REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
22-02-2011	Reprint	Jan 2005 - Sept 2010

4. TITLE AND SUBTITLE

Improved detection and false alarm rejection using FLGPR and color imagery in a forward-looking system

5a. CONTRACT NUMBER

W911NF-05-1-0069

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)

Havens, Tim, C.
Spain, Christopher, J.
Ho, Dominic, K.C.
Keller, James, M.
Ton, Tuan, T.
Wong, David, C.
Soumekh, Mehrdad

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

University of Missouri
Office of Research
205 Jesse Hall
Columbia, MO 65211

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Department of the Army
US Army Research, Development and Engineering Command Acquisition Center
Research Triangle Park Contracting Division
P.O. Box 12211
Research Triangle Park, NC 27709-2211

10. SPONSOR/MONITOR'S ACRONYM(S)

ARO Terrestrial Sciences

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release; federal purpose rights.

13. SUPPLEMENTARY NOTES

14. ABSTRACT

Forward-looking ground-penetrating radar (FLGPR) has received a significant amount of attention for use in explosive-hazards detection. A drawback to FLGPR is that it results in an excessive number of false detections. This paper presents our analysis of the explosive-hazards detection system tested by the U.S. Army Night Vision and Electronic Sensors Directorate (NVESD). The NVESD system combines an FLGPR with a visible-spectrum color camera. We present a target detection algorithm that uses a locally-adaptive detection scheme with specrum-based features. The remaining FLGPR detections are then projected into the camera imagery and image-based features are collected. A one-class classifier is then used to reduce the number of false detections, We show that our proposed FLGPR target detection algorithm, coupled with our camera-based false alarm (FA) reduction method, is effective at reducing the number of FAs in test data collected at a US Army test facility.

15. SUBJECT TERMS

Sensor Fusion, Forward-Looking Explosive Hazards Detection, Ground-Penetrating Radar, False Alarm Rejection

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Jim Keller
UU	UU	UU	UU	12	19b. TELEPHONE NUMBER *(Include area code)* (573) 882-7339

Reset

Standard Form 298 (Rev. 8/98)
Prescribed by ANSI Std. Z39.18

Improved Detection and False Alarm Rejection Using FLGPR and Color Imagery in a Forward-Looking System

Timothy C. Havens*[a], Christopher J. Spain[a], K.C. Ho[a], James M. Keller[a],
Tuan T. Ton[b], David C. Wong[b], and Mehrdad Soumekh[c]
[a]Dept. of Electrical and Computer Engineering, University of Missouri, Columbia, MO, USA 65211;
[b]U.S. Army REDCOM CERDEC NVESD, Fort Belvoir, Virginia, USA 22060;
[c]Dept. Of Electrical Engineering, University of New York at Buffalo, Amherst, NY, USA 14260

ABSTRACT

Forward-looking ground-penetrating radar (FLGPR) has received a significant amount of attention for use in explosive-hazards detection. A drawback to FLGPR is that it results in an excessive number of false detections. This paper presents our analysis of the explosive-hazards detection system tested by the U.S. Army Night Vision and Electronic Sensors Directorate (NVESD). The NVESD system combines an FLGPR with a visible-spectrum color camera. We present a target detection algorithm that uses a locally-adaptive detection scheme with spectrum-based features. The remaining FLGPR detections are then projected into the camera imagery and image-based features are collected. A one-class classifier is then used to reduce the number of false detections. We show that our proposed FLGPR target detection algorithm, coupled with our camera-based *false alarm* (FA) reduction method, is effective at reducing the number of FAs in test data collected at a US Army test facility.

Keywords: Sensor fusion, forward-looking explosive hazards detection, ground-penetrating radar, false alarm rejection

1. INTRODUCTION

Remediation of the threat of explosive hazards is an extremely important goal, as these hazards are responsible for uncountable deaths and injuries to both civilians and soldiers throughout the world. Systems that detect explosive hazards have included *ground-penetrating-radar* (GPR), *infrared* (IR) cameras, and acoustic technologies.[1-3] Both handheld and vehicle-mounted GPR-based systems have been examined in recent research and much progress has been made in increasing detection capabilities.[4,5] *Forward-looking* synthetic aperture GPR (FLGPR) is an especially attractive technology because of its ability to detect hazards before they are encountered; standoff distance can range from a few to tens of meters. FLGPR has been applied to the detection of side-attack mines[6], and mines in general.[7,8] A drawback to these systems is that FLGPR is not only sensitive to objects of interest, but also to other objects, both above and below the ground. This results in an excessive number of false detections.

The FLGPR images we present in this paper were collected by a system called ALARIC. This system is an FLGPR system that is composed of a physical array of sixteen receivers and two transmitters. In the past decade, FLGPR systems have primarily used their physical arrays (aperture) as well as their radar bandwidth for imaging (resolution); conventional backprojection or time domain correlation imaging has been used for this purpose. Those FLGPR systems rarely tried to exploit imaging information that is created by the motion of the platform. The ground-based FLGPR community has referred to imaging methods that leverage platform motion as multi-look imaging though in the airborne radar community this is better known as *synthetic aperture radar* (SAR) imaging. SAR has been shown to be an effective tool for airborne *intelligence, surveillance* and *reconnaissance* (ISR) applications.

The ALARIC system is equipped with an accurate GPS system. As a result, we are capable of processing both physical and synthetic aperture imaging even when the platform moves along a nonlinear path with variations in its heading. To create the FLGPR images we use a nonlinear processing technique called *Adaptive Multi-Transceiver Imaging*. This method exploits a measure of similarity among the 32 T/R images which adaptively suppresses artifacts such as sidelobes and aliasing ghots.

