Radio Science, Volume ???, Number , Pages 1–28,

Statistical Classification of Buried Objects from Spatially Sampled Time or Frequency Domain EMI Data

Ashley B. Tarokh

Center for Subsurface Sensing and Image Systems, Dept. of Electrical and Computer Engineering, Northeastern University, Boston, MA

Eric L. Miller

Center for Subsurface Sensing and Image Systems, Dept. of Electrical and Computer Engineering, Northeastern University, Boston, MA

I. J. Won

Geophex, Ltd., Raleigh, NC

Haoping Huang

Geophex, Ltd., Raleigh, NC

Ashley B. Tarokh and Eric L. Miller are with the Center for Subsurface Sensing and Image Systems, Dept. of Electrical and Computer Engineering, Northeastern University, Boston, MA 02115, {btarokh,elmiller}@ece.neu.edu

I. J. Won and Haoping Huang are with Geophex, Ltd., 605 Mercury St., Raleigh, NC 27603

Abstract.

Methods for classifying objects based on spatially sampled electromagnetic induction data taken in the time or frequency domain are developed and analyzed. To deal with nuisance parameters associated with the position of the object relative to the sensor as well as the object orientation, a computationally tractable physical model explicit in these unknowns is developed. The model is also parameterized by a collection of decay constants (or equivalently Laplace-plane poles) whose values in theory are independent of object position and orientation. These poles are used as features for classification. The overall algorithm consists of two stages. First we estimate the values of the unknown parameters and then we do classification. 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. The library can be constructed using either simulated or calibration data. A maximum likelihood method is developed and analyzed for the problem of joint pole, location, and orientation parameter determination. Here, we examine and compare two classification schemes. The first classification method is based on data residuals generated from estimated signal parameters. This scheme performs well in low SNR cases. The second is based on estimated pole values themselves, which performs well in high SNR cases. We validate our methods on both simulated and field data taken from frequency and time domain sensors.

1. Introduction

The problem of detecting and classifying buried objects using electromagnetic induction (EMI) based sensing technologies has received considerable attention in recent years in a range of application areas including unexploded ordinance (UXO) and landmine remediation. In the last decade, 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 requires the development of advanced signal processing methods which are tied directly to the physical model of the sensor. In this paper, 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 [*Collins et al.*, 2000] 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 [*Tantum and Collins*, 2001]. The library itself is built from data signature vectors [*Riggs et al.*, 2001; *Tantum and Collins*, 2001] or feature vectors from all targets of interest in a given application [*Tantum and Collins*, 2001; *Bell et al.*, 2001; *Barrow*

and Nelson, 2001]. The classifier takes as the selected object that element of the library which in a sense "best fits" the data or the features.

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 [*Barringer Research Ltd.* 1976; *Barringer Research Ltd.* 1979; *Das et al.* 1984; *Das et al.*, 1990], 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 or 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 and the manner in which the dipole model is employed in the processing. Aside from [*Bell et al.*, 2001], classification algorithms have been concerned strictly with either time-domain EMI sensors [*Barrow and Nelson*, 2001; *Tantum and Collins*, 2001] or frequency domain [*Norton and Won*, 2001;

Riggs et al., 2001] but not both. In [*Riggs et al.*, 2001; *Tantum and Collins*, 2001], the authors employ a parametric model for the PAPFs which takes the form of a sum of decaying exponentials [*Sower et al.*, 1999; *Collins et al.*, 2000; *Carin et al.*, 2000] 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 constants are then used as the features for classification. Alternatively, the methods of [*Bell et al.*, 2001; *Barrow et al.*, 2001; *Norton and Won*, 2001] 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 [*Bell et al.*, 2001; *Barrow and Nelson*, 2001], for example, the location of the object is estimated as part of the processing, while in [*Riggs et al.*, 2001; *Norton and Won*, 2001], 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 [*Tantum and Collins*, 2001] while in [*Riggs et al.*, 2001], a different target signature is used for each orientation of each target in the library. Finally, in [*Bell et al.*, 2001; *Barrow and Nelson*, 2001; *Norton and Won*, 2001], the rotation matrix is in fact determined as part of the eigen-analysis of the polarizability tensor, however the authors do not map this rotation matrix back to an explicit orientation of the object in space.

