Feature-Level and Decision-Level Fusion of Noncoincidentally Sampled Sensors for Land Mine Detection

Ajith H. Gunatilaka, *Member, IEEE*, and Brian A. Baertlein, *Member, IEEE*

Abstract—We present and compare methods for feature-level (predetection) and decision-level (postdetection) fusion of multisensor data. This study emphasizes fusion techniques that are suitable for noncommensurate data sampled at noncoincident points. Decision-level fusion is most convenient for such data, but it is suboptimal in principle, since targets not detected by all sensors will not obtain the full benefits of fusion. A novel algorithm for feature-level fusion of noncommensurate, noncoincidentally sampled data is described, in which a model is fitted to the sensor data and the model parameters are used as features. Formulations for both feature-level and decision-level fusion are described, along with some practical simplifications. A closed-form expression is available for feature-level fusion of normally distributed data and this expression is used with simulated data to study requirements for sample position accuracy in multisensor data. The performance of feature-level and decision-level fusion algorithms are compared for experimental data acquired by a metal detector, a ground-penetrating radar, and an infrared camera at a challenging test site containing surrogate mines. It is found that fusion of binary decisions does not perform significantly better than the best available sensor. The performance of feature-level fusion is significantly better than the individual sensors, as is decision-level fusion when detection confidence information is also available (“soft-decision” fusion).

Index Terms—Land mines, sensor fusion, infrared, ground penetrating radar, metal detectors.

1 INTRODUCTION

BURIED land mines are an extremely hazardous fact of life for civilians in many countries. Various estimates put the number of buried landmines worldwide at over 80 million and growing. The extreme risk posed by mines implies that a mine sensor should have essentially no missed detections. This requirement and the clutter-rich environment in which mine sensors operate leads to high false alarm rates. Nonetheless, low false alarm rates are required because of the large land areas where mine clearance is necessary and the cost of clearing a putative mine. Several sensor technologies have been considered for the detection of mines, including electromagnetic induction (EMI), ground penetrating radar (GPR), infra-red (IR) imaging, quadrupole resonance (QR), various chemical detectors (e.g., electronic “sniffers”), and sensors of acoustically induced surface vibrations. At present, none of these technologies has proven to be acceptably effective and reliable for the wide range of mines and environmental conditions encountered in the field.

Because of the diverse physical properties measured by different sensor technologies, multisensor fusion is attractive as a means of improving mine detection.

Unfortunately, demining sensor fusion is hindered by several factors. First, sensors used for mine detection typically produce noncommensurate data, i.e., data that are not of comparable form and, hence, are not easily combined. By way of example, three commonly used mine sensors are IR cameras, which produce surface imagery; ground penetrating radars which produce depth profiles along a line scan; and metal detectors, which produce scalar measurements along a line scan. Even when the data are commensurate, the sensors may operate on different platforms, leading to noncoincident sampling and problems with data registration. Since fusion comprises combining information from multiple sources regarding a specified phenomenon (e.g., the presence of a mine) at a specific location, accommodating noncoincident samples is a necessity.

Some of these concerns can be alleviated by decision-level fusion, in which the operator performs detection at the sensor level and then combines (fuses) detections from individual sensors. This process reduces diverse sensor data to a common format (binary decisions or detection probabilities), which are readily combined using conventional statistical techniques. A variety of techniques have been developed for this approach [1], [2].

In principle, however, decision-level fusion is suboptimal since if a target is not detected by all sensors, it will not experience the full benefits of fusion. This situation and stressing performance requirements, led us to explore other methods of fusion. In this paper, we discuss a feature-level fusion approach to noncoincident sensor sampling. The data are represented by a model with unknown parameters (features) and random additive clutter. Optimization

• A.H. Gunatilaka is with Lucent Technologies, 6200 E. Broad Street, Columbus, OH 43213.

• B.A. Baertlein is with The Ohio State University, ElectroScience Laboratory, 1320 Kinnear Road, Columbus, OH 43212.
E-mail: baertlein.1@osu.edu.

Manuscript received 8 Feb. 2000; revised 28 Nov. 2000; accepted 14 Dec. 2000.

Recommended for acceptance by R. Beveridge.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number 111425.

techniques are used to determine features from the available data, and classification is performed on the basis of the features. The effects of sample position uncertainty are also examined for the case of normally distributed features.

Fusion of demining sensor data has been investigated previously. Chauduri et al. [3] developed a demining system in which coregistered GPR and EMI sensors were fused using several decision-level fusion approaches. Brumark et al. [4] demonstrated the decision-level fusion of coincidentally sampled GPR and EMI sensor data collected over targets buried in a sand box. Clark et al. [5] fused two IR sensors of different wavelengths, which were registered using a set of fiducial markers. A neural network-based feature-level fusion approach was used in that work. Miao et al. considered the fusion of six coregistered IR images corresponding to different wavelengths [6]. Each image was separately classified with neural networks and the resulting decisions were fused using both majority voting and a consensus theory approach. Weissensteil et al. [7] studied decision-level fusion of GPR and EMI data using synthetic data. The problem was greatly simplified by assuming that the mines were detectable by both sensors, metallic clutter by EMI alone and nonmetallic clutter by GPR alone. Only an AND rule was considered for fusion. Breejen et al. [8] compared several methods for decision-level fusion of EMI, GPR, and IR data. They used the same data for both training and testing the classifiers, which often yields optimistic results. The performance of a vehicle mounted mine detector employing GPR, EMI, and video and IR cameras, developed by EG&G Management Systems, Incorporated, was reported in [9]. While surface emplaced antitank mines and mines with significant metal content were easily detected, buried plastic mines were found to be challenging. Development and testing of a vehicle-mounted multisensor demining platform that uses EMI, GPR, and IR for detection and a TNA sensor for target confirmation is underway at the Canadian Defence Research Establishment [10]. Gunatilaka and Baertlein [11] compared predetection (feature-level) and postdetection (hard and soft decision-level) fusion of EMI, GPR, and IR data. A benefit was observed for feature-level fusion, and fusion of soft decisions was found to be better than fusion of hard decisions. Gunatilaka and Baertlein [12] compared several methods for feature extraction and several fusion strategies. In all cases, feature-level fusion was found to be superior to decision-level fusion. Baertlein [13] recently investigated the problem of constructing fusion algorithms with limited training data. Using approaches based on Statistical Learning Theory for feature-level fusion and joint optimization of local decision thresholds for hard decision fusion, it was shown that algorithm performance could be improved. For the specific training data used in that work, superior performance was noted for hard decision fusion. Marble et al. [14] demonstrated the benefits of using hard decision-level fusion to combine EMI, GPR, and magnetometer data for detection of mines with large metal content. In that work, significant problems with misregistration of noncommensurate sensors were overcome by converting all

sensor outputs to a scalar, image-like format and interpolating the result.

The work reported here complements and extends the prior studies in several respects. First, we consider both feature-level and decision-level fusion of experimental data from noncommensurate (EMI, GPR, and IR) sensors. Second, we investigate the benefits of fusion under more realistic conditions by using a test site containing large amounts of clutter and having an irregular surface. Finally, each sensor was used on a different platform, which led us to consider the effects of noncoincident sampling and positioning errors.

The work is organized in three major sections. In Section 2 we describe the theoretical basis for feature-level fusion of noncoincident samples. The basis for decision-level fusion is presented in Section 3. Experimental data and examples of fusion are given in Section 4. For a list of symbols referenced in this paper, see Table 1. Finally, concluding remarks appear in Section 5.

2 FEATURE-LEVEL FUSION OF RANDOMLY SAMPLED MULTISENSOR DATA

2.1 Problem Definition

Suppose that a suite of N_S countermeasure sensors are used to acquire data in a region. Sensor i acquires J_i data samples $\mathbf{d}_i$ at locations $\mathbf{R}_{ij}$, which we denote $\mathbf{d}_{ij} \equiv \mathbf{d}_i(\mathbf{R}_{ij})$. Fusion is to be performed on the data set $\mathcal{D} = \{\mathbf{d}_i(\mathbf{R}_{ij}), i = 1, \dots, N_S; j = 1, \dots, J_i\}$. We assume that different sensors produce samples that differ in number J_i , sample positions $\mathbf{R}_{ij}$, format, and the dimensionality of their data.

