Journal of Applied Geophysics 85 (2012) 59–67

Contents lists available at SciVerse ScienceDirect

Journal of Applied Geophysics

journal homepage: www.elsevier.com/locate/jappgeo

Regularizing dipole polarizabilities in time-domain electromagnetic inversion

Laurens Beran [a,b,*], Stephen Billings [b], Douglas Oldenburg [a]

[a] Department of Earth and Ocean Sciences, University of British Columbia, 6339 Stores Road, Vancouver, B.C., Canada V6T 1Z4
[b] Sky Research Inc., Suite 112A, 2386 East Mall Vancouver, B.C., Canada V6T 1Z3

ARTICLE INFO

Article history:
Received 6 February 2012
Accepted 26 June 2012
Available online 5 July 2012

Keywords:
Unexploded ordnance
Discrimination
Electromagnetics
Regularized inversion

ABSTRACT

Recent advances in time-domain electromagnetic (TEM) sensors have dramatically improved discrimination of buried unexploded ordnance (UXO). In contrast to commercial standard mono-static sensors, the multi-static, multi-component geometries of next generation TEM sensors provide diverse excitations of a detected target. Inversion of observed data using the parametric TEM dipole model typically produces well-constrained estimates that can subsequently be inputted into a discrimination algorithm. In particular, the principal dipole polarizabilities provide information about target size and shape. Shape is represented by two transverse polarizabilities orthogonal to a target's axis of symmetry.

Equality of transverse polarizabilities is diagnostic of an axisymmetric body of revolution and so has been proposed as a useful feature to discriminate between axisymmetric UXO and non-axisymmetric metallic clutter. Here we show that estimated transverse polarizabilities can sometimes be poorly constrained in an inversion of multi-static TEM data. This motivates our development of a regularized inversion algorithm that penalizes the deviation between transverse polarizabilities. We then develop an extension of the support vector machine (SVM) classifier that uses all models obtained via regularized inversion to make discrimination decisions. This approach achieves the best performance of all candidate discrimination algorithms applied to a number of real data sets.

© 2012 Elsevier B.V. All rights reserved.

1. Introduction

The 2003 Defense Science Board report on unexploded ordnance (UXO) projected that a reduction in false alarm rates from 100:1 to 10:1 would save $36 billion on remediation projects within the United States (Delaney and Etter, 2003). This cost reduction was expected to be achieved by improvements in sensor and data processing technologies. These goals have been met, and sometimes exceeded, in recent demonstration projects conducted by the Environmental Security Technology Certification Program (ESTCP) (Billings et al., 2010, Prouty et al., 2011, Shubitidze et al., 2011, Steinhurst et al., 2010). Advances in electromagnetic (EM) sensors have been crucial to these successes: the data provided by multi-static, multi-component EM platforms are much improved inputs into the inversion and discrimination algorithms applied to this problem. Fig. 1 compares the geometry and time channels of the commercial standard Geonics EM-61 with two multi-static EM instruments designed for UXO discrimination. The Time-domain Electro-Magnetic Multi-sensor Towed Array Detection System (TEMTADS, shown in Fig. 1b) is composed of an array of 25 horizontal transmitter loops arranged in a 5×5 grid, with horizontal receivers measuring the vertical field arranged concentrically to these transmitters (Steinhurst et al., 2010). The transmitters are fired sequentially and the secondary field response is recorded in all receivers simultaneously. This configuration provides a diverse data set which is better able to constrain target parameters. The MetalMapper sensor (Fig. 1c) has also greatly improved the reliability of estimated parameters by transmitting orthogonal primary fields and measuring all components of the secondary field in multiple tri-axial receivers (Prouty et al., 2011). Both MetalMapper and TEMTADS systems are deployed in a static (or cued) mode: previously-detected targets are interrogated with a stationary sensor. A mono-static sensor such as the EM61 has to be moved to several positions to illuminate a detected target and success of this process is critically dependent on accurate geolocation (Tantum et al., 2008, Tarokh and Miller, 2007). The single static acquisition required when deploying TEMTADS or MetalMapper sensors thereby removes the need for accurate geolocation.

In this paper we study parameter estimation with multi-static sensor data. We show that while these data generally support inversion and discrimination, in some cases parameter estimates can be poorly constrained. This motivates regularization of the inverse problem and here we seek models corresponding to targets with axial symmetry. This property is diagnostic of many UXO and so provides a useful feature for classifying targets of interest. We investigate methods for model selection and develop a technique that uses all models from regularized inversion to make discrimination decisions. We also extend our regularized inversion technique to multi-object scenarios to deal with overlapping target anomalies. Finally, we apply our techniques to data sets from ESTCP live-site demonstrations and compare discrimination performance.

* Corresponding author at: Tel.: +1 541 525 5149; fax: +1 541 488 4606.
E-mail address: lberan@eos.ubc.ca (L. Beran).

0926-9851/$ – see front matter © 2012 Elsevier B.V. All rights reserved.
doi:10.1016/j.jappgeo.2012.06.011

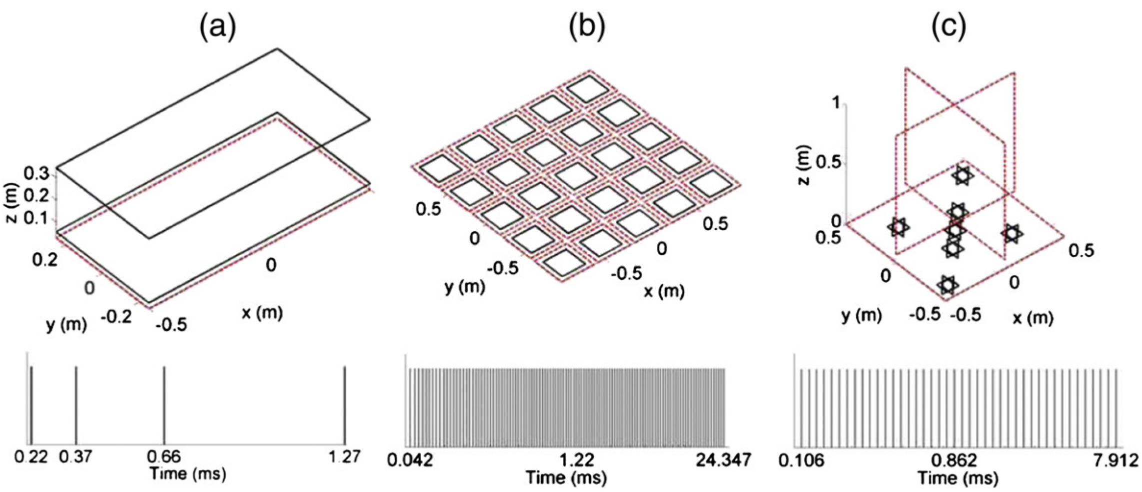

Fig. 1. (a) EM-61, (b) TEMTADS, and (c) MetalMapper sensors for unexploded ordnance detection and discrimination. Top row shows sensor geometry, with solid and dashed lines indicating receiver and transmitter coils, respectively. Bottom row shows logarithmically-spaced time channels.

