

A Unified Approach to the Processing and Fusion of Time and Frequency Domain EMI Data for UXO Discrimination

Final Report for SERDP SEED Project CU-1217

Eric L. Miller
Center for Subsurface Sensing and Imaging Systems
Department of Electrical and Computer Engineering
315 Stearns Center
Northeastern University
Boston, MA 02115

January 7, 2003

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 07 JAN 2003		2. REPORT TYPE		3. DATES COVERED 00-00-2003 to 00-00-2003	
4. TITLE AND SUBTITLE A Unified Approach to the Processing and Fusion of Time and Frequency Domain EMI Data for UXO Discrimination				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Northeastern University,Center for Subsurface Sensing and Imaging Systems,315 Stearns Center,Boston,MA,02115				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 31	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Contents

1	Introduction and Background	3
2	Physical Model	5
3	Processing When Sensor Locations Are Known	8
3.1	Library Construction	9
3.2	Parameter Estimation	9
3.3	Classification	10
3.4	Numerical Examples	11
3.5	Field Data Examples	16
3.5.1	Geophex Pipe Data	17
3.5.2	Geophex JPG-IV Data	17
3.5.3	TEMIDS Data	21
4	Processing When Sensor Locations Are Uncertain	24
4.1	General Problem Formulation	24
4.2	A Min-Max Solution Approach	26
4.2.1	Boxes	27
4.2.2	Ellipsoids	27
4.3	Numerical Examples	28
4.4	Field Data Examples	29
5	Conclusions and Future Work	29

1 Introduction and Background

The problem of detecting and classifying buried objects using electromagnetic induction (EMI) sensing technologies has received considerable attention in recent years in a range of application areas including unexploded ordnance (UXO) and landmine remediation. In the last decade so, considerable advances have been made in the area of EMI instrumentation yielding sensors capable of providing data both in the time and frequency domains which convey far more information concerning the structure of buried objects than is the case with older metal detectors. Extracting information such as size, shape, orientation, and type of target though requires the development of advanced signal processing methods which are tied directly to physical models of the sensor. In this report, we consider a number of options for the classification of buried objects given EMI data obtained at multiple points in space in the vicinity of an already-detected object with particular attention paid to UXO and demining applications. The problem of object *detection* from EMI data has received considerable attention in recent years [4] and is more-or-less solved. Thus, we concentrate on the related *classification problem*; i.e. declaring the type of object in the sensor field of view.

Generally, classification is done by comparing either the raw data or some low-dimensional collection of features extracted from the data to entries in a library [11]. The library itself is built from data signature vectors [9, 11] or feature vectors from all targets of interest in a given application [1, 2, 11]. The classifier takes as the selected object that element of the library which in a sense “best fits” the data or the features. If a sufficiently good fit cannot be obtained, the classifier declares that the object under consideration is in fact not in the library. This last step is quite crucial for the UXO and demining problems where there is a strong desire to correctly reject so-called “clutter” items (objects not of immediate threat or interest) as the unnecessary removal of each such items consumes both time and resources.

Most all of the recent work in the area of EMI classification has been based on a simplified physical model for the interaction of the fields with the unknown target. As we describe in greater mathematical depth in § 2, assuming that the target scatters the incident energy like a dipole [5], information concerning the class and orientation of the object in space is encoded in the magnetic polarizability tensor (MPT) which is independent of the location of the object relative to the sensors. This location information is contained in two 3×1 field vectors. Mathematically the MPT is a 3×3 matrix which has a functional dependence on time or frequency depending on the sensor being used. In theory, this matrix can be diagonalized by a time/frequency independent rotation matrix indicating the orientation of the object in space. Each element of the resulting 3×3 diagonal matrix (which carries all of the time or frequency dependence) provides the scattering characteristics of the object along each of its three principal axes and are used for classification purposes. We refer to these as the principal axis polarizability functions (PAPFs).

The various EMI classification methods developed to date differ according to factors related to the sensors being studied, the way clutter is handled, and the manner in which the dipole-model is employed in the processing. Aside from [2], classification algorithms have been concerned strictly with either time-domain EMI sensors [1, 11] or frequency domain [8, 9] but not both. When clutter rejection is considered it is typically in the context of binary problems (Is the object of interest or not?) and not the more challenging classification problem (Which of N objects are we looking at or is the object not in our library?) In [9, 11], the authors employ a parametric model for the PAPFs which takes the form of a sum of decaying exponentials [3, 4, 10] in the time domain or a sum of one-pole rational system functions each with a zero at DC in frequency. The poles and/or decay constant are then used as the features for classification. Alternatively, the methods of [1, 2, 8] use

no such models and treat the time or frequency samples of the individual PAPFs as independent quantities using the entire time or frequency domain signals for purposes of object determination.

The current collection of processing methods also differ in how they treat the fact that, in addition to unknown object class, the orientation and location of the target may not be known or may be known imprecisely. In [1, 2] for example the location of the object is estimated as part of the processing while in [8, 9], the unknown location effects are included in an overall scaling of the data and not treated explicitly. A similar approach is taken to the orientation issue in [11] while in [9] a different target signature is used for each orientation of each target in the library. Finally, in [1, 2, 8], the rotation matrix is in fact determined as part of the eigenen-analysis of the polarizability tensor, however the authors do not map this rotation matrix back to an explicit orientation of the object in space.

Finally, a tacit assumption in most all of these methods is that the positions of the sensor are known precisely at least relative to some fixed reference point, a condition which is often not met in practice. For handheld systems especially, the sensor may not be equipped with a global positioning sensor (GPS) in which case some sort of dead reckoning may be employed to get approximate locations. In the case of vehicular or cart mounted sensors, even with GPS, the effects of positional uncertainties have not been extensively studied for problems where one requires high resolution localization of buried objects.

Given this as background, the goal of the work in this report is the development, analysis, and validation of a general purpose approach to the EMI sensing classification problems which explicitly accounts for the types of uncertainties described in the previous paragraphs. Initially we concentrate on the problem where the sensor positions are assumed known. The insight and experience from that work then forms the basis for approaching the far more complex problem of classification in the presence of unknown sensor location.

All of the classification methods described in this report are based on a physical model that is the fusion of the dipole scattering model in [5] and a parametric PAPF model elucidated in [3]. This model is analytical in the parameters of the PAPF, the (x, y, z) location of the object, the three Euler angles [7] describing the rotation matrix, and the (x, y, z) locations of all sensors. The relatively simple closed form nature of the model with respect to these parameters leads us to an initial classification method in which the sensor positions are assumed known and we explicitly estimate the orientation and location of the object along with the parameters needed for classification. Thus, our approach provides information regarding these geometric characteristics of the object. Also, the closed form nature of the PAPF model allows our approach to be applied with equal ease to both time or frequency domain sensor data.

The model of [6] indicates that in theory the PAPF are comprised of an infinite number of decaying exponentials in time or one pole transfer functions in frequency. Unfortunately, it is both impossible and not necessary to estimate an infinite number (or even a large number) of poles/decay rates. First, in the presence of noise, this is know to be a very delicate signal processing problem. Second, because objects do not scatter *exactly* as dipoles, it makes sense to consider reduced order models for purposes of processing. Indeed, experimentally it has been shown that one or two poles can typically be used to match the model to measured data [6]. Thus in this report, we posit a model in which each PAPF is represented as a single decaying exponential or one pole transfer function. Since there are three principal axes, we estimate three poles/decay constants as part of the classification routine. This then is similar to the one or two pole methods of e.g., [6] however where these approaches disregard the information in the coefficients multiplying the e.g. decaying exponentials in time, we are able to relate these coefficients directly and analytically to the location and orientation of the object in space.

Moreover, our method explicitly accounts for modeling error introduced by the fact that we are *not* using an exact scattering model. More specifically, because a three pole model cannot (in general) exactly represent the data, the pole values will not be independent of object position and orientation. Rather for a given type of target, there will be a “spread” of pole values as a function of these geometric nuisance parameters. Thus, we introduce a simple quadratic-form classifier which compensates for this effect of model mismatch. Classifiers based in data residuals and a fusion of pole estimates and data residuals are also constructed. Finally, all of our classification schemes explicitly contain a null hypothesis that the object is in fact clutter. Thus, we provide a quantitative means of rejecting such items in the context of a multi-object classification procedure.

Having formulated and tested a method for comprehensively solving the classification and characterization problem when the sensor positions are known, we next turn our attention to the case where there may be uncertainty in these quantities. The method we propose is motivated by the observation that real data sets for which positions of the sensor are recorded are often accompanied by the caveat that these locations are accurate to within some tolerances in x , y , and z .