Let $\mathbf{R}$ be a point where we wish to determine the presence of a mine. We refer to $\mathbf{R}$ as the "interrogation point." We form the K hypotheses $H_k(\mathbf{R})$, $k = 1, 2, \dots, K$ for the fused suite regarding the presence or absence of various types of mines at $\mathbf{R}$. We may also include in the set $\{H_k(\mathbf{R})\}$ discrete clutter objects, e.g., buried rocks, metallic clutter, etc. The explicit dependence of these hypotheses on the interrogation point is omitted hereafter in the interest of brevity.

The decision criterion for the fusion processor is the Bayes risk. At each point $\mathbf{R}$, we determine the hypothesis H_k that minimizes the risk $\mathcal{R}$ conditioned on the data $\mathcal{D}$, viz:

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

where $\mathcal{R}_j(\mathcal{D})$ is the risk of selecting hypothesis H_j given data $\mathcal{D}$, and C_{ij} is the cost of choosing H_i when H_j is true. If all costs are equal or for the special case $K = 2$ (binary detection), it is sufficient to compute the likelihood ratio and to compare it to a threshold. In general, however, multiple hypotheses may be involved and equal costs are inappropriate for the mine-detection problem. It is difficult to assign costs in the mine-detection problem but, fortunately, fusion can be explored without knowledge of the costs. Results are reported here in a manner independent of cost, specifically, a posteriori probabilities $\Pr(H_k|\mathcal{D})$ and ROC curves are computed. When suitable costs are known, one can use the general formulation given above to

TABLE 1
List of Symbols

Symbol	Description
N_S	The number of demining sensors
$\mathbf{R}$	The location at which fusion is being used to determine the presence of a mine
H_k	The k th hypothesis regarding the object present at location $\mathbf{R}$
K	The number of possible hypotheses H_k
$\mathbf{R}_{ij}$	The location at which the j th sample of sensor i is acquired
$\mathbf{d}_{ij}$	The data sampled by sensor i at the j th sampling position
$\mathbf{d}_i$	The data acquired by sensor i
J_i	The number of data samples acquired by sensor i
M_i	The number of data samples from sensor i used in fusion at $\mathbf{R}$
N_{D_i}	The number of elements in each sample from sensor i
N_E	The number of elements from all sensors and all samples used in fusion [= $\sum_i M_i N_{D_i}$]
$\mathcal{D}$	The set of data acquired by all sensors at all sampling points
$\mathcal{R}_j(\mathcal{D})$	The risk of choosing hypothesis H_j given data $\mathcal{D}$
C_{jk}	The cost of choosing H_j when H_k is true
$\Pr(H_k \mathcal{D})$	The probability that H_k is true given data $\mathcal{D}$
θ_{ik}	The feature vector for sensor i given hypothesis H_k
$\mathbf{R}_{0,ik}$	The nominal center of the mine signature for sensor i under H_k
Θ_{ik}	$= [\theta_{ik} \mathbf{R}_{0,ik}]^T$
Θ_k	$= [\Theta_{1k} \Theta_{2k} \dots \Theta_{N_S k}]^T$
$\mathbf{g}_i(\mathbf{R}_{ij}; \Theta_{ik})$	Noise-free (and clutter-free) signal acquired by sensor i at $\mathbf{R}_{ij}$ when H_k is true
$\mathbf{n}_i(\mathbf{R}_{ij})$	The contribution of noise and clutter to the signal measured by sensor i at $\mathbf{R}_{ij}$
$\mathbf{D}$	A vector comprising the subset of $\mathcal{D}$ relevant to making the target decision at $\mathbf{R}$
$\mathbf{N}$	The noise and clutter vector corresponding to $\mathbf{D}$
$\mathbf{G}_k(\Theta_k)$	The signal vector corresponding to $\mathbf{D}$
h_{iq}	The q -th hypothesis of sensor i
Q_i	The number of distinct hypotheses that can be reported by sensor i
$\mathbf{H}$	$= [h_1 h_2 \dots h_{N_S}]^T$
t_i	The decision threshold used by sensor i
$\mathbf{T}$	$= [t_1 t_2 \dots t_{N_S}]^T$
u_i	The decision reached by sensor i in response to the data
u_0	The decision resulting from decision-level fusion of $\mathbf{U}$
$\mathbf{U}$	$= [u_1 u_2 \dots u_{N_S}]^T$
$P_{F_i}(t_i), P_{M_i}(t_i)$	The probabilities of false and missed detection respectively for sensor i
P_D, P_F, P_M	The respective probabilities of detection, false detection and missed detection after fusion
$\Lambda(\mathbf{D})$	The likelihood ratio
$\Lambda_G(\mathbf{D})$	The likelihood ratio for the generalized likelihood ratio test

derive the minimum risk decision from the a posteriori probabilities. Similarly, knowledge of the costs permits one to readily determine minimum risk operating points on a ROC curve.

2.2 Fusion of Noncoincidently Sampled Features

We assume that sensor i acquires at a point $\mathbf{R}_{ij}$ data $\mathbf{d}_i(\mathbf{R}_{ij})$, which lies in the vicinity of a putative target of type k at $\mathbf{R}$. This data is modeled as a signal $\mathbf{g}_i(\mathbf{R}_{ij}; \Theta_{ik})$ corrupted by additive clutter (and noise) $\mathbf{n}_i(\mathbf{R}_{ij})$, viz:

$$\mathbf{d}_i(\mathbf{R}_{ij})|H_k = \mathbf{g}_i(\mathbf{R}_{ij}; \Theta_{ik}) + \mathbf{n}_i(\mathbf{R}_{ij}), \quad (2)$$

where Θ_{ik} is a feature vector (i.e., a set of model parameters) that describes the signature of targets of type k for sensor i (see Fig. 1). We assume that the clutter $\mathbf{n}_i(\mathbf{R})$ is independent of the parameter Θ_{ik} . Note that, if the clutter discretizes are also assigned specific hypotheses H_j , the phenomena modeled by $\mathbf{n}_i(\mathbf{R})$ can be relatively smooth.

For each sensor, the presence of a mine produces a measurable response over a finite spatial region and, hence, only a subset of $\mathcal{D}$ may be relevant to interrogation of the region around $\mathbf{R}$. Let the relevant data for sensor i be given

by subsequences $j_1, j_2, \dots, j_{M_i}$ with length M_i . It is convenient to assemble this reduced data set into the vectors

$$\mathbf{D} = \begin{bmatrix} \mathbf{d}_1(\mathbf{R}_{1j_1}) \cdots \mathbf{d}_1(\mathbf{R}_{1j_{M_1}}) \\ \mathbf{d}_2(\mathbf{R}_{2j_1}) \cdots \mathbf{d}_{N_S}(\mathbf{R}_{N_S j_{M_{N_S}}}) \end{bmatrix}^T \quad (3)$$

$$\mathbf{N} = \begin{bmatrix} \mathbf{n}_1(\mathbf{R}_{1j_1}) \cdots \mathbf{n}_1(\mathbf{R}_{1j_{M_1}}) \\ \mathbf{n}_2(\mathbf{R}_{2j_1}) \cdots \mathbf{n}_{N_S}(\mathbf{R}_{N_S j_{M_{N_S}}}) \end{bmatrix}^T \quad (4)$$

$$\mathbf{G}_k(\Theta_k) = \begin{bmatrix} \mathbf{g}_{1k}(\mathbf{R}_{1j_1}; \Theta_{1k}) \\ \cdots \mathbf{g}_{N_S k}(\mathbf{R}_{N_S j_{M_{N_S}}}; \Theta_{N_S k}) \end{bmatrix}^T. \quad (5)$$

If a single datum $\mathbf{d}_i(\mathbf{R})$ from sensor i has length N_{D_i} , then the length of these column vectors is $N_E = \sum_i N_{D_i} M_i$. Using this notation, we have for the combined signal model $\mathbf{D}|H_k = \mathbf{G}_k(\Theta_k) + \mathbf{N}$, where the combined (concatenated) feature vector for all sensors is

$$\Theta_k = [\Theta_{1k} \Theta_{2k} \cdots \Theta_{N_S k}]^T. \quad (6)$$

Fig. 1. The data model for three hypotheses $k = 0, 1, 2$. For each sensor i data $d_i(R_{ij})$ are available. For hypothesis k , the sensor data is represented by a known function $g_{ik}(R; \Theta_{ik})$ of position R and random parameters Θ_{ik} . The reference position of signature g_{ik} is given by $R_{0,ik}$. The signature parameters Θ_{ik} , which include the reference position $R_{0,ik}$, are estimated from the data. In this example, $k = 0$ corresponds to the no-target case, $k = 1$ is a target with a unimodal signature, and $k = 2$ is a target with a bimodal signature. The objective of fusion is to determine if a target exists at any point R on the basis of data from multiple sensors.

