This paper originally appeared in *Detection and Remediation Technologies for Mines and Mine-Like Targets IV*, A. C. Dubey, J.F. Harvey, J. Broach, R.E. Dugan editors, Proc. SPIE V 3710, SPIE, Orlando, (1999)

Statistical fusion of GPR and EMI data

Robert A. Weisenseel[a*], W. Clem Karl[a], David A. Castañon[a],
Eric Miller[b], Carey Rappaport[b], Charles A. Dimarzio[b]

[a]Boston University, Electrical and Computer Engineering Department
[b]Northeastern University, Electrical and Computer Engineering Department

ABSTRACT

In this paper, we develop a statistical detection system exploiting sensor fusion for the detection of plastic A/P mines. We design and test the system using data from Monte Carlo electromagnetic induction (EMI) and ground penetrating radar (GPR) simulations. We include the effects of both random soil surface variability and sensor noise. In spite of the presence of a rough surface, we can obtain good results fusing EMI and GPR data using a statistical approach in a simple clutter environment.

More generally, we develop a framework for simulation and testing of sensor configurations and sensor fusion approaches for landmine and unexploded ordinance (UXO) detection systems. Exploiting accurate electromagnetic simulation, we develop a controlled environment for testing sensor fusion concepts, from varied sensor arrangements to detection algorithms. In this environment, we can examine the effect of changing mine structure, soil parameters, and sensor geometry on the sensor fusion problem. We can then generalize these results to produce mine detectors robust to real-world variations.

Keywords: mine, sensor fusion, statistical, radar, induction, model, UXO, detection

1. INTRODUCTION

The need for mine and unexploded ordnance (UXO) detection and removal is growing in both military and humanitarian applications. Civilian casualties from landmines are on the order of hundreds per week, and in some locations around the world, landmines are emplaced faster than they can currently be removed. The current U.N. standard for humanitarian mine detection probability is set at 99.97%. But with any single current buried object detection technology, such a high P_d results in an unacceptably large number of false alarms, particularly when searching for plastic anti-personnel (A/P) mines. However, through information fusion, we can combine the best aspects of multiple sensor technologies to achieve the goals of landmine removal.

To achieve an optimal information fusion system, we must start from fundamental, controlled studies of system design tradeoffs, taking into account every aspect of the problem, from sensor physics and configuration to fusion and detection. Unfortunately, little work has been done to date on such a comprehensive approach to system design. While some researchers have studied individual sensors and detection schemes extensively, sensor fusion researchers have not conducted necessary controlled studies of the effects of design choices on the fusion and ultimate detection problems.

This paper represents a first effort toward filling in this gap in our knowledge. Our goal is to develop a framework for mine detection system research and design that includes the critical aspects of a sensor fusion scenario. We have chosen to focus on detecting plastic A/P mines because their detection represent a particularly challenging problem for which there is currently no satisfactory fielded solution. Information fusion appears to be the solution for this difficult detection problem.

In this work, we concentrate on fusing information from two sensors, a Ground Penetrating Radar (GPR) and an Electro-Magnetic Induction Spectroscopy (EMIS) sensor. We demonstrate that even a simple sensor fusion design can be quite successful over either individual sensor when the physics of their operation provides complementary information. However, we show that individual sensors can be strongly affected by variations in soil parameters and sensor configuration, and this in turn can affect our design choices sensor fusion and detection, which supports our contention that more extensive controlled studies are necessary.

*Correspondence: Email:lorax@bu.edu; WWW: http://mdsp.bu.edu/lorax_html

2. BACKGROUND

Our current approach to detecting plastic A/P mines is based on fusing data from two sensors, a Ground Penetrating Radar (GPR) array similar to that operated by Geocenters[1] and an Electro-Magnetic Induction Spectroscopy (EMIS) sensor such as those developed by Geophex.[2] Each type of sensor has its advantages and drawbacks, but neither alone is sufficient to the problem of reliably detecting plastic A/P mines with sufficiently low false alarm rates. A visualization of the basic GPR sensing concept appears in Figure 1(a) and a photo of one kind of EMIS sensor appears in Figure 1(b).

