Information-based sensor management for landmine detection using electromagnetic induction, ground-penetrating radar, and seismic sensors

Mark Kolba[a], Peter Torrione[a], Waymond Scott[b] and Leslie Collins[a]

[a]Duke University, Department of Electrical and Computer Engineering, Box 90291, Durham, NC 27708, USA
[b]Georgia Institute of Technology, School of Electrical and Computer Engineering, 777 Atlantic Dr., Atlanta, GA 30332-0250, USA
lcollins@ee.duke.edu

An information-based sensor management framework is discussed that enables the automated tasking of a suite of sensors when detecting static targets. The sensor manager chooses the sensors to use and the grid-based locations to observe in order to maximize the expected information gain that will be obtained with each new sensor observation. Initially, sensor probabilities of detection and false alarm, P_d and P_f, are assumed to be known by the sensor manager. In a field setting, however, P_d and P_f cannot be known exactly, and so uncertainty modeling for P_d and P_f is also discussed. The sensor manager is tested on real landmine data using electromagnetic induction (EMI), ground-penetrating radar (GPR), and seismic sensors. A matched subspace detector is used to process the EMI data, an adaptive pre-screening algorithm based on the least mean squares (LMS) adaptive filter is used to process the GPR data, and whitening followed by an energy detector is used to process the seismic data. The sensor manager is able to detect the landmines more quickly and more effectively than an unmanaged, blind-search approach. Using all three sensor modalities also results in superior detection performance to that achieved by only a single sensing modality.

1 Introduction

In recent years, many remote sensing applications have seen a dramatic increase in the number and quality of available sensors. A human operator can be overwhelmed by the number of sensor tasking decisions that need to be made when many sensors are present or when many conflicting mission goals must be best fulfilled given a constrained amount of resources. Furthermore, in many military applications, a human operator is placed in harm's way when operating a suite of sensors. Automated sensor management techniques have therefore been explored for the purpose of assisting or replacing a human operator in a complex or dangerous operational environment [1-3].

In previous work, an information-based sensor management framework has been presented that directs a suite of sensors in a search for static targets within a grid of cells [4, 5]. The sensor manager uses the Kullback-Leibler divergence as an information measure and functions by tasking the sensors to make the new observation that will produce the largest expected information gain. When multiple sensing modalities are considered, the sensor manager determines whether it is beneficial to observe with all of the available sensing modalities or only with some of them. This paper briefly reviews the sensor management framework. In the following development, the sensor probabilities of detection and false alarm, P_d and P_f, will first be assumed to be known [5]. In a field setting, however, P_d and P_f cannot be known exactly, and so uncertainty modeling will also be discussed [4].

The sensor manager will be tested using a set of real landmine detection data that has been collected with three sensors: an electromagnetic induction (EMI) sensor, a ground penetrating radar (GPR) sensor, and a seismic sensor. The data from each of these sensing modalities is processed in order to obtain decision statistic outputs that are then operated on by the sensor manager. A matched subspace detector is used to process the EMI data [6], an adaptive pre-screening algorithm based on the least mean squares (LMS) adaptive filter is used to process the GPR data [7], and whitening followed by an energy detector is used to process the seismic data.

The remainder of the paper is organized as follows. Section 2 reviews the information-based sensor management framework and also discusses the use of uncertainty modeling for modeling unknown sensor P_d and P_f. Section 3 discusses the signal processing algorithms that were used to process the data obtained from each of the individual sensing modalities. Section 4 will present simulation results from using the sensor manager on the real landmine data. It will be demonstrated that the sensor manager outperforms an unmanaged, blind sweep technique and is able to detect the landmines more quickly. It will be further demonstrated that the use of uncertainty modeling is able to improve the sensor manager performance over that obtained when the sensor P_d and P_f are modeled as certain. Finally, Section 5 offer conclusions and a summary discussion.

