

FINAL REPORT

Demonstration and Validation of Statistical Analysis Techniques
for TOI Discrimination Using Advanced EMI Sensor Systems

ESTCP Project MR-201156

JUNE 2012

Lawrence Carin
Levi Kennedy
Todd Jobe
Xianyang Zhu
Signal Innovations Group, Inc.

This document has been cleared for public release



REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

1. REPORT DATE (DD-MM-YYYY) 14/06/12		2. REPORT TYPE Final Report		3. DATES COVERED (From - To)	
4. TITLE AND SUBTITLE UXO Discrimination Study Camp Butner, NC: Final Report				5a. CONTRACT NUMBER MR-201156	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Larry Carin (PI), Levi Kennedy, Todd Jobe, Xianyang Zhu				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Signal Innovations Group, Inc. 4721 Emperor Blvd., Suite 330 Durham, NC 27703				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT This report details the application of the SIG statistical learning approach to UXO discrimination for Camp Butner, North Carolina. This technology has been developed and validated under previous SERDP/ESTCP efforts by SIG and Duke University. Specific core technologies were used in this discrimination. These technologies fall broadly into the four analysis categories: the sensor/target model, feature selection, classification, and active label selection. The non-linear classifier outperformed the linear classifier. Both linear and non-linear classifiers would have left more than 75% of the clutter in the ground. The stopping point for both classifiers left UXO in the ground, however. Two of these anomalies could have been captured earlier by selecting additional features. This study validated the robustness of key SIG technologies for target/sensor models, feature selection, classification, and active learning.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT SAR	18. NUMBER OF PAGES 21	19a. NAME OF RESPONSIBLE PERSON Levi Kennedy
a. REPORT u	b. ABSTRACT u	c. THIS PAGE u			19b. TELEPHONE NUMBER (include area code) 919-323-3456

Table of Contents

List of Acronyms	ii
List of Figures	ii
List of Tables	iii
Executive Summary	1
1. Introduction.....	2
1.1. Background	2
1.2. Objective of the Demonstration	2
2. Technology	2
2.1. Technology Description	2
2.2. Technology Development	5
3. Performance Objectives	7
3.1. Maximize correct classification of targets of interest	7
3.2. Maximize correct classification of non-UXO	7
3.3. Specification of no-dig threshold	8
3.4. Minimize the number of anomalies that cannot be analyzed	8
4. Site Description.....	8
5. Test Design	8
6. Data Analysis and Products	8
6.1. Parameter Estimates	8
6.2. Feature Selection	9
6.3. Training and Classification	10
6.4. Excavation.....	11
7. Performance Assessment	14
7.1. Maximize correct classification of targets of interest	14
7.2. Maximize correct classification of non-UXO	14
7.3. Specification of no-dig threshold	15
7.4. Minimize the number of anomalies that cannot be analyzed	15
8. Cost Assessment	15
8.1. Cost Model	15
8.2. Cost Drivers.....	16
8.3. Cost Benefit.....	16
9. Implementation Issues	16
10. References.....	16

11.	Appendices.....	16
11.1.	Appendix A: Points of Contact.....	16

List of Acronyms

BENet	Bayesian Elastic Net
EMI	Electromagnetic Induction
MTL	Multi-task learning
NMAL	Non-myopic active learning
PNBC	Parametric Neighborhood-Based Classifier
ROC	Receiver Operating Characteristic
SIG	Signal Innovations Group, Inc.
SLO	San Luis Obispo Demonstration Site
STL	Single-Task Learning
TEMTADS	Time-domain Electromagnetic Multi-sensor Tower Array Detection System
UXO	Unexploded Ordnance