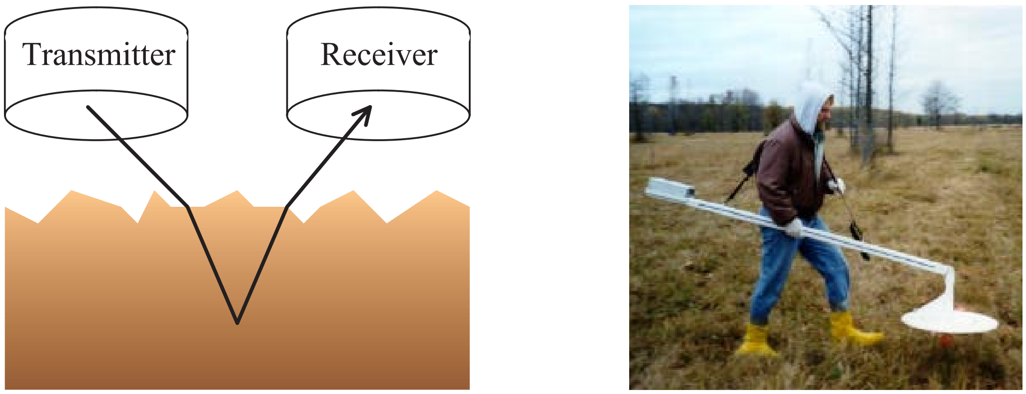

(a) GPR visualization (b) EMIS photo

Figure 1. Ground penetrating radar and electromagnetic induction spectroscopy sensors

To date, GPRs have been the favored electromagnetic technology for detecting plastic A/P mines.[3] GPRs operate by emitting wideband electromagnetic pulses, which propagate into the earth and reflect from subsurface structures. The reflections are then measured and, in the case of arrays, fused. GPRs are good for detecting shallow dielectric cavities, of which plastic A/P mines are examples. Unfortunately, plastic A/P mines are not the only examples of dielectric cavities, with rocks representing a significant potential confusing class of objects. This potential confusion is especially true since random reflections from a rough air/ground interface can swamp the mine or discrete clutter signatures, making distinguishing between these two classes difficult. Similarly, soil inhomogeneity, such as soil compaction or moisture content, can lead to numerous false alarms. Also, very small metallic objects, such as the firing pin found in many plastic A/P mines, are almost invisible to GPR, which would remove a significant potential characteristic for discerning plastic landmines if GPR were used alone.

An EMIS sensor can complement a GPR by detecting and recognizing the small metal content of many plastic A/P mines. By emitting sinusoidal signals and sweeping across multiple frequencies, an EMIS sensor can obtain information that depends on metal material properties and spatial structure, even for very small amounts of metal.[2] Also, EMIS sensors are less likely to be affected by soil variations, unlike GPR. However, at gains high enough to see the small metal content of many plastic A/P mines, such an inductive sensor is also highly sensitive to, and frequently confused by, discrete metallic clutter such as small shell fragments, pop-tops from soda cans, or many other kinds of metallic clutter. Additionally, EMIS sensors cannot distinguish landmines with zero metallic content. Thus, an EMIS sensor alone also cannot successfully discriminate plastic A/P mines. But many potential false alarms in the GPR domain are not the same as those in the EMIS domain, and vice versa, i.e. we expect that in most cases, dielectric cavities and metallic objects that coincide in space are likely to be plastic A/P mines and not random clutter.

3. METHOD

As stated in Section 1, our goal is to develop a framework for testing a wide range of possible sensor and fusion scenarios for landmine detection, far more cases than is possible with currently available landmine data sets. By providing an environment in which we can control every aspect of the problem, from rough soil surfaces and soil inhomogeneity to variations in mine pose and sensor geometry to different clutter environments, we can conduct fundamental tradeoff studies for sensor fusion. With this approach, we will be able to find optimal solutions that account for the wide range of possible situations that can arise in real, practical problems. In this paper, we present an initial study that uses this approach to examine array GPR and EMIS data fusion.

3.1. Sensors

Our detection system consists of two sensors, one a GPR array and the other an electromagnetic induction spectroscopy (EMIS) sensor, and a set of detection and fusion algorithms. Figure 2 shows the basic sensor and landmine geometry.