2 Sensor management framework

2.1 Sensor manager

This section will review the sensor management framework presented in [4, 5]. The sensor management framework uses M sensor platforms, each with D sensing modalities, to search for N targets in a grid of C cells. Each of the grid cells has a binary state, either containing or not containing a target, which is denoted $S_c = 0$ or $S_c = 1$, respectively. The state probabilities are initialized with a spatially uniform distribution so that $P(S_c = 1) = N/C$ and $P(S_c = 0) = (C - N)/C$ for each cell. Observation k in cell c is written $x_{c,k}$; it may also be written as $x_{c,k,m,d}$, where m and d denote the sensor platform and sensing modality used to make observation k. A sequence of observations $x_{c,1}, x_{c,2}, \ldots, x_{c,k}$ is written $X_{c,k}$. The sensors make binary observations—either "target present" or "no target present"—with a known probability of detection and false alarm that is specific to the particular sensor making the observation:

$$\begin{aligned} P\left(x_{c,k,m,d} = 1 \middle| S_c = 1\right) &= P_{d,m,d} \\ P\left(x_{c,k,m,d} = 0 \middle| S_c = 1\right) &= 1 - P_{d,m,d} \\ P\left(x_{c,k,m,d} = 1 \middle| S_c = 0\right) &= P_{f,m,d} \\ P\left(x_{c,k,m,d} = 0 \middle| S_c = 0\right) &= 1 - P_{f,m,d} \end{aligned} \quad (1)$$

After a sensor observation is made, the state probability of the observed cell is updated using Bayes's rule:

$$P\left(S_c = s \middle| X_{c,k}\right) = \frac{P\left(x_{c,k,m,d} \middle| S_c = s\right) P\left(S_c = s \middle| X_{c,k-1}\right)}{\sum_{j=0}^{1} P\left(x_{c,k,m,d} \middle| S_c = j\right) P\left(S_c = j \middle| X_{c,k-1}\right)} . \quad (2)$$

The measure of information used by the sensor manager is the Kullback-Leibler divergence, which is defined for probability mass functions p and q as

$$D_{KL}\left(p\|q\right)=\sum_i p_i \log\left(\frac{p_i}{q_i}\right). \qquad (3)$$

The expected Kullback-Leibler divergence after one additional sensor observation may in fact be computed analytically, without requiring the sensor to physically make the observation, yielding

$$E\left[D_{KL}\left(P_{c,k+1}\|Q_c\right)|\, c,m,d\right]= \sum_{j=0}^{1} D_{KL}\left(P_{c,k+1}\|Q_c\right)P\left(x_{c,k+1,m,d}=j\,|X_{c,k}\right) \qquad (4)$$

where $P_{c,k}$ represents the state probabilities in cell c after k observations and Q_c represents the prior state probabilities in cell c. The expected information gain obtained with a new sensor observation is then given straightforwardly as

$$\Delta D_{KL}\left(c,m,d\right)=E\left[D_{KL}\left(P_{c,k+1}\|Q_c\right)|\, c,m,d\right] - D_{KL}\left(P_{c,k}\|Q_c\right). \qquad (5)$$

The notion of sensor cost of use was introduced in [5]; in the formulation of the sensor manager utilized in this paper, it will be assumed that the cost of use for each sensing modality is one. The sensor platforms will then be tasked in the following way. Each sensor platform will move to the cell that will produce the largest expected information gain for the first modality for that sensor platform. Upon entering the cell, the first modality will be used to make an observation. Subsequent modalities d' will then be used to make an observation if they satisfy

$$\Delta D_{KL}\left(c,m,d'\right)>\Delta D_{KL}\left(c',m,1\right), \qquad (6)$$

where c' is the cell that will produce the largest expected information gain when observed with the first modality.

2.2 Uncertainty modeling

In a real-world sensing environment, the sensor probabilities of detection and false alarm shown in Eq.(1) are not certain quantities; they will in fact vary from day to day and from location to location, both locally and globally. It is therefore of interest to model the sensor probabilities of detection and false alarm as being uncertain. Uncertainty modeling for the sensor manager has been presented in [4] and will be reviewed here.