One approach proposed for directly incorporating this information into a processing scheme is based on Bayesian statistics. Specifically, the perturbations in the positions of the sensor are modeled as random variables. Typically, these random variables are taken to be independent and, in the absence of any other information, uniformly distributed between bounds related to the tolerances previously mentioned. The resulting statistical estimation or decision problem is then solved via Monte-Carlo integration methods to “average away” the effects of these unknowns. Such methods are computationally intensive and also have the tendency to average away some of the distinguishing signal as well.

Rather than using a statistical approach, here we examine a worst-case (or min-max) formulation of the problem. That is, we look for those parameters which minimize the maximum misfit to the data where the maximum is taken over all possible sensor locations each of which is restricted to lie within some bounded region of space. The geometric structure of these uncertainty regions is derived again from the prescribed tolerances. In particular, we consider both box shaped as well as ellipsoidal regions in this report. Our motivation for this overall approach is twofold. First, to the best of our knowledge a min-max formulation of the problem of uncertain sensor locations has not been proposed or examined to date in the context of subsurface sensing. Thus, our effort provides a new and potentially useful way of approaching this issue. Second, while min-max problems are generally quite computationally intensive, under a small perturbation assumption we obtain two very tractable algorithms (one for boxes and the second for ellipsoids) for solving the resulting inverse problem. The methods exploits some convenient underlying structure of the problem and are very amenable to a parallel implementation.

The remainder of the report is organized as follows. In § 4.1 the mathematical structure of the basic problem is formulated. § 3 is devoted to a description of the processing method for known sensor positions. Experiments using both simulated as well as real data from time and frequency domain sensors is presented in § 3.4 and § 3.5. Two general solution methods for approaching the problem where the sensor positions are not certain are detailed in § 4.2. § 4.3 and § 4.4 is devoted to a detailed examination these algorithms. Finally conclusions and future work are provided in § 5.

2 Physical Model

We consider a combination of the physical EMI model in [5] describing the scattering of low frequency electromagnetic radiation by spherical or spheroidal objects with the model in [6] which




Figure 1: One sensor comprising sensor coils and target object.

rigorously justifies the use of decaying exponentials in time or one-pole models in frequency for problems of this type. As seen in Fig. 1 the transmitters and receivers are taken to be rectangular¹ coils (not necessarily co-located) with sides of length $2A$. The target center is located at $r_0 = (x_0, y_0, z_0)$ in the $x-y-z$ coordinate system. We are concerned with processing methods based on time or frequency domain sampled data obtained from multiple transmitter/receiver locations. Assuming we collect M time or frequency samples from each of N combinations of transmitters and receivers positions then under the model the k th sample at the n th position is

$$y_{n,k} = g_n^T R^T \Lambda_k R f_n + \sigma w_{n,k} \equiv s_{n,k} + \sigma w_{n,k} \quad (1)$$

where Λ_k is the polarizability tensor for the k th frequency or time sample, R is a rotation matrix which orients the object in space, g is a 3×1 vector holding the x , y , and z components of the magnetic field produced at r_0 by a current I flowing in the receive coil, g^T indicates the transpose of g , f is the excitation field vector evaluated at the dipole position and has a similar functional form to that of g , ω is the operating frequency, σ is taken as the standard deviation of the assumed-additive white Gaussian measurement noise, and $w_{n,k}$ is a zero mean, unit variance normal random variable. We note that g is also explicitly a function of the (x, y, z) coordinates of the measuring sensor which f is a function of the coordinates of the transmitting sensor. Functional forms for f and g are provided in Appendix A of [5].

In the frequency domain the matrix Λ takes the form:

$$\Lambda(\omega) = \begin{bmatrix} \lambda_1(\omega) & & \\ & \lambda_2(\omega) & \\ & & \lambda_3(\omega) \end{bmatrix} \quad (2)$$

with ω replaced by t for time domain sensors. The three λ 's each are associated with one of the principal axes of the object. Here we consider a form of that model provided by [6]

$$\lambda_i(\omega) = \sum_{l=1}^{\infty} \frac{a_{i,l} j \omega}{p_{i,l} + j \omega} \quad i = 1, 2, 3 \quad (3)$$

¹With a little work, the model could be generalized to circular coils.

where $j = \sqrt{-1}$, $p_{i,l}$ is the l th pole for the i th axis, and $a_{i,l}$ is the expansion coefficient. An inverse Fourier transform yields the time-domain version of λ

$$\lambda_i(t) = - \sum_{l=1}^{\infty} a_{i,l} p_{i,l} e^{-p_{i,l} t} u(t) \quad (4)$$

with $u(t)$ the unit step function. For cylinders and disks, Carin *et al* recently provided a fast numerical method for computing the $p_{i,l}$ [3]. The model in (3) and (4) strictly holds for non-ferrous objects. In the case of ferrous objects, one must add a DC offset in frequency or a Dirac δ function in time [3]. For notational ease, in what follows we concentrate on the non-ferrous case with the understanding that these small changes need to be made for ferrous objects.

In (1), R represents a rotation matrix used to transform field quantities between a global frame of reference and the local frame of the object. Here R is parameterized by 3 Euler angles [7] and explicitly takes the form

$$R = \begin{bmatrix} \cos \phi \cos \psi - \sin \phi \cos \vartheta \sin \psi & -\cos \phi \sin \psi - \sin \phi \cos \vartheta \cos \psi & \sin \phi \sin \vartheta \\ \sin \phi \cos \psi + \cos \phi \cos \vartheta \cos \psi & -\sin \phi \sin \psi + \cos \phi \cos \vartheta \cos \psi & -\cos \phi \sin \vartheta \\ \sin \vartheta \sin \psi & \sin \vartheta \cos \psi & \cos \vartheta \end{bmatrix} \quad (5)$$

Gathering the data together from all sensors, we write the overall model in compact notation as

$$d(p, a, \theta) = s(p, a, \theta) + \sigma w. \quad (6)$$

where d is the vector comprised of the data from all sensor locations and time/frequency samples, s is the signal vector, w the noise vector, p the vector of all poles in the model a the vector of expansion coefficients and θ is the six dimensional vector composed of the object coordinates and Euler angles.

This model assumes that the object behaves electromagnetically like a dipole. The three λ_i 's then fully summarize the scattering behavior of the object and are dependent only on the size, shape, and material of the object *not* on the orientation and position of the object relative to the sensor. Thus, the pole and expansion coefficients make good candidates for use in a classification routine. The orientation information is explicitly contained in the matrix R while the field vectors f and g convey position information. Due to the simple, analytical nature of this model, it is quite well suited for use in a signal processing routine where operations like pole fitting and parameter estimation are accomplished using optimization routines. The complexity of these routines is the substantially reduced due to our ability to use the model to compute closed form sensitivity information; essentially the derivative of the data with respect to any of the unknowns: poles, expansion coefficients, Euler angles or location coordinates. Such calculations are at the heart of any parameter fitting scheme employing e.g. a gradient decent, conjugate gradient, or Newton type of optimization scheme.

While the utility of the model described here has been validated using real sensor data [5, 9, 10], generally objects do not behave exactly as dipoles. Moreover, one cannot practically use an infinite number of poles for each λ_i . Rather, a single pole per axis is the most that is typically supported by the data [3, 9, 10]. In such a case, the "effective" pole for each axis *will* be dependent on the object position and orientation. The end result is that for all practical purposes model mismatch or required model reduction will force us to consider pole-based classifiers which explicitly account for variations in the feature values. If such variations are small, then one expects success in using poles (really effective poles) for classification.

Before moving on, to gain a more intuitive feeling the model we are using, in Figure 2 we plot the structure of the in-phase and quadrature components of a single, one-pole transfer function of the form

$$\frac{j\omega}{50 + j\omega}.$$

Based on Figure 2, our ability to distinguish UXO from clutter and one type of UXO from another using EMU data rests on two assumptions:

1. **Assumption 1:** The salient structure in the data varies across objects. Relative to Figure 2 this means in part that the “peak” in the quadrature component, or equivalently the inflection point of the in-phase component, differs from one object to another. Note that the location of the peak is defined directly by the value of the pole in the model (50 in this case). Hence the “salient” structure in the data which we mention here is basically the same as the value of the parameter in the model upon which we shall base our classification processing. Of course, the true situation is a bit more complex as the model we posit has three poles. Nonetheless, the basic intuition is the same.
2. **Assumption 2** The data collected fact reflect the salient structure of the objects. Again, referring to Figure 2, this means that the data should span a range of frequencies allowing us to clearly resolve the rise, peak, and fall of the quadrature structure (or equivalently, the inflection of the in-phase component.)