To support both decision-level fusion and feature-level fusion with uncertain sample locations, we include in Θ_{ik} a position offset $R_{0,ik}$ that describes the nominal "center" of the mine signature with respect to R and we write $\Theta_{ik} = [\theta_{ik} \ R_{0,ik}]^T$. In general, we take θ_{ik} and $R_{0,ik}$ to be independent in what follows.

2.3 Approximate Forms

Feature-level fusion of the foregoing data can be formulated along two lines. Consider first the direct approach using the a posteriori probabilities, which avoids the problem of estimating the a priori probabilities. Minimizing the risk in (1) leads us to consider the a posteriori probabilities as an integral over the classifier $\Pr(H_k|\Theta_k)$

$$\Pr(H_k|\mathbf{D}) = \int d\Theta_k \Pr(H_k|\Theta_k) f_{\Theta|\mathbf{D}}(\Theta_k), \quad (7)$$

in which we have assumed that H_k is determined by Θ_k alone and, hence, $\Pr(H_k|\Theta_k, \mathbf{D}) = \Pr(H_k|\Theta_k)$. This expression is exact, but it requires knowledge of the density $f_{\Theta|\mathbf{D}}$ and that we perform an integral over a feature space of possibly high dimensions. To avoid these issues, it is attractive to employ the maximum a posteriori (MAP) approximation. If the data $\mathbf{D}$ strongly imply the value $\hat{\Theta}$ and the features Θ_k are not too specific (i.e., the maximum of $\Pr(H_k|\Theta_k)$ is relatively broad), then $f_{\Theta|\mathbf{D}}$ will be strongly peaked about $\hat{\Theta}$ and we have

$$\Pr(H_k|\mathbf{D}) \approx \Pr(H_k|\hat{\Theta}) \int d\Theta_k f_{\Theta_k|\mathbf{D}}(\Theta_k) = \Pr(H_k|\hat{\Theta}). \quad (8)$$

Conversely, if $\mathbf{D}$ provides no information about Θ , then $f_{\Theta|\mathbf{D}}(\Theta) = f_{\Theta}(\Theta)$ and $\Pr(H_k|\mathbf{D})$ reduces to the a priori probability.

Feature-level fusion can also be formulated using class-conditioned densities, and this approach will be used in our study of sensor position errors. Bayes' rule leads to

$$\mathcal{R}_k(\mathbf{D}) = \frac{1}{f_{\mathbf{D}}} \sum_{j=1}^K C_{jk} \Pr(H_j) f_{\mathbf{D}|H_j}(\mathbf{D}). \quad (9)$$

The a priori probabilities $\Pr(H_j)$ are required but, since those probabilities are seldom known, this approach is somewhat less attractive in practice. The factor $[f_{\mathbf{D}}]^{-1}$ is common to all $\mathcal{R}_k(\mathbf{D})$ and, hence, our primary task is to evaluate the class-conditioned probabilities $f_{\mathbf{D}|H_j}(\mathbf{D})$. Since $\mathbf{G}_k$ is presumed known, the data $\mathbf{D}$ conditioned on the parameter Θ_k has the same density as the clutter $\mathbf{N}$, which leads to $f_{\mathbf{D}|\Theta_k, H_k}(\mathbf{D}) = f_{\mathbf{N}}(\mathbf{D} - \mathbf{G}_k(\Theta_k))$ and

$$f_{\mathbf{D}|H_k}(\mathbf{D}) = \int d\Theta_k f_{\Theta_k|H_k}(\Theta_k) f_{\mathbf{N}}(\mathbf{D} - \mathbf{G}_k(\Theta_k)). \quad (10)$$

We approximate this exact result as done in (8). Assuming that $f_{\mathbf{N}}(\mathbf{D} - \mathbf{G}_k(\Theta_k))$ has a well-defined maximum for $\Theta_k = \hat{\Theta}_k$ and that $f_{\Theta_k|H_k}$ is not strongly peaked leads to

$$\begin{aligned} f_{\mathbf{D}|H_k}(\mathbf{D}) &\approx f_{\Theta_k|H_k}(\hat{\Theta}_k) \int d\Theta_k f_{\mathbf{N}}(\mathbf{D} - \mathbf{G}_k(\Theta_k)) \\ &\approx E_k f_{\Theta_k|H_k}(\hat{\theta}_k) f_{\mathbf{R}_{0,k}|\mathbf{D}}(\hat{\mathbf{R}}_{0,k}), \end{aligned} \quad (11)$$

where we have approximated θ_{ik} and $\mathbf{R}_{0,ik}$ as conditionally independent, and the constant factor

$$E_k = \int d\Theta_k f_{\mathbf{N}}(\mathbf{D} - \mathbf{G}_k(\Theta_k)) \quad (12)$$

is related to the residual error.

For the two class problem ($K = 2$), a likelihood ratio formulation is possible. The likelihood ratio for a known signal with unknown parameters is given by

$$\Lambda(\mathbf{D}) = \frac{f_{\mathbf{D}|H_1}(\mathbf{D})}{f_{\mathbf{D}|H_0}(\mathbf{D})} = \frac{\int d\Theta f_{\Theta|H_1}(\Theta) f_{\mathbf{N}}(\mathbf{D} - \mathbf{G}_1(\Theta))}{\int d\Theta f_{\Theta|H_0}(\Theta) f_{\mathbf{N}}(\mathbf{D} - \mathbf{G}_0(\Theta))}. \quad (13)$$

If the same model is used for both hypotheses, then approximating the results as proposed above yields

$$\Lambda(\mathbf{D}) \approx \frac{f_{\Theta|H_1}(\hat{\Theta}_1) \int d\Theta f_{\mathbf{N}}(\mathbf{D} - \mathbf{G}_1(\Theta))}{f_{\Theta|H_0}(\hat{\Theta}_0) \int d\Theta f_{\mathbf{N}}(\mathbf{D} - \mathbf{G}_0(\Theta))} = \frac{f_{\Theta|H_1}(\hat{\Theta}_1) E_1}{f_{\Theta|H_0}(\hat{\Theta}_0) E_0}, \quad (14)$$

where $\hat{\Theta}_k$ is the estimate of Θ that maximizes $f_{\mathbf{N}}(\mathbf{D} - \mathbf{G}_k(\Theta_k))$ under H_k . Detection is performed on the basis of the class-conditioned densities of Θ .

It is interesting to compare this formulation to the more conventional generalized likelihood ratio test (GLRT) [15, p. 38]. In the GLRT formulation, the model parameters Θ are estimated using the maximum likelihood technique. This leads to the expression

$$\Lambda_G(\mathbf{D}) \approx \frac{\max_{\Theta_1} f_{\mathbf{D}|\Theta_1, H_1}}{\max_{\Theta_0} f_{\mathbf{D}|\Theta_0, H_0}} = \frac{f_{\mathbf{N}}(\mathbf{D} - \mathbf{G}_1(\hat{\Theta}_1))}{f_{\mathbf{N}}(\mathbf{D} - \mathbf{G}_0(\hat{\Theta}_0))}. \quad (15)$$

Thus, GLRT detection is based primarily on the distribution of the model error. The model parameters are considered indirectly, via their contribution to the error. The estimator-correlator method [15, p. 79] uses a similar approach, in which the parameters are estimated and the resulting model signature is correlated with the data. For Gaussian statistics,

detection is based on the correlation. The feature-level fusion algorithm described above could be referred to as an “estimator-classifier” technique in that detection (or classification) is based on the parameters rather than the residual error in the model. We will also find it useful to include in the parameter vector Θ_k , a factor related to the integrated residual error density E_k , typically, in the form of the error energy.