In this work we construct a two classification methods based on a physical model that is the fusion of the dipole scattering model in [*Das et al.*, 1990] and a parametric PAPF model elucidated in [*Carin et al.*, 2000]. This model is analytical in the parameters of the PAPF, the (x, y, z) location of the object, and the three Euler angles [*Hassani*, 1991] describing the rotation matrix. The relatively simple closed form nature of the model with respect to these parameters leads us to classification methods in which the orientation and location of the object are explicitly estimated 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 [*Geng et al.*, 1999] indicates that in theory the PAPF are comprised of an infinite number of decaying exponentials in time domain, which, in the frequency domain, translate to an infinite number of one pole transfer functions. Unfortunately, it is both impossible and unnecessary to estimate an infinite number (or even a large number) of poles or decay rates. First, pole estimation in the presence of noise is known to be a very delicate signal processing problem [*Apollo*, 1991]. 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 [*Geng et al.*, 1999]. Thus in this paper, 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 is similar to the one or two pole methods of [*Geng et al.*, 1999]; however where these approaches disregard the information in the coefficients multiplying the decaying exponentials in time (or poles in the frequency domain), 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. 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. For the purpose of comparison, classifiers based on data residuals are also constructed.

The physical model employed in this work was introduced in [*Miller*, 2001], where a preliminary set of results demonstrated the effectiveness of the model for a limited set of simulated data. In this work, we study the model and the processing schemes more extensively and develop an improved understanding of the physical model. In particular, we investigate the effect of SNR on the performance of our residual-based and pole-based classifiers. We find that the residual-based classifiers perform well under low SNR, whereas under high SNR the pole-based classifiers perform well. We validate all claims using simulated as well as field data.

The remainder of this paper is organized as follows. In § 2 the scattering model is discussed. A description of the processing method is given in § 3. Experiments using both simulated as well as real data from time and frequency domain sensors is presented in § 4. Conclusions and directions for future work are the subjects of § 5.

2. Physical Model

We consider a combination of the physical EMI model in [*Das et al.*, 1990] describing the scattering of low frequency electromagnetic radiation by spherical or spheroidal objects with the model in [*Geng et al.*, 1999] which 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 coils (not necessarily co-located) with sides of length $2A$ and $2B$ respectively [1]. 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 j-th position is

$$y_{j,k} = {g_j}^T R^T \Lambda_k R f_j + \sigma n_{j,k} \equiv s_{j,k} + \sigma n_{j,k}. \tag{1}$$

Here, g is a 3×1 vector holding the (x, y, z) components of the magnetic field produced at r_0 by a current I flowing through the receiver coil. Also, f represents the excitation

fields vector evaluated at the dipole position. The variable $n_{j,k}$ is a zero mean, unit variance random variable and σ is the standard deviation of the assumed additive white Gaussian measurement noise. Functional forms for f and g are provided in Appendix A of [*Das et al.*, 1999]. The quantity Λ_k is the complex-valued polarizability tensor for the kth frequency and has the following form:

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

where $\lambda_1, \lambda_2, \lambda_3$ are associated with one of each of the principle axes of the object and ω is the operating frequency. Replacing ω by t will give the time domain version of this equation. Here we consider a form of that model provided by [*Geng et al.*, 1999]:

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

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:

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

with $u(t)$ the unit step function. For cylinders and disks, [*Carin et al.*, 2000] provides a fast numerical method for computing the $p_{i,l}$. 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 delta function in time. 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 is a rotation matrix which orients the object in the space and is 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 [*Hassani*, 1991] and explicitly takes the form:

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

The magnetic polarizability tensor Λ can be diagonalized by R. Each element of the resulting matrix holds the scattering characteristics of the object along each of the three-principle axis.