2. Parameter estimation with the dipole model

The time (or frequency) dependent dipole model is essential to most electromagnetic data processing for UXO discrimination (Bell et al., 2001b, Pasion and Oldenburg, 2001, Zhang et al., 2003). While Shubitidze and collaborators have also achieved excellent discrimination results with more "physically complete" models (e.g. Shubitidze et al., 2011), our focus in this article is on improving discrimination with the dipole model in challenging scenarios.

The dipole model provides a simple parametric representation of the response of a confined conductor. The rate of change of the secondary magnetic field is computed as

$$\frac{\partial \mathbf{B_s}}{\partial t}(\mathbf{r},t) = \frac{p(t)}{r^3}(3(\hat{\mathbf{p}}(t)\cdot\hat{\mathbf{r}})\hat{\mathbf{r}} - \hat{\mathbf{p}}(t)) \tag{1}$$

with $\mathbf{r} = r\hat{\mathbf{r}}$ the separation between target and observation location, and $\mathbf{p}(t) = p(t)\hat{\mathbf{p}}(t)$ a time-varying dipole moment

$$\mathbf{p}(t) = \frac{1}{\mu_o}\mathbf{P}(t)\cdot\mathbf{B}_o. \tag{2}$$

The induced dipole is the projection of the primary field $\mathbf{B_o}$ onto the target's polarizability tensor $\mathbf{P}(t)$ (Bell et al., 2001b). Here the elements of the polarizability tensor ($P_{ij}(t)$) represent the convolution of the target's B-field impulse response ($\tilde{\mathbf{P}}(t)$) with the transmitter waveform $i(t)$ (Wait, 1982)

$$P_{ij}(t) = \frac{\partial}{\partial t}\int_{-\infty}^{\infty} \tilde{P}_{ij}\left(t^{'} - t\right) i\left(t^{'}\right) dt^{'}. \tag{3}$$

The polarizability tensor is assumed to be symmetric and positive definite and so it can be decomposed as

$$\mathbf{P}(t) = \mathbf{A}^T\mathbf{L}(t)\mathbf{A} \tag{4}$$

with $\mathbf{A}$ an orthogonal matrix which rotates the coordinate system from geographic coordinates to a local, body centered coordinate system. The diagonal eigenvalue matrix $\mathbf{L}(t)$ contains the principal polarizabilities $L_i(t)$ $(i=1,2,3)$, which are assumed to be independent of target orientation and location.

Features derived from the dipole model have been successfully used to discriminate between targets of interest (TOI) and non-hazardous metallic clutter. In particular the amplitude and decay of the principal polarizabilities provide a simple parameter set for discrimination (Beran et al., 2011). For a sensor with N channels, these target features can be computed as

$$\begin{aligned} \text{amplitude} &= \sum_{j=1}^{N} L_{total}\left(t_j\right) \\ \text{decay}\left(t_k, t_j\right) &= \frac{L_{total}(t_k)}{L_{total}\left(t_j\right)} \end{aligned} \tag{5}$$

with the total polarizability $L_{total}(t_j)$ defined as the sum of the polarizabilities at each time channel

$$L_{total}\left(t_j\right) = \sum_{i=1}^{3} L_i\left(t_j\right). \tag{6}$$

The decay parameter is a ratio of total polarizabilities at selected channels. For $t_k > t_j$ we have decay$(t_k, t_j) < 1$, so that a larger decay parameter is diagnostic of a slow decaying total polarizability.

The amplitude and decay parameters are physically meaningful because, to first order, a confined conductor can be modeled as a simple LR loop which is inductively coupled to transmitters and receivers on the surface. The current response of this loop is a decaying exponential which is fully described by an amplitude and time constant (West and Macnae, 1991). In practice, many larger UXO (e.g. 105 mm projectiles, 81 mm mortars) produce large amplitude, slow decaying polarizabilities relative to metallic debris. However, at more challenging sites with smaller items (e.g. 37 mm projectiles, fuzes), amplitude and decay parameters alone may not be sufficient to reliably discriminate UXO from clutter of similar size.

The TEM dipole model generalizes the simple circuit model to account for target shape. A ferrous, prolate (rod-like) target with rotational symmetry about its principal axis will produce equal transverse (secondary and tertiary) polarizabilities (Bell et al., 2001b). For ferrous items, internal demagnetization is strongly affected by the shape of the object. The result is that the strength of the induced dipoles along the semi-minor axes is reduced, so that transverse polarizabilities are smaller in magnitude than axial polarizabilities.

Most ordnance are composed primarily of steel and can be treated as bodies of revolution (Bell et al., 2001a, Shubitidze et al., 2002, Zhang et al., 2003) and so equality of transverse polarizabilities has been proposed as a useful feature for discriminating between TOI and irregularly-shaped clutter. However, in practice it has been difficult to reliably estimate target shape from mono-static TEM data. This is because single loop, vertical-component transmitters and receivers often cannot adequately interrogate the transverse response of buried targets. Fig. 2 illustrates this effect for a spherical target illuminated by a mono-static Geonics EM-61 sensor (geometry and time channels are shown in Fig. 1).

In the case of a sphere all polarizabilities are equal, and so here we define the primary, axial polarizability as aligned along the z axis. The corresponding induced field is maximal when the sensor is directly over the target. However to excite the transverse (x and y) responses of the target the sensor must be positioned with a horizontal stand-off from the target. Assuming an approximately dipolar field radiated from transmitter and target, the secondary field (Eq. (1)) decays as $1/r^6$ with increasing sensor–target separation r. For this reason, the axial polarizability response dominates the measured data in Fig. 2. Data which are sensitive to transverse polarizabilities therefore tend to have low signal to noise, particularly for vertically-oriented targets. This geometric effect is exacerbated by the reduced amplitude of transverse polarizabilities for ferrous ordnance. These factors confounded early attempts to estimate target shape from mono-static sensors (e.g. Bell et al., 2001b).

Multi-static TEM sensors designed for UXO detection have helped address these limitations. Fig. 3 shows the components of the dipole response for the TEMTADS sensor over the same target as in Fig. 2. The data received for transmitters immediately adjacent to the center transmitter are primarily sensitive to a combination of the transverse excitations. Inversion of these data therefore produces better constrained estimates of transverse polarizabilities than can be obtained with mono-static sensor data. This is illustrated in Fig. 4, which compares polarizabilities estimated from EM-61, MetalMapper, and TEMTADS data acquired over the same 37 mm projectile. Here we quantify the discrepancy between transverse polarizabilities with the asymmetry parameter

$$\varsigma = \frac{1}{N}\sum_{i}^{N} (L_2(t_i) - L_3(t_i))/L_2(t_i). \tag{7}$$

The estimated polarizabilities at each channel are sorted so that $L_1 \geq L_2 \geq L_3$, implying $\varsigma < 1$. An ideal axisymmetric target with equal transverse polarizabilities will have $\varsigma = 0$. For the example in Fig. 4 the cued sensors produce a significantly smaller asymmetry parameter than the EM-61 (paired t-tests, 95% confidence level).

