{\mathbf{x}}_N\}$ to characterize the mean and covariance of the data under the H_1 hypothesis only. This information will suffice in the development of the detection algorithm we propose.

A. Classical Neyman–Pearson Detector

For the sake of completeness and to establish a frame of reference against classical methods, consider the case when the pdfs can be associated with both the H_1 and H_0 hypotheses. In this case, the NP theorem [18] determines the optimal test statistic, yielding the highest probability of detection for a given probability of false alarm. The NP theorem yields the likelihood ratio test

$$L(\mathbf{x}) = \frac{p(\mathbf{x};H_1)}{p(\mathbf{x};H_0)} \underset{H_0}{\overset{H_1}{\gtrless}} \gamma \tag{1}$$

where $p(\mathbf{x};H_1)$ and $p(\mathbf{x};H_0)$ are the pdfs under the H_1 and H_0 hypotheses, respectively, and γ is a threshold chosen to achieve the desired probability of false alarm. As the logarithm is a monotonic function, the log-likelihood function $\log(L(\mathbf{x}))$ is an alternative test statistic with the same performance.

By assuming a (complex) Gaussian pdf $\mathcal{CN}(\boldsymbol{\mu}, \mathbf{C}+\mathbf{I})$ for the H_1 hypothesis, and recalling from (1) that $\mathbf{e} \sim \mathcal{CN}(\mathbf{0}, \mathbf{I})$, the log-likelihood function is given by

$$T_{\mathrm{NP}}(\mathbf{x}) = -(\mathbf{x}-\boldsymbol{\mu})^H(\mathbf{C}+\mathbf{I})^{-1}(\mathbf{x}-\boldsymbol{\mu}) + \mathbf{x}^H\mathbf{x} \tag{2}$$

where $\boldsymbol{\mu}$ is the mean of the signal $\mathbf{s}$, $\mathbf{C}$ is the covariance matrix of the signal $\mathbf{s}$, and the superscript H denotes the conjugate transpose operation.

If the parameters $\boldsymbol{\mu}$ and $\mathbf{C}$ are unknown (the case of interest to us), then the above detector structure cannot be immediately useful. But there do exist some well-known approaches to this problem. On the one hand, we can assume a prior on these parameters and integrate them out. This approach is quite difficult, as it requires the calculation of complicated multidimensional

[1]The former option is certainly more desirable, but the latter can also be effective if care is taken to ensure that the estimates are not corrupted by any target-bearing pixels. A procedure for estimating the clutter mean and covariance matrix from training data, or from a given image under test, has been described in the [26] (available by request from the author).

integrals. Alternatively, one can estimate the unknown parameters, (say using the maximum-likelihood (ML) principle) and use the estimated values in the NP test.

Taking the latter view, commonly known as the *generalized likelihood* approach, we can attempt to apply the NP detector (2) to our mine detection problem by first computing ML estimates of $\boldsymbol{\mu}$ and the matrix $\mathbf{C}$ from the training data. Assuming a known distribution for the underlying data under H_1, this approach is attractive, but in some instances impractical.

First, while the clutter statistics may arguably be modeled as complex Gaussian after the whitening process, the variations in the signal statistics, owing to the complicated underlying physical phenomena, cannot be reliably modeled (certainly not using a Gaussian pdf). Even if such a pdf is assumed for the data under H_1, from a relatively small set of training data we can only hope to accurately estimate a low-rank approximation to the matrix $\mathbf{C}$. This can be particularly problematic, since as (2) shows, the matrix $\mathbf{C}$ indeed appears in inverted form. With these observations, we conclude that the generalized likelihood approach is not suitable directly, and in this paper, we suggest an alternative approach.

B. Deflection-Optimal Detector

Detector performance, in terms of probability of detection versus probability of false alarm, is a measure of the separation between the pdfs associated with the H_1 and H_0 hypotheses. The NP detector (2), which is a linear-quadratic form for Gaussian data, maximizes this measure. However, under some circumstances such as the present application, at worst one cannot assume knowledge of the pdf for the data under H_1; and at best, the estimation of the parameters of such a pdf from a small set of training data can be very difficult. We seek a practical solution to the detection problem when the pdf is unknown under H_1 and when a relatively small training dataset is available. We consider an alternative criterion called *deflection* for determining the parameters of a linear-quadratic detector. This criterion maximizes a reliable measure of the separation between the two hypotheses and is analogous to the SNR at the output of the detector. Although this measure is not in general optimal with respect to probability of detection versus probability of false alarm, it does coincide with the optimal NP criterion when the signal is known to be deterministic.