The beta distribution, which is the natural conjugate prior distribution for a binomial process (the process generating the observed data), is used to model uncertain probabilities of detection and false alarm. The beta distribution is parameterized by r and k, which may be thought of as the number of successes and the number of trials, respectively, in a binomial process. The beta distribution is defined as

$$f_\beta\left(z\right)=\frac{\Gamma\left(k\right)}{\Gamma\left(r\right)\Gamma\left(k-r\right)}z^{r-1}\left(1-z\right)^{k-r-1}. \qquad (7)$$

With P_d and P_f described by densities, integration must be performed to determine the probability that a sensor will make a specific new observation:

$$P\left(x_{c,k,m,d}\,|S_c\right)=\iint P\left(x_{c,k,m,d}\,|S_c,P_d,P_f\right)\cdot f\left(P_d,P_f\right)dP_d dP_f \qquad (8)$$

A derivation, detailed in [4], then shows that

$$\begin{aligned} P\left(x_{c,k,m,d}=1|S_c=1\right)&=E_{P_d}\left\{f_\beta\left(P_{d,m,d}\right)\right\}\\ P\left(x_{c,k,m,d}=0|S_c=1\right)&=1-E_{P_d}\left\{f_\beta\left(P_{d,m,d}\right)\right\}\\ P\left(x_{c,k,m,d}=1|S_c=0\right)&=E_{P_f}\left\{f_\beta\left(P_{f,m,d}\right)\right\}\\ P\left(x_{c,k,m,d}=0|S_c=0\right)&=1-E_{P_f}\left\{f_\beta\left(P_{f,m,d}\right)\right\} \end{aligned} \qquad (9)$$

Uncertainty modeling changes the computation of the state probabilities and expected Kullback-Leibler divergence from Section 2.1; the observation probabilities given in Eq.(9) are used in both Eq.(2) and Eq.(4) when uncertainty is being modeled. Since the beta distribution is the natural conjugate prior distribution for a binomial process, posterior densities for P_d and P_f are also beta distributions. For a beta distribution with prior parameters r' and k' and observed data with r successes (that is, "target present" observations) out of k observations, the parameters of the posterior beta distribution after collecting the observed data will be r'' and k'', with $r'' = r' + r$ and $k'' = k' + k$.

3 Signal processing algorithms

The sensor observations for the sensor management framework presented in Section 2 may be considered to be thresholded decision statistic outputs from an arbitrary signal processing algorithm. The sensor manager optimizes sensor tasking at the decision level, not at the data level. In other words, the sensor manager is not attempting to optimize where the next piece of raw data should be observed in order to maximize detection performance for a specific signal processing strategy. Instead, the sensor manager assumes that a sensor, when making an observation, will collect data over a point or a small region and make a decision about whether a target is present or not in that observed region. The specific signal processing algorithms used may be designed completely independently from the sensor management framework. As mentioned previously, the sensor manager will be tested using a set of real landmine data that was collected with three sensing modalities: an EMI, a GPR, and a seismic sensor. The signal processing algorithms used to process data from each of these three sensors will now be detailed.

3.1 EMI sensor

The EMI sensor used to collect data has dipole transmit and receive coils, and an auxiliary bucking transformer is used to cancel the mutual coupling between the transmitter and receiver [8]. Modeling of wideband EMI responses [9] motivates the use of a matched subspace detector to process the received EMI data [6]. The matched subspace detector functions in the following way. Each type of landmine has a signature template that is learned from training data and is defined to span the subspace $\langle \boldsymbol{H} \rangle$. It is assumed that received EMI data, $\boldsymbol{x}$, is a scaled version of the template signature that has been corrupted by additive Gaussian noise. Under these conditions, the decision statistic d, given by

$$d=\frac{\boldsymbol{x}^T\boldsymbol{P_H}\boldsymbol{x}}{\boldsymbol{x}^T\boldsymbol{x}}, \qquad (10)$$

is the optimal detector. The matrix $\boldsymbol{P_H}$ in Eq.(10) is the projection matrix onto the subspace $\langle \boldsymbol{H} \rangle$.

Each type of landmine has a differently shaped EMI signature. Since the data under consideration contains several types of landmines, a template—and correspondingly, a subspace—is created for each type of landmine that might be present in the data. Contingent on an energy-based template-specific pre-screening process, the received EMI data $\boldsymbol{x}$ is then processed through each of the matched subspace detectors, and the maximum of the individual matched subspace detector outputs is selected as the observed decision statistic. The maximum matched subspace detector output is chosen rather than the sum of the matched subspace detector outputs in order to eliminate false alarms that occur when a clutter signature partially matches several of the different mine templates; this procedure is the equivalent of maximum a posteriori classification [6].

