Frequency Subband Processing and Feature Analysis of Forward-Looking Ground-Penetrating Radar Signals for Land-Mine Detection

Tsaipei Wang, *Member, IEEE*, James M. Keller, *Fellow, IEEE*, Paul D. Gader, *Senior Member, IEEE*, and Ozy Sjahputera, *Member, IEEE*

***Abstract*—There has been significant amount of study on the use of ground-penetrating radar (GPR) for land-mine detection. This paper presents our analysis of GPR data collected at a U.S. Army test site using a new approach based on frequency subband processing. In this approach, from the radar data that have over 2.5 GHz of bandwidth, we compute separate radar images using the one wide (2 GHz) and four narrow (0.6 GHz) frequency subbands. The results indicate that signals for different frequency subbands are significantly different and give very different performance in land-mine detection. In addition, we also examine a number of features extracted from the GPR data, including magnitude and local-contrast features, ratio between copolarization and cross-polarization signals, and features obtained using polarimetric decomposition. Feature selection procedures are employed to find subsets of features that improve detection performance when combined. Results of land-mine detection, including performance on blind test lanes, are presented.**

***Index Terms*—Feature extraction, feature selection, frequency subbands, ground-penetrating radar (GPR), land-mine detection.**

I. Introduction

LAND mines around the world pose a serious problem as they continue to cause serious injuries and loss of lives for both civilians and soldiers. Detection and removal of land mines is a serious issue to be addressed. Over the past several years, a variety of technologies have been applied to land-mine detection. One technology that has been studied extensively is ground-penetrating radar (GPR), which is able to detect both metal mines and mines of low metal content.

Manuscript received February 21, 2005; revised March 17, 2006. This work was supported in part by the U.S. Army CECOM under Contract DAAB15-00-C-1021/3.

T. Wang was with the Department of Electrical and Computer Engineering, University of Missouri–Columbia, Columbia, MO 65211 USA. He is now with the Department of Computer Science, National Chiao Tung University, Hsinchu 300, Taiwan, R.O.C. (e-mail: wangts@cs.nctu.edu.tw).

J. M. Keller is with the Department of Electrical and Computer Engineering, University of Missouri–Columbia, Columbia, MO 65211 USA (e-mail: kellerj@missouri.edu).

P. D. Gader is with the Department of Computer and Information Science and Engineering, University of Florida, Gainesville, FL 32611 USA (e-mail: pgader@cise.ufl.edu).

O. Sjahputera was with the Department of Electrical and Computer Engineering, University of Missouri–Columbia, Columbia, MO 65211 USA. He is now with the Ellis Fischel Cancer Research Laboratory, University of Missouri–Columbia, Columbia, MO 65211 USA (e-mail: sjahputerao@missouri.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TGRS.2006.888142

Most studies on GPR-based land-mine detection, so far, have focused on either hand-held or vehicle-mounted downward-looking systems. In such systems, the radar signal propagates in a direction close to vertical to the ground surface. A good performance has been reported for such systems, especially given recent progress in radar technology [1], [2].

Compared with downward-looking systems, forward-looking GPR (FLGPR) systems are quite new. Currently, there exist forward-looking systems using IR imaging [3], [4], acoustic [5], and GPR [6]–[9] mine-detection technologies. These systems employ sensors that look ahead of the vehicle, with a standoff distance (the distance between the vehicle and the mine when it is detected) of up to a few tens of meters. Forward-looking GPR has also been applied to the detection of side-attack mines [6]. Existing forward-looking GPR systems are vehicle-mounted synthetic aperture radar (SAR) systems, with antennas that either move along a rail [7] or arranged in an array [9].

This paper presents the analysis of forward-looking GPR data collected by SRI International. The raw radar data are made available to us by SRI. The data set is divided into "runs." During each run, the vehicle is driven continuously along one of the test mine lanes at a U.S. Army facility that contain a variety of metal and plastic antitank mines. The data set includes "calibration" and "blind" runs. For calibration runs, the locations of targets (the ground truth), including actual mines and mine simulants (SIMs), are made available to us. This allows us to design the detection algorithm with labeled data. For blind runs, the ground truth is not known to the authors nor to SRI. The blind data runs are processed using the same algorithm that was generated and tested on the calibration data. Our results on the blind data, which consist of lists of suspected mine locations with associated confidences, are sent to an independent agency, the Institute for Defense Analyses (IDA), for scoring. We received from IDA the receiver operating characteristic (ROC) curves of the blind test, which we included in the results section of this paper.

There exist a number of algorithms for mine detection using forward-looking GPR data [8], [10]–[13]. A number of them focus on exploiting high-resolution spatial features [8], [10], [11], which requires a wide-frequency-bandwidth signal to give the desired range resolution. These can be considered generally as "spatial-domain" or "time-domain" techniques since the range information is derived from the timing between the transmitted and received signals.

In this paper, we present a different approach for analyzing wide-frequency-band GPR data for land-mine detection that can be considered "quasi-frequency-domain." This work was introduced in [13] but is expanded here. Frequency dependence has been identified in downward-looking systems [14], where mines are found to give stronger responses at certain frequency ranges. For the forward-looking system, our approach is to divide the whole frequency range into narrower subbands and create separate SAR images for each individual frequency subband. The motivation here is that signal levels from mines and clutter may very likely have different frequency dependence and that we can exploit these differences by focusing on one or more frequency subbands that give the best contrast between mine and clutter signals. On the other hand, the use of narrower frequency subbands results in lower range resolution and makes it less likely to extract detailed spatial features of mine or clutter signatures. In this regard, we can consider our approach and existing "time-domain" approach to be complementary to each other. The term "quasi-frequency-domain" is used because, although we explore frequency dependence, we are not truly analyzing individual frequency components in the radar signal.

Our algorithm consists of components that create SAR images for specified frequency subbands, preprocessing that removes background, an energy-based prescreener, feature extraction and selection, and the classifier itself. A total of ten features, including two different approaches of polarimetric decomposition, are analyzed. Our results are very good on both calibration and blind lanes for metal mines and provide the best detection of plastic mines to date.

II. Description of Algorithm

The radar system consists of two transmitting and 18 receiving antennas mounted at the front of the vehicle at about 5 m above the ground. During each run, while the vehicle is driven continuously along a mine lane, data are recorded at intervals of 1–3 m; each of such data recording is called a "scan." The quad-ridge horn antennas can operate in either horizontal (H) or vertical (V) polarization. We define a polarization mode as one combination of transmitter and receiver polarizations. For example, the polarization mode "HV" means transmitting in H and receiving in V polarizations. In each scan, 36 sets of values (one for each combination of the two transmitting and 18 receiving antennas) are recorded for each of the three polarization modes: HH, VV, and HV. Every one of the 36 sets contains received radar signals at 1024 frequencies (442.5 MHz to 3 GHz in 2.5-MHz steps). A picture of the system is shown in Fig. 1. The radar antennas are tilted toward the ground by 15°. The rectangle on the ground is just a visual aid to indicate the approximate imaged area. The closest range is approximately 7 m and is determined by the need to avoid overlap between transmitted and received pulses. Readers can refer to the study in [7] for more information regarding the antennas.

A scan image is a SAR reconstruction for a given scan and polarization mode, formed by the coherent combination of its 36 sets of data. We use the array $\nu_k(x, y, P, F)$ to represent complex pixel values in all of the scan images of a given run.

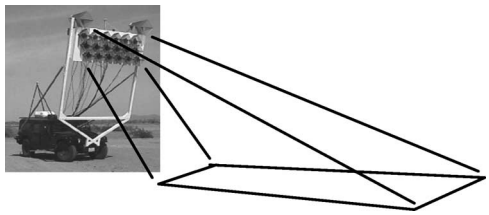

Fig. 1. Schematic of the radar system and imaged area.

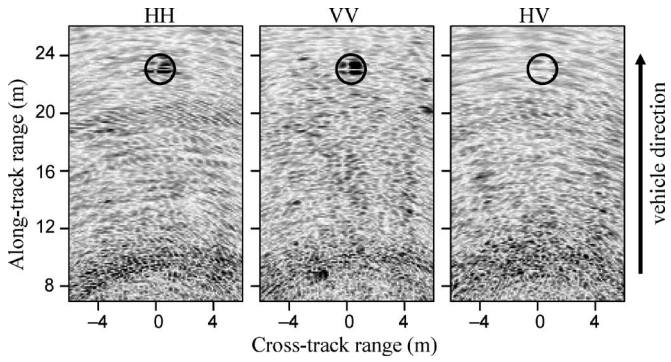

Fig. 2. Example of scan images for the same location and frequency subband (F_0), but different polarization modes.