Going back to (1), by gathering the data together from all sensors, the overall model in compact notation can be written as

$$y(p, a, \theta) = s(p, a, \theta) + \sigma n, \tag{6}$$

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

As stated, this model assumes that the object behaves electromagnetically like a dipole. The three λ_i's fully summarize the scattering behavior of the object and only depend on the size, shape, and material of the object and *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 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 descent, conjugate gradient, or Newton type of optimization scheme.

While the utility of the model described here has been validated using real sensor data [*Das et al.*, 1990; *Sower et al.*, 1999; *Riggs et al.*, 2001], as described in § 1, 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 [*Riggs et al.*, 2001; *Carin et al.*, 2000; *Sower et al.*, 1999]. 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 for the physical model described above 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.

3. Processing Method

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 two 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. 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

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

where θ_0 holds the true position and orientation information, y_k is the true data vector, and p and p^{eff} are vectors of three pole parameters, one per λ_i. The symbol "$\hat{}$" 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. 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

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

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

where the index i ranges over a grid of points in θ space. Thus, 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, $y - \hat{s}_k$, where $\hat{s}_k$ is an 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:

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

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 a nonlinear least squares solver that 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

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

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 problem of reaching a local minimum, which is often associated with non-linear parameter estimation problems.

3.3. Classification

Given the model-based approach we have developed in this paper, there are two natural classification schemes that can be employed. The first is based on the idea that if the kth model is the true model then according to (6) $\rho_k = y - s(\hat{p}_{1,k}, \hat{a}_{1,k}, \hat{\theta}_{1,k})$ should be a zero mean, uncorrelated Gaussian random vector. The second compares the pole estimates, $\hat{p}_2$ to the elements of the library in a way which accounts for the known "spread" in the pole estimates for a given target

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, $\delta s_k = s(p_0, a_0, \theta_0) - s(\hat{p}_{1,k}, \hat{a}_{1,k}, \hat{\theta}_{1,k})$ 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:

$$\epsilon_{1,k} = \frac{\|\rho_k\|_2^2}{\sqrt{N}\sigma} - \sqrt{N} \tag{12}$$

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

$$\epsilon_{2,k} = (\hat{p}_k - \bar{p}_k)^T R_k^{-1} (\hat{p}_k - \bar{p}_k). \tag{13}$$

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 for which the magnitude of ϵ_{k^*} is minimum. Here k^* is selected in one of two 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.

4. Numerical Examples

Here we consider numerical tests of the classification methods described in § 3 using two different object libraries. The first sensing system 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 spaced 5×5 grid of measurement points. Frequency domain versions of the sensor collected complex valued data (in-phase and quadrature) at 30 logarithmically spaced frequencies between 10 Hz and 30 kHz. For time domain, we employed a sensor that collected 60 equally spaced samples between 10^{-6} s and 10^{-3} s.

The simulated object library is comprised of four objects: a 3 inch long by one inch diameter stainless steel cylinder (S1), a 6 inch long by one inch diameter stainless

steel cylinder (S2), a 3 inch long by one inch diameter aluminum cylinder (A1), and a 6 inch long by one inch diameter aluminum steel cylinder (A2). The "ground truth" model for the scattering characteristics of these objects was obtained using the method of [*Carin et al.*, 2000] in which the dipole model was taken to be exact, four terms were kept in each of the λ_i summations in (4), and all expansion coefficients were taken to be one. As cylinders are symmetric about their primary axis, these objects have two unique $\lambda_i^{'}$ 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 the object position and orientation are plotted in pole-space in Fig. 2. 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 plots in this figure show the effective pole distributions for the steel targets while the bottom plots do the same but for the aluminum targets. Note that the axes for the top and bottom 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 material type better than precise object.