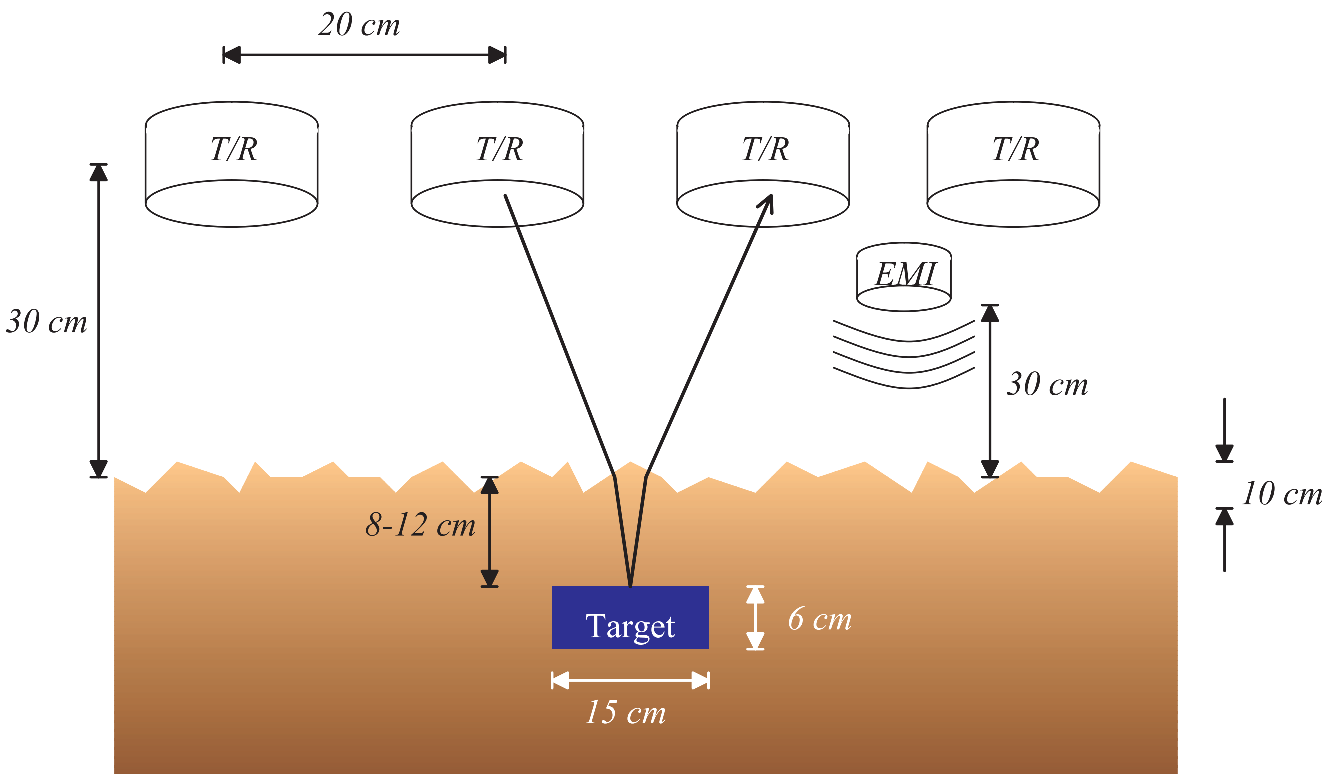

Figure 2. Geometry of the sensor/mine configuration

Our GPR array uses four antennae with transmit/receive capability. We locate these sensors approximately 30cm from the nominal soil surface, corresponding to a vehicular mounting arrangement. The antennae are spaced at 20cm intervals. We choose to use a short, wideband pulse. Our pulse width is about 0.8ns, corresponding to a 1.25GHz band width. The observed signal from the GPR consists of 16 time traces, one for each transmitter/receiver pair. It is important to note here that we are not explicitly focusing the array. We simply transmit a pulse from a single antenna while receiving at all four antennae, and then transmit from the next antenna while again receiving with all four, and so on. This approach maximizes the amount of information available to the detection algorithm. We show a noiseless example set of time traces for one transmit/receive pair in Figure 3. The two overlaid traces correspond to the cases with and without a mine in dispersive soil for a single example of a rough ground surface.

Our EMIS sensor is a single square induction coil located approximately 30cm above the nominal soil surface. The frequency is swept over a 30Hz-20kHz range, which we have sampled logarithmically. Multiple such sensors could be place in an array, but here we have only considered the single coil case. Our observed data is a vector of samples of the observed EMF signal, logarithmically spaced over the frequency range.

3.2. Physical Models

To generate our data, we use accurate electromagnetic software tools. In the case of the GPR array, we use a Finite Difference Time-Domain (FDTD) algorithm.[4] With FDTD, we can generate data corresponding to a variety of

