

A Comparison of the Performance of Statistical and Fuzzy Algorithms for Unexploded Ordnance Detection

Leslie M. Collins, *Member, IEEE*, Yan Zhang, Jing Li, *Student Member, IEEE*, Hua Wang, Lawrence Carin, *Fellow, IEEE*, Sean J. Hart, Susan L. Rose-Pehrsson, Herbert H. Nelson, and Jim R. McDonald

Abstract—In most field environments, unexploded ordnance (UXO) items are found among extensive surface and subsurface clutter and shrapnel from ordnance. Traditional algorithms for UXO remediation experience severe difficulty distinguishing buried targets from anthropic clutter. Furthermore, naturally occurring magnetic geologic noise often adds to the complexity of the discrimination task. These problems render site remediation a very slow, labor-intensive, and inefficient process. While sensors have improved significantly over the past several years in their ability to detect conducting and/or permeable targets, reduction of the false alarm rate has proven to be a significantly more challenging problem. Our work has focused on the development of signal processing algorithms that incorporate the underlying physics characteristic of the sensor and of the anticipated UXO target, in order to address the false alarm issue. In this paper, we describe several algorithms for discriminating targets from clutter that have been applied to data obtained with the multisensor towed array detection system (*MTADS*). This sensor suite has been developed by the U.S. Naval Research Laboratory (NRL), and includes both electromagnetic induction (EMI) and magnetometer sensors. We describe four signal processing techniques that incorporate features derived from simple physics-based sensor models: a generalized likelihood ratio technique, a maximum likelihood estimation-based clustering algorithm, a probabilistic neural network, and a subtractive fuzzy clustering technique. These algorithms have been applied to the data measured by *MTADS* in a magnetically clean test pit and at a field demonstration. We show that overall the subtractive fuzzy technique performs better than the alternative techniques when the training and testing data sets are separate. The results also allow us to quantify the utility of fusing the magnetometer and the EMI data, and we show that performance is improved when both EMI and magnetometer features are utilized. The results indicate that the application of advanced signal processing algorithms could provide up to a factor of two reduction in false alarm probability for the UXO detection problem.

Index Terms—Electromagnetic induction, magnetometer, sensor fusion, statistical detection, unexploded ordnance detection.

Manuscript received June 23, 2000; revised November 11, 2000 and November 27, 2000. This work was supported by the Strategic Environmental Research and Development Program (SERDP) by Grant DACA72-99-C-0012-CU-1123 and CU-1092.

L. M. Collins, Y. Zhang, J. Li, H. Wang, and L. Carin are with the Department of Electrical and Computer Engineering, Duke University, Durham, NC 27708 USA (e-mail: lcollins@ee.duke.edu).

S. J. Hart, S. L. Rose-Pehrsson, H. H. Nelson, and J. R. McDonald are with the Naval Research Laboratory, Chemistry Division, Code 6110, Washington, DC 20375 USA (e-mail: j.mcdonald@nrl.navy.mil).

Publisher Item Identifier S 1063-6706(01)01505-3.

I. INTRODUCTION

THERE are many areas in the United States and many more outside the continental United States that are contaminated or potentially contaminated with unexploded ordnance (UXO). Several million acres of Department of Defense (DOD) ranges have been returned to civilian control without adequate UXO clearance, and several hundred thousand additional acres on ranges are returned to civilian control each year. In the United States alone there are 1900 formerly used defense sites (FUDS) and 130 base realignment and closure (BRAC) installations that need to be cleared of UXO. Using current technologies, the cost of identifying and disposing of UXO in the United States is estimated to range up to \$500 billion. Site specific clearance costs vary from \$400/acre for surface UXO to \$1.4 million/acre for subsurface UXO. Thus, there is a clear need to effectively and cost-efficiently remediate UXO contaminated lands, rendering them safe for their current or intended uses. Development of new UXO detection technologies with improved data analysis has been identified as a high priority requirement for over a decade.

Several sensor modalities are currently being explored for the detection and identification of surface and buried UXO. These include electromagnetic induction (EMI), magnetometers, radar, and seismic sensors. These sensors generally experience little difficulty detecting the UXO, thus detection does not create the bottleneck that results in the high cost of remediating sites. The primary contributor to the costs and time associated with remediating a UXO-contaminated site is the high false-alarm rate associated with each of the sensors when operated individually. This high cost results since the false alarms require digging numerous nonordnance targets for each intact ordnance item recovered. Recently, the joint use of towed arrays of magnetometer and time-domain electromagnetic sensors in conjunction with GPS navigation has generated high quality, digital georeferenced mapped data. With the availability of such high quality data, advanced physics-based computer data analysis tools have been developed and have demonstrated the ability to detect essentially all DOD ordnance up to its maximum self-burial depths on a wide variety of military ranges [1]–[5]. The false-alarm rate has, however, continued to be an impediment to efficient range clearance. Other research efforts have indicated that the application of statistical, fuzzy, and neural approaches to the detection of landmines and other ordnance-like objects has shown considerable promise [6]–[35]. Performance improvements have been shown in both

a theoretical context as well as on field-collected data. The goal of this effort is to apply an optimized set of algorithms to data collected in the field with a modern multisensor system in order to investigate to what extent the false-alarm rate could be mitigated.

The Jefferson Proving Ground (JPG) site was a 55 000 acre live fire testing range used from World War II until 1994 [1], [36]. Various estimates exist as to the amount of unexploded ordnance on the site, but a minimum of more than 1 million rounds is thought to still exist. Given the cleanup needs of this and other sites, the development and demonstration of emerging technologies for the detection, identification, and cleanup of sites contaminated by UXO has been pursued over the last several years. Several test sites have been prepared at JPG to evaluate sensor performance. These sites have been seeded with ordnance placed at particular positions, depths, and orientations. There have been four demonstrations of technology at JPGs controlled test sites to establish a baseline for performance.

Based on the success of the detection technologies demonstrated in the first three JPG tests, the goal of the fourth JPG demonstration (JPG IV) was not to detect ordnance, but to determine individual contractor's ability to discriminate ordnance items from nonordnance items [36]. In this test, demonstrators knew the location of each test item, but not its identity, depth, or orientation. After the completion of the demonstration, ground truth for the test site along with raw data for some of the demonstrators was released to the public, allowing the research community to develop and test discrimination algorithms. In this paper, we consider the application of a variety of statistical and fuzzy signal processing techniques to data collected during the JPG IV demonstration with a fielded system, the multisensor towed array detection system (*MTADS*) [1]–[5], [30]–[35]. Specifically, we have investigated the performance of four algorithms whose performance has been optimized for the task of discriminating UXO objects from the clutter objects emplaced during the JPG IV test. These algorithms include a generalized likelihood ratio approach, a maximum likelihood estimation-based clustering algorithm, a probabilistic neural network, and a subtractive fuzzy clustering technique. The *MTADS* sensor suite has been developed by the Naval Research Laboratory (NRL), and includes both EMI and magnetometer sensors. The implementation of the UXO detection algorithms is described as well as their robustness and performance. We also consider the impact of fusing the EMI and magnetometer data on performance. Furthermore, the impact of the specific training set used on overall performance is investigated.

II. THE *MTADS* SYSTEM AND THE JPG IV DATA SET

The *MTADS* technology has been described in detail previously [1]–[5], however the most salient features are summarized here. The system hardware includes a low-magnetic-signature vehicle that is used to tow linear arrays of magnetic and EMI sensors to conduct surveys of large areas. The *MTADS* tow vehicle is a custom-built off-road vehicle specifically modified to have an extremely low magnetic self-signature. Most ferrous components have been removed from the body, drive train, and engine, and have been replaced by nonferrous alloys.

The *MTADS* magnetic sensors are Cs vapor full-field magnetometers (a variant of the Geometrics 822 sensor, designated as the Model 822ROV). An array of eight sensors is deployed as a magnetometer array. The time-dependence of Earth's background field is measured by a ninth sensor, deployed at a static site during survey operations. The EMI sensors are deployed as an array of three pulsed induction sensors and are a variant of the Geonics EM-61 instrument. These sensors, configured as an overlapping horizontal array, transmit a tailored electromagnetic pulse into the earth. The received EMI signal is time-sampled by six detection coils that are colocated with and above the three transmission coils. The Geonics EM-61 sensors have been extensively modified. These modifications include changing the time position and time width of the sampling window monitoring the return signal. The power of the transmitted pulse has been increased, as has the pulse repetition rate. The amplifier gain of the detectors has been increased and the time constant applied to the signal has been significantly reduced. The overall detection signal strength has been increased by a factor of 3–6 depending upon the composition, size, and depth of the target [37].

The sensor positions on the surface of Earth (latitude, longitude, and height) are determined using GPS navigation. This technology provides a real-time position update (at 5 Hz) with an accuracy of about 5 cm. GPS satellite clock-derived time is used to time-stamp both position and sensor data information for subsequent correlation. All navigation and sensor data are provided through electronic interfaces to a data acquisition computer in the tow vehicle.