Figure 2: Generic Structure of One-Pole Model

3 Processing When Sensor Locations Are Known

Our approach to classification starts with the construction of a target signature library which will be used in the actual processing. For each target of interest, this library will be comprised of the three effective pole and expansion coefficients which define the PAPFs. Given that library, classification is a two step process. First, for each target in the library, the data are used to estimate the unknown parameters associated with that model: poles, expansion coefficients, object location and object orientation. Second, using these estimates, we examine three classification schemes. The first classifier is based on using the pole estimates alone and is expected to work well in high

signal to noise cases when we can get accurate estimates of these quantities. The second classifier is based on the fit of the k th model in the library to the available data. As explained more fully in § 3.3, this approach is likely to be of use when the noise level is relatively high. Finally, a simple hybrid classifier is proposed which combines both the pole as well as the data-based methods. We begin by discussing the construction of this library.

3.1 Library Construction

As discussed in § 2, the pole estimates which we use for classification will have some orientation and position dependence which should be accounted for in the construction of the library and in the processing. Let us suppose that we have data from a known target in a known position and orientation which either has been computed using an exact computational model or measured using an actual sensor. For the k -th target in the library, we are going to use one effective pole per λ_i defined in a best fit manner as the solution to the following optimization problem

$$p_k^{eff}(\theta_0), \hat{a}(\theta_0), \hat{\theta}(\theta_0) = \arg \min_{p,a,\theta} \|d_k(\theta_0) - s(p, a, \theta)\|_2^2 \quad (7)$$

where θ_0 holds the true position and orientation information, d_k is the true data vector, and p and p^{eff} are vectors of three pole parameters, one per λ_i . The symbol “ $\hat{\cdot}$ ” above quantities indicate that these are fitted to data. We note that to be consistent with the estimation scheme developed in § 3.2, here we do fit a and θ ; however we care only about the effective pole values in constructing the library. Additionally, the effective pole parameters are implicitly dependent on the specifics of the sensing system we use including frequencies of operation, time gates measured, and spatial sampling strategy. Hence in theory each sensing configuration will require a separate library.

While we could construct a library holding $p_k^{eff}(\theta)$ for a dense sampling of points in θ space, here we choose a simpler approach. Specifically, for the classifiers considered in § 3.3, we look only at the first two moments of the effective pole vector averaged over θ . Mathematically, we define the mean pole vector and the associated covariance matrix respectively via

$$\bar{p}_k = \frac{1}{Q} \sum_{i=1}^Q p_k^{eff}(\theta_i) \quad (8)$$

$$R_k = \frac{1}{Q} \sum_{i=1}^Q (p_k^{eff}(\theta_i) - \bar{p}_k)(p_k^{eff}(\theta_i) - \bar{p}_k)^T \quad (9)$$

where the index i ranges over a grid of points in θ space. For the examples in § 3.4, we considered cases where θ was discretized into a 4D grid corresponding to five possible object depths, seven possible values for each of the three Euler angles and no horizontal variation of the target resulting in $Q = 3 \times 7 \times 7 \times 7 = 1715$. Thus, to summarize, the feature library we employ for classification based on pole estimates is comprised of one three dimensional vector and one three by three matrix for each target of interest and each sensing system under investigation.

3.2 Parameter Estimation

The first stage of processing is to estimate the parameters of our model for each target in the library. We actually do this twice: the first time to obtain parameter estimates to be used in a classifier based on data fit and the second time to obtain estimates for a pole-based classifier.

As explained more fully in § 3.3, the classifier based on the data vector directly makes use of the k th residual vector, $d - \hat{s}_k$, where $\hat{s}_k$ is some estimate of the signal vector for the k -th object in the

library. A first approach to generating $\hat{s}_k$ is to solve a problem similar to that of (7). We could then make use of the fact that *if* the data did in fact come from the k -th object (which is what we will ultimately be testing), *then* the poles should be $\bar{p}_k$. Thus under this scheme, we would not need to estimate the poles (and the expansion coefficients if we were to keep track of these as well) and we would only need to determine the elements of θ . While such an approach is feasible, it ignores the fact that we have information concerning the behavior of the pole estimates in the form of a mean vector *and* a covariance matrix. Hence, rather than fixing the poles in the estimation scheme we let them float but impose some bounds on their values in recognition of the fact that since we are going to be testing the fitness of the data to the k -th model the poles should be constrained to be close to the average pole value for this model. Specifically we solve the constrained optimization problem:

$$\begin{aligned} \hat{p}_{1,k}, \hat{a}_{1,k}, \hat{\theta}_{1,k} &= \arg \min_{p,a,\theta} \|y - s(p, a, \theta)\|_2^2 \\ \text{subject to } [p_k]_i &\in [[\bar{p}_k]_i - 2[\sigma_k]_i, [\bar{p}_k]_i + 2[\sigma_k]_i] \end{aligned} \quad (10)$$

with $[\bar{p}_k]_i$ the i th element of the vector $\bar{p}_k$ and $[\sigma_k]_i$ the square root of the i -th element along the diagonal of R_k . Hence $[\sigma_k]_i$ is the estimated standard deviation of $[\bar{p}_k]_i$. The above optimization problem essentially restricts the estimates of the poles to stay within plus or minus two standard deviations of their expected value. Again, the philosophy underlying this choice is that since we will be using these estimates to test the goodness of fit of the data vector to the k -th model we should encourage the parameter estimates to stay “close” to the model.

To solve the problem in (10), we use the nonlinear least squares solver available in Matlab’s Optimization Toolbox. This code makes use of a constrained Gauss-Newton algorithm for finding the a local minimizer of the objective function in the neighborhood of an initial guess. We initialize the algorithm as follows. For the poles we use $\bar{p}_k$ and we initially take the expansion coefficients to be equal to 1.0. The initial (x, y) location of the object is taken to be that point in space with the largest magnitude response in the data (a heuristic, but one which seems to work well) while the initial depth is 1.0m from the sensor. The Euler angles are initialized all to 0.0.

The second classifier discussed in § 3.3 is based on estimates of the poles. To allow the maximum flexibility in determining these quantities, we use the following, second estimation scheme in which the bound constraints are lifted

$$\hat{p}_{2,k}, \hat{a}_{2,k}, \hat{\theta}_{2,k} = \arg \min_{p,a,\theta} \|y - s(p, a, \theta)\|_2^2. \quad (11)$$

Again, a nonlinear least squares solver is used. This time, the algorithm is initialized with $\hat{p}_{1,k}, \hat{a}_{1,k}$, and $\hat{\theta}_{1,k}$. We have found that by constraining the poles in the first estimation stage, we obtain high quality estimates of the position and orientation parameters. These estimates are then used to obtain strong overall estimates of all relevant parameters in the second estimation step. Thus, this appears to provide an effective means to avoid the local minimum problem associated with non-linear parameter estimation problems.

Finally, we also employ bound constraints in (10) and (11) for the estimates of the elements of θ . The precise values of these constraints are problem dependent and described in § 3.4.

3.3 Classification

To motivate the different classification schemes we examine in this report, we define the residuals associates with the estimate of the k -th object in the library as:

$$\begin{aligned} \rho_k &= y - s(\hat{p}_{1,k}, \hat{a}_{1,k}, \hat{\theta}_{1,k}) = \delta s_k + \sigma w \\ \delta s_k &\equiv s(p_0, a_0, \theta_0) - s(\hat{p}_{1,k}, \hat{a}_{1,k}, \hat{\theta}_{1,k}) \end{aligned} \quad (12)$$

These residuals are comprised of two terms. The first, δs_k arises from errors in the signal vector caused by inaccuracies in the estimation of the model parameters. The second, σw , is the additive measurement noise.

Consider the low SNR case where ρ_k is dominated by the additive white noise. For a fixed σ , this scenario would arise when the object under investigation is deeply buried in which case the signal strength is significantly lowered due to the $1/r^3$ -type one-way amplitude loss seen in the field strength for problem of this type. In such situations, when the k -th object is in fact the true object, the statistical distribution of the residuals is dominated by the zero mean, white σw term and goodness of fit classification tests based on the zero mean nature of the residuals are expected to do well. Alternatively, when the k -th target is not the correct object, δs_k will be significant thereby adding to the mean of the ρ_k . Hence a classifier constructed to test that the mean of ρ_k is in fact zero would correctly reject this hypothesis.