The generalized SNR at the output of the detector T, otherwise known as the generalized deflection [22], [27], is defined as

$$d = \frac{(E_1(T) - E_0(T))^2}{\alpha \mathrm{Var}_0(T) + (1-\alpha)\mathrm{Var}_1(T)} \tag{3}$$

where $E_i(T)$ and $\mathrm{Var}_i(T)$ are the mean and variance of T under the H_i hypothesis. In this paper, we use a specialized version of the above definition where $\alpha = 1$. The resulting measure is referred to as simply *deflection* [22] from now on. Having set when $\alpha = 1$, the deflection-optimal detector is then defined as the detector that maximizes (3).

For the complex linear-quadratic form

$$T(\mathbf{x}) = \mathbf{x}^H \mathbf{A}\mathbf{x} + \mathbf{x}^H \mathbf{b} + \mathbf{b}^H \mathbf{x} \tag{4}$$

which should be compared with (2), the Hermitian matrix $\mathbf{A}$ and the vector $\mathbf{b}$ that maximize the deflection d can be found in [20]. We specialize the results in [20] for the present case where we assume $\mathbf{x} = \mathbf{e} \sim \mathcal{CN}(\mathbf{0}, \mathbf{I})$ under H_0. Furthermore, if under the H_1 hypothesis, we assume *only* the moments $E_1(\mathbf{x}) = \boldsymbol{\mu}$, and $\mathrm{Cov}_1(\mathbf{x}) = \mathbf{C} + \mathbf{I}$ to be given, we obtain

$$T_d(\mathbf{x}) = \mathbf{x}^H(\mathbf{C} + \boldsymbol{\mu}\boldsymbol{\mu}^H)\mathbf{x} + \mathbf{x}^H\boldsymbol{\mu} + \boldsymbol{\mu}^H\mathbf{x} \tag{5}$$

where

$$\mathbf{b} = E_1(\mathbf{x}) = \boldsymbol{\mu} \quad \mathbf{A} = E_1(\mathbf{x}\mathbf{x}^H) - \mathbf{I} = \mathbf{C} + \boldsymbol{\mu}\boldsymbol{\mu}^H. \tag{6}$$

A key idea explored and used in this paper is that the deflection-optimal detector is advantageous for two complementary reasons. First, it is well defined when no specific pdf is available under H_1. Second, when a small number of training samples is given, the deflection-based detector involves the unknown detector parameters more directly. Namely, a comparison of (5) to (2) reveals that unlike the NP test statistic, the deflection approach does not involve the inversion of the matrix $\mathbf{C}$. In either case, of course, the detector parameters are unknown and must be estimated from training data. However, the parameters $\mathbf{A}$ and $\mathbf{b}$ of the deflection-based detector can be estimate from the training data even if only a few training samples are available, as we demonstrate below. The same cannot be said of the alternative approach. In Section IV-A, we will describe specifically how we estimate these parameters to form the "learned" detector based on the deflection criterion in a practical mine detection scenario.

III. Buried Mines in SAR Images and Choice of a Feature Vector

To apply the algorithm to detecting buried mines (of approximately known size) in SAR image, s we must decide what image features (derived from the pixels) to provide as inputs to the detector. An obvious choice is to directly use the pixels in a mine-sized region of the image (an image chip). The full image can then be scanned pixel by pixel, taking the input data from the chip centered on the current pixel. [If desired, a prescreener can be used to reduce the search region (e.g., see [9]).] For reasons to be explained below, only the pixels in a one-dimensional (1-D) range cut through the current pixel will be used as input data. We will refer to these collections of data from 1-D cuts as feature vectors below.

A. Image Characteristics

The SAR technique involves a coherent combination of radar returns from a large number of closely spaced looks at the scene to be imaged. In our experiment, the radar is mounted on a van and driven down the test lane, with periodic stops to collect data. The collection geometry is shown in Fig. 1. At each stop, the antennas are scanned through 72 horizontal positions, transverse to the lane, across the top of the van. A SAR image is formed from the 72 datasets, using a filtered backprojection method algorithm, a standard SAR imaging algorithm [24]. Separate complex valued images are formed in the horizontal (HH) and vertical (VV) polarization modes. A sample (HH polarization magnitude) image containing three mines is shown in Fig. 2.

