

Comparison of Pre-Detection and Post-Detection Fusion for Mine Detection

A. Gunatilaka and B. A. Baertlein

The Ohio State University ElectroScience Laboratory

1320 Kinnear Road, Columbus, OH 43212

ABSTRACT

We present and compare methods for pre-detection (feature-level) and post-detection (decision-level) fusion of multi-sensor data. This study emphasises methods suitable for data that are non-commensurate and sampled at non-coincident points. Decision-level fusion is most convenient for such data, but this approach is sub-optimal in principle, since targets not detected by all sensors will not achieve the maximum benefits of fusion. A novel feature-level fusion algorithm for these conditions is described. The optimal forms of both decision-level and feature-level fusion are described, and some approximations are reviewed. Preliminary results for these two fusion techniques are presented for experimental data acquired by a metal detector, a ground-penetrating radar, and an infrared camera.

Keywords: land mines, sensor fusion, infrared, ground penetrating radar, metal detectors

1. INTRODUCTION

Fusion of multi-sensor data for mine detection is complicated by several factors. The sensors used are diverse and produce non-commensurate data. The need for essentially perfect detection leads to high false alarm rates, and the clutter-rich environment in which these sensors operate leads to additional false alarms. Finally, the sensors may reside on different platforms, leading to problems with data registration.

One approach to these problems is to perform detection at the sensor level and to then combine the detections from individual sensors. This process, also known as “decision-level” fusion, is convenient since it reduces diverse sensor data to a common format (binary decisions or detection probabilities) that are readily combined using a variety of techniques developed for this purpose.^{1,2} In principle, decision level fusion is suboptimal, since if a target is not detected by more than one sensor, it will not experience the benefits of fusion.

In this paper we discuss a feature-level fusion approach to non-coincident sensor sampling. The data are represented by a model with unknown parameters (features) and random additive clutter. Optimization techniques are used to determine the features from the available data, and classification is performed on the basis of the features. A technique for dealing with position uncertainty is also described. This approach is capable of detecting targets even when the features derived from any single sensor are insufficient for detection.

The work is organized in four major sections. In Section 2 we describe the theoretical basis for feature-level fusion. The basis for decision-level fusion is presented in Section 3. Descriptions of our sensors and techniques for sensor data processing and feature estimation are described in Section 4. Experimental data and examples of fusion are given in Section 5. Finally, concluding remarks appear in Section 6.

2. FEATURE-LEVEL FUSION OF RANDOMLY SAMPLED MULTI-SENSOR DATA

2.1. Problem Definition

Suppose that N_S countermine sensors are used to acquire data in a region. Sensor i acquires J_i data samples at locations \mathbf{R}_{ij} which we denote by the row vector $\mathbf{d}_i(\mathbf{R}_{ij})$. Detection and/or classification is to be performed using the data set

$$\mathcal{D} = \{\mathbf{d}_i(\mathbf{R}_{ij}), i = 1, \dots, N_S; j = 1, \dots, J_i\} \quad (1)$$

We assume that different sensors produce samples that differ in number J_i , sample positions \mathbf{R}_{ij} , and the format and dimensionality of their data (i.e., they are non-commensurate).

Send correspondence to B.A.B. (614) 292-0076 (voice), (614) 292-7297 (fax), baertlein.1@osu.edu

Let \mathbf{R} be any point at which we wish to determine the presence (or identify) of a mine. We refer to \mathbf{R} as the “interrogation point.” We form the K hypotheses $H_k(\mathbf{R})$, $k = 1, 2, \dots, K$ regarding the presence or absence of various types of mines at \mathbf{R} . We may also include in the set $\{H_k\}$ discrete clutter objects, e.g., buried rocks, metallic clutter, etc., that might be false alarms.

At each point \mathbf{R} we wish to determine the hypothesis $H_k(\mathbf{R})$ that minimizes the Bayes risk conditioned on the data \mathcal{D}

$$k = \arg \min_j \sum_{\ell=1}^K C_{j\ell} \Pr(H_\ell(\mathbf{R})|\mathcal{D}) \quad (2)$$

where $C_{j\ell}$ is the cost of deciding hypothesis H_j when H_ℓ is true. In the special case of $K = 2$ (binary detection) or when all the cost functions are equal, it is sufficient to compute likelihood ratios and to compare them to a common threshold value. In this work, however, we will use the more general criterion noted above, since multiple hypotheses may be involved and equal costs are inappropriate for the mine-detection problem.

We now suppose that data acquired by sensor i at a point \mathbf{R}_{ij} over a target of type H_k at \mathbf{R} can be modeled as a signal $\mathbf{G}_i(\mathbf{R}_{ij}; \Theta_{ik})$ and additive noise $\mathbf{N}_i(\mathbf{R}_{ij})$

$$\mathbf{d}_i(\mathbf{R}_{ij})|H_k(\mathbf{R}) = \mathbf{G}_i(\mathbf{R}_{ij}; \Theta_{ik}) + \mathbf{N}_i(\mathbf{R}_{ij}) \quad (3)$$

where Θ_{ik} is a (column) vector parameter that describes the size, shape, amplitude, and position of targets of type k for sensor i . Since the discrete clutter objects have been assigned hypotheses H_k , the phenomena modeled by $\mathbf{N}_i(\mathbf{R})$ consist of instrument noise and large-scale clutter. We assume that the clutter $\mathbf{N}(\mathbf{R})$ is independent of the parameter Θ .

