

PROCEEDINGS OF SPIE

SPIDigitalLibrary.org/conference-proceedings-of-spie

Comparison of hand-held WEMI target detection algorithms

Connor H. McCurley, James Bocinsky, Alina Zare

Connor H. McCurley, James Bocinsky, Alina Zare, "Comparison of hand-held WEMI target detection algorithms," Proc. SPIE 11012, Detection and Sensing of Mines, Explosive Objects, and Obscured Targets XXIV, 110120U (10 May 2019); doi: 10.1117/12.2519454

SPIE.

Event: SPIE Defense + Commercial Sensing, 2019, Baltimore, Maryland, United States

Comparison of Hand-held WEMI Target Detection Algorithms

Connor H. McCurley, James Bocinsky, and Alina Zare

Electrical and Computer Engineering, University of Florida, Gainesville, FL 32611, USA

ABSTRACT

Wide-band Electromagnetic Induction Sensors (WEMI) have been used for a number of years in subsurface detection of explosive hazards. While WEMI sensors have proven effective at localizing objects exhibiting large magnetic responses, detecting objects lacking or containing very low amounts of conductive materials can be challenging. In this paper, we compare a number of target detection algorithms in the literature in terms of detection performance. In the comparison, methods are tested on two real-world data sets: one containing relatively low amounts of ground noise pollution, and the other demonstrating highly-magnetic soil interference. Results are quantitatively evaluated through receiver-operator characteristic (ROC) curves and are used to highlight the strengths and weaknesses of the compared approaches in hand-held explosive hazard detection.

Keywords: Hazard detection, WEMI, ACE, SMF, JOMP, Multiple-instance, Mean-shift

1. INTRODUCTION

Subsurface explosive hazards pose a significant threat to soldiers and civilians alike. While such threats are typically buried during wartime, they can remain active for years afterward. Timely, and effective, discovery and disposal of explosive ordinance is imperative to ensuring safe traveling conditions for troops and limiting the number of civilian casualties. Consequently, extensive efforts have been spent in developing *explosive hazard detection* (EHD) systems which are capable of uncovering this type of threat.

Subsurface object detection systems typically employ one or more sensors to extract information about an underground scene. This paper focuses on detection methods utilizing a hand-held sensor platform outfitted with *wide-band electromagnetic induction* (WEMI) sensors, also called metal detectors, and a *Universal Transverse Mercator* (UTM) positioning system which provides accurate spatial information. WEMI sensors have been studied extensively for use in detecting buried, metallic objects.¹⁻⁵ These sensors operate by sending a time-varying electromagnetic field through the ground via a transmit coil. Conductive subsurface objects are energized by this field and quickly return to resting states. The emitted relaxation energy is captured as a shift in magnitude and phase across one or more of the receive coils. An array of operating frequencies is used to observe inherent responses from a wide variety of conductive objects. While traditional WEMI approaches

Further author information: (Send correspondence to C.H.M)

C.H.M: E-mail: cmccurley@ufl.edu

J.B.: E-mail: jbocinsky@ufl.edu

A.Z.: E-mail: azare@ece.ufl.edu

rely on observing high-magnitude responses to flag potential threats,⁶ Scott, et al. proposed that the relaxation energy emitted by conductive objects can be modeled as a function of operating frequency and used to create a dictionary of expected target responses.⁷⁻⁹ Elaboration on the *discrete spectrum of relaxation frequencies* (DSRF) target dictionary and its exploitation for use in WEMI hazard detection is provided in Section 3.1.

While there have been many successes in algorithm development for detection of high-metal targets, EHD is certainly not a solved problem. The first obstacle making EHD difficult is that there exists a wide assortment of targets to be detected, ranging from small anti-personal to large anti-tank. Additionally, metal content can range from high to low, or even be absent in the case of plastic targets. While it has long been assumed that the inherent properties exhibited by metal detectors make discovery of lowly and non-conductive objects difficult to achieve, recent work has shown that the utilization of a novel WEMI processing procedure may provide discriminative information for unmasking difficult targets. Further description of this processing as well as an initial investigation to the use of its generated features for hazard discovery is presented in Sections 3.6 and 4.4.3, respectively. In addition to the range of targets, non-target objects (clutter) and natural variations in the soil's electromagnetic properties can bring rise to *false alarms* (FA) where non-target objects are labeled as potential threats. A principal challenge in EHD is mitigating false alarms while not masking true detections, or *true positives* (TP).

In this paper, the authors compare an assortment of WEMI prescreeners which have shown considerable success as subsurface object detectors on alternative WEMI hazard discovery systems. Experiments were conducted to 1.) gauge the performance of target detectors across a variety of WEMI sensors, 2.) compare prescreeners utilizing pre-defined versus learned target concepts, and 3.) evaluate two pre-processing methods for interference removal. Quantitative evaluation of performance is given by receiver-operating characteristic (ROC) curves which provide a measure for the probability of detection (PD) against false alarm rates (FAR).

The remainder of this paper is organized as follows. Section 2 describes the hand-held WEMI data used in this paper and Section 3 elaborates on the methods explored. Section 3.1 explains the DSRF target dictionary utilized by the algorithms outlined in 3.2, 3.3 and 3.4. Section 3.5 illustrates the algorithms which learn target representatives through training. Section 3.6 describes the novel WEMI interference removal and feature extraction proposed by Hayes, et al. Experiment outlines and results are presented in Section 4 and Section 5 concludes the paper.

2. DESCRIPTION OF DATA

2.1 Measurement Matrix

The WEMI sensor functions by moving over regions of interest and taking measurements at discrete positions along its path. Two hand-held WEMI sensors were utilized in this study. The sensors, hereby denoted as Sensor A and Sensor B, employ different numbers of operating frequencies to capture information at each position measurement. It follows that the WEMI data collection can be described as a function of frequency and position. Following the notation outlined by Hayes, et al., the WEMI measurement data can be written as $\hat{M}(\omega, p)$, where ω describes the operating frequency at which a measurement was taken, and p is the corresponding spatial location at the time of measurement.^{10, 11}

A measurement matrix can be created by combining a series of frequency samples through position such that

$$\hat{\mathbf{M}}(\boldsymbol{\omega}, \mathbf{p}) = [\hat{M}(\omega_i, p_j)]_{i,j} \quad (1)$$

where ω_i is the i^{th} operating frequency and p_j is the j^{th} position. This combination produces a measurement matrix $\hat{\mathbf{M}} \in \mathfrak{C}^{K \times N}$, where K frequencies were sampled at N positions.

Each frequency sample collected by the sensor is stored as a complex vector representing the phasor difference between the emitted and received signals at each of the operating frequencies. Since the complex plane \mathfrak{C} is isomorphic to \mathfrak{R}^2 , the real and imaginary components for each sample can be concatenated to form real-valued feature vectors

$$\mathbf{M}(\boldsymbol{\omega}, \mathbf{p}) = \begin{bmatrix} \Re(\hat{\mathbf{M}}(\boldsymbol{\omega}, \mathbf{p})) \\ \Im(\hat{\mathbf{M}}(\boldsymbol{\omega}, \mathbf{p})) \end{bmatrix} \quad (2)$$

such that $\mathbf{M} \in \mathfrak{R}^{M \times N}$, with $M = 2K$, becomes the data matrix used throughout this work.

2.2 Measurement Model

With the measurement matrix defined, the underlying data model can be described. As previously delineated, the measurement matrix $\mathbf{M}(\boldsymbol{\omega}, \mathbf{p})$ can be decomposed as a superposition of four major components¹¹

$$\mathbf{M}(\boldsymbol{\omega}, \mathbf{p}) = \mathbf{S}(\boldsymbol{\omega}, \mathbf{p}) + \mathbf{G}(\boldsymbol{\omega}, \mathbf{p}) + \mathbf{R}(\boldsymbol{\omega}, \mathbf{p}) + \mathcal{E} \quad (3)$$

where $\mathbf{S}(\boldsymbol{\omega}, \mathbf{p})$ is the desired target response, $\mathbf{G}(\boldsymbol{\omega}, \mathbf{p})$ is a signal resulting from soil interference, $\mathbf{R}(\boldsymbol{\omega}, \mathbf{p})$ is the WEMI sensor's self-response, and \mathcal{E} is zero-mean additive random noise following an *i.i.d.* normal random variable $\mathcal{N}(0, \sigma^2)$.