In the JPG IV demonstration, a group of inert ordnance items and a second group of nonordnance items were buried in known positions at unknown depths and orientations. Demonstrators were provided with the positions of the flagged targets and asked to determine the identity of each object by scoring it in terms of its probability of being ordnance. Demonstrators were given the identities, though not the locations, of the ten types of inert ordnance (20, 57, 76, 90, 105, 152, and 155 mm projectiles, and 60-, 81-, and 108-mm mortars) and some were provided with examples of the nonordnance items included. The nonordnance objects (scrap simulants) consisted of steel items welded primarily from half-inch sheet, rods, I-beams, pipes, channels, and assorted scrap items. These items included about two dozen different plates, cubes, rods, cylinders, and esoteric welded objects. The dimensions, shapes, sizes, and weights of many of the nonordnance items closely approximated the ordnance items emplaced for the test. No information was provided regarding the relative proportions of the ordnance and nonordnance items, or the number of items of each type that were emplaced.

NRL conducted the *MTADS* surveys during the JPG IV demonstration using the magnetometer array in a single continuous mapping survey and using the EMI array to conduct two mapping surveys in orthogonal directions. NRL was not provided access to the UXO and non-UXO items in advance. Their classification performance on the demonstration was based on prior experience with ordnance signatures as opposed to training *in situ*. After the demonstration was completed, ground truth for the site was released. The number of targets and clutter items listed in the JPG IV Ground Truth is 50 and 110, respectively. Five of each of the ten types of ordnance items were emplaced.




Fig. 1. Measured (a), (b) and modeled (c), (d) magnetometer responses for a 152 mm projectile buried at a depth of 0.57 m (a), (c) and a 4.1 kg nonordnance, or scrap, item (b), (d) buried at a depth of 0.48 m. The vertical axis plots sensor output and the two horizontal axes indicate distance in the north and east directions.

III. PHENOMENOLOGICAL MODEL AND FEATURE EXTRACTION

As the name implies, the *MTADS* was designed as a detection system. In the early testing of the system [1], it was discovered that several features of the *MTADS* survey data allowed the development of some classification ability. Because of the low noise and precise sensor positioning associated with *MTADS* magnetometer data, the depths and sizes obtained from fits of magnetic anomalies to a dipole model are precise enough to allow these parameters to be used as features for discrimination. A survey was conducted at the Badlands Bombing Range [2], a site where the targets consist of practice bombs, rockets, and fins and clutter associated with the bombs and rockets. The results of this survey indicated that an impressive degree of classification was possible based on these parameters from the magnetometer model and the judgement of a skilled analyst concerning the anomaly shape and model goodness of fit.

In the same data set, it was also observed that there was significant target shape-dependence of the observed EMI anomalies [33]. Methods have been developed to take advantage of this shape-dependence of the EMI response, and these methods have recently been applied at a live test site [34]. Although these methods showed some promise, there was still progress needed to reach an acceptable classification performance. These features, among others, have been investigated in the context of this paper.

An example of the magnetometer data measured by the *MTADS* system at the JPG IV demonstration for an ordnance and a nonordnance item is shown in Fig. 1 in panels A and B. Sensor output (z axis) is plotted as a function of distance in the North (x axis) and East (y axis) directions. The data in panel A was measured from a 152-mm projectile buried at a depth of 0.57 m and the data in panel B was measured from a 4.1-kg nonordnance item buried at a depth of 0.48 m. Fig. 2 plots the output of the EMI sensor for these same two objects. Again, the

data in panel A was measured from the 152-mm projectile and the data in panel B was measured from the nonordnance item.

Features are extracted from magnetometer data by fitting the data to a dipole magnetic model [33]. This approach to model-based feature extraction is consistent across all of the algorithms considered. The features that can be extracted from a magnetic dipole model include

- 1) Object position.
- 2) Object depth.
- 3) Object size.
- 4) Magnetic moment.
- 5) Dipole azimuth.
- 6) Dipole inclination.
- 7) Fit quality.

Each algorithm described in the next section employed a particular subset of these parameters. Each feature set was chosen using statistical performance criteria that are different for each algorithm. In this analysis, a particular feature set was not imposed on all of the algorithms as our goal was to optimize the performance capability for each algorithm and compare their relative performance. Panels C and D of Fig. 1 plot the magnetometer model fits obtained for the 152-mm projectile and the 4.1-kg nonordnance item, respectively.

A *baseline* EMI model has been developed previously for use with the *MTADS* system [33]. This modeler utilizes a dipole model for the transmitter coil and a point approximation for the receive coil. The target is represented as a magnetic polarizability dyadic $\{M_t, M_z\}$ [38], [39], and the secondary field is approximated as target dipole moment. Based on these assumptions the modeler reports the following features:

- 1) Object position.
- 2) Object depth.
- 3) Object size.
- 4) Object orientations.




Fig. 2. Measured (a), (b) and modeled (c), (d) EMI responses for a 152 mm projectile buried at a depth of 0.57 m (a), (c) and a 4.1 kg nonordnance, or scrap, item (b), (d) buried at a depth of 0.48 m. The vertical axis plots sensor output and the two horizontal axes indicate distance in the north and east directions.

- 5) Magnetostatic model coefficients $\{M_x, M_t\}$ (polarizabilities).
- 6) Fit quality.

Panels C and D of Fig. 2 plot the EMI model fits obtained for the 152-mm projectile and the 4.1-kg nonordnance item, respectively.

The approximations associated with a dipole model for the excitation coil, the point model for the receive coil, and the colocated dipole model for the target may be tenuous, particularly for shallow objects. Two modifications were made to this model in order to improve feature estimation. First, the excitation field has been calculated accurately for the known transmit coil shape via the Biot-Savart law [40]. Second, the model has been modified so that different effective centers for the transverse and axial target dipoles $\{M_x, M_t\}$ are included, and the center locations are determined during the fitting procedure. A modification to the model to incorporate the shape of the receive coil was considered, but the slight changes in the parameter estimates were not reflected by a change in algorithm performance. This modified model was used for parameter estimation for all of the algorithms discussed below with the exception of the probabilistic neural network. Utilization of the features obtained with the modified EMI model did not substantially affect performance, so the trends in the results discussed in Section V are valid regardless of the model that was utilized. The modified EMI model did provide substantially better estimates of the depth of the object, especially for the shallow-buried objects. However, as object depth is not a useful discriminator of targets and clutter for this particular data set, performance was not affected by utilization of the features from the modified model.

The EMI model fitting is implemented by the Levenberg-Marquardt method [41]. The model parameters are fit to the data multiple times with randomly selected initial conditions, and the fit quality is monitored. In general, multiple local minima

are obtained. Only the fitted model estimates with $M_z/M_t > 1$ are retained, as, in general, UXO objects have an aspect ratio that is greater than unity. Often, the fitted depths of the multiple local minima are different. To select between the various local minima, the depth information determined from the magnetometer data is utilized. When multiple local minima with $M_z/M_t > 1$ are obtained in the training phase, a minimum-within-class-distance criteria is used, i.e., the fit that minimizes the distance of the fitted parameters between different targets within the same class is selected. In the testing phase, each local minimum is processed and the final output is the maximum of the processed values. The features generated by this *modified* EMI model are the same as those listed for the baseline model, with the addition of the location of the two dipole centers. These additional features were not considered in this effort, however, but will be considered in future work.

IV. SIGNAL PROCESSING ALGORITHMS

Several signal processing algorithms are considered in this analysis. First, we consider a statistical signal processing approach based on the generalized likelihood ratio test [42]. Traditionally, such statistical processors are less robust when limited training data are available, or when the statistics of the parameters that affect the signal cannot be adequately characterized, as was hypothesized to be the case for the JPG IV demonstration. Therefore, suboptimal but potentially more robust approaches were considered. These include a maximum likelihood estimation-based clustering algorithm [41], [42] whose clusters are initially determined using the ISODATA algorithm [43], [44], a probabilistic neural network [47], and a subtractive fuzzy clustering technique [48]. The various approaches are described below along with the method by which each of the algorithms was trained.

As aforementioned, the feature set that provided the optimum performance on a training data set was chosen independently for each algorithm. Each algorithm initially used the same feature set, but performed a down-select to obtain the subset that provided the best performance. Thus, the final feature set was not necessarily consistent across algorithms. This approach was taken since the eventual goal of this effort is to develop algorithms that provide the best performance in a field environment, and because this approach mimics the mechanisms by which a variety of algorithms would be developed and tested for this system.

A. Generalized Likelihood Ratio Test Algorithm

The two hypotheses to be tested are H_1 : UXO, or target; and H_0 : clutter, or nonordnance. The observation \mathbf{x} is an N -dimensional feature vector and contains a subset of the features described above. Assuming each UXO item is equally likely, the likelihood ratio (LR) or Bayesian processor is the ratio between the probability density functions (pdfs) describing the data under the two hypotheses [42]. In the JPG IV data set there were ten targets which fell into two classes: mortars and projectiles. An analysis of the data indicated that the pdfs for the features for the objects within a class were fairly similar, however the pdfs across classes differed, thus we assumed that uncertainty regarding target class would affect the underlying pdf. With this assumption, the LR is given by (1), shown at the bottom of the page, where $p(\cdot)$ denotes a probability density function. Since the pdfs are not known *a priori*, each of the pdfs must be estimated from the training data. The use of an estimate of the parameters in the likelihood ratio formulation results in a generalized likelihood ratio (GLR). To execute the generalized likelihood ratio test (GLRT), the GLR is compared to a threshold, and if the GLR is greater than a threshold, an ordnance object is declared, otherwise a nonordnance object is declared. This approach is not guaranteed to provide the optimal decision performance, but does provide a statistical means of treating the data. In some cases the performance of a GLR can approach that of the LR, particularly when the signal to noise ratio is high enough to provide accurate estimates of the parameters of interest [42]. The GLRT can then be performed for each feature vector in the test set, and then the GLRT can be compared to a sequence of thresholds to evaluate the True Positive and False Positive rates.

