

Results from a forward-looking GPR mine detection system

Joel Kositsky^{*a}, Russell Cosgrove^a, Charles Amazeen^b, Peyman Milanfar^c
^aSRI International; ^bNVESD, Fort Belvoir, VA 22060;
^cDept. of Electrical Engineering, University of California, Santa Cruz, CA

ABSTRACT

In previous papers, we reported on the high-resolution ground-penetrating radar system designed, built, and deployed by SRI under contract to the Night Vision and Electronic Sensors Directorate at Fort Belvoir. Here, we report on some of the latest test results from the field demonstrations performed at government test sites, carefully designed to produce statistically significant results by employing many samples of a few representative metal and plastic mine types, buried at several depths. Significant improvements in performance above the baseline have been realized by using a number of statistically optimal image processing algorithms based on sound mathematical principles and techniques, guided by electromagnetic models of both mines and clutter. Principal component analysis was employed to define empirical models of buried mines using the polarimetric, complex data. A detector based on the generalized likelihood ratio test was then used on each image. Finally, multilook processing combined the results from several independent views of the same mines (at various ranges). For buried metal mines, at a probability of detection of 94%, the probability of false alarm per m² decreased by an average factor of about four orders of magnitude over the baseline.

Keywords: radar, GPR, SAR, land mine, wideband, forward-looking, standoff, GLRT, PCA

1. INTRODUCTION

In 1999, 2000, and 2001, we reported on a synthetic aperture based, forward-looking ground-penetrating radar (FLGPR) system built as part of the Night Vision and Electronic Sensors Directorate (NVESD) forward-looking antitank (AT) mine detection efforts.^{1,2} This system was designed with a goal of assessing the capability of a FLGPR in detecting plastic- and metal-cased surface and buried AT mines on roadways. In this paper, we review some of the main parameters of this system and report on the results from the demonstrations and tests at government facilities. We will also address the significant work done on clutter suppression and target recognition using algorithms that improve the detection and false alarm statistics.

2. HARDWARE SYSTEM

The data collection FLGPR system has been described in some detail in previous papers;^{1,2} we will quickly review some of the highlights of the hardware.

The system was designed and built as a multipolarization, wideband, high-resolution research and development tool. Calculations and experience showed that the clutter returns from the ground surface would be the most challenging impediment to obtaining good performance statistics. As such, we incorporated those features that would be the most useful in maximizing the signal-to-clutter ratio (SCR). In particular, excellent spatial resolution was considered key. Range resolution of 5 cm was incorporated by using a very large frequency bandwidth of 0.3 to 3.0 GHz. Azimuthal resolution of 15 cm at 7 m was achieved using a (synthetic) aperture of 4 m. This was accomplished by mechanically stepping antennas in 5 cm intervals, perpendicular to the direction of travel, using a computer-controlled, linear drive.

The received signals were mixed with the transmitted frequency in a homodyne receiver to produce in-phase (I) and quadrature-phase (Q) signals at baseband. The transmit and receive signals were gated in such a way as to match the

* Contact joel.kositsky@sri.com; phone 1 650 859-6035; fax 1 650 850 5149; SRI International, 333 Ravenswood Avenue, Menlo Park, CA 94025

required range extent, and to boost the signals from the far ranges relative to those from near ranges. In this way we were able to decrease dramatically the dynamic range requirements of the receiver. In the GPR system tests, we concentrated on a range extent of 7 to 30 m with a nominal cross range extent of 8 m.

The antennas are bilinearly polarized, quad-ridged horns. High-speed switches connect the radar electronics to the two feeds of each horn, so that all linear polarization combinations can be addressed to produce fully polarimetric data products.

The radar parameters are summarized in Table 1, and a cartoon of the radar implementation is shown in Figure 1.

Table 1 Radar Parameters

| | |
|------------------------------|---------------------------|
| Radar approach: | Stepped frequency SAR |
| Transmitter waveform: | Chopped CW |
| Peak power: | 0.5 W |
| Duty cycle: | 6% |
| Frequency span: | 0.3 to 3.0 GHz |
| Frequency step size: | Variable, 2.5 MHz typical |
| Range: | 7 to 30 m |
| Cross range: | 8 m |
| Azimuthal resolution at 7 m: | 15 cm (range dependent) |
| Range resolution: | 5 cm (range independent) |
| Receiver: | Homodyne/IQ |
| Digitization: | 12 bits (nominal) |
| Polarization: | HH, VV, HV, VH |
| Antenna elements: | Quad-ridged horns |
| Antenna gain: | 7 to 12 dBi |
| Antenna arrays: | 1 Tx, 1 Rx (independent) |
| Array orientation: | 4 m horizontal |
| Antenna scanning: | Mechanical and electrical |