We note, however, that at late times the transverse polarizabilities estimated from TEMTADS and MetalMapper data begin to diverge due to decreased signal to noise. This effect is exacerbated when the cued sensors are not properly positioned directly over the target. In Fig. 5 we show two inversion results for MetalMapper data acquired over a 37 mm projectile. In the first data collection the target is near the edge of the sensor and the resulting transverse polarizabilities are poorly constrained. Repositioning the sensor over the target in Fig. 5(b) significantly reduces the estimated asymmetry.

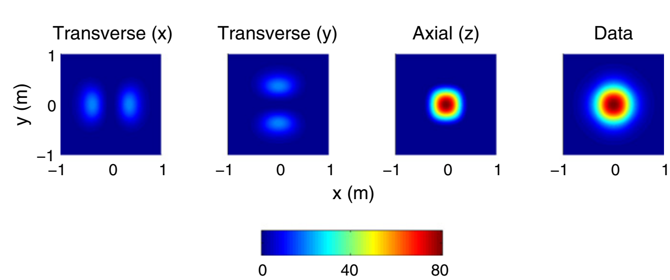

Fig. 2. Components of the dipole response over a target positioned at $\mathbf{r} = [0, 0, -0.3]$ m for the EM-61. Predicted data are a linear combination of axial and transverse responses, here for a spherical target with polarizabilities $L_i = 1$, $i = 1,2,3$. Excitation of transverse responses requires a horizontal standoff, resulting in a lower SNR than for axial excitation.

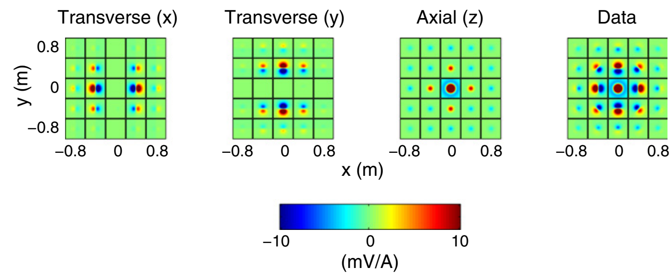

Fig. 3. Components of the dipole response over a spherical target (same as in 2) for the 5×5 TEMTADS array. Each subplot shows the received field at all receivers excited by the corresponding transmitter in the array. Data units for the TEMTADS result from normalization of the received EMF by the transmitter current.

While multi-static sensors can greatly improve the reliability of estimated polarizabilities, we conclude from the preceding examples that challenging scenarios with low SNR can still result in poorly constrained secondary polarizabilities that will confound an algorithm that incorporates these parameters to make discrimination decisions.

Here we investigate techniques for explicitly constraining transverse polarizability estimates when inverting multi-static sensor data. An obvious and viable approach to this problem is to simply reparameterize the dipole model so that secondary and tertiary polarizabilities are equal. Practical use of this axisymmetric dipole model is motivated by the analytic response of a spheroid and by successful fits to high-fidelity test stand data acquired over real axisymmetric targets (Pasion, 2007). Of course, this model will not provide a good fit to data acquired over a non-axisymmetric target with three unique polarizabilities.

A data processing approach which has been proposed to handle this ambiguity is to fit each target using both axisymmetric and non-axisymmetric (i.e. unequal transverse) dipole parameterizations and then to compare the fits to the observed data (Pasion, 2007). The non-axisymmetric parameterization has more degrees of freedom that can be used to fit observed data and so generally provides a lower misfit than the axisymmetric dipole parameterization. The problem is then to determine what constitutes a significant difference in data misfit for the two competing parameterizations. Model selection criteria can be used to select the most parsimonious model parameterization which can explain the data (Hastie et al., 2001).

In this work we instead apply regularization techniques to constrain the polarizabilities estimated from TEMTADS and MetalMapper data sets. Constraints on model parameters are typically applied in the form of parameter bounds, here we will impose an additional constraint in the form of a penalty on unequal transverse polarizabilities. This approach provides us with a continuum of possible models between constrained and unconstrained models (or, equivalently, between axisymmetric and non-axisymmetric dipole parameterizations). We investigate methods for selecting a model, or set of models, from regularized inversion. Finally, we show applications to data sets from ESTCP demonstrations at San Luis Obispo (SLO), CA, and Camp Butner, NC.

3. Regularized inversion

When solving parametric inverse problems, it is often sufficient to minimize a data norm quantifying the misfit between observed and predicted data

$$\phi_d = \mathbf{W}_d \left(\mathbf{d}^{obs} - \mathbf{d}^{pred}\right)^2 \tag{8}$$

with $\mathbf{d}^{pred} = F(\mathbf{m})$ generally a nonlinear functional of the model $\mathbf{m}$, and $\mathbf{W}_d$ a weighting matrix accounting for estimated errors on the data. Assuming Gaussian errors on the observed data, minimization of Eq. (8) yields a maximum likelihood estimate of the model parameters (Menke, 1989). Additional prior information can be incorporated

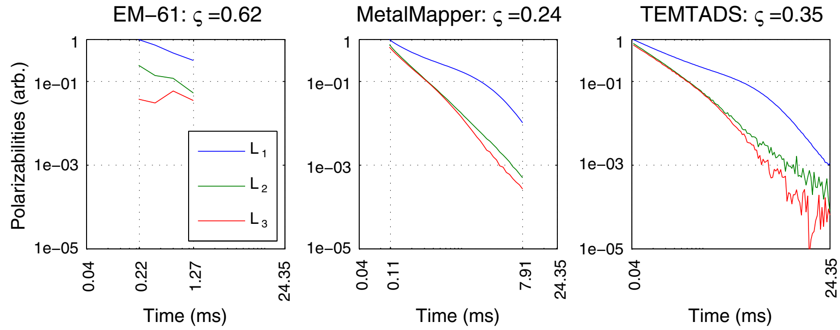

Fig. 4. Estimated dipole polarizabilities for a 37 mm projectile buried at 10 cm. The primary polarizability at channel 1 is normalized to 1 for each sensor. Abscissa limits in all plots correspond to the time range of the TEMTADS sensor. The asymmetry parameter ς is defined in Eq. (7).

in the inversion via parameter bounds (e.g. positivity) or by constructing a model which has specified properties. In the latter case, the optimization problem can be solved by minimizing the norm (Oldenburg and Li, 2005)

$$min\phi = \phi_d + \beta\phi_m. \tag{9}$$

where the regularization parameter β controls the trade-off between data and model norms. The model norm ϕ_m is a regularizer that ensures that the recovered model has, for example, a minimum deviation from some prior reference model. Aliamiri et al. (2007) employed a regularization of this form for estimation of dipole model parameters, with the regularization parameter fixed a priori. An even stronger regularization of the inverse problem is developed in Pasion et al. (2007): they fix polarizabilities at library values for each ordnance class. They then fit the observed data for target location and orientation and use the data misfit as a metric for ranking targets. This requires multiple inversions of each anomaly (one for each ordnance class) and cannot be easily generalized to multi-object scenarios.