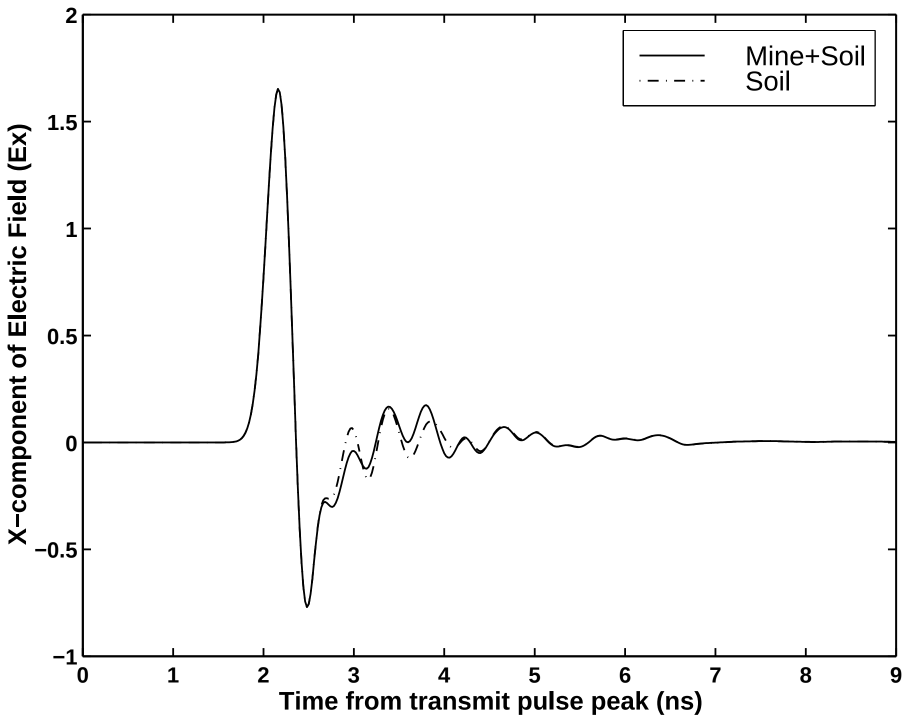

Figure 3. Noiseless GPR time traces; transmitter 2, receiver 3

situations that we can control explicitly: varying rough ground surfaces, different soil parameters, and alternative sensing geometries, for cases with plastic A/P landmines or discrete clutter such as rocks or no subsurface objects at all. For the EMIS sensor, we use a quasi-static magnetic moment physical model, for which we can control sensing geometry and the properties of mine or clutter.[5] For this study, we choose to assume that our sensors are centered over the section of ground we are testing, and that we are merely trying to test whether or not a landmine is located directly below the sensors at a known depth.

For the GPR, we include a random, but correlated, air/ground interface as a source of uncertainty, as well as white, Gaussian sensor noise of about 20dB per channel signal to noise ratio. In our physical GPR model, we construct three cases. In the first, we include a plastic landmine of known shape, using a relative permittivity of 2.9, corresponding to TNT. In the second, we include a discrete clutter object, with a relative permittivity of 2.8, that is approximately an ellipsoid of roughly the same size as the landmine. In the last case, we modeled no discrete object at all, measuring the response from the empty, homogeneous soil. The soil's relative permittivity is 2.6, and we studied both dispersive, lossy and non-dispersive, lossless cases.

For the EMIS sensor, we model the plastic landmine as having a small, steel firing pin centered in the mine, approximately 1cm in size, with a relative permeability of 5000 and a conductivity of 10^7 S/m. We assume the earth is essentially invisible to the sensor at such low frequencies. We also model a discrete metal fragment as approximately 4cm in size, with the same material properties as the firing pin. For sensor noise, we include white, Gaussian noise for 14dB signal to noise ratio. Note here that, at this noise level, the resulting clutter object looks indistinguishable from the firing pin. In addition to the mine and clutter situations, we include empty soil by simply using the sensor noise with zero mean.

In modeling this system, we only include single objects within the field of view. Thus, the landmine and discrete clutter cases are mutually exclusive. To stress this physical model of clutter versus target, we include the chart of Figure 4, showing how the GPR and EMI clutter characteristics relate. We also never vary the pose of the objects; we assume they lie flat. We will explore more complicated models in future work.