The resolution of the images formed in this way is about 10 cm in the range direction (down the lane) and 40 cm in the cross-range direction (although the latter is a function of range).

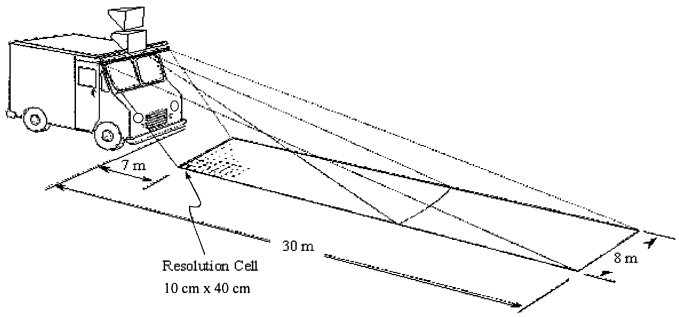

Fig. 1. Collection geometry.

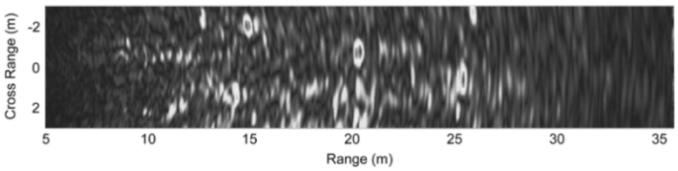

Fig. 2. Sample of horizontally polarized raw image. Three mines are readily visible at 15, 20, and 25 m.

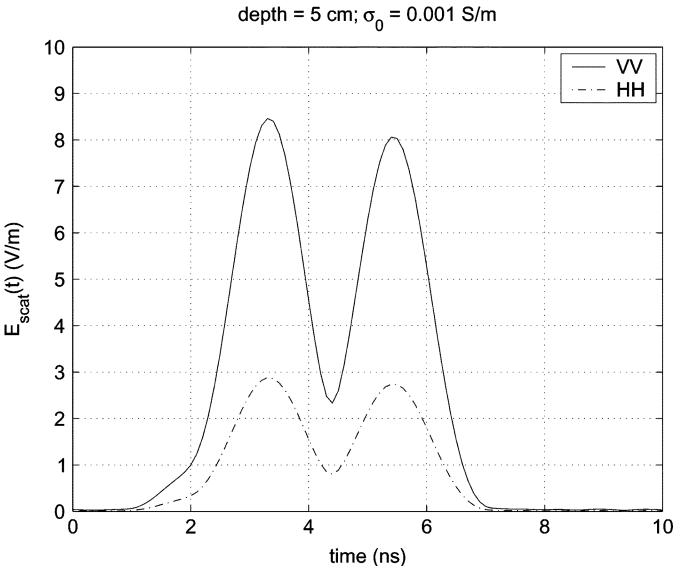

Fig. 3. Magnitudes of simulated mine reflection signatures for (dashed line) horizontally and (solid line) vertically polarized radars.

The low cross-range resolution arises because the distance traversed by the antennas when they scan across the top of the van is limited to the width of the van. The range resolution depends on the system bandwidth. The hardware is a stepped frequency design, capable of operating from about 0.3–3 GHz. However, experimentation has shown that the mines are most easily detected in the 0.3–1.9-GHz range. Only data in this range have been used to make the images, and this limits the range resolution.

B. Mine Signature

Physics-based simulations (e.g., [17]) have shown that buried mines should exhibit a "double humped" radar signature in the range direction. Fig. 3 shows a simulated range cut generated using the physical optics approximation, for both the HH and VV polarizations. Fig. 4 shows range cuts through image chips

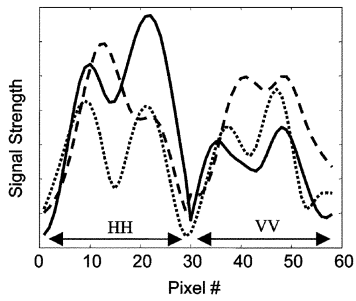

Fig. 4. Range cuts through imaged mines, for horizontally and vertically polarized radiation. The magnitude of the complex image is plotted.

for three different mines (of the same type), with the HH and VV polarized images stacked back to back. (The mine diameter is about 30 cm, and the images have been oversampled to 2-cm range bins.) The double humped signature is evident in the data.