Based on an analysis of the feature data, it was assumed that $p(\mathbf{x} | H_1)$ and $p(\mathbf{x} | H_0)$ follow multivariate Gaussian distributions. This assumption was not necessarily valid in all cases, but allowed the correlation between features to be incorporated into the processing. When more accurate pdfs of the features were utilized, and independence was assumed in order to allow the formulation of the joint pdf as the product of the individual fea-

ture pdfs, performance was degraded. Thus, in the remainder of this paper it is assumed that $p(\mathbf{x} | H_1) = N_N(\mathbf{u}_1, \Sigma_1)$ and $p(\mathbf{x} | H_0) = N_N(\mathbf{u}_0, \Sigma_0)$, where $\mathbf{u}_1, \mathbf{u}_0, \Sigma_1, \Sigma_0$ are the mean vectors and covariance matrices of \mathbf{x} under H_1 and H_0 and $N_N(\cdot)$ represents an N -dimensional normal, or Gaussian, pdf. There were ten types of UXO buried in the JPG IV test, and we assumed each type was equally likely. Three of the UXO items belong to the mortar class and the remaining items belong to the projectile class so $p(\text{mortar})$ and $p(\text{projectile})$ were defined appropriately. The estimates, $\hat{\mathbf{u}}_1, \hat{\mathbf{u}}_0, \hat{\Sigma}_1, \hat{\Sigma}_0$, are obtained from the training data. Without considering the different target classes (mortar and projectile), the GLRT becomes:

$$\begin{aligned} \lambda(x) &= \frac{p(\mathbf{x} | H_1)}{p(\mathbf{x} | H_0)} \\ &= \frac{\frac{1}{(2\pi)^{N/2} |\Sigma_1|^{1/2}} \exp\left(-\frac{1}{2} (\mathbf{x} - \mathbf{u}_1)^T \Sigma_1^{-1} (\mathbf{x} - \mathbf{u}_1)\right)}{\frac{1}{(2\pi)^{N/2} |\Sigma_0|^{1/2}} \exp\left(-\frac{1}{2} (\mathbf{x} - \mathbf{u}_0)^T \Sigma_0^{-1} (\mathbf{x} - \mathbf{u}_0)\right)}. \end{aligned} \quad (2)$$

Taking the natural logarithm and incorporating the constant term into the threshold, (2) becomes

$$\begin{aligned} \ln(\lambda(\mathbf{x})) &= (\mathbf{x} - \mathbf{u}_0)^T \Sigma_0^{-1} (\mathbf{x} - \mathbf{u}_0) \\ &\quad - (\mathbf{x} - \mathbf{u}_1)^T \Sigma_1^{-1} (\mathbf{x} - \mathbf{u}_1). \end{aligned} \quad (3)$$

It is a simple extension of the GLRT to incorporate the different target classes, although the succinct form of the previous equation is not obtained in this case.

B. Maximum Likelihood Estimation-Based Clustering Algorithm

A clustering algorithm was investigated since it is well documented that this type of algorithm performs well for cases involving limited training data in which the goal is to decide to which class a particular object belongs. The difference between the standard approach and the approach we have pursued lies in both how the clusters are determined and how the iteration that is used to assign objects to clusters is performed. The clusters for the maximum likelihood estimation-based clustering algorithm are initialized using the ISODATA algorithm [43], [44]. This algorithm is a modified K -means approach that was utilized to initially divide the training data into clusters associated with UXO and nonordnance. A traditional K -means algorithm assumes that the data can be modeled as K clusters. The modification considered here assumes that K is unknown, and is then estimated from the data. The training data is a set of training feature vectors as described above. In the algorithm, K' cluster centroids are initialized, and then each training fea-

$$\begin{aligned} \lambda(x) &= \frac{p(\mathbf{x} | H_1)}{p(\mathbf{x} | H_0)} = \frac{\sum_{\text{target class}} p(\mathbf{x} | H_1, \text{target class}) p(\text{target class})}{p(\mathbf{x} | H_0)} \\ &= \frac{p(\mathbf{x} | H_1, \text{mortar}) p(\text{mortar}) + p(\mathbf{x} | H_1, \text{projectile}) p(\text{projectile})}{p(\mathbf{x} | H_0)} \end{aligned} \quad (1)$$

ture vector is associated with the closest of the K' centroids. The K' centroids are then reestimated, and clusters are merged if their centroids are within a predefined distance measure, or split if enough members are present and the cluster has a high average distance between its mean and its members. The process of re-associating the feature vectors with the various centroids and re-estimating the centroids is continued until all feature vectors do not change their association from iteration to iteration. A detailed discussion of this algorithm can be found in [43]–[46].

The ISODATA algorithm assumes that the clusters associated with the data do not overlap. Furthermore, it assumes that the variance of the data within each of the clusters is the same. These two assumptions are violated for the problem of discriminating UXO from nonordnance. To mitigate this problem, a maximum likelihood classification rule is implemented. The conditional pdfs $p(i/X_j)$, which is the probability that the feature vector X_j belongs to cluster i , and new estimates of the cluster centers V_i , are determined iteratively using the training data. Briefly, we define $p(i/X_j)$ as above, P_i to be the *a priori* probability of selecting the i th cluster, Σ_i as the covariance of the i th cluster, V_i to be the center of the i th cluster, M to be the number of training vectors, and K to be the number of clusters. Prior to estimating the conditional densities, the parameters are initialized as

$$p(i/X_j) = 1, \quad i = \arg \min_{i=1 \dots K} d(X_j, V_i)$$

$$P_i = \frac{\sum_{j=1}^M p(i/X_j)}{M}$$

$$\Sigma_i = \frac{\sum_{j=1}^M p(i/X_j)(X_j - V_i)(X_j - V_i)^T}{\sum_{j=1}^M p(i/X_j)} \quad (4)$$

where d is an exponential distance metric. The updated versions of the conditional pdfs, $\hat{p}(i/X_j)$ are calculated using the iterative procedure outlined below:

Step 1)

$$\hat{p}(i/X_j) = p(i/X_j)$$

$$V_i = \frac{\sum_{j=1}^M \hat{p}(i/X_j) X_j}{\sum_{j=1}^M \hat{p}(i/X_j)}$$

Step 2)

$$p(i/X_j) = \frac{d_e(V_i, X_j)}{\sum_{k=1}^K d_e(V_k, X_j)} \quad (5)$$

$$d_e(V_i, X_j) = \frac{P_i}{\det(\Sigma_i)^{0.5}} \exp[-(X_j - V_i)^T \Sigma_i^{-1} (X_j - V_i)]$$

$$P_i = \frac{\sum_{j=1}^M p(i/X_j)}{M} \quad (6)$$

Step 3)

$$\Sigma_i = \frac{\sum_{j=1}^M p(i/X_j)(X_j - V_i)(X_j - V_i)^T}{\sum_{j=1}^M p(i/X_j)} \quad (7)$$

Step 4)

$$\max_{i,j} |p(i/X_j) - \hat{p}(i/X_j)| < \varepsilon,$$

go to step 1 if false, else terminate. (8)

The exponential distance metric may render this algorithm unstable when the initial conditions are not properly chosen, thus the need to implement the ISODATA algorithm as the initial cluster estimation procedure.

To classify a particular test vector y , the probability that y belongs to the i th cluster is calculated as

$$p(i/y) = \frac{d_e(V_i, y)}{\sum_{k=1}^K d_e(V_k, y)}$$

$$d_e(V_i, y) = \frac{P_i}{\det(\Sigma_i)^{0.5}} \exp[-(y - V_i)^T \Sigma_i^{-1} (y - V_i)] \quad (9)$$

and the decision rule utilized is

$$\max_i p(i/y) > \gamma \quad (10)$$

where the threshold γ determines the relationship between the detection and false-alarm probabilities.

C. Probabilistic Neural Network (PNN)

The probabilistic neural network (PNN) is based upon Bayes' classification method [47], [49]–[51]. The basis of the classification method is given in the following equation, where P_i and P_j are the prior probabilities for class i and j , c_i and c_j are the costs of misclassification, and $p_i(x)$ and $p_j(x)$ are the true probability density functions:

$$P_i c_i p_i(x) > P_j c_j p_j(x). \quad (11)$$

The difficulty with this relationship is that the prior probabilities (the probability that a sample will come from a given population distribution) are generally unknown and must be estimated from training data. This is done using Parzen's method of pdf estimation [52]. Bayes' classification will be more likely to group a new sample x into class i if the prior probability or the cost of misclassification is high. This is especially important for classifications where the cost of misclassification is not equal among the classes. In our application, UXO detection, the potential cost of mis-classifying UXO is much more serious than the cost for nonordnance. In the PNN method, different costs can be set for each class, thus producing a better classification for those classes that demand higher performance. In the analysis performed here, however, these costs were not adjusted. Finally, if the probability density of a given class is large in the region of the new sample x then that class is favored. This allows for multimodal distributions to be dealt with appropriately when a

nearest neighbor-based classifier might fail. It has proven convenient and practical to implement the Parzen pdf estimator in a neural network format [52].