Here, x and y are pixel coordinates within a scan image and are relative to the vehicle, with x and y being the directions perpendicular and parallel to the vehicle's motion, respectively. The parameter k represents the index of the scan within a run, P the polarization mode, and F the frequency range (subband) of the data used in creating the image. We investigate five different frequency subbands in this paper, represented by F_0 (0.8−2.8 GHz), F_1 (0.75−1.35 GHz), F_2 (1.25–1.85 GHz), F_3 (1.75−2.35 GHz), and F_4 (2.25−2.85 GHz). The study of different frequency subbands was originally inspired by previous observations of frequency dependence in downward-looking GPR data. The subbands F_1 to F_4 approximately divide the whole frequency range into four parts while allowing some overlap. Their frequency bandwidth of 0.6 GHz is selected such that the range resolution estimated through the simulation is approximately 25 cm, which is around the typical size of antitank mines. The choice of these four subbands involves the tradeoff between two considerations: more narrower subbands to give more detailed information on frequency dependence and wider subbands for better spatial resolution.

Each scan image we generate covers a spatial region 19 m long (7–26 m from the front of the vehicle, in the direction of the vehicle's motion) and 12 m wide (6 m on each side from the center of the vehicle) at 5-cm resolution. Figs. 2–4 contain examples of scan images. Fig. 2 displays scan images of the same spatial region and frequency subband, but with different polarization modes. In Fig. 3, there are scan images of the same spatial region and polarization mode, but with different frequency subbands. Finally, Fig. 4 shows the images of five consecutive scans for the same frequency subband and polarization mode. Each scan image contains a surface-laid plastic mine indicated by the circle (2-m diameter).

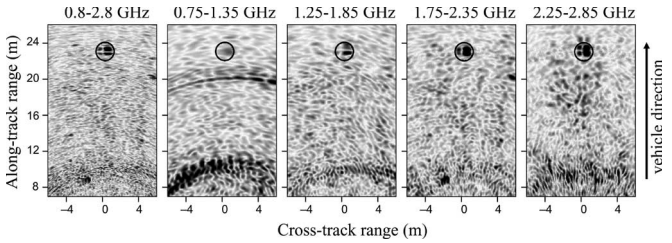

Fig. 3. Example of scan images for the same location and polarization mode (VV), but different frequency subbands.

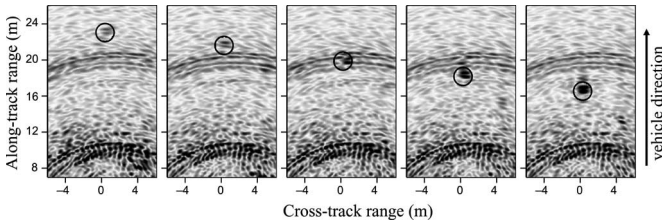

Fig. 4. Example of five consecutive images for the same frequency subband (F_1) and polarization mode (HH).

A. Background Removal and Normalization

This section describes the procedure of removing image features that result only from properties of the radar system and are independent of the vehicle's location. This procedure is further divided into two steps: the estimation and subtraction of self-signatures, and the estimation and normalization of signal magnitude variations (gain background). To simplify notation, we use $\nu_k(x, y)$ instead of $\nu_k(x, y, P, F)$ to represent pixel values in all the scan images, as all background removal and normalization procedures involve only data of the same polarization mode and frequency subband. The parameter k represents the index of the scan within a run. We use $m_k(x, y)$ and $\theta_k(x, y)$ to represent the magnitude and phase of $\nu_k(x, y)$

$$\nu_k(x, y) = m_k(x, y) \exp\left[i\theta_k(x, y)\right]. \tag{1}$$

Self-signatures are complex patterns that likely result from interaction among radar antennas and other components of the vehicle. These patterns appear to be mostly identical in all scan images of the same polarization and frequency band, regardless of the motion of the vehicle. This is evident in Fig. 4, where we see that the mine signature moves between images due to vehicle motion while the background patterns are stationary. We estimate self-signatures by averaging all the scan images of the same polarization and frequency band in the same run. This is done coherently, i.e., we use scan images containing complex numbers. However, to reduce the effect on self-signature estimation by other strong signals, such as those from surface-laid mines, we apply an upper limit on the magnitude of pixel values. The estimated self-signature is given by

$$B_1(x, y) = \frac{1}{N_{\text{scan}}} \sum_{k=1}^{N_{\text{scan}}} \frac{1}{r_k} \min\left[m_k(x, y), \mu_k + 5\sigma_k\right] \exp\left[i\theta_k(x, y)\right]. \tag{2}$$

Here, N_{scan} is the number of scans in the run, μ_k and σ_k are the mean and standard deviation of $m_k(x, y)$, respectively, and

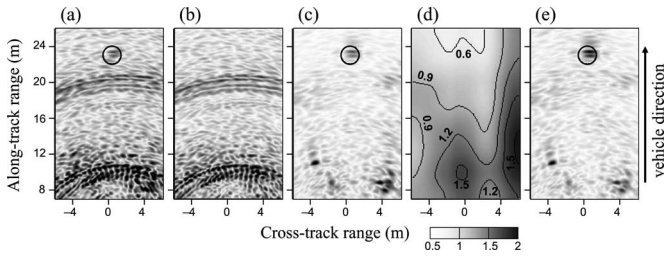

Fig. 5. Effect of background removal and normalization. (a) Original scan image (F_1, HH). (b) Estimated self-signatures. (c) Same as (a) after subtraction of (b). (d) Estimated gain background. (e) Same as (c) after normalization of gain background.

r_k is defined as the rms of the lowest 75% of $m_k(x, y)$ values. The magnitude of each scan image is divided by r_k before the averaging in (2) to remove scan-to-scan gain variations. Possible causes of such variations include noise in the raw data as well as ground properties such as moisture content. Since many of our features are obtained by aggregating their values in individual scan images using intensity-based weights, this normalization is necessary to avoid bias.

To estimate the "gain background," we first perform self-signature subtraction for each scan image

$$m'_k(x, y) = \left| \frac{1}{r_k} \nu_k(x, y) - B_1(x, y) \right|. \tag{3}$$

Next, we smooth $m'_k(x, y)$ by first applying a 21 × 21 (1 m × 1 m) median filter and then a 51 × 51 (2.5 m × 2.5 m) averaging filter; for simplicity, we continue to call this smoothed function $m'_k(x, y)$. The gain background is then given by

$$B_2(x, y) = \frac{1}{N_{\text{scan}}} \sum_k m'_k(x, y). \tag{4}$$

After we have estimated both the self-signatures and the gain background, we subtract the self-signatures from the scan images, followed by division with the gain background

$$m''_k(x, y) = \left| \frac{\frac{1}{r_k} \nu_k(x, y) - B_1(x, y)}{B_2(x, y)} \right|. \tag{5}$$

We illustrate in Fig. 5 the effect of background removal and normalization. Fig. 5(b) is the magnitude of the estimated self-signatures, i.e., $|B_1(x, y)|$, of a certain run. Fig. 5(a) is a scan image within this run, and Fig. 5(c) is the result of applying self-signature subtraction on Fig. 5(a), but before smoothing is applied. Fig. 5(d) is $B_2(x, y)$, the estimated gain background. Fig. 5(e) shows the corresponding $m''_k(x, y)$, which is Fig. 5(c) after the normalization of the gain background.

B. Prescreener

For our data set, the GPR vehicle typically moves forward by 1–3 m between successive scans. This distance is small compared with the size of computed scan images, resulting in the fact that a spatial location can appear in multiple scan images. It is therefore necessary to have a common coordinate system that allows us to identify the multiple occurrences, each

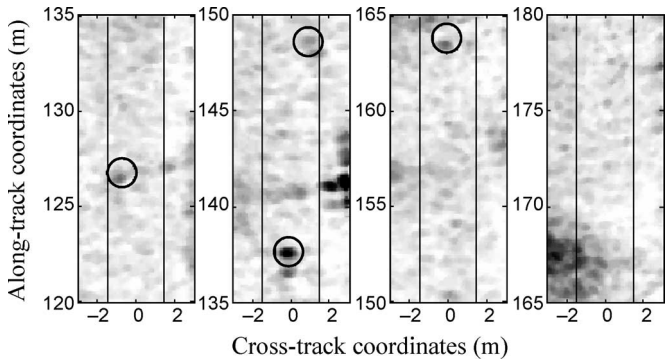

Fig. 6. Partial run-view image of a calibration lane for HH polarization mode and frequency subband F_1. The individual images are 6-m cross-track and 15-m down-track. The vertical lines represent lane boundaries, and the circles indicate targets in the ground truth.

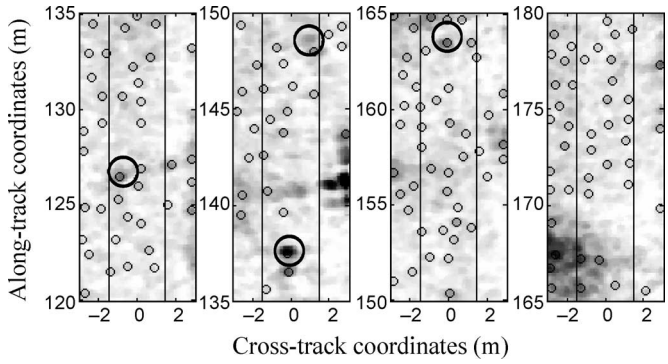

Fig. 7. Alarms (small circles) found in the same 60-m-long segment of a calibration-lane run-view image as Fig. 6. The vertical lines are lane boundaries.

in a different scan image, of a given spatial location. We choose to use the coordinate system of the mine lane being surveyed for convenience because these lanes are rectangular.