3.2 GPR sensor

The GPR used to collect data uses resistively loaded vee-dipole antennas that are well-suited to the landmine detection problem [10]. GPR data collected with the sensor is processed using an adaptive technique based on the least mean square (LMS) adaptive filter [7]. Before processing with the LMS adaptive filter, the raw data is pre-processed in a number of ways. The data is first aligned at the location of the ground bounce, and the ground bounce is subsequently removed by time-clipping the data. A median filter helps remove any interference that might be present in the data. Finally, the data is depth-segmented so that processing may be performed on individual depth segments in order to mitigate the effects of energy attenuation due to propagation of the electromagnetic pulse through the ground.

After the pre-processing steps have been completed, a two-dimensional version of the LMS adaptive filter is applied to each depth segment of the data. The filter is used to estimate the response at each individual pixel, beginning with the first down-track pixel and moving in the down-track direction. The LMS adaptive filter maintains a vector of weights, $\boldsymbol{w}$, which are multiplied by an input signal, $\boldsymbol{u}$, to produce an estimate, y, of the response that will result from input signal $\boldsymbol{u}$:

$$y_n = \boldsymbol{w}_n^T \boldsymbol{u}_n , \tag{11}$$

where n is the time index. The estimated response, y, and the desired response, d, are then used to determine the estimation error, e:

$$e_n = y_n - d_n . \tag{12}$$

Finally, the estimation error and the learning rate, μ, are used to update the weight vector:

$$\boldsymbol{w}_{n+1} = \boldsymbol{w}_n + \mu e_n \boldsymbol{u}_n . \tag{13}$$

The learning rate is a parameter that may be tuned to affect the rapidity with which the LMS weights adapt. The weight vector and input signal are defined to be a group of pixel locations both in front of and behind the location of the pixel whose intensity is to be estimated [7]. When the LMS error as computed by Eq.(12) is large, that indicates the presence of a subsurface anomaly. The output decision statistics are the sums of the LMS error energy across the various levels of depth segmentation, resulting in a cross-track and down-track matrix of decision statistics over the ground surface.

3.3 Seismic sensor

The seismic sensor used to collect data consists of an electrodynamic shaker that generates seismic waves and a separate, radar-based non-contact displacement sensor that measures the displacement created predominantly by the propagating seismic surface (Rayleigh) waves [11]. When a seismic wave encounters a mine or other buried object, the structural and elastic properties of the object will cause perturbations in the surface displacement that will be measured by the displacement sensor. The received seismic data is processed by whitening the data and then using an energy detector to produce the output decision statistics.

4 Simulation results

Now that the sensor management framework and signal processing algorithms have been presented, simulation results will be shown for testing the sensor management framework on the landmine data. The landmine data consists of a 1.8 m by 1.8 m region containing six mines and twenty-one clutter objects representing a variety of sizes of metallic and non-metallic clutter. The layout of the mines and clutter objects may be seen in Fig.1. Since the sensor manager operates on a cell grid and since the sensor manager has further been designed assuming that only one object exists per cell and that objects do not straddle cells, a nine by nine grid has been manually positioned over the data collection region so that each cell contains at most one object and so that objects straddle cells as little as possible. Each cell in this grid represents an 18 cm by 18 cm region of ground.

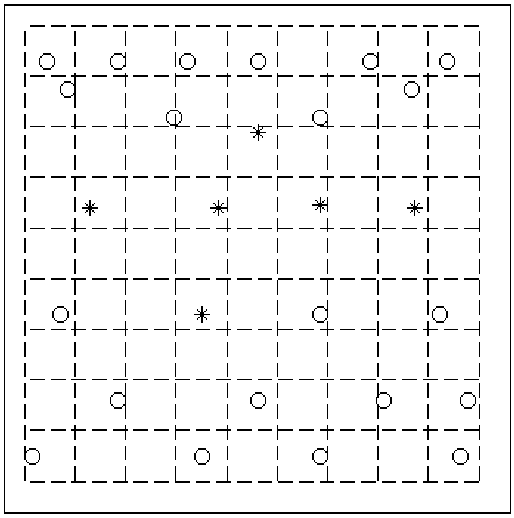

Fig.1 Nine by nine cell grid and object layout for landmine data. Targets are represented by asterisks and clutter objects by circles. The outer box represents the entire region of collected data.