The PNN is a nonlinear, nonparametric pattern recognition algorithm that operates by defining a pdf for each data class based on the training set data and the optimized kernel width parameter. For ordnance discrimination, the inputs are the physics-based modeler outputs or pattern vectors. The outputs of the PNN are the Bayesian a posteriori probabilities (i.e., measures of confidence in the classification) that the input pattern vector is a member of one of the possible output classes, for example, UXO or non-UXO.

The hidden layer of the PNN is the core of the algorithm. During the training phase, the pattern vectors in the training set are simply copied to the hidden layer of the PNN. Thus, there are the same number of feature sets at the hidden level as there are training vectors. Unlike other types of artificial neural networks, the basic PNN only has a single adjustable parameter. This parameter, termed sigma (σ), or kernel width, along with the members of the training set, define the pdf for each data class. In a PNN, each pdf is composed of Gaussian-shaped kernels of width σ located at each pattern vector. The pdf essentially determines the boundaries for classification. The kernel width is critical because it determines the amount of interpolation that occurs between adjacent pattern vectors. As the kernel width approaches zero, the PNN essentially reduces to a nearest neighbor classifier. A large kernel width has the advantage of producing a smooth pdf, which exhibits good interpolation properties for predicting new pattern vectors. Small kernel widths reduce the amount of overlap between adjacent data classes. The optimized kernel width is a compromise between an overly small or large σ .

Prediction of new targets using a PNN is more complicated than the training step. Each member of the training set of pattern vectors (i.e., the patterns stored in the hidden layer of the PNN and their respective classifications), and the optimized kernel width are used during each prediction. As new pattern vectors are presented to the PNN for classification, they are serially propagated through the hidden layer by computing the Euclidean distance d between the new pattern and each pattern stored in the hidden layer. The Euclidean distance scores are then processed through a nonlinear transfer function (the Gaussian kernel) given in the following equation:

$$\text{Hidden_Neuron_Output} = \exp\left(\frac{-d}{\sigma^2}\right). \quad (12)$$

Because each pattern in the hidden layer is used during each prediction, the execution speed of the PNN is considerably slower than some other algorithms. The mass data storage requirements can also be quite large since every pattern in the hidden layer is needed for prediction. Several researchers have developed modified PNN algorithms to overcome this limitation, but they were not deemed necessary for this application [53], [54].

The summation layer consists of one neuron for each output class and simply collects the outputs from all hidden neurons of each respective class. The products of the summation layer are forwarded to the output layer, where we compute the estimated probability of the new pattern being a member of each class. In the PNN, the sum of the output probabilities equals 100%. The pattern being predicted is then assigned to the class

with the highest output probability. One of the main advantages associated with using a PNN is the ability to output a probability for each of its classifications. For critical applications, such as ordnance detection and remediation, such an indicator of confidence is extremely useful in assisting the decision making process and reducing the likelihood that individual ordnance items are missed by reducing the detection probability.

As discussed earlier, the calculation of the optimum kernel width σ is imperative for high classification rates to be achieved. For the work described herein, an iterative algorithm for σ optimization was employed [55]. The approach by which the training data was used to determine this parameter is described in Section IV-E.

D. Subtractive Fuzzy Clustering Algorithm

In contrast to the maximum likelihood classification algorithm described above, the subtractive fuzzy clustering algorithm allows a feature vector of interest to be associated with multiple clusters, i.e., it is not restricted to have membership in only one cluster. In addition, once a feature vector's level of membership in each of the possible clusters is determined, a set of fuzzy logic rules can be imposed to determine whether a feature vector should be associated with an ordnance or nonordnance class.

In the UXO discrimination problem, the input is the N -dimensional feature vector \mathbf{x} and the system output is related to the likelihood that the data is a target, although it is not a probability in the mathematical sense. In the training data, the output z is set to 1 for an ordnance item or 0 for a nonordnance item. The training of the clusters is implemented via a modification of the mountain clustering method described by Yager and Filev [56]. One shortcoming of this procedure is that the computation time grows exponentially with the dimensionality of the feature space. The subtractive clustering method used here [48] is a modification to the mountain clustering method that is computationally more efficient. As implemented in this application, the N features are normalized in each dimension so that their coordinate ranges in each dimension are equal, and a standard subtractive clustering algorithm is applied. The clusters are determined by applying this approach to the training data.

The cluster validation problem, or the criteria for accepting and rejecting cluster centers is important for any clustering method, and the problem is essentially unsolved [57]. In the subtractive clustering algorithm, the rule that was adopted is to set an upper threshold $\bar{\epsilon}$ and a lower threshold $\underline{\epsilon}$. If $m_{c_k} > \bar{\epsilon}$, the cluster is accepted definitely, if $m_{c_k} < \underline{\epsilon}$, the cluster is rejected definitely, where m_{c_k} is the mountain function value of the cluster candidate. Otherwise, if the mountain function value falls between the two thresholds, the shortest distance to the existing cluster centers is calculated. This value is used to determine whether the candidate cluster center provides a good tradeoff between having a large mountain function value and being far away from existing cluster centers. Although the number of clusters is determined by the algorithm automatically, the user-specified parameters strongly affect the number of clusters. In our implementation, the default values from the original algorithm were used.

Instead of determining the absolute membership of a test vector in one of the clusters, a fuzzy approach is implemented where each cluster center defines a rule. Thus, the subtractive

clustering training algorithm establishes the number of fuzzy rules (number of clusters) as well as the rule premises (cluster centers). Given a test data point \mathbf{x} in the N -dimensional feature space (\mathbf{x} is a $1 \times N$ vector), the degree to which the i th rule is fired, i.e., its firing strength, is defined as

$$u_i = \exp\left(-\frac{\|\mathbf{x} - \mathbf{V}_i\|^2}{2\sigma^2}\right) \quad (13)$$

where \mathbf{V}_i is the i th cluster center and σ is the application-specific constant that determines the height and smoothness of the mountain function. The output z is computed via

$$z = \frac{\sum_{i=1}^K u_i z_{c_i}}{\sum_{i=1}^K u_i} \quad (14)$$

where z_{c_i} is the consequence of the i th rule and K is the number of rules.

The rule structure adopted was a first-order Sugeno rule [58], since it has been established that this model can accurately represent complex behavior with only a few rules [59]. The rules thus have the form

$$z_{c_i} = \mathbf{g}_i \mathbf{x}^T + h_i \quad (15)$$

where \mathbf{g}_i a $1 \times N$ vector and h_i is a constant. Given that we have a set of rules with fixed premises, optimizing the parameters in the consequent equations with training data reduces to a least squares problem [58]. Letting

$$\rho_i = \frac{u_i}{\sum_{j=1}^K u_j} \quad (16)$$

the equation for the output of the system can be written

$$z = \sum_{i=1}^K \rho_i z_{c_i} = \sum_{i=1}^K \rho_i (\mathbf{g}_i \mathbf{x}^T + h_i). \quad (17)$$

Given M training vectors, the set of outputs can be written

$$\begin{bmatrix} z_1 \\ \vdots \\ z_M \end{bmatrix} = \begin{bmatrix} \rho_{1,1} \mathbf{x}_1^T & \rho_{1,1} & \cdots & \rho_{K,1} \mathbf{x}_1^T & \rho_{K,1} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \rho_{1,M} \mathbf{x}_M^T & \rho_{1,M} & \cdots & \rho_{K,M} \mathbf{x}_M^T & \rho_{K,M} \end{bmatrix} \times \begin{bmatrix} \mathbf{g}_1^T \\ h_1 \\ \vdots \\ \mathbf{g}_K^T \\ h_K \end{bmatrix} \quad (18)$$

where $\rho_{i,j}$ denotes ρ_i evaluated at \mathbf{x}_j . Again, an output z is set to 1 or 0 depending on whether the corresponding input feature vector \mathbf{x} is associated with an ordnance or nonordnance item. Given the set of training vectors, the first matrix on the right-hand side (RHS) of the above equation is constant while the second one consists of the parameter vectors to be estimated. Thus, it takes the form of a standard linear problem with the format of $\mathbf{A}\mathbf{P} = \mathbf{B}$ that can be solved by least squares. To solve for the parameters of interest, the recursive least squares (RLS) estimation algorithm is used in order to avoid both the computational complexity as-

sociated with calculating the inverse of \mathbf{A} , as well as the numerical stability issues. RLS has been fully described [60], and determines the parameters \mathbf{P} via the following iterative formula

$$\begin{aligned} \mathbf{P}_{i+1} &= \mathbf{P}_i + \Sigma_{i+1} \mathbf{a}_{i+1} (\mathbf{b}_{i+1}^T - \mathbf{a}_{i+1}^T \mathbf{P}_i) \\ \Sigma_{i+1} &= \Sigma_i - \frac{\Sigma_i \mathbf{a}_{i+1} \mathbf{a}_{i+1}^T \Sigma_i}{\mathbf{1} + \mathbf{a}_{i+1}^T \Sigma_i \mathbf{a}_{i+1}} \quad i = 0, 1, \dots, n-1 \end{aligned} \quad (19)$$

where

- \mathbf{P}_i parameter estimate at the i th iteration;
- Σ_i covariance matrix;
- \mathbf{a}_i^T i th row vector of \mathbf{A} ;
- \mathbf{b}_i^T i th row vector of \mathbf{B} ;
- \mathbf{P}_n final least squares estimate.

The initial condition of the iteration can be set to $\mathbf{P}_0 = \mathbf{0}$ and $\Sigma = \mu \mathbf{I}$, where \mathbf{I} is an identity matrix and μ is a large positive number.