The purpose of the prescreener is to select a list of alarm locations for further analysis and classification. As each run of our data covers a single mine lane, we first construct a run-view image over a grid of spatial locations along the mine lane. More specifically, this grid covers an area 6 m wide and of the same length as the mine lane itself at 10-cm resolution; the width is selected to extend 1.5 m from each side of the 3-m-wide lane. For each run, we compute a run-view image for each combination of polarization mode and frequency subband.

To obtain the pixel value of a given spatial location $\boldsymbol{p}$ in the run-view image, we first identify all individual scan images that contain $\boldsymbol{p}$. We use only those scan images in which $\boldsymbol{p}$ appears within $Y_{\max}$ from the front of the vehicle. This is because we find that we can achieve better detection performance by focusing on the part of the scan images closer to the vehicle than when the entire down-track range is used. We then set the pixel value corresponding to $\boldsymbol{p}$ in the run-view image to be the second largest pixel value (magnitude only) corresponding to $\boldsymbol{p}$ in those scan images. This usage of "second largest value" is just a simple way to remove the effect of spurious bright pixels with minimal computational load. Currently, $Y_{\max}$ is set to 16 m to ensure that $\boldsymbol{p}$ is between 7 m (the closest distance to the vehicle in the scan images) and $Y_{\max}$ from the front of the vehicle in at least two to three scan images when the vehicle moves by 3 m between scans, which is the largest interscan distance in our data set. This is because two occurrences of $\boldsymbol{p}$ are needed to get the "second largest value." However, in rare cases, a spatial location $\boldsymbol{p}$ might still appear in only one scan image. This occurs when the vehicle moved out of the mine lane to avoid obstacles, causing an unusually large distance between the two scans just before and after the segment when the vehicle is outside of the lane. In these cases, we just use the pixel value corresponding to $\boldsymbol{p}$ in that one scan image as the pixel value of $\boldsymbol{p}$ in the run-view image.

Fig. 6 shows the run-view image of a 60-m-long segment in one calibration lane for HH polarization mode and frequency subband F1. The individual images are 6-m cross-track and 15-m down-track. The vertical lines represent the lane boundaries, and the circles indicate targets in the ground truth.

From a run-view image, we generate a list of alarms consisting of all pixels in the run-view image that are local maximas within a radius of r_{alarm} (currently set to 75 cm). The location and confidence value of an alarm are just the location and magnitude of its corresponding pixel in the run-view image. Pixels on the border of the run-view image are excluded. This can be considered an energy-based prescreener.

Further pruning of alarms is done with an adaptive thresholding scheme. The goal is to reduce the total number of alarms without having to manually select a threshold. Instead, the threshold is selected such that the total number of alarms retained after thresholding is $\rho A/(\pi r_{\text{alarm}}^2)$, where A is the total area under consideration (the area of the run-view image here) and ρ is a positive density factor. The result is that, on average, there remain ρ alarms per unit area of (πr_{alarm}^2). Currently, we use $\rho = 0.75$. Fig. 7 displays the same region as in Fig. 6, with the alarms shown as small circles.

For the calibration runs, since we know the locations of all targets, we are able to label the detected alarms. Alarms located within a radius of r_{label} (currently set to 1 m) from the centers of known targets (including mines and SIMs) are labeled as "hits." We associate each hit with its closest actual target. Alarms that are not hits but fall within the lane's boundary are labeled as "FAs"; "FA" stands for false alarm. Alarms that are neither hits nor FAs are not included in the subsequent analysis.

C. Analysis of Different Frequency Subbands

Up to this point, we have generated scan images and run-view images for individual frequency subbands. In Fig. 8, scan images of different frequency subbands are displayed side by side for the purpose of qualitative comparison. Each row in Fig. 8 is a set of five images of the same scan and polarization mode, but with different frequency subbands (from left to right, F_0, F_1, F_2, F_3, and F_4). The scan images here contain only regions that are 6 m wide (centered at the mine lane) and 8 m long (8–16 m in front of the vehicle). It appears that frequency subband F_1 exhibits the lowest level of clutter, in general.

For each run, we generate a run-view image for each combination of different frequency subbands and polarization modes. Every run-view image allows us to obtain a list of alarms. At this point, we are faced with the question of which frequency subband and polarization mode to use for this purpose. A simple method to compare different combinations of

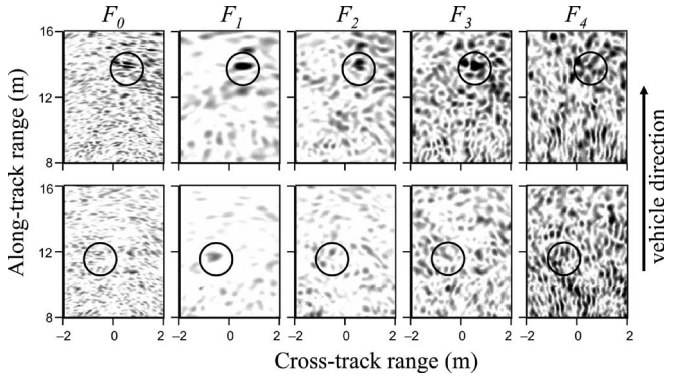

Fig. 8. Comparison of scan images at different frequency subbands. Each row contains 8-m-long and 6-m-wide scan images (after background removal and normalization) of the same scan and polarization mode, but different frequency subbands. Each scan image contains a buried mine indicated by the circle. (Top row) VV polarization mode with a 15-cm-deep metal mine. (Bottom row) HH polarization mode with a 5-cm-deep plastic mine.

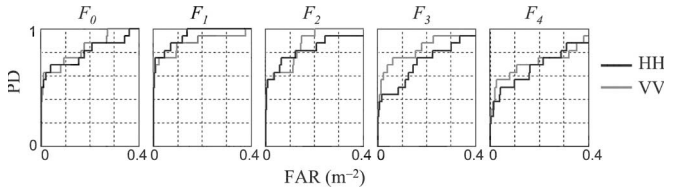

Fig. 9. ROC curves generated at the prescreener stage for the alarms in a calibration run.

polarization settings and frequency bands is to create ROC curves for the calibration-lane alarms generated with the various combinations. Hits associated with mines are considered detections, and FAs are considered false alarms; alarms outside of the lanes and hits associated with SIMs are excluded. In Fig. 9, we plot the ROC curves for different frequency band and polarization mode combinations. We can clearly see the trend of lower probability of detection (PD) versus false alarm rate (FAR) curves from F_1 to F_4. Currently, we choose HH polarization and frequency subband F_1 (0.75–1.35 GHz) for the prescreener because this combination appears to perform best according to the ROC curves in Fig. 8 and also gives 100% detection in all calibration runs with FAR ≤ 0.15 m^{-2}.

SIMs are designed to be similar to actual mines in terms of: 1) metal content; 2) thermal properties; and 3) dielectric properties of explosives (readers interested in SIMs are referred to the study in [20]). However, they usually do not match actual mines in terms of what actually makes FLGPR work, such as air gaps. As a result, FLGPR does not perform well when detecting SIMs, and they are usually excluded during scoring.

D. Feature Computation

1) Extraction of Patch Images: Before we construct the classifier to separate alarms corresponding to mines and false alarms, we have to obtain a set of features for each alarm. We compute the features of an alarm from extracted image patches centered at that alarm's location. Because of different requirements of different features, two sets of image patches are generated: 1) the incoherent set, which consists of 25 $\times$ 25, magnitude-only image patches (at 10-cm spatial resolution), and 2) the coherent set, which consists of 11 $\times$ 11, complex-

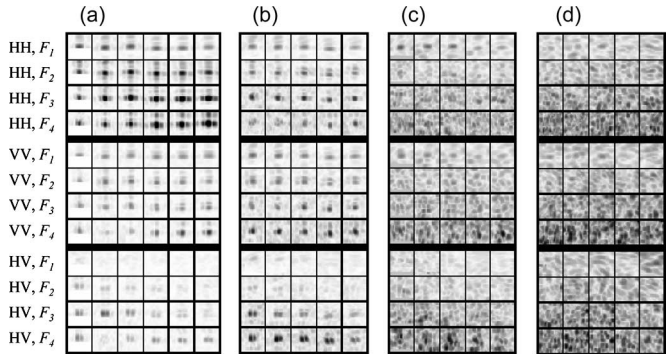

Fig. 10. Example image patches in the incoherent set. (a)–(d) are image patches that belong to different alarms. Each small rectangle is a 25 $\times$ 25 image patch. For each alarm, its image patches are arranged into 12 rows, each corresponding to a combination of polarization mode and frequency subband, and a number of columns, each corresponding to a distance between the alarm and the vehicle (from left to right, nearest to farthest). Only image patches for which the alarms are within $Y_{\max}$ (16 m) of the vehicle are shown, so the number of columns depends on how many times the alarms appear within that distance over multiple scan images. (a) Hit for a surface-laid mine. (b) Hit for a 5-cm-deep buried metal mine. (c) Hit for a 5-cm-deep buried plastic mine. (d) False alarm.

number image patches (at 5-cm spatial resolution). We use lower spatial resolution for the incoherent set simply to reduce the amount of computation, since the image patches in this set contain more pixels due to the need to compute region-based features (local contrasts, to be discussed later). Currently, the coherent set is used solely for the generation of features through polarimetric decomposition, also discussed later. These image patches are extracted from all the individual scan images that contain the alarm's location. Fig. 10 contains the image patches (the incoherent set) for four alarms. The first three alarms are associated with a surface-laid mine, a 5-cm-deep buried metal mine, and a 5-cm-deep buried plastic mine, respectively. The last one is not associated with any known target and is therefore a FA. For each alarm, the image patches are arranged from top to bottom according to the frequency-subband/polarization combination and from left to right according to their distances to the vehicle.