Since for each sensor a mine signature has a finite spatial extent, only a portion of the samples in \mathcal{D} are relevant to interrogation in the region around \mathbf{R} . Let the relevant data be given by subsequences j_1, j_2, \dots, j_{M_i} of each sensor data set where M_i is the number of relevant samples for sensor i . It is convenient to assemble this reduced data set into the column vectors

$$\mathbf{D} = \left[\mathbf{d}_1(\mathbf{R}_{1j_1}) \cdots \mathbf{d}_1(\mathbf{R}_{1j_{M_1}}) \mathbf{d}_2(\mathbf{R}_{2j_1}) \cdots \mathbf{d}_{N_s}(\mathbf{R}_{N_s j_{M_{N_s}}}) \right]^T \quad (4)$$

$$\mathbf{G}_k = \left[\mathbf{G}_{1k}(\mathbf{R}_{1j_1}; \Theta_{1k}) \cdots \mathbf{G}_{1k}(\mathbf{R}_{1j_{M_1}}; \Theta_{1k}) \mathbf{G}_{2k}(\mathbf{R}_{2j_1}; \Theta_{2k}) \cdots \mathbf{G}_{N_s k}(\mathbf{R}_{N_s j_{M_{N_s}}}; \Theta_{N_s k}) \right]^T \quad (5)$$

$$\mathbf{N} = \left[\mathbf{N}_1(\mathbf{R}_{1j_1}) \cdots \mathbf{N}_1(\mathbf{R}_{1j_{M_1}}) \mathbf{N}_2(\mathbf{R}_{2j_1}) \cdots \cdots \mathbf{N}_{N_s}(\mathbf{R}_{N_s j_{M_{N_s}}}) \right]^T \quad (6)$$

If $\mathbf{d}_i(\mathbf{R})$ has length N_{Di} , then the length of these vectors is $N_D = \sum_i N_{Di} M_i$. Using this notation, we have for the vector signal model

$$\mathbf{D}|H_k(\mathbf{R}) = \mathbf{G}_k(\Theta_k) + \mathbf{N} \quad (7)$$

To support decision-level fusion and feature-level fusion with uncertain sample locations we can include in Θ_{ik} a position offset \mathbf{R}_0 that describes the nominal “center” of the mine signature with respect to \mathbf{R} . We write

$$\Theta_{ik} = [\theta_{ik} \ \mathbf{R}_{0,ik}]^T \quad (8)$$

In general, we will take θ and \mathbf{R}_0 to be independent in what follows.

2.2. General Formulation

Feature-level fusion can be formulated in two ways. First consider the direct approach using the a posteriori densities. To minimize the risk given by equation (2) we are led to consider

$$\Pr(H_k(\mathbf{R})|\mathbf{D}) = \int d\Theta \Pr(H_k(\mathbf{R})|\Theta) f_{\Theta|\mathbf{D}}(\Theta) = \int d\theta \int d\mathbf{R}_0 \Pr(H_k(\mathbf{R})|\theta, \mathbf{R}_0) f_{\Theta|\mathbf{D}}(\Theta) \quad (9)$$

The expression in equation (9) is exact, but it requires the densities $\Pr(H_k|\Theta)$ (a classifier), $f_{\Theta|\mathbf{D}}(\Theta)$ (related to an estimator of $\Theta|\mathbf{D}$) and an integral over a feature space of possibly high dimensions. In many applications it is

attractive to employ approximations to this result. If the data \mathbf{D} strongly imply the value $\hat{\Theta}$, then $f_{\Theta|\mathbf{D}}$ will be strongly peaked about $\hat{\Theta}$ and we have a maximum a posteriori approximation

$$\Pr(H_k(\mathbf{R})|\mathbf{D}) \approx \Pr(H_k(\mathbf{R})|\hat{\Theta}) \quad (10)$$

Conversely, if \mathbf{D} provides no information about Θ then $f_{\Theta|\mathbf{D}}(\Theta) = f_{\Theta}(\Theta)$ and we recover the a priori probabilities.

A further approximation will permit us to separate the effects of features and position offsets. If we take the features θ and the position \mathbf{R}_0 to be independent when conditioned on $H_k(\mathbf{R})$, then we obtain

$$\Pr(H_k(\mathbf{R})|\hat{\Theta}) = \Pr(H_k(\mathbf{R})|\hat{\theta}) \frac{f_{\mathbf{R}_0|H_k(\mathbf{R})}(\mathbf{R}_0)}{f_{\mathbf{R}_0}(\mathbf{R}_0)} \quad (11)$$

In the absence of any a priori information about mine position, we take $f_{\mathbf{R}_0}$ to be uniform over the region of interest.