Next, the (first-order) Sugeno type fuzzy inference system is constructed. There are K fuzzy IF-THEN rules; each rule follows the format

$$\begin{aligned} \text{If } x(1) \text{ is } A_{i1} \text{ and } x(2) \text{ is } A_{i2} \text{ and } \dots x(n) \text{ is } A_{in} \\ \text{then } z = \mathbf{g}_i \mathbf{x}^T + h_i \end{aligned} \quad (20)$$

where

- $x(j)$ j th element of the vector under test, \mathbf{x} ;
- z system output;
- \mathbf{g}_i, h_i defined above.

A_{ij} is an exponential, or Gaussian, membership function given by

$$A_{ij}(q) = \exp\left(-\frac{(q - x_{c_i}(j))^2}{2\sigma^2}\right) \quad (21)$$

where $x_{c_i}(j)$ is the j th element of \mathbf{x}_{c_i} (i.e., the i th rule premise), and σ is defined in the firing strength expression [(13)]. The membership functions are completely defined by the cluster centers. In effect, aggregation is performed by adding the outputs after multiplication by the firing strength and defuzzification is implemented by weighting the average of the outputs. A Gaussian membership function was chosen since it is completely specified by two parameters, it is invariant under multiplication and Fourier transformation, it is smooth, and the notation is concise. Multiplication is used as the AND operator. The fuzzy inference system weights the output of each rule by the rule's firing strength, and computes the output value as a weighted average of the output of each rule.

E. Training of the Algorithms

There was no separate training area available to NRL prior to or during the JPG IV demonstration, so a round-robin training method, also known as a leave-one-out method, was one of the approaches used to determine the parameters required by the various algorithms. This approach attempted to provide a baseline as to the best possible performance attainable by these algorithms for the particular data set. The method uses the data set circularly, so that the object being tested is not used to develop the statistics that are used to test it. For example, when attempting to determine whether the object associated with measurement $\hat{\mathbf{x}}$ is a target or clutter item for the GLRT processor, we

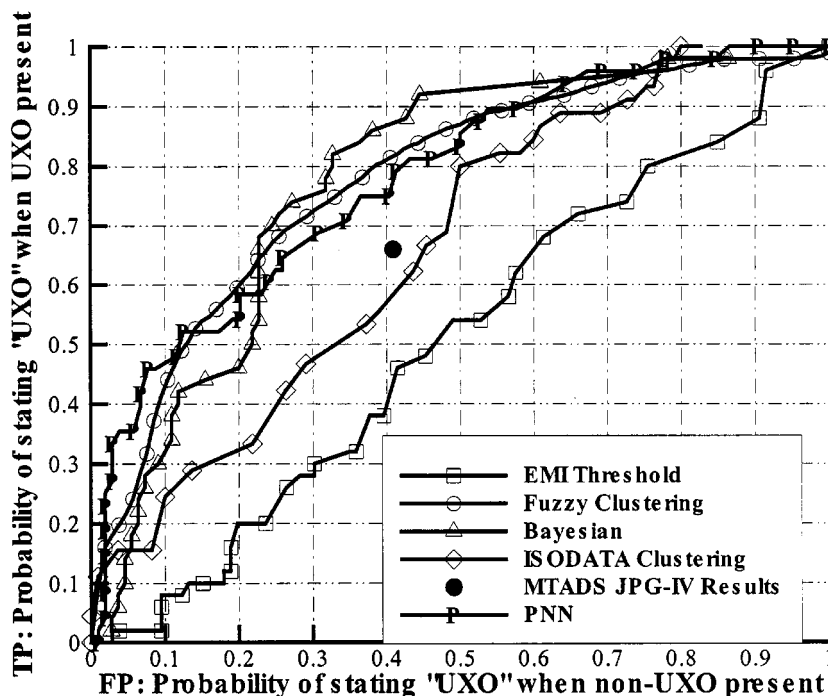


Fig. 3. Performance for the four algorithms and a baseline measure (threshold on the EMI data) for cross-validation testing on the JPG IV data. *MTADS* system performance at the JPG IV demonstration is also shown.

train the model pdfs, $p(\mathbf{x} | H_1, \bullet)$ and $p(\mathbf{x} | H_0)$, with the entire data set with the exception of $\hat{\mathbf{x}}$. In this way the training and detection data are separate, and the largest possible training set is used. This approach results, for example, in 109 feature vectors at the hidden level of the PNN.

After the ground truth was released, NRL acquired the nonordnance items used at JPG IV. Measurements at several depths and orientations of 20 of these items, as well as 9 inert ordnance items were made in the test pit located at the NRL test facility in Blossom Point, MD, USA using the magnetometer array and the EMI arrays to conduct orthogonal surveys. The test pit data thus formed a comprehensive, but separate, “training data set.” This test pit data was used to train a subset of the algorithms described earlier, and then the JPG IV data was used as a test set. In this scenario, no training was performed on the JPG IV data. Such an approach is often preferred to evaluate the robustness of the candidate algorithms. Other data collected with the *MTADS* system was not utilized for training as the nonordnance items employed for the JPG IV test were substantially different than nonordnance items commonly encountered in field exercises.

In the leave-one-out training approach, 110 different training sets were used each composed of 109 training vectors, and thus the discrimination algorithm for each of the four discrimination approaches change slightly as the training sets change. For example, the cluster centers for the subtractive fuzzy clustering algorithm change slightly across each of the 110 training sets; however, the changes are not dramatic. Although this is not an approach that would be used in practice, it is a standard mechanism by which to evaluate performance when the amount of data available is limited. When the test pit data was used for training, each algorithm was trained once and the resulting algorithm was applied to the JPG IV data. For the fuzzy-clustering algorithm, the number of cluster centers, which is the same as

the number of fuzzy rules, was four in the leave-one-out training approach. This relatively small number of cluster centers attests to the particular advantage of using the Sugeno type fuzzy rules in that complex behavior can be represented with only a few rules. Three clusters were generated when the test pit data was used for training the algorithm.

V. RESULTS

The ROC curves for the four algorithms described in the previous section are shown in Fig. 3 for the case where leave-one-out training was applied in conjunction with the JPG IV data. ROC performance is plotted as the “Probability of stating UXO is present when UXO is present,” or True Positives, as a function of the “Probability of stating UXO is present when UXO is not present,” or False Positives. All processors utilize both magnetometer and EMI data, thus the results are essentially “sensor fusion” results. For the GLRT, or Bayesian, processor, the feature set that provides the best performance consists of magnetometer-moment azimuth, transverse-EMI-polarizability, and the ratio between the two EMI polarizability moments. The remaining processors utilized all of the features—performance was not dramatically affected by exclusion of the features not included in the GLRT processor. Fig. 3 plots the ROC performance of each of the four processors considered, as well as a ROC obtained for a simple threshold processor operating on the EM data alone. The EM threshold is indicative of the performance obtained by a cueing processor that might be used as a preprocessor. In addition, the performance of the *MTADS* system using “man in the loop” processing as scored for the JPG IV demonstration is shown.

Clearly, a simple threshold on the EMI data is not an effective discriminator of UXO from non-UXO items. The maximum-