2.4 Gaussian Approximation

The case of Gaussian statistics is of interest because it provides insight into the problems of feature-level fusion and because it can be used to efficiently simulate the performance of certain fusion strategies. With certain assumptions, the expression for $f_{D|H_k}$ in (10) can be evaluated in closed form.

Consider first, the case in which only the clutter $\mathbf{N}$ has a Gaussian density. That is, the N_s random variables $\mathbf{N}_i(\mathbf{R}_{ij})$ (each of which are evaluated at M_i locations) have a density that is jointly normal with mean μ_N and covariance matrix $\mathbf{C}_N$.

Now, make the additional assumption that the mine signature shape matrix $\mathbf{G}_k$ is known (including the underlying mine positions $\mathbf{R}_{0,ik}$) up to an unknown signature amplitude. Taking θ_k to be this amplitude, we make the linear approximation $\mathbf{G}_k(\Theta_k) = \tilde{\mathbf{G}}_k \theta_k$ in which $\tilde{\mathbf{G}}_k$ is a known constant vector independent of θ_k . If the features θ_k are also Gaussian distributed with density $\theta_k|H_k \sim \mathcal{N}(\mu_{\theta_k}, \mathbf{C}_{\theta_k})$, we obtain the Bayesian linear model [16, p. 326], which leads to

$$f_D(\mathbf{D}|H_k) = \frac{(2\pi)^{-N_E/2} e^{-Q_k/2}}{\det(\mathbf{C}_N)^{1/2} \det(\mathbf{C}_{\theta_k})^{1/2} \det(\mathbf{C}_{\hat{\theta}_k}^{-1} + \mathbf{C}_{\theta_k}^{-1})^{1/2}}, \quad (16)$$

where N_E is the length of these vectors (defined previously) and

$$Q_k = (\mathbf{D} - \tilde{\mathbf{G}}_k \hat{\theta}_k - \mu_N)^T \mathbf{C}_N^{-1} (\mathbf{D} - \tilde{\mathbf{G}}_k \hat{\theta}_k - \mu_N) + (\hat{\theta}_k - \mu_{\theta_k})^T (\mathbf{C}_{\hat{\theta}} + \mathbf{C}_{\theta_k})^{-1} (\hat{\theta}_k - \mu_{\theta_k}). \quad (17)$$

In this result, the quantity $\hat{\theta}_k$ is the maximum-likelihood estimate (MLE) for the feature vector θ_k , given by

$$\hat{\theta}_k = \mathbf{C}_{\hat{\theta}} \tilde{\mathbf{G}}_k^T \mathbf{C}_N^{-1} (\mathbf{D} - \mu_N) \quad (18)$$

and $\mathbf{C}_{\hat{\theta}_k}^{-1} = \tilde{\mathbf{G}}_k^T \mathbf{C}_N^{-1} \tilde{\mathbf{G}}_k$ is its covariance. The terms of (17), we interpret as errors due to random clutter and to feature variability, respectively. Since the first term in (17) is identical to that arising in $f_{D|\Theta, H_k}$, it is clear that the GLRT approximation of (15) ignores the second term in (17), which describes the feature distribution.

2.5 Sample Position Errors

We noted earlier that to fuse multisensor data, it is necessary to know what locations in the environment are represented by the data, so that one can make a statement about the presence of a target the “interrogation” point. In practice, errors in sample position are largely unavoidable, especially when the sensors are located on different platforms. Some common types of sample errors include

uniform offsets produced by a baseline error, linearly increasing offsets produced by a drift in position baseline, isotropically distributed errors due to sample-to-sample positioning inaccuracy, and combinations of the above.

When position errors are present, the sample position $\mathbf{R}_{ij}$ has the form $\hat{\mathbf{R}}_{ij} + \mathbf{r}_{ij}$, where $\hat{\mathbf{R}}_{ij}$ is the intended position and $\mathbf{r}_{ij}$ is an unknown error. Thus,

$$\mathbf{d}_i(\mathbf{R}_{ij})|H_k = \mathbf{g}_k(\hat{\mathbf{R}}_{ij} + \mathbf{r}_{ij}; \Theta) + \mathbf{n}(\hat{\mathbf{R}}_{ij} + \mathbf{r}_{ij}).$$

Using a Bayesian approach and making the additional assumption that the sample positions are independent of the mine signature Θ_k and of the hypothesis H_k , these random variables are integrated out as follows:

$$\begin{aligned} f_{D|H_k}(\mathbf{D}) = & \int d\mathbf{r}_{1j_1} \int d\mathbf{r}_{1j_2} \cdots \int d\mathbf{r}_{N_s j_{M_{N_s}}} f_{\mathbf{r}_{1j_1}, \mathbf{r}_{1j_2}, \dots, \mathbf{r}_{N_s j_{M_{N_s}}}(\mathbf{r}_{1j_1}, \mathbf{r}_{1j_2}, \dots, \mathbf{r}_{N_s j_{M_{N_s}}})} \\ & \int d\Theta_k f_{\Theta_k|H_k}(\Theta_k) f_{N_{1j_1}, \dots, N_{N_s j_{M_{N_s}}}}[\mathbf{d}_{1j_1} - \mathbf{g}(\hat{\mathbf{R}}_{1j_1} + \mathbf{r}_{1j_1}; \Theta_{1k}), \dots, \\ & \mathbf{d}_{N_s j_{M_{N_s}}} - \mathbf{g}(\hat{\mathbf{R}}_{N_s j_{M_{N_s}}} + \mathbf{r}_{N_s j_{M_{N_s}}}; \Theta_{N_s k})]. \end{aligned} \quad (19)$$

The effect of position errors on system performance for feature-level fusion was studied using this result and the Bayesian linear model described in (16). If modest numbers of sensor data points are being fused at each interrogation point, the resulting integrals can be numerically evaluated using Monte Carlo techniques. Gaussian-distributed features were simulated for three sensors with isotropically distributed errors around an intended sample point. The position errors were normally distributed with variance σ^2 . The sensor signature models $\mathbf{g}_{ik}$ were taken to be zero when the mine was absent and Gaussian functions with variance Δ_i^2 when the mine was present. A single feature was simulated for each sensor and a common value $\Delta_i \equiv \Delta$ was used for all sensors. The “feature-to-noise” ratio $\sqrt{\text{var}(\theta_{ik})/\text{var}(N_{ik})}$ (an analog of the signal-to-noise ratio for this problem) was set to two.

It was found that system performance in the presence of isotropic position errors could be roughly characterized by the ratio σ/Δ . Figs. 2, 3, and 4 show the results for $\sigma/\Delta = 0.5, 1$, and 1.5 , respectively. (The case of a random offset was also examined. The results were similar and are not shown.) We see that for $\sigma/\Delta < 1$ the performance of the classifier is not significantly affected by position offsets, but for $\sigma/\Delta > 1$ there is a significant degradation in performance. These results suggest the intuitively reasonable finding that the performance of feature-level fusion systems will not be adversely effected if the positioning errors are smaller than the extent of the measurable target signature. When this condition is violated, fusion becomes less effective since all sensors may not see the target.

It is worth noting that position-related errors can sometimes be overcome through ingenuity on the part of the analyst and, hence, these results should be viewed as suggestive only. For example, when the position errors comprise constant offsets and the individual sensor performance is good (i.e., high probability of detection

Fig. 2. The ROC curve for simulated data with normalized position errors $\sigma/\Delta = 0.5$. Performance is essentially perfect.

and low false alarm rates), the offset might be estimated by registering the detections with those of other sensors. Such conditions, however, are uncommon in mine detection.

3 DECISION-LEVEL FUSION

The formulation of an optimal decision-level fusion algorithm is similar to that defined in Section 2 for feature-level fusion. As before, each sensor i collects data $\{\mathbf{d}_i(\mathbf{R}_{i1}), \dots, \mathbf{d}_i(\mathbf{R}_{iM_i})\}$. On the basis of these data, it makes an intermediate declaration u_i regarding the truth of intermediate sensor-specific hypotheses $h_{iq}(\mathbf{R})$ for $q = 1, 2, \dots, Q_i$. For example, a sensor suite that comprises EMI, GPR, and IR sensors might use $Q_i = 2$ for all i with $h_{11/12} = \text{"metallic object present/not present,"}$ $h_{21/22} = \text{"dielectric discontinuity present/not present,"}$ and $h_{31/32} = \text{"thermal anomaly present/not present"}$ rather

Fig. 3. The ROC curve for simulated data with normalized position errors $\sigma/\Delta = 1$. Performance degrades as a result of sensor position errors.