List of Figures

Figure 1.	Flow diagram of the SIG Isolate process.	4
Figure 2:	ROC curves for UXO classifier at SLO site with features selection using the BENet algorithm (red line) and without feature selection (blue line). The number of false alarms is lower for the classifier where feature selection was used.	5
Figure 3:	A comparison between supervised and semi-supervised classifiers for a 2. Labeled data from both classes (red and green circles) are shown, along with unlabeled data (black dots). The supervised classifier is trained on only the labeled data and the decision boundary is shown (dotted line). The semi-supervised classifier is trained on both the labeled and unlabeled data and the decision boundary (solid line) makes the two classes linearly separable.	5
Figure 5.	BENet weights for the polarizabilities of the TEMTADS dataset. Selected feature names are shown (tx is time gate x and ax is axis x).	8
Figure 4.	Histograms of fit errors for one (left), two (middle) and three (right) dipole inversion results for the Camp Butner TEMTADS sensor.	8
Figure 6.	Initial basis selection of Butner TEMTADS data along two primary features. The two features are the log polarizability of the primary axis at the first time-gate. Test pit UXO are also shown. M1 is the primary axis. M2 is the secondary axis. T1 is the first time gate.	9
Figure 7.	Labeled data at the end of training. UXO (red) and clutter (blue) training labels are shown along with the last round of actively learned labels (pink squares). Obvious clusters of munitions are highlighted.	11
Figure 8.	Non-linear PNBC classification boundary along with test and training data. Contour intervals are 0.1 posterior predicted probability of being UXO.	12
Figure 9.	ROC curve for the linear classification of the Camp Butner TEMTADS data.....	13
Figure 10.	ROC curve for the non-linear classification of the Camp Butner TEMTADS data....	13

Figure 11. Polarizabilities for anomalies that were difficult to discriminate for both the linear and non-linear models. IDs for each anomaly are shown in the title of the plots. 14

List of Tables

Table 1. Program Office Performance Objectives for Discrimination Analysis 7

Table 2. Cost Model for the SIG Discrimination at Camp Butner 15

Executive Summary

Signal Innovations Group, Inc. (SIG) has previously demonstrated the effectiveness of site-specific statistical learning for smartly selecting labeled training data to maximize target discrimination. This report details the application of the SIG statistical learning approach to unexploded ordnance (UXO) discrimination for Camp Butner, North Carolina. This technology has been developed and validated under previous SERDP/ESTCP efforts by SIG and Duke University. Specific core technologies were used in this discrimination. These technologies fall broadly into the four analysis categories: the sensor/target model, feature selection, classification, and active label selection. Feature selection was performed using the Bayesian Elastic Net which has the benefit of retaining correlated and informative features for classification. Classification was performed using two approaches: one a linear semi-supervised Bayesian classifier, and a non-linear semi-supervised Bayesian classifier.

The objectives of the study were to maximize correct classification of UXO and non-UXO, specify a no-dig threshold, and minimize the number of anomalies that could not be analyze. Most of the UXO items were detected and, generally, a substantial number of non-UXO were left unexcavated. Usable features were extracted for 98% of the anomalies. Feature selection significantly improved the performance of the classifiers. The non-linear classifier outperformed the linear classifier. Both linear and non-linear classifiers would have left more than 75% of the clutter in the ground. The stopping point for both classifiers left UXO in the ground, however. Two of these anomalies could have been captured earlier by selecting additional features. The goal of the SIG discrimination process is to provide a significant degree of automation for UXO discrimination problems. This study validated the robustness of key SIG technologies for target/sensor models, feature selection, classification, and active learning. These technologies are broadly applicable, and scalable to production level UXO remediation.

1. Introduction

1.1. Background

Signal Innovations Group, Inc. (SIG) has previously demonstrated the effectiveness of site-specific statistical learning for smartly selecting labeled training data to maximize target discrimination. This report details the application of the SIG statistical learning approach to UXO discrimination for Camp Butner, North Carolina. This technology has been developed and validated under previous SERDP efforts by SIG and Duke University.