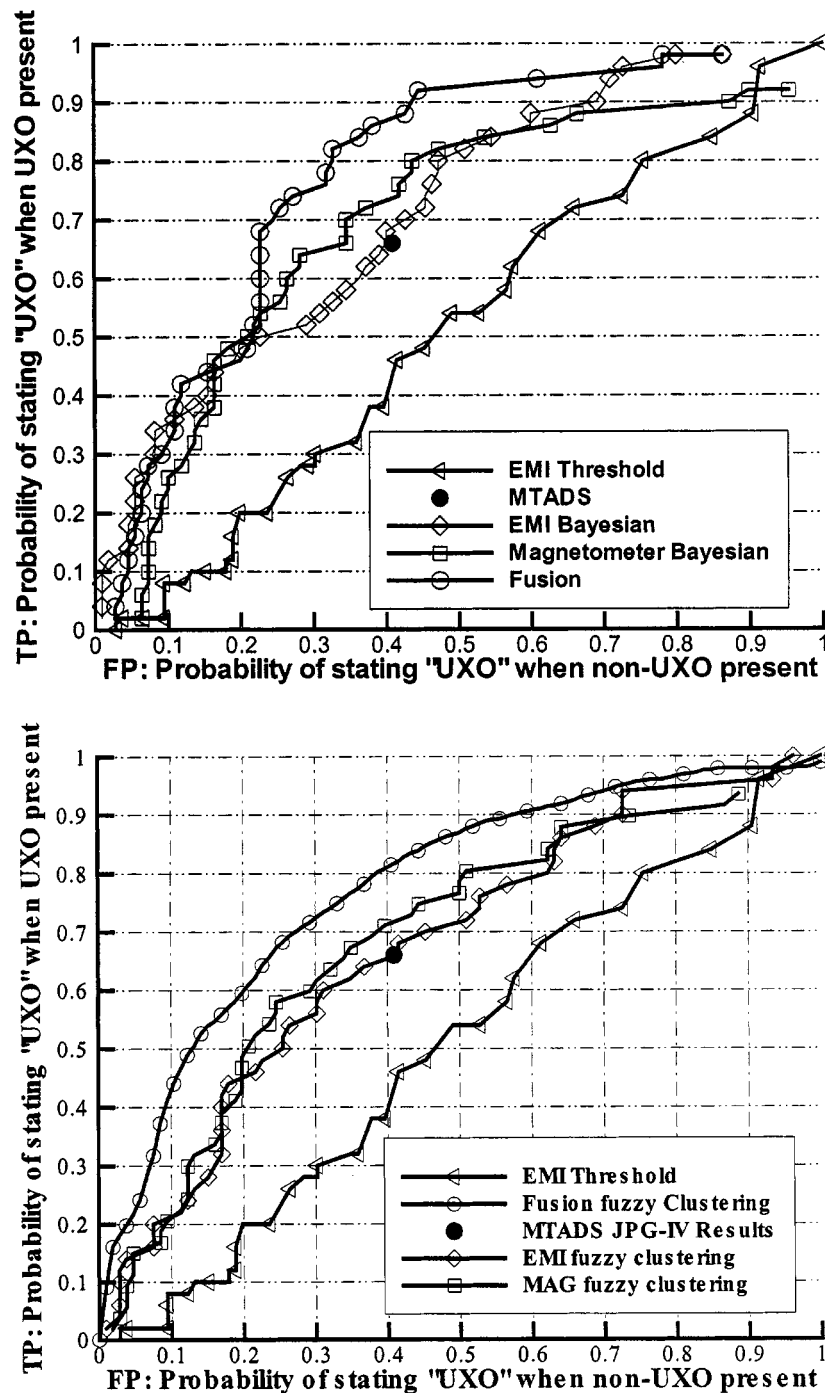


Fig. 4. Comparison of single sensor performance versus fused performance for the Bayesian/GLRT algorithm (top panel) and the subtractive fuzzy clustering algorithm (bottom panel).

likelihood clustering algorithm improves performance somewhat, and its performance is near that achieved by the NRL "man-in-the-loop" approach. The PNN, subtractive fuzzy clustering technique, and the Bayesian approach provide substantial performance improvements over that obtained by the maximum likelihood clustering technique. This may be a result of the additional information incorporated by the fuzzy rule base and the PNN, or the correlation structure incorporated into the Bayesian processor. Although the limited amount of data does not allow strict statistical comparisons across algorithms, it appears that the PNN performance is better than the other algo-

ritms at the lower true positive rates and that the GLRT performance is better than the other algorithms at the higher true positive rates. The subtractive fuzzy technique tends to perform at the average of the PNN and GLRT algorithms. No attempt has been made to date to fuse the processor outputs using a technique such as algorithm ensembles [61] (i.e., joint decisions made using many diverse classification algorithms), although this and other techniques will be considered in future work.

ROC curves for the GLRT processor and the subtractive fuzzy processor, where the EMI and magnetometer features have been considered separately are shown in Fig. 4. These




Fig. 5. Cross-site testing and training results.

processors were considered so that the relative performance associated with each sensor could be gauged. Again, the performance of a simple EMI thresholding algorithm and the NRL “man in the loop” processing results are shown for comparison. The GLRT (top panel) operating on the magnetometer features alone provides results comparable to NRLs man-in-the loop processing. The features used for the GLRT/magnetometer-only processor consisted of dipole-azimuth and dipole-orientation. Utilization of these features has been suggested and validated in previous studies that suggested that discriminating UXO based on induced versus remnant magnetization was possible [62]. The GLRT operating on the EMI features alone, which consisted of the depth estimate, transverse-EMI-polarizability, and the ratio between the two EMI polarizability moments, results in better performance than that obtained with the magnetometer alone. Fusion of the data features between the EMI and the magnetometer provides an additional performance improvement over that obtained with the EMI features alone, especially at the higher true positive probabilities. The fuzzy technique (bottom panel) operating on the EMI features alone provide results comparable to NRLs man-in-the loop processing, and the fuzzy technique applied to the magnetometer features results in slightly better performance. As observed with the GLRT, fusion of the data features across sensors provides an additional performance improvement over either sensor alone. This performance improvement is observed for all true-positive levels.

Fig. 5 shows ROC curves with the sensor features fused for the Bayesian, fuzzy clustering, and PNN for the case where the algorithms were trained on the test pit data collected at the Blossom Point site and tested on the data collected during the JPG IV exercise. In this case, the training data set and the test data set were completely separate. As in the previous figures, the

performance of an algorithm which simply thresholds the EMI data as well as the NRL “man in the loop” performance is shown for comparison. As expected, the performance for each of the algorithms degrades over that obtained with the leave-one-out training on the JPG IV data set. The PNN performance is degraded the most, although its performance continues to exceed that of the simple EMI threshold. The subtractive fuzzy technique provides the best performance overall, although the GLRT performance is similar at the higher true positive rates. In the UXO classification problem, the 90%–95% true positive region or higher on the ROC curve is considered the ultimate discrimination goal, given the risks associated with lower UXO clearance rates. Comparing the results at a 90% true positive rate shows that the PNN and GLRT results are equivalent (approximately 76% false positives) while the fuzzy clustering algorithm has higher performance (approximately 64% false positives).

One possible reason that the performance is degraded so substantially for the cross-site training results is that the EMI data at JPG IV were collected using an older version of the data acquisition system than was used to obtain the training data at the Blossom Point site. In the JPG IV collection with the older data acquisition system, there was a variable and drifting offset between the three EMI sensors. As a result, what should have been symmetric signals from vertical objects often were not as symmetric as theory would predict, and this lack of symmetry affected the estimation of the features. Thus, the statistical properties of the physics-based features associated with the Blossom Point training set could be different from those encountered at JPG IV. The magnitude of this effect will be considered in future work using data from other sites. An additional source of error may arise from the presence of naturally occurring magnetic noise at JPG IV which was not present in the Blossom Point training data.

VI. CONCLUSION

The performance of a variety of signal processing techniques has been evaluated on a set of data collected by the *MTADS* sensor suite with the goal of decreasing the false-alarm rate associated with discriminating UXO from clutter. Prior to this analysis, data was collected with the *MTADS* sensor during the JPG IV technology demonstration exercise. This exercise was designed to test the ability to discriminate between ordnance and nonordnance. The techniques that were applied to this data included a maximum likelihood clustering technique, a probabilistic neural network, a generalized likelihood ratio test, and a subtractive fuzzy clustering technique. These approaches were selected based on their hypothesized robustness to uncertainty or their robustness when faced with a limited training set. A set of physics-based features obtained from the simple magnetometer and EMI models was available to each algorithm. The algorithms were compared based on their ability to discriminate nonordnance, or clutter, items from ordnance items. Performance of the various algorithms was evaluated under two training scenarios: a leave-one-out approach where the JPG IV data was used for both testing and training, and a cross-site approach where data from a collection separate from the JPG IV exercise was utilized. This approach allowed both an assessment of performance bounds and an estimate of the robustness of the various algorithms.

All of the signal processing algorithms considered, with the exception of the maximum likelihood clustering technique, appear to hold promise for processing EMI and magnetometer data for UXO discrimination when limited training data is available. The *MTADS* false-alarm probability was 42% during the demonstration, and the detection probability was 66%. The PNN, GLRT, and subtractive fuzzy technique achieved false-alarm probabilities (respectively) of 27%, 22%, and 24% in the leave one out training approach and 47%, 30%, and 27% in the cross-site training approach at a detection probability of 66%. Thus, the false alarm rate can be reduced by as much as a factor of 2 using these approaches on this data set. Alternatively, at a fixed false-alarm probability, the improvement in the detection rate can also be considered. The PNN, GLRT, and subtractive fuzzy technique achieved detection probabilities (respectively) of 75%, 87%, and 81% in the leave one out training approach and 50%, 76%, and 71% in the cross-site training approach at a false alarm probability of 42%. Again, the algorithms considered here provided substantial improvements in performance.

The success of the GLRT and the PNN may be at least partially due to the fact that these approaches model the decision process statistically, and thus deal effectively with problem uncertainty. The success of the subtractive fuzzy technique is probably also due, at least partially, to the ability of such techniques to robustly handle uncertainty in the data. Such approaches must be tested on more extensive data sets prior to implementing them for a system such as *MTADS*.

The analysis of the JPG IV data is predicated on the assumption that a model can be used to extract features from data that has been collected in an area where either a target or clutter item is located. The features obtained from this region of interest (ROI)

are then processed by one of the algorithms to determine whether the object is ordnance or nonordnance. During field operations, ground truth is not available, and an algorithm or operator must cue the algorithms to process specific data segments. Clearly, it is important that the cueing algorithm maintain a very high detection rate. During the third JPG demonstration, the *MTADS* was reported to achieve 95% detection at a false-alarm rate of approximately 42 per hectare [63]. Currently, the *MTADS* uses three cueing approaches: an operator (man-in-the-loop), an EM-based cue-er, and a magnetometer-based cue-er. The performance of such cueing algorithms without the subsequent signal processing algorithms was simulated in this work by analyzing the performance of a simple EM threshold algorithm. The cross-site training and testing results indicate that the false-alarm rate was reduced by about 30% by the best performing algorithm for most detection rates over this simple cueing algorithm.

The performance of all of the algorithms considered exceeded that obtained by a trained human operator processing the JPG IV data "in the loop," although this operator was not allowed access to training data from this particular site. When the data was re-evaluated by a separate operator, the "man in the loop" performance improved [31]. However, the performance of the automated algorithms presented here continues to out-perform the man in the loop approach. The performance also exceeded that of a simple EMI threshold. It is important to acknowledge that in addition to the classification performance improvement, all of the algorithms remove the human expert from the detection and remediation crew while at the same time significantly reducing the time required for analysis, thus reducing the associated clean-up costs. Future work will consider the utilization of features from more sophisticated EMI and magnetometer models.