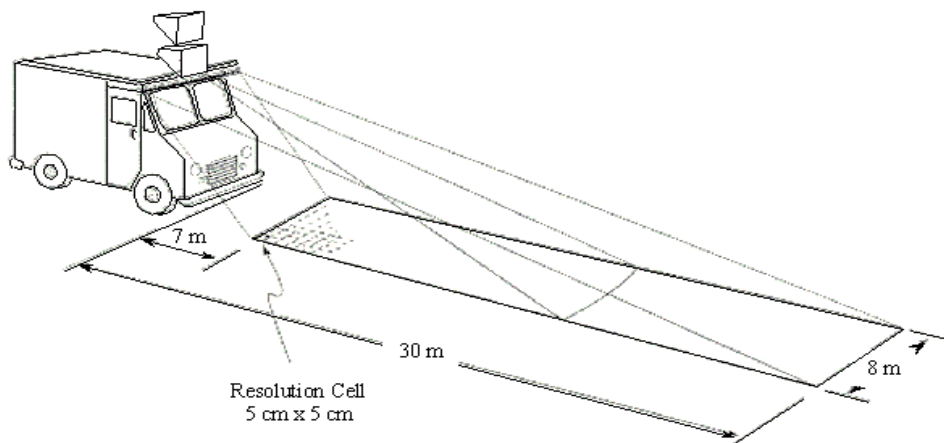


Figure 1: The forward-looking ground penetrating radar concept and major parameters.

3. FIELD TESTS

In 1999, we collected data with our radar at a government mine lane test site. These data showed mixed results, typically demonstrating good detection of metal mines and spotty detection of plastic mines. Some plastic mines were missed altogether, while others were detected with good SCRs. Because the number of any particular type of mine was quite small (two or three at best, and these at different burial depths), the results must be considered anecdotal rather than statistically significant, and are not described in this paper other than to show a sample image. Figure 2 is a gray-scale image from the first test lanes showing three buried plastic mines: (1) VS1.6 at 6" depth and 10 m range, (2) simulation mine #25 at 2" depth and 17 m range, and (3) an M19 at 5" depth and 20 m range.




Figure 2: Gray-scale image showing buried plastic mine detections at first test site.

In an effort to collect statistically significant data, several mine lanes were set up at two government test sites in 2000 and 2001. Only three types of AT mines were selected for these tests, but 60 of each type were deployed, so that about 12 of each type would be deployed at each of 5 depths. As these mines would be viewed from two directions, and from many ranges as the FLGPR moved along the test lane, a total of over 120 encounters would be made with a given mine type at a given depth.[†]

The three mines selected were the TM-62M, the TM-62P, and the VS2.2. The TM-62 mines were selected because of their availability in metal (TM-62M) and plastic (TM-62P) versions, differing in casing material, but having very similar size and shape. The third mine type, the VS2.2, is also a plastic-cased mine, but it has a significantly different profile from the TM-62P. The main features of the three mine types are listed in Table 2. The high explosive was not removed from the mines, as its dielectric properties are potentially important in radar detection of plastic-cased mines. For safety reasons, the boosters were removed and replaced with RTV molded to shape and containing the same small amount of metal as in the active mine.

Table 2 Three AT Mines Used in the Main Tests

| <u>Feature</u> | <u>TM-62M</u> | <u>TM-62P</u> | <u>VS2.2</u> |
|----------------|---------------|---------------|--------------|
| Case material | Metal | Plastic | Plastic |
| Diameter (cm) | 32 | 32 | 18 |
| Height (cm) | 10.2 | 12.5 | 21.8 |
| Weight (kg) | 8.5 | 8.2 | 3.54 |
| Metal content | Case | In fuze | 5.0 g |
| Explosive type | TNT/RDX/AL | TNT | TNT/RDX |

[†] The number of strictly independent looks at a given mine type at a given depth is, of course, less.

The mines were emplaced in prepared lanes at one of five depths (surface emplaced; flush buried; or buried 5, 10, or 15 cm from the ground surface to the top of the mines). Several extra holes were dug in the same manner as for mine burial, but were left empty and simply backfilled. These refilled “holes” were used as controls to separate the detection of mines from the detection of soil disturbance.

The lanes were lightly graded following the flat to gently undulating profile of the areas. The three different types of lanes consisted of gravel, dirt, and native desert soils. Several areas in each lane were left free of targets, so that these could be used for collecting clutter statistics. The mines were buried at 2 to 5 m intervals in range, at three cross ranges—0 and ± 1.5 m from the nominal centerline of the lanes.