Detection and Sensing of Mines, Explosive Objects, and Obscured Targets XV,
edited by Russell S. Harmon, John H. Holloway Jr., J. Thomas Broach, Proc. of SPIE
Vol. 7664, 76641U · © 2010 SPIE · CCC code: 0277-786X/10/$18 · doi: 10.1117/12.852274

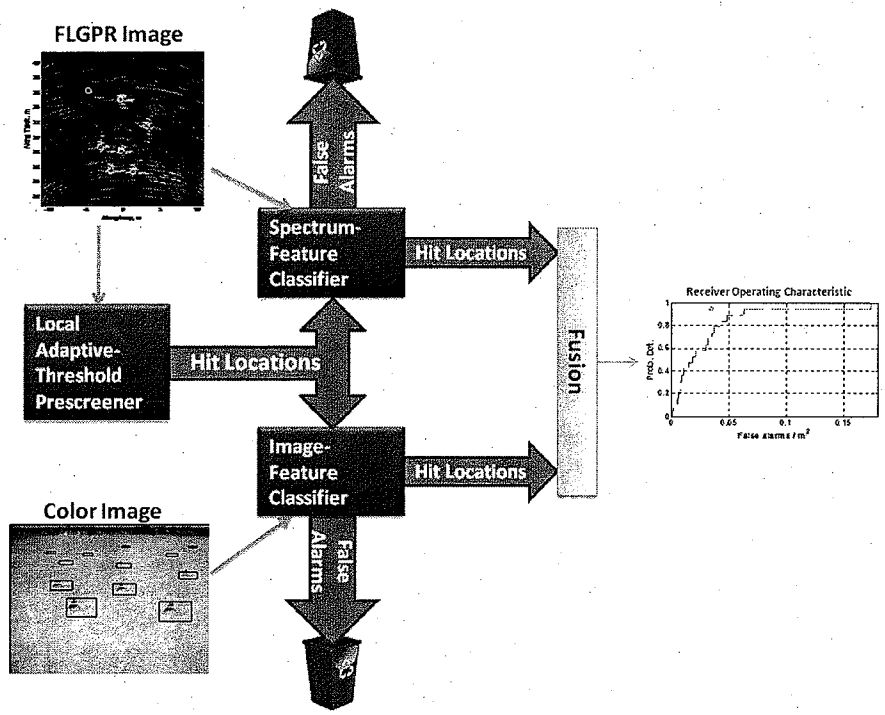

Fig. 1. Block diagram of forward-looking explosive hazards detection algorithms.

This paper presents a sensor-fusion algorithm that detects explosive hazards in FLGPR scans and uses associated color imagery to reduce the number of FAs. Figure 1 illustrates our system in block-diagram. A locally-adaptive threshold detector is used on the FLGPR image to produce alarm locations. These alarms are passed through two classifiers that reject FAs: one classifier uses a FLGPR spectrum-based feature and one uses camera-based features. The outputs of these classifiers are fused into candidate target locations. Section 1.1 briefly describes our previous work on the co-registration of the imagery[9] and FLGPR. Section 1.2 describes our locally-adaptive threshold detection algorithm[10,11]. Our FLGPR spectrum-based feature and the classifier are presented in Section 1.3. In Section 2 we describe the method by which we combine the FLGPR data with the color imagery in order to further reduce the number of false detections. Section 3 presents test results of both classifiers independently and fused. The *receiver-operating-characteristic* (ROC) curves both from training and test data show that this is an effective method for reducing FAs. Section 4 concludes this paper.

1.1 Camera registration

To effectively use the color imagery for screening FLGPR alarm locations, an accurate transformation between two-dimensional world coordinates, as reported by the FLGPR, and camera image coordinates is needed. No information about the camera, specifically the internal camera parameters, or the camera's pose or location on the vehicle is assumed. The only available information is the ground-truth locations of several calibration targets that are visible in the color images, and the location of the vehicle when each color image was taken. Below is a brief description of the camera registration method; reference [9] has a more detailed description of this method.

A generalized perspective projection model based on an ideal pinhole camera is used to represent the transformation from camera reference frame coordinates (X_c, Y_c, Z_c) to two-dimensional image coordinates (X_i, Y_i). In this model, each point in the image corresponds to the intersection of a line with the image plane, running from a point in the camera

reference frame through the center of projection. In homogeneous coordinates, the projection onto the image plane can be represented as $\begin{bmatrix} X_i \\ Y_i \\ 1 \end{bmatrix} = P \begin{bmatrix} X_c \\ Y_c \\ Z_c \\ 1 \end{bmatrix}$, where P is the 4x4 projection matrix

The projection matrix P allows transformation from camera reference frame coordinates to a pixel position in the image plane. The full projective model cannot be used in our particular case because the Z-coordinate of the calibration objects is unknown. Hence, we assume a flat earth, where $Z_c = 0$. This assumption reduces the projection model to

$$\begin{bmatrix} X_i \\ Y_i \\ 1 \end{bmatrix} = P \begin{bmatrix} X_c \\ Y_c \\ 1 \end{bmatrix},$$

where P is now a 3x3 matrix. Additionally, we can project pixel coordinates in the image plane to the camera reference plane with the inverse transformation

$$\begin{bmatrix} X_c \\ Y_c \\ 1 \end{bmatrix} = P^{-1} \begin{bmatrix} X_i \\ Y_i \\ 1 \end{bmatrix}.$$

This model assumes that the two-dimensional coordinates (X_c, Y_c) are in the camera reference frame, i.e. situated relative to the position and heading of the camera. However, the FLGPR alarm locations are reported as two-dimensional world coordinates. These world coordinates must first be transformed into the camera reference frame before projection into the image plane can occur. This transformation is possible since the heading and location of the vehicle when each image was taken is known. The matrix equation

$$\begin{bmatrix} X_c \\ Y_c \end{bmatrix} = \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix} \begin{bmatrix} X_w - X_v \\ Y_w - Y_v \end{bmatrix}$$

can be used to transform two-dimensional world coordinates (X_w, Y_w) into camera reference frame coordinates. The point (X_v, Y_v) is the location of the vehicle, and θ is the vehicle heading.