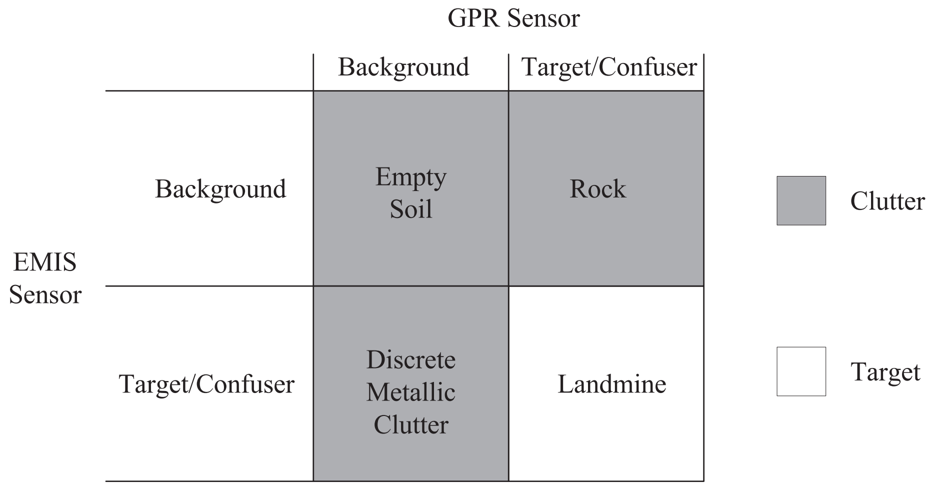

Figure 4. Clutter versus target domains

3.3. Statistical Models

Our statistical models are related to, but quite distinct from, our physical models. Our goals in doing statistical modeling are threefold. First, by allowing statistical error into the problem, we can make our approach robust to modeling error. Thus, the statistics can absorb any differences between our physical model and reality, particularly difficult-to-manage nonlinearities. Second, statistical approaches usually can provide measures of error, so that we can estimate how well we are doing. Third, by using a statistical approach, we open a wealth of available techniques and knowledge for sensor fusion and detection.

For the GPR, our statistical model assumes that the data comes from three additive vector sources: the rough ground return, $\underline{g}$; white Gaussian sensor noise, $\underline{w}$; and possibly a buried landmine, $\underline{s}$. We formulate our two detection hypotheses by the presence or absence of the landmine signal, $\underline{s}$:

$$\begin{array}{rr} H_{r,0}: & \underline{y}_r = \underline{g} + \underline{w} \\ H_{r,1}: & \underline{y}_r = \underline{s} + \underline{g} + \underline{w} \end{array} \tag{1}$$

where $\underline{s}$, $\underline{g}$, and $\underline{w}$ are simply the time signals corresponding to the mine signal, the ground return, and the white sensor noise, fused at the data level by vectorization:

$$\underline{s} = [s_{j,k}(n)] \quad \underline{g} = [g_{j,k}(n)] \quad \underline{w} = [w_{j,k}(n)] \quad \left\{ \begin{array}{l} j,k = 1,2,3,4 \\ n = 1,\ldots,750 \end{array} \right. \tag{2}$$

where j and k denote the specific transmit/receive antenna pair, and n is the time sample. The sum of these signals, $\underline{y}_r$, is what we can observe as output from the antenna array. We assume that these signals are independent Gaussian random vectors, distributed as follows:

$$\underline{s} \sim N(\underline{\mu}_s, \Sigma_s) \quad \underline{g} \sim N(\underline{\mu}_g, \Sigma_g) \quad \underline{w} \sim N(\underline{0}, \Sigma_w) \tag{3}$$

where $\underline{\mu}_s$ and $\underline{\mu}_g$, and Σ_s and Σ_g, are the means and covariances, respectively, of their distributions, and the elements of w are zero-mean and identically distributed, so $\Sigma_w = \sigma_w^2 I$. For computational tractability, we currently assume Σ_s and Σ_g are diagonal, though in future work we may introduce tractable correlation structure.

For EMIS, we have a single sensor, with data from a sweep over frequency, and this data is corrupted by Gaussian noise. Going immediately to vector notation, we assume that the signal $\underline{c}$ is a known deterministic signal and that $\underline{u}$ is independent and identically distributed, zero-mean, Gaussian noise, so our two cases are simply:

$$\begin{array}{rr} H_{e,0}: & \underline{y}_e = \underline{u} \\ H_{e,1}: & \underline{y}_e = \underline{c} + \underline{u} \end{array} \tag{4}$$

where

$$\underline{u} \sim N(0, \Sigma_u) \quad \Sigma_u = \sigma_u^2 I \tag{5}$$

It is important to note that in this preliminary study, we elected not to explicitly model the discrete clutter in our statistical detection problem, but we do include such clutter in the physical model for generating the data. Nowhere in any of the statistical models do we include a discrete clutter case, i.e. a rock or metallic clutter. We did this because we wanted to explore the problem of model mismatch and the advantages of detection-level fusion when discrete clutter models are unknown or ignored.

3.4. Detectors

Given the statistical models above, it is relatively easy to formulate an "optimal" maximum likelihood detector, based on a likelihood ratio test. In the GPR case, we assume the three Gaussian signals are independent, so we can simply add their means and covariances

$$\begin{array}{ll} H_{r,0}: & \underline{y}_r \sim N(\underline{\mu}_g, \Sigma_g + \Sigma_w) = N(\underline{\mu}_0, \Sigma_0) \\ H_{r,1}: & \underline{y}_r \sim N(\underline{\mu}_s + \underline{\mu}_g, \Sigma_s + \Sigma_g + \Sigma_w) = N(\underline{\mu}_1, \Sigma_1) \end{array} \tag{6}$$

The detector resulting from this model and a likelihood ratio test is the full quadratic classifier:

$$\left(\underline{y}_r - \underline{\mu}_0\right)^T \Sigma_0^{-1} \left(\underline{y}_r - \underline{\mu}_0\right) - \left(\underline{y}_r - \underline{\mu}_1\right)^T \Sigma_1^{-1} \left(\underline{y}_r - \underline{\mu}_1\right) \underset{H_0}{\overset{H_1}{\gtrless}} \gamma_r \tag{7}$$

We estimate the parameters from a large number of Monte Carlo runs by estimating $\underline{\mu}_g$ and Σ_g for the case with empty soil, and $\underline{\mu}_s + \underline{\mu}_g$ and $\Sigma_s + \Sigma_g$ for the landmine case. These parameters are averaged over many different, random rough soil/air interfaces. Because we control the sensor noise, we can add the known white noise covariance to each case ourselves.