However, Fig. 4 also shows considerable variation in the mine signature. The variation can be explained by several sources. First, a phase randomization process occurs when the signal is transmitted through the rough air–ground interface [17]. Another source of this variation is the inexact positioning in both depth and angle of the buried mines (even those of the same type). Due to these variations, under the H_1 hypothesis, we regard the mine signature as a deterministic physical signature that is transmitted through a random medium and measured subject to further additive noise at the sensor. With this in mind, we will use a set of training data to characterize the statistics under the H_1 hypothesis in Section IV.

Because the cross-range resolution of the radar is only about 40 cm, there is little observable cross-range structure in the mine images. In fact, even the known diameter of the mine does not provide a clearly identifiable feature, because any focused target will appear at least as wide as the cross-range resolution. For this reason, we will use a 1-D range cut through the image chip, which we will call a feature vector, as the input to our detector (instead of the full two-dimensional image chip). A range cut will contain the double hump mine signature, which appears to be the distinguishing feature of a buried mine in a SAR image and will result in a computationally simpler detection framework.[2]

C. Combination of Polarizations

The clutter return is at least partially uncorrelated across polarizations, while the return from a buried mine is well correlated. Therefore, using both polarizations should allow better discrimination between mine and clutter.

To use both polarizations we simply append the range cut from the VV channel to the range cut from the HH channel, as shown in Fig. 5. By applying the results of Section II we obtain the deflection-optimal fusion of the available data from both polarizations HH and VV.

As a preview, we mention here that the standard algorithm for combining radar data collected with different polarizations is the

[2]Note that even though we do not employ a cross-range signature, it is still important to have the (limited) cross-range resolution provided by the synthetic aperture. The focusing obtained from the synthetic aperature concentrates the target energy into a small region of the image. If this energy were spread out across the whole cross range of the image, it would be almost indistinguishable from the clutter.

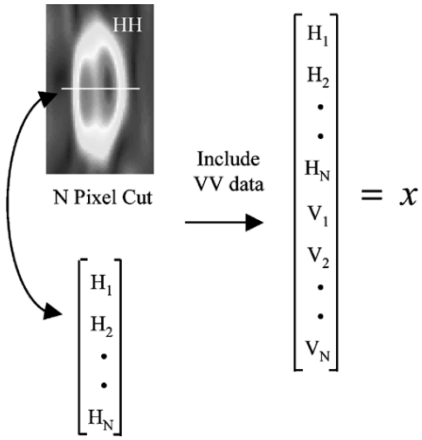

Fig. 5. Illustration of extraction of a feature vector $\mathbf{x}$. In the application, $N = 28$.

polarimetric whitening filter (PWF) [5]. Polarimetric whitening essentially combines the two polarizations by applying a cross-polarization whitening step followed by a quadratic form which does not use any information about the target signature. The performance of the PWF method is used as a benchmark against which our proposed algorithm can be compared.

IV. Field Test

A. Application of Algorithm

As mentioned earlier, buried mine detection field tests have been performed at two government test sites. These sites provide three distinct soil conditions: The soil at site one is dry and sandy, while site two has two test lanes, one of gravel, and the other of heavy clay soil. Conditions were wet in both lanes at site two due to recent rains, although, of course, the gravel provided more drainage than the clay. At each site, mines were buried at four depths: slightly below flush, 5, 10, and 15 cm deep.

For each of the three soil types, a set of training mines was chosen and used to form estimates $\hat{\mathbf{A}}$ and $\hat{\mathbf{b}}$ of the deflection-optimal detector parameters (details below). At site one, the training mines were chosen to be the 10-cm-deep mines, and the algorithm was tested on mines buried at the three remaining depths. This 10-cm depth was chosen for training because it was the intermediate depth between deep and shallow mines. At site two, the training mines were chosen to be the flush buried mines instead (due to practical constraints on identifying and labeling the training data in this lower SNR environment), and again the algorithm was tested on mines buried at the remaining depths. Below, as a specific example, we will consider the case of site one.

Using the H_1 training set[3] $S_1 = \{\tilde{\mathbf{x}}_1, \tilde{\mathbf{x}}_2, \ldots, \tilde{\mathbf{x}}_N\}$, we form the estimates of $\mathbf{A}$ and $\mathbf{b}$ as

$$\hat{\mathbf{A}} = \frac{1}{N-1}\left(\sum_{i=1}^{N} \tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^H\right) - \mathbf{I} \tag{7}$$

$$\hat{\mathbf{b}} = \frac{1}{N}\sum_{i=1}^{N} \tilde{\mathbf{x}}_i. \tag{8}$$

[3]We note that naturally, this training data was prewhitened by the clutter covariance matrix, estimated from H_0 training data, prior to the estimation of the parameters of the detector.