Here we instead apply a regularizer which addresses the problem of poorly constrained transverse polarizabilities. In the case of the dipole model, an appropriate model norm which penalizes differences in secondary and tertiary polarizabilities is

$$\begin{aligned}\phi_m &= (L_2 - L_3)^2 \\ &= \|\mathbf{W}_m\mathbf{m}\|^2\end{aligned} \tag{10}$$

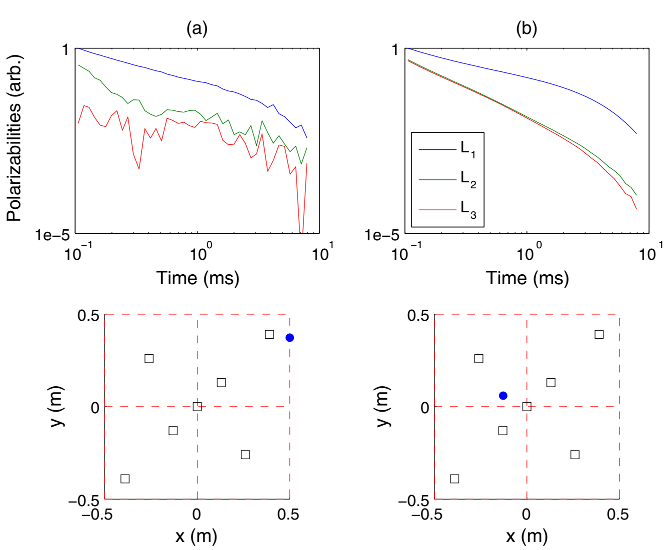

Fig. 5. Top row: estimated polarizabilities for two soundings with the MetalMapper over the same 37 mm projectile (a different 37 mm than in Fig. 4). Bottom row: estimated target locations (blue circles) for these soundings, relative to the sensor. Black squares indicate receivers and red dashed lines are transmitters.

with $\mathbf{W}_m$ a model weighting matrix acting as a differencing operator on the appropriate elements of the model vector $\mathbf{m}$. For a linear forward modeling ($\mathbf{d}^{pred} = \mathbf{Gm}$) with fixed β, the model estimate is obtained by solving the system (Oldenburg and Li, 2005)

$$\left(\mathbf{G}^T\mathbf{W}_d^T\mathbf{W}_d\mathbf{G} + \beta\mathbf{W}_m^T\mathbf{W}_m\right)\mathbf{m} = \mathbf{G}^T\mathbf{W}_d^T\mathbf{W}_d\mathbf{d}^{obs}. \tag{11}$$

In practice, the inverse problem can be solved by minimizing Eq. (9) over a range of β values, beginning with a large value of β and progressively decreasing (or "cooling") this parameter (Oldenburg and Li, 2005). When regularizing overdetermined problems the model has limited degrees of freedom with which to fit the data and so the β cooling procedure will stall at a model corresponding to the unconstrained ($\beta = 0$) solution. This is in contrast to the underdetermined inverse problem, where continued decrease of the regularization parameter eventually introduces spurious model structure.

Following on the works of Shubitidze et al. (2007) and Song et al. (2011), we use a sequential inversion approach that decouples estimation of target location and dipole polarizabilities. The predicted data can be expressed as

$$\mathbf{d}^{pred} = \mathbf{G}(\mathbf{r})\mathbf{m}_\mathbf{P}. \tag{12}$$

Here the model $\mathbf{m}_\mathbf{P}$ is composed of the six unique elements of the symmetric polarizability tensor $\mathbf{P}$ at a single time channel

$$\mathbf{m}_\mathbf{P} = \left[P_{xx},\ P_{xy},\ P_{xz},\ P_{yy},\ P_{yz},\ P_{zz}\right]^T. \tag{13}$$

The forward modeling matrix is

$$\mathbf{G}(\mathbf{r}) = \begin{bmatrix} B_s^x B_p^x \\ B_s^x B_p^y + B_s^y B_p^x \\ B_s^x B_p^z + B_s^z B_p^x \\ B_s^y B_p^y \\ B_s^y B_p^z + B_s^z B_p^y \\ B_s^z B_p^z \end{bmatrix}^T \tag{14}$$

with B_p the primary field at the target and B_s the secondary field at the receiver, with all fields implicitly dependent upon target ($\mathbf{r}$) and sensor location. Superscripts denote the x,y,z components of the respective fields. We then solve the regularized inverse problem as follows:

1 Solve an inverse problem for target location $\mathbf{r}$. The model is related to the predicted data via the nonlinear functional

$$\mathbf{d}^{pred} = F[\mathbf{r}] = \mathbf{G}(\mathbf{r})\mathbf{G}^\dagger(\mathbf{r})\mathbf{d}^{obs}. \tag{15}$$

with $\mathbf{G}^\dagger$ denoting the pseudo-inverse. We estimate $\mathbf{r}$ by minimization of Eq. (8) using an iterative Gauss–Newton algorithm.

2 Solve the linear inverse problem for the unique elements of the polarizability tensor ($\mathbf{m_P}$) at all time channels using the location obtained at the previous step.
3 Compute the eigenvalues of the polarizability tensor at each time channel using joint diagonalization (Cardoso, 1996). This algorithm returns a single eigenvector matrix for all channels, corresponding to a fixed target orientation. The eigenvalues at each time channel are then an initial estimate of principal polarizabilities. The eigenvectors correspond to the columns of the Euler rotation matrix $\mathbf{A}$ in Eq. (4). To estimate the orientation angles we then minimize the least squares difference (Frobenius norm) between the eigenvector matrix and a Euler rotation matrix parameterized by orientation angles (ϕ,θ,ψ).
4 We now address the problem stated at the beginning of this section where we solve for polarizabilities using a constrained approach. At each time channel we solve a linear regularized inverse problem for principal polarizabilities, at a fixed orientation. The model at each time channel is

$$\mathbf{m}_L\left(t_j\right) = \left[L_1\left(t_j\right),\ L_2\left(t_j\right),\ L_3\left(t_j\right)\right]^T. \tag{16}$$

At each time channel we obtain a set of models corresponding to the solution of Eq. (11) over a range of βs.

We minimize Eq. (9) starting from an initial, large value of β. The regularization parameter is initialized so that the term $\beta\phi_m \gg \phi_d$ at the initial model obtained in step 3. We then lower the regularization parameter by a cooling factor κ (e.g. $\kappa=0.5$) and solve Eq. (11) at the new β value. This procedure is repeated until the relative change in the model parameters achieves some tolerance ϵ.

We remark that the sequential inversion approach is very fast relative to an "all-at once" algorithm that tries to estimate polarizability parameters at all time channels simultaneously. Because the time channels are inverted separately in the sequential method, the relative weighting of time channels via estimated errors in the data weighting matrix ($\mathbf{W}_d$) is less critical. In this implementation we assume a constant standard deviation for all data at each time channel, so that $\mathbf{W}_d$ is just a linear scaling of the data misfit. Since each time channel is inverted separately at the estimated source location, parallel processing can greatly reduce computation time.