Many current analysis approaches rely on expert scientists to make educated decisions at multiple points in the discrimination analysis process. This situation is not scalable, transferable, or cost effective. The SIG approach standardizes the options and creates a documented process flow that can be explicitly followed.

1.2. Objective of the Demonstration

The main technical objective of the Camp Butner demonstration is to validate and substantially automate the SIG learning process using next-generation electromagnetic induction (EMI) sensor data for discriminating targets-of-interest. All elements of human interpretation and intuition are being incrementally constrained or removed from the process, resulting in an automated process, where all algorithm parameters and thresholds will either be determined by specified site parameters (i.e., expected or inferred munitions types) or by data-driven inferences (i.e., cross-validated operating threshold). SIG applied and matured each of the three key process phases that constitutes the SIG statistical learning approach to UXO discrimination - called the “SIG Isolate” process. The three phases of Isolate include: Phase I - feature extraction, Phase II – site learning, and Phase III – excavation. Each of the phases is described in detail below

2. Technology

SIG applied the Isolate discrimination process in the Camp Butner demonstration for the TEMTADS sensor. The SIG Isolate process involves the following key technologies including: Bayesian feature selection, semi-supervised classifier training, and non-myopic active selection of labeled data. These methods are described briefly in the following subsections.

2.1. Technology Description

The SIG Isolate process laid out in [5] can be summarized in the following ‘recipe’ (Figure 1):

- **Data Conditioning** - First, raw, unlabeled anomaly data are received.
- **Subspace Denoising** - The anomaly data is denoised to ensure robust performance for discriminating late time-gate features.
- **Feature Extraction** - A robust multi-anomaly dipole model is fitted to the data. The polarizability parameters from this fitting become the set from which features are drawn for classifier training. In addition to the time-domain polarizabilities, a set of 9 ‘rate’ features were calculated. These features were the calculated by fitting the time-domain polarizabilities of each axis to an exponential-decay model:

$$p_i = r_{1i} + r_{2i}e^{\frac{-t}{\tau_{3i}}}$$

where $i \in \{x, y, z\}$ is the current axis, p is the polarizability, t is time and $\{\tau_1, \tau_2, \tau_3\}$ are the fitted rate parameters. Though r_{1i} is unphysical, it is useful for adjusting for noise at late time gates and where odd responses would make the optimization difficult. The optimized values of the rate parameters were found using non-linear least squares.

- **Basis Selection** - A few of the many possible features are selected based on their physical interpretation as they relate to the anomaly, and, using these features, the most informative set of anomalies are selected via an information metric to begin classifier training.

- **Feature Set Augmentation** - The feature set is then augmented by adding early, mid and late time polarizabilities values.
- **Automated Features Selection** - For the now larger feature set, the most relevant set of features is selected using BENet.
- **Semi-supervised PNBC Training (STL or MTL)** - When the PNBC is trained only using data from the current site of interest, it is called Single Task Learning (STL). When the PNBC is trained for multiple sites simultaneously it is called MTL. For the Camp Butner demonstration only STL was used.
- **Non-myopic Active Learning** - Based on the estimates made with the PNBC classifier, a new set of anomalies will be selected for labeling using NMAL. The goal at this step is to maximize the information gain from new labels requested from the set of unlabeled anomalies. The process is repeated as the PNBC classifier adequately learns data manifold. The stopping criterion for the learning process is apparent when the remaining unlabeled data points have approximately equal information for improving the classifier. At which point, labeling any one anomaly is no better than any other.
- **Excavation Adapted Threshold Selection** - At this point, the highest probability UXO are selected for excavation and labels. The classifier continues to be retrained when new labels are revealed. This process continues until the highest probability UXO items excavated are all found to be clutter at which point digging stops.