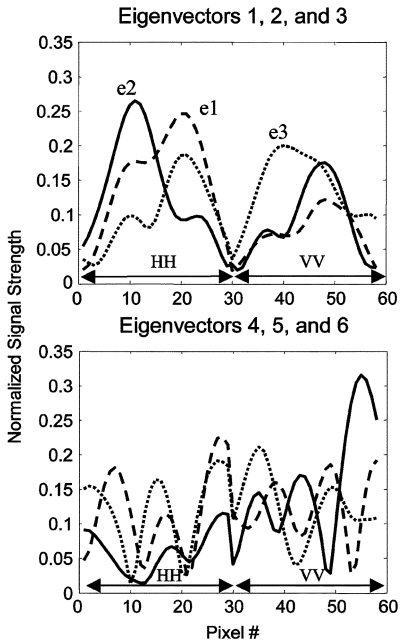

Fig. 6. First six eigenvectors for 10-cm-deep mines at site one.

It is worth mentioning here that this procedure for estimating $\mathbf{A}$ is not without risk. Namely, it is possible that the estimate should fail to be positive definite. In our experience in this particular case, the estimator never produced a negative-definite result. However, it is possible, but rather more complicated, to find better estimators for $\mathbf{A}$ by constraining the estimate to be positive definite. This important and nontrivial issue should be kept in mind, and more robust procedures for the estimation of the mean and covariance matrix [28] should be studied and applied within the context of any particular application.

With the above estimates in hand, we can write the detector structure by considering an eigendecomposition of $\hat{\mathbf{A}}$ as follows. Let

$$\hat{\mathbf{A}} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{U}^H = \sum_{j=1}^{M} \boldsymbol{\sigma}_j \mathbf{u}_j \mathbf{u}_j^H \tag{9}$$

where M is the dimension of the feature vector, and where we assume the eigenvalues $\boldsymbol{\sigma_1} \geq \boldsymbol{\sigma_2} \geq \cdots \geq \boldsymbol{\sigma_M} > 0$ are ordered.

The deflection-optimal detector derived from the training data can then be written as

$$T_d(\mathbf{x}) = \sum_{j=1}^{M} \boldsymbol{\sigma}_j \mathbf{x}^H \mathbf{u}_j \mathbf{u}_j^H \mathbf{x} + \mathbf{x}^H \hat{\mathbf{b}} + \hat{\mathbf{b}}^H \mathbf{x} \tag{10}$$

$$= \sum_{j=1}^{M} \boldsymbol{\sigma}_j |\mathbf{x}^H \mathbf{u}_j|^2 + 2\mathrm{Re}(\mathbf{x}^H \hat{\mathbf{b}}). \tag{11}$$

The above formulation shows one informative interpretation of this detector. Specifically, the eigendecomposition of the estimated $\hat{\mathbf{A}}$ amount to performing a principal components analysis on the set of training data S_1. The principal components then essentially identify a signal subspace, and (the first term of) the detector measures a weighted sum of the projection of any given data vector $\mathbf{x}$ onto the subspace basis vectors.