Next consider cases of shallowly buried objects where the SNR is high so σw is fundamentally small. In these cases, even when the k -th object is in fact correct, the errors caused by even slight inaccuracies in the estimates of p , a and θ will dominate the noise so that a classifier wishing to exploit the expected zero mean nature of ρ_k under the true hypothesis will fail. These high SNR situations however are exactly those where we anticipate the ability to obtain good estimates of the pole structure. Hence, classification schemes based on the pole estimates themselves are expected to perform well here.

Motivated by the considerations outlined in the above discussion, here we define two statistics to be used for classification. The first is based on the data residuals and is taken as:

$$\epsilon_{1,k} = \frac{\|\rho_k\|_2^2}{\sqrt{N}\sigma} - \sqrt{N} \quad (13)$$

where N is the dimensionality of ρ_k . The normalization of the residuals in this way ensures that $\epsilon_{1,k}$ is asymptotically distributed as a zero mean, unit variance Gaussian random variable when the k corresponds to the true object. Thus classifiers based on tests of the closeness of ϵ_1 to zero are appropriate here. To generate a classifier based on the estimates of the poles we use

$$\epsilon_{2,k} = (\hat{p}_k - \bar{p}_k)^T R_k^{-1} (\hat{p}_k - \bar{p}_k). \quad (14)$$

This Mahalanobis-type distance metric is expected to be close to zero when k is true and larger than zero when the true object is not the k -th. Using $\epsilon_{1,k}$ and $\epsilon_{2,k}$ the classification rule is defined as follows: Choose the k^* object in the library if the magnitude of ϵ_{k^*} is less than a threshold τ else say that the object under investigation is clutter. Here k^* is selected in one of three ways. If we just want a classifier based on the residuals, we let k^* be the index of the smallest $\epsilon_{1,k}$. For a classifier based only on the pole estimates, k^* is that index minimizing $\epsilon_{2,k}$ over all k . Finally, a hybrid classifier is obtained by letting k^* be the minimizer of the union of $\epsilon_{1,k}$ and $\epsilon_{2,k}$. In the following section, we examine the performance of all three classification schemes.

3.4 Numerical Examples

Here we consider numerical tests of the classification methods described in § 3.3 using two different object libraries. The sensing system we simulated was comprised of co-located, square transmit and receive coils one half meter on a side. These coils sampled a one meter square area on an equally space 5×5 grid of measurement points. Frequency domain versions of the sensor collected complex valued data (in-phase and quadrature) at 20 logarithmically spaced frequencies between

	min p_1 (kHz)	max p_1 (kHz)	min p_2 (kHz)	max p_2 (kHz)
3 inch steel	3.4	5.6	3.6	10.3
6 inch steel	3.4	4.0	3.5	6.5
3 inch aluminum	0.11	0.19	0.12	0.34
6 inch aluminum	0.11	0.13	0.12	0.21

Table 1: Pole Characteristics for Objects in First Library


10Hz and 20 kHz. For time domain simulated we looked at a sensor which collected 40 equally spaced samples between 10e-6 and 1e-3 seconds.

The first object library is comprised of four objects: a three inch long by one inch diameter stainless steel cylinder, a six inch long by one inch diameter stainless steel cylinder, a three inch long by one inch diameter aluminum cylinder and a six inch long by one inch diameter aluminum cylinder. The “ground truth” model for the scattering characteristics of these objects was obtained using the method of [3] in which the dipole model was taken to be exact and four terms were kept in each of the λ_i summations in (3) and all expansion coefficients were taken to be one. As cylinders are symmetric about their primary axis, these objects have only two unique λ_i ’s. The range of values for the minimum and maximum poles for each of the four objects in the library are given in Table 1. The effective pole parameters for the frequency domain version of the sensing system as a function of objects position and orientation are plotted in pole-space in Fig. 3. Each point on the plot corresponds to the $(p_1^{eff}, p_2^{eff}, p_3^{eff})$ value computed from (7) for the i -th term in the summation of (8) or (9). The top plot in Fig. 3 shows the effective pole distributions for the steel targets while the bottom plot does the same but for the aluminum. Note that the axes for the two plots are distinctly different and that the points for the different objects cluster reasonably well in pole space. Thus as is evident, from the table as well as the figure, the pole characteristics of the steel objects are quite distinct from those of the aluminum; however the differences between the six and three inch versions of the same material are a bit more subtle. Hence it is anticipated that we will be able to distinguish shape better than precise object.


The second library of objects is comprised of M15, M21, TM62M mines and an aluminum disk. Their scattering characteristics were obtained from Table 1 in [10]. In that work, the authors extracted one pole and expansion coefficient per axis from actual sensor measurements. Because multiple measurements were made, they were able to compute the standard deviations of their estimates of the poles and expansion coefficients. Hence for our purposes, the target library was constructed by taking their published average pole values to build $\bar{p}_k$ and letting R_k be a diagonal matrix with the squares of the published standard deviations along the diagonal.

Unlike detection problem where the primary means of evaluating the algorithm is a plot of detection probability against false alarm rate [12], for the classification problem, there are other relevant outcomes of the test which need to be considered. Hence the following four probabilities are defined:

1. The detection probability which is the likelihood that we correctly identify the true object from the data: P_d .
2. The false detection probability which is the likelihood that we call a clutter item a target in the library: P_{fd}
3. The miss probability which is the likelihood that we call a target in the library clutter: P_m



(a) Effective Pole Distribution for Steel Objects. 'o' = 3 inch and 'x' = 6 inch



(b) Effective Pole Distribution for Aluminum Objects. 'o' = 3 inch and 'x' = 6 inch

Figure 3: Distributions of Effective Poles as a Function of Object Location and Orientation. All axes in kHz.

	Library 1	Library 2
Min. x coordinate	-.2 m	-.2 m
Max. x coordinate	-.2 m	-.2 m
Min. y coordinate	-.2 m	-.2 m
Max. y coordinate	-.2 m	-.2 m
Min. depth	-0.3 m	-0.3 m
Max. depth	-2.0 m	-.75 m
Min. Euler angle	0 rad.	$-\pi/6$ rad
Max. Euler angle	2π rad.	$\pi/6$ rad

Table 2: Bounds for Monte Carlo Analysis

4. The misclassification probability which is the likelihood that we call a true target in the library a different target in the library: P_{mc} .


For the examples described below, these quantities were estimated via Monte Carlo analysis. Specifically 250 separate runs were generated where we randomized over object type (four classes per library plus clutter), object location, orientation, sensor and noise. The bounds on the various quantities are provided in Table 2. For any Monte Carlo run, the value of the associated parameter was selected from a uniform distribution between the minimum and maximum bounds for that quantity. As is seen from Table 2, for the objects in the first library, we consider deeper burial depths and a wide range of orientations. This is to crudely model the UXO problem. For the second, mine-like library, we consider shallower burial depths and a more restricted range of orientations. The poles for clutter items were randomly generated using a uniform distribution whose upper and lower bounds were comparable to the upper and low bounds of the pole values across the library currently under test. Finally, we note that the bounds used in the Monte Carlo analysis are also used as constraints on the values of the location and rotation angle parameters for the optimization problems in (10) and (11).

To test the robustness of the approach we randomly perturbed the values of the true poles used to generate the data. For the first library each of the exact poles computed from the Carin model of [3] was perturbed by the addition of a zero-mean Gaussian random variable whose standard deviation was ten percent of the true pole value. For the second model to each of the mean pole values we added a zero mean Gaussian variable whose standard deviation was equal to the standard deviation computed in [10].


The results of the classifier using the first library are shown in Fig. 4. Frequency domain results are displayed in the top row and time domain results in the bottom. The first column displays our ability to exactly classify the target while the second has curves indicating our ability to tell aluminum from stainless steel. For visualization purposes, we plot P_d versus P_m and P_{fd} on the hopefully correct assumption that misclassification is a less serious error than either a miss or a false detection. In analogy to a more traditional receiver-operating characteristic, strong performance is obtained when P_d approaches 1.0 and the other two quantities go to 0.0. For the plots in Fig. 4, this implies that ideal performance corresponds to points in the upper back corner of the grid.

We see from the first row of Fig. 4 that for the frequency domain system under investigation the pole-only classifier is superior to the other two. For this classification scheme, we can classify the correctly determine the exact object about 80% of the time with a false detect rate below 1% and a miss rate also in the 1% range. In terms of classifying material type, detection performance jumps to around 90% with the other factors staying about the same. The second row of Fig. 4 shows


that for the time domain system we simulated, the pole only and the hybrid approach function comparably with detection, false detection and miss rates slightly worse than those seen in the frequency domain.