ACKNOWLEDGMENT

The authors would like to thank Dr. C. Baum for the helpful discussions regarding the modifications to the dipole EMI model.

REFERENCES

- [1] H. Nelson, J. McDonald, and R. Robertson, "MTADS TECHEVAL Demonstration, October 1996," NRL/PU/6110-97-348, 1997.
- [2] J. McDonald, H. Nelson, J. Neece, R. Robertson, and R. Jeffries, "MTADS unexploded ordnance operations at the Badlands Bombing Range on the Pine Ridge Reservation, Cundy Table, SD," NRL/PU/6110-98-353, 1998.
- [3] J. McDonald, H. Nelson, and R. Robertson, "MTADS Live Site Survey, Bombing Target #2 at the Former Buckley Field, Arapahoe County, CO," NRL/PU/6110-99-379, 1999.
- [4] J. McDonald, H. Nelson, R. Robertson, and R. Jeffries, "MTADS Live Site Demonstration at the Pueblo of Laguna, NM," NRL/PU/6110-00-398, July 6–Aug. 7 2000.
- [5] J. McDonald, H. Nelson, R. Jeffries, and R. Robertson, "Results of the MTADS technology demonstration #3," NRL Tech. Rep., NRL/PU/6110-99-375, 1999.
- [6] S. Tantum and L. Collins, "A comparison of algorithms for subsurface object detection and identification using time domain electromagnetic induction data," *IEEE Trans. Geosc. Remote Sensing*, to be published.
- [7] L. Collins, S. Tantum, P. Gao, J. Moulton, L. Makowsky, D. Reidy, and D. Weaver, "Improving detection of low metallic content landmines using EMI data," in *Proc. SPIE*, Apr. 2000, pp. 14–24.
- [8] T. Dogaru, L. Collins, and L. Carin, "Optimal time-domain detection of a deterministic target buried under a randomly rough interface," *IEEE Trans. Antennas Propagat.*, to be published.

- [9] P. Gao and L. Collins, "A 2-dimensional generalized likelihood ratio test for landmine detection," *Signal Processing*, vol. 80, pp. 1669–1686, July 2000.
- [10] —, "A theoretical performance analysis and simulation of time-domain EMI sensor data for land mine detection," *IEEE Trans. Geosci. Remote Sensing*, vol. 38, pp. 2042–2055, July 2000.
- [11] P. Gao, L. Collins, P. Garber, N. Geng, and L. Carin, "Classification of landmine-like metal targets using wide-band electromagnetic induction," *IEEE Trans. Geosci. Remote Sensing*, vol. 38, pp. 1352–1361, May 2000.
- [12] P. Gao and L. Collins, "A comparison of optimal and sub-optimal processors for classification of buried metal objects," *IEEE Signal Processing Lett.*, vol. 6, Aug. 1999.
- [13] L. Collins, P. Gao, and L. Carin, "An improved Bayesian decision theoretic approach for land mine detection," *IEEE Trans. Geosci. Remote Sensing*, vol. 37, pp. 811–819, Mar. 1999.
- [14] P. Gader, M. Mystkowski, and Y. Zhao, "Application of hidden Markov models to landmine detection using ground penetrating radar," *IEEE Trans. Geosci. Remote Sensing*, to be published.
- [15] P. Gader, J. Keller, and B. Nelson, "Recognition technology for the detection of buried land mines," *IEEE Trans. Fuzzy Syst.*, vol. 9, pp. 31–43, Feb. 2001.
- [16] P. Gader, B. Nelson, H. Frigui, G. Vaillette, and J. Keller, "Landmine detection in ground penetrating radar using fuzzy logic," *Signal Processing—Special Issue Fuzzy Logic Signal Processing*, vol. 80, pp. 1069–1084, June 2000.
- [17] P. Gader, B. Nelson, A. Hocaoglu, S. Auephanwiriyakul, and M. Khabou, "Neural versus heuristic development of choquet fuzzy integral fusion algorithms for land mine detection," in *Neuro-Fuzzy Pattern Recognition*, H. Bunke and A. Kandel, Eds. Singapore: World Scientific Co., submitted for publication.
- [18] P. Gader, Y. Zhao, M. Mystkowski, M. Khabou, Y. Zhang, and A. Hocaoglu, "Hidden Markov models and morphological neural networks for GPR-based landmine detection," in *Proc. SPIE Conf. Detection and Remediation Technology Mines and Minelike Targets IV*, Orlando, FL, to be published.
- [19] J. Keller, S. Auephanwiriyakul, and P. Gader, "New fuzzy set tools to aid in predictive sensor fusion," in *Proc. SPIE Conf. Detection and Remediation Technology for Mines and Minelike Targets IV*, to be published.
- [20] A. Hocaoglu, P. Gader, E. Gelenbe, and T. Kocak, "Optimal linear combination of order statistics filters and their relationship to the delta-operator," in *Proc. SPIE Conf. Detection and Remediation Technology Mines and Minelike Targets IV*, Apr. 1999, pp. 1323–1329.
- [21] M. McClure and L. Carin, "Matched pursuits with a wave-based dictionary," *IEEE Trans. Signal Processing*, vol. 45, pp. 2912–2927, Dec. 1997.
- [22] T. Dogaru and L. Carin, "Time-domain sensing of targets buried under a rough air-ground interface," *IEEE Trans. Antennas Propagat.*, vol. 46, pp. 360–372, Mar. 1998.
- [23] M. McClure and L. Carin, "Wave-based matched-pursuits detection of submerged elastic targets," *J. Acoustical Soc. Amer.*, vol. 104, pp. 937–946, Aug. 1998.
- [24] L. Carin, R. Kapoor, and C. E. Baum, "Polarimetric SAR imaging of buried landmines," *IEEE Trans. Geosci. Remote Sensing*, vol. 36, pp. 1985–1988, Nov. 1998.
- [25] D. Wong and L. Carin, "Analysis and processing of ultra-wideband SAR imagery for buried landmine detection," *IEEE Trans. Antennas Propagat.*, vol. 46, pp. 1747–1748, Nov. 1998.
- [26] L. Carin, N. Geng, M. McClure, J. Sichina, and L. Nguyen, "Ultra-wide-band synthetic aperture radar for mine field detection," *IEEE Antennas Propagat. Mag.*, vol. 41, pp. 18–33, Feb. 1999.
- [27] P. Runkle, P. Bharadwaj, and L. Carin, "Hidden Markov model multi-aspect target classification," *IEEE Trans. Signal Processing*, vol. 47, pp. 2035–2040, July 1999.
- [28] P. K. Bharadwaj, P. R. Runkle, and L. Carin, "Target identification with wave-based matching pursuits and hidden Markov models," *IEEE Trans. Antennas Propagat.*, vol. 47, pp. 1543–1554, Oct. 1999.
- [29] A. Sullivan, R. Damarla, N. Geng, Y. Dong, and L. Carin, "Ultra-wide-band synthetic aperture radar for detection of unexploded ordnance: Modeling and measurements," *IEEE Trans. Antennas Propagat.*, to be published.
- [30] S. Hart, H. Nelson, R. Grimm, S. Rose-Pehrsson, and J. McDonald, "Probabilistic neural networks for unexploded ordnance (UXO) classification using data fusion of magnetometry and EM physics-derived parameters," in *Proc. UXO Forum*, Anaheim, CA, May 2000.
- [31] B. Barrow and H. Nelson, "Model-based characterization of EM induction signatures for UXO/clutter discrimination using the MTADS platform," in *Proc. UXO Forum*, Atlanta, GA, May 1999.
- [32] H. Nelson and J. McDonald, "Pulsed induction and magnetic sensor fusion results using the MTADS," in *UXO Forum Proceedings*, Atlanta, GA, May 1999.
- [33] B. Barrow and H. Nelson, "Collection and analysis of multi-sensor ordnance signatures," *J. Environ. Eng. Geophys.*, vol. 3, no. 2, pp. 71–79, 1998.
- [34] —, "Model-based characterization of EM induction signatures obtained with the MTADS EM array," *IEEE Trans. Geosci. Remote Sensing*, to be published.
- [35] S. Hart, R. Shaffer, S. Rose-Pehrsson, and J. McDonald, "A feasibility study on using physics-based modeler outputs to train probabilistic neural networks for UXO classification," NRL/PU/6110-99-385, 1999.
- [36] "UXO Technology demonstration program and jefferson proving ground: Phase IV," U.S. Army Environmental Ctr., Rep. SFIM-AEC-ET-CR-99051, May 1999.
- [37] B. Barrow, "Testing of the Modified Geonics EM61 for the MTADS System," AETC Tech. Rep. VA-074-063-TR, 1996.
- [38] C. E. Baum, Ed., *Detection and Identification of Visually Obscured Targets*: Taylor Francis, 1998.
- [39] L. Carin, H. Yu, Y. Dalichaouch, and C. Baum, "On the wide-band EMI response of a rotationally symmetric permeable and conducting target," *IEEE Trans. Geosci. Remote Sensing*, to be published.
- [40] C. Balanis, *Advanced Engineering Electromagnetics*. New York: Wiley, 1989.
- [41] P. Gill, W. Murray, and M. Wright, *Practical Optimization*. New York: Academic, 1981.
- [42] H. L. Van Trees, *Detection, Estimation, and Modulation Theory*. New York: Wiley, 1968.
- [43] G. Ball and D. Hall, "Isodata—A novel method of data analysis and pattern classification," Tech. Rep., (NTIS AD699 616), Stanford Res. Inst., Stanford, CA, 1965.
- [44] —, "ISODATA, an iterative method of multivariate data analysis and pattern classification," *IEEE Trans. Int. Comm. Conf.*, June 1966.
- [45] A. Jain and R. Dubes, *Algorithms for Clustering Data*. Englewood Cliffs, NJ: Prentice-Hall, 1988.
- [46] C. Therrien, *Decision, Estimation and Classification—An Introduction to Pattern Recognition and Related Topics*. New York: Wiley, 1989.
- [47] D. Specht, *Neural Networks*, 1990, vol. 3, pp. 109–118.
- [48] S. L. Chiu, "Fuzzy model identification based on cluster estimation," *J. Intell. Fuzzy Syst.*, vol. 2, no. 3, pp. 267–278, 1994.
- [49] R. Shaffer, S. Rose-Pehrsson, and R. McGill, "A comparison study of chemical sensor array pattern recognition algorithms," *Anal. Chim. Acta.*, vol. 384, pp. 305–317, 1999.
- [50] —, "Probabilistic neural networks for chemical sensor array pattern recognition: Comparison studies, improvements and automated outlier rejection," Naval Research Lab., Formal Rep. 6110-98-9879, 1998.
- [51] R. Shaffer and S. Rose-Pehrsson, "Improved probabilistic neural network algorithm for chemical sensor array pattern recognition," *Analytical Chem.*, to be published.
- [52] T. Masters, *Advanced Algorithms for Neural Networks a C++ Sourcebook*. New York: Academic, 1995.
- [53] Y. Chtioui, D. Bertrand, and D. Barba, "Reduction of the size of the learning data in a probabilistic neural network by hierarchical clustering—Application to the discrimination of seeds by artificial vision," *Chemom. Int. Lab. Syst.*, vol. 35, pp. 175–186, 1996.
- [54] Y. Chtioui, D. Bertrand, M. F. Devaux, and D. Barba, "Comparison of multivariate perceptron and probabilistic neural networks in artificial vision. Application to the discrimination of seeds," *J. Chem.*, vol. 11, pp. 111–129, 1997.
- [55] D. F. Specht, *Computational Intelligence: A Dynamic System Perspective*, M. Palaniswami, Y. Attikiouzel, R. J. Marks, D. Fogel, and T. Fokuda, Eds. New York: IEEE Press, 1995.
- [56] R. R. Yager and D. P. Filev, *Essentials of Fuzzy Modeling and Control*. New York: Wiley, 1995.
- [57] D. Langan, J. Modestino, and J. Zhang, "Cluster validation for unsupervised stochastic model-based image segmentation," *Proc. IEEE Trans. Image*, vol. 7, pp. 180–195, 1998.
- [58] T. Takagi and M. Sugeno, "Fuzzy identification of systems and its applications to modeling and control," *IEEE Trans. Syst., Man, Cybern.*, vol. SMC-15, pp. 116–132, 1985.
- [59] M. Sugeno and K. Tanaka, "Successive identification of a fuzzy model and its applications to prediction of a complex system," *Fuzzy Sets Syst.*, vol. 42, pp. 315–334, 1991.