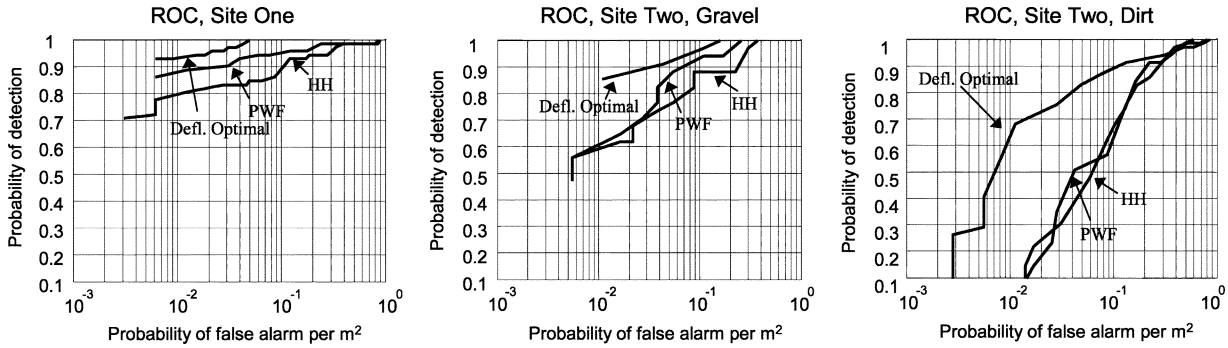

Fig. 7. ROC curves comparing the deflection-optimal processing results with polarimetric whitening, and with the HH polarized image alone, at a range of 15 m. For site one, mines buried just below the surface (5 and 15 cm deep) are included. For site two, mines buried 5, 10, and 15 cm deep are included. The mines not included (10 cm deep for site one, and just below the surface for site two) were used as training data.

An advantage of viewing the detector this way is that we may also consider a low-rank approximation to $\hat{\mathbf{A}}$ if the estimated eigenvalues $\boldsymbol{\sigma}_j$ are dominated by the first few largest ones. In effect, considering such a low-rank approximation is equivalent to retaining only the first few dominant principal components. This can lead to a more computationally efficient detector structure, while eliminating the effect of some possibly spurious and insignificant variations observed in the target signature measured in the training set S_1. Furthermore, the dominant principal components provide a least squares optimal [14] representation of essentially all possible variations in which a mine has appeared in the radar training images.

We found the above low-dimensional approximation approach to be very attractive in practice. To be more specific, considering the eigendecompoition of $\hat{\mathbf{A}}$, we implemented

$$T_d = \sum_{j=1}^{M'} \boldsymbol{\sigma}_j |\mathbf{x}^H \mathbf{u}_j|^2 + 2\mathrm{Re}(\mathbf{x}^H \hat{\mathbf{b}}) \tag{12}$$

where M' is much smaller than the dimension of the feature vector ($M = 56$ in this case). Specifically, we found the ratio $\boldsymbol{\sigma}_j/\boldsymbol{\sigma}_1$ of the six largest eigenvalues ($j = 1, \ldots, 6$) to the maximum eigenvalue to be approximately 1, 0.089, 0.075, 0.028, 0.019, and 0.014, respectively. (The corresponding eigenvectors are shown in Fig. 6.) Hence, we concluded that the first two or three eigenvalues, and the corresponding eigenvectors, of $\hat{\mathbf{A}}$ were dominant (i.e., we picked $M' = 2$ or 3).

We note that the structure of the first three dominant eigenvectors indicates that they should be sufficient to represent the signature of a buried mine in the training set accurately. Indeed, the very first eigenvector is an excellent candidate by itself.

B. Presentation of Results

Receiver operating characteristic (ROC) curves are plots of the probability of detection versus the probability of false alarm. ROC curves for the application of the trained detector to buried mines not in the training set, imaged at a range of 15 m, are shown in Fig. 7. The ROC curves were compiled using survey data for the buried mines (the ground truth). After a detector was applied to a given image, a "likelihood map" consisting of the values of the detector output at each pixel location was produced. From this (real-valued) image, the peak pixels in 1-m square boxes around expected mine locations were compiled in histograms, and associated with the H_1 hypothesis. The peak pixels[4] in 1-m square boxes offset from the expected mine locations were also compiled into histograms, and associated with the H_0 hypothesis. As one might expect, it was possible to compile far more data for the H_0 hypothesis than for the H_1 hypothesis. The probability of detection and probability of false alarm per square meter were computed by stepping a threshold through the histograms bin by bin, and computing the number of entries above the threshold, divided by the total number of entries, for each histogram. The jumps in some of the curves are due to the relatively small number of mines in the sample. There were 72 mine images in the sample at site one, and 34 (each, for dirt and gravel) at site two.[5] The curves end at the lowest probability of false alarm, greater than zero, that could be computed given the number of test samples.

For comparison, ROC curves are shown also for the case where the magnitude of horizontally polarized images was thresholded directly, and for the so-called polarimetric whitening filter [5] algorithm for combining the horizontally and vertically polarized images. Polarimetric whitening involved performing a linear transformation that decorrelates the clutter across polarizations (hence the term polarimetric whitening) and then combines the resulting components using a sum of squares, which amounts to an energy computation.