(a) Frequency Domain, Classification of Exact Object



(b) Frequency Domain, Classification of Object Material



(c) Time Domain, Classification of Exact Object



(d) Time Domain, Classification of Object Material

Figure 4: Performance Curves for First Library. 'o' = pole-based classifier, 'x' = residual-based classifier, '+' = classifier based on poles and residuals

Results for the frequency domain system and the second library are shown in Fig. 5. In this case, the mixed pole-residual classifier slightly outperforms the other two approaches. The correct classification rate here is again about 90% with a false detect probability below 1% but a slightly higher miss rate of about 4%.




Figure 5: Performance Curves for Second Library. 'o' = pole-based classifier, 'x' = residual-based classifier, '+' = classifier based on poles and residuals

3.5 Field Data Examples

The remainder of this section is devoted to an exposition of the results we obtained in applying our methods to three field data sets collected by two different SERDP contractors:

1. Data collected by Geophex Inc., of Raleigh NC using a frequency domain GEM-3 sensor to classify a collection pipes of differing lengths and materials buried on the premises of the company.
2. Data collected by Geophex Inc. of Raleigh NC using a frequency domain GEM-3 sensor during the Jefferson Proving Ground IV (JPG-IV) evaluation procedure.
3. Data collected over a limited number of UXO-type objects at Blossom Point MD by the Johns Hopkin's Applied Physics Laboratory using a version of their time domain, TEMIDS sensor.

It is important to note in evaluating the results presented in this report that not one of the three data sets were collected with the goal of validating the algorithms under consideration in this work. Rather, each of the three data sets was obtained by the vendor under separate contract either to validate their own algorithms or their sensor hardware. The primary impact of this situation is that the quantity of data which required to obtain the level of performance seen in simulation was not available in practice. In particular, the following two deficiencies were observed:

1. The training data required to build libraries which were robust to unknown orientation and location libraries was not present. For example, in the case of the GEM pipe data, a spectral signature taken at a single point in space for the object in only one orientation was all that was provided for building a library. This stands in stark contrast to the quantity of data used in the simulation study of the previous section.

2. The data were not rich enough. In the case of the JPG-IV GEM 3 data, only eight frequencies were collected at each of eight points in space. In many cases, these frequencies failed to cover the most relevant portions of the spectra of the UXO objects thereby making discrimination quite difficult. Moreover, it is far from clear that the sampling positions chosen for the JPG IV data collection provide a dense enough coverage to adequately resolve one target from another. In the case of the TEMIDS data, we were only really provided with test stand data for a couple of objects in a couple of orientations each. Each such data set was collected over a single line as opposed to a grid. There really was not any field data.

3.5.1 Geophex Pipe Data

Geophex buried a number of lengths of pipes of varying materials and sizes in a variety of orientations in a 10 meter by 10 meter area on the site of their offices. A GEM 3 sensor held about 20 cm above the ground collected data about once every 5 cm along lines spaced approximately 25 cm apart using dead reckoning to ascertain position.

A total of nine classes of objects were buried in the field. To construct the library, Geophex placed one example of each class oriented vertically in a test stand and took one spectral measurement with the GEM 3 centered over the top of the object. Thus, while the field data were sampled fairly densely, not nearly enough training data was available to construct a pole library as thoroughly as we would have liked. In Table 3, we summarize the classes of objects as well as the estimated poles. This table indicates that, for example, targets L1 and L2 are both steel cylinders of diameter 15.5 cm and length 50.8 cm. The three poles estimated for this class of objects was 4.51 kHz, 0.42 kHz, and 0.05 kHz.

An immediate indication that a single spectra of data is not sufficient for building the library is indicated by the pole values in Table 3. Since these are all cylindrical objects, we would expect only two distinct values for the poles. While most of the objects do have one pole significantly larger than the others, the remaining two are generally not particularly close. In effect, providing only one “look” at the target in terms of spatial location of the sensor as well as the orientation of the object does not provide a rich enough data set to really allow us to reliably distinguish three poles.

Despite these difficulties, a plot of these poles in 3-space as shown in Figure 6 does indicate that we can expect some discriminability based on these pole estimates. Specifically, the aluminum targets, S2 and M6, are well separated in pole-space from most of the steel objects. Also, similarly shaped targets are clustered together and far from other, dissimilar objects; e.g., (S4 is close to S3, M3 is close to M11, but S4/S3 and M3/M1 are relatively far from one another).

Due in large part to the scarcity of training data, we simplified the classification method. First, we used only the pole-base statistic, ϵ_2 in (14). Second, the covariance matrix, R , was set to the identity in this case as we did not have enough training data to estimate this matrix. The results of processing are shown in Table 4 which indicates that we were in fact able to successfully classify all of the objects in this study.

3.5.2 Geophex JPG-IV Data

The JPG IV GEM 3 set was comprised of data collected from 25 spatial locations arranged in a 5 by 5 grid over and around each of 164 different targets. In all cases, only eight frequencies were used (30, 90, 210, 510, 1350, 3570, 9210, and 23970 Hz). Of the 164 objects, 45 were UXO and the remainder were fragment. Within the UXO class, there were a total of 7 types of projectiles (20 mm HE, 57 mm, 76mm HE, 90 mm AP, 105mm, 152 mm, 155 mm) and three classes of mortars

Target ID	Description	Average Poles (KHz)
L1, L2	15.5 × 50.8 Steel Pipe	[0.0820 , 1.3656 , 0.0820]
M1, M2, M5	7.90 × 15.7 Steel Pipe	[0.0734 , 0.9229 , 0.0734]
M3, M4, M7	6.40 × 30.5 Steel Pipe	[0.2076 , 2.9342 , 0.2076]
S1, S10	2.00 × 10.2 Steel Pipe	[0.5544 , 6.1501 , 0.5544]
S3, S5, S6, S8	4.10 × 15.2 Steel Pipe	[0.3934 , 0.3934 , 4.2486]
S4, S7	4.10 × 10.2 Steel Pipe	[0.7471 , 0.7471 , 3.1921]
M6	6.4 × 30.5 Aluminum Pipe	[0.4897 , 0.0524 , 0.4897]
S2, S9	2.30 × 15.2 Aluminum Pipe	[0.9770 , 0.3042 , 0.9770]
S11, S12	2.30 × 15.2 Copper Pipe	[0.2922 , 0.6531 , 0.2922]

Table 3: GEM 3 Pipe Data Pole Library: All dimensions in the "Description Column" are in cm. All poles are in units of kHz.





Figure 6: Pole Locations for Pipe Data Library

True Target	Orientation	Estimated	Result	Highest Score	Next Highest
L1	Horizontal	L2	Correct	0.0001	0.2954
L2	45 Degrees	M1	Misclassification	0.0220	0.0847
M1	Vertical	M1	Correct	0.0083	0.0861
M2	45 Degrees	M1	Correct	0.0870	0.1513
M5	Horizontal	M1	Correct	0.0076	0.0172
M3	Horizontal	M3	Correct	0.0002	0.0008
M4	45 Degrees	M4	Correct	0.0991	0.3370
M7	Horizontal	M1	Misclassification	0.0827	0.1086
M6	Horizontal	M6	Correct	0.0082	0.0873
S1	Horizontal	M3	Misclassification	0.0027	0.0071
S10	Horizontal	S1	Corrcet	0.1246	0.2152
S9	Horizontal	M6	Misclassification	0.0001	0.0010
S2	45 Degrees	S2	Correct	0.0188	0.0232
S3	Vertical	M1	Misclassification	0.0339	0.0494
S5	Horizontal	M1	Misclassification	0.0083	0.0813
S6	Horizontal	M1	Misclassification	0.0819	0.8554
S8	Horizontal	S8	Correct	0.0083	0.0183
S7	Horizontal	S7	Correct	0.0603	0.0782
S4	Horizontal	S4	Correct	0.0654	0.0824
S11	45 Degrees	S11	Correct	0.0011	0.0044
S12	Horizontal	S11	Correct	0.0004	1.2615

Table 4: GEM 3 Pipe Data Results

(4.2, 60mm, 81 mm) each buried in a number of different orientations and configurations (e.g. with and without fuse). To build the models, we were provided with data collected from a sensor in only one spatial location directly above each of 72 targets. Within this training set were all of the UXO objects as well as examples of the clutter in various orientations.