Fig. 4. The ROC curve for simulated data with normalized position errors $\sigma/\Delta = 1.5$. Sensor position errors are so large that fusion is ineffective.

than simply $H_{1/2} = \text{"mine present/not present."}$ In what follows, we restrict our attention to the case of equal numbers of intermediate hypotheses $Q_i \equiv Q$.

3.1 General Formulation

The goal of fusion is to combine these declarations into a fused declaration u_0 that addresses hypotheses H_k , $k = 1, 2, \dots, K$. Each detection u_i also involves a threshold t_i , which must be specified by the operator. To simplify the notation, form the vector of individual declarations given by $\mathbf{U} \equiv [u_1 \ u_2 \ \dots \ u_{N_s}]^T$ and the threshold vector $\mathbf{T} \equiv [t_1 \ t_2 \ \dots \ t_{N_s}]^T$. The declarations $\mathbf{U}$ and thresholds $\mathbf{T}$ are assumed to be the only data on which the decision u_0 is based. In particular, it is assumed that the fusion processor has no knowledge of the sensor data $\mathbf{D}$. It is further assumed that, for each sensor i , the detection confidences $\Pr(u_i|h_{iq})$ and the a priori probabilities $\Pr(H_k)$ are known for all k .

For optimal decision-level fusion, we seek a rule for fusing $\mathbf{U}$ into u_0 such that the Bayes risk is minimized. We can write this risk in the form

$$\mathcal{R} = \sum_{\mathbf{U}} \left[\sum_{j=1}^K \sum_{k=1}^K C_{jk} \Pr(H_k) \Pr(u_0 = j|\mathbf{U}) \Pr(\mathbf{U}|H_k; \mathbf{T}) \right], \quad (20)$$

which makes explicit the fusion rule $\Pr(u_0 = j|\mathbf{U})$. In this result, the summation on $\mathbf{U}$ is over all possible combinations of sensor outcomes and the decision u_0 is based only on the local decisions u_i . The individual decisions u_i address the hypotheses h_{iq} , which are related to H_k in a manner that implies the following:

$$\Pr(\mathbf{U}|H_k; \mathbf{T}) = \sum_{q=1}^Q \Pr(\mathbf{U}|\mathbf{H}_q; \mathbf{T}) \Pr(\mathbf{H}_q|H_k), \quad (21)$$

where $\mathbf{H}_q = [h_{1q} \ h_{2q} \ \dots \ h_{N_s q}]$ and we have used the fact that decision u_i depends on h_{iq} and not on H_k . It remains to

specify the fusion rule $\Pr(u_0 = j|\mathbf{U})$ that will minimize the bracketed quantity in (20). Some approaches are discussed in the following sections.

3.2 An Optimal Hard-Decision Fusion Rule

For sensors that produce “hard” decisions (i.e., all-or-nothing declarations), the fusion rule is a mapping from the N_s -fold product of the space $\{1, 2, \dots, Q\}$ to one of K output hypotheses. There exist $K^{Q^{N_s}}$ such mappings and, from this (possibly large) set, one must determine the most effective rule. In general, the most effective mapping will depend on the effectiveness of the individual sensors. Ad hoc mappings (AND, OR, majority voting, etc.) have been used, some with considerable success.

In some situations, there exist optimal mappings for decision-level fusion [1, §§ 3.3, 3.4]. 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 and the sensors output hard decisions. For $K = Q = 2$, the minimum in (20) reduces to a well-known likelihood ratio test, given by

$$\frac{\Pr(\mathbf{U}|\mathbf{H}_1; \mathbf{T})}{\Pr(\mathbf{U}|\mathbf{H}_0; \mathbf{T})} \underset{u_0=0}{\overset{u_0=1}{>}} \eta, \quad (22)$$

where η is the threshold for the fused decision and is independent of $\mathbf{T}$. If the declarations u_i are conditionally independent for each sensor,¹ then we can factor the joint probabilities into a product. Defining miss and false alarm probabilities for each sensor as follows:

$$P_{Mi}(t_i) \equiv \Pr(u_i = 0|\mathbf{H}_1; t_i) \quad (23)$$

$$P_{Fi}(t_i) \equiv \Pr(u_i = 1|\mathbf{H}_0; t_i) \quad (24)$$

and taking the logarithm of the likelihood-ratio test, we can write this optimum test as

$$\sum_{i=1}^{N_s} \left[u_i \log \left(\frac{1 - P_{Mi}(t_i)}{P_{Fi}(t_i)} \right) + (1 - u_i) \log \left(\frac{P_{Mi}(t_i)}{1 - P_{Fi}(t_i)} \right) \right] \underset{u_0=0}{\overset{u_0=1}{>}} \eta. \quad (25)$$

It is easy to see that the logarithms in this expression have a large magnitude for reliable sensors and, hence, the optimal decision fusion rule makes fused decisions that are dominated by the reliable decisions.

Since the Bayes risk can also be written in terms of the a posteriori probabilities $\Pr(H_k|\mathbf{U})$, we can use another approach to derive the fusion rule. A simple (but not necessarily optimal) approach is to select u_0 such that the a posteriori probability $\Pr(H_k|\mathbf{U})$ is maximized [17, § 6.3]. We write the a posteriori probability as follows:

$$\Pr(H_k|\mathbf{U}; \mathbf{T}) = \left[\prod_{i=1}^{N_s} \Pr(u_i|\mathbf{H}_k; t_i) \right] \Pr(H_k) / \Pr(\mathbf{U}; \mathbf{T}), \quad (26)$$

where we have invoked the conditional independence of the sensors. This result can be evaluated using known values for $\Pr(u_i|\mathbf{H}_k; t_i)$, which are specified by $P_{Fi}(t_i)$ and $P_{Mi}(t_i)$.

1. Some physical phenomena (e.g., a small depression in the soil) can have an effect on several sensors and, hence, this assumption can be difficult to justify. Nonetheless, it is commonly used in fusion studies.

Note that in both formulations $\Pr(H_k|\mathbf{U})$ depends implicitly on the thresholds $\mathbf{T}$. Optimal fusion requires that we select $\mathbf{T}$ to produce the best performance. For modest numbers of sensors, a search over all combinations of $\mathbf{T}$ is possible, but, in general, the search may be impractical. In such cases, it is more attractive to select a value of t_i that produces “good” sensor performance and to use that threshold for the calculations described above.

3.3 Soft Decision-Level Fusion

Consider the situation in which sensor i reaches decision u_i on the basis of features Θ_i computed from its data using a deterministic transformation, e.g., $\mathbf{F}_i(\mathbf{D}_i) = \Theta_i$. The hard decision-level fusion rules in (25) and (26) use global sensor reliability information in the form of $\Pr(u_i|\mathbf{H}_k; t_i)$ and the prior probabilities $\Pr(H_k)$, but they do not utilize information about the confidence each sensor places in its individual decisions. Such information is quantified by $\Pr(u_i|\Theta_i; t_i)$. When the sensors output this local confidence information, the outputs are equivalent to “soft” decisions and we can perform fusion using an approach similar to that described above for feature-level fusion. Treating the $\mathbf{U}$ as features returned by each sensor, a discrete analog of (7) yields

$$\Pr(H_k|\Theta) = \sum_{u_1=1}^Q \dots \sum_{u_{N_s}=1}^Q \Pr(H_k|\mathbf{U}) \Pr(u_1|\Theta_1; t_1) \Pr(u_2|\Theta_2; t_2) \dots \Pr(u_{N_s}|\Theta_{N_s}; t_{N_s}), \quad (27)$$

where we have assumed that the fused hypothesis H_k depends only on the individual sensor decisions u_i . The quantity $\Pr(H_k|\mathbf{U})$ is computed as described in the hard decision case.

4 EXAMPLE RESULTS: FUSION OF EXPERIMENTAL DATA

To test the foregoing algorithms, a sensor suite was used to acquire experimental data over a surrogate mine field, and the resulting data were fused using feature-level and decision-level techniques. In this section, we describe the sensors, the mine field, and the fused system performance.