Fig. 7 shows that the deflection-optimal detector significantly outperforms the standard PWF algorithm and direct thresholding of the the HH images in all three soil types. The dry sandy soil at site one gave the best overall radar performance. The heavy clay soil at site two (which appeared to have a high moisture content) gave the worst overall radar performance.

V. Conclusion

Based on a limited set of training data, we developed a framework for the detection of targets against a noise/clutter background using the deflection-optimal linear-quadratic detector. A key idea explored and used in this paper is that the deflection-optimal detector is advantageous for two complementary reasons. First, it is well defined when no specific pdf is available under H_1. Second, when a small number of training samples is given,

[4]As observed by a reviewer, using a more stable order statistics such as a high percentile would perhaps be better, though we observed satisfactory results with the peak pixel as well.

[5]Due to the random nature of the ground–air interface, images of the same mine from opposite directions were considered to be distinct.

the deflection-based detector involves the unknown detector parameters more directly in that unlike the NP test statistic, the deflection approach does not involve the inversion of an estimate of the signal covariance matrix.

While the methods developed in this paper are more broadly applicable, the specific case studied here uses observations collected from a forward-looking ground-penetrating synthetic aperture radar developed by SRI International. The deflection-optimal detector allows the use of the singular value decomposition applied to collected training data to model the statistical variations of target signatures using a small set of spanning basis vectors. The resulting detection algorithm was tested on a complementary set of collected data, not in the training set. It was demonstrated that the proposed approach is effective for the detection of buried metal mines, realizing detection rates of as high as 90% for false-alarm rates of $0.01/\text{m}^2$.

Several important topics of research remain open. Combining physics-based modeling with collected training data should provide a more practical and accurate way to estimate the target statistics. Such a procedure can involve calibrating the physics-based model to the training data and then using the calibrated model to extrapolate to other soil conditions. Beyond the deflection-optimal detector, there exist other approaches to robust detector design (e.g., [25]), and these should also be considered. From a broader perspective, it remains to be seen if the proposed methods can be extended to apply to the detection of buried *plastic* mines. The signature of plastic mines is significantly weaker than that of metal mines, so a direct application of the techniques presented here may not be practical.

REFERENCES

[1] J. Kositsky, R. Cosgrove, C. Amazeen, and P. Milanfar, "Results from a forward-looking GPR mine detection system," *Proc. SPIE*, vol. 4742, Apr. 2002.

[2] M. L. Skolnik, *Introduction to Radar Systems*, 2nd ed. New York: McGraw-Hill, 1980.

[3] D. M. Silva, I. E. Abdou, and R. E. Warren, "Optimum detection of small targets in a cluttered background," *Opt. Eng.*, vol. 37, no. 1, pp. 83–92, Jan. 1998.

[4] R. Pagnet, J. Homer, and D. Crisp, "Automatic target detection in synthetic aperture radar imagery via terrain recognition," in *Proc. 2001 Int. Con. Image Processing*, Thessaloniki, Greece.

[5] L. M. Novak, M. C. Burl, and W. W. Irving, "Optimal polarimetric processing for enhanced target detection," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 29, pp. 234–244, Jan. 1993.

[6] C. W. Therrien, T. F. Quatieri, and D. E. Dudgeon, "Statistical model-based algorithms for image analysis," *Proc. IEEE*, vol. 74, pp. 532–551, Apr. 1986.

[7] K. C. Ho and P. D. Gader, "A linear prediction land mine detection algorithm for hand held ground penetrating radar," *IEEE Trans. Geosci. Remote Sensing*, vol. 40, pp. 1374–1385, June 2002.

[8] A. van der Merve and I. J. Gupta, "A novel signal processing technique for clutter reduction in GPR measurements of small, shallow land mines," *IEEE Trans. Geosci.Remote Sensing*, vol. 38, pp. 2627–2637, Nov. 2002.

[9] L. M. Novak, G. J. Owirka, and C. M. Netishen, "Radar target identification using spatial matched filters," *Pattern Recognit.*, vol. 27, no. 4, pp. 607–617, 1994.

[10] L. M. Novak, G. J. Owirka, and A. L. Weaver, "Automatic target recognition using enhanced resolution SAR data," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 35, pp. 157–175, Jan. 1999.

[11] M. D. DeVore and J. A. O'Sullivan, "Performance complexity study of several approaches to automatic target recognition from SAR images," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 38, pp. 632–648, Apr. 2002.