4. DATA COLLECTIONS

The forward-looking ground-penetrating radar van was driven along the (nominal) centerline of each lane during data collection. At 2 to 5 m intervals the van was stopped, and a synthetic aperture radar (SAR) scan was performed. During each scan, all frequencies and polarization combinations were used before the antennas were mechanically translated in 5 cm steps. Once the antennas spanned the entire 4 m aperture, the scan was completed and the van was driven forward before stopping for the next scan.[‡] Differential global positioning system receivers operating in real-time kinematic mode registered the position and orientation of the antennas during the scans. During the last collections, each complete scan took about 2 minutes, with the data imaging running concurrently.[§]

Data were collected with the radar moving in both lane directions to produce hundreds of data files, each containing HH, VV, HV, and VH polarization data over the 0.3 to 3.0 GHz bandwidth in dry, moist, and wet conditions. For calibration purposes, data were also collected of calibration targets (retroreflectors, spheres, and hemispheres on ground planes), loop-back, ambient RFI levels, and receiver noise levels.

5. PROCESSING AND RESULTS

5.1 Basic processing

After being corrected by the independently measured system transfer function, the frequency domain data are changed to equivalent time domain returns, for each antenna position. These returns are then coherently combined, using the appropriate geometrical factors to each pixel in the scene, to produce a focused geo-rectified complex image for each polarization, using a standard SAR back-projection algorithm.

A visual inspection of the imagery gave us a “feel” for the data as we compared and contrasted amplitude imagery as a function of polarization, frequency band, and range. We also compared the data products with physics-based simulations using both MOM code (supplied by Duke University) and analytical expressions derived from the physical optics (PO) approximation to validate the proper functioning of the radar as well as our understanding of the results.

Further processing included implementing several standard and mathematically well-understood algorithms, as well as extensions to these standard techniques. In developing our more sophisticated algorithms, we would state the necessary approximations and assumptions, and carefully build the mathematics on sound methodology, avoiding ad hoc processes whenever possible. As we applied the different techniques, we compared results using common automatically generated performance metrics such as ROC curves, probability of false alarm (PFA) at a given probability of detection (Pd), divergence, and SCR. We have thus constructed a set of algorithms that have a strong mathematical basis, and are tested against the real data to validate the approximations and assumptions.

[‡] We are presently building a test system that collects data while moving, using a real (as opposed to a synthetic) antenna array.

[§] In the real-time system presently being built, a scan will take about 12 ms and the data will be imaged in real time.

In the following sections we will present some of these processing schemes and the resulting detection improvements. The baseline for the comparisons are the results generated by the amplitude of the back-projection HH images without any further processing except some frequency filtering and whitening to improve the baseline results.**

Another performance benchmark was established by combining the complex data from the two main polarizations using the standard polarimetric whitening filter (PWF) approach.³ Our most successful processing scheme, which consistently outperformed PWF for the FLGPR data, is described in the sections that follow.

5.2 Baseline results

In previous papers, we discussed the baseline data results as functions of the major parameters of interest. These included:

- mine placement, surface and buried (also depth);
- mine type, plastic and metal;
- range (standoff distance);
- radar polarization; and
- spectral bands.

The high-frequency band was very useful in detecting surface mines, particularly for the low-profile plastic TM-62P mine. Detection performance of surface mines, including plastic-cased ones was excellent.

The flush-buried metal-cased mines were quite detectable by the low-frequency band, but were only poorly detected by the high-band frequencies. Being covered by only a small layer of soil, the flush-buried mines showed the best high-frequency response of all buried mines. However, because of the poor response to high frequencies, we used a low-pass (2 GHz) filter in the detectors for buried mines.

The HH and VV baseline results for mines buried at the four depths (flush with the surface, and at 5, 10, and 15 cm) showed the expected monotonic decrease with burial depth. This trend was quite consistent at each range. It is worth mentioning that the HH data consistently outperformed the VV data. The result is explained by the larger buried target returns for VV (due to better penetration near the Brewster angle) being more than offset by the increased ground clutter returns in VV polarization. With the baseline processing, the plastic mines were poorly detected at all burial depths.

5.3 Image signature

In Figure 3, a small section of an amplitude HH image is displayed, showing a typical signature of a buried mine. A cut through the mine image in the range direction (dotted line) yields a characteristic “double-hump” (shown in Figure 4a), which is present more or less consistently in all collected data of buried metal mines at varying levels of strength. Variation in this “signature” are to be expected due to possible tilt in the mine, inhomogeneities in the soil over and around the mine, and the soil surface, which acts like a phase screen that degrades the image coherence. Significantly, we have shown analytically that burial depth will not dramatically affect the form of the mine signature: This has been borne out by the data.