The majority of WEMI preprocessing procedures in the literature follow a two-stage system where the self-response component is addressed through some method of downtrack filtering before a separate processing step is utilized on the frequency information. For now, we only address removal of the sensor self-response component as a method of data pre-processing. Mitigation of soil-interference is addressed in Section 3.6.

2.3 Down-track Filtering

Self-response interference results from induced coupling between the transmit and receive coils of the WEMI sensor. While the interference is a strong function of frequency, it can essentially be viewed as constant with respect to position. This allows the self-response to be considered as a DC component in the downtrack dimension of the data. The standard approach to self-response removal is to apply a high-pass downtrack filter. A variety of filtering approaches have been investigated in the literature, including: Weiner, matched, and adaptive filtering, as well as convolution with zero-mean sine waves and application of Discrete Fourier Transform (DFT) filter banks.¹² In this work, however, the method described by Scott and Hayes is adopted which uses an orthonormal *discrete cosine transform* (DCT) basis matrix to remove both the DC and high-frequency spatial noise from the measurement matrix.^{10,12} The DCT method operates on real cosines and extends the signal with mirroring before filtering. These qualities are desirable as sinusoidal shapes are ideal for

downtrack filtering and mirroring the signal before filtering avoids generation of potentially harmful edge effects. Additionally, the DCT operates using real-valued coefficients which inherently allows for faster matrix computations than complex-valued methods such as the DFT.

To apply downtrack filtering to the measurement matrix, a projector \mathbf{P}_R is developed as

$$\mathbf{P}_R = \mathbf{D}\mathbf{I}_S \quad (4)$$

where \mathbf{D} is defined as the orthonormal DCT basis matrix with $\mathbf{D} \in \mathfrak{R}^{N \times N}$ and $\mathbf{I}_S \in \mathfrak{R}^{N \times N}$ is a selection matrix defined to keep the \hat{d} desired cosine frequencies. As described, the sensor self-response is mostly contained within the DC term of the spatial frequencies. High frequencies corresponding to short spatial wavelengths are expected only to contain noise. For these reasons, \mathbf{I}_S is chosen to keep the mid-spatial frequencies of the DCT which are believed to correspond mostly to the wavelengths of targets.

The downtrack projection is applied to the column space of \mathbf{M} by

$$\mathbf{M}_S = \mathbf{M}\mathbf{P}_R \quad (5)$$

so that $\mathbf{M}_S \in \mathfrak{R}^{M \times N}$ is the measured data matrix with the self-response and high-frequency noise removed.

After downtrack filtering has been applied, the mean of the real data components are subtracted from each sample to supply the final data matrix for prescreening.

3. METHODOLOGY

3.1 Discrete Spectrum of Relaxation Frequencies Dictionary

The foundation for success in a large majority of recent WEMI target detection algorithms can be attributed to the utilization of a *discrete spectrum of relaxation frequencies* (DSRF) dictionary.^{3,13} The DSRF is a discrete case of the *Distribution of Relaxation Times* (DRT) model.⁹ In contrast to the DRT which is formulated in terms of relaxation times, the DSRF utilizes relaxation frequencies to create a dictionary basis for the expected electromagnetic responses of targets. A benefit to using the DSRF is that it is directly related to the physical properties of potential targets and is thus invariant to the relative position and orientation of the WEMI sensor.^{3,4,14} The dictionary \mathbf{A} can be created by selecting a range of possible relaxation frequency terms ζ , and defining each element as

$$a(\omega, \zeta) = \frac{j\omega/\zeta}{1 + j\omega\zeta} \quad (6)$$

It follows that the atom for the k^{th} relaxation frequency can be constructed as

$$\mathbf{a}_k = \left[\frac{j\omega_1/\zeta_k}{1+j\omega_1\zeta_k} \quad \frac{j\omega_2/\zeta_k}{1+j\omega_2\zeta_k} \quad \cdots \quad \frac{j\omega_M/\zeta_k}{1+j\omega_M\zeta_k} \right]^T \quad (7)$$

where sensor operating frequency ω_m and the relaxation frequency ζ_k are both in radians. In this work, ζ was created from 100 equally log-spaced frequencies ranging from just below the minimum to just above the maximum sensor operating frequency. A visual depiction of an 18-frequency DSRF dictionary is presented in the work by Alvey, et al. and Cook, et al.^{3,4}

When using the DSRF dictionary for target detection, it is assumed that a target sample can be represented by a linear combination of the dictionary atoms of \mathbf{A} with a small amount of additive zero-mean Gaussian noise. Given the fixed dictionary, algorithms can be developed to search for the expected target responses defined in the atoms of \mathbf{A} .

3.2 Joint Orthogonal Matching Pursuits (JOMP)

One method which has adopted the DSRF for use in target detection is a well-known technique for sparse signal representation, *Matching Pursuits* (MP).^{15,16} MP is a greedy algorithm which finds linear approximations of signals by iteratively projecting them onto an over-complete dictionary. In MP, a sample \mathbf{x} is estimated as a sparse linear combination of p dictionary atoms, $\mathbf{D} = \{\mathbf{g}_d\}_{d=1}^K$ where $\|\mathbf{g}_d\| = 1$, described by

$$\hat{\mathbf{x}} = \sum_{j=0}^{p-1} \omega_j^{(x)} \mathbf{g}_{d_j}^{(x)} \quad (8)$$

where $\hat{\mathbf{x}}$ represents the approximation of \mathbf{x} , $\omega_j^{(x)}$ is the scaling factor for the contribution of the dictionary atom $\mathbf{g}_{d_j}^{(x)}$ to the full signal, and $p \ll K$. MP iteratively selects dictionary elements for signal estimation by first finding the dictionary element \mathbf{g}_{d_0} which produces the largest absolute value when \mathbf{x}_n is projected onto it

$$\mathbf{g}_{d_0}^{(x)} = \mathbf{g}_{d_0}, \quad \text{where } d_0 = \underset{k}{\operatorname{argmax}} |\langle \mathbf{x}, \mathbf{g}_k \rangle| \quad (9)$$

The residual error vector between the initial estimate $\hat{\mathbf{x}}$ and \mathbf{x} is given by

$$\mathbf{r}_1^{(x)} = \mathbf{x} - (\omega_0^{(x)} \mathbf{g}_{d_0}^{(x)}) \quad (10)$$

where $\omega_0^{(x)} = \langle \mathbf{x}, \mathbf{g}_{d_0}^{(x)} \rangle$ is the scaling factor for $\mathbf{g}_{d_0}^{(x)}$.

The signal estimation process continues iteratively by projecting the residual vectors $\mathbf{r}_i^{(x)}$ onto the dictionary atoms of \mathbf{D} and updating the following residuals $\mathbf{r}_{i+1}^{(x)}$ for p iterations. It follows that \mathbf{x} can be reconstructed completely by

$$\mathbf{x} = \hat{\mathbf{x}} + \mathbf{r}_p^{(x)} = \sum_{j=0}^{p-1} \omega_j^{(x)} \mathbf{g}_{d_j}^{(x)} + \mathbf{r}_p^{(x)} \quad (11)$$

with $\omega_j^{(x)} = \langle \mathbf{r}_j^{(x)}, \mathbf{g}_{d_j}^{(x)} \rangle$ and $\mathbf{r}_0^{(x)} \equiv \mathbf{x}$. Figure 1 demonstrates the residual calculation for a three-dimensional sample reconstructed with two dictionary elements.