A Monte Carlo approach is used to analyze the performance of our algorithms for this library. Two sets of simulations were carried out to compare the performance

of our classifiers for high and low SNR cases. In each case 100 separate data sets were generated where we randomized uniformly over object type, object location, orientation and additive sensor noise (by selecting a new set of parameters according to a uniform distribution at each Monte Carlo round). The bounds on the various quantities are provided in Table 2. The corresponding values of the Euler angles were chosen uniformly over their full range of definition (either 0 to 2π or 0 to π depending on the angle [*Hassani*, 1991]). For our purposes SNR is calculated according to:

$$\text{SNR} = 10\log_{10}\frac{||y||_2^2}{\sigma^2 l_y} \tag{14}$$

where y denotes the signal vector, l_y is its length, and σ^2 denotes the noise variance. At each Monte Carlo run, we select the burial depths of the targets randomly from a specific interval, while keeping the noise variance σ^2 fixed. For the low SNR case this interval consists of burial depths in the range of $1.00m$ to $2.00m$, while for the high SNR case this interval consists of burial depths in the range of $0.1m$ to $0.3m$.

The classification results of this example, for each of the above cases, using the two classification methods, are summarized in the confusion matrices of Tables 3 through 10. The i, j-th element of each matrix demonstrates the number of times that object i was the true target and object j was declared by our processing scheme. The results presented in these tables verify our claim with respect to the two different classifiers. As shown here, when operating under high SNR (on the order of 20 dB or higher), pole-based classifiers operate well, whereas residual-based classifiers operate well under low SNR (on the order of 0 dB).

The misclassification presented in our confusion matrices can be attributed to the fact that the scattering characteristics of the two aluminum (or steel) targets are very similar. Note that the target material is rarely misclassified under high SNR, while under low SNR, the residual-based algorithm, though misclassifying target type more often, performs significantly better than the pole-based method. It is also clear that similar results are obtained from our algorithm regardless of the characteristics of the sensor employed (time or frequency domain).

The second object library is comprised of 9 targets. Their scattering characteristics were obtained from the work of *Huang and Won* [2003]. The data for this sensing system comes from the GEM-3 sensor developed by Geophex. This sensor has been used successfully in many environmental sites and can detect small targets, such as UXO and landmines, providing high spatial resolution [*Huang and Won*, 2003]. The current GEM-3 operates in a bandwidth from 30 Hz to 24 kHz. The Geophex test site in Raleigh, NC, specially designed by Geophex Inc., is a $10m \times 10m$ test site. This test site, detailed in [*Huang and Won*, 2003] contains a total of 21 metal pipes of various types, lengths and diameters. The authors of [*Huang and Won*, 2003] obtained the ground truth data by placing the targets at two different orientations at a single known position. Because of the lack of multiple measurements we were not able to generate the standard deviations of the estimates of the poles. Hence for our purposes R_k was assumed to be the identity matrix. The effective pole parameters of these objects are plotted in pole-space in Fig. 3. Each point on the plot reflects the $(\bar{p}_1, \bar{p}_2, \bar{p}_3)$ value computed from the corresponding target in the library.

We have separated the aforementioned nine targets into three sets of targets, designated by letters L, M and S. The letter L corresponds to targets that are deeply buried at a depth ranging from 90 cm to 110 cm. The letter M corresponds to targets buried at a depth ranging from 50 cm to 80 cm. Finally, the letter S corresponds to shallowly buried targets at a depth ranging from 10 cm to 30 cm.

The scattering characteristics of these targets were also acquired by the authors [*Huang and Won*, 2003] at a line spacing of 25 cm using the dead reckoning method at a height of about 20 cm above the ground. The GEM-3, in this case collected about 8 to 10 data points per second, which resulted in a data interval of about 15 cm [*Huang and Won*, 2003]. The position error for such data could be as big as 20 cm due to uneven walking speed and incorrect walking path. Also the errors associated with the sensor height could be more than 5 cm [*Huang and Won*, 2003].