Although the camera is located on the vehicle, the reported heading and location of the vehicle is not the exact heading and location of the camera itself; the camera is pointed towards the side of the road. However, the camera is fixed on the vehicle; hence, the transformation between the vehicle heading and location and the camera heading and location is static. This transformation can be modeled by a static 3x3 transformation matrix R. The final transformation model is

$$\begin{bmatrix} X_i \\ Y_i \\ 1 \end{bmatrix} = PR \begin{bmatrix} X_c \\ Y_c \\ 1 \end{bmatrix}.$$

The elements of both P and R are unknown and can simply be combined into a single 3x3 matrix. We refer to this projection matrix, aptly, as PR.

The parameters PR are approximated by using an evolutionary optimization algorithm called CMA-ES.[12-14] The training data are composed of images of several targets whose ground-truth locations, as well as the vehicle headings and locations, are known. The CMA-ES algorithm chooses candidate projection matrices and then computes error by computing the difference between the projected ground-truth coordinates and the pixel coordinates of the training targets. The candidate partitions are then evolved by the CMA-ES algorithm at each iteration until an acceptable

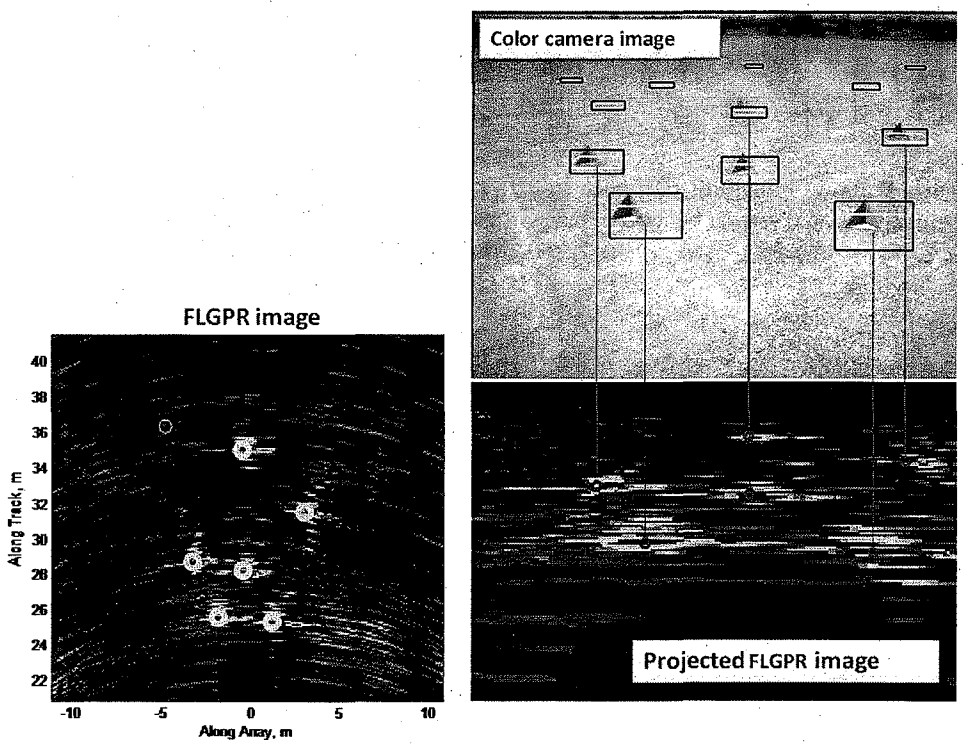

Fig. 2. Example of FLGPR and camera co-registration. Left image is an FLGPR scan with target ground-truth locations shown by yellow circles. Lower-right image shows FLGPR scan projected into the camera reference frame (shown in the upper-right). Red lines indicate FLGPR-camera correspondence.

solution is found. Figure 2 illustrates the results of approximating *PR* by showing the correspondence between the FLGPR image and the IR image for several targets. In the absence of ground-truth information for the target locations, this method could also be used to approximate *PR* by gleaning the target locations directly from the FLGPR.

1.1 FLGPR locally-adaptive threshold detector

The FLGPR images are created for the area -11m to 11m in the cross-range direction (although, in practice, only a sub-region of this is used in our detection algorithms), where negative numbers indicate to the left of the vehicle. Coherent integration of radar scans is performed in an area 9m to 25m in front of the vehicle. The pixel-resolution of the FLGPR image is 0.05m x 0.05m. The nominal center frequency is 1.2GHz and the bandwidth is 1.5GHz. We chose a detection region 9m wide; if the targets are on the left side of the road, relative to the vehicle, this region is positioned from -7m to +2m, if the targets are on the right side of the road this region is positioned from -2m to +7m.

References [15] describes our previous efforts detecting land mines in FLGPR data. The algorithm we present in this paper is an adaptation from our work in [15] and the locally-adaptive threshold detector is described in detail in [11].