Pixel values in the extracted image patches are represented as $\nu_{(p)}(x, y, Y, P, F)$, with the subscript "(p)" indicating that we are referring to image patches. Here, similar to the case of scan images, they are identified by five parameters: x and y for the pixel coordinates within the patch in directions that are perpendicular and parallel to the vehicle, respectively, Y for the distance from the front of the vehicle to the alarm, P for the polarization mode ("HH," "VV," or "HV"), and F for the frequency subband. Since the feature calculation always involve only one frequency subband, we skip the parameter F in the following equations to make them more concise. In addition, each image patch of the incoherent set is divided into three regions (see Fig. 11): the inner region (R_{in}) consists of the central 5 $\times$ 3 pixels, the outer region (R_{out}) consists of pixels outside of the central 9 $\times$ 7 pixels, and the pixels between the inner and outer regions form a guard region.

2) Description of Features: We compute the following ten features for each alarm from the extracted image patches:

1) M_{VV}: magnitude in VV polarization mode;
2) C_{VV}: local contrast in VV polarization mode;

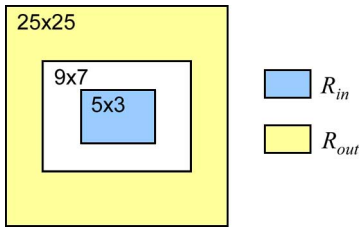

Fig. 11. Separate regions of an image patch in the incoherent set.

3) $R_{\mathrm{VV/HV}}$: ratio between VV and HV polarization modes;
4) M_{HH}: magnitude in HH polarization mode;
5) C_{HH}: local contrast in HH polarization mode;
6) $R_{\mathrm{HH/HV}}$: ratio between HH and HV polarization modes;
7) P_{H}: the feature H (entropy) in H/α polarimetric decomposition;
8) P_{α}: the feature α in H/α polarimetric decomposition;
9) P_{τ}: the feature "degree of scatterer symmetry" in A/S polarimetric decomposition;
10) P_{ψ}: the feature "scatterer rotation angle" in A/S polarimetric decomposition.

The first six (magnitude, local contrast, and ratio) features are derived from the incoherent set of image patches. In the following, we only describe the procedures for the HH polarization mode; the corresponding features for the VV polarization mode are obtained by simply replacing "HH" with "VV" in the equations. Let us define $m_{\mathrm{in}}(Y,P)$ and $\mu_{\mathrm{in}}(Y,P)$ as the maximum and mean value of $\nu_{(p)}(x,y,Y,P)$ for all $(x,y)\in R_{\mathrm{in}}$, respectively, and $\mu_{\mathrm{out}}(Y,P)$ and $\sigma_{\mathrm{out}}(Y,P)$ as the mean and standard deviation values of $\nu_{(p)}(x,y,Y,P)$ for all $(x,y)\in R_{\mathrm{out}}$, respectively. The magnitude feature is given by

$$M_{\mathrm{HH}}=\max_{Y\le Y_{\max}} m_{\mathrm{in}}(Y,\text{“HH”}). \tag{6}$$

The local contrast of an image patch is the difference between the inner and outer regions divided by the standard deviation in the outer region. This is actually what is commonly referred to as the constant FAR (CFAR) value. The local contrast feature of an alarm is the weighted average of the local contrast values in all the patches associated with that alarm

$$C_{\mathrm{HH}}=\frac{\sum_{Y\le Y_{\max}} m_{\mathrm{in}}(Y,\text{“HH”})c_{\mathrm{HH}}(Y)}{\sum_{Y\le Y_{\max}} m_{\mathrm{in}}(Y,\text{“HH”})} \tag{7}$$

where

$$c_{\mathrm{HH}}(Y)=\frac{m_{\mathrm{in}}(Y,\text{“HH”})-\mu_{\mathrm{out}}(Y,\text{“HH”})}{\sigma_{\mathrm{out}}(Y,\text{“HH”})}.$$

The ratio features $R_{\mathrm{HH/HV}}$ and $R_{\mathrm{VV/HV}}$ are motivated by the observation that clutter objects exhibit larger cross-polarization signals (i.e., HV) than actual mines. They are calculated as

$$R_{\mathrm{HH/HV}}=\frac{\sum_{Y\le Y_{\max}} m_{\mathrm{in}}(Y,\text{“HH”})\, r_{\mathrm{HH/HV}}(Y)}{\sum_{Y\le Y_{\max}} m_{\mathrm{in}}(Y,\text{“HH”})} \tag{8}$$

where

$$r_{\mathrm{HH/HV}}(Y)=\frac{\mu_{\mathrm{in}}(Y,\text{“HH”})}{\mu_{\mathrm{in}}(Y,\text{“HV”})}.$$

The four features involving polarimetric decomposition are derived from the coherent set of image patches. Polarimetric decomposition is the process of extracting physics-inspired features from the 2×2 scattering matrix

$$S=\begin{pmatrix} s_{\mathrm{hh}} & s_{\mathrm{hv}} \\ s_{\mathrm{vh}} & s_{\mathrm{vv}} \end{pmatrix}. \tag{9}$$

Each element in S represents the complex-valued backscattered signal for a specific combination of receiver and transmitter polarizations. For example, we use s_{hv} to represent signals with H transmitter and V receiver polarizations. However, we only have data for three polarization settings: HH, VV, and HV. It is common to assume $s_{\mathrm{vh}}=s_{\mathrm{hv}}$ based on reciprocity.

3) Computing Features From Polarimetric Decomposition: We implemented two different approaches for polarization decomposition, H/α and A/S. The following describes the procedure of computing features of H/α decomposition (please see [15] for more details of this technique). In H/α decomposition, two numbers, H and α, are obtained from the eigenvalues and eigenvectors of Q, which is a 3×3 matrix defined as

$$Q=\frac{1}{2}\langle \boldsymbol{q}^T\boldsymbol{q}^*\rangle,\ \text{where } \boldsymbol{q}=[s_{\mathrm{hh}}+s_{\mathrm{vv}}\quad s_{\mathrm{hh}}-s_{\mathrm{vv}}\quad 2s_{\mathrm{hv}}]. \tag{10}$$

We use the average over a 5×5 window for the $\langle\bullet\rangle$ operator.

In our case, the values of H and α are calculated for each (x,y,Y). We average them locally at the alarm's location by defining $H_{\mathrm{in}}(Y)$ and $\alpha_{\mathrm{in}}(Y)$ as the mean value of $H(x,y,Y)$ and $\alpha(x,y,Y)$ over R_{in}, respectively. Here, R_{in} also represents the central 5×3 region of the image patches. The feature values are obtained by weighted average

$$P_H=\frac{\sum_{Y\le Y_{\max}} m_{\mathrm{in}}(Y,\text{“HH”})H_{\mathrm{in}}(Y)}{\sum_{Y\le Y_{\max}} m_{\mathrm{in}}(Y,\text{“HH”})} \tag{11}$$

and

$$P_{\alpha}=\frac{\sum_{Y\le Y_{\max}} m_{\mathrm{in}}(Y,\text{“HH”})\alpha_{\mathrm{in}}(Y)}{\sum_{Y\le Y_{\max}} m_{\mathrm{in}}(Y,\text{“HH”})}. \tag{12}$$

The A/S polarimetric decomposition technique (please see [16] and [17] for more details) involves the following steps. The first task is to estimate S_{rec}, the part of the scattering matrix S that corresponds to a reciprocal scatterer. Since we have assumed reciprocity, we always have $S=S_{\mathrm{rec}}$. The second task is to estimate S_{sym}, the largest component of S_{rec} that corresponds to a symmetric scatterer. The characteristic of a scattering matrix corresponding to a symmetric scatterer is that it is diagonalizable by a 2×2 rotation matrix R_{ψ}, with ψ being the rotation angle. S_{sym} is given by

$$S_{\mathrm{sym}}=(S_{\mathrm{rec}}\bullet S_a)S_a+(S_{\mathrm{rec}}\bullet S')S' \tag{13}$$

where

$$S_a=\frac{1}{\sqrt{2}}\begin{bmatrix}1 & 0\\ 0 & 1\end{bmatrix}\quad S'=\frac{1}{\sqrt{2}}\begin{bmatrix}\cos\theta & \sin\theta\\ \sin\theta & -\cos\theta\end{bmatrix}$$
$$\text{with } \theta=\arg\max_{\theta'}\left|S_{\mathrm{rec}}\bullet\begin{bmatrix}\cos\theta' & \sin\theta'\\ \sin\theta' & -\cos\theta'\end{bmatrix}\right|. \tag{14}$$

The "inner product" operator between two 2×2 matrices, "$\bullet$," is defined as the inner product when the two matrices are treated as 4×1 vectors.