In the subsequent simulations, the observed sensor data will be generated by randomly selecting one of the ten largest decision statistic outputs in the observed cell for the sensing modality making the observation. The selected decision statistic will then be thresholded to produce a binary observation of either "target present" or "no target present." Such a strategy is intended to emulate results that might be observed were a sensor to make multiple collections over the same region of interest. Thresholds are selected for the three sensing modalities so that the overall operating characteristics of the modalities are P_d = 0.850 and P_f = 0.323 for the EMI, P_d = 0.850 and P_f = 0.085 for the seismic sensor, and P_d = 0.950 and P_f = 0.056 for the GPR. The cost of use for all sensors is assumed to be one.

In each iteration of the simulation, three identical sensor platforms will move through the cell grid guided by either discrimination-directed search (the sensor manager described in this paper) or direct search, which is an unmanaged technique in which the sensors blindly sweep through the grid in a predefined pattern, making one observation in each cell with each available sensing modality as they move. The performance metric used in the results that follow is the probability of error, P_e. To produce each plot, one thousand full realizations of the simulation will be performed and the results averaged.

Fig.2 shows performance results for discrimination-directed and direct searches using different combinations of sensing modalities. The EMI, seismic, and GPR sensing modalities are denoted S1, S2, and S3 in the figure legend. Both discrimination-directed and direct search using all three sensing modalities obtain a lower probability of error than a search with only a single modality. This result reflects the performance improvements that are typically available through multimodal sensing. Furthermore, discrimination-directed search using all three modalities outperforms direct search using all three modalities. Use of the sensor manager presented in this paper allows the targets to be detected more quickly than they can be detected with a direct search technique.

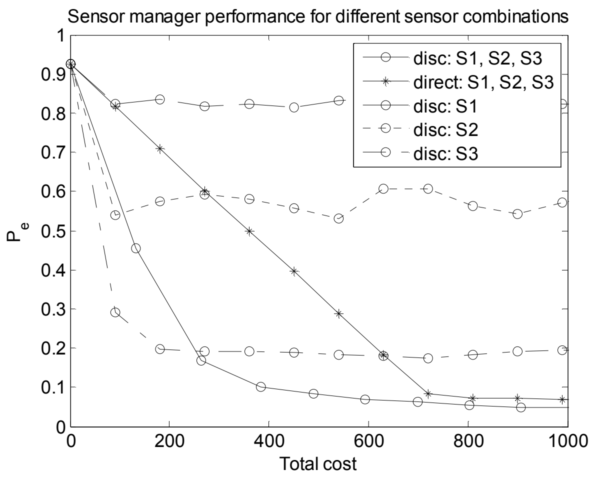

Fig.2 P_e vs. cost performance for discrimination-directed and direct searches on the landmine data using the indicated sensing modalities.

Uncertainty modeling as discussed in Section 2 is intended to model the fact that sensor P_d and P_f cannot be known a priori in a real-world environment and that P_d and P_f may furthermore vary on a local scale, from grid cell to grid cell. Simulations are now performed that incorporate uncertainty modeling into the sensor management framework. Three different sets of parameters are used to model the beta prior densities on P_d and P_f: k = 100, k = 10, and k = 5, which correspond to increasing levels of uncertainty. In each case, the value of r is set to ensure that the expected value, r/k, of the resulting beta density equals the P_d or P_f value for the certain case. Example beta densities for these three levels of uncertainty are shown in Fig.3.