4.1 Sensors

The sensor suite used in this work comprised a commercial electromagnetic induction (EMI) sensor, a ground penetrating radar (GPR), and a commercial infrared (IR) camera. The sensors and the signal processing used in feature extraction are documented below.

4.1.1 Electromagnetic Induction (EMI) Sensor

Virtually all land mines in use today contain some metal, which may be detected by a metal detector. The quantity of metal used varies widely from several kilograms to a few hundred milligrams. Excessive false alarm rates arise when detecting mines with small metal content because of the large amount of metallic clutter of anthropic origin found in most minefields.

The EMI sensor used in this work was a commercially available pulsed-induction metal detector, the Schiebel AN-19/2 (Schiebel Instruments, Incorporated, Washington,

D.C.) Internal signals within the device were sampled with a digital oscilloscope. The sensor was fitted to a linear scanner, which was used to acquire 56 samples over a 55 inch scan. Scans were performed at regularly spaced intervals (nominally every six inches) over the region of interest to obtain a grid of sample points.

Signal processing for the EMI sensor is a multistep process. The sensor response acquired over soil known to be metal-free, is subtracted from the waveforms. The integral of the resulting difference signal is an effective detector of buried metal. The integrated values acquired along a scan form spatial signatures. For targets not too near the sensor head and not too large compared to the loop radius, the sensor's spatial response has a monopolar shape that is reasonably well approximated by the function

$$G_1(\mathbf{R}; \Theta_1) = B + \frac{S}{1 + (|\mathbf{R} - \mathbf{R}_0|/a)^b}, \quad (28)$$

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. A more sophisticated physics-based model of the response of a buried sphere has been presented by Das et al. [18] and could be used as an alternative model, but the simple empirical form above was adequate to approximate the limited set of targets considered here.

The second step in EMI signal processing involves the estimation of the feature vector Θ_1 given by $\Theta_1 = [a \ b \ S \ E \ \mathbf{R}_0]^T$. Note that the background level B is not used in Θ_1 and we have added the quantity $E = |\mathbf{D} - \mathbf{G}|^2/|\mathbf{D}|^2$, which is the residual error in the estimate normalized to the signal energy. To compute Θ_1 at each sample position $\mathbf{R}$, a spatial window was defined and the data $\mathbf{D}_1$ within this window were fitted to the model using a nonlinear optimization technique. Since stronger signals tend to persist over larger distances, the size of the window was adjusted on the basis of the signal amplitude.

4.1.2 Ground Penetrating Radar

A ground penetrating radar detects the presence of reflecting boundaries in the soil. Such boundaries occur at the surface of the ground and at most buried objects. A GPR will detect both metallic and metal-free objects and, hence, it is a useful complement to the metal detector. Unfortunately, voids, stones, changes in soil density and moisture content, and many other subsurface natural objects are also detected, leading to a high false alarm rate.

The GPR used in this effort was developed at The Ohio State University (OSU), ElectroScience Laboratory (ESL) [19]. It employs a novel dielectric rod antenna that was scanned horizontally over the earth at a fixed height (nominally two inches). A network analyzer was used to measure the complex reflection coefficient at 51 frequencies between 1 and 6 GHz. Approximately 100 samples were acquired on each scan, which had a length of 55 inches. The antenna and its supporting platform were advanced in two inch intervals, resulting in a grid of samples.

Calibration of the radar was performed via a two step process. A "background" signature (the response of the radar when the antenna points into free space) was acquired and

subtracted from all data. The temporal impulse response of the radar was then determined by measuring the (background-free) field scattered from a reference target (a short cylinder) in free space. The ideal scattered field for this target is known from numerical calculations and the ratio of the measured and ideal frequency-domain responses yields the impulse response of the radar. This response is then deconvolved from the measured data using frequency-domain division. By design, the spectral response of the reference target contains no zeroes, and regularization was not found to be necessary in this work.

Processing of the GPR data begins with calibration, in which the response of the system is deconvolved from the measured data. A Fourier transform of the calibrated data is then used to produce a map of the reflection coefficient as a function of along-scan position and time (or depth). The data is contaminated by a surface-reflected waveform, which we remove using an OSU-developed technique [11] based on a local, spatially varying maximum-likelihood estimate of the ground position and reflection coefficient. The resulting data are modeled as a point scatterer plus clutter, the signature of which at a typical depth in homogeneous soil is known and defines $\mathbf{G}_2(\mathbf{R})$. Target depth has an effect on both the signature shape and time delay, but, because the range of mine burial depths is restricted, the effect of depth errors on target shape was found to be insignificant. The GPR feature vector is given by $\Theta_2 = [\rho \ d \ \mathbf{R}_0]^T$, where ρ is the peak correlation coefficient for the model signature and data, and d and $\mathbf{R}_0$ are, respectively, the depth and horizontal position indicated by the correlation peak. Although false targets are often seen in GPR scans, they tend not to persist across adjacent scans and, hence, the correlation was computed over a three-dimensional (depth, along-scan, and cross-scan) data window. GPR data tend to be highly correlated along a scan, because of returns from plane-stratified media and antenna ringing not eliminated by calibration. A modest amount of whitening was performed prior to correlation.

4.1.3 Infrared Camera

The natural heating and cooling of soil that occurs over a diurnal cycle causes thermal energy to flow into and out of the earth. The presence of a buried object changes this thermal flow, leading to temperature anomalies above the buried object at certain times in the diurnal cycle. These anomalies can be detected in soil temperature maps created by a suitable infrared (IR) camera. Since a mine's thermal properties are only weakly related to its metal content and electromagnetic properties, passive IR imaging is presumed to be a good adjunct to EMI and GPR sensors. In addition to buried mines, naturally occurring buried objects (e.g., stones), variations in soil mineral and water content and variations in solar illumination (due to shadowing by foliage, for example) also produce temperature anomalies. Our sensor was a commercially available MWIR camera, the IRRIS 160ST, (Cincinnati Electronics, Cincinnati, OH).

The IR data was processed by first remapping the imagery to ground coordinates using a standard perspective transformation [20], followed by bilinear interpolation of the distorted pixels and resampling to a uniform pixel size. An image chip was extracted at each interrogation point $\mathbf{R}$.

Fig. 5. Plan view of the surrogate mine and clutter test grid. The surface of the mine grid is bare soil. The surrounding area is covered with grass.

Based on empirical observations, we employed a signature model identical to that used for the EMI sensor, namely,

$$G_3(\mathbf{R}; \Theta) = B + \frac{S}{1 + (|\mathbf{R} - \mathbf{R}_0|/a)^b}. \quad (29)$$

The feature vector is $\Theta_3 = [a \ b \ S \ E \ \mathbf{R}_0]^T$, where E is the normalized residual error energy.

4.2 Test Site

A surrogate mine field was created to test our sensors under conditions that approximate those found in field situations. The site location, adjacent to our laboratory, is former farm land. A total of 40 mine-like and clutter-like targets were emplaced in a 4×10 grid, as shown in Fig. 5. An abbreviation identifying the buried object and the depth in inches is shown near each object in the map. The identities of the buried objects are given in Table 2.

This site presented a surprising number of challenges to mine detection and data from it are valuable in that they illustrate problems that could be encountered in real demining operations. The target set was deliberately selected to be stressing to our sensors. In contrast to real mines, essentially all of which contain some metal, roughly half the mine surrogates were metal-free. In addition, some of the surrogate mines are quite small (only 1.5 inches in diameter) and undetectable to all of the sensors. These objects were redesignated as clutter. The site has been intermittently used

for various engineering experiments for more than 40 years, which produced a significant amount of small metallic debris. Consistent with our goal of achieving realistic environmental conditions, the test area was not cleared prior to emplacing the targets. As a result, during two years of testing, roughly 100 metallic fragments and a number of stones were located in (and, subsequently, removed from) the region. When the targets were initially emplaced, a lush grass cover was present in the area, which was subsequently removed to improve the performance of the IR camera. The soil contains a large amount of clay, which has a negative effect on GPR performance. The surface of the mine field contains a number of topological irregularities but it is planar to within an estimated variation of ± 2 inches. The objects in this field had been in situ for more than 18 months at the time these data were acquired.