One of the two features derived from A/S polarimetric decomposition is the "degree of scatterer symmetry"

$$\tau = \cos^{-1}\left[\frac{|S_{\rm sym} \bullet S_{\rm rec}|}{\|S_{\rm sym}\| \times \|S_{\rm rec}\|}\right]. \tag{15}$$

Here, the $\|\cdot\|$ operator of a 2×2 matrix is its length when treated as a 4×1 vector.

The other feature is the rotation angle ψ that diagonalizes $S_{\rm sym}$. Due to the existence of angular ambiguity, we apply the following restrictions following [16]: 1) $-\pi/2 < \psi \leq \pi/2$ and 2) ψ is selected such that the (1,1) element of $S_{\rm sym}$, after being diagonalized, is not smaller than its (2,2) element.

Similar to the case of H/α decomposition, the values of τ and ψ are calculated for each (x, y, Y). We also average them locally at the alarm's location by defining $\tau_{\rm in}(Y)$ and $\psi_{\rm in}(Y)$ as the mean value of $\tau(x, y, Y)$ and $\psi(x, y, Y)$ over $R_{\rm in}$, respectively. The corresponding features are then obtained through weighted averaging over Y

$$P_\tau = \frac{\sum_{Y \leq Y_{\max}} m_{\rm in}(Y, \text{"HH"})\tau_{\rm in}(Y)}{\sum_{Y \leq Y_{\max}} m_{\rm in}(Y, \text{"HH"})} \tag{16}$$

and

$$P_\psi = \frac{\sum_{Y \leq Y_{\max}} m_{\rm in}(Y, \text{"HH"})\psi_{\rm in}(Y)}{\sum_{Y \leq Y_{\max}} m_{\rm in}(Y, \text{"HH"})}. \tag{17}$$

E. Data Duplication and Feature Normalization

One property of our training data is that there are many more data points labeled as FAs than those labeled as hits. We handle this by duplicating alarms that are considered hits when training the classifier. We define a function $W(\boldsymbol{x})$ to represent the ratio of duplication for an alarm (feature vector) $\boldsymbol{x}$. We always have $W(\boldsymbol{x}) = 1$ for FAs as they are not duplicated. For hits, it is calculated as

$$W(\boldsymbol{x}) = W_0 \beta(\boldsymbol{x}) w(\boldsymbol{x}) \tag{18}$$

where

$$w(\boldsymbol{x}) = \frac{w'(\boldsymbol{x})}{\sum_{T(\boldsymbol{x})=T(\boldsymbol{x}')} w'(\boldsymbol{x}')} \text{ with } w'(\boldsymbol{x}) = \exp\left[-\frac{d_{\rm tgt}(\boldsymbol{x})^2}{r_{\rm label}^2}\right]. \tag{19}$$

W_0 is the ratio between the number of FAs and the number of actual targets. Its value is approximately 30 for the training data. $\beta(\boldsymbol{x}) = 1$ if $\boldsymbol{x}$ is associated with an actual mine and 0.5 if $\boldsymbol{x}$ is associated with a SIM. The purpose of the term $\beta(\boldsymbol{x})$ is to differentiate between actual mines and SIMs because of possible differences between features of alarms associated with these two types of targets. Even though we are only interested in detecting actual mines, due to the small number of hits available for training purpose, we choose to give smaller weights to alarms associated with SIMs instead of excluding them altogether.

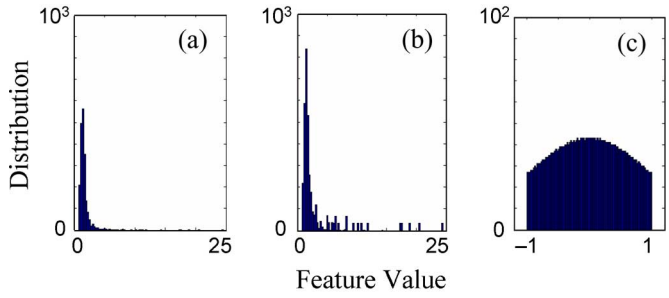

Fig. 12. Example of the effect of data duplication and feature value normalization on feature value distribution. (a) Original histogram of $M_{\rm HH}$ values (all calibration runs) before data duplication and feature normalization. (b) Same as (a) but after data duplication. (c) Histogram after feature normalization.

An assumption made for the training data is that a hit actually results from a target in the mine lane. However, a complication is that sometimes there are two or more alarms associated with the same target. In such cases, we are less confident about any of the individual alarms actually resulting from the target. We use the weighting term $w(\boldsymbol{x})$ in (18) and (19) to handle this reduced confidence. In (19), we make the $w(\boldsymbol{x})$ values of all the alarms associated with the same target add to one. Here, we use $T(\boldsymbol{x})$ to represent the actual target associated with a hit $\boldsymbol{x}$. When more than one alarm are associated with the same target, the $w(\boldsymbol{x})$ values for these alarms are determined by their Euclidean spatial distance to the ground truth location of the target, $d_{\rm tgt}(\boldsymbol{x})$. An alarm with smaller $d_{\rm tgt}(\boldsymbol{x})$ gets larger $w(\boldsymbol{x})$ and vice versa. For example, assume $r_{\rm label} = 1$ m, and there are two alarms associated with the same target, with $d_{\rm tgt}(\boldsymbol{x})$ being 0.2 and 0.9 m, respectively. Using (19), we calculate $w(\boldsymbol{x})$ of these two alarms to be 0.68 and 0.32, respectively. After we have computed $W(\boldsymbol{x})$ for all alarms, we replace each $W(\boldsymbol{x})$ with the larger value of its closest integer and one. Now, we can create a new training set in which there are $W(\boldsymbol{x})$ instances of $\boldsymbol{x}$.

Many of the features have very skewed distribution. For example, in Fig. 12(a) and (b), we plot the histogram of $M_{\rm HH}$ in all the calibration runs before and after data duplication, respectively. We see that the distribution is concentrated at smaller $M_{\rm HH}$ in both plots. We apply a feature normalization procedure that is similar to histogram specification. The goal is to rearrange the values of a given feature such that the distribution is similar to a given function. We choose to use a Gaussian between -1 and 1, with a standard deviation of one. The resulting distribution is shown in Fig. 12(c). The pseudocode of the procedure is given below. It produces a list of (a_i, b_i) pairs, i.e., the pairs of pre- and postnormalization feature values.

Let $a_1, a_2, \ldots, a_N$ be the sorted values
  (small to large) of a single feature
Let $f(a)$ be the new distribution
$A \leftarrow$ area under $f(a)$ between -1 and 1
$d_A \leftarrow A/(N-1)$
$b_1 \leftarrow -1$
for $k = 2$ to N
  $b_k \leftarrow b_{k-1} + d_A/(f(a_k) + f(a_{k-1}))$
end for
return $b_1, b_2, \ldots, b_N$ as the new feature values

When we have separate training and testing/validation data, both data duplication and feature normalization involves only data in the training set. The "normalized" values of a feature in the testing or validation data are obtained through interpolation/extrapolation using the (a_i, b_i) pairs obtained from the training data.

F. Classifier Design, Scoring, and Feature Selection

To use as few free parameters as possible in our classifier so that it will not be overtrained, we choose to use Fischer's linear discriminant (FLD) for the classifier design. The basic idea is to project all data points onto a single direction in the feature space, giving one number (the confidence) for each data point, such that a class separation measure is optimized for data points in the training set. This is described in the following.

Let there be two classes C_1 and C_2, and feature vectors $\boldsymbol{x}_k$ $(1 \leq k \leq N)$ that have been labeled as either C_1 or C_2. FLD selects a vector $\boldsymbol{y}_{\mathrm{FLD}}$ as the projection direction in the feature space

$$\boldsymbol{y}_{\mathrm{FLD}} = (\boldsymbol{u}_2 - \boldsymbol{u}_1) S_w^{-1}. \tag{20}$$

Here, $\boldsymbol{u}_1$ and $\boldsymbol{u}_2$ are the mean vectors of all $\boldsymbol{x}_k$ in C_1 and C_2, respectively. S_w is the within-class scatter matrix [note: do not confuse S_w with the scattering matrix in (9)]

$$S_w = \frac{1}{N} \times \left[\sum_{x_k \in C_1} (\boldsymbol{x}_k - \boldsymbol{u}_1)(\boldsymbol{x}_k - \boldsymbol{u}_1)^{\mathrm{T}} + \sum_{x_k \in C_2} (\boldsymbol{x}_k - \boldsymbol{u}_2)(\boldsymbol{x}_k - \boldsymbol{u}_2)^{\mathrm{T}} \right]. \tag{21}$$

The confidence value of an alarm is given by the projection of its feature vector on $\boldsymbol{y}_{\mathrm{FLD}}$. We use C_1 and C_2 to represent "FA" and "TGT" alarms, respectively, so that the mean score of hits is larger than that of FAs (see, for example, [18] for more detailed description of FLD).