Table 11 summarizes the classification results for this target library. For the measurement data corresponding to each true target, the column labeled "Depth" corresponds to the burial depth of each target in cm. Clearly, the results presented in this table suggest a correlation between object depth and the dominating classification method. Specifically, we find that our L-series targets are classified perfectly using the residual-based method (and poorly using the pole-based method); the M-series targets show mixed results, with correct classification under the residual-based method alone for particularly deeply buried targets, both methods classifying for an object of moderate depth, and correct classification under the pole-based method alone for a particularly shallowly buried target; for the majority of the S-series targets, correct

classification is achieved under the pole-based method alone. These results closely correspond to our claim that the pole-based classifier is better suited for shallowly buried objects, while the residual-based classifier is better suited for deeply buried targets.

5. Conclusions and Future Work

In this paper, 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. Extending the classifier algorithm to be able to distinguish objects that are not in the library (clutter items). This step can be very important for the UXO and demining problems where there is a strong desire to correctly reject clutter items.

2. Testing the approach on field data from other EMI sensors.

3. Analytical performance evaluation of the proposed methods. By obtaining closed form or computable values or bounds on the relevant probabilities, we can begin to explore issues such as optimizing sensor configuration (time gates, frequencies collected, spatial sampling rate) to maximize performance.

4. More rigorous error analysis. It would be useful to take a more careful, quantitative look at the role of parameter estimation errors as well as sensor noise in the ultimate performance of the classification scheme.

5. Alternate classification techniques. The pole-based classification statistic, $\epsilon_{2,k}$ in (13) implicitly assumes that the effective pole parameters as a function of object location and orientation cluster into an ellipsoid. The plots in Fig. 2 show that while the effective poles for different objects do cluster, they are far from ellipsoidal. Thus it may be useful to consider alternate classification schemes such as a linear discriminant function.

Acknowledgments. This work was supported by funding from the Strategic Environmental Research and Development Program under Project CU-1217.

The authors also wish to thank Dr. Lawrence Carin of Duke University for providing us with the data used in this work.

Notes

1. We assume square coils for convenience; with a little work, the model could be generalized to circular coils.

References

Steven J. Apollo. Maximum likelihood estimation of exponentials contained in signal dependent noises Ph.D. thesis, Univ. of Texas, Arlington, 1991.

Barringer Research Ltd. Report on studies into magnetic and electro-magnetic range clearance techniques. Contract Report 8SR5-0069, Defense Research Establishment Suffield, P.O. 4000, Station Main, Medicine Hat AB, Canada T1A8K6, December 1976.

Barringer Research Ltd. User's manual DRES range clearance prototype system. Contract Report 8SR77-00150, Defense Research Establishment Suffield, P.O. 4000, Station Main, Medicine Hat AB, Canada T1A8K6, March 1979.

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.

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.

Lawrence Carin, Haitao Yu, Yacine Dalichaouch, Alexander R. Perry, Peter V. Czipott and Carl Baum. On the wideband EMI response of a rotationally symmetric permeable and conducting target. *IEEE Trans. on Geoscience and Remote Sensing*, 39(6):1206–1213, June 2001.

R. Chesney, Y. Das, J. E. McFee, M. Ito Identification of metallic spheroids by classification of their electromagnetic responses. *IEEE Trans. on Pattern Recognition and Machine Intelligence*, PAMI-6:809–820, 1984.

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.

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.

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.

Sadri Hassani. *Foundations of Mathematical Physics*. Allyn and Bacon, 1991.

H. Huang and I. J. Won. Characterization of uxo-like targets using broadband electromagnetic induction sensors. *IEEE Trans. Geoscience and Remote Sensing*, 41(3):640–651, March 2003.

E. L. Miller. On some options for statistical classification of buried objects from spatially sampled time or frequency domain EMI data. *SPIE Conference on Detect. of Mines and Minelike Targets: VI*, Vol 4394, April 2001.