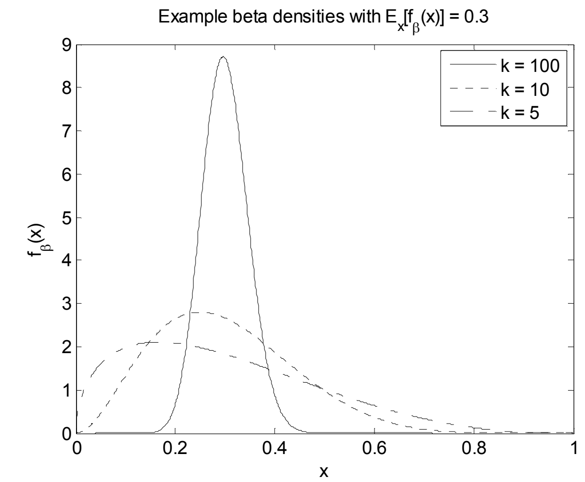

Fig.3 Example beta densities for three different levels of uncertainty. Each of the three beta density has an expected value of 0.3.

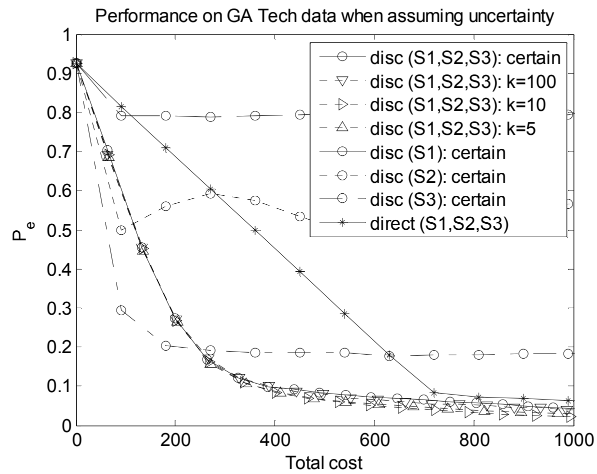

Fig.4 P_e vs. cost performance for discrimination-directed search with uncertainty modeling compared to performance for discrimination-directed and direct searches without uncertainty modeling. All uncertainty modeling curves are clustered with the certain discrimination-directed curve.

Performance results for the sensor manager using uncertainty modeling are shown in Fig.4. Notice that all of the uncertainty modeling curves cluster relatively closely with the discrimination-directed search curve without uncertainty modeling, meaning that discrimination-directed search with uncertainty modeling is outperforming direct search. In fact, the uncertainty modeling curves actually demonstrate improved performance over discrimination-directed search performance without uncertainty modeling, as may be seen more clearly in Fig.5. The smallest amount of uncertainty, k = 100, provides only a slight gain over the performance that is obtained with no uncertainty modeling. However, the two higher levels of uncertainty modeling, k = 10 and k = 5, both provide a more substantial gain in performance. After a total cost of 1000, for example, uncertainty modeling with k = 10 provides a 50% reduction in the probability of error.

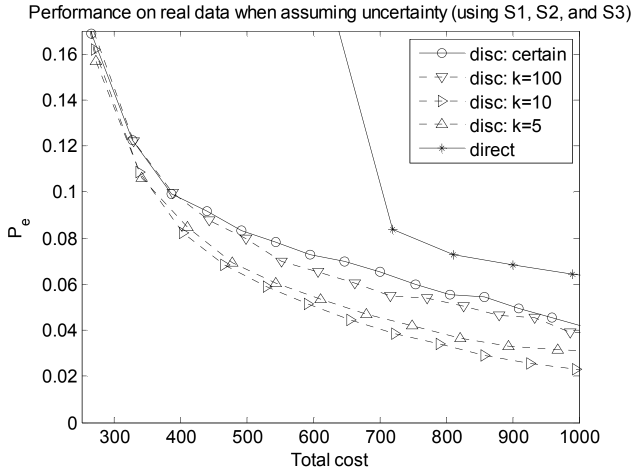

Fig.5 P_e vs. cost performance for discrimination-directed search with uncertainty modeling compared to performance for discrimination-directed and direct searches without uncertainty modeling.

Previously, the results in [4] demonstrated that sensor manager performance on uncertain simulated data will be improved by properly modeling the uncertainty that is present, and the results in Fig.4 and Fig.5 demonstrate a similar effect for real landmine data. Even though the real data does not exactly follow the modeled distributions in the way that the simulated data did in [4], the beta distributions for P_d and P_f model the uncertainty present in the real landmine data with sufficient fidelity that a noticeable performance improvement is obtained through uncertainty modeling.