In order to evaluate the performance of a classifier, confidence values of all alarms are thresholded at different levels to produce ROC curves. Although FLD provides a class separation measure, we are more interested in the ROC curves for judging the merits of various feature combinations for land-mine detection. We generate ROC curves using the following rules, with r_{halo} currently set to 1.0 m.

1) An alarm is considered a detection if it is within r_{halo} from the center of a known target.
2) When multiple alarms are within r_{halo} from the center of the same known target, only the one with the highest confidence value is considered a detection, and the other alarms are ignored.
3) An alarm is considered a false alarm if it is not within r_{halo} from the center of any known target and it is within the lane boundary.
4) An alarm is considered as neither a detection nor a false alarm if it is not within r_{halo} from the center of any known target and it is outside of the lane boundary.
5) When we generate the ROC curve for a specific subset of targets (e.g., buried metal mines), alarms within r_{halo} from the center of other known targets are not considered as either detections or false alarms.
6) The total area used in FAR calculation is the total lane area minus $N_{\mathrm{tgt}}(\pi r_{\mathrm{halo}}^2)$, where N_{tgt} is the total number of known targets, regardless of whether we are generating the ROC curve for a specific subset of targets.

Please note that the rules stated above apply to our computation of ROC curves for the calibration data only. The ROC curves for blind test data are generated by the independent agency IDA.

Specifically, to compare different classifiers, we use $\mathrm{PD}_{\mathrm{mean}}$, which is defined as the average PD for FAR between 0 and 0.1 m^{-2}, as the performance measure of a classifier. We limit the consideration to FAR ≤ 0.1 m^{-2} because, at higher FAR, there might be substantial number of "chance" detections.

We employ the sequential-forward-search (SFS) technique to search for the best feature combinations (see, for example, [18]). In this approach, we start the feature collection with the best individual feature. We then evaluate all feature combinations involving the current feature collection and one additional feature. This is repeated until we have the desired number of features in the collection.

III. Experimental Results

Our data set includes three calibration and four blind runs, each over a different mine lane at the same U.S. Army facility. All the mine lanes are 3 m wide. Their lengths are 300–340 m for the three calibration lanes and 1200 m each for the four blind lanes. Targets in the calibration lanes include a total of 15 surface-laid metal or plastic mines, 12 buried metal mines at 5–15 cm of depth, 20 buried plastic mines at 5–10 cm of depth, and 16 SIMs. Since we are most interested in the performance for actual mines, SIMs are not included in the results below. We have the ground truth for the calibration lanes, but have no information as to the number and placement of land mines in the blind lanes.

Before we examine the various feature combinations, we first look at the classification performance of the ten individual features. The $\mathrm{PD}_{\mathrm{mean}}$ values obtained using a single-feature FLD are listed in Table I. For the "resubstitution" method, we include all the alarms of the three calibration runs in both training and testing sets. For the "cross-validation" method, we repeat the procedure three times, each time calculating the confidence values of the alarms in one run (the validation set), with the alarms in the other two runs being the training set. In addition, we also show the results with and without the data-duplication procedure discussed in the previous section. We can see that, with the exception of feature #10 (P_ψ), the results for resubstitution and cross-validation are very close when we apply data duplication. We also find that data duplication does not produce much difference with resubstitution, but its effect

TABLE I
PD_{mean} VALUES OBTAINED USING INDIVIDUAL FEATURES

Feature / Method	#1	#2	#3	#4	#5	#6	#7	#8	#9	#10
(A)	0.75	0.82	0.71	0.69	0.93	0.77	0.66	0.19	0.76	0.30
(B)	0.75	0.80	0.68	0.69	0.90	0.77	0.65	0.14	0.76	0.11
(C)	0.75	0.82	0.71	0.69	0.93	0.77	0.66	0.19	0.76	0.21
(D)	0.73	0.81	0.65	0.60	0.77	0.59	0.66	0.12	0.65	0.13

Methods: (A) Resubstitution with data duplication; (B) Resubstitution without data duplication; (C) Cross-validation with data duplication; (D) Cross-validation without data duplication.

Features: #1:M_{VV}; #2:C_{VV}; #3:$R_{VV/HV}$; #4:M_{HH}; #5:C_{HH}; #6:$R_{HH/HV}$; #7:P_H; #8:P_α; #9:P_τ; #10:P_ψ

TABLE II
FEATURE COMBINATIONS SELECTED USING SFS

#Features / Method	1	2	3	4	5	6	7	8	9	10
(A)	(5) 0.93	(8) 0.93	(1) 0.94	(4) 0.94	(3) 0.94	(2) 0.94	(6) 0.94	(10) 0.94	(7) 0.93	(9) 0.85
(B)	(5) 0.90	(4) 0.91	(8) 0.90	(10) 0.90	(7) 0.89	(2) 0.88	(3) 0.87	(9) 0.86	(1) 0.84	(6) 0.82
(C)	(5) 0.93	(10) 0.93	(8) 0.92	(2) 0.89	(4) 0.90	(1) 0.89	(7) 0.89	(6) 0.88	(9) 0.88	(3) 0.87
(D)	(2) 0.81	(5) 0.86	(4) 0.89	(9) 0.88	(6) 0.88	(10) 0.87	(1) 0.87	(3) 0.85	(7) 0.84	(8) 0.82

Methods: (A) Resubstitution with data duplication; (B) Resubstitution without data duplication; (C) Cross-validation with data duplication; (D) Cross-validation without data duplication.

In each cell, the first number (in parenthesis) indicates the feature added in the given iteration of sequential forward search, and the second number indicates the corresponding PD_{mean} value.

is much more significant with cross validation. C_{HH} is the one feature giving the best performance, except for cross validation without data duplication, in which case it is second to C_{VV}. In either case, it appears that local contrast is the best single feature.

Next, we apply SFS to our ten features in order to identify the best feature combination. Table II lists the one- to ten-feature combinations selected using this approach, as well as the corresponding PD_{mean} values. The results using resubstitution or cross-validation and with or without data duplication are listed separately. We also see that the difference between PD_{mean} values obtained with resubstitution and cross validation generally increase with the number of features used. Furthermore, except for cross validation without data duplication, additional features, compared with using only feature #5 (local contrast in HH), hardly improve the detection rate. We also find that the effect of data duplication is more significant with resubstitution than with cross validation. It appears that data duplication gives somewhat better results. In addition, the results indicate that the best feature accounts for most of the classification performance, and hence, there is little need to use classifiers with a large number of features. However, this finding might be specific to this data set, as the ability of various features to differentiate mines and clutter likely depends on the type of mines and clutter objects present, soil condition, and radar configuration, etc.

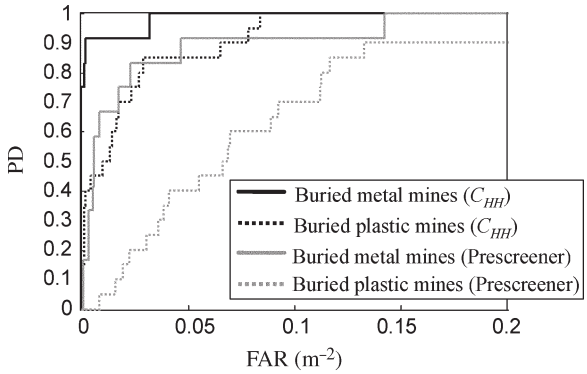

Fig. 13. Comparison of ROC curves for both the prescreener stage and for the best feature selected, C_{HH}.

TABLE III
FEATURE COMBINATIONS SELECTED USING SFS FROM EACH INDIVIDUAL FREQUENCY SUBBANDS

Frequency Subband	F_1	F_2	F_3	F_4
3-Feature Combinations Selected	5,8,1	2,7,10	5,1,6	1,2,3
PD_{mean}	0.94	0.72	0.74	0.64

TABLE IV
FEATURE COMBINATIONS SELECTED USING SFS FROM ALL FOUR FREQUENCY SUBBANDS

# Features	1	2	3	4	5
Feature Added	5	8	29	31	18
PD_{mean}	0.93	0.93	0.94	0.94	0.94

Fig. 13 displays the calibration data ROC curves generated using our own scoring algorithms for both the prescreener stage (light lines) and for the best feature selected C_{HH} (dark lines). The improvement for both buried metal and plastic mines is very significant.