Feature-level fusion can also be formulated using the class-conditioned densities. This approach has advantages when models for such densities exist. From Bayes' rule we can write

$$k = \arg \min_j \sum_{\ell=1}^K C_{j\ell} \Pr(H_\ell(\mathbf{R})) f_{\mathbf{D}|H_\ell(\mathbf{R})}(\mathbf{D}|H_\ell(\mathbf{R})) \quad (12)$$

Note that this formulation requires knowledge of the priors $\Pr(H_k(\mathbf{R}))$, which are sometimes difficult to obtain. Again using Bayes' rule, we find

$$f_{\mathbf{D}|H_k(\mathbf{R})}(\mathbf{D}) = \int d\Theta f_{\Theta|H_k(\mathbf{R})}(\Theta) f_{\mathbf{D}|\Theta, H_k(\mathbf{R})}(\mathbf{D}) \quad (13)$$

Since \mathbf{G} is presumed known, the data \mathbf{D} conditioned on the parameter Θ have the same density as the clutter \mathbf{N} , which leads to

$$f_{\mathbf{D}|\Theta, H_k(\mathbf{R})}(\mathbf{D}) = f_{\mathbf{N}}(\mathbf{D} - \mathbf{G}_k(\Theta)) \quad (14)$$

and

$$f_{\mathbf{D}|H_k(\mathbf{R})}(\mathbf{D}) = \int d\Theta f_{\Theta|H_k(\mathbf{R})}(\Theta) f_{\mathbf{N}}(\mathbf{D} - \mathbf{G}_k(\Theta)) \quad (15)$$

Once again we have an exact result, but the cost of evaluating equation (15) can be high. As done above, we obtain an approximate expression by assuming that $f_{\mathbf{N}}(\mathbf{D} - \mathbf{G}_k(\Theta))$ has a well defined maximum for $\Theta = \hat{\Theta}$. We find

$$f_{\mathbf{D}|H_k(\mathbf{R})}(\mathbf{D}) \approx f_{\Theta|H_k(\mathbf{R})}(\hat{\Theta}) \int d\Theta f_{\mathbf{N}}(\mathbf{D} - \mathbf{G}_k(\Theta)) = (const) f_{\Theta|H_k(\mathbf{R})}(\hat{\Theta}) = (const) f_{\theta|H_k(\mathbf{R})}(\hat{\theta}) f_{\mathbf{R}_0|H_k(\mathbf{R})}(\hat{\mathbf{R}}_0) \quad (16)$$

where the constant term depends only on the data \mathbf{D} and in the last step we have made the assumption that θ and \mathbf{R}_0 are conditionally independent.

3. DECISION-LEVEL FUSION

The formulation of an optimal decision-level fusion algorithm is similar to that described in Section 2 for feature-level fusion. As before, each sensor i collects data $\{\mathbf{d}_i(\mathbf{R}_{i1}), \dots, \mathbf{d}_i(\mathbf{R}_{iJ_i})\}$. On the basis of these data, it makes a declaration u_i regarding the truth of intermediate hypotheses $h_{iq}(\mathbf{R})$ for $q = 1, 2, \dots, Q_i$. The intermediate hypotheses may be selected on the basis of the sensors involved. For example, for a sensor suite that comprises EMI, GPR and IR we might use h_{11} = "metallic object present," h_{12} = "metallic object not present," h_{21} = "dielectric discontinuity present," h_{22} = "dielectric discontinuity not present," h_{31} = "thermal discontinuity present," and h_{32} = "thermal discontinuity not present," rather than simply "mine present" and "mine not present." To simplify the discussion we will only treat the case in which all the Q_i s are equal. The single-sensor detections u_i can be obtained by any process, including that described in Section 2, if in that discussion the data set is restricted to one sensor only. To simplify the notation we form the vector of local declarations given by

$$\mathbf{u} = [u_1 \ u_2 \ \dots \ u_{N_s}]^T \quad (17)$$

We wish to combine these declarations into a fused declaration u_0 which describes the truth of hypotheses $H_k(\mathbf{R})$, $k = 1, 2, \dots, K$. We will assume the following: (1) the fusion processor has no knowledge of the sensor data \mathbf{D} and u_0 is based only on the declarations \mathbf{u} , (2) the association problem has been solved so that a putative detection $H_k(\mathbf{R})$ is not confused with another detection $H_k(\mathbf{R}')$, (3) for each sensor we know the detection confidences $\Pr(u_i|H_k(\mathbf{R}))$ for all k , and (4) the a priori probabilities $\Pr(H_k(\mathbf{R}))$ are known. For optimal fusion we must minimize the risk of the decision, which for a fixed value of \mathbf{u} is equivalent to minimizing

$$\arg \min_j \sum_{\ell=1}^K C_{j\ell} \Pr(H_\ell|\mathbf{u}) = \arg \min_j \sum_{\ell=1}^K C_{j\ell} \Pr(\mathbf{u}|H_\ell) \Pr(H_\ell) \quad (18)$$

There are several approaches to decision fusion. The relation between u_0 and the individual declarations u_i , which is expressed by $\Pr(u_0 = j|\mathbf{u})$, is a mapping from the N_s -fold product of the space $\{1, 2, \dots, Q\}$ to one of K output hypothesis values. There are $K^{Q^{N_s}}$ such functions, from which we must determine the most effective rule. Even for the case of binary decisions with $K = Q = 2$, this leads to a large number of possible mappings. Several ad hoc mappings from the u_i to u_0 have been used (some with considerable success) in performing fusion. These include logical AND and OR operations, as well as majority voting.