Similarly, we develop the maximum likelihood detector for the EMIS case. However, because the signal is a known deterministic signal and we control the noise, we can simplify both the equations and the detector:

$$\begin{array}{ll} H_{e,0}: & \underline{y}_e \sim N(0, \Sigma_u) \\ H_{e,1}: & \underline{y}_e \sim N(\underline{c}, \Sigma_u) \end{array} \tag{8}$$

$$\underline{c}^T \Sigma_u^{-1} y_e \underset{H_0}{\overset{H_1}{\gtrless}} \gamma_e \tag{9}$$

While this approach is adequate for our current simulations, in practice, our detector and noise parameters will need to be estimated from data.

From these models, we perform three detection experiments to demonstrate the advantages of fusion. First we test the GPR alone, then we test the EMIS sensor alone. Lastly, we fuse the two sensing modalities with detection-level fusion, using an AND detector. To build an appropriate optimal detector, we need to explicitly model the confusing clutter classes statistically. The advantage of the AND detector is that it can approximate the optimal multi-class detector without explicit clutter models. While it is always better to obtain the optimal detector by explicitly modeling clutter, in cases where this is infeasible, an AND detection rule will win. In future work, we intend to formulate an optimal detector by determining and using appropriate multi-class clutter models.

4. RESULTS

We conducted three major experiments, over which we varied soil parameters and mine depth. In each of our experiments, we ran 400 Monte Carlo runs with different rough air/soil interface and sensor noise realizations. In 100 cases, we placed a landmine. In another 100 cases we placed the GPR clutter object (the "rock"). In the next 100 cases, we used the EMIS clutter object (the small metal fragment). In the last 100 cases, we included no discrete clutter or mine at all, only empty soil. In each of the three experiments, we used each of the three detection methods: GPR alone, EMIS alone, and their detection-level fusion.

For each of the 400 Monte Carlo runs, we generated the random soil surface using a correlated Gaussian random model with approximately ±5cm of deviation from nominal, maximum, and an 8cm correlation length. In every case, we set the soil's relative permittivity to 2.63, but in two cases we used a lossless, non-dispersive soil model, while in the third, we implemented a lossy, dispersive model based on data from the Bosnia Steele Castle site. In this case, the moisture content was 4.7% and the density was 1.181 g/cm^3, and we modeled the loss and dispersion using a 2-1 Padé approximation.

In one of the two lossless, non-dispersive cases, we placed a landmine with its upper surface at a depth of 8cm below the nominal soil surface, while in the other we placed it 12 cm below the nominal soil surface. In the lossy, dispersive case we placed the landmine 8cm below the surface. Our landmine was 6cm in thickness, with a diameter of 15cm. Except for the firing pin, we assumed the mine was filled with TNT, with a relative permittivity of 2.9, relative permeability 1, and conductivity 0. We also assumed that the mine contained a small, 1cm firing pin with a relative permeability of 5000 and a conductivity if 10^7S/m.

For our clutter, we assumed that the rock was roughly the same size and shape as the mine, but approximately ellipsoidal with major axis 17cm and minor axis 7cm, oriented similarly to the landmine. We assumed the rock had a relative permittivity of 2.8, with relative permeability of 1 and no conductivity. We assumed that the metal fragment clutter about 4cm in size with the same material properties as the landmine firing pin.

Figure 5 shows the ROC curves for the case of a mine at a depth of 8cm in lossless, non-dispersive soil. First, note that the ROC curve for the EMIS sensor alone shows good correspondence to what we expect for the case that the metal clutter object is nearly indistinguishable from the landmine. The straight line ROC is indicative that guessing is as good as any other technique when trying to distinguish the metal clutter from the firing pin when using the EMIS sensor. The GPR is better at distinguishing between the rock and the landmine than the EMIS sensor is for the metal clutter, but it still has a significant number of false alarms for $P_d \approx 0.99$. However, the detection-level fusion scheme has a much lower false alarm rate for comparable P_d.