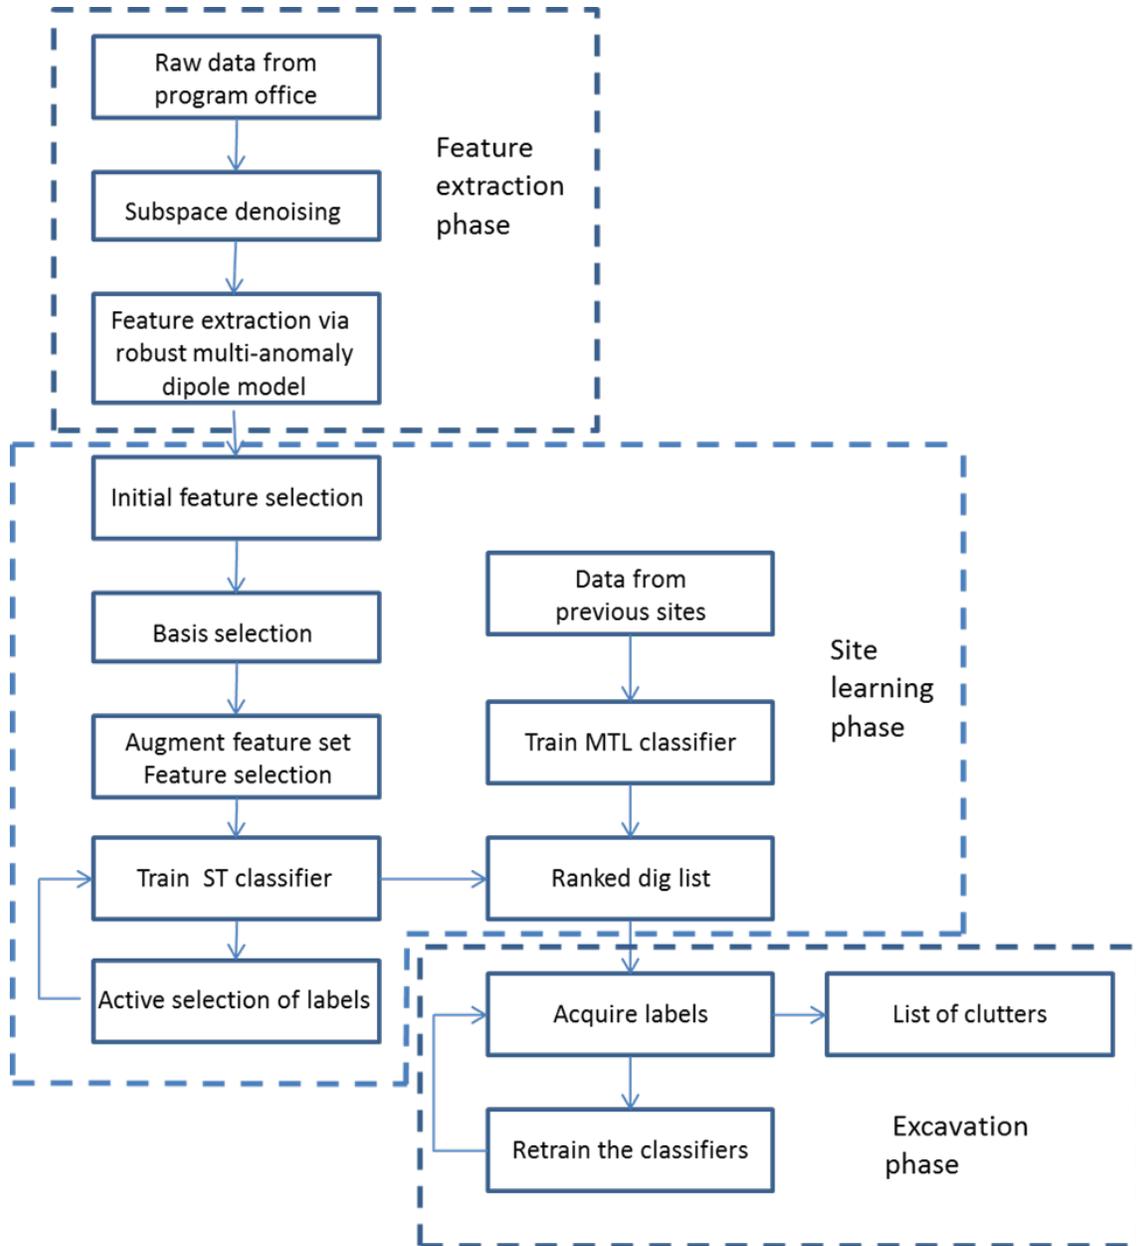


Figure 1. Flow diagram of the SIG Isolate process.