S. J. Norton and I. J. Won. Identification of buried unexploded ordinance from broadband electromagnetic induction data. *IEEE Trans. Geoscience and Remote Sensing*, 39(10):2253–2261, October 2001.

L.S. Riggs, J. E. Mooney, and D. E. Lawrence. Identification of metallic mine-like objects using low frequency magnetic feilds. *IEEE Trans. on Geoscience and Remote Sensing*, 39(1):56–66, January 2001.

Gary Sower, John Endsley, and Ed Christy. Discrimination of metal land mines from metal clutter: Results of field tests. *SPIE Conference on Detect. of Mines and Minelike Targets: IV*, pages 78–88, April 1999.

S. L. Tantum 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.

Target	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

Min x coordinate	0.25 m
Max x coordinate	0.75 m
Min y coordinate	0.25 m
Max y coordinate	0.75 m
Min Depth	-0.10 m
Max Depth	-2.00 m

Table 2. Bounds for Monte Carlo Analysis

	Estimated A1	Estimated A2	Estimated S1	Estimated S2
True A1	81	19	0	0
True A2	16	84	0	0
True S1	0	0	92	8
True S2	0	0	22	78

Table 3. Frequency Domain Classification Results for Monte Carlo Analysis Based on Pole Estimates Under High SNR.

	Estimated A1	Estimated A2	Estimated S1	Estimated S2
True A1	56	32	7	5
True A2	29	57	3	11
True S1	2	0	89	9
True S2	2	0	35	63

Table 4. Frequency Domain Classification Results for Monte Carlo Analysis Based on Data Residual Under High SNR.

	Estimated A1	Estimated A2	Estimated S1	Estimated S2
True A1	34	19	45	2
True A2	34	11	43	12
True S1	15	19	21	45
True S2	21	16	54	9

Table 5. Frequency Domain Classification Results for Monte Carlo Analysis Based on Pole Estimates Under Low SNR.

	Estimated A1	Estimated A2	Estimated S1	Estimated S2
True A1	59	27	8	6
True A2	33	52	9	6
True S1	13	11	64	12
True S2	8	14	20	58

Table 6. Frequency Domain Classification Results for Monte Carlo Analysis Based on Data Residual Under Low SNR.

	Estimated A1	Estimated A2	Estimated S1	Estimated S2
True A1	79	21	0	0
True A2	18	82	0	0
True S1	0	0	77	23
True S2	0	0	28	72

Table 7. Time Domain Classification Results for Monte Carlo Analysis Based on Pole Estimates Under High SNR.

	Estimated A1	Estimated A2	Estimated S1	Estimated S2
True A1	54	25	9	12
True A2	37	41	7	15
True S1	4	1	61	39
True S2	3	0	33	67

Table 8. Time Domain Classification Results for Monte Carlo Analysis Based on Data Residual Under High SNR.

	Estimated A1	Estimated A2	Estimated S1	Estimated S2
True A1	20	16	60	4
True A2	32	17	38	13
True S1	23	21	32	24
True S2	28	18	47	7

Table 9. Time Domain Classification Results for Monte Carlo Analysis Based on Pole Estimates Under Low SNR.

	Estimated A1	Estimated A2	Estimated S1	Estimated S2
True A1	54	22	19	5
True A2	17	57	15	11
True S1	1	3	62	34
True S2	1	4	40	55

Table 10. Time Domain Classification Results for Monte Carlo Analysis Based on Data Residual Under Low SNR.