In some situations we can determine the optimal mapping.¹ Here we consider the simplest case, in which the local decision rules (i.e., the mappings from \mathbf{D} to the local decisions u_i) are fixed. For $K = Q = 2$ the minimum in equation (18) reduces to a likelihood ratio test, given by

$$\frac{\Pr(u_1, u_2, \dots, u_{N_s}|H_1)}{\Pr(u_1, u_2, \dots, u_{N_s}|H_0)} > \eta \quad (19)$$

where η is a constant threshold and the cases $>$ and $<$ correspond to the decisions $u_0 = 1$ and $u_0 = 0$ respectively.

If the declarations u_i are conditionally independent for each sensor, then we have

$$\frac{\Pr(u_1, u_2, \dots, u_{N_s}|H_1)}{\Pr(u_1, u_2, \dots, u_{N_s}|H_0)} = \prod_{i=1}^{N_s} \frac{\Pr(u_i|H_1)}{\Pr(u_i|H_0)} \quad (20)$$

If we define miss and false-alarm probabilities for each sensor as follows:

$$P_{Mi} = \Pr(u_i = 0|H_1) \quad (21)$$

$$P_{Fi} = \Pr(u_i = 1|H_0) \quad (22)$$

then we can write the optimum test as

$$\sum_{i=1}^{N_s} \left[u_i \log \left(\frac{1 - P_{Mi}}{P_{Fi}} \right) + (1 - u_i) \log \left(\frac{P_{Mi}}{1 - P_{Fi}} \right) \right] \geq \log(\eta) \quad (23)$$

which is a weighted form of voting.

4. SENSORS AND SIGNAL PROCESSING

The foregoing algorithms are being used with a sensor suite that comprises a ground penetrating radar (GPR), a commercial infrared (IR) camera, and a Schiebel electromagnetic induction (EMI) sensor. Brief descriptions of these devices and their supporting signal preprocessing algorithms are presented below.

4.1. Electromagnetic Induction Sensor

4.1.1. Hardware

The EMI sensor used in this work was a Schiebel AN-19/2 pulsed-induction metal detector. The standard output signal, an audio tone, is not suitable for the modeling work described here. Instead, data were acquired with a digital oscilloscope at an internal signal within the sensor. The sensor was fitted to a linear scanner, which was programmed to acquire 56 samples over a 55 inch path. Scans were performed at six inch intervals along a track to obtain a grid of sample points over the region of interest.



Figure 1. The GPR dielectric rod antenna and scanning system.

4.1.2. Signal Processing

Signal processing for the EMI sensor is a two-step process. First, the sensor response over metal-free soil was acquired, and this background was subtracted from the waveforms collected over other locations. In principle, the response of a pulse-induction EMI sensor should be well approximated by a sum of decaying exponentials. The resulting difference signal, however, is not well approximated in this manner (especially for small mines), but its integral is a very effective detector of buried metal.

The second step involves estimation of the feature vector Θ . For targets not too near the sensor head and not too large compared to the loop radius, the integrated sensor response described above has a spatial variation that is approximated by the function

$$\mathbf{G}_1(\mathbf{R}; \Theta_1) = B + \frac{S}{1 + (\|\mathbf{R} - \mathbf{R}_0\|/a)^b} \quad (24)$$

where \mathbf{R}_0 is the location of the target's centroid, B is the response of the background, S is the amplitude of the target response and a and b are shape parameters. At each sample position \mathbf{R} we fit all data within a window to this model. The size of the window is selected on the basis of the signal amplitude. A nonlinear optimization technique is then used to estimate the values of these parameters from samples within the window. The resulting feature vector is given by

$$\Theta_1 = [a \ b \ S \ B \ E \ \mathbf{R}_0]^T \quad (25)$$

where $E = |\mathbf{D} - \mathbf{G}_1|^2 / |\mathbf{D}|^2$ is the normalized residual error in the estimate.

4.2. Ground Penetrating Radar

4.2.1. Hardware

The GPR used in this effort was developed at The Ohio State University (OSU) ElectroScience Laboratory (ESL)³ and is shown in Figure 1. It employs a novel dielectric rod antenna, which is scanned horizontally over the earth at a fixed height. The antenna and its supporting platform are moved along a linear track in two inch intervals, and a horizontal scan of 55-inch length and 101 samples is taken in the cross track dimensions at each position. A network analyzer is used to acquire complex reflection coefficient measurements at 51 frequencies between 1 and 6 GHz. The frequency regime from 5 to 6 GHz was found to be dominated by clutter and was not used in this analysis.