We further note that the structure of the target signature is quite a bit less interesting in the cross-range direction primarily because of the poorer resolution in cross range. Therefore, we confined our modeling and estimation efforts to capturing the more pronounced target signature in the range direction. Figure 4a shows range cuts through buried mine image chips while Figure 4b shows PO model-generated image cuts. The similarities between the results confirm that the “double-hump” signature is both predicted by theory and validated by the data. The method we used to capture and exploit this buried-mine signature is described in the next section.

** Typically, the high-frequency band added more clutter than signal for buried mines.




Figure 3: Image chip showing a buried metal mine.

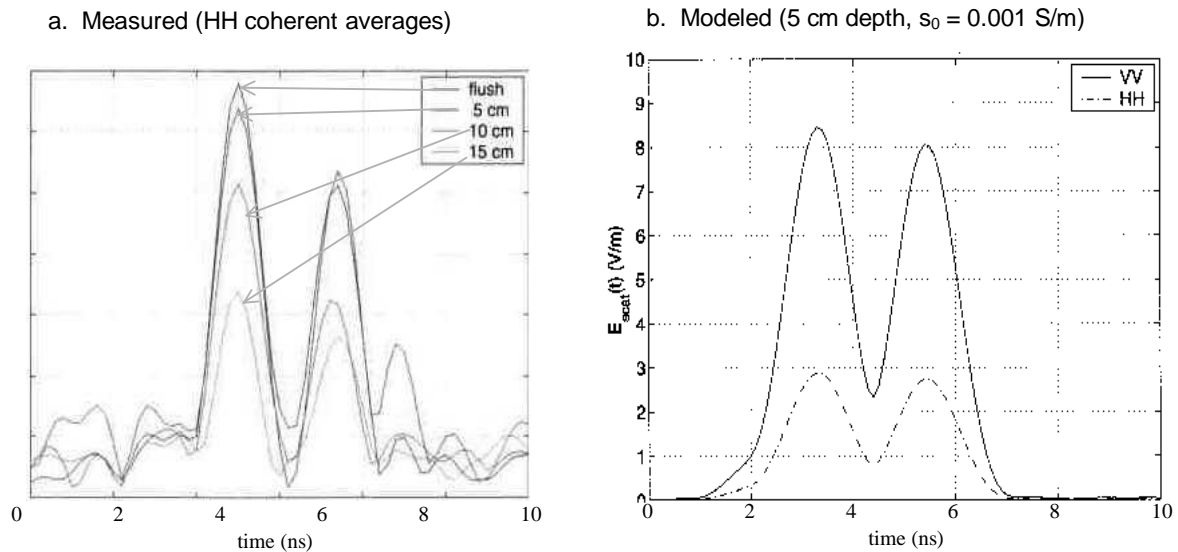


Figure 4: Amplitude range cuts through buried metal mine images.

Because our generalized likelihood ratio test (GLRT) detector required a data-based model for implementation, we wanted to test the processing on mine data that exhibited a reasonably strong “signature.” The surface mines give very consistent signatures with high SCRs, but are in some sense too “easy” a case, as the baseline results are already excellent. As an exercise, we did form signatures for the three kinds of surface mines (TM-62M, TM-62P, and VS2.2). We used the models in an automated test that placed the unknown targets into the correct mine type category in almost 100% of the cases. In future tests, we will employ these models to measure how well we can differentiate surface mines from surface clutter objects like rocks and shrapnel. On the other extreme, the buried plastic-cased mines were too difficult a case for this approach since the SCR of the baseline imagery was usually too low to allow a reasonable extraction of signatures.^{††} Therefore, we tried this processing using buried metal mine data, which showed a fairly consistent “signature,” as well as significant SCRs.

^{††} For buried plastic mines, we plan to use signatures derived from physics-based models.

5.4 A model-learning approach to generalized likelihood ratio testing (GLRT)^{‡‡}

First, we divided the experimental data into a training set and a testing set. Using the training set, we built a model to effectively describe the target data using a variant of the well-known Principal Components Analysis (PCA), which, under Gaussian noise assumptions, provides the optimal estimate (in the maximum likelihood sense) of the subspace best describing the data. We then modeled each measured target response \mathbf{x}_k in the training set as a linear combination of *modes*, plus noise,

$$\mathbf{x}_k = \mathbf{M} \boldsymbol{\theta} + \boldsymbol{\omega}_k \quad (1)$$

where \mathbf{M} , the matrix of modes, and $\boldsymbol{\theta}$, the mode weights, are unknown. Assuming $\boldsymbol{\omega}$ is zero-mean, complex Gaussian noise with covariance \mathbf{C} , the maximum likelihood (ML) estimate for $\boldsymbol{\theta}$ is⁴

$$\hat{\boldsymbol{\theta}} = (\mathbf{M}^H \mathbf{C}^{-1} \mathbf{M})^{-1} \mathbf{M}^H \mathbf{C}^{-1} \mathbf{x}_k \quad (2)$$