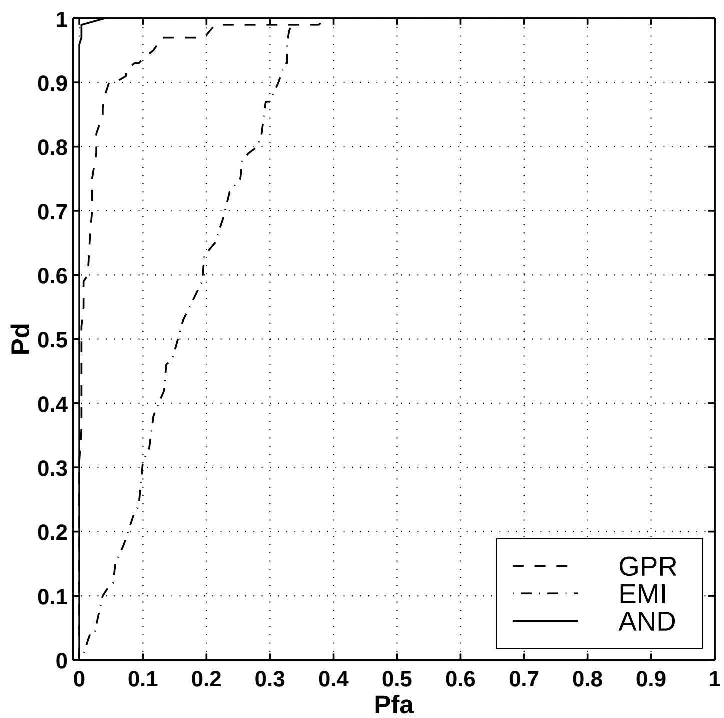

Figure 5. ROC curves for mine at 8cm depth, non-dispersive soil

In Figure 6, we see the ROC curves for the mine at a depth of 12cm, with the same soil parameters. As we expect, the increased distance makes the EMIS sensor perform slightly worse than before. The GPR alone, however, gets significantly improved performance. This improvement comes from the time delay of the mine signal. While the mine signal is attenuated due to the increased distance, the distance attenuation of the reflection from the rough ground surface is much stronger, resulting in a much smaller interfering signal relative to the mine signal at this time delay. This improvement would be mitigated somewhat by lossy soil. Also, as expected, the fused detector performs better than either sensor alone, despite its simplicity.

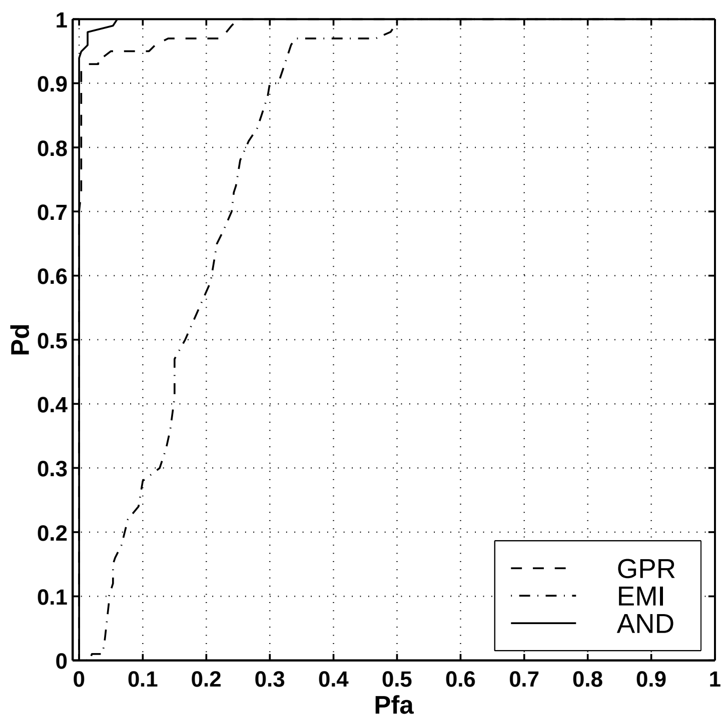

Figure 6. ROC curves for mine at 12cm depth, non-dispersive soil

Lastly, in Figure 7, we show the ROCs for the case of a mine buried at a depth of 8cm in lossy, dispersive soil. The EMIS ROC is essentially the same as the first case, but the GPR curve is much different. Here we begin to see problems with model mismatch. Because we have not explicitly modeled clutter statistically, we have found a situation where the GPR clutter signal is much stronger than the mine signal. Thus, despite the differences in the two signals, the clutter signal is detected as a mine more often than the mine is. This poor GPR detector performance suggests that we can do much better if we model the clutter explicitly in our statistics. The AND detection rule again does quite well; in fact it appears to do perfectly. This is an artifact of ROC estimation from discrete, experimental data and the small number of Monte Carlo runs. More Monte Carlo samples, or another approach, such as importance sampling, would help rectify this issue. With more samples, our ROC estimates and confidence in them would improve.