An extension to MP for WEMI target detection is called *Joint Orthogonal Matching Pursuits* (JOMP).^{3,4,13,17} Instead of estimating a single sample signal with a given dictionary, the JOMP algorithm considers multiple signals simultaneously. Goldberg, et al. used the DSRF dictionary and a “two-tapped” JOMP (one sample positioned ℓ points ahead of the query and one sample ℓ behind) for explosive hazard detection. Instead of finding the reconstruction atoms for the query sample, the dictionary elements are selected jointly between the two points. The motivation for




Figure 1: Geometric interpretation of MP for $\mathbf{x} \in \mathfrak{R}^3$ and \mathbf{D} as a dictionary of two vectors, g_1 and g_2 . It can be observed that the residual vector is the subtraction of the projection of \mathbf{x} onto the space spanned by \mathbf{D} . Note that the residual vector is orthogonal to the projection onto \mathbf{D} .




Figure 2: Example of the data used in two-tapped JOMP. A confidence value is assigned to point $\mathbf{x}^{(i)}$ by jointly considering the reconstruction errors for samples $\mathbf{x}^{(i-\ell)}$ and $\mathbf{x}^{(i+\ell)}$.

this method is to reduce the number of false detections attributed to single anomalous, coherent signals by only flagging signatures large enough to extend over multiple scans. The residual error for each signal is found by $r_p^{(i)} = \left\| \hat{\mathbf{x}}^{(i)} - \mathbf{x}^{(i)} \right\|_2^2$, where i is the signal index. The final residual error is the average between the final residuals of the two signals ℓ points away from the center sample. The confidence value c is assigned to the location of \mathbf{x} at the center of the WEMI antenna by

$$c(\mathbf{x}^{(i)}, p) = \frac{1}{1 + \frac{1}{2}(r_p^{(i-\ell)} + r_p^{(i+\ell)})} \quad (12)$$

This value provides a measure of the representation ability of the DSRF for the two points being considered. If a target is present the confidence value should be close to 1, otherwise it should be near zero.

3.3 Spectral Matched Filter (SMF)

The *spectral matched filter* (SMF) is a statistical method for which the target and background signatures are modeled as random variables distributed according to some underlying probability distribution.¹⁸ The detection process can be formulated as a binary hypothesis test between the null hypothesis \mathcal{H}_0 (target absent) and \mathcal{H}_1 (target present). A detector can be modeled using the *generalized likelihood ratio test*. The hypotheses for SMF are

$$\begin{aligned}\mathcal{H}_0 : \mathbf{x} &\sim \mathcal{N}(0, \mathbf{\Sigma}_b) \\ \mathcal{H}_1 : \mathbf{x} &\sim \mathcal{N}(\mathbf{s}, \mathbf{\Sigma}_b)\end{aligned}\quad (13)$$

where $\mathbf{\Sigma}_b$ is the covariance of the background distribution and \mathbf{s} is a known target signature. The SMF detector resulting from the square root of the GLRT for Equation (13) is defined by

$$D_{SMF}(\mathbf{x}, \mathbf{s}) = \frac{\mathbf{s}^T \mathbf{\Sigma}_b^{-1} (\mathbf{x} - \boldsymbol{\mu}_b)}{\sqrt{\mathbf{s}^T \mathbf{\Sigma}_b^{-1} \mathbf{s}}}\quad (14)$$

where $\boldsymbol{\mu}_b$ is the mean of the background which is subtracted to ensure a zero-mean distribution defined by \mathcal{H}_0 .

The spectral matched filter can be interpreted as a target detector which measures the inner product between a sample \mathbf{x} and known target signature \mathbf{s} , scaled by the Mahalanobis distance of \mathbf{s} in the background space. This can be understood geometrically as the projection of an (unnormalized) test point onto a target vector in a whitened coordinate space. Since test points are not normalized, data samples with larger magnitude (that are not orthogonal to the test signature) are provided larger-valued detection statistics. This quality is desirable in WEMI target detection since highly conductive buried objects will inherently induce greater magnitude responses in the sensor.

In test, the SMF statistic is computed with each element in the DSRF target dictionary. The maximum statistic across all target representatives is assigned as the sample's final detection score.

3.4 Adaptive Cosine Estimator (ACE)

The *adaptive cosine estimator* (ACE) is an alternative detection statistic which closely resembles SMF.^{3, 18} The hypotheses utilized by ACE are provided by

$$\begin{aligned}\mathcal{H}_0 : \mathbf{x} &\sim \mathcal{N}(0, \sigma_0^2 \mathbf{\Sigma}_b) \\ \mathcal{H}_1 : \mathbf{x} &\sim \mathcal{N}(\mathbf{s}, \sigma_1^2 \mathbf{\Sigma}_b)\end{aligned}\quad (15)$$

where $\sigma_0^2 = \frac{1}{M} \mathbf{x}^T \mathbf{\Sigma}_b^{-1} \mathbf{x}$ and $\sigma_1^2 = \frac{1}{M} (\mathbf{x} - \mathbf{s})^T \mathbf{\Sigma}_b^{-1} (\mathbf{x} - \mathbf{s})$ to add scale-invariance with M as the number of features included in the data.

The square-root of the GLRT for ACE follows as

$$D_{ACE}(\mathbf{x}, \mathbf{s}) = \frac{\mathbf{s}^T \mathbf{\Sigma}_b^{-1} (\mathbf{x} - \boldsymbol{\mu}_b)}{\sqrt{\mathbf{s}^T \mathbf{\Sigma}_b^{-1} \mathbf{s}} \sqrt{(\mathbf{x} - \boldsymbol{\mu}_b)^T \mathbf{\Sigma}_b^{-1} (\mathbf{x} - \boldsymbol{\mu}_b)}}\quad (16)$$

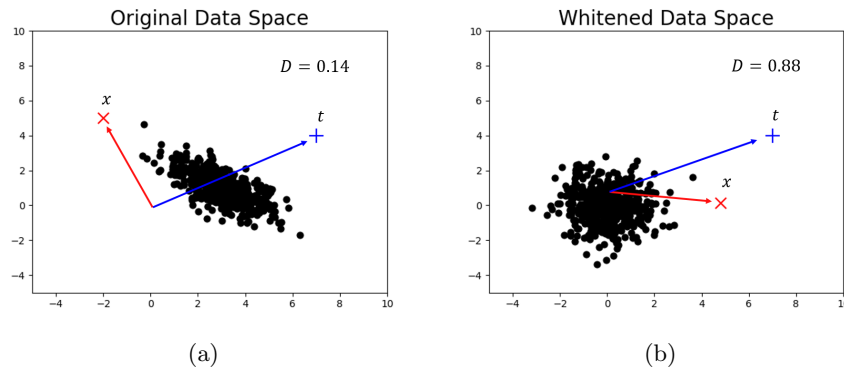


Figure 3: Example of target detection using both unwhitened (a) and whitened (b) data. The cosine similarity between a known target signature, t , and a data sample x is given as $D = 0.14$ in the original data space. After whitening by the mean and covariance of the background samples (black) the statistic increased to $D = 0.88$.

with notation matching the SMF detector described in Section 3.3.

Similar to SMF, the ACE detector can be interpreted as a measure of the inner product between a test sample and a known target signature in a whitened space. The most prominent difference between ACE and SMF is that ACE first normalizes test samples before measuring dissimilarity. The detection statistic then reduces simply to a measure of the vector angle between a test sample and a known target signature in a whitened space which takes values on the interval $[-1, 1]$. This measure can be viewed as a confidence representing the belief that a test sample comes from the target class distribution. Contrary to SMF, the magnitude of a test sample does not influence the ACE detection statistic value.

As with SMF, the ACE statistic is computed with each element in the DSRF target dictionary during testing. The maximum statistic across all target representatives is assigned as the sample's final detection confidence.

It should be noted that accurate background estimation is paramount to both ACE and SMF's success as detection statistics. Alvey, et al. explored a global background estimation and Woodbury update approach for use in ACE detection.³ Both methods assume that background samples follow a multivariate Gaussian random variable parameterized by a mean μ_b and covariance Σ_b . In an ideal environment, background samples would all be generated from the same stationary distribution. Detection approaches known as ACE and SMF Global leverage this idea by utilizing a single training set to empirically estimate the background statistics in one pass. While this method is non-causal, it is simple to implement and quick to execute. The Woodbury update could be employed to allow for online training, if necessary.