A good estimate for the covariance matrix, \mathbf{C} , can be obtained from the clutter data collected by the FLGPR. Once \mathbf{M} is constructed it can be used on the testing set to find the estimate of $\boldsymbol{\theta}$ by using Equation 2.^{§§}

It can be shown that the maximum likelihood estimate for \mathbf{M} can be obtained by computing the Singular Value Decomposition (SVD) of the data matrix \mathbf{X} formed by the column vectors \mathbf{x}_k . The process is illustrated schematically in Figure 5. First, column vectors are formed from the complex values of range cuts through mines in the HH and VV images, which are then phase-normalized and concatenated into a single column vector \mathbf{x}_1 . Other column vectors \mathbf{x}_k are formed from images of other mines in the training set and combined to form the “data matrix” \mathbf{X} .

After removal of the mean of $\{\mathbf{x}_k\}$, the SVD of this matrix is performed:

$$\mathbf{X} = \mathbf{U} \mathbf{S} \mathbf{V}^H, \quad (3)$$

The first few columns of \mathbf{U} give \mathbf{M} , the ML estimates of the modes. Note that these modes form an orthogonal basis describing the model (i.e., PCA). Note, too, that since the data vectors \mathbf{x}_k contain components from both the H and V polarizations, the procedure is more general (and in fact encompasses) the polarimetric whitening procedure. Finally, we note that we select only a few (1 to 3) of the columns of \mathbf{U} to represent the model, as explained below.

Additionally, when data from other sensors are available, the new data can be concatenated with the columns of the matrix \mathbf{X} to provide a natural framework for fusion of multisensor data. While we have not specifically done this, it would be rather straightforward to implement the concept.

Having estimated a model matrix $\hat{\mathbf{M}}$ from the “training” data $\{\mathbf{x}_k\}$, we apply this model to detect targets in the “testing” data set $\{\mathbf{y}_k\}$ by computing the *generalized* log likelihood function and comparing with a threshold γ (chosen to satisfy an acceptable false alarm or detection rate).

$$\mathbf{y}^H \hat{\mathbf{M}} \mathbf{D} \hat{\mathbf{M}}^H \mathbf{y} \begin{matrix} \geq \\ \leq \end{matrix} \gamma \quad (4)$$

where \mathbf{D} is a diagonal matrix related to \mathbf{S} (in Equation 3). This detector has a fairly simple structure and lends itself to some pleasing interpretations. In particular, the detector is, in effect, a weighted *energy detector*. However, the detected energy is computed only in the *signal subspace* defined by \mathbf{M} . In addition, the *signal subspace* given by \mathbf{M} is itself

^{‡‡} A journal article, being prepared by R. Cosgrove and P. Milanfar, will expand this abbreviated presentation.

^{§§} In practice, explicit estimates of the parameter vector $\boldsymbol{\theta}$ are not needed.

learned from the training data. Furthermore, the weights in each subspace are assigned by D in accordance with the SNR within the subspace.