Consider a FLGPR image $G(u,v)$, where u is the cross-range coordinate and v is the down-range coordinate. We first filter G with a locally-adaptive standard deviation filter. This computes the local standard deviation in a variable-size rectangular halo around each pixel. Figure 3 shows the region in which the local standard deviation is calculated. The standard-deviation filtered image is calculated by

$$F(x,y) = I(x,y) / \sigma_{local}(x,y).$$

We then use the MUFL prescreener, described in references [10,15], on this standard deviation filtered image. Essentially, the threshold calculated by the prescreener is adaptive to the local standard deviation. Figure 4 shows the

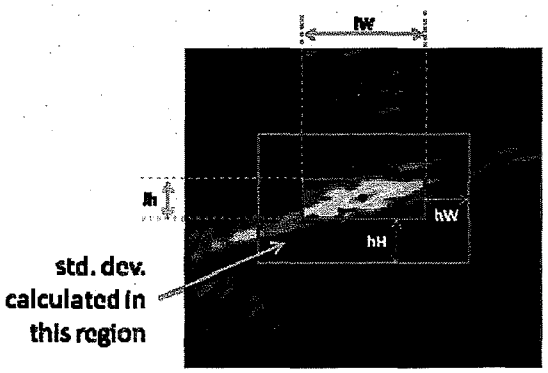

Fig. 3. Local adaptive-threshold prescreener calculates standard deviation in rectangular halo around each radar image pixel.

ROC curves of our MUFL prescreener both with and without the adaptive threshold. Note that these data are for the MUFL prescreener *only*; these do not include the side-lobe rejection, linking, and video-based false-alarm rejection methods previously presented in [10]. This plot shows that the locally-adaptive threshold prescreener, using a 5x5 / 5x20 sized window, results in >60% reduction in FAs compare to the non-adaptive prescreener. For the rest of the results in this paper, we will use a 5x5 / 5x20 sized window in the locally-adaptive prescreener algorithm.

1.3 Spectrum-feature classifier

In reference [11], we also describe a FA rejection method that is based on characterizing the spectrum of the clutter. The spatial spectrum of FAs in training data is computed and the spectrum elements are used as training features for a one-class classifier. Figure 6(a) in Section 3 shows the training results of using the spectrum-based classifier on the alarm locations following the locally-adaptive threshold prescreener. The training data is Test Run A. These results show that the classifier is able to reduce the FA rate from 0.06 FA / m^2 to 0.04 FA / m^2 – a 33% reduction. We note, however, that these are resubstitution results and represent the best performance that would be expected from this classifier. Later we will show results of the spectrum-based classifier on test data.

2. CAMERA-BASED FALSE ALARM REJECTION

Section 1.1 briefly described our method for projecting the FLGPR detections – computed by the method in Section 1.2 – into the corresponding camera images. With this method, we are able to find the areas in the camera images that correspond to each FLGPR detection. Hence, we can use the information in the IR images to classify the types of detections from the FLGPR, assuming that the image pixels corresponding to a false detection (e.g. bushes, rocks,

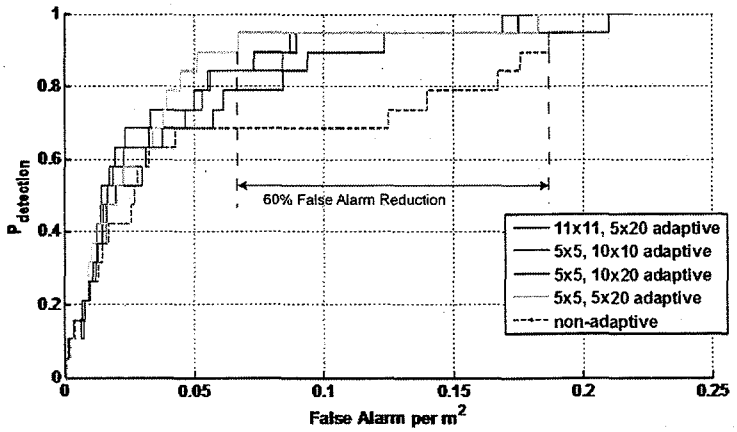

Fig. 4. ROC curve of MUFL prescreener for non-filtered radar image and three different sized locally-adaptive filter halos. The size of the rectangular halo is denoted as **iW**x**iH**, **hW**x**hH**, as shown in Fig. 3.

Fig. 5. Example of camera image taken by system.

garbage, etc.) are different from the pixels corresponding to an explosive hazard. The camera used on the NVESD system is a 1024x768 visual-spectrum color camera. The camera is aimed forward such as to image the same portion of the scene at which the FLGPR is radiating. Figure 5 shows an example of one of these images. For this paper, we focused on developing a robust and simple method for using the camera images to classify FLGPR detections as either true or false detections.

3.1 Color Feature Extraction

Each FLGPR detection can be projected into a camera pixel location using the method described in Section 1.1 (assuming that the detection is within the camera field-of-view). Generally, there are multiple frames, between 15 and 30, for each FLGPR detection. The distance to the detection location differs in each frame, and, therefore, the number of pixels that targets comprise in a corresponding camera image differs. We are interested in examining a fixed area, in meters, around each detection location; thus, an adaptive-sized window around each detection in the image is selected. The projection matrix PR allows us to compute the size of each image pixel, in meters, by using the inverse transformation from pixel positions to camera reference frame coordinates. Hence, it is possible to determine the appropriate window size to use for each image position, which corresponds to a chosen real world distance. We use a window size corresponding to a side length of one meter in the horizontal direction (cross-range) and two meters in the vertical direction (down-range), as we discovered that this is large enough to contain all targets present in our data. We denote these sub-images as W.