3.5 Multiple Instance ACE (MI-ACE)/ SMF (MI-SMF)

The WEMI hazard detection algorithms described in Sections 3.2-3.4 utilized prior target information through employment of the DSRF dictionary. However, it is intuitive that a target concept could be learned to take advantage of inherent properties exhibited by the data. Traditional supervised learning tasks are difficult to implement in WEMI detection because uncertainty in the

expanse of true target response often provides imprecise ground-truth. For this reason, WEMI detection tasks adhere well to *multiple instance learning* (MIL) models.^{4,18-23}

Multiple instance learning is a variation on supervised learning for problems with uncertain or imprecisely-labeled training data.¹⁸ Instead of pairing each training sample with a class label, MIL methods learn from a set of labeled concepts called “bags”. Each bag is defined as a multi-set of data points. A bag is labeled as “negative” if it is composed entirely of non-target training samples. Alternatively, “positive” bags are known to contain *at least one* positive target sample, or *instance*, although the true number of target samples within the bag is unknown.

The MIL problem is characterized by the following. Let $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_N] \in \mathfrak{R}^{M \times N}$ be training data where M is the dimensionality of instance \mathbf{x}_i , and N is the total number of training instances. The training data is divided into K bags, $\mathbf{B} = \{\mathbf{B}_1, \dots, \mathbf{B}_K\}$ with corresponding binary bag-level labels, $L \in \{L_1, \dots, L_K\}$ where $L_j \in \{0, 1\}$ and $\mathbf{x}_{ji} \in \mathbf{B}_j$ is the i^{th} instance in bag \mathbf{B}_j . Positive bags known to contain at least one positive target instance are given a bag-label $L_j = 1$ and are denoted by \mathbf{B}_j^+ . Alternatively, negative bags do not contain any positive target samples. These bags are provided labels $L_j = 0$, and are denoted \mathbf{B}_j^- . The positive and negative bag formulations can be described succinctly by

$$L_j = \begin{cases} 1, & \exists \mathbf{x}_{ji} \in \mathbf{B}_j^+ \ni \mathbf{x}_{ji} \sim \mathcal{N}(\mathbf{s} + \boldsymbol{\mu}_b, \sigma_1^2 \Sigma_b) \\ 0, & \mathbf{x}_{ji} \sim \mathcal{N}(\boldsymbol{\mu}_b, \sigma_0^2 \Sigma_b) \forall \mathbf{x}_{ji} \in \mathbf{B}_j^- \end{cases} \quad (17)$$

with notation matching Sections 3.3 and 3.4.

In contrast to ACE/ SMF Global which require an *a priori* signature dictionary, the objective of MI-ACE and MI-SMF is to learn a target signature, \mathbf{s} , that maximizes the appropriate detection statistic for the target instances in each of the positive bags while also minimizing the statistic across all negative instances. This task is realized by maximizing the objective:

$$\operatorname{argmax}_{\mathbf{s}} \frac{1}{N^+} \sum_{j:L_j=1} D(\mathbf{x}_j^*, \mathbf{s}) - \frac{1}{N^-} \sum_{j:L_j=0} \frac{1}{N_j^-} \sum_{\mathbf{x}_i \in \mathbf{B}_j^-} D(\mathbf{x}_i^*, \mathbf{s}) \quad (18)$$

where N^+ and N^- are the number of positive and negative bags, respectively, N_j^- is the number of instances in the j^{th} negative bag, and \mathbf{x}_j^* is the instance from positive bag \mathbf{B}_j^+ which is most likely a target signature. The most likely target signature from each positive bag is determined by

$$\mathbf{x}_j^* = \operatorname{argmax}_{\mathbf{x}_i \in \mathbf{B}_j^+} D(\mathbf{x}_i, \mathbf{s}) \quad (19)$$

implying that \mathbf{x}_j^* is the training sample with the maximum detection statistic given a target signature concept, \mathbf{s} . The training process is carried out by alternating between candidate signature selection from each of the positive bags and a statistic-dependent update step for \mathbf{s} given the chosen candidates, \mathbf{x}^* 's. Detailed derivations of the signature concept update equations for both SMF and ACE detection statistics are provided in the work by Zare, et al.¹⁸

MI-SMF and MI-ACE have both shown considerable success in the areas of hyperspectral target and explosive hazard detection.^{18,24} This can likely be attributed to the ability of MIL methods to handle uncertainty in training. Primary advantages of MI-ACE, as compared to alternative MIL target estimation methods, are that it does not require algorithm-specific hyper-parameters to be

set and also demonstrates comparatively fast run-times. These attributes facilitate learning of new target concepts in reliable and timely fashions. Following target concept estimation, test samples are classified in the same manner as Sections 3.3 and 3.4.

3.6 Soil Interference Removal

The algorithms previously described only utilized downtrack filtering as a method for interference removal. However, it has been shown that in some sensing environments, particular soil varieties can induce magnetic responses which are detrimental to the desired signal quality.^{1,25-27} An empirical soil model was constructed to define the expected signal presented by a variety of soil types.^{25,28} It was determined that the soil response can mostly be characterized as a function of frequency in a two-dimensional subspace

$$\hat{\Psi}(\omega) = \left[\mathbf{1} \quad \log \frac{\omega}{\omega_c} + \frac{j\pi}{2} \right] \quad (20)$$

where $\hat{\Psi}(\omega) \in \mathcal{C}^{K \times 2}$ and ω_c is a constant valued as the geometric mean of ω . The real and complex components can be concatenated as before to create a real-valued matrix such that $\Psi \in \mathfrak{R}^{M \times 2}$.

At a given position

$$\Xi(p) = [\xi_1(p) \quad \xi_2(p)] \quad (21)$$

where $\xi_1(p)$ and $\xi_2(p)$ are the strengths of the soil response. It follows that the soil interference can be modeled over a full sweep as

$$\mathbf{G}(\omega, p) = \Psi(\omega)\Xi^T(p) \quad (22)$$

where $\Psi(\omega) \in \mathfrak{R}^{M \times 2}$ and $\Xi(p) \in \mathfrak{R}^{N \times 2}$. It can be interpreted that the expected soil response at a given position can be decomposed as a linear combination of a DC term and a single frequency-dependent component.

While it has been assumed that WEMI sensors should only be effective at discovering targets demonstrating highly-conductive attributes, a recent work by Hayes, et al. showed that the empirical soil model may be further exploited for removing soil-response interference.^{10,11} Given appropriate conditions, this method may even be able to detect targets supplying a small conductive response or even lacking metal completely. The method uses a projection matrix P_G to transform the data into a space which is void of soil interference.

The projection matrix P_G takes advantage of the known frequency dependencies in the measured signal M . It was shown that the soil model G can be described in terms of a two-dimensional matrix in the frequency space. Additionally, Section 3.1 demonstrated that target signals show a frequency dependence which can be modeled as a sparse collection of atoms in a DSRF dictionary, A . Hayes showed that careful combination of A and G can be exploited to project the data into a space which is mostly isolated from the soil.²⁹ The creation of P_G is completed in two steps. First, a projection matrix into the soil subspace is developed as $P_{GG} = \Psi(\Psi^T\Psi)^{-1}\Psi^T$. It follows that A can be decomposed by

$$\mathbf{A} = P_{GG}\mathbf{A} + P_{GG}^\perp\mathbf{A} \quad (23)$$

with $P_{GG}^\perp = I - P_{GG}$ defined as the projection into the space orthogonal to the soil model. Next, a *singular value decomposition* (SVD) can be performed on the components of \mathbf{A} such that

$$\mathbf{A} = U_G^A \Sigma_G^A V_G^{A^T} + U_G^A \Sigma_G^A V_G^{A^T} \quad (24)$$

A target SNR threshold can be enforced to allow further division between signal and noise. In this manner, any components with a corresponding singular value below the threshold $\lambda_{G\mathcal{E}}$ can be considered as noise. Using this signal/ noise separation on $P_{GG}^\perp \mathbf{A}$, the DSRF can be reformulated as

$$\mathbf{A} = \begin{bmatrix} U_{GS}^A & U_{G\mathcal{E}}^A & U_G^A \end{bmatrix} \begin{bmatrix} \Sigma_{GS}^A & 0 & 0 \\ 0 & \Sigma_{G\mathcal{E}}^A & 0 \\ 0 & 0 & \Sigma_G^A \end{bmatrix} \begin{bmatrix} V_{GS}^{A^T} \\ V_{G\mathcal{E}}^{A^T} \\ V_G^{A^T} \end{bmatrix} \quad (25)$$

which resembles an SVD but has no guarantees of orthogonality between the right singular vectors corresponding to the soil subspace and the right singular vectors representing the space orthogonal to the soil model.