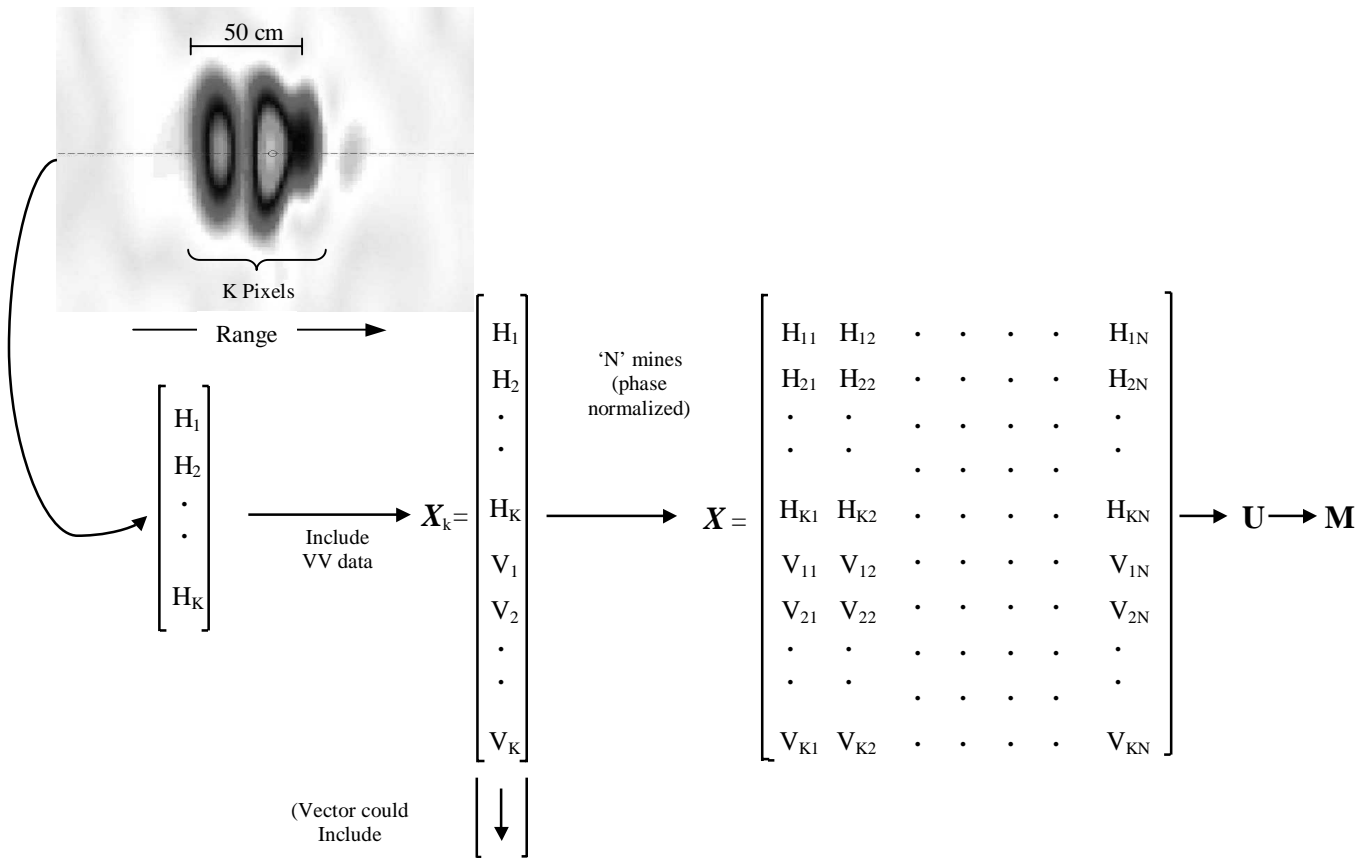


Figure 5: Finding the principal components from the data images.

An important question arises about the number of principal components that should optimally be included in the matrix M (i.e., to what rank should we limit M to get the best performance (high Pd, low PFa). Basically, we are assessing how well a given data vector matches a linear combination of a set of vectors (the principal components). If we increase the number of component vectors, we will increase how well the data from the mines match the model; but we will also increase the number of clutter vectors that match well. In the extreme case of a complete set of principal components that span the data space, all data will match, and the Pd and PFa will *both* approach 100%.

We wish to choose the rank of M so as to maximize the SCR. Empirically, we have found that two principal components do the job best, producing slightly better results than a detector based on a single principal component. This approach relies on knowing the ground truth or having an extensive training set. In practice, the data imply that we could elect to use a single component, for computational efficiency, without sacrificing much in the way of detector performance.

Using the approach outlined in this section, we can construct a database of various targets in different environments for operational reference. In the field, the local adaptation to the given environment can then be effected by making a sparse set of measurements and deciding which of the many models learned earlier and stored in the database are most appropriate for the construction of the matched subspace detectors.

Finally, it is worth stating that it is possible to generate the estimated models from physical first principles, a priori, using the modeling techniques such as those described previously.² In fact, it will probably be impractical to derive models for all AT mines under various conditions using experimental data. Physics-based modeling must eventually take over this task.

Figure 6a is a linear scale ROC curve for the 24 metal mines buried to depths of 15 cm in a wet gravel lane. Figure 6b is the same ROC curve, but presented as a semi-log plot. These figures show the typical improvements in detection statistics as the processing increases in sophistication from HH baseline processing (green trace), to PWF processing (blue trace) and finally to PCA/GLRT processing (black trace at 100% Pd).




Figure 6: ROC curves of detection of buried TM-62M mines at 15 m range using different processing algorithm.

5.5 Multilook processing

Multilook processing refers to combining the data derived from the same patch of ground as seen from various viewpoints as the radar drives down the lane. The hope is that the mine target will correlate well from one view to another, while the ground clutter will decorrelate more rapidly. However, if the clutter, which is the bulk of the image, does decorrelate significantly, registering the various images to a precision that would permit coherent processing will probably be impossible. So, rather than adopting a processing scheme that relies on finely registering the complex images taken at different ranges, we adopted a more forgiving approach.