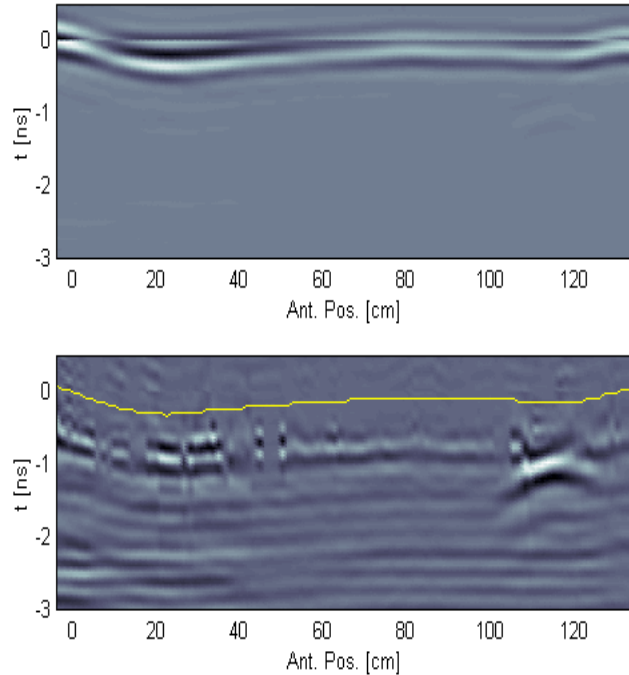


Figure 2. Example GPR data acquired over the 3.5 inch metal disk mine surrogate. The upper figure shows the time-domain data. The lower figure shows the same data with the surface reflection suppressed. The thin white line in the lower figure is the estimated position of the surface.

4.2.2. Signal Processing

A background signature (the response of the radar when the antenna points into free space) was acquired and subtracted from all data. The impulse response of the radar was determined by measuring scattering from a reference target (a short cylinder) in the absence of soil and comparing the result with a calculated scattered response. This impulse response was then deconvolved from the background-corrected measured data.

The resulting data are processed by first locating and removing the ground-reflected wave. This wave, a bandlimited approximation to an impulse response, is removed using an OSU-developed TLS-Prony technique.⁴ An example of the process is shown in Figure 2 for a 3.5 inch metal mine surrogate buried at 1 inch depth. Mine signatures are detected by performing a three-dimensional convolution with a matched filter, which comprises \mathbf{G}_2 . The filter is formed by simulating the response of a point scatterer in the soil medium at the depth of interest. GPR clutter tends to be highly correlated in the horizontal dimension (primarily because of the return from plane-stratified media), and whitening is performed prior to using this filter. The sensor feature vector Θ_2 used in classification is the position and maximum value of the matched filter within a window several samples wide at each position \mathbf{R} .

4.3. Infrared Camera

4.3.1. Hardware

We used a commercially available MWIR camera in this work. The sensor is a InSb photovoltaic array of 160 by 120 pixels. The optics provide a field of view of 9.1° by 6.8° for an IFOV of 1 mrad. The noise equivalent temperature difference for the sensor is 0.025 K. More information on the sensor and our IR measurements is available in a companion paper.⁵

The camera was positioned on the roof of the ESL building adjacent to the test site, and data were acquired from this vantage point. The camera’s field of view does not permit us to sample the entire region simultaneously, and a sequence of images was necessary. Fiducial markers were placed over the region to permit registration.




Figure 3. An example IR image after remapping to ground coordinates.

4.3.2. Signal Processing

Processing the IR data begins by remapping the imagery to ground coordinates using a standard perspective transformation. This step is followed by bilinear interpolation and resampling to a uniform pixel size. An example image is shown in Figure 3. An image chip is extracted at each interrogation point \mathbf{R} . Based on empirical observations, we employ a signature model identical to that used for the EMI sensor. A nonlinear optimization process is used to determine parameter values within the chip.

5. RESULTS

The fusion procedures described above in Sections 2 and 3 have been tested on experimental data acquired by the sensors described in Section 4. These data are preliminary and require further analysis before a complete comparison of the methods is possible. The test procedures and results are described below.

5.1. Test Site

Our experimental work was performed using a simulated mine field of 40 mine-like and clutter-like surrogates. The layout of this mine field is shown in Figure 4 and the identities of the buried objects are given in Table 1. The objects in this field had been in situ for more than 18 months at the time these data were acquired.

Our objective in creating this mine field was to provide an environment for testing sensors, validating fusion algorithms, and identifying problems that arise in realistic conditions. The target set is quite challenging. Roughly half the mine surrogates contain no metal (in contrast to real mines, essentially all of which contain some metal). In addition, some of the targets are quite small (only 1.5 inches in diameter), which present detection problems for all of the sensors. Those targets (all P15 objects in Figure 4) were designated clutter in what follows.

Consistent with our initial goal of simulating realistic field conditions, the test area was not specially conditioned prior to emplacing the targets. As a result, there exists a variety of clutter and buried objects in this area. Since our tests began about two years ago, roughly 100 metallic fragments, and a number of stones have been located in (and subsequently removed from) the region. The four concrete pilings shown in the figure were concealed by a lush