5. CONCLUSIONS

We have developed a framework for studying the system design tradeoffs in sensor fusion for landmine detection applications. We can now conduct controlled experiments that examine the effects that soil or landmine parameters, sensor geometries, clutter, and fusion algorithms have on the landmine detection problem. With this approach, we aim to identify the optimal sensor configurations and fusion algorithms for general, real-world landmine problems.

We have used this framework to study the effect of variations in soil parameters and mine depth on GPR and EMIS data fusion, particularly for the case of plastic A/P mines. We also focused on the issue of model mismatch and the advantages of accurate clutter modeling for detection. We found that even simple detection-level sensor fusion techniques can provide a large improvement in our ability to detect plastic A/P mines.

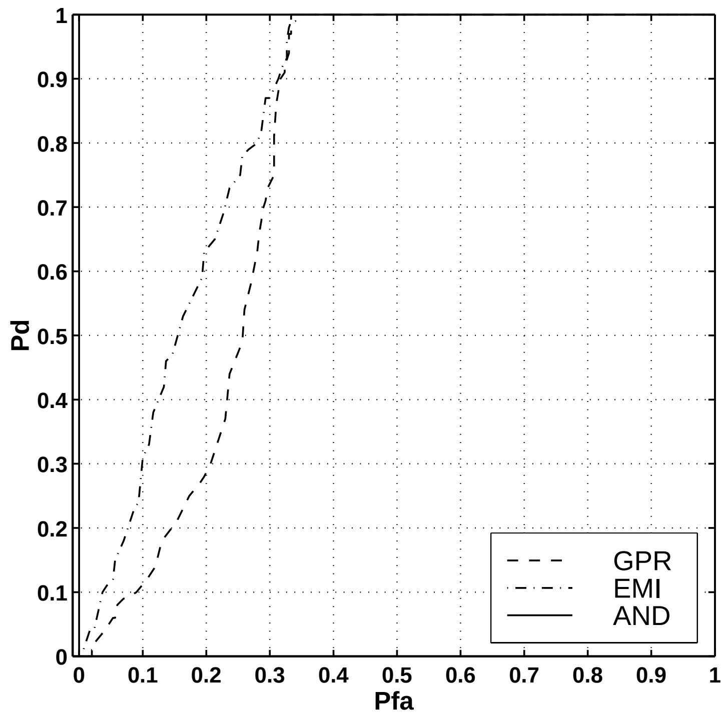

Figure 7. ROC curves for mine at 8cm depth, dispersive soil

6. ACKNOWLEDGMENTS

This work is partially supported OSD MURI program, the Air Force Office of Scientific Research under Grant F49620-96-1-0028, and the National Institutes of Health under Grant NINDS 1 R01 NS34189. We would also like to thank Scott Winton an Mustafa Özdemir of Northeastern University for their assistance with the electromagnetic simulations.

REFERENCES

1. C.M. Rappaport, D.M. Reidy, "Focused array radar for real-time imaging and detection", *Radar Sensor Technology*, Gerald S. Ustach, editor, Proc. SPIE V 2747, p202-213, SPIE, Orlando (1996)
2. I.J. Won, D. Keiswetter, "Electromagnetic Induction Spectroscopy", *Detection and Remediation Technologies for Mines and Mine-Like Targets III*, A.C. Dubey, J.F. Harvey, J. Broach, editors, Proc. SPIE V 3392, p14-22, SPIE, Orlando, (1998)
3. T.R. Witten, "Present state-of-the-art in ground penetrating radars for mine detection", *Detection and Remediation Technologies for Mines and Mine-Like Targets III*, A.C. Dubey, J.F. Harvey, J. Broach, editors, Proc. SPIE V 3392, p576-585, SPIE, Orlando, (1998)
4. G.A Tsihrintzis, C.M. Rappaport, S.C. Winton, P.M. Johansen, "Statistical modeling of rough surface scattering for ground-penetrating radar applications", *Detection and Remediation Technologies for Mines and Mine-Like Targets III*, A.C. Dubey, J.F. Harvey, J. Broach, editors, Proc. SPIE V 3392, p735-744, SPIE, Orlando, (1998)
5. M. Özdemir, E. Miller, S. Norton, "Localization and characterization of buried objects from multi-frequency, array inductive data", *Detection and Remediation Technologies for Mines and Mine-Like Targets IV*, A.C. Dubey, J.F. Harvey, J. Broach, R.E. Dugan, editors, Proc. SPIE V 3710, SPIE, Orlando, (1999)