We calculate a set of features from the pixels in the windows corresponding to each FLGPR detection. First, the intensity, local standard deviation, Laplacian, and Sobel images are calculated. The Laplacian is calculated using the convolution kernel

$$\begin{bmatrix} 1/3 & 2/3 & 1/3 \\ 2/3 & -4/3 & 2/3 \\ 1/3 & 2/3 & 1/3 \end{bmatrix}.$$

The local standard deviation is calculated in a 5x5 window around each pixel. The Sobel image is calculated as

$$S = (W * S_x)^2 + (W * S_y)^2,$$

where * indicates convolution and the squares are calculated element-wise. We use the standard Sobel gradient operators, denoted as S_x and S_y.[17] We also create three other images, one each of the red, blue, and green channels of the image.

The set of features calculated on the target detections in each of the seven images (intensity, local standard deviation, Laplacian, Sobel, red, green, and blue) are the average, minimum, maximum, median, standard deviation, skewness, and kurtosis. For example, consider the red-channel image. The seven features corresponding to a sub-image W would be the average red pixel-value in W, the minimum red pixel-value in W, the maximum red pixel-value in W, etc. In total, 49 features are calculated from each window W, which is the sub-image where an FLGPR detection is visible. Recall that each detection location can appear in multiple images (usually 15-30); thus, each detection is represented by 15 to 30 sets of the 49 camera-based features. The median of these 15-30 sets of features is calculated so that each detection is represented, finally, by 49 aggregate feature values. We have experimented with other feature aggregation methods, including mean (both conventional and alpha-trimmed), min, and max, and we discovered that median was the most effective aggregation operator for combining the features from the multiple camera frames. In the future we hope to examine methods by which all sets of features can be used.

We then train a one-class classifer to reject FAs based on the 49 aggregate features.

2.2 One-class classifier

The 49 camera-based features and the FLGPR confidence value for each detection are used to classify the detection as either true (an explosive hazard) or false. We train a classifier by first calculating the multivariate normal distribution that best represents the feature values of the false detections for a given set of training data. Hence, the values of the false detections are assumed to be accurately represented by

$$f(x_1,\ldots,x_{49}) = \frac{1}{(2\pi)^{49/2}|\Sigma|^{0.5}} \exp\left[-0.5(x-\mu)^T \Sigma^{-1}(x-\mu)\right],$$

where μ is the mean vector and Σ is the covariance matrix. We fit the distribution parameters to the training data using the well-known maximum-likelihood estimator.[16] Once we have trained the classifier, we can use the Mahalanobis-metric to determine how well a new feature vector X fits the false detection distribution, where this distance is calculated by

$$D(X) = \sqrt{(X-\mu)^T \Sigma^{-1}(x-\mu)}\,.$$

If the Mahalanobis-metric $D(X)$ is large-valued, this indicates that the detection does not fit the false detection distribution and is, most likely, a true detection. Hence, a threshold T must be chosen such that a $D(X) > T$ indicates a true detection and a $D(X) \leq T$ indicates a false detection. The advantage of this method is that the threshold T can be tuned to offer an optimal tradeoff between true and false detections. Also, the distribution is trained on false detection data, of which there are many, rather than true detection data, of which there are few. Furthermore, the true detection features can be drastically different for different types and configurations of the explosive hazards, whereas the false detection features tend to more generalized. In practice, if one is using $D(X)$ to produce a threshold detector, then the square-root does not need to be included.

2.3 Feature and Threshold Selection

There are a total of 49 camera-based features for each FLGPR detection. It is unlikely that all of these features are necessary or effective for training an optimal classifier. Additionally, given a set of features we must choose the threshold T which determines whether an input feature vector is classified as a true or false detection. We use an *exhaustive search* to find the four best features. In [10], we used a forward sequential search to determine the best N features. However, we have since discovered that an exhaustive search can be performed relatively quickly and produces more generalized classification results. At each iteration of the exhaustive feature selection, the threshold T is set such that each target in the training data has at least one associated detection. In this manner, the optimal T eliminates the most false detections while maintaining a $P_D > 90\%$. Thus, the exhaustive search determines the four best features and associated classifier parameters, μ, Σ, and T. For comprehensive results on this classification scheme in regards to FLGPR and IR imagery, please refer to [11].

Figure 8(a) shows the training results of using the camera-based classifier on the alarm locations following the locally-adaptive threshold prescreener. The training data is Test Run A. These results show that the classifier is able to reduce the FA rate from 0.06 FA / m^2 to 0.05 FA / m^2 – a 16% reduction in FAs. Again, we note, that these are resubstitution results which result in the best performance that would be expected from this classifier. In the next section, we present test results of both the spectrum-feature and camera-feature classifiers, and the resulting fusion of these classifiers.

3. RESULTS

3.1 Spectrum-feature classifier test results

Figure 6 outlines the FA rejection results for the one-class classifier trained with the spectrum features. A confidence threshold was chosen from the training data that resulted in a >90% classification rate with the least number of FAs. This is shown as the blue dot in view (a) – this is the expected performance using just the locally-adaptive prescreener. As this figure shows, the expected FA rate at 95% probability of detection is 0.06 FA/m^2. The red dot in view (a) shows the FA rate after the spectrum-feature classifier is used. As this shows, the FA rate was reduced by 33% to 0.04 FA/m^2.

The same confidence threshold was then applied to Test Run B. View (b) shows that the locally-adaptive prescreener, with the threshold chosen from the training results in view (a), results in 90% probability of detection with 0.11 FA/m^2 (shown by the blue dot). If we apply the trained spectrum-feature classifier to Test Run B, we only achieve a probability of detection of 75% with a FA rate of 0.06 FA/m^2. This is clearly undesirable. However, recall that we use only 4 of the 50 spectrum features in the training of the classifier. Thus, we examined other combinations (of 4 features) of the 50 spectrum features to see if we could find features that would better generalize across the data sets.