This decomposition of \mathbf{A} provides the desired orthonormal matrix for P_G as

$$P_G = \begin{bmatrix} P_{GS} \\ P_{G\mathcal{E}} \\ P_{GG} \end{bmatrix} \begin{bmatrix} U_{GS}^A & U_{G\mathcal{E}}^A & U_G^A \end{bmatrix}^T \quad (26)$$

It follows that the soil interference removal can be applied along with the self-response filtering described in Section 2.3 by

$$M_S = P_G M P_R^T \quad (27)$$

As described by Hayes, application of these projections creates a data measurement matrix which can be defined in subsets demonstrating highly unique attributes.¹⁰ This decomposition is shown in Figure 4.

Additional work by Hayes, et al. explored the value of individual sub-matrices resulting from application of the projection matrices.¹¹ In this work, however, only the subspace M_{GS}^{RS} is utilized, which is believed to contain the purest portion of the signal which is void of both soil and self-response interference. The feature generation in this work was completed by first dividing each sweep into overlapping windows centered at query samples, and applying the projection operations to each window, individually. This process provides local information to be exploited through the SVD process. The projection operations were performed on each data sub-matrix for the ℓ^{th} position of the sliding window

$$M_{GS}^{RS} = P_{GS} M_\ell P_{RM}^T \quad (28)$$

where M_ℓ is the ℓ^{th} window sub-matrix of the collected data, P_{GS} is the signal portion of the soil projection and P_{RM} is the mid-spatial frequency portion of the downtrack projection matrix.

An additional SVD was performed on each data window

$$M_{GS}^{RS} = U_S \Sigma_S V_S^T \quad (29)$$

and the left singular vectors U_S (corresponding to frequency) were scaled by the square root of the ratio between the corresponding and first singular values.^{10,11} The scaled representations were concatenated to create feature vectors.

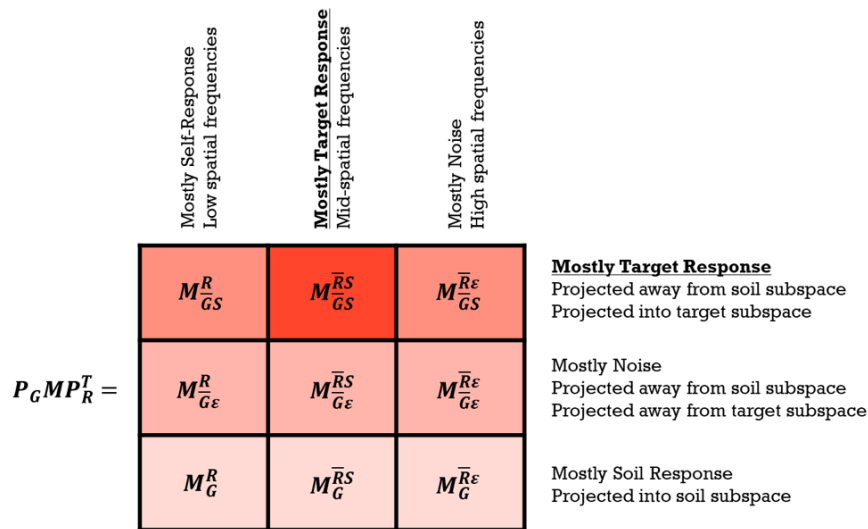


Figure 4: Data subspaces created through application of P_G and P_R .¹¹

4. EXPERIMENTS AND RESULTS

4.1 Data Set Description

Experiments were conducted on wideband electromagnetic induction data taken from two collection sites. Data collected from test site A has historically provided more ideal conditions for WEMI buried hazard detection. Although the challenges detailed in Section 1 would still be present, it was anticipated that targets from this site would be detected more easily than the alternative collection location. The site disclosed as B was known to provide highly challenging conditions for WEMI hazard detection. Expert knowledge determined that it was unlikely for low and non-metal targets to be unmasked at this location due to the strong soil-response interference emitted from this collection environment. Test sites were divided into physical lanes where explosive hazards were buried. Each lane was subdivided into approximately 1-2m² grids, local to each buried target. Operators swept the WEMI sensors coarsely in sinusoidal patterns across the grids. This is typically performed to locate potential alarms before performing more refined sweeps meant to center a discovered target. Four operators collected at each test site over the same grids using two WEMI sensors, A and B. A UTM position coordinate was recorded with each WEMI sample. Both data sets consist of collections across a wide assortment of explosive hazards, as well as metal and non-metal clutter. Targets were buried at various depths and are classified by their purpose, such as anti-tank (AT) and anti-personnel (AP). Targets are also classified by their metal content: metal target (MT), low-metal target (LMT), non-metal target (NMT), and clutter (CL). The number of targets considered in each subset for this experimentation is shown in Table 1. Blank data believed to represent background signature void of target response was also collected at locations between grids. Blank data was needed for both training and algorithmic scoring.

Number of Targets from Site A	
Data Subset	Number of Targets
Sensor A (Metal)	7
Sensor A (Low)	10
Sensor A (All)	167
Sensor B (Metal)	9
Sensor B (Low)	20
Sensor B (All)	41

Number of Targets from Site B	
Data Subset	Number of Targets
Sensor B (Metal)	57
Sensor B (Low)	196

Table 1: Number of target class instances contained in specific data subsets for test sites A and B. Subsets are denoted by the sensor and metal content.

4.2 Alarm Generation

Each prescreener output a *confidence map* where every UTM coordinate was associated with a value of belief that a target had been detected. Given a confidence map, *alarms* signifying points of interest were generated to mark spatial regions requiring further analysis. Alarm generation condenses groupings of high confidence samples into a single location as a way to mitigate false alarms caused by erratic, high-confidence samples. A variety of methods for alarm generation have been explored in the literature, including: non-maximal suppression, mean-shift clustering, morphological clustering, and connected components.^{3,24} A mean-shift algorithm with a radial basis function (RBF) kernel was used to condense the data points into alarms for scoring. Each test sample \mathbf{x}_i , was assigned a confidence value, c_i . The corresponding UTM coordinate vector is denoted by \mathbf{p}_i . The final alarm confidence was set as the mean confidence of all points within a 0.25m radius from the converged UTM location weighted by their Euclidean distances to the center. This weighting was performed to help mitigate spurious alarms caused by sensor cupping, centered on high-density edges of the collection. The width of the RBF kernel was empirically determined as $\sigma = 0.075$. The convergence threshold τ was set as 0.001. Pseudo-code for the mean-shift algorithm is provided in Algorithm 1.