Each potential target in a PCA/GLRT analyzed image was assigned to a 1 m² area in the lane, centered on the pixel with the largest amplitude. This same square-meter area was then examined in other radar images, formed from data collected at other van positions. The pixel with the largest amplitude in the 1 m² area of the new image was then taken as being from the same potential target. The peak pixels, p_1 , p_2 , and p_3 , in each one meter square region from three separate images are combined according to

$$p = k_1 p_1 + k_2 p_2 + k_3 p_3, \quad (5)$$

where k_1 , k_2 , and k_3 are weights chosen to emphasize the better detection statistics of the radar at closer ranges. Equation 5 represents an interim (admittedly ad hoc) approach to the problem of how best to combine the several radar looks, while we pursue a more rigorous solution.

Registration error only arises as a problem if the peak pixel in one of the cells is very close to the boundary—within the registration error. In this case a target’s peak may be in different regions for different looks. However, these cases may be screened for and an additional cell centered on the actual peak added. Hence, this approach to multilook is robust with respect to registration error.

The probability density function (PDF) under the H1 hypothesis (mines present) for Equation 5 turned out to be much the same as for a single look. This is because the mines appear very similar to the radar, regardless of range. Clutter areas arising from unfocused radar returns were far less consistent from look to look. Hence, the PDF under the H0 hypothesis (mines absent) for Equation 5 showed a substantially reduced variance, which translates to an improved detector. This is illustrated in Figure 7, where the red and black traces represent the clutter and target plus clutter PDFs, respectively; and Figures 7a and 7b show the curves before and after multilook processing, respectively. The ROC curves that result from the PDFs shown in Figures 7a and 7b are shown in Figures 8a and 8b, respectively.

a. PDFs: Single look PCA/GLRT processing

b. PDFs: Multilook after PCA/GLRT processing

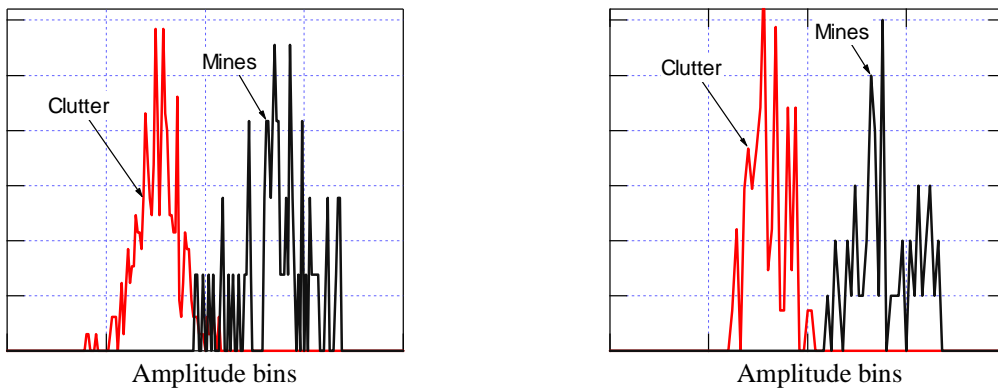


Figure 7: Probability density functions for clutter and buried TM -62M mines (plus clutter).

a. Linear scale

b. Semi-log scale




Figure 8: ROC curves of multilook and single look performance with several algorithms.

Some false alarm rates for TM-62M metal antitank mines at a 94% detection rate are presented in Table 3. The results are presented separately as a function of mine burial depth. The first three rows show false alarm rates derived from single look images after PCA/GLRT processing, for mines at ranges of 20, 15, and 10 m. The false alarm rates at each

burial depth are shown as a PFa (a number between 0 and 1), and a dB measure ($10 \times \log[\text{PFa}]$) for clarity; both of these are statistics per square meter using a $1 \text{ m} \times 1 \text{ m}$ halo. The single row marked “Multilook” shows the results for our multilook processing that combines the information from the images of the mine at the three ranges.

Table 3 False Alarm Rates/m² at 94% Detection Rates for Single and Multilook Processing

| Range (m) | Flush-Buried (dB) | | 5 cm Deep (dB) | | 10 cm Deep (dB) | | 15 cm Deep (dB) | |
|-----------------------------------|-------------------|-----|----------------|-----|-----------------|-----|-----------------|-----|
| 20 | 1.8 E-03 | -27 | 2.2 E-04 | -37 | 5.0 E-06 | -53 | 1.9 E-02 | -17 |
| 15 | 3.5 E-06 | -55 | 5.0 E-03 | -23 | 1.8 E-06 | -57 | 1.2 E-03 | -29 |
| 10 | 2.5 E-03 | -26 | 7.0 E-05 | -42 | 6.5 E-05 | -42 | 1.1 E-04 | -40 |
| Multilook | 9.5 E-09 | -80 | 4.0 E-07 | -64 | 1.8 E-09 | -87 | 5.0 E-07 | -63 |
| Improvement over best single look | | 26 | | 22 | | 30 | | 23 |

The final row in the table shows the improvement achieved using multilook over the best single look false alarm rate. The improvement (decrease) in false alarm rates is consistently between two and three orders of magnitude, with an average improvement of about 25 dB.