[12] D. Waagen, J. Pecina, and R. Pickens, "Evolving spatially-localized projection filters for SAR automatic target recognition," in *Proc. 7th Int. Conf. Evolutionary Programing*, 1998.

[13] L. A. Chan, N. M. Nasrabadi, and D. Torrieri, "Discriminative eigen targets for automatic target recognition," *Proc. SPIE*, vol. 3371, pp. 334–345, 1999.

[14] A. K. Shaw and V. Bhatnager, "Automatic target recognition using eigen-templates," *Proc. SPIE*, vol. 3370, pp. 448–459, 1998.

[15] Q. Zhao and J. C. Principe, "Support vector machines for SAR automatic target recognition," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 37, pp. 643–654, Apr. 2001.

[16] X. Yu, L. E. Hoff, I. S. Reed, A. M. Chen, and L. B. Stotts, "Automatic target detection and recognition in multiband imagery: A unified ML detection and estimation approach," *IEEE Trans. Image Processing*, vol. 6, pp. 143–156, Jan. 1997.

[17] K. F. Casey and G. N. Oetzel, "Physical-optics models for scattering from buried mines," in *Proc. 2nd Austral.–Amer. Joint Conf. Technologies of Mine Countermeasures*, Sydney, Australia, 2001.

[18] S. M. Kay, *Fundamentals of Statistical Signal Processing: Detection Theory*. Uppser Saddle River, NJ: Prentice-Hall, 1998.

[19] C. R. Baker, "Optimum quadratic detection of a random vector in Gaussian noise," *IEEE Trans. Commun.*, vol. COM-14, pp. 802–805, Dec. 1966.

[20] P. Chevalier and B. Picinbono, "Complex linear-quadratic systems for detection and array processing," *IEEE Trans. Signal Processing*, vol. 44, pp. 2631–2633, Oct. 1996.

[21] M.-R. Shikh-Bahaei, A. H. Aghvami, A. Ghorashi, and N. Ali-Akbarian, "A statistical processing approach to interference cancellation in W-CDMA systems," *IEICE Trans. Commun.*, vol. E83-B, no. 8, pp. 1619–1630, Aug. 2000.

[22] B. Picinbono and P. Duvaut, "Optimal linear-quadratic systems for detection and estimation," *IEEE Trans. Inform. Theory*, vol. 34, pp. 304–311, Mar. 1988.

[23] B. Picinbono, "On deflection as a performance criterion in detection," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 31, pp. 1072–1081, July 1995.

[24] W. G. Carrara, R. S. Goodman, and R. M. Majewski, *Spotlight Synthetic Aperture Radar Signal Processing Algorithms*. Norwood, MA: Artech House, 1995.

[25] H. V. Poor, "Robust decision design using a distance criterion," *IEEE Trans. Inform. Theory*, vol. IT-26, pp. 575–587, Sept. 1980.

[26] R. Cosgrove, "Background clutter whitening for forward-looking ground-penetrating radar images," SRI Int., Menlo Park, CA, Tech. Rep., 2003.

[27] J. D. Taft and N. K. Bose, "Quadratic-linear filters for signal detection," *IEEE Trans. Signal Processing*, vol. 39, pp. 2557–2559, Nov. 1991.

[28] R. A. Maronna and R. H. Zamar, "Robust estimates of location and dispersion for high-dimensional datasets," *Technometrics*, vol. 44, no. 4, pp. 307–317, Nov. 2002.

Russell B. Cosgrove, photograph and biography not available at the time of publication.

Peyman Milanfar (SM'98) received the B.S. degree in electrical engineering and mathematics from the University of California, Berkeley, and the S.M., E.E., and Ph.D. degrees in electrical engineering from the Massachusetts Institute of Technology, Cambridge, in 1988, 1990, 1992, and 1993, respectively.

Until 1999, he was a Senior Research Engineer at SRI International, Menlo Park, CA. He is currently an Associate Professor of electrical engineering, University of California, Santa Cruz. He was a Consulting Assistant Professor of computer science at Stanford University, Stanford, CA, from 1998 to 2000, where he was also a Visiting Associate Professor from June to December 2002. His technical interests are in statistical signal and image processing and inverse problems.

Dr. Milanfar won a National Science Foundation CAREER Award in 2000, and he was an Associate Editor for the IEEE SIGNAL PROCESSING LETTERS from 1998 to 2001.

Joel Kositsky, photograph and biography not available at the time of publication.