Algorithm 1: Mean Shift Alarm Generation

Required: $\{\mathbf{p}_j\}_{j=1}^N, \{c_j\}_{j=1}^N$
for $j = 1, \dots, N$ **do**
 $\mathbf{p} \leftarrow \mathbf{p}_j$
 while $\|\Delta\mathbf{p}\| > \tau$ **do**
 $\forall i : p(i|\mathbf{p}) \leftarrow \frac{c_i \exp\left(-\frac{1}{2}\frac{\|\mathbf{p}-\mathbf{p}_i\|^2}{\sigma}\right)}{\sum_{i'=1}^N c_{i'} \exp\left(-\frac{1}{2}\frac{\|\mathbf{p}-\mathbf{p}_{i'}\|^2}{\sigma}\right)}$
 $\mathbf{p} \leftarrow \sum_{i=1}^N p(i|\mathbf{p})\mathbf{p}_i$
 end
 $\mathbf{z}_j \leftarrow \mathbf{p}$
end
Connected-components($\{\mathbf{z}_j\}_{j=1}^N, \epsilon$)

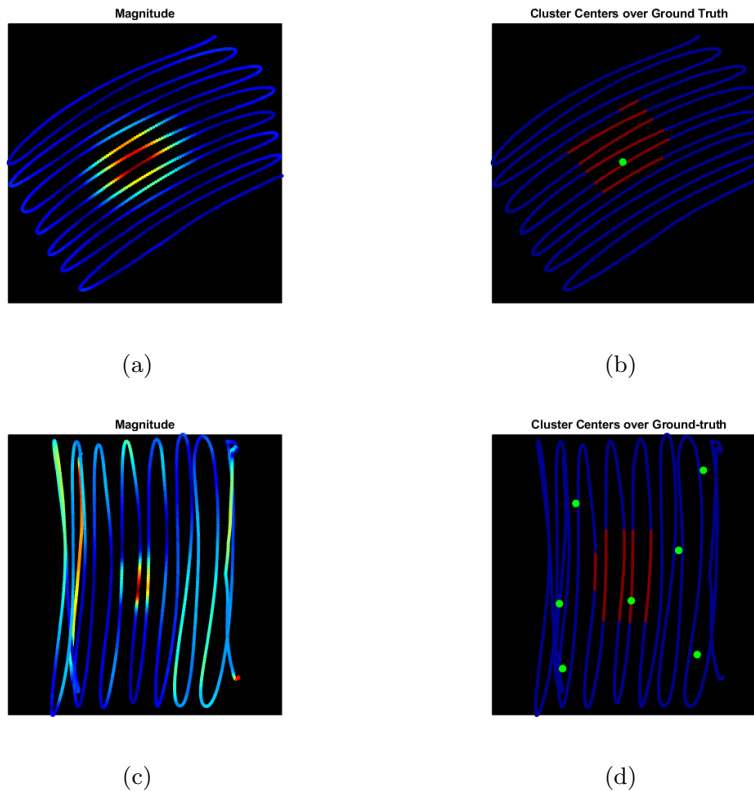


Figure 5: Examples of mean shift alarm generation over confidence maps. Plots (a) and (c) show confidences produced by taking the vector magnitudes of each sample at corresponding UTM locations. Plots (b) and (d) denote the final cluster center locations determined by the mean-shift algorithm in relation to the ground-truth regions. The mean-shift algorithm placed cluster centers at local maxima in the confidence maps.

For mean-shift, each data point begins as its own cluster center. In each iteration of the algorithm, centers are updated towards dense regions of high confidence in the map. The final positions of centers are located at local maxima.³⁰⁻³² Connected components are found on the map to quantify the converged data points and the generated alarms are passed along for scoring. Examples of alarm generation using mean-shift are shown in Figure 5.

4.3 Experiment Descriptions

In this section, the authors consider three target detection methodologies utilizing the DSRF dictionary: JOMP, ACE Global, SMF Global, as well as two methods trained and evaluated with lane-based cross-validation: MI-ACE, and MI-SMF. A simple magnitude detector was also tested to provide a baseline. The purpose of experimentation was three-fold. First, comparison was made between the WEMI sensors on Site A to determine algorithmic performance across a variety of platforms. Next, the trained and untrained detectors were evaluated for sensor B on Site B to

determine the benefits of multiple-instance learning on difficult data. Finally, a comparison of pre-processing procedures was performed with sensor B on test site B to discover their effectiveness in highly-contaminated soil environments.

In experimentation, each algorithm produced an alarm set signified by easting and northing UTM locations as well as confidence values. These confidence values are not all on the same scale. For example, JOMP is constrained to provide values between 0 and 1, ACE between -1 and 1, and SMF and magnitude greater than 0. Since the algorithms compared are all deterministic, tests were only run once. Mean-shift clustering was used to aggregate local regions of high-confidence in order to reduce the number of false alarms.

4.4 Experimental Results

4.4.1 Comparison of Sensors

The ROC curves shown in Figure 6 demonstrate the detection performance of the JOMP, ACE, SMF, and magnitude detectors for both sensors on test site A. Data subsets included high-metal only, low-metal only, and a combination of high, low, and non-metal targets. Figures 6a and 6d demonstrate that each high-metal target was discovered by every algorithm at relatively-low false alarm rates. This result was expected as test site A had previously supplied WEMI data showing relatively low amounts of soil response interference. Across the two sensors, SMF generally provided the best performance while the magnitude detector supplied the worst. The performance of the magnitude detector was, again, expected. Unlike the alternatives, the magnitude detector does not provide any mechanism for reducing false alarms. Figures 6b and 6e show detection accuracy for low-metal targets on linear and logarithmic scales, respectively. Figures 6c and 6f show performance on subsets containing all target classes. As can be gleaned from the plots, SMF typically demonstrated better performance than the alternatives on both sensors A and B. The magnitude detector followed SMF, as determined by Figure 6f. While each of the detectors employing the DSRF exploit spectral shape in their decision making, only SMF considers the magnitude of a sample when assigning its confidence. Thus, the result that SMF outperforms the alternatives implies that not only spectral shape, but also magnitude, are important features for hazard discovery. Performance evaluation for the magnitude detector is likely less reliable for low and non-metal targets as they were not scored on the same scale as the high-metal examples. Improved evaluation could likely be performed using a lane-based collection where data would be accumulated in a continuous stream without breaks for fiducial marking. Each prescreener performed poorly on the test set containing high, low, and non-metal classes. Results, however, are likely skewed due to an imbalance of target types represented in the test set.

4.4.2 Comparison of Target Representatives

A comparison of detectors utilizing the DSRF versus methods which were trained was conducted with the sensor B for test site B on high-metal and low-metal subsets. Sensor B was chosen for this experimentation based on its frequency structure in order to facilitate comparison of preprocessing methodologies in Section 4.4.3. Test site B was used to evaluate performance in a high-interference environment. The ROC curves shown in Figure 7 demonstrate that hazard detection was fairly less accurate in comparison with the results shown for test site A in Section 4.4.1. This was expected due to the anticipated amounts of soil-interference presented by Site B. Relative algorithmic performance was unchanged between test sites on the high-metal subset. In observation, prescreeners run on test Site B typically produced confidence halos with larger radii around the expected target centers than

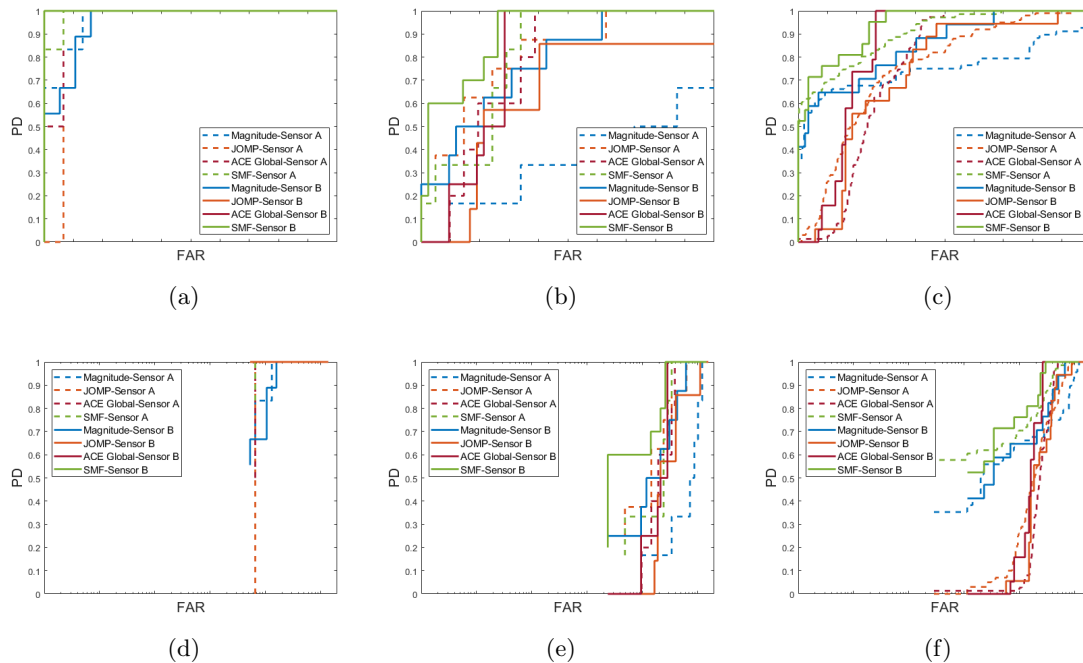


Figure 6: ROC curves demonstrating the detection performance of the untrained algorithms for both sensors on test site A. The ROCs are divided into the following subsets: (a) high-metal, (b) low-metal, (c) all target types, (d) high-metal plotted on a logarithmic scale, (e) low-metal on a logarithmic scale, and (f) all target classes on a logarithmic scale.