5 Conclusion

This paper has reviewed a sensor management framework that has been proposed for the efficient detection of static targets in a gridded region of interest and has considered the implementation of this framework on a set of real landmine data collected with EMI, GPR, and seismic sensing modalities. The signal processing algorithms used to process each type of sensor data have also been presented. Simulation results demonstrate that the sensor manager, operating on the decision statistic outputs produced by the aforementioned signal processing algorithms, outperforms a direct search technique and is able to obtain a lower probability of error after a fixed total cost. The presented simulation results further demonstrate that modeling the sensor P_d and P_f values as uncertain quantities allows the sensor manager to perform even better. These results both demonstrate the utility of the proposed sensor manager in operation on real data and reinforce the importance of modeling the uncertainty that will be present in real-world problems. A further important note is that the sensor manager is not designed for use with a specific signal processing strategy; rather, as new and better-performing processing techniques become available, they may be easily incorporated into the overall system. The sensor manager will simply operate on the decision statistic outputs of the new signal processing algorithms.

Acknowledgments

This work was supported by a MURI sponsored by DARPA and ARO.

References

[1] M. Kalandros and L. Y. Pao, "Multisensor covariance control strategies for reducing bias effects in interacting target scenarios," *IEEE Trans. on Aerospace and Electronic Systems*, vol. 41, no. 1, pp. 153-173, 2005.

[2] C. Kreucher, K. Kastella, and A. O. Hero, "Multitarget tracking using the joint multitarget probability density," *IEEE Trans. on Aerospace and Electronic Systems*, vol. 41, no. 4, pp. 1396-1414, 2005.

[3] R. Mahler, "Objective functions for bayesian control-theoretic sensor management, I: multitarget first-moment approximation," *Proc. IEEE Aerospace Conf.*, vol. 4, pp. 4/1905-4/1923, 2002.

[4] M. P. Kolba and L. M. Collins, "Information-based sensor management in the presence of uncertainty," *IEEE Trans. on Signal Processing*, vol. 55, no. 6, pp. 2731-2735, 2007.

[5] M. P. Kolba, P. A. Torrione, and L. M. Collins, "Information-based sensor management for landmine detection using multimodal sensors," *Proc. SPIE*, vol. 5794, pp. 1098-1107, 2005.

[6] P. A. Torrione and L. M. Collins, "Performance comparison of automated induction-based algorithms for landmine detection in a blind field test," *Subsurface Sensing Technologies and Applications*, vol. 5, no. 3, pp. 121-150, 2004.

[7] P. A. Torrione, C. S. Throckmorton, and L. M. Collins, "Performance of an adaptive feature-based processor for a wideband ground penetrating radar system," *IEEE Trans. on Aerospace and Electronic Systems*, vol. 42, no. 2, pp. 644-658, 2006.

[8] W. R. Scott, Jr., K. Kim, G. D. Larson, A. C. Gurbuz, and J. H. McClellan, "Combined seismic, radar, and induction sensor for landmine detection," *Proc. IEEE International Geoscience and Remote Sensing Symposium*, vol. 3, pp. 1613-1616, 2004.

[9] L. Carin, H. Yu, Y. Dalichaouch, A. R. Perry, P. V. Czipott, and C. E. Baum, "On the wideband EMI response of a rotationally symmetric permeable and conducting target," *IEEE Trans. on Geoscience and Remote Sensing*, vol. 39, no. 6, pp. 1206-1213, 2001.

[10] K. Kim and W. R. Scott, Jr., "Design of a resistively-loaded vee dipole for ultra-wideband ground-penetrating radar applications," *IEEE Trans. on Antennas and Propagation*, vol. 53, no. 8, pp. 2525-2532, 2005.

[11] W. R. Scott, Jr., J. S. Martin, and G. D. Larson, "Experimental model for a seismic landmine detection system," *IEEE Trans. on Geoscience and Remote Sensing*, vol. 39, no. 6, pp. 1155-1164, 2001.