In a second experiment, we examined other sets of spectrum-features to determine if we could find a set of 4 features that would result in better generalized performance. Figure 7 illustrates the results of this experiment. We first trained a spectrum-feature classifier on Test Lane A (the training lane) for all possible sets of 4 spectrum-based features. We then examined the resulting performance on Test Lane B (the testing lane). View (b) shows the resulting detection characteristics for the classifier using bins [22, 29, 39, 42] of the spatial FFT. As this plot shows, by using these features the FA rate on the test lane was reduced from 0.11 FA/m^2 to 0.06 FA/m^2 while maintaining a 90% probability of detection. View (a) shows that the training lane performance is slightly degraded as compared to the results in Fig. 6(a); however, we stress that there is still a 15% reduction in FAs. The results shown in Fig. 7 are promising as this shows that by choosing a different set of features, we can train a classifier that performs better for both the training data and the testing data.

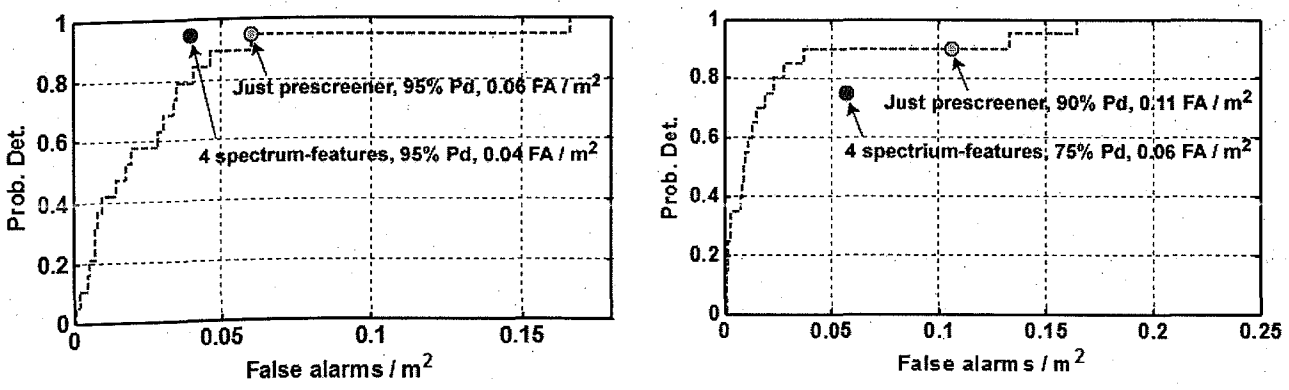

(a) Training (resubstitution) results on Test Run A

(b) Test results on Test Run B

Fig. 6. Training and testing results of one-class classifier with 4 spectrum-based features – bins [23,32,33,50] of FFT. Feature selection based on best training (resubstitution) results.

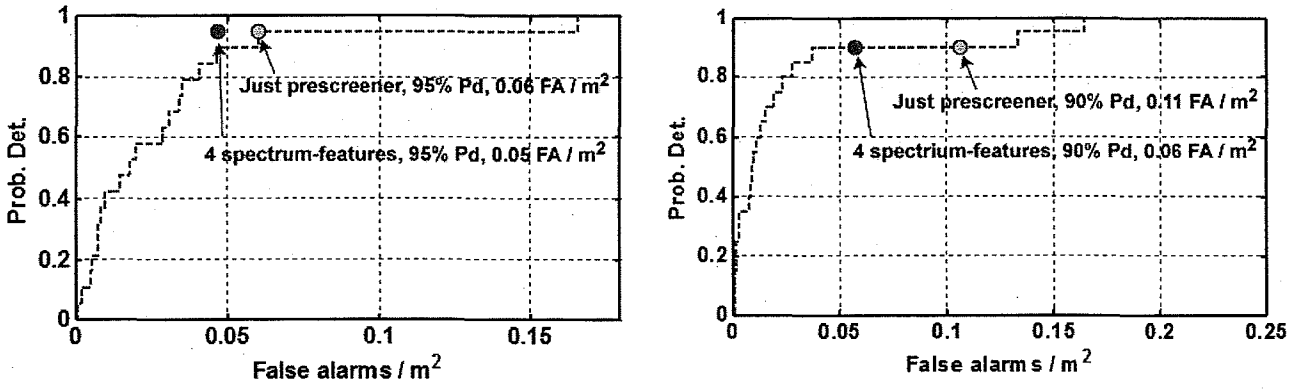

(a) Training results on Test Run A

(b) Test results on Test Run B

Fig. 7. Test results of one-class classifier with 4 spectrum-based features – bins [22,29,39,42] of FFT. Feature selection based on best test results. This feature selection method results in a more generalized classifier.

3.2 Image-feature classifier test results

Figure 8 illustrates the performance of the image-feature classifier. The red dot in view(a) indicates the performance using the set of 4 camera-based features that minimize the FA rate while maintaining at least 90% probability of detection on the training data, Test Run A. The 4 features selected by our exhaustive search were skewness of the pixel intensity, the minimum of the Laplacian, the mean of the Laplacian, and the median of the Laplacian. View (b) shows the resulting performance of the trained image-feature classifier on the test data, Test Run B. As this plot shows, the probability of detection was not reduced (as was seen with the spectrum-feature in Fig. 6); however, the FA rate was negligibly reduced. Note that the results in this section do not include the spectrum-feature classifier described in Section 3.1. In Section 3.3 we specifically discuss fusing the two classifiers.