4. Multi-object regularized inversion

Any practical inversion algorithm for UXO discrimination must consider overlapping responses from multiple targets. Our focus in this paper is on estimation of regularized models and their subsequent application in identifying targets of interest, and so here we only briefly describe our regularized multi-object inversion algorithm.

Song et al. (2011) augments the model vector with multiple dipole sources and simultaneously estimates locations for all sources, followed by estimation of all polarizabilities. Applying regularized inversion at this second step would require separate model norm terms for each object, and careful balancing of these terms via separate regularization parameters. We pursue a more straightforward route and decouple the regularized inversions into a series of single object inversions, as follows.

We first estimate locations for two dipole sources simultaneously as in Song et al. (2011). The predicted data are then a superposition of the data predicted by dipoles at locations $\mathbf{r}_i$

$$\mathbf{d}^{pred} = \mathbf{d}^{pred}(\mathbf{r}_1) + \mathbf{d}^{pred}(\mathbf{r}_2). \tag{17}$$

Fig. 6 compares (unregularized) single and two-object inversion results for a TEMTADS data set acquired over 37 mm projectile and clutter items at Camp Butner. Note that the 115 raw TEMTADS channels have been averaged in windows of length 5 to speed processing.

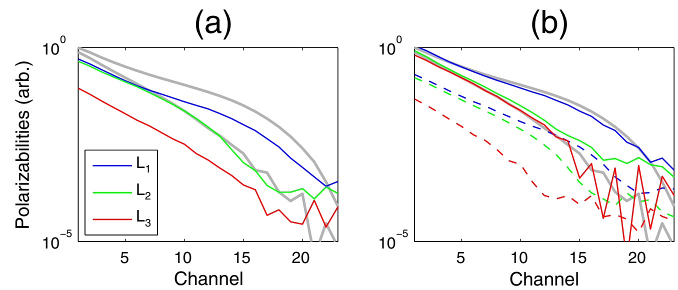

Fig. 6. (a) Estimated polarizabilities obtained from unregularized single object inversion of TEMTADS data acquired over a 37 mm projectile at Camp Butner. Gray lines are library polarizabilities for 37 mm item. (b) Estimated polarizabilities obtained from unregularized two-object inversion over the same data.

In this example the two-object inversion recovers a set of polarizabilities that are in much better agreement with 37 mm library values than the single object result.

Despite this improvement, we note that the estimated secondary polarizabilities for this target are poorly constrained at late times; subsequent regularized inversion at the corresponding source locations can address this ill-conditioning. We carry out separate regularized inversions at these locations, with each regularized inversion fitting the residual data that cannot be predicted by the other dipole. That is, the observed data sets for estimation of target polarizabilities for the two objects at fixed locations $\mathbf{r}_i$ are

$$\begin{aligned}\mathbf{d}^{obs}(\mathbf{r}_1) &= \mathbf{d}^{obs} - \mathbf{d}^{pred}(\mathbf{r}_2)\\ \mathbf{d}_2^{obs}(\mathbf{r}_2) &= \mathbf{d}^{obs} - \mathbf{d}^{pred}(\mathbf{r}_1).\end{aligned} \tag{18}$$

This is similar to the "iterative residual fit" (IRF) technique investigated in Bell (2006). This method alternates between single object inversions, decoupled as in Eq. (18). They estimate a target's model parameters (location and polarizabilities) with the other target's parameters held constant. The alternating inversions terminate when both model vectors stop changing over successive iterations. Bell (2006) found that the IRF method was less reliable than a "double happiness" approach that estimated all target parameters simultaneously. Our implementation is a hybrid of the sequential multi-object inversion in Song et al. (2011) and the IRF: we estimate target locations simultaneously, then decouple estimation of polarizabilities using Eq. (18). Our discrimination results, presented in Section 7, indicate that this is a viable approach.

5. Model selection

A final model that balances the trade-off between model and data norms can be selected based upon achieving some target data misfit (discrepancy principle). Alternatively, the L-curve criterion identifies a model corresponding to the point of maximum positive curvature on a logarithmic plot of data misfit versus model norm (Hansen, 1997). This method is less sensitive to estimates of measurement errors and Farquharson and Oldenburg (2004) found it to be quite robust when inverting real electromagnetic data sets.

Fig. 7 shows a regularized inversion result for channel 1 of MetalMapper data acquired over the 37 mm projectile presented in Fig. 5(a) (recall that the sensor is poorly positioned over the target for this sounding). As the regularization parameter β is decreased, the objective function ϕ in Eq. (9) is dominated by the data misfit term and the secondary and tertiary polarizabilities diverge to their unregularized values, with the primary polarizability unaffected by the regularization. Applying the standard L-curve criterion in Fig. 5(a) yields a model with unequal transverse polarizabilities, similar to the unregularized model. In order to bias our model selection toward more axisymmetric objects,

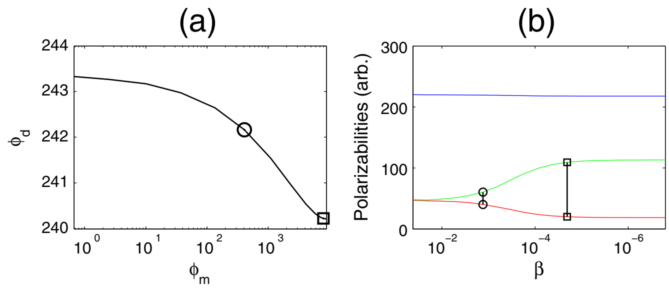

Fig. 7. Regularized inversion result for channel one of MetalMapper data acquired over a 37 mm projectile, same data as in Fig. 5(a). (a) Tikhonov curve. Markers indicate points of maximum negative (circle) and positive (square) curvature, respectively. (b) Estimated polarizabilities as a function of regularization parameter β. Markers indicate transverse polarizabilities for points identified on the L-curve.

we instead select the point of maximum *negative* curvature on the L-curve, as shown in Fig. 7.

Applying this modified criterion across all time channels yields a significant reduction in estimated asymmetry for this target (Fig. 8), relative to an unregularized inversion. Also shown in Fig. 8 is a comparison of unregularized and regularized models for a non-axisymmetric clutter item. Regularization reduces the estimated asymmetry, but at early time channels the secondary and tertiary polarizabilities are not equal, as expected for this target. At late times, however, the SNR is reduced and the data cannot constrain target shape. In this case the regularized inversion defaults to axisymmetric models. We also remark that because regularization is applied to each channel separately, there is still some jitter in the estimated polarizabilities, especially at late times. This can be addressed by adding an additional term to the model norm to enforce smoothness between channels. However, a computational advantage of our inversion algorithm is parallel processing of time channels and a smoothness term will preclude this strategy.