location A. This artifact could have potentially caused mean-shift to create additional alarms for the detected targets which fell outside the estimated ground-truth region, thus degrading the ROC evaluation. This inaccuracy in scoring stems from known difficulties in providing accurate sample-level ground-truth labels for WEMI data. Careful selection of the mean-shift kernel bandwidth was taken in attempt to mitigate this undesirable effect; although, selection of this parameter is difficult due to uncertainty in the radius of response provided by alternative targets in differing soil environments.

Because of the ground-truth uncertainty, it was hypothesized that the multiple-instance approaches taken by MI-ACE and MI-SMF would outperform the alternatives. While this was not the case overall, multiple-instance methods showed promising results on low-metal targets. Degradation of performance for MI-ACE and MI-SMF can most likely be attributed to the fact that training and evaluation of the algorithms was performed by learning single target representatives through lane-based cross validation. Compared to the DSRF dictionary which was composed of 100 signatures, MI-ACE and MI-SMF each learned only a single target signature to represent all target classes. However, this approach to training is likely sub-optimal, as experience has shown that even targets belonging to the same class can exhibit highly varying spectral signatures. Moreover, evaluation of performance was left slightly to chance since learned target signatures had potential to be tested on specific target classes which were absent from their corresponding training folds.

Given the inherent nature of uncertainty presented in the buried hazard detection problem, MI-ACE and MI-SMF would likely outperform the alternative approaches if multiple target signatures were estimated from target-specific subsets.

Figures 7a and 7c show that each algorithm had difficulties discovering low-metal targets, although MI-SMF performed better than the rest. This was anticipated due to the interference qualities provided by the test site. MI-ACE and JOMP provided the next best detection performances on the low-metal subset.

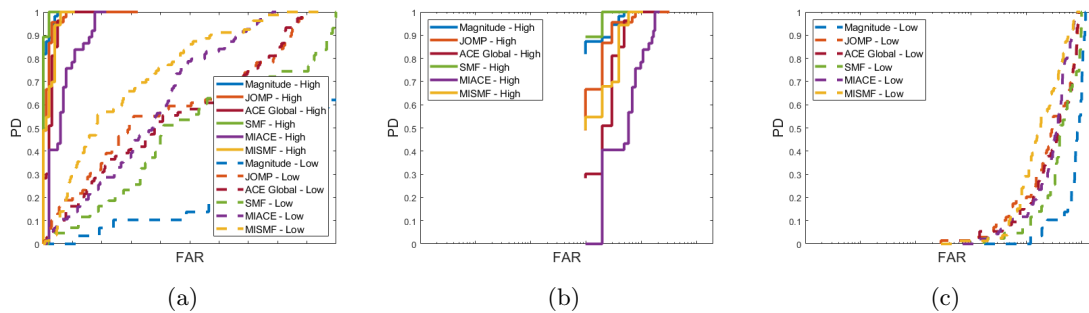


Figure 7: ROC curves demonstrating the detection performance of both trained and untrained algorithms for sensor platform B on test site B. The ROCs are divided into the following subsets: (a) high and low-metal, (b) high-metal only on a logarithmic scale, (c) low-metal only on a logarithmic scale. Solid lines correspond to tests run on the high-metal subset, while dashed lines represent low-metal experiments.

4.4.3 Comparison of Pre-processing

Following the results presented in Section 4.4.2, an investigation to the effectiveness of the soil interference removal procedure presented by Hayes, et al. was performed on the highly-conductive test environment, Site B. In their method, a projection is applied in addition to DCT filtering in order to transform the data into a space which is orthogonal to an empirically-defined soil model. The first three frequency singular vectors produced by this processing were weighted by ratios of their corresponding singular values and concatenated to form feature vectors. Only MI-ACE and MI-SMF were evaluated in this test, although the soil projection could potentially be applied to the DSRF as well. Performance results are exhibited in Figure 8. While the same complications discussed in Section 4.4.2 arose in training and evaluation, it can be observed that the DCT pre-processing for self-response removal significantly outperformed the soil projection method for both high and low-metal subsets. Upon further investigation, the soil projection method provided high confidence to the majority of the 1-2 m^2 grids under test. The authors believe this attribute is a result of the windowing needed to perform the SVD estimation. Because of this spatial smoothing artifact, performance of the projection processing should be evaluated on data collected in a ‘localization’ mode where there would exist a long, continuous stream of data with well-separated targets. This type of data would be conducive to showing if there is, in fact, a difference in detection across targets, as well as validate the authors’ hypothesis that the target signature gets spread throughout the grids in the current collection. Further testing is needed to fully evaluate the effectiveness of the soil projection method.




Figure 8: ROC curves demonstrating the detection performance of MI-ACE and MI-SMF for DCT and Projection pre-processing. Tests were conducted on high and low-metal subsets collected with sensor platform B at test site B. Solid lines correspond to tests on data employing DCT pre-processing while dashed lines demonstrate the projection approach.

5. CONCLUSION

This paper compared six alternative WEMI target detection methodologies in terms of effectiveness for buried, explosive hazard discovery. As expected, high-metal targets were distinguishable by each of the algorithms compared while low and non-metal targets proved more challenging. Four primary points were determined through experimentation. First, it was concluded that both magnitude and spectral shape provide valuable information for exploitation in hazard discovery. Next, results demonstrated that prescreening methods which utilize training to learn positive target concepts would likely benefit from estimating signatures for individual target types, rather than a single, global target representative. Additionally, performance of the pre-processing designed to remove soil-interference, which did not prove efficacious in this work, is likely attributed to the the grid-based collection scheme employed at the test sites. Investigating the use of this method in a ‘localization’ setting where the averaging effects of windowing would be diminished would likely provide an avenue for more accurate analysis. Finally, evaluation of WEMI target detection methods is difficult due to ground-truth uncertainty. Development of more precise scoring metrics would greatly benefit this area of research.

ACKNOWLEDGMENTS

This work was funded by Army Research Office grant number W911NF-17-1-0213 to support the US Army RDECOM CERDEC NVESD. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies either expressed or implied, of the Army Research Office, Army Research Laboratory, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

REFERENCES

- [1] Vijayakumar, S. and Thomas, M. J., “Study of propagation of ultra-wide band electromagnetic pulse through a dispersive soil medium for the detection of buried unexploded ordnance,” in [2015 13th Int. Conf. on Electromagnetic Interference and Compatibility (INCEMIC)], 81–86 (July 2015).