As in the previous section, we ran a second experiment in which we examined other sets of 4 image features, with the intention of finding a set that better generalized. Thus, we trained the classifier on all possible sets of 4 image features from the training data, Test Run A, and then examined the performance of these classifiers on Test Run B. Figure 9 shows that using the skewness of the pixel intensity, the skewness of the Laplacian, the median of the local standard deviation, and the minimum of the red channel results in a more generalized classifier. The FA rate on the test data was reduced from 0.11 FA/m^2 to 0.08 FA/m^2 at 90% probability of detection. Notice, however, that the FA rate in the training data was only slightly reduced. However, we believe that this method of selecting the features results in a more generalized classifier, which is essential in an operational system.

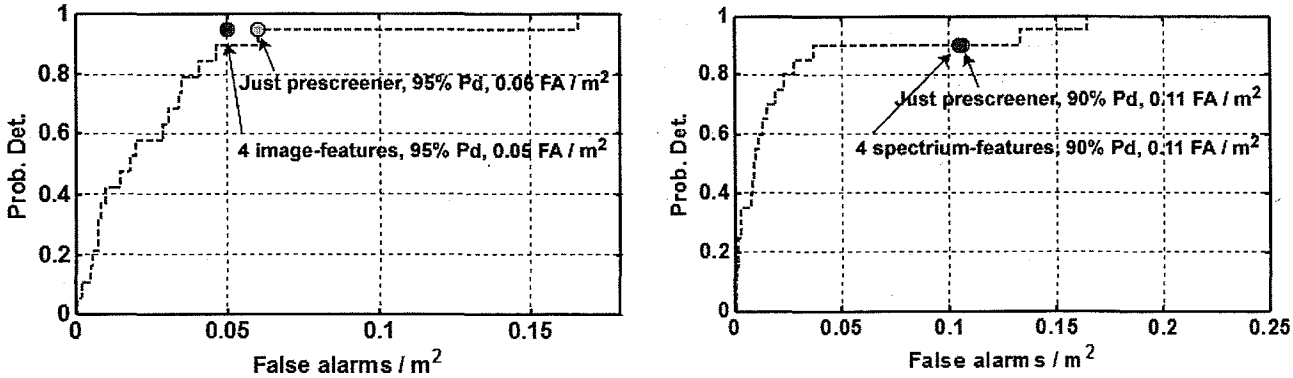

(a) Training (resubstitution) results on Test Run A

(b) Test results on Test Run B

Fig. 8. Training and testing results of one-class classifier with 4 image-based features – skewness(intensity), minimum(Laplacian), mean(Laplacian), median(Laplacian). Feature selection based on best training (resubstitution) results.

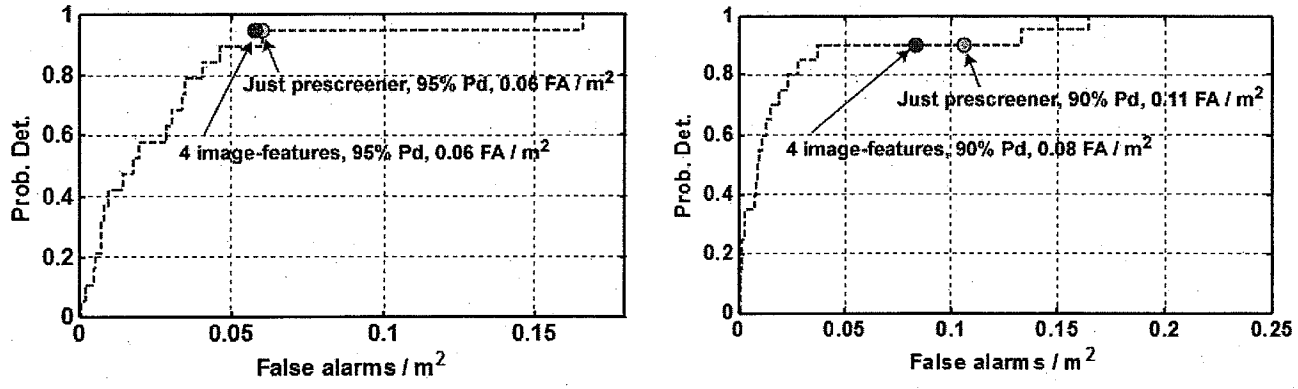

(a) Training results on Test Run A

(b) Test results on Test Run B

Fig. 9. Training and testing results of one-class classifier with 4 image-based features – skewness(intensity), skewness(laplacian), median(local standard deviation), minimum(red channel). Feature selection based on best test results.

3.3 Fusion test results

Sections 3.1 and 3.2 showed the test results for the spectrum- and image-feature classifiers, respectively. We now show the performance of the system when these two classifiers are fused. The first step in our detection algorithm is to apply the locally-adaptive threshold detector.[11] The ROC curve of this detector is shown as the blue dotted line in all the figures in this section. Thus, we first choose a threshold that gives the least number of FAs with at least 90% probability of detection. This is shown as the blue dots in Fig. 10. Second, we fuse the spectrum- and image-feature classifiers using a logical OR. If either classifier determines that an alarm is a FA then the fused result is a FA.

Figure 10 shows the results of our fusion experiment. View (a) shows the resulting FA rate on the training data and view (b) shows the resulting FA rate on the testing data. For these results, we used the set of features that resulted in the best generalized classifier performance – these features are listed in the captions of Figs. 7 and 9. As Fig. 10 shows, the fusion of the spectrum- and image-features classifiers causes significant reduction in FAs in both the training data and the testing data. The training data FA rate was reduced from 0.06 FA/m^2 to 0.03 FA/m^2, a 50% reduction, while maintaining a 95% probability of detection. The FA rate in the test data was reduced from 0.11 FA/m^2 to 0.05 FA/m^2 while maintaining a 90% probability of detection. These results show that our FA rejection method is very effective.