While the diversity of orientation is an improvement relative to the training data for the Pipes Data Set, the lack of spatial diversity still hampers the ability to truly determine the performance limits of the proposed algorithms. Moreover, for many targets, the frequencies used to probe the subsurface failed to capture the salient structure of the spectral profiles of the objects. That is, Assumption 2 on page 8 was violated. For example, in Figure 7, we plot the in-phase and quadrature data collected directly over the center of two different objects each in two different orientations. Relative to the structure indicated in Figure 2, for these data sets we capture little of those characteristics which we expect would be of use in telling UXO from clutter or one type of UXO from another. The in-phase portions of these signals show only the rising edge of the response and never the inflection. Similarly, we do not observe the peak in the quadrature. While this lack of detail was not universal in the data set, it was observed in many of the signatures collected over the different targets. Given that such data do not resolve those features which we expect to provide distinguishability, it is not likely that classification performance will be strong for this set of data.



(a) Solid line: Horizontal orientation rotated 315 degrees. Dashed line: Horizontal orientation rotated 135 degrees.

(b) Solid line: Horizontal orientation rotated 45 degrees. Dashed line: Horizontal orientation rotated 215 degrees.

Figure 7: Samples of JPG-IV GEM-3 Data for two objects in different orientations.

In using this data set, we needed to modify the processing routine in a number of ways to obtain any type of reasonable results:

1. Of the 25 spectral responses taken in the vicinity of each target, only the one over the center proved to be of any significant value. Thus only eight complex values data points were available for estimating the parameters of the model.
2. Because the data did not, in general, resolve the relevant structure in the data and since we had only a few data samples, it was exceedingly difficult to stably estimate the model parameters (poles, location, amplitudes, and angles). Thus, in the model, we fixed the rotation angles to

zero degrees. This meant that for this data set an each target-orientation pair was treated as a separate “object” both for building the library and for purposes of classification. For example, an 105 mm at 45 degrees was treated as being distinct from a 105 mm oriented vertically.

Even with these simplifications it was basically impossible to distinguish one class of UXO from another. Even separating clutter from UXO was rather difficult when the whole data set was taken into consideration. For example, we considered the following simple, somewhat *ad hoc* strategy designed to allow us to declare an object as a UXO in a way which provides some flexibility in terms of goodness of fit between the estimated poles for that object and those poles in the library:

1. Estimate poles, $\hat{p}$ and other parameters
2. For each of the 64 sets of poles in the library, p_i compute $d_i = \|\log_{10} \hat{p} - \log_{10} p_i\|_2^2$ where the logarithm of a vector is meant to indicate the logarithm for each of the components.
3. Rank order the d_i from smallest to largest.
4. If target item appears within the first N elements of the rank ordered list declare that object a UXO item.

For a given N , we compute the fraction of correct declarations and the fraction of false declarations (false digs essentially) out of the 164 available data sets. The results as N is varied provides a ROC-like curve which is plotted in Figure 8. The dashed line in this plot is the chance diagonal and N is increasing as we move from lower left to upper right. This plot indicates that even for the simple decision rule we chose to employ, our chances of successfully telling UXO from non-UXO are no better than 50%.

Subsequent to obtaining these results, we were told that the GEM is known to be appropriate for classifying only certain objects buried over a well defined range of depths. Thus, the evaluation of the performance of our approach over the full 164 set of signatures is in a sense overly pessimistic. Table 5 summarizes the reduced set of targets we used for a more constrained classification exercise as well as the estimated pole structures. The 3D plots of this pole library is shown in Figure 9 where we see that unfortunately, many of the objects tend to cluster tightly in pole-space. Hence we expect discrimination to be challenging. This intuition is born out by examining the results presented in Table 6. This table indicates first that it is possible to use this GEM data set to do limited classification successfully. Comparing the highest and next highest scores here with those seen in the pipe data set in Table 4, we see that the margins here are much smaller than they were for the pipes. In other words, the separability of targets from the JPG-IV data is far smaller than was the case for the pipes. So, while we *can* do well, the small distance to the next score is indicative of a case where the results are not likely to be terribly robust. Whether this difficulty is related to the sparse nature of the training data or is inherent to our pole-based classification method can only be determined pending the collection, processing, and analysis of a more complete set of training and testing data.

3.5.3 TEMIDS Data

Finally we examine the results of our algorithms for the processing of time domain data collected at Blossom Point by the TEMIDS sensor from John’s Hopkins Applied Physics Laboratory. Three sets of data from this extensive campaign were the most relevant for use with our algorithms. Each




Figure 8: ROC-like curve for JPG-IV data set

Target ID and Description	Average Poles (KHz)
20 MM	[3.6280, 6.8990, 3.6280]
57 MM	[21.725, 4.6114, 4.6114]
60 MM	[9.5678, 4.4538, 5.0825]
81 MM	[7.0851, 7.0851, 2.9756]
90 MM	[2.9232, 2.9232, 3.6104]
105 MM	[6.1350, 6.1350, 1.1696]

Table 5: JPG-IV GEM 3 Data Pole Library: All dimensions in the "Description Column" are in cm. All poles are in units of kHz.




Figure 9: Pole Locations for JPG-IV Data Library

Number	True Target	Estimated	Result	Highest Score	Next Highest
51	60MM	60MM	Correct	0.6572	0.7736
3	81MM	81MM	Correct	0.0034	0.0321
33	20MM	20MM	Correct	0.0182	0.2898
141	90MM	90MM	Correct	0.0057	0.0058
148	105MM	105MM	Correct	0.0167	0.0506
118	57MM	90MM	Misclassification	0.0001	0.0004
124	81MM	81MM	Correct	0.0078	0.0148
104	90MM	90MM	Correct	0.0038	0.3720

Table 6: JPG-IV Data Results

set was obtained by translating the sensor in 2 inch increments from 4 inches to 0 inches along a line that was displaced from the object by y inches where y was set to 0, 2, 4, 6, 8, and 10.

Three such objects were interrogated in this manner. The first was an 81 mm mortar where the sensor was first location 23 cm above the target and then a second set of data were acquired with the sensor at a height of 28 cm. Next, data were taken for a 61 mm mortar with the sensor 16 cm above the target. Again, two passes were made. The first was with the mortar in a vertical up position and the second in a vertical down position. Finally, a “steel egg” target was used in three configurations: the egg held vertically with the sensor 6 cm above target, the egg held vertically with the sensor 9 cm above the target, and the egg held horizontally with the sensor 9 cm above the object.

As there was no clear way of distinguishing “training” from “test” data in this case (since we had no independent field measurements), we used the TEMIDS data to examine the stability of our pole estimates for the three different objects as we move the sensor in y . In theory, the model says these poles estimates should be the same regardless of the relative locations of the object and sensor. Due to modeling errors and approximations, in practice they will not be the same, but rather cluster in hopefully distinct regions of pole-space. Thus, for each displacement y (a defined at the beginning of this section), we use the data at 0, 2, and 4 inches to estimate to parameters of the model. For the mortars, this gives us a total of 12 parameters sets (6 values of y by two configurations of either the sensor or the object). For the egg, we have 18 parameter sets. As with the JPG-IV data, here again, in no case do we have a sufficient number of “looks” to justify the estimations of the rotation angles. Hence our processing protocol for the TEMIDS data is to extract object location, poles, and pole amplitudes with the rotation angles fixed.

The results of this exercise are shown in Figure 10. The two types of yellow markers are the pole estimates for the 61 mm mortar with the sensor 16 cm and 28 cm above the target. Similarly, the two collections of red markers are the results for the 81 mm mortar while the steel egg estimates are shown in black. Generally, these results are fairly encouraging. The tight clustering of the 81 mm poles indicate that the stability of these estimates. The fact that this cluster and the 61 mm cluster are so well separated bodes well for the distinguishability of these objects under more stressing tests.

4 Processing When Sensor Locations Are Uncertain

4.1 General Problem Formulation

We begin with a general formulation of the problem we wish to consider. For simplicity we assume that a mono-static sensor stops at $i = 1, 2, \dots, N$ nominal locations in space. At each location, $r_i = [x_i \ y_i \ z_i]^T$, a length M vector of data (e.g. samples in a time trace or the in-phase and quadrature components in a frequency sweep), d_i is acquired. The sensing problem is to extract from the d_i the values of a collection of parameters such as poles, orientation angles, etc. of the buried objects needed for characterization and classification. These parameters are assembled into a column vector θ and are mathematically related to the data via a sensor model, $s(\theta, r_i) \equiv s_i(\theta)$, which as indicated is also dependent on r_i for $i = 1, 2, \dots, N$. For what follows, we assume here that s is in fact differentiable with respect to the elements of θ and r_i .