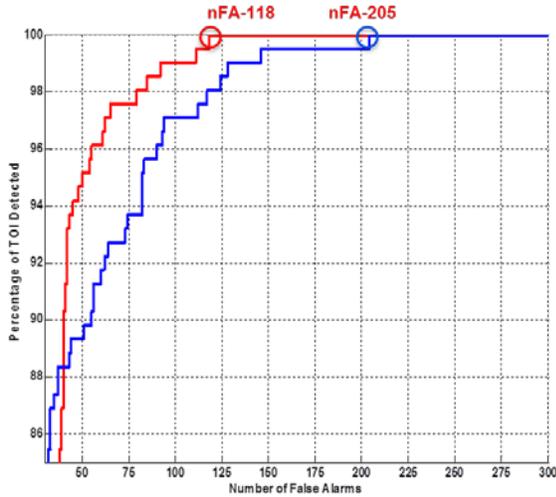


Figure 3: ROC curves for UXO classifier at SLO site with features selection using the BENet algorithm (red line) and without feature selection (blue line). The number of false alarms is lower for the classifier where feature selection was used.

2.2. Technology Development

Feature Selection with BENet

Adaptive learning of a classifier *in situ* benefits from refining the appropriate set of extracted features for the targets under test. This occurs because of the ‘curse of dimensionality’ where the number of data points required to cover the breadth of a features space grows exponentially with the number of features considered. If the amount of training data does not sufficiently sample the feature space, then the learned classifier will lack statistical support and class estimate uncertainty is large. At the San Luis Obispo (SLO) demonstration site in particular, feature selection played a key role in classifier performance (Figure 2). Bayesian classification models perform feature selection by placing a sparseness prior on the inferred feature weights. The Bayesian elastic net (BENet) regression model used for feature selection employs a sparseness prior equivalent to a convex combination of $L1$ -norm and $L2$ -norm penalties in a least squares optimization formulation [3], [2]. The sparseness prior of the BENet model jointly infers the essential subset of relevant features, including correlated features, for a given classification task. Rather than encouraging the selection of a single feature in a set of correlated important features (like similar approaches such as RVM), the BENet model encourages the selection of all correlated important features. By performing sparse and grouped feature selection, the BENet algorithm provides a more robust approach to feature adaptability and the interpretation of important features, requiring fewer training data samples to achieve robust statistical support.

Semi-Supervised Classification

Semi-supervised learning is applicable to any sensing problem for which all of the labeled and unlabeled data are available at the same time, and therefore, particularly for the current demonstration study. In most practical applications (including the recent demonstration at Camp SLO), semi-supervised learning has been found to yield superior performance relative to widely