While SIMs are not included in the scoring so far, it is also interesting to look at the detection performance when only SIMs are considered. Using only the feature C_{HH}, we achieve a detection rate of 43% (7 of 16) at FAR ≈ 0.02 m^{-2} and 56% (9 of 16) at FAR ≈ 0.07 m^{-2}. This is much worse than what we obtained for buried plastic mines with the same feature (65% at FAR ≈ 0.02 m^{-2} and 90% at FAR ≈ 0.07 m^{-2}). Adding additional features does not give better results for SIMs. This is consistent with our discussion above regarding why SIMs are not suitable for FLGPR algorithm development.

So far, we have focused our attention on features computed from the lowest frequency-subband data based on the prescreener results. It is also interesting to see whether similar features extracted from other frequency subbands can improve detection performance. Two experiments are conducted for this purpose. First, we used SFS to select the best three-feature combination separately for each frequency subband, with resubstitution and data duplication. The results are listed in Table III. It is clear that the lowest frequency subband gives much better results than the other three. Second, SFS is used to select the best combination from all 40 features (ten for each of the four subbands). Table IV lists the results up to five feature

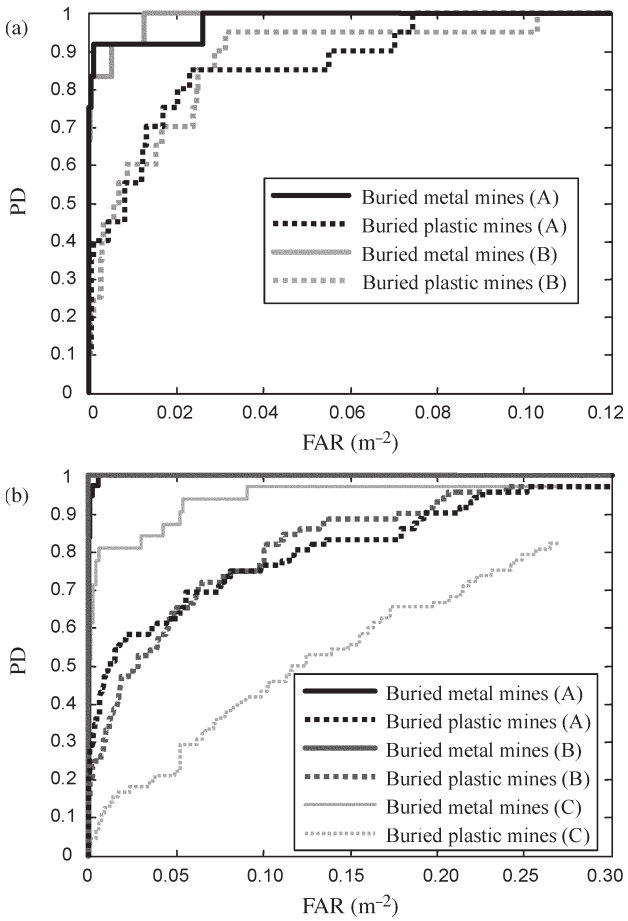

Fig. 14. ROC curves of our algorithm for (a) calibration data and (b) blind data. Both were produced by IDA from lists of alarms generated by our algorithm. Classifiers (A) and (B) are our algorithm using one and three features, respectively. Classifier (C) is the algorithm of a different research group for comparison purpose.

combinations using resubstitution with data duplication. Here, features numbered 1–10, 11–20, 21–30, and 31–40 correspond to the ten features for F_1, F_2, F_3, and F_4, respectively. While features from subbands other than F_1 do get included starting with the third feature selected by SFS, they still do not add much to the performance.

In Fig. 14, we show the ROC curves obtained from our algorithm, with separate curves for buried metal and plastic mines. Fig. 14(a) and (b) is for calibration and blind data, respectively. Both sets of ROC curves are generated and supplied to us by IDA, and for the blind data, the quantity and location of mines have remained unknown to the authors, although we are informed that these blind-data ROC curves include only mines buried at 7.5-cm depth. Two sets of results are generated from our algorithm: classifier A, which uses only a single feature $(C_{\rm HH})$, and classifier B, which uses a three-feature combination ($C_{\rm HH}$, $M_{\rm VV}$, and P_α) selected by SFS using resubstitution with data duplication. It appears that the performance of classifiers A and B is similar for both calibration and blind data. We can also see that we have near perfect results (100% PD at FAR < 0.01 m^{-2}) for buried metal mines in the blind data (for which the ROC curves for classifiers A and B mostly overlap). For the blind data, we also include the next best ROC curves, supplied to us by IDA, of results obtained with another existing algorithm (classifier C) for the same data set. A description of classifier C, which used features based on time–frequency analysis and employed sophisticated multistage classifier design, can be found in [19]. It is evident that our approach, which uses straightforward features and classification algorithms while focusing on one subband of the whole frequency range, gives significantly better performance, with 25%–30% higher detection probability for the same FAR for plastic mines. Other algorithms use a variety of complicated features and classification algorithms but do much worse. Therefore, we conclude that, at present, the use of frequency subbands has shown more value than complicated features and algorithms for forward-looking mine detection with radar and should be investigated in more details.

While the focus of this paper is not on the investigation of the underlying physical process involved in FLGPR, we would like to point out some qualitative physics-based interpretations of our observations. The finding that the lowest of the four frequency subbands yields the best performance is related to the following two points: 1) low-frequency radar signal penetrates the ground better and therefore is better suited for detecting buried objects and 2) a low-frequency radar signal is less prone to surface clutter (clutter signals from reflected radar waves by objects on or above the ground surface), resulting in the cleaner signal for lower frequency subbands, as evidenced in Fig. 8. The magnitude and local contrast features simply reflect the expectation that regions with buried mines should have stronger radar return signals. The features $R_{\rm VV/HV}$ and $R_{\rm HH/HV}$ originated from an observation made by researchers at SRI that cross-polarization signals from mines tend to disappear. In Fig. 10, this appears to be more valid for lower frequency bands. Additional modeling will be needed to better understand this phenomenon. As of the two polarimetric features that exhibit significant discriminating power individually, we find that mines tend to have lower P_H (less entropy or randomness) and higher P_τ (more like symmetric scatters) compared to false alarms. These are intuitively consistent with what we will usually expect from man-made objects. Again, more modeling work is necessary to understand these phenomena.

IV. Conclusion

In summary, we have described an approach for land-mine detection using forward-looking GPR data. We also presented the analysis of a number of features, including those based on polarimetric data, as well as the procedure for selecting subsets of features. Weighted averages were used to aggregate features of 2-D image patches to features of alarms, which usually appear in multiple scan images as the vehicle moves forward. The algorithm was applied to both calibration and blind data with good performance, considerably better than other algorithms that were tested on the same data in the same time frame. A major contribution is the finding that, by dividing up the wide-frequency-band radar signals into narrower frequency subbands, we are able to find a frequency subband with good target-to-clutter contrast. By focusing on this frequency subband, we obtained significantly improved detection

performance over existing results for the same data set. Good results are obtained with just one of the features, a CFAR-like feature from the HH polarization mode at the lower frequency subband. It turns out that the addition of other features gives little improvement over the single-feature classifier, as evidenced in ROC curves for both calibration and blind data.

The fact that we gain little performance improvement by adding additional features to the best one (C_{HH}) seems to argue for the use of only this one feature. However, we believe that such a conclusion would be premature because, so far, the analysis uses only a single data set at a single site. The ability of various features to differentiate mines and clutter likely depends on the type of mines and clutter objects present, soil condition, radar configuration, etc. All these can change between data collections. In addition, we can see from Table I that the other features (with the exception of P_α and P_ψ), when used individually, also possess significant discriminating power. It is quite possible that, at a new data collection, the best feature combination is different from those presented in this paper. We believe it is desirable to have all these features available for the analysis of future data sets. Our main conclusion, therefore, is not in the selection of one particular feature, but rather in the fact that good performance can be obtained by combining the use of frequency subbands and a set of simple features.

Forward-looking mine detection using GPR is a relatively new technology with lots of questions needing investigation in the future. For example, much knowledge can be gained by more extensive study, through modeling or experiments, of how frequency response depends on parameters including radar antenna configuration, standoff distance, properties of soil, and target and clutter characteristics. We also plan to investigate how we can leverage the strengths of both existing "time-domain" approaches using detailed spatial features and our "quasi-frequency-domain" approach for the creation of more robust systems. Overall, we believe that the approach described in this paper, especially the analysis of frequency subbands, shows very good promise for continuously advancing research in this area.

ACKNOWLEDGMENT

The authors would like to thank K. Sherbondy, I. McMichael, and E. Rosen for their support.

REFERENCES

[1] L. M. Collins, P. A. Torrione, C. S. Throckmorton, X. Liao, Q. E. Zhu, Q. Liu, L. Carin, F. Clodfelter, and S. Frasier, "Algorithms for landmine discrimination using the NIITEK ground penetrating radar," *Proc. SPIE*, vol. 4742, pp. 709–718, 2002.