Data were acquired along the paths shown in Fig. 6. The small circles are the target grid points shown previously. The GPR samples were acquired at ~ 0.5 inch intervals along north-south (N-S) lines spaced every two inches east-west (E-W). The EMI sensor acquired samples at ~ 1 inch intervals along N-S lines spaced six inches E-W. The IR camera was positioned on the roof of a two-story building adjacent to the test site. The camera's field of view did not permit us to sample the entire mine field with an acceptable pixel density and a sequence of eight images was necessary as indicated by the trapezoids. Styrofoam fiducial markers

TABLE 2
Description of the Objects Buried in the Mine Test Grid

Abbr.	Object	Dia. (in.)	Ht. (in.)	Metal Content?	Identity	Quantity
CC	Soda can	2.5	5	Y	clutter	4
CX	Crushed soda can*	3	5	Y	clutter	1
PT	Pop top*	0.5	1	Y	clutter	2
M35	Aluminum disk	3.5	5/8	Y	surrogate	1
P15	Plexiglas disk	1.5	2	N	clutter	4
R4	Rock	4	2	N	clutter	2
V4	Void (Styrofoam)	4	2	N	clutter	2
H4	Refilled hole	6	2	N	clutter	2
N35	Nylon disk	3.5	5/8	N	surrogate	2
N3P	Nylon disk with steel pin	3	1	Y	surrogate	4
T7	Teflon disk	7	1	N	surrogate	2
N7	Nylon disk	7	1	N	surrogate	2
B8	Concrete pylon*	8	8	N	clutter	4
-	Blank	-	-	?	clutter	12

For the planar objects indicated by an asterisk, dimensions are length and width.

were placed on the region to permit later image warping and registration, which are indicated by squares in Fig. 6.

4.3 Performance of Individual Sensors

For each sensor i , feature vectors Θ_i were extracted using the signal processing algorithms discussed above and supplied to sensor-specific classifiers $\Pr(u_i|\Theta_i, t_i)$. Results for the binary detection problem involving hypotheses $H_1 = \text{"nonmine"}$ and $H_2 = \text{"mine"}$ are reported here. The prior probabilities for this site are $\Pr(H_0) = 28/40$ and $\Pr(H_1) = 12/40$. It is important to note that the number of mine surrogates being sampled is small. The results obtained must be interpreted in light of this fact and, hence, it is difficult to draw general conclusions.

To compute P_D and P_F , the test area was divided into a large number of square cells (280). A cell was considered to contain a detection if at least four adjacent detections were present. We computed the false alarm probability as the number of false detections divided by the total number of cells not occupied by mines.

To permit comparison between sensors, we have taken all intermediate sensor-specific hypotheses h_{ki} to be H_k , even though some surrogates are undetectable to some sensors. It is known that with sufficient training data

backpropagation neural networks will approximate an a posteriori density [21], and the outputs of such a network (normalized to a unit sum) were used to approximate the required a posteriori probabilities. The number of input and output nodes for these networks were defined by the sizes of Θ_i and K , respectively. The number of hidden nodes were 3, 7, and 6 for the GPR, EMI, and IR sensors, respectively. Target detections u_i were produced by thresholding the network outputs at level t_i . Leave-one-out training and validation was used in all of the results. Because networks trained with random weight initializations may converge to spurious local error minima, several trials were conducted and the best results were retained.

Figs. 7, 8, and 9 present binary detection maps for each sensor, which were generated by choosing appropriate detection thresholds t_i and discarding regions with small number of adjacent detections. The dark masses on these maps indicate detections, while circles represent the locations of the declared mine surrogates. Detection cells

Fig. 6. Locations of sample points for the sensor suite.

Fig. 7. Detection map obtained using EMI sensor.

Fig. 8. Detection map obtained using IR sensor.

Fig. 9. Detection map obtained using the GPR sensor.

for the targets are indicated by squares. It is clear from these results that the EMI and IR sensors have a high false alarm rate when all targets are detected. The poor performance of the EMI sensor was anticipated because most of the mine surrogates in our test grid are nonmetallic and the site contains a large quantity of metallic clutter. The poor performance of the IR sensor was not anticipated, but may be due to the relatively large number of surface irregularities. For the threshold chosen here, the GPR sensor shows reasonably good performance, with only a single missed detection and a moderate false alarm rate. Detection of the missed mine, however, is problematic. The performance of these sensors can be compared quantitatively using receiver operating characteristic (ROC) curves, as shown in Fig. 10. These results confirm the findings shown in the detection maps. In particular, the inability of the GPR to detect all mines at a reasonable false alarm rate is clear.

4.4 Performance of Fusion Algorithms

For decision-level fusion, we used the “good threshold” approximation described in Section 3.2 to avoid the search over all possible sensor thresholds t_i implied in (25). The EMI and IR sensor thresholds were selected such that all expected targets (i.e., targets with metal for the EMI sensor and targets that present a sizable thermal discontinuity for the IR sensor) were detected while minimizing clutter detections. The criterion used for selecting the GPR threshold was the “knee” in the ROC curve in Fig. 10. The resulting individual sensor P_{M_i} and P_{F_i} values are shown in

Fig. 10. Receiver operating characteristic (ROC) curves of individual sensors.

TABLE 3
Individual Sensor P_F and P_M Using “Good” Thresholds

Sensor	P_{F_i}	P_{M_i}
GPR	0.0929	0.0909
EMI	0.3755	0.4545
IR	0.5613	0.0909

Table 3. The detections u_i for these thresholds comprise the inputs u_i for the decision-level fusion algorithms and are shown in the foregoing individual sensor detection maps. The local confidences $\Pr(u_i|\Theta_i, t_i)$ required for the soft-decision fusion scheme of (27) were set to the a posteriori probability estimates produced by the corresponding classifier output. The ROC curve for decision-level fusion appears in Fig. 11.

Comparing Figs. 10 and 11, we see that hard-decision fusion based on (25) offers little benefit over the GPR sensor alone. Based on the form of (25), which most strongly weights the most reliable sensors, this result was expected. The performance of soft-decision fusion, however, is somewhat better. It is clear that with soft-decision fusion, $P_D = 1$ can be achieved with a smaller P_F , which is a significant improvement over the best individual sensor performance. The detection map for soft-decision fusion is shown in Fig. 12. The detection map for hard-fusion is not shown, because it is nearly identical to the GPR detection map.

Fig. 11. Receiver operating characteristic (ROC) curves for three types of fusion.

Fig. 12. Detection map obtained using soft-decision fusion.

Fig. 13. Detection map obtained using feature-level fusion.

For feature-level fusion, the sensor feature vectors Θ_i were concatenated and used in a single classifier (cf. (8)), resulting in the ROC curve shown in Fig. 11. As done for the individual sensors, the classifier $\Pr(H_k|\hat{\Theta})$ was approximated by a backpropagation neural network. A total of 14 hidden nodes were used in this network.

Using as a figure of merit, the smallest P_F at which $P_D = 1$, we see that the feature-fused system is significantly better than the individual sensors and marginally better than the soft-decision fusion case. The detection map produced by the feature-fused system, corresponding to the largest threshold that can achieve $P_D = 1$, is given in Fig. 13. Comparing the soft-decision case in Fig. 12 and the feature-level case in Fig. 13, the reduction in false alarm rate indicated by the ROC curve (roughly 40 percent) is evident.

5 SUMMARY AND CONCLUSIONS

We have developed a feature-level fusion technique for noncommensurate sensor data sampled at noncoincident locations. The features used for fusion were extracted from spatial signatures of targets by applying nonlinear optimization techniques to data collected in a neighborhood of the location of interest. The maximum a posteriori approximation was used to simplify the exact fusion formulation to a form that was practical to implement. For the special case of normal clutter and normal linear features, fusion can be expressed as the Bayesian linear model for which a compact closed form expression exists. This result was used to demonstrate the importance of accurate sensor position information. We also discussed several decision-level fusion schemes that can be used for these data.

We tested the fusion algorithms on noncommensurate noncoincidently sampled data acquired by EMI, GPR, and IR sensors over a test field containing mine surrogates and both deliberate and unknown clutter. The targets and the clutter present in the test site were both very challenging. Although the number of mine surrogates sampled was small, the results presented here indicate that the performance of feature-level fusion and soft-decision fusion can be significantly better than that of the individual sensors. This performance enhancement is encouraging, especially when we recognize that the IR and EMI sensors were only marginally effective due to both the heavy clutter at the test site and our selection of surrogate targets. As expected, hard-decision fusion showed little improvement over the performance of the best individual sensor since the algorithm most strongly weights the most effective sensors.