In order to emphasize the operational significance of this result, we could retabulate the dB false alarm rates for the same situations illustrated by Table 3 and include the resulting mean distances and times that the vehicle could travel before registering a false alarm. If we assumed a 3 m wide roadway, a 1 m^2 halo size around nominated targets, a 94% probability of detection, and a travel speed of 30 kph, the results in Table 3 would indicate that we could travel a kilometer or so using the single-look statistics. Using the multilook processing, we could travel false-alarm-free 100s of kilometers and for many hours. False alarm rates using the PWF alone or the thresholded baseline image were not tabulated, but would support shorter false-alarm-free travel distances of about one and two orders of magnitudes, respectively.

While the results in Table 3 are truly excellent, we need to point out several caveats. Firstly, these results were generated for metal TM-62M mines buried in the fairly clean and flat test lanes at government prepared sites. One of the three types of lanes was, however, prepared by simply and crudely grading a natural site. The results at that lane, which could be representative of many rough roads, were also excellent. Another important caveat is that actual false alarm rates and false-alarm-free distances, are dependent upon the structure of the tail of the distributions several standard deviations from the mean. Our sample size of 10s of buried mines does not allow us to extrapolate with confidence to several hours of travel.

We would expect that, because of attenuation in the ground, the system performance should degrade with burial depth. Table 3, however, shows potentially anomalously good results for the 10 cm deep mines. We believe that this is because the principal component model we used was derived from mines at that depth, thus skewing the results somewhat. If this is true, it also implies that we could do better at different burial depths if we had several models in our analysis from which to choose, using a simple extension of our detection procedure.

Plastic buried mines have yet to be tested using these techniques, but we know that the results will not match those for the buried metal mines. However, we are confident that *surface* plastic mines, will be very highly detectable with very low false alarm rates, since we believe that the techniques described here will do an excellent job of differentiating between surface mines and clutter. This thesis will be tested in the near future.

6. CONCLUSIONS

- Buried metal mines can be detected with the SRI FLGPR, at high Pd and low PFa, using multilook processing on PCA/GLRT processed images.
- Surface plastic and metal mines can be readily detected with the FLGPR, and have specific and robust signatures that may support a low false alarm rate even on cluttered roadways.
- Buried plastic mines are still presenting a great challenge to robust detection by stand-off radar.
- Modeling efforts show great promise in defining mine signatures a priori.

7. ONGOING AND FUTURE WORK

Ongoing and future work include the following:

- Implementing a real aperture radar that will
 - collect data on the move,
 - image the data, in real time,
 - provide some automatic target recognition in real time, and
 - be mounted on a Hummer
- Publishing a journal article to present our processing algorithms in great detail
- Extending our processing to surface mines and buried plastics mines
- Measuring discrimination of targets in more severe clutter environments
- Incorporating a priori knowledge about ranges of soil and surface parameters into our models
- Employing physics-based models for processing

ACKNOWLEDGMENTS

This work was supported by the U.S. Army, CECOM RDE Center, Night Vision and Electronic Sensors Directorate under Contract No. DAAB15-01-C-0023. Thanks to the Ft. Belvoir and SRI team members especially, Murray Baron, Kendall Casey, George Oetzel, and Gary Price, who all contributed material for this paper.

REFERENCES

1. J. Kositsky and P. Milanfar, "A forward-looking high-resolution GPR system," *SPIE Conference on Detection and Remediation Technologies for Mines and Minelike Targets IV*, A.C. Dubey, J.F. Harvey, J.T. Broach, and R.E. Dugan, eds., **3710**, pp. 1052-1062, SPIE, Orlando, 1999.
2. J. Kositsky, "Results from a forward-looking GPR mine detection system," *SPIE Conference on Detection and Remediation Technologies for Mines and Minelike Targets V*, A.C. Dubey, J.F. Harvey, J.T. Broach, and R. Dugan, eds., **4038**, pp. 1077-1087, SPIE, Orlando, 2000.
3. L. M. Novak, M. C. Burl, and W. W. Irving, "Optimal polarimetric processing for enhanced target detection," *IEEE Trans. on Aero and Elect Syst.*; **Vol. 29**, No. 1, January 1993.
4. L. Sharf, "Statistical Signal Processing," Wiley, 1993.