[2] P. D. Gader, R. Grandhi, W.-H. Lee, J. N. Wilson, and D. K. Ho, "Feature analysis for the NIITEK ground penetrating radar using order weighted averaging operators for landmine detection," *Proc. SPIE*, vol. 5415, pp. 953–962, 2004.

[3] F. Cremer, J. G. Schavemaker, W. de Jong, and K. Schutte, "Comparison of vehicle-mounted forward-looking polarimetric infrared and downward-looking infrared sensors for landmine detection," *Proc. SPIE*, vol. 5089, pp. 517–526, 2003.

[4] N. Playle, D. M. Port, R. Rutherford, I. A. Burch, and R. Almond, "Infrared polarization sensor for forward-looking mine detection," *Proc. SPIE*, vol. 4742, pp. 11–18, 2002.

[5] R. D. Costley, J. M. Sabatier, and N. Xiang, "Forward-looking acoustic mine detection system," *Proc. SPIE*, vol. 4394, pp. 617–626, 2001.

[6] M. R. Bradley, T. R. Witten, M. Duncan, and R. McCummins, "Anti-tank and side-attack mine detection with a forward-looking GPR," *Proc. SPIE*, vol. 5415, pp. 421–432, 2004.

[7] J. Kositsky, R. Cosgrove, C. A. Amazeen, and P. Milanfar, "Results from a forward-looking GPR mine detection system," *Proc. SPIE*, vol. 4742, pp. 206–217, 2002.

[8] R. B. Cosgrove, P. Milanfar, and J. Kositsky, "Trained detection of buried mines in SAR images via the deflection-optimal criterion," *IEEE Trans. Geosci. Remote Sens.*, vol. 42, no. 11, pp. 2569–2575, Nov. 2004.

[9] M. R. Bradley, T. R. Witten, M. Ducan, and R. McCummine, "Mine detection with a forward-looking ground-penetrating synthetic aperture radar," *Proc. SPIE*, vol. 5089, pp. 334–347, 2003.

[10] G. Liu, Y. Wang, J. Li, and M. R. Bradley, "SAR imaging for a forward-looking GPR system," *Proc. SPIE*, vol. 5089, pp. 322–333, 2003.

[11] Y. Sun and J. Li, "Plastic landmine detection using time-frequency analysis for forward-looking ground-penetrating radar," *Proc. SPIE*, vol. 5089, pp. 851–862, 2003.

[12] T. Wang, O. Sjahputera, J. M. Keller, and P. D. Gader, "Landmine detection using forward-looking GPR with object-tracking," *Proc. SPIE*, vol. 5794, pp. 1080–1088, 2005.

[13] ——, "Feature analysis for forward-looking landmine detection using GPR," *Proc. SPIE*, vol. 5794, pp. 1233–1244, 2005.

[14] K. C. Ho, P. D. Gader, and J. N. Wilson, "Improving landmine detection using frequency domain features from ground penetrating radar," in *Proc. IEEE Int. Geosci. and Remote Sens. Symp.*, Anchorage, AK, Sep. 2004, vol. 3, pp. 1617–1620.

[15] S. R. Cloude and E. Pottier, "An entropy based classification scheme for land applications of polarimetric SAR," *IEEE Trans. Geosci. Remote Sens.*, vol. 35, no. 1, pp. 68–78, Jan. 1997.

[16] W. L. Cameron and L. K. Leung, "Feature motivated polarization scattering matrix decomposition," in *Proc. IEEE Int. Radar Conf.*, 1990, pp. 549–557.

[17] W. L. Cameron, N. N. Youssef, and L. K. Leung, "Simulated polarimetric signatures of primitive geometrical shapes," *IEEE Trans. Geosci. Remote Sens.*, vol. 34, no. 3, pp. 793–803, May 1996.

[18] S. Theodoridis and K. Koutroumbas, *Pattern Recognition*, 2nd ed. San Diego, CA: Academic, 2003.

[19] Y. Sun and J. Li, "Landmine detection using forward-looking ground penetrating radar," *Proc. SPIE*, vol. 5794, pp. 1089–1097, 2005.

[20] Joint Research Centre, European Commission. (2001, Aug. 7). *MIMEVA: Study of Generic Mine-Like Objects for R&D in Systems for Humanitarian Demining*. [Online]. Available: http://humanitarian-security.jrc.it/demining/final_reports/mimeva/part_4.pdf

Tsaipei Wang (M'05) received the B.S. degree in physics from the National Tsing Hua University, Hsinchu, Taiwan, R.O.C., in 1989, the Ph.D. degree in physics from the University of Oregon, Eugene, Oregon, in 1999, and the M.S. degree in computer science and the Ph.D. degree in computer engineering and computer science from the University of Missouri–Columbia, Columbia, in 2002 and 2005, respectively.

He is currently an Assistant Professor with the Department of Computer Science, National Chiao Tung University, Hsinchu. He was a Postdoctoral Fellow with the University of Missouri–Columbia. He was previously involved in research topics including frequency-selective optical memory, coherent transient phenomena, signal/image processing for land-mine/tripwire detection, and automated vision development assessment with video photoscreening. His current research interests include fuzzy systems, image and video processing, medical image analysis, and visualization.

James M. Keller (M'79–SM'92–F'00) received the Ph.D. in mathematics from the University of Missouri-Columbia, Columbia, in 1978.

He holds the University of Missouri Curators' Professorship in the Electrical and Computer Engineering and Computer Science Departments on the Columbia campus. He is also the R. L. Tatum Professor in the College of Engineering. His research interests center on computational intelligence: fuzzy set theory and fuzzy logic, neural networks, and evolutionary computation, with a focus on problems in computer vision, pattern recognition, and information fusion, including bioinformatics, spatial reasoning in robotics, sensor and information analysis in technology for eldercare, and land-mine detection. He has been supported by several industrial and government institutions, including the Electronics and Space Corporation, Union Electric, Geo-Centers, NSF, the Administration on Aging, NASA/JSC, the Air Force Office of Scientific Research, the Army Research Office, the Office of Naval Research, and the Army Night Vision and Electronic Sensors Directorate. He has coauthored over 250 technical publications. He is an Associate Editor of the *International Journal of Approximate Reasoning*, and is on the editorial board of *Pattern Analysis and Applications*, *Fuzzy Sets and Systems*, *International Journal of Fuzzy Systems*, and the *Journal of Intelligent and Fuzzy Systems*.

Prof. Keller has presented live and video tutorials on fuzzy logic in computer vision with the IEEE, is a national lecturer for the Association for Computing Machinery (ACM), is an IEEE Computational Intelligence Society Distinguished Lecturer, and was the past President of the North American Fuzzy Information Processing Society (NAFIPS). He finished a full six-year term as an Editor-in-Chief of the IEEE TRANSACTIONS ON FUZZY SYSTEMS. He is currently the Vice President for Publications of the IEEE Computational Intelligence Society. He was the Conference Chair of the 1991 NAFIPS Workshop, program cochair of the 1996 NAFIPS meeting, Program Cochair of the 1997 IEEE International Conference on Neural Networks, and the Program Chair of the 1998 IEEE International Conference on Fuzzy Systems. He was the General Chair for the 2003 IEEE International Conference on Fuzzy Systems.

Paul D. Gader (M'87–SM'99) received the Ph.D. degree in mathematics from the University of Florida, Gainesville, in 1986.

He has worked as a Senior Research Scientist with Honeywell's Systems and Research Center, as a Research Engineer and Manager with the Environmental Research Institute of Michigan, and as a Faculty Member with the University of Wisconsin, the University of Missouri, and the University of Florida, where he is currently a Professor of computer and information science and engineering. He has led teams involved in real-time handwritten address recognition systems for the U.S. Postal Service, developing algorithms for handwritten digit recognition and segmentation, numeric field recognition, word recognition, and line segmentation. He has led teams that devised and tested several real-time algorithms in the field of mine detection. He served as the Technical Director of the University of Missouri MURI on Humanitarian Demining for two years. He is currently involved in land-mine-detection projects investigating hand-held and ground-based detection systems, acoustic detection, EO/IR detection of mines and trip-wires, hyperspectral detection, and multisensor fusion. He has over 165 technical publications in the areas of image and signal processing, applied mathematics, and pattern recognition, including over 55 refereed journal articles.

Ozy Sjahputera (M'05) received the B.S. degree in computer and electrical engineering, the M.S. degree in electrical engineering, and the Ph.D. degree in computer engineering and computer science from the University of Missouri–Columbia, Columbia, in 1994, 1996, and 2004, respectively.

He is currently a Postdoctoral Fellow in bioinformatics with the Ellis Fischel Cancer Research Laboratory in Columbia, MO. His research experiences include scene analysis with spatial reasoning, supported by the U.S. Office of Naval Research, and automatic land-mine detection using forward looking GPR, supported by the U.S. Army Night Vision Laboratory. His current research interests include scene analysis and pattern recognition in geospatial intelligence and bioinformatics in cancer research.

Dr. Sjahputera is a member of the IEEE Computational Intelligence Society.