Overall, we have shown that feature-level and soft-decision level fusion of noncoincidently sampled EMI, IR, and GPR data can result in significant improvements in mine detection performance. We found a small additional benefit in using feature-level fusion, which provided a 40 percent reduction in the false alarm rate beyond the soft-decision fusion case. The feature-level fusion algorithm also provides an explicit method for dealing with position errors, which is not available through the decision-level fusion algorithms. The computational costs of decision-level and feature-level fusion algorithms are approximately equal since the computational expense involved in fusion is dominated by features estimation, and all methods of detection must perform this processing. There is, however, a small penalty for feature-level fusion since the classifier must be somewhat more complex to account for the greater number of inputs.

These results are encouraging but the small sample size of our data set suggests that a larger study be done to confirm these results. A data collection has been done at a larger minefield operated by the US Army at Fort A.P. Hill, Virginia. That study is now complete, and the results will be reported in the future.

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. A preliminary version of this paper was presented at the 1999 SPIE AeroSense Conference, Orlando, Florida.

REFERENCES

- [1] P.K. Varshney, *Distributed Detection and Data Fusion*. New York: Springer-Verlag, 1997.
- [2] B.V. Dasarthy, *Decision Fusion*. Los Alamitos, Calif.: IEEE CS, 1994.
- [3] S. Chauduri, A. Crandall, and D. Reidy, "Multisensor Data Fusion for Mine Detection," *Proc. SPIE: Sensor Fusion III*, R.C. Harney, ed., vol. 1306, pp. 187-203, 1990.
- [4] B. Brusmark, A. Christiansen, P. Jägerbro, and A. Lauberts, "Combination of Ground Penetrating Radar and Metal Detector Data for Mine Detection," *Proc. Seventh Int'l Conf. Ground Penetrating Radar*, 1998.
- [5] G.A. Clark, S.K. Sengupta, R.J. Sherwood, J.E. Hernandez, M.R. Buhl, P.C. Schaich, R.J. Kane, M.J. Barth, and N.K. Del Grande, "Sensor Feature Fusion for Detecting Buried Objects," *Proc. SPIE: Underground and Obscured-Object Imaging and Detection*, N.K. Del Grande, I. Cindrich, and P.B. Johnson, eds., vol. 1942, pp. 178-188, 1993.
- [6] X. Miao, M.R. Azimi-Sadjadi, B. Tian, A.C. Dubey, and N.H. Witherspoon, "Detection of Mines and Minelike Targets Using Principal Component and Neural Methods," *IEEE Trans. Neural Networks*, vol. 9, no. 3, pp. 454-463, May 1998.
- [7] R.A. Weisenseel, W.C. Karl, D.A. Castanon, E. Miller, C. Rappaport, and C.A. Dimarzio, "Statistical Fusion of GPR and EMI Data," *Proc. SPIE: Detection and Remediation Technologies for Mines and Minelike Targets IV*, A.C. Dubey, J.F. Harvey, J.T. Broach, and R.E. Dugan, eds., vol. 3710, pp. 1179-1187, 1999.
- [8] E. den Breejen, K. Schutte, and F. Cremer, "Sensor Fusion for Anti-Personnel Landmine Detection, a Case Study," *Proc. SPIE: Detection and Remediation Technologies for Mines and Minelike Targets IV*, A.C. Dubey, J.F. Harvey, J.T. Broach, and R.E. Dugan, Eds., vol. 3710, pp. 1235-1245, 1999.

- [9] P.G. Johnson and P. Howard, "Performance Results of the EG&G Vehicle Mounted Mine Detector," *Proc. SPIE: Detection and Remediation Technologies for Mines and Minelike Targets IV*, A.C. Dubey, J.F. Harvey, J.T. Broach, and R.E. Dugan, eds., vol. 3710, pp. 1149-1159, 1999.
- [10] J.E. McFee, V.C. Aitken, R. Chesney, Y. Das, and K.L. Russell, "Multisensor Vehicle-Mounted Teleoperated Mine Detector with Data Fusion," *Proc. SPIE: Detection and Remediation Technologies for Mines and Mine-Like Targets*, A.C. Dubey, J.F. Harvey, and J. Broach, eds., vol. 3392, pp. 1082-1093, 1998.
- [11] A. Gunatilaka and B.A. Baertlein, "Comparison of Pre-Detection and Post-Detection Fusion for Mine Detection," *Proc. SPIE: Detection and Remediation Technologies for Mines and Mine-Like Targets*, A.C. Dubey, J.F. Harvey, J.T. Broach, and R.E. Dugan, eds., vol. 3710, pp. 1212-1223, 1999.
- [12] A. Gunatilaka and B.A. Baertlein, "A Subspace Decomposition Technique to Improve GPR Imaging of Anti-Personnel Mines," *Proc. SPIE: Detection and Remediation Technologies for Mines and Minelike Targets V*, A.C. Dubey, J.F. Harvey, J.T. Broach, and R.E. Dugan, eds., vol. 4038, 2000.
- [13] B.A. Baertlein, "Neural-Network Based Multi-Sensor Fusion for Mine Detection," Technical Report N61331-93-C-0050, Ballena Systems Corporation, 1994.
- [14] J.A. Marble, J.G. Ackenhusen, J.W. Wegrzyn, J. Mancuso, and C. Dwan, "Sensor Fusion Performance Gain for Buried Mine/uxo Detection Using GPR, EMI, and MAG Sensors," *Proc. SPIE: Detection and Remediation Technologies for Mines and Mine-Like Targets V*, A.C. Dubey, J.F. Harvey, J.T. Broach, and R.E. Dugan, eds., vol. 4038, pp. 1473-1484, Apr. 2000.
- [15] H.V. Poor, *An Introduction to Signal Detection and Estimation*, second ed. New York: Springer-Verlag, 1994.
- [16] S.M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory*. Englewood Cliffs, N.J.: Prentice-Hall, 1993.
- [17] D.L. Hall, *Mathematical Techniques in Multisensor Data Fusion*. Boston: Artech House, 1992.
- [18] Y. Das, J.E. McFee, J. Toews, and G.C. Stuart, "Analysis of an Electromagnetic Induction Detector for Real-Time Location of Buried Objects," *IEEE Trans. Geoscience Remote Sensing*, vol. 28, no. 3, pp. 278-287, May 1990.
- [19] S. Nag, L. Peters Jr., I.J. Gupta, and C.C. Chen, "Ramp Response for the Detection of Anti-Personnel Mines," *Proc. SPIE: Detection and Remediation Technologies for Mines and Minelike Targets IV*, A.C. Dubey, J.F. Harvey, J.T. Broach, and R.E. Dugan, eds., vol. 3710, pp. 1313-1322, 1999.
- [20] I.K. Sendur and B.A. Baertlein, "Techniques for Improving Buried Mine Detection in Thermal IR Imagery," *Proc. SPIE: Detection and Remediation Technologies for Mines and Minelike Targets IV*, A.C. Dubey, J.F. Harvey, J.T. Broach, and R.E. Dugan, eds., vol. 3710, pp. 1272-1283, 1999.
- [21] M.D. Richard and R.P. Lippmann, "Neural Network Classifiers Estimate Bayesian a Posteriori Probabilities," *Neural Computation*, vol. 3, pp. 461-483, 1991.

from 1994 to 2000. He is currently a member of technical staff at Lucent Technologies in Columbus, Ohio. He is a member of the IEEE.

Brian A. Baertlein received the PhD in electrical engineering from the University of Arizona in 1988. His professional career has included analyses of sensor systems of various types, development of detection and estimation algorithms, and work in sensor fusion, scattering and propagation phenomena, antennas, and electromagnetic compatibility. He is currently a research scientist and adjunct associate professor of electrical engineering at The Ohio State University (OSU) Electro-Science Laboratory (ESL). Before joining OSU, he was a senior scientist with several small businesses doing work for the US Department of Defense and the Department of Energy. He is a member of the IEEE.

► For further information on this or any computing topic, please visit our Digital Library at <http://computer.org/publications/dlib>.