6. Using regularized models in discrimination

As an alternative to conventional model selection with the L-curve criterion, we also consider discrimination using *all* models obtained in regularized inversion. To account for the relative quality of the fit to the data, for the jth model at channel t_i (here denoted $\mathbf{m}_j(t_i)$), we compute the likelihood

$$p\left(\mathbf{m}_j(t_i)\right) = exp\left(-\frac{\phi_d\left(\mathbf{m}_j(t_i)\right) - \phi_d^{min}(t_i)}{\phi_d^{min}(t_i)}\right) \tag{19}$$

with $\phi_d^{min}(t_i)$ the minimum data misfit obtained over all models at channel t_i, i.e.

$$\phi_d^{min}(t_i) = \min_k \phi_d(\mathbf{m}_k(t_i)). \tag{20}$$

The model likelihood is a monotonic transformation of the data misfit ϕ_d to a probability.

We remark that since the argument in Eq. (19) is small, the model likelihood is approximately

$$\begin{aligned} p\left(\mathbf{m}_j(t_i)\right) &\approx 1 - \frac{\phi_d\left(\mathbf{m}_j(t_i)\right) - \phi_d^{min}(t_i)}{\phi_d^{min}(t_i)} \\ &= 1 - \text{relative misfit} \end{aligned} \tag{21}$$

with relative misfit used to compare models in Lhomme et al. (2008). This form of model likelihood is somewhat arbitrary; rigorous calculation of model likelihoods (or posterior probabilities) requires a full uncertainty appraisal. However, we will show in Section 7 that incorporating the relative quality of the data fit via the model likelihood is beneficial to discrimination.

We now consider discrimination with a nonlinear support vector machine (SVM) using our regularized inversion models and corresponding likelihoods (see Hastie et al., 2001 or Burges, 1998 for a full description of the SVM). The decision function for the SVM is

$$y^{test} = w^T K\left(x^{train}, x^{test}\right) + b_0 \tag{22}$$

with K a kernel matrix, w a (sparse) weight vector, and b_0 a constant bias term. For a radial basis function

$$K_{ij} = exp\left(-\nu\left\|\mathbf{x}_i^{train} - \mathbf{x}_j^{test}\right\|^2\right) \tag{23}$$

with $\mathbf{x}_i^{train}$ and $\mathbf{x}_j^{test}$ the ith training and jth test vectors, respectively. The parameter ν controls the width of the kernels. For discrimination with polarizabilities, the elements of these feature vectors are log-transformed primary, secondary, and tertiary polarizabilities at a subset of the available time channels, i.e.

$$\mathbf{x}_j^{test} = F[\mathbf{m}]. \tag{24}$$

Expanding the norm in Eq. (23) in terms of the elements of training and test vectors, the kernel matrix can be expressed as

$$K_{ij} = \prod_{k=1}^{N} exp\left(-\nu\left(x_{ik}^{train} - x_{jk}^{test}\right)^2\right). \tag{25}$$

When considering multiple models from a regularized inversion, the kth element of the jth test vector has multiple values, denoted x_{jkl}^{test}, with corresponding likelihoods $p(x_{jkl}^{test})$ from Eq. (19). For discrimination, we select the element for which the term

$$\kappa_{ijkl} = \left[exp\left(-\nu\left(x_{ik}^{train} - x_{jkl}^{test}\right)^2\right)p\left(x_{jkl}^{test}\right)\right] \tag{26}$$

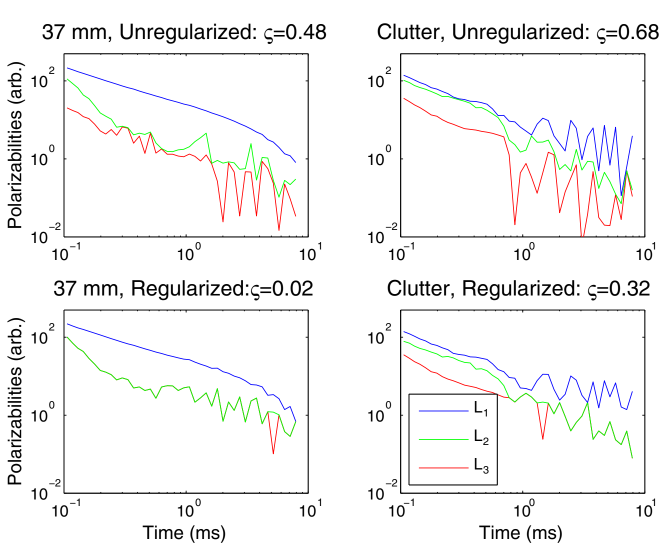

Fig. 8. Comparison of regularized and unregularized polarizability models for an axisymmetric 37 mm projectile and a non-axisymmetric clutter item. Polarizabilities are estimated from MetalMapper data sets acquired at Camp Butner. Regularized models are selected using L-curve criterion.

is maximized. The elements of the kernel matrix are then

$$K_{ij} = \prod_{k=1}^{N} \max_{l} \kappa_{ijkl} = \prod_{k=1}^{N} \max_{l} \left[exp\left(-\nu\left(x_{ik}^{train} - x_{jkl}^{test}\right)^2\right) p\left(x_{jkl}^{test}\right)\right]. \quad (27)$$

The preceding computations compare all possible values of a test vector with a given training vector and retain the test vector elements which are "closest" (in the sense of the radial basis function) to that training vector. The weighting by model likelihoods acts to penalize vectors (models) which may agree with a training vector but which do not fit the observed data.

7. Results

We processed cued MetalMapper and TEMTADS data sets acquired during ESTCP demonstrations at San Luis Obispo (SLO) and Camp Butner. All targets in each data set were inverted with unregularized and regularized algorithms and with single and two-object models. We then classified each target on the basis of the model from single or two-object inversions which are predicted as most likely to be a UXO by a given discrimination algorithm.

At SLO, the discrimination task was to identify seeded targets ranging in size from 4.2 in. mortars down to 60 mm mortars (Fig. 9a). For this site training data were provided by ESTCP and comprised a random sample of 174 detected targets. Targets of interest at Camp Butner ranged from large 105 mm projectiles down to 37 mm projectiles (Fig. 9b). Beyond test pit measurements of individual items from each class of TOI, no training data were initially available for classifier training. For each data set at Camp Butner, we requested ground truth for a small number (<50) of targets in order to characterize the distributions of TOI and non-TOI polarizabilities. A description of the procedure used to select training data is given in Pasion et al. (2011). Note that for (labeled) training feature vectors we use only the minimum misfit models for classifier training and prediction with Eq. (27).

Fig. 9. Targets of interest for ESTCP demonstrations at SLO and Camp Butner.

In Fig. 10 we compare the performance of a number of different algorithms applied to the SLO MetalMapper data:

- Primary polarizabilities (pols): SVM using estimated primary polarizabilities from unregularized inversions.
- All pols (unregularized): SVM using all polarizabilities (primary, secondary and tertiary) from unregularized inversions.
- All pols (L-curve): SVM using all polarizabilities from regularized inversions. A single model is selected using the modified L-curve criterion described in Section 5.
- All pols (unweighted): SVM using all polarizabilities from regularized inversions. All models are inputted into discrimination, regardless of fit quality.
- All pols (weighted): SVM using all polarizabilities from regularized inversions. All models are inputted into discrimination, with model likelihoods used to weight the SVM decision function (Eq. (27)).