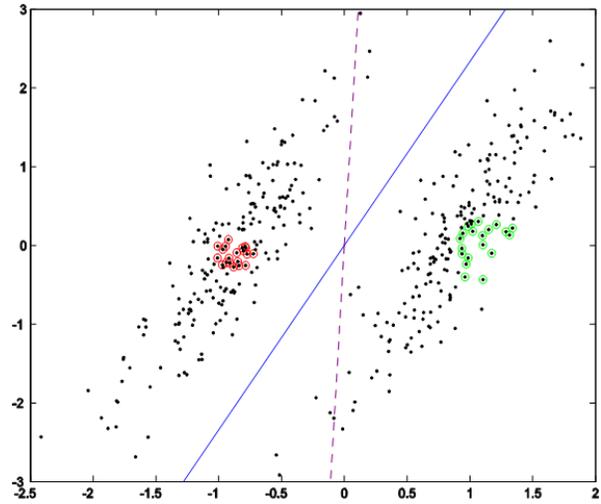


Figure 2: A comparison between supervised and semi-supervised classifiers for a two feature dataset. Labeled data from both classes (red and green circles) are shown, along with unlabeled data (black dots). The supervised classifier is trained on only the labeled data and the decision boundary is shown (dotted line). The semi-supervised classifier is trained on both the labeled and unlabeled data and the decision boundary (solid line) makes the two classes linearly separable.

applied supervised algorithms. Figure 3 depicts the advantage of a semi-supervised approach to classification over its supervised counterpart. A classifier trained purely on labeled data (depicted as red and green circles) is shown as a purple dashed line and generates classification errors. In contrast, a semi-supervised classifier trained on both labeled and unlabeled data will generate perfect classification (depicted by the blue line). Note that the context provided by the unlabeled data was crucial in improving the classification performance in this case, since the labeled data were not representative of the two class distributions. As the number of training samples increases, the supervised classifier should approximate the semi-supervised classifier. Semi-supervised formulation treats the dataset (labeled and unlabeled) as a set of connected nodes, where the affinity w_{ij} between any two feature vectors (nodes) \mathbf{f}_i and \mathbf{f}_j is defined in terms of a radial basis function [4]. Based on the above formulation, one can design a Markov transition matrix $\mathbf{A} = [a_{ij}]_{N \times N}$ that represents the probability of transitioning from node \mathbf{f}_i to \mathbf{f}_j . Assuming $\mathcal{L} \subseteq \{1, 2, \dots, N_L\}$ represents the set of labeled data indices, the likelihood functional can be written as

$$(\{y_i, i \in \mathcal{L}\} | \mathcal{N}(\mathbf{f}_i), \boldsymbol{\theta}) = \prod_{i \in \mathcal{L}} p(y_i | \mathcal{N}(\mathbf{f}_i), \boldsymbol{\theta}) = \prod_{i \in \mathcal{L}} \sum_{j=1}^{N_i} a_{ij} p(y_i | \mathbf{f}_j, \boldsymbol{\theta})$$

where $\mathcal{N}(\mathbf{f})$ defines the neighborhood of \mathbf{f} . Estimation of classifier parameters $\boldsymbol{\theta}$ can be achieved by maximizing the log-likelihood via an Expectation-Maximization algorithm [5]. To enforce sparseness of $\boldsymbol{\theta}$ (enforcing most of the components of the parameter vector $\boldsymbol{\theta}$ to be zero), one may impose a zero-mean Gaussian prior on $\boldsymbol{\theta}$. A zero-mean Gaussian prior with appropriate variance can strongly bias the algorithm in choosing parameter weights that are most likely very small (close to zero). The algorithm we have used for this semi-supervised learning is termed a parameterized neighborhood-based classifier (PNBC).

Non-myopic Active Learning (NMAL)

Given that available training data labels at the beginning of a demonstration are not available and that excavations must be performed to reveal training data labels, one may ask in which order anomalies should be excavated to maximally improve the performance of the classifier algorithm. One useful criterion is to use the confidence on the estimated identity of the anomalies that are yet to be excavated. Specifically, one may ask which unlabeled anomaly label would be most informative to improve classifier performance if the associated label could be made available. It has been shown [6] that this question can be answered in a quantitative information-theoretic manner.