- [60] L. Ljung, *System Identification: Theory for the User*. Englewood Cliffs, NJ: Prentice-Hall, 1987.
- [61] D. Opitz and R. Maclin, "Popular ensemble methods: An empirical study," *J. Artificial Intell. Res.*, vol. 11, pp. 169–198, 1999.
- [62] H. Nelson, T. Altshuler, E. Rosen, J. McDonald, B. Barrow, and N. Khadr, "Magnetic modeling of UXO and UXO-like targets and comparison with signatures measured by *MTADS*," in *Proc. UXO Forum*, Anaheim, CA, May 1998.
- [63] S. Park and M. Mander, "Evaluation of the multi-sensor towed array detection system (*MTADS*) performance at Jefferson proving ground, January 14–24, 1997," Institute for Defense Analysis, IDA Document D-2174, Feb. 1999.

Leslie M. Collins (M'96) was born in Raleigh, NC in 1963. She received the B.S.E.E. degree from the University of Kentucky, Lexington, in 1985 and the M.S.E.E. and Ph.D. degrees in electrical engineering, both from the University of Michigan, Ann Arbor, in 1986 and 1995, respectively.

She was a Senior Engineer with the Westinghouse Research and Development Center, Pittsburgh, PA, from 1986 to 1990. In 1995, she became an Assistant Professor with the Electrical and Computer Engineering Department, Duke University, Durham, NC. Her current research interests include incorporating physics-based models into statistical signal processing algorithms. She is pursuing applications in subsurface sensing, as well as enhancing speech understanding by hearing impaired individuals.

Dr. Collins is a member of the Tau Beta Pi, Eta Kappa Num, and Sigma Xi honor societies.

Yan Zhang was born in Changchun, China, in 1972. He received the B.S., M.S., and Ph.D. degrees in electrical engineering all from Jilin University of Technology, Changchun, China, in 1993, 1996, and 1998, respectively.

Since January 1999, he has been a Postdoctoral Research Associate with the Department of Electrical and Computer Engineering, Duke University, Durham, NC. His present research interests include statistical signal processing, sensor fusion, and related applications.

Jing Li (S'97) was born in Xian, China, in July 1974. He received the B.S. and M.S. degrees in automatic control and system engineering from Xian Jiao-Tong University, China, in 1993 and 1996, respectively. He is currently working toward the Ph.D. degree in the Department of Electrical and Computer Engineering, Duke University, Durham, NC.

His research interests include fuzzy control, LPV control, signal processing, sensor fusion, and pattern recognition. His particular interest is in advanced control design for industrial systems with high performance and reliability.

Mr. Li was a recipient of the Sunshine Fellowship in 1996. He is a member of SIAM.

Hua Wang received the Ph.D. degree from the University of Maryland, College Park.

He is currently an Assistant Professor of Electrical and Computer Engineering, University of Maryland. He is an Associate Editor of the conference editorial board of the IEEE Control Systems Society. His interests include nonlinear control systems, nonlinear dynamics, intelligent systems and control, sensor fusion, control network, biomedical systems, robotics and applications.

Lawrence Carin (S'85–M'89–SM'96–F'00) was born in Washington, DC, on March 25, 1963. He received the B.S., M.S., and Ph.D. degrees in electrical engineering from the University of Maryland, College Park, in 1985, 1986, and 1989, respectively.

In 1989, he joined the Electrical Engineering Department, Polytechnic University, Brooklyn, NY, as an Assistant Professor and became an Associate Professor there in 1994. In September 1995, he joined the Electrical Engineering Department, Duke University, Durham, NC, where he is an Associate Professor. He is the Principal Investigator on a Multidisciplinary University Research Initiative (MURI) on demining. His current research interests include short-pulse scattering, subsurface sensing, and wave-based signal processing.

Dr. Carin is a member of the Tau Beta Pi and Eta Kappa Nu honor societies. He is currently an Associate Editor of the IEEE TRANSACTIONS ON ANTENNAS AND PROPAGATION.

Sean J. Hart received the Ph.D. degree in analytical chemistry in 1998 from Tufts University, Medford, MA.

He then accepted a position as a Senior Staff Scientist with Nova Research, Inc. working on-site at the Naval Research Laboratory (NRL). In June 2000, he accepted a permanent position as a Research Chemist at NRL. He has been involved with the analysis of multivariate data from a variety of chemical and physical sources and the development of optically based chemical sensor systems.

Susan L. Rose-Pehrsson received the Ph.D. degree in analytical chemistry from Pennsylvania State University, University Park, in 1988.

She is a Research Chemist in the Chemistry Division, Naval Research Laboratory. During the past 17 years, she has worked on chemical sensors and multivariate analysis for the detection of hazardous species in air and water.

Herbert H. Nelson was born in Yokosuka, Japan, on September 15, 1953. He received the B.S. degree in 1975 from Tulane University, New Orleans, LA, and the Ph.D. degree in physical chemistry from the University of California, Berkeley, in 1980.

From 1980 to 1983, he was an NRC Postdoctoral Associate at the Naval Research Laboratory (NRL), Washington, DC, where he is now Head of the Molecular Dynamics Section. For the past ten years, he has been involved in the area of UXO detection, most recently as a Member of the *MTADS* team. His research at NRL involves investigation of gas-phase reaction dynamics focusing on chemistry important in combustion systems and the atmosphere.

Dr. Nelson is a member of Sigma Xi, the American Chemical Society, and the American Physical Society.

Jim R. McDonald was born in Austin, TX, on August 20, 1942. He received the B.S. degree from Southwestern University, Georgetown, TX, and the Ph.D. degree from Louisiana State University, Baton Rouge, in 1964 and 1969, respectively.

After a Postdoctoral Fellowship, he joined the Naval Research Laboratory in 1970, where he is now the head of the Chemical Dynamics and Diagnostics Branch in the Chemistry Division. Since the mid 1980s, he has headed programs sponsored by both the Department of Defense and other government agencies to develop automated technologies for detection, location, and characterization of buried unexploded ordnance associated with military ranges and associated adjacent lands. These research and development efforts have resulted in the transition of the *MTADS* technology to the commercial sector, which is now available to perform commercial UXO survey services for the Department of Defense and other agencies.