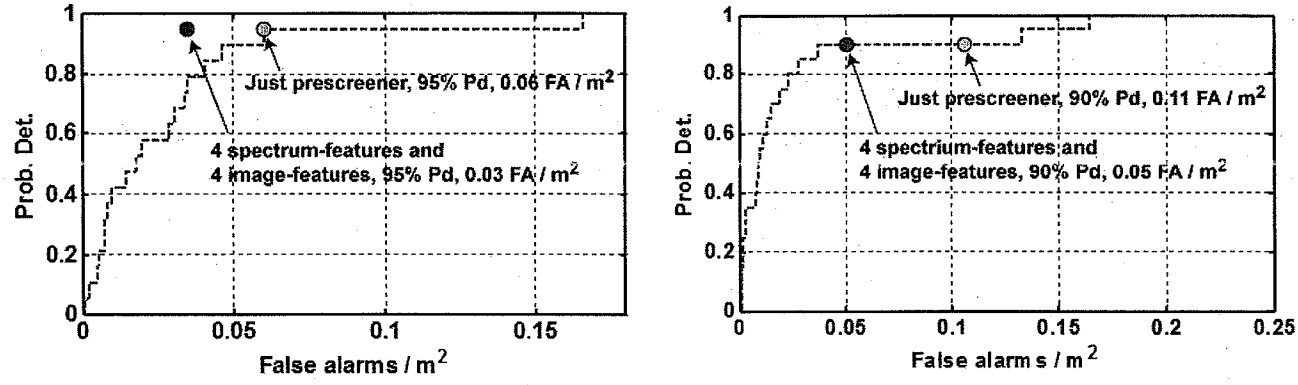

(a) Training results on Test Run A

(b) Test results on Test Run B

Fig. 10. Test results and training results of fusion of spectrum- and image-based false alarm rejection methods. Feature selection based on best test results.

4. CONCLUSION

The fusion of FLGPR- and image-based classifiers resulted in a significant reduction in FAs in both training and testing data. As Fig. 10 showed, there was a 50% reduction of FAs in the training data and a 55% reduction in the testing data. These results are promising as they show that we can build a FA-rejection classifier that is effective both in training *and* testing scenarios. Thus, we believe that the features we have chosen will generalize well to an operational environment. Additionally, the features and detection methods we have employed are computationally inexpensive and robust. Hence, in practice, our methods could be implemented in a real-time system.

In the future we will examine ways in which our algorithm can be tuned to different types of explosive hazards. For example, different FLGPR center frequencies and bandwidths, image features, or spectrum features may be optimal for different types of targets. We are also experimenting with different camera-based features, such as more complex texture-based measures, Zernike moments, and fractal dimension. Finally, the methods described in this paper used the FLGPR to detect the targets and the spectrum and images to reduce the FAs. We believe that the images could be used in tandem with the FLGPR to detect targets, and we have already begun work in this realm. We are also examining the fusion of cross-platform sensors to improve the detection / FA rate performance. Overall, the fusion of the FLGPR- and image-based explosive hazards detection approaches shows promise for significantly contributing to the remediation of the explosive hazards threat.

ACKNOWLEDGEMENTS

The authors would like to thank Justin Farrell for his contributions. This work was supported by grants from the Leonard Wood Institute (LWI 181-222) and Army Research Office (48343-EV) in support of the U.S. Army RDECOM CERDEC NVESD.

REFERENCES

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

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

[3] Costley, R.D., Sabatier, J.M., and Xiang, N., "Forward-looking acoustic mine detection system", Proc. SPIE 4394, 617-626 (2001).

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

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

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

[7] Cosgrove, R.B., Milanfar, P., and Kositsky, J., "Trained detection of buried mines in SAR images via the deflection-optimal criterion", IEEE. Trans. Geoscience and Remote Sensing 42(11), 2569-2575 (2004).

[8] Sun, Y., and Li, J., "Plastic landmine detection using time-frequency analysis for forward-looking ground-penetrating radar", Proc. SPIE 5089, 851-862 (2003).

[9] Stone, K., Keller, J.M., Ho, K.C., and Gader, P.D. "On the registration of FLGPR and IR data for the forward-looking landmine detection system and its use in eliminating FLGPR false alarms," Proc. SPIE 6953, (2008).

[10] Havens, T.C., Stone, K., Keller, J.M., and Ho, K.C. "Sensor-fused detection of explosive hazards", Proc. SPIE 7303, 73032A (2009).

[11] Havens, T.C., Ho, K.C., Farrell, J., Keller, J.M., Popescu, M., Ton, T.T., Wong, D.C., and Soumekh, M., "Locally adaptive detection algorithm for forward-looking ground-penetrating radar", Proc. SPIE (2010).

[12] Auger, A. and Hansen, N., "A Restart CMA Evolution Strategy with Increasing Population Size", Evolutionary Computation 2, 1769-1776 (2005).
[13] Hansen, N., Ostermeier, A., "Completely Derandomized Self-Adaptation in Evolution Strategies", Evolutionary Computation 9(2), 159-195 (2001).
[14] Hansen, N., "The CMA Evolution Strategy: A Tutorial", http://www.bionik.tu-berlin.de/user/niko/cmatutorial.pdf (2007).
[15] Wang, T., Sjahpetura, O., Keller, J.M., and Gader, P.D., "Landmine detection using forward-looking GPR with object-tracking," Proc. SPIE 5794, 1080-1088 (2005).
[16] Theodoridis, S., Koutroumbas, K., [*Pattern Recognition* (3rd ed.)], Academic Press, San Diego, CA (2006).
[17] Duda, R.O, and Hart, P.E., [*Pattern Classification and Scene Analysis*], John Wiley & Sons Inc. (1973).