True Target	Depth	Pole-Based Classifier		Residual-Based Classifier	
		Estimated Type	Result	Estimated Type	Result
L1	100	M6	Misclassification	L1L2	Correct
L2	92-110	M3M4M7	Misclassification	L1L2	Correct
M1	47-70	S3S5S6S8	Misclassification	M1M2M5	Correct
M2	64-80	S4S7	Misclassification	M1M2M5	Correct
M3	50	M6	Misclassification	M3M4M7	Correct
M4	49-60	L1L2	Misclassification	M3M4M7	Correct
M5	70	M1M2M5	Correct	M1M2M5	Correct
M6	50	L1L2	Misclassification	M6	Correct
M7	60	M3M4M7	Correct	M6	Misclassification
S1	10	S1S10	Correct	S3S5S6S8	Misclassification
S2	10-15	S2S9	Correct	M3M4M7	Misclassification
S3	22-30	S3S5S6S8	Correct	M1M2M5	Misclassification
S4	30	S4S7	Correct	S4S7	Correct
S5	25-30	S3S5S6S8	Correct	M3M4M7	Misclassification
S6	30	S3S5S6S8	Correct	M6	Misclassification
S7	30	S4S7	Correct	S4S7	Correct
S8	15-20	S3S5S6S8	Correct	S4S7	Misclassification
S9	15	S2S9	Correct	S4S7	Misclassification
S10	10	S1S10	Correct	M6	Misclassification
S11	15-20	S11S12	Correct	M6	Misclassification
S12	15	S11S12	Correct	S1S10	Misclassification

Table 11. GEM 3 Pipe Data Results

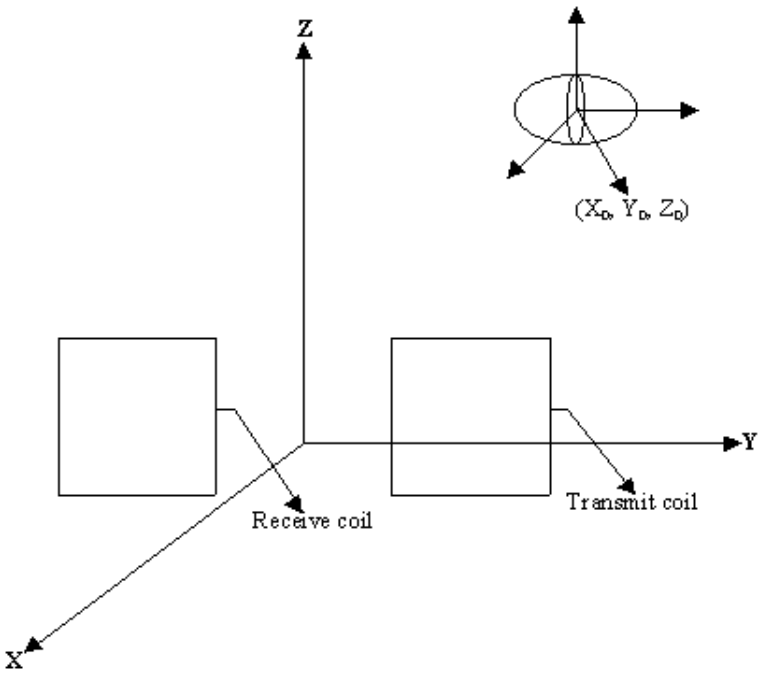

Figure 1. One sensor comprising sensor coils and target object.