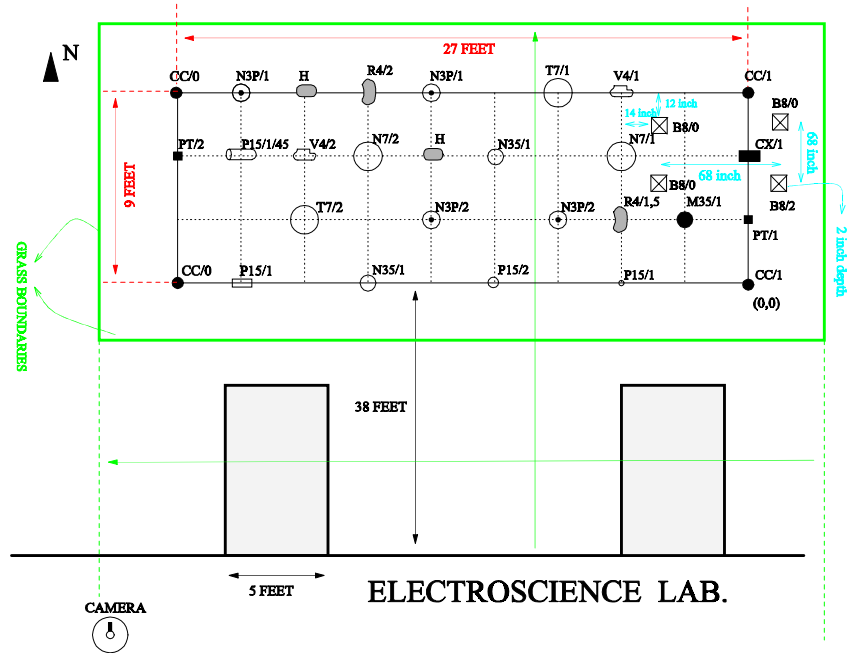


Figure 4. The layout of the ESL mine field.

grass cover that was present when the test objects were emplaced. Surface vegetation on the mine field has since been completely removed. In addition to these objects, the soil contains a large amount of clay, which affects GPR performance. The surface of the mine field contains a number of irregularities, but it is planar to within an estimated variation of ± 2 inches. As a result of this surface roughness, scanned sensors such as the EMI and GPR cannot maintain a constant height, and some undesirable variations in signal amplitude occur.

Position errors in the data arise from several sources. The positions of the targets shown in Figure 4 are accurate to within approximately one inch. Each sensor is operated on a different platform, and the platform position is known to within approximately one inch. As a result of roll, pitch, and yaw by the sensor platform, surface roughness produce errors in the sensor field of view on the order of one to two inches, for a total maximum position uncertainty of four inches. Processing of the individual sensor data sets has confirmed this estimate.

Data were acquired along the paths shown in Figure 5. The GPR samples were acquired at roughly 0.5 inch intervals along N-S lines spaced every two inches E-W, and the EMI sensor samples at a 1 inch interval along N-S lines spaced six inches E-W. The infrared camera acquired eight images as shown by the trapezoids. After remapping to ground coordinates, the pixel dimensions are approximately 0.5 inches. The small squares are IR fiducial markers, and the small circles are the target grid points shown in Figure 4.

5.2. Results for Feature-Level Fusion

Features were extracted from each of the sensors as described above, and those features were used in a classifier as indicated in equation (10). For feature-level fusion, in estimating the features Θ we did not permit the optimization process to use a variable shift \mathbf{R}_0 . Features samples were computed on a two inch grid of interrogation points \mathbf{R} over the entire mine field. The classifier $\Pr(H_k(\mathbf{R})|\Theta)$ was approximated by a backpropagation neural network. We attempted only the detection problem, for which the output hypotheses were $H_1 = \text{“mine”}$ and $H_2 = \text{“non-mine.”}$ The result of this process was compared to a threshold, leading to a binary map. Regions with small numbers of adjacent detections were discarded. The performance of the fused system is shown in Figure 6. Targets detectable by the fused sensor suite are indicated by circles, and we see that all of the mine surrogates have been detected with nine false alarms.

Table 1. Description of the objects buried in the MURI mine test grid

Abbreviation	Object	Diameter (inches)	Height (inches)
CC	Coke can		
CX	Crushed coke can		
PT	Pop top		
M35	Aluminum disk	3.5	5/8
P15	Plexiglass disk	1.5	2
R4	Rock	4	2
V4	Void(styrofoam)	4	2
H4	Refilled hole	6	2
N35	Nylon disk	3.5	5/8
N3P	Nylon disk with .078 diam 1 long steel pin	3	1
T7	Teflon disk	7	1
N7	Nylon disk	7	1
B8	Concrete pilings	8	varies

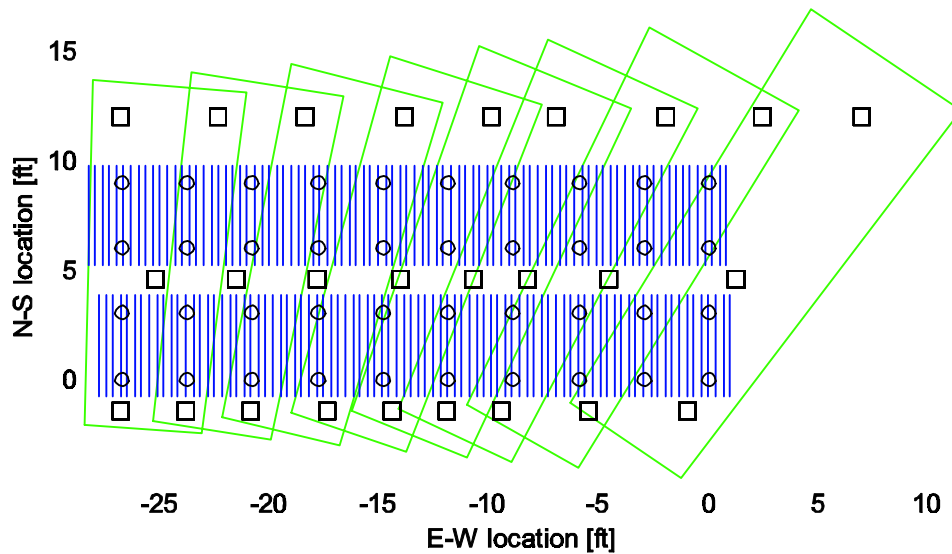


Figure 5. Locations of sample points for the sensor suite.