For active label selection, posterior distribution of the classifier is approximated as a Gaussian distribution centered on the maximum *a posteriori* estimate. The uncertainty of the classifier is quantified in terms of the posterior precision matrix. The objective of NMAL is to choose a feature vector for labeling that maximizes the mutual information (I) between the classifier $\boldsymbol{\theta}$ and the new data point to be labeled. The mutual information can be quantified as the expected decrease of the entropy of $\boldsymbol{\theta}$ after new sample \mathbf{f}_{i^*} and its label y_{i^*} are observed.

$$I = \frac{1}{2} \log \frac{|H'|}{|H|} = \frac{1}{2} \log \{1 + p(y_{i^*} | \mathbf{f}_{i^*}, \boldsymbol{\theta}) \times [1 - p(y_{i^*} | \mathbf{f}_{i^*}, \boldsymbol{\theta})] \mathbf{f}_{i^*}^T H^{-1} \mathbf{f}_{i^*}\}$$

It is important to note that the mutual information I is large when $p(y_{i^*} | \mathbf{x}_{i^*}, \boldsymbol{\theta}) \approx 0.5$. Hence, the NMAL prefers to acquire labels on those unlabeled samples for which the current classifier is most confused or uncertain. In this fashion the classifier learns quickly by not excavating anomalies that reveal redundant information. The process continues as new labels are revealed until the expected information gain for the remaining anomalies is approximately uniformly low.

At that point the classifier is adequately trained and target inference on the remaining unlabeled anomalies can be reliably performed. By invoking the principle of submodularity in the algorithm optimization, the approach has been adapted to allow for the selection of multiple simultaneous labels at one time, making the technique operationally practical.

3. Performance Objectives

The Performance objectives of the demonstration are summarized in Table 1. Specific descriptions of each objective follow.

Table 1. Program Office Performance Objectives for Discrimination Analysis

Performance Objective	Metric	Data Required	Success Criteria	Results
Analysis and Classification Objectives				
Maximize correct classification of targets of interest	Number of targets-of-interest retained.	<ul style="list-style-type: none"> • Prioritized anomaly lists • Scoring reports from the IDA 	Approach correctly classifies all targets-of-interest	Retained 163 and 164 targets out of 170.
Maximize correct classification of non-UXO	Number of false alarms eliminated.	<ul style="list-style-type: none"> • Prioritized anomaly lists • Scoring reports from IDA 	Reduction of false alarms by > 30% while retaining all targets of interest	Reduced false alarms by 75% with all targets retained
Specification of no-dig threshold	Probability of correct classification and number of false alarms at demonstrator operating point.	<ul style="list-style-type: none"> • Demonstrator - specified threshold • Scoring reports from IDA 	Threshold specified by the demonstrator to achieve criteria above	Operating point: approx. 230 false alarms for both methods
Minimize number of anomalies that cannot be analyzed	Number of anomalies that must be classified as “Unable to Analyze.”	<ul style="list-style-type: none"> • Demonstrator target parameters 	Reliable target parameters can be estimated for > 98% of anomalies on each sensor’s detection list.	Approx. 55 (%2) of targets labeled “can’t analyze”

3.1. Maximize correct classification of targets of interest

A non-linear and a linear classifier were trained based on training labels requested from the program office. The objective was to predict all remaining UXO using the trained classifiers. This is measured by comparing the number of UXO captured from the dig list against the total number of UXO in the dataset. The necessary data are the dig lists and the scoring reports from the IDA. Some UXO were missed, and so the performance objective was evaluated in the context of how many additional digs would have been necessary to actually capture all the UXO.

3.2. Maximize correct classification of non-UXO

For both classifiers, a secondary objective is to capture all the UXO while keeping much of the clutter in the ground. Success was measured by keeping at least 70% of the clutter in the ground.