As indicated in § 1, it is often the case that the locations of the sensor are not known precisely. To model this situation, r_i is taken to be the sum of two components: $r_{0,i}$, the nominal or expected position of the sensor, and δr_i , the perturbation to $r_{0,i}$. The perturbation is such that $r_i = r_{0,i} + \delta r_i \in \mathcal{S}_i$ where $\mathcal{S}_i$ is a region of space whose size is dictated by the predefined tolerances




Figure 10: Pole estimates for TEMIDS data. Yellow markers are for the 61 mm mortar. Red markers are for the 81 mm mortar and the black represent data for the steel egg.

provided with the data. For simplicity, here we assume that these tolerances are the same for all positions. In this report, we consider algorithms based on two choices for the $\mathcal{S}_i$: boxes and ellipsoids. Mathematically, a box of size X in the x direction, Y in y , and Z in z centered at a point $r_0 = [x_0 \ y_0 \ z_0]^T$ is defined as

$$\mathcal{B}(r_0) = \left\{ r = [x \ y \ z]^T \left| |x - x_0| < \frac{X}{2}, |y - y_0| < \frac{Y}{2}, |z - z_0| < \frac{Z}{2} \right. \right\}. \quad (15)$$

Similarly, an ellipsoid centered at r_0 with axes of length X , Y and Z is:

$$\mathcal{E}(r_0) = \left\{ r = [x \ y \ z]^T \left| \left(\frac{x - x_0}{X} \right)^2 + \left(\frac{y - y_0}{Y} \right)^2 + \left(\frac{z - z_0}{Z} \right)^2 \leq 1 \right. \right\}. \quad (16)$$

Given these definitions, the problem we face is to determine the “optimal” choice of θ given the data, d_i , and knowledge of X , Y , and Z .

4.2 A Min-Max Solution Approach

A mathematical formulation of this problem requires first a definition of optimality. As is typical for problems such as these, here we take optimality to be defined in the sense of the minimum (over all θ) norm of the square error difference between the data observed and the prediction of the data provided by the model s . That is an iterative algorithm is constructed which adjusts θ to find a (local) minimum of

$$J(\theta, r_i) = \sum_i \|d_i - s(\theta, r_i)\|_2^2 \quad (17)$$

where, for a vector x , $\|x\|_2^2 \equiv x^T x$.

The problem now is how to specify the r_i in the s_i in (17). In most all cases, the r_i are set equal to their nominal values, $r_{0,i}$ and a least squares estimate of θ is found. Here though we seek to augment the minimum error notion of optimality in a manner which easily incorporates the additional knowledge we have concerning the geometry of the $\mathcal{S}_i$. Essentially, the problem we have here is to determine a set of primary parameters, θ , in the presence of a collection of “nuisance” parameters, the r_i , which are restricted to exist in known regions of space. One common methods for solving such problems is via a “min-max” formulation. This amounts to selecting that θ which minimizes the worst error, as measured by $J(\theta, r_i)$, as r_i ranges over $\mathcal{S}_i$. Formally, $\hat{\theta}$, our estimate of θ , is defined as

$$\hat{\theta} = \arg \min_{\theta} \max_{r_1 \in \mathcal{S}_1, \dots, r_N \in \mathcal{S}_N} \sum_i \|d_i - s(\theta, r_i)\|_2^2 \quad (18)$$

To obtain a tractable solution to the problem, we now make the assumption that the δr_i are sufficiently small so that we can write

$$s(\theta, r_i) = s(\theta, r_{i,0}) + A_i(\theta) \delta r_i \quad (19)$$

where A_i is the matrix whose (j, k) element is the partial derivative of the j th component of the vector s_i with respect to the k th element of δr_i . Essentially, (19) is a Taylor series expansion of the model about the nominal location of the sensor. Substitution of (19) into (18) yields

$$\hat{\theta} = \arg \min_{\theta} J_1(\theta) \quad (20)$$

$$J_1(\theta) = \max_{r_1 \in \mathcal{S}_1, \dots, r_N \in \mathcal{S}_N} \sum_i \|d_i(\theta) - \mathbf{A}_i(\theta) \delta r_i\|_2^2 \quad (21)$$

with $d_i(\theta) \equiv d_i - s(\theta, r_{i,0})$.

Were it not for the inner maximization, the typical approach to determine $\hat{\theta}$ requires the use of an iterative optimization routine such as steepest decent, non-linear conjugate gradient, or a Gauss-Newton type method which exploits the underlying least squares structure of the problem. These techniques all require that the cost function, J_1 is differentiable at least once in the elements of θ . For our problem, since J_1 itself is the solution to an optimization problem we cannot guarantee its smoothness. Even if we could, computing gradients, Hessians and such is not particularly feasible. Thus we adopt a two level approach to determining $\hat{\theta}$. First, a general purpose non-smooth optimization scheme (Matlab's Nelder-Mead method in fact) is used to solve the "outer" minimization problem. This leaves us with having to solve the "inner" maximization problem at each iteration of the Nelder-Mead method.

In general, (21) represents a numerically intensive optimization problem. Here however, there is significant structure which greatly simplifies the situation. Two observations serve to ease the burden of the inner maximization:

1. First, because the perturbations at one sensor location are independent of those at any other location, the inner maximization in (20) separates and we obtain

$$\max_{r_1 \in \mathcal{S}_1, \dots, r_N \in \mathcal{S}_N} \sum_i \|d_i(\theta) - \mathbf{A}_i(\theta) \delta r_i\|_2^2 = \sum_i \max_{r_i \in \mathcal{S}_i} \|d_i(\theta) - \mathbf{A}_i(\theta) \delta r_i\|_2^2 \quad (22)$$

Hence rather than having to solve one large maximization problem involving $3N$ variables, we can, in parallel, solve N small problems each with only 3 variables.

2. Second, both classes of $\mathcal{S}_i$ (boxes and ellipsoids) are convex shapes. Because $\|\cdot\|_2^2$ is a convex function of its argument, we are guaranteed that the maxima for each of our N problems is achieved on the boundary of $\mathcal{S}_i$ [13, Chapter2]. The implication of this fact for the box and ellipsoidal regions lead then to the two basic algorithms developed in this report.

4.2.1 Boxes

When the regions $\mathcal{S}_i$ are boxes, the existence of the solution to the basic maximization problem on the boundary of the box can in fact be strengthened. Specifically, the solution must exist on one of the corners of the box. Thus, each of the maximization problems in (22) can be solved simply by evaluating the cost at each of the eight corners of the associated box and finding the maximum value.

4.2.2 Ellipsoids

The case of ellipsoids is a bit more involved and a bit more interesting. Since each of the maximization problems in (22) is structurally identical, we drop the subscript i as well as the explicit dependence on θ and here consider the basic problem

$$\max_{\delta r \in \mathcal{S}} \|y - A\delta r\|_2^2 \quad (23)$$

First, we observe from (16) that the boundary of set $\mathcal{S}_i$ can be written as $(\delta r)^T D^T D(\delta r) = 1$ with D the diagonal matrix containing X , Y , and Z on the main diagonal. Letting $x = D\delta r$ and $C = AD^{-1}$ yields the following normalized form of the constrained optimization problem we are considering:

$$\max_{x^T x = 1} \|y - Cx\|_2^2 = \max_x \{ \|y - Cx\|_2^2 + \lambda (x^T x - 1) \} \quad (24)$$

where we have introduced the Lagrange multiplier, λ .

To make some headway here, we factor C using the singular value decomposition as $U\Sigma V^T$. In our case C is an $M \times 3$ matrix so that U is an $M \times M$ orthonormal matrix, V is a 3×3 orthonormal matrix and Σ is an $M \times 3$ matrix of all zeros except the first three elements on the main diagonal which are the non-zero singular values, σ_i $i = 1, 2, 3$, of C . Now, introducing the variables $z = V^T x$ and $w = U^T y$, making use of the structure of Σ and the fact that an orthonormal change of variables leaves the two norm of the resulting vector invariant allows us to write (24) as

$$\max_{z, \lambda} \left\{ \sum_{i=1}^3 (w_i - \sigma_i z_i)^2 + \sum_{i=4}^N (w_i)^2 + \lambda (z_1^2 + z_2^2 + z_3^2 - 1) \right\} \quad (25)$$

The first order necessary conditions for a solutions for a solution to (25) are obtained by setting the derivatives of (25) with respect to z_1, z_2, z_3 , and λ equal to zero. Following this procedures yields:

$$w_i - \sigma_i z_i + 2\lambda z_i = 0 \quad i = 1, 2, 3 \quad (26)$$

$$\sum_{i=1}^3 z_i^2 = 1. \quad (27)$$

Solving for z_i in terms of λ in (26) gives

$$z_i = \frac{w_i}{\sigma_i - 2\lambda}. \quad (28)$$

substituting this result into (27) gives the optimal λ as the solution to

$$\sum_{i=1}^3 \frac{w_i^2}{(\sigma_i - 2\lambda)^2} = 1. \quad (29)$$

A simple root finding scheme is used to solve (29) for λ . Substituting the result into (28) gives the optimal z_i from which we obtain the optimal δr as $D^{-1}Vz$ with V the matrix of left singular vectors of C and D defined after (23). From this we can easily find the value of the cost function in (23).