Here discrimination with only primary polarizabilities is the baseline algorithm for processing MetalMapper data: no shape information is included in this algorithm. Also shown for comparison in Fig. 10 is the ROC derived from EM-61 data acquired over the same targets. For this sensor our discrimination algorithm is a threshold on the decay parameter (Eq. (5)). This represents the optimal classification performance attained for the EM-61 data at this site (Nelson et al., 2010).

Discrimination with the MetalMapper primary polarizabilities provides a dramatic improvement in discrimination performance relative to the EM-61. Late-time information (>5 ms) supplied by the MetalMapper (and TEMTADS) is particularly useful for discriminating between small ordnance (e.g. 60 mm) and clutter of similar size.

The shape information encoded in the transverse polarizabilities can provide an improvement in discrimination performance over the baseline approach. However, these parameters are sensitive to noise and without regularization the last target of interest occurs relatively late in the dig list (Fig. 10). Selecting a single regularized model using the L-curve criterion reduces the false alarm rate, but this method also has difficulty finding the last few targets of interest in the dig list. For low

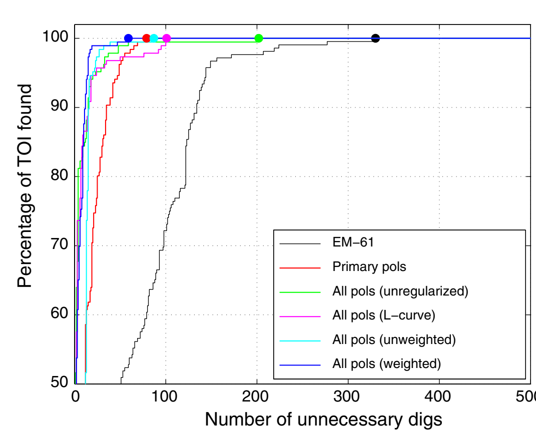

Fig. 10. Receiver operating characteristics (ROCs) for discrimination algorithms applied to SLO MetalMapper data, see text for an explanation of each algorithm. ROC derived from EM-61 data acquired over the same targets is also shown. Markers indicate the point on each ROC at which all targets of interest are found.

SNR targets the Tikhonov curve has a small curvature and it becomes difficult to accurately identify the L-curve criterion. The weighted SVM achieves the best discrimination performance for algorithms applied to the SLO MetalMapper data. The benefit of weighting by model likelihoods is illustrated in Fig. 10 by the reduction in false alarm rate relative to using unweighted models.

Fig. 11 summarizes the performance of nonlinear support vector machines using regularized and unregularized test polarizabilities on both SLO and Camp Butner data sets. For clarity we show a subset of algorithms presented in the previous figure, the results in Fig. 10 are representative of all test cases. For the data sets considered in Fig. 11, regularization consistently decreases the false alarm rate relative to using all (primary, secondary and tertiary) unregularized polarizabilities. For SLO and Camp Butner MetalMapper data sets the regularized method achieves the maximal area under the ROC (AUC) and a false alarm rate (FAR) comparable to discrimination using only primary polarizabilities. For the TEMTADS data sets, discrimination with regularized inversion achieves a FAR comparable to discrimination with primary polarizabilities. The performance for the Camp Butner TEMTADS data is relatively insensitive to the discrimination method, with the regularized algorithm producing a slightly lower FAR than the other approach. These data were acquired with careful in-field quality control so that sensor positioning and data SNR consistently support reliable estimation of dipole model parameters and no additional benefit is derived from model regularization.

8. Conclusions

We have applied a regularized inversion algorithm that penalizes the deviation between transverse polarizabilities in the TEM dipole model. Rather than selecting a single model from this inversion process, we input all models into a support vector machine classifier. This corresponds to test feature vectors with multiple values for each element. We compare the elements of each test vector with the training data and retain the model value which best corresponds to a given training vector. We also penalize the match of test and training vectors by the likelihood that the model fits the observed data. In applications to real data sets this technique outperforms discrimination using models selected from regularized inversion using an L-curve criterion.

We find that the regularized method with model weighting can improve initial performance on high SNR targets with well-constrained transverse polarizabilities. The method also prevents the occurrence of outlying TOI that arise when we use unregularized parameters to rank all targets. This produces a receiver operating characteristic with a large area under the curve and a final false alarm rate that is comparable to that obtained when we rely on primary polarizabilities throughout the dig list. While we have focused on support vector machines for classification, the model weighting approach developed here can be readily extended to any kernel based classifier, e.g. the relevance vector machine (Tipping, 2001).

A viable alternative to the approach presented here is to begin ranking well-constrained models using all polarizabilities, and then to switch to primary polarizabilities for ranking of lower SNR targets. This can produce similar results to the algorithms developed here, but requires judgment from an analyst to decide when to switch between feature sets. In contrast, discrimination with weighted models from regularized inversion introduces no additional tuning parameters into the discrimination process.

Acknowledgments

We would like to thank Dr. Lin-Ping Song of UBC and Dr. Len Pasion of Sky Research, Inc. for their forward modeling codes. We also thank all reviewers for the helpful feedback. This work was funded by the Strategic Environmental Research and Development Program (SERDP) project MR-1629. Data preprocessing was carried out under ESTCP projects MR-0504 and MR-1004.

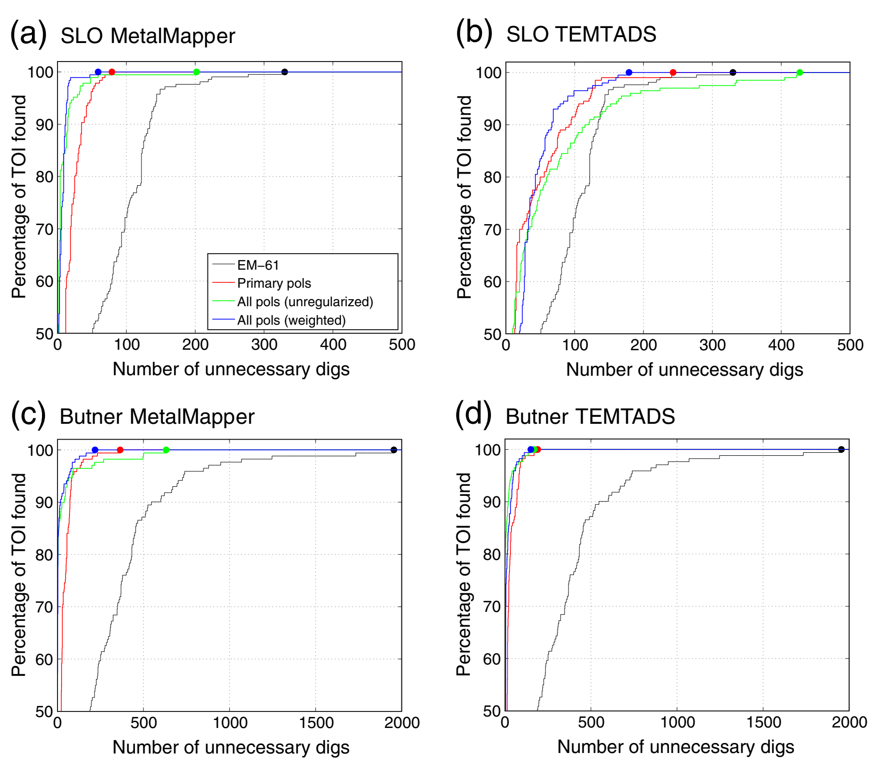

Fig. 11. Receiver operating characteristics (ROCs) for support vector machines applied to cued data sets at San Luis Obispo (SLO) and Camp Butner. ROCs derived from EM-61 data acquired over the same targets are also shown. Markers indicate the point on each ROC at which all targets of interest are found.