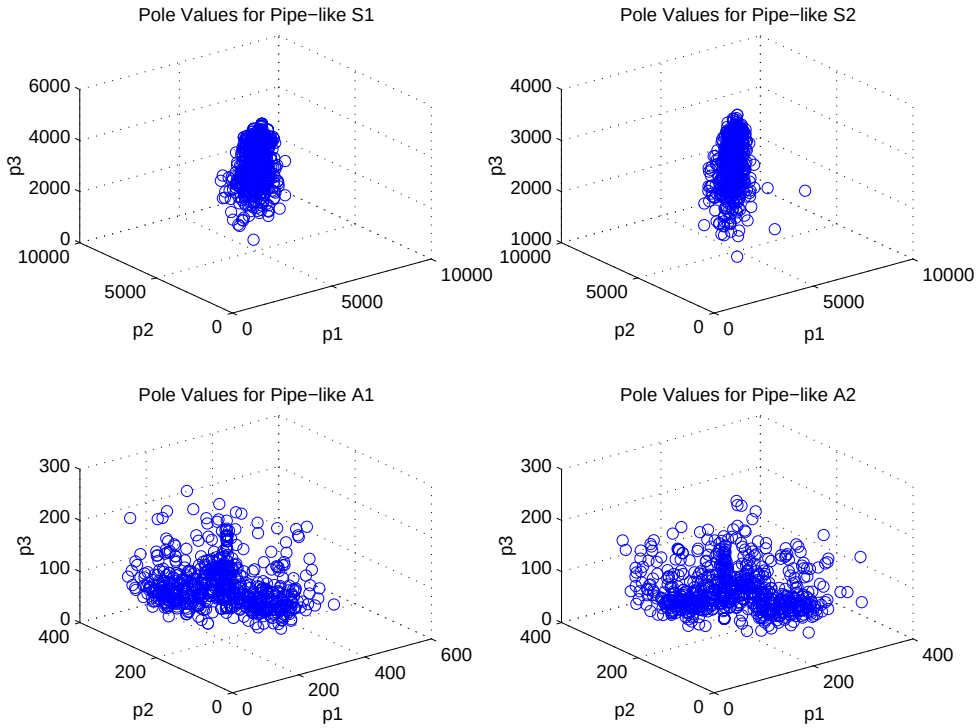

Figure 2. Effective Pole Distribution for Steel and Aluminum Objects in Frequency Domain.

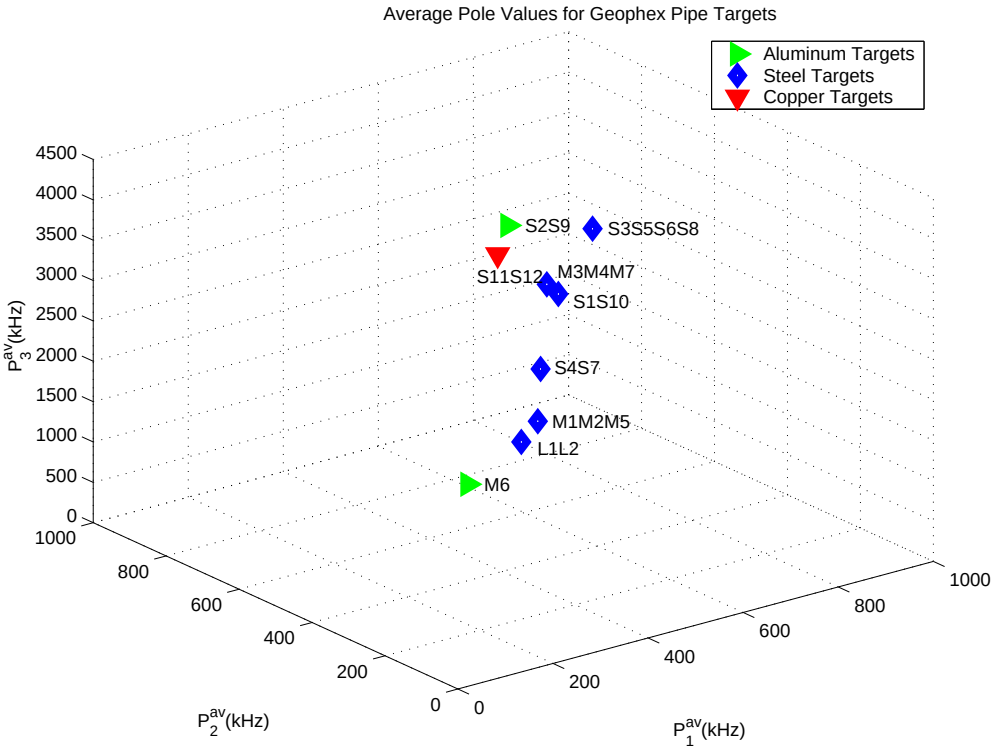

Figure 3. Pole Locations for GEM-3 Data Library.