4.3 Numerical Examples

Under this contract we were able only to begin the analysis of this approach for dealing with sensor positional uncertainties. A simulation Monte Carlo experiment was run where four targets were to be distinguished given frequency domain data. As with the simulation in § 3.4, we used a four-pole-per-axis model to generate the data and a one-pole-per-axis model for identification. Two steel-type and two aluminum-type targets were simulated. The average values for the four poles for each of the three axes are provided in Table 7. Ten frequencies were used: 30, 90, 150, 270, 570, 1230, 2610, 5430, 11430, and 20070 Hz. We simulated a scenario where 200 of each of the four types of objects were buried in a 45 m $\times$ 45 m site and data were acquired on a nominal 25 cm grid. The position of the sensor at each location was randomly perturbed in a rectangle with shape parameters given by $X = 5$ cm, $Y = 4$ cm, and $Z = 3$ cm. Variations in each of the three dimensions were selected according to independent uniform probability distributions.

In Table 8 we display the confusion matrix for the case where we processed the data using our original algorithm while the results of the ellipsoidal min-max approach are shown in Table 9. In

Object	Avg. pole 1	Avg. pole 2	Avg. pole 3
Steel 1	4.246	8.922	11.179
Steel 2	4.688	9.477	11.315
Aluminum 1	0.132	3.923	5.753
Aluminum 2	0.152	4.101	5.712

Table 7: Pole structure for Uncertain Position Simulation. All average pole values are in units of kHz.

each of these tables, the number in row i and column j is the number of times the object at the top of row j was declared when in fact the true object was that given at the left hand side of row i . These results show first that for the level of error induced in the sensor position, the original approach does in fact perform well. Still, in most cases where there were mistakes, the min-max approach reduced the error rate by around 50%.

	S1	S2	A1	A2
S1	191	9	0	0
S2	5	195	0	0
A1	0	0	187	13
A2	0	0	6	194

Table 8: Confusion Matrix for Original Algorithm

	S1	S2	A1	A2
S1	194	6	0	0
S2	2	198	0	0
A1	0	0	193	7
A2	0	0	3	197

Table 9: Confusion Matrix for Min-Max Approach

4.4 Field Data Examples

Finally, we tested both the box and ellipsoid min-max methods on the Geophex pipe data. In both cases, discrimination results were the same as with the original algorithm (i.e. 100% classification). Improvements however were seen in our ability to localize the objects when positional errors were taken into consideration. The results are summarized in Table 10.

5 Conclusions and Future Work

In this report, we have explored a number of statistically-motivated, model-based options for object classification using spatially sampled time and frequency domain EMI data. Preliminary results using synthetic data are promising and indicate there is much work to be done in the future. Of specific interest are the following items:

1. **Model calibration.** To support the acquisition and processing of sensor data for algorithm training and for the processing of field data, the sensor model used by our methods must be

Object	True position (x, y, z) m	Original Algorithm (x, y, z) m	Ellipsoid Min-Max (x, y, z) m	Box Min-Max (x, y, z) m
L2	(2.75,2.75,0.99)	(2.81,2.81,0.99)	(2.78,2.76,0.99)	(2.78,2.83,0.93)
S11	(0.75,0.75,0.14)	(0.77,0.79,0.10)	(0.76,0.77,0.13)	(0.75,0.76,0.12)
M3	(8.75,1.25,0.50)	(8.69,1.21,0.50)	(8.76,1.23,0.50)	Results not run

Table 10: True and Estimated Positions for Objects in Geophex Pipe Dataset

properly calibrated to the systems collecting the data. For example, the exact fields produced by the sensing coils need to be properly modeled. Here we assume field produced by co-located rectangular coils which was not the case for either the GEM or TEMIDS systems.

2. **Training data collection** Not one of the three data sets were collected with the goal of validating the algorithms under consideration in this work. Rather, each was obtained by the vendor under separate contract either to validate their own algorithms or their sensor hardware. Issues included (a) an insufficient quantity of training data, (b) the collection of data over a frequency band inappropriate for discriminating targets from clutter or one target class from another, (c) or the collection of test data for a small number of targets over only a few usable points in space. More comprehensive and conclusive validation of the methods described in this report require the collection of detailed spatial “scan” data for a number of known UXO and clutter objects in a variety of orientations and distances from the scan plane of the sensor.
3. **Development and validation of a processing-driven theory of sensor optimization.** Tools from statistical estimation and decision theory such as Cramer-Rao and Chernoff bounds could be used with our calibrated sensor models to obtained bounds on the quantity of data as well as how finely sampled such data needs to be in space, time, and frequency to meet classification requirements or to determine that such requirements are infeasible.
4. **Test data collection and validation** An extensive data collection campaign should be undertaken to fully validate the methods developed in this work. In addition to extending the training library as required, we would be quite interested in collecting data from known calibration sites as well as unknown blind test sites and submit the results of our methods for scoring by SERDP, JUXOCO, or the cognizant authority.
5. **Algorithmic Extensions** A number of technical extensions and generalizations could be pursued to improve the performance of the methods developed here. These issues include the most effective means of parameterizing and estimating object orientation; the use of kernel density techniques for improved classification performance; and the development of more effective optimization approaches for dealing with the problem of uncertain sensor location.

References

- [1] B. Barrow and H. H. Nelson. Model-based characterization of electromagnetic induction signatures obtained with the MTADS electromagnetic array. *IEEE Trans. Geoscience and Remote Sensing*, 39(6):1279–1285, June 2001.
- [2] T. H. Bell, B. J. Barrow, and J. T. Miller. Subsurface discrimination using electromagnetic induction sensors. *IEEE Trans. Geoscience and Remote Sensing*, 39(6):1286–1293, June 2001.

- [3] Lawrence Carin, Haitao Yu, Yacine Dalichaouch, and Carl Baum. On the wideband EMI response of a rotationally symmetric permeable and conducting target. *IEEE Trans. on Geoscience and Remote Sensing*, preprint, 2000. In review.
- [4] L. Collins, P. Gao, and L. Carin. An improved Bayesian decision theoretic approach for landmine detection. *IEEE Trans. on Geoscience and Remote Sensing*, 38:1352–1361, May 2000.
- [5] Y. Das, J. E. McFee, J. Toews, and G. C. Stuart. Analysis of an electromagnetic induction detector for real-time localization of buried objects. *IEEE Trans. on Geoscience and Remote Sensing*, 28:278–287, May 1990.
- [6] Norbert Geng, Carl E. Baum, and Lawrence Carin. On the low-frequency natural response of conducting and permeable targets. *IEEE Trans. on Geoscience and Remote Sensing*, 37(1):347–359, January 1999.
- [7] Sadri Hassani. *Foundations of Mathematical Physics*. Allyn and Bacon, 1991.
- [8] S. J. Norton and I. J. Won. Identification of buried unexploded ordinance from broadband induction data. *IEEE Trans. Geoscience and Remote Sensing*, 39(10):2253–2261, October 2001.
- [9] L. S. Riggs, J. E. Mooney, and D. E. Lawrence. Identification of metallic mine-like objects using low frequency magnetic fields. *IEEE Trans. Geoscience and Remote Sensing*, 39(1):56–66, January 2001.
- [10] Gary Sower, John Endsley, and Ed Christy. Discrimination of metal land mines from metal clutter: Results of field tests. In *SPIE Conference on Detect. of Mines and Minelike Targets: IV*, pages 78–88, April 1999.
- [11] S. L. Tatum and L. M. Collins. A comparison of algorithms for subsurface target detection and identification using time-domain electromagnetic induction data. *IEEE Trans. Geoscience and Remote Sensing*, 39(6):1299–1306, June 2001.
- [12] Harry L. Van Trees. *Detection, Estimation and Modulation Theory: Part I*. John Wiley, New York, 1968.
- [13] Willard I. Zangwill. *Nonlinear Programming: A Unified Approach*. Prentice Hall, 1969.