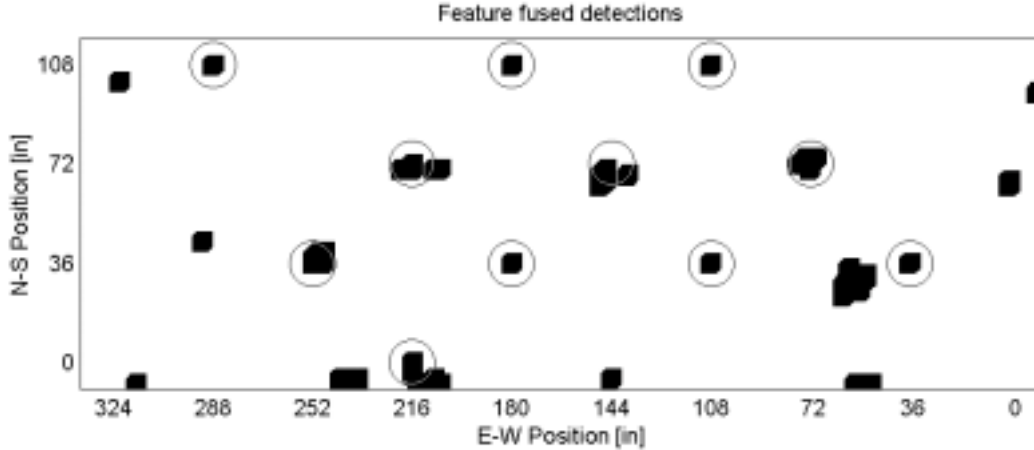


Figure 6. The results of feature-level fusion. The known targets are indicated by circles.

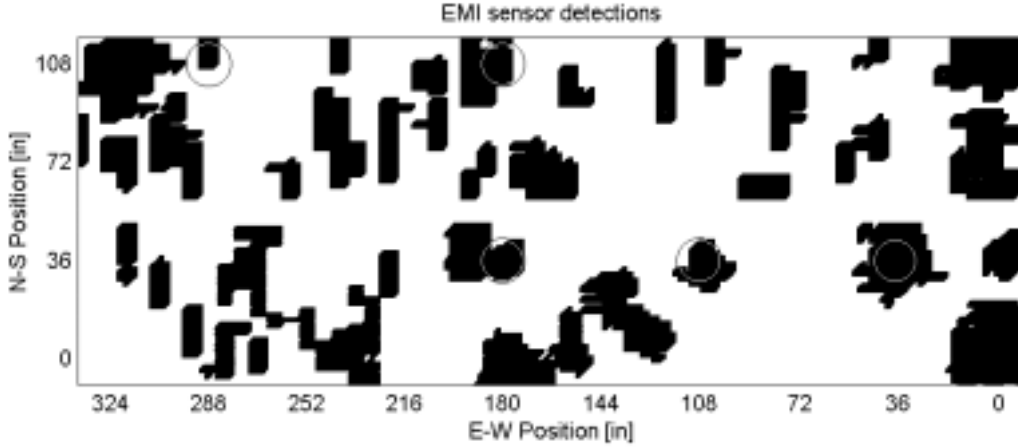


Figure 7. Detections by the EMI sensor.

5.3. Results for Decision-Level Fusion

For decision-level fusion we used the same grid of feature values described above. The resulting features θ_i for each sensor i are applied to sensor-specific classifiers (backpropagation networks) $\Pr(h_k(\mathbf{R})|\theta_i)$, and a target detection u_i is produced at each point. The target position information $\mathbf{R}_{0,i}$ was used to weight the classification results u_i as given in equation (16). The density $f_{\mathbf{R}_0|H_k}(\mathbf{R})$ was taken to be uniform over $|\mathbf{R} - \mathbf{R}_0| < 4$ inches and zero otherwise. This process leads to the three detection maps shown in Figure 7, 8, and 9.

Because our mine-field contains roughly equal numbers of metallic and metal-free mines, we used the EMI and GPR sensors to recognize metal-bearing objects and the IR and GPR sensors to recognize metal-free objects. A two-level fusion approach was employed. For the EMI-GPR sensor suite the intermediate hypotheses h_{11} = “metallic mine” and h_{12} = “metallic clutter” were used, while for the IR-GPR suite we used h_{21} = “non-metallic mine” and h_{22} = “non-metallic clutter.” The second-level fusion involved the trivial mapping $H_1 = h_{11} \cup h_{21}$ and $H_2 = h_{12} \cup h_{22}$. The set of all mine detections is shown in Figure 10. We find that the decision-level fusion system has produced a detection rate equal to that of the feature-level fused system (all targets were detected), and its false-alarm rate is comparable. The reduction in false alarms compared with the individual sensors is significant.