- [2] Scott Jr., W. R., “Broadband electromagnetic induction sensor for detecting buried landmines,” in [*2007 IEEE International Geoscience and Remote Sensing Symposium*], 22–25 (July 2007).
- [3] Alvey, B., Zare, A., Cook, M., and Ho, D. K. C., “Adaptive coherence estimator (ACE) for explosive hazard detection using wideband electromagnetic induction (WEMI),” in [*Proc. SPIE 9823, Detection and Sensing of Mines, Explosive Objects, and Obscured Targets XXI, 982309*], *Proc. SPIE* **9823**, 9823 – 9823 – 7 (Apr. 2016).
- [4] Cook, M., Zare, A., and Ho, D. K. C., “Buried object detection using handheld wemi with task-driven extended functions of multiple instances,” in [*Proc. SPIE 9823, Detection and Sensing of Mines, Explosive Objects, and Obscured Targets XXI, 98230A*], *Proc. SPIE* **9823**, 9823 – 9823 – 9 (Apr. 2016).
- [5] Dula, J., Zare, A., Ho, D. K. C., and Gader, P., “Landmine classification using possibilistic k-nearest neighbors with wideband electromagnetic induction data,” in [*Proc. SPIE 8709, Detection and Sensing of Mines, Explosive Objects, and Obscured Targets XVIII, 87091F*], *Proc. SPIE* **8709**, 8709 – 8709 – 12 (June 2013).
- [6] Fails, E. B., Torrione, P. A., Scott Jr., W. R., and Collins, L. M., “Performance comparison of frequency domain quadrupole and dipole electromagnetic induction sensors in a landmine detection application,” in [*Detection and Sensing of Mines, Explosive Objects, and Obscured Targets XIII, 695304*], *Proc. SPIE* **6953**, 6953 – 6953 – 11 (Apr. 2008).
- [7] Scott Jr., W. R., “Broadband array of electromagnetic induction sensors for detecting buried landmines,” in [*IGARSS 2008 - 2008 IEEE Int. Geoscience and Remote Sensing Symp.*], **2** (July 2008).
- [8] Scott Jr., W. R. and Larson, G. D., “Measured dipole expansion of discrete relaxations to represent the electromagnetic induction response of buried metal targets,” in [*Detection and Sensing of Mines, Explosive Objects, and Obscured Targets XV, 76640E*], *Proc. SPIE* **7664**, 7664 – 7664 – 11 (Apr. 2010).
- [9] Wei, M., Scott Jr., W. R., and McClellan, J. H., “Robust estimation of the discrete spectrum of relaxations from multiple-electromagnetic induction responses,” in [*Detection and Sensing of Mines, Explosive Objects, and Obscured Targets XVII, 83570A*], *Proc. SPIE* **8357**, 8357 – 8357 – 9 (May 2012).
- [10] Hayes, C. E., McClellan, J., and Scott Jr., W. R., “Novel model based EMI processing framework,” in [*Detection and Sensing of Mines, Explosive Objects, and Obscured Targets XXII, 1018218*], *Proc. SPIE* **10182**, 10182 – 10182 – 17 (May 2017).
- [11] Hayes, C. E., Scott Jr., W. R., and McClellan, J. H., “Exploiting measurement subspaces for wideband electromagnetic induction processing,” in [*Detection and Sensing of Mines, Explosive Objects, and Obscured Targets XXIII, 106280M*], *Proc. SPIE* **10628**, 10628 – 10628 – 20 (Apr. 2018).
- [12] Hayes, C. E., McClellan, J. H., Scott Jr., W. R., and Kerr, A. J., “Improved electromagnetic induction processing with novel adaptive matched filter and matched subspace detection,” in [*Detection and Sensing of Mines, Explosive Objects, and Obscured Targets XXI, 98230E*], *Proc. SPIE* **9823**, 9823 – 9823 – 14 (May 2016).
- [13] Goldberg, S., Glenn, T., Wilson, J. N., and Gader, P. D., “Landmine detection using two-tapped joint orthogonal matching pursuits,” in [*Detection and Sensing of Mines, Explosive Objects, and Obscured Targets XVII, 83570B*], *Proc. SPIE* **8357**, 8357 – 8357 – 8 (May 2012).

- [14] Smith, R., Anderson, D., Ball, J., Zare, A., and Alvey, B., “Aggregation of choquet integrals in GPR and EMI for handheld platform-based explosive hazard detection,” in [*Detection and Sensing of Mines, Explosive Objects, and Obscured Targets XXII, 1018217*], *Proc. SPIE* **10182**, 10182 – 10182 – 15 (May 2017).
- [15] Mallat, S. G. and Zhang, Z., “Matching pursuits with time-frequency dictionaries,” *IEEE Transactions on Signal Processing* **41**, 3397–3415 (Dec 1993).
- [16] Mazhar, R., Gader, P. D., and Wilson, J. N., “Matching-pursuits dissimilarity measure for shape-based comparison and classification of high-dimensional data,” *IEEE Transactions on Fuzzy Systems* **17**, 1175–1188 (Oct 2009).
- [17] Cai, T. T. and Wang, L., “Orthogonal matching pursuit for sparse signal recovery with noise,” *IEEE Transactions on Information Theory* **57**, 4680–4688 (July 2011).
- [18] Zare, A., Jiao, C., and Glenn, T. C., “Multiple instance hyperspectral target characterization,” *CoRR* **abs/1606.06354** (2016).
- [19] Jiao, C. and Zare, A., “Multiple instance hybrid estimator for learning target signatures,” in [*2017 IEEE Int. Geoscience and Remote Sensing Symp. (IGARSS)*], 988–991 (July 2017).
- [20] Jiao, C. and Zare, A., “Functions of multiple instances for learning target signatures,” *IEEE Transactions on Geoscience and Remote Sensing* **53**, 4670–4686 (Aug 2015).
- [21] Maron, O. and Lozano-Pérez, T., “A framework for multiple-instance learning,” in [*Advances in Neural Information Processing Systems 10*], Jordan, M. I., Kearns, M. J., and Solla, S. A., eds., 570–576, MIT Press (1998).
- [22] Shrivastava, A., Pillai, J. K., Patel, V. M., and Chellappa, R., “Dictionary-based multiple instance learning,” in [*2014 IEEE Int. Conf. on Image Processing (ICIP)*], 160–164 (Oct 2014).
- [23] Zhang, Q. and Goldman, S. A., “EM-DD: An improved multiple-instance learning technique,” in [*In Advances in Neural Information Processing Systems*], 1073–1080, MIT Press (2001).
- [24] Alvey, B., Ho, D. K. C., and Zare, A., “Fourier features for explosive hazard detection using a wideband electromagnetic induction sensor,” in [*Detection and Sensing of Mines, Explosive Objects, and Obscured Targets XXII, 101820E*], *Proc. SPIE* **10182**, 10182 – 10182 – 9 (May 2017).
- [25] Das, Y., “Effects of soil electromagnetic properties on metal detectors,” *IEEE Trans. on Geoscience and Remote Sensing* **44**, 1444–1453 (June 2006).
- [26] Cross, G., “Soil electromagnetic properties and metal detector performance: Theory and measurement,” (Nov. 2008).
- [27] Takahashi, K., Preetz, H., and Igel, J., “Influence of soil properties on the performance of metal detectors and GPR,” *Journal of ERW and Mine Action* **17**, 52–57 (01 2013).
- [28] Dabas, M., Jolivet, A., and Tabbagh, A., “Magnetic susceptibility and viscosity of soils in a weak time varying field,” *Geophysical Journal International* **108**, 101 – 109 (Apr. 2007).
- [29] Hayes, C. E., McClellan, J. H., and Scott Jr., W. R., “Sparse recovery using an svd approach to interference removal and parameter estimation,” in [*2015 IEEE Signal Processing and Signal Processing Education Workshop (SP/SPE)*], 202–207 (Aug 2015).
- [30] Fukunaga, K. and Hostetler, L., “The estimation of the gradient of a density function, with applications in pattern recognition,” *IEEE Trans. on Info. Theory* **21**, 32–40 (Jan. 1975).
- [31] Carreira-Perpiñán, M. Á., “A review of mean-shift algorithms for clustering,” *CoRR* **abs/1503.00687** (2015).
- [32] Comaniciu, D. and Meer, P., “Mean shift: A robust approach toward feature space analysis,” *IEEE Trans. Pattern Anal. Mach. Intell.* **24**, 603–619 (2002).