References

Aliamiri, A., Stalnaker, J., Miller, E.L., 2007. Statistical classification of buried unexploded ordnance using nonparametric prior models. IEEE Transactions on Geoscience and Remote Sensing 45, 2794–2806.

Bell, T., 2006. Adaptive and Iterative Processing Techniques for Overlapping Signatures: Technical Summary Report, ESTCP Project MM-200415. Technical Report. ESTCP.

Bell, T., Barrow, B., Miller, J., Keiswetter, D., 2001a. Time and frequency domain electromagnetic induction signatures of unexploded ordinance. Subsurface Sensing Technologies and Applications 2, 153–175.

Bell, T., Barrow, B., Miller, J.T., 2001b. Subsurface discrimination using electromagnetic induction sensors. IEEE Transactions on Geoscience and Remote Sensing 39, 1286–1293.

Beran, L.S., Billings, S.D., Oldenburg, D., 2011. Incorporating uncertainty in unexploded ordnance discrimination. IEEE Transactions on Geoscience and Remote Sensing 49, 3071–3080.

Billings, S.D., Pasion, L.R., Beran, L., Lhomme, N., Song, L., Oldenburg, D.W., Kingdon, K., Sinex, D., Jacobson, J., 2010. Unexploded ordnance discrimination using magnetic and electromagnetic sensors: case study from a former military site. Geophysics 75, B103–B114.

Burges, C.J.C., 1998. A tutorial on support vector machines for pattern recognition. Data Mining and Knowledge Discovery 2, 955–974.

Cardoso, J.F., 1996. Jacobi angles for simultaneous diagonalization. SIAM Journal on Matrix Analysis and Applications 17, 161–163.

Delaney, W.P., Etter, D., 2003. Report of the Defense Science Board on unexploded ordnance. Technical Report. Office of the Undersecretary of Defense for Acquisition, Technology and Logistics.

Farquharson, C.G., Oldenburg, D.W., 2004. A comparison of automatic techniques for estimating the regularization parameter in non-linear inverse problems. Geophysical Journal International 156, 411–425.

Hansen, P.C., 1997. Rank-deficient and discrete ill-posed problems: numerical aspects of linear inversion. SIAM.

Hastie, T., Tibshirani, R., Friedman, J., 2001. The Elements of Statistical Learning: Data Mining, Inference and Prediction. Spring-Verlag.

Lhomme, N., Oldenburg, D.W., Pasion, L.R., Sinex, D., Billings, S.D., 2008. Assessing the quality of electromagnetic data for the discrimination of UXO using figures of merit. Journal of Engineering and Environmental Geophysics 13, 165–176.

Menke, W., 1989. Geophysical Data Analysis: Discrete Inverse Theory. Academic Press.

Nelson, H.H., Andrews, A., Kaye, K., 2010. ESTCP pilot program: classification approaches in munitions response, San Luis Obispo, California. Technical Report. ESTCP.

Oldenburg, D., Li, Y., 2005. Investigations in Geophysics No. 13: near-surface. Society of Exploration Geophysics. Chapter Inversion for Applied Geophysics: A Tutorial, pp. 89–150.

Pasion, L.R., 2007. Inversion of time-domain electromagnetic data for the detection of unexploded ordnance. Ph.D. thesis. University of British Columbia.

Pasion, L.R., Oldenburg, D.W., 2001. A discrimination algorithm for UXO using time domain electromagnetic induction. Journal of Environmental and Engineering Geophysics 6, 91–102.

Pasion, L.R., Billings, S.D., Oldenburg, D.W., Walker, S., 2007. Application of a library-based method to time domain electromagnetic data for the identification of unexploded ordnance. Journal of Applied Geophysics 61, 279–291.

Pasion, L.R., Beran, L.S., Zelt, B.C., Kingdon, K.A., 2011. Feature extraction and classification of EMI data, Camp Butner, NC, ESTCP project MR-201004. Technical Report. ESTCP.

Prouty, M., George, D.C., Snyder, D.D., 2011. MetalMapper: a multi-sensor TEM system for UXO detection and classification, ESTCP project MR-200603. Technical Report. ESTCP.

Shubitidze, F., O'Neill, K., Haider, S.A., Sun, K., Paulsen, K.D., 2002. Application of the method of auxiliary sources to the wide-band electromagnetic induction problem. IEEE Transactions on Geoscience and Remote Sensing 40, 928–942.

Shubitidze, F., O'Neill, K., Barrowes, B., Shamatava, I., Fernández, J., Sun, K., Paulsen, K., 2007. Application of the normalized surface magnetic charge model to UXO discrimination in cases with overlapping signals. Journal of Applied Geophysics 61, 292–303.

Shubitidze, F., Barrowes, B., Shamatava, I., Fernandez, J.P., O'Neill, K., 2011. The ortho normalized volume magnetic source technique applied to live-site UXO data: inversion and classification studies. SEG Technical Program Expanded Abstracts 30, 3766–3770.

Song, L.P., Pasion, L.R., Billings, S.D., Oldenburg, D., 2011. Nonlinear inversion for multiple objects in transient electromagnetic induction sensing of unexploded ordnance: technique and applications. IEEE Transactions on Geoscience and Remote Sensing 49, 4007–4020.

Steinhurst, D.A., Harbaugh, G.R., Kingdon, J.B., Furuya, T., Keiswetter, D.A., George, D.C., 2010. EMI array for cued UXO discrimination, ESTCP project MM-0601. Technical Report. ESTCP.

Tantum, S.L., Li, Y., Collins, L.M., 2008. Bayesian mitigation of sensor position errors to improve unexploded ordnance detection. IEEE Geoscience and Remote Sensing Letters 5, 103–107.

Tarokh, A.B., Miller, E.L., 2007. Subsurface sensing under sensor positional uncertainty. IEEE Transactions on Geoscience and Remote Sensing 45, 675–688.

Tipping, M.E., 2001. Sparse Bayesian learning and the relevance vector machine. Journal of Machine Learning Research 1, 211–244.

Wait, J.R., 1982. Geo-electromagnetism. Academic Press.

West, G.F., Macnae, J.C., 1991. Electromagnetic methods in applied geophysics. SEG. Chapter physics of the electromagnetic exploration method, pp. 5–45.

Zhang, Y., Collins, L.M., Yu, H., Baum, C.E., Carin, L., 2003. Sensing of unexploded ordnance with magnetometer and induction data: theory and signal processing. IEEE Transactions on Geoscience and Remote Sensing 41, 1005–1015.