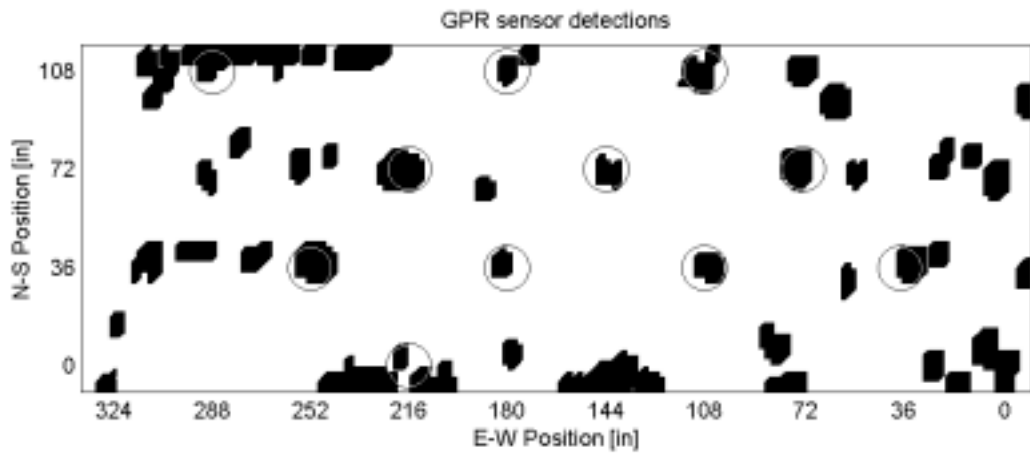


Figure 8. Detections by the GPR sensor.

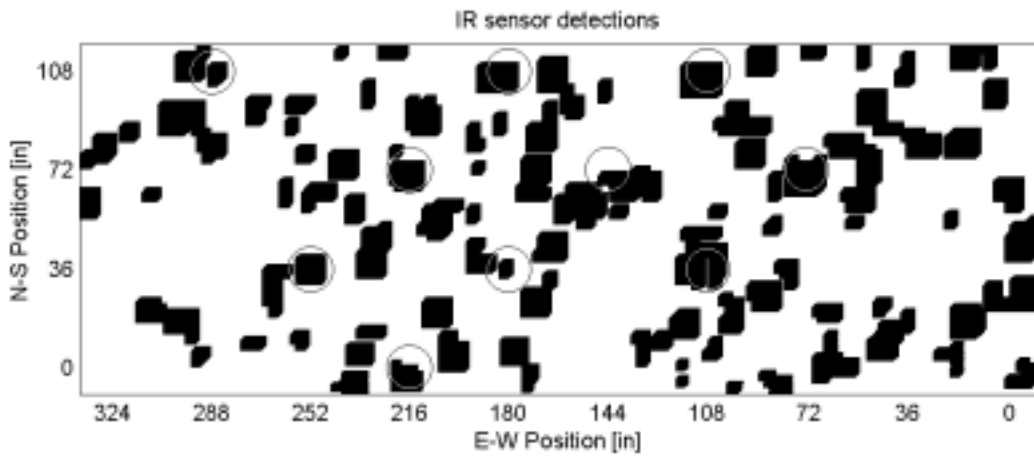


Figure 9. Detections by the IR sensor.




Figure 10. Detection of all mine types obtained by decision-level fusion of the EMI, GPR and IR sensors.

6. CONCLUDING REMARKS

Mine detecting sensors tend to produce data that are non-commensurate and sampled at non-coincident positions. We have developed and implemented a feature-level fusion technique for such data. The method has a rigorous basis, and with some approximations it can be reduced to form that is implemented efficiently and with a modest amount of information. The signal-processing tools required to use EMI, GPR and IR sensors for fusion have been developed.

We have conducted some preliminary tests of the algorithm and compared it to a decision-level fusion algorithm. Comparable performance was found for both forms of fusion. The theoretical work suggests that feature-level fusion may have some advantages, but further study is necessary to confirm these findings. To be more convincing, future tests of the algorithm must involve larger data sets. Development of more effective feature sets is also desirable to get maximum information from the available sensors.

ACKNOWLEDGMENTS

This project was supported by funds from Duke University under an award from the ARO (the OSD MURI program). The findings, opinions and recommendations expressed therein are those of the authors and are not necessarily those of Duke University or the ARO.

REFERENCES

1. P. K. Varshney, *Distributed Detection and Data Fusion*, Springer-Verlag, New York, NY, 1997.
2. B. V. Dasarathy, *Decision Fusion*, IEEE Computer Society Press, Los Alamitos, CA, 1994.
3. S. Nag, L. Peters, Jr., I. J. Gupta, and C.-C. Chen, "Ramp response for the detection of anti-personnel mines," in *Detection and Remediation Technologies for Mines and Minelike Targets IV*, J. T. B. A. C. Dubey, J. F. Harvey and R. E. Dugan, eds., *SPIE* **3710-137**.
4. I. J. Gupta, A. van der Merwe, and C.-C. Chen, "Extraction of complex resonances associated with buried targets," in *Detection and Remediation Technologies for Mines and Minelike Targets III*, J. T. B. A. C. Dubey, J. F. Harvey, ed., *SPIE* **3392**, pp. 1022–1032, 1998.
5. I. K. Sendur and B. A. Baertlein, "Techniques for improving buried mine detection in thermal ir imagery," in *Detection and Remediation Technologies for Mines and Minelike Targets II*, J. T. B. A. C. Dubey, J. F. Harvey and R. E. Dugan, eds., *SPIE* **3710-130**.