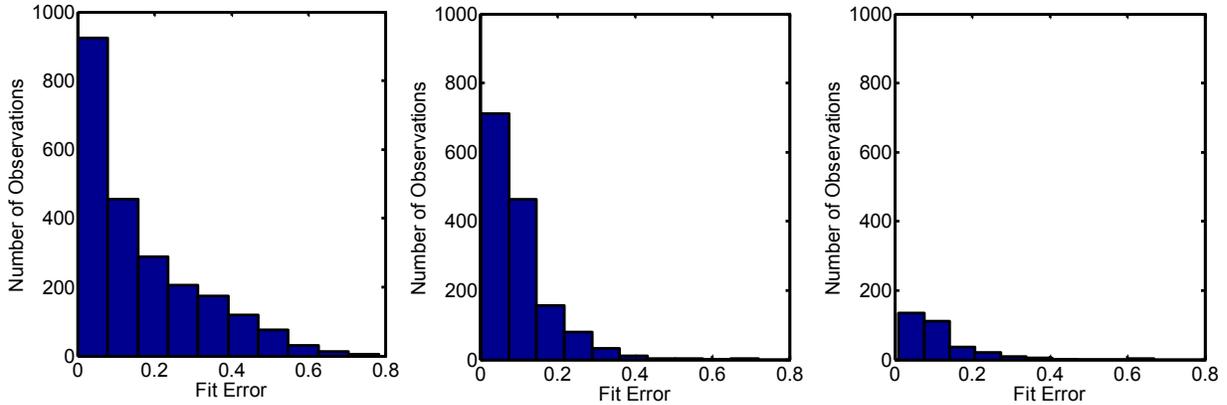


Figure 5. Histograms of fit errors for one (left), two (middle) and three (right) dipole inversion results for the Camp Butner TEMTADS sensor.

Since, some UXO were left in the ground given the no-dig threshold, the number of false alarms was smaller than it should have been. This objective was re-evaluated in terms of how many false alarms would have been necessary were the digging thresholds set to capture all the UXO.

3.3. Specification of no-dig threshold

The objective was to give a reasoned operating point for splitting the dataset into anomalies that should be dug and those that should not be dug. The decision for this objective influenced the performance of the values in the first two objectives. The decision to stop digging was based on the separation between the posteriors predicted probabilities of the anomalies not used for training. The selected operating point based on this criterion left UXO in the ground.

3.4. Minimize the number of anomalies that cannot be analyzed

The objective was to have a minimal number of anomalies where the dipole inversion model gave poor results. This is a function of the data quality, something that was not controlled in this study, and a function of the efficacy of the inversion model. The decision to place anomalies in the ‘can’t analyze’ category was based on the residual error of the least-squares model used for the dipole inversion. Anomalies with high residual error were removed. Success in this objective was defined as a creating effective parameterizations for >98% of the anomalies. This objective was achieved.

4. Site Description

All raw sensor data were provided to SIG directly. So there were no in-field components to the SIG discrimination.

5. Test Design

All raw sensor data were provided to SIG directly. So there were no in-field components to the SIG discrimination.

6. Data Analysis and Products

6.1. Parameter Estimates

SIG performed feature extraction and discrimination for the Time-domain Electromagnetic Multi-sensor Tower Array Detection System (TEMTADS) sensor at Camp Butner. There were 2291 total flags that required a

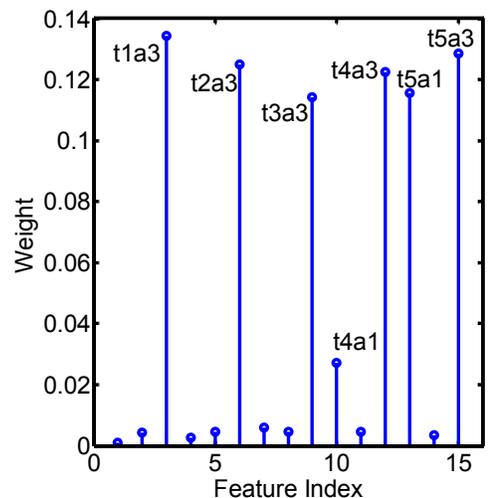


Figure 4. BENet weights for the polarizabilities of the TEMTADS dataset. Selected feature names are shown (tx is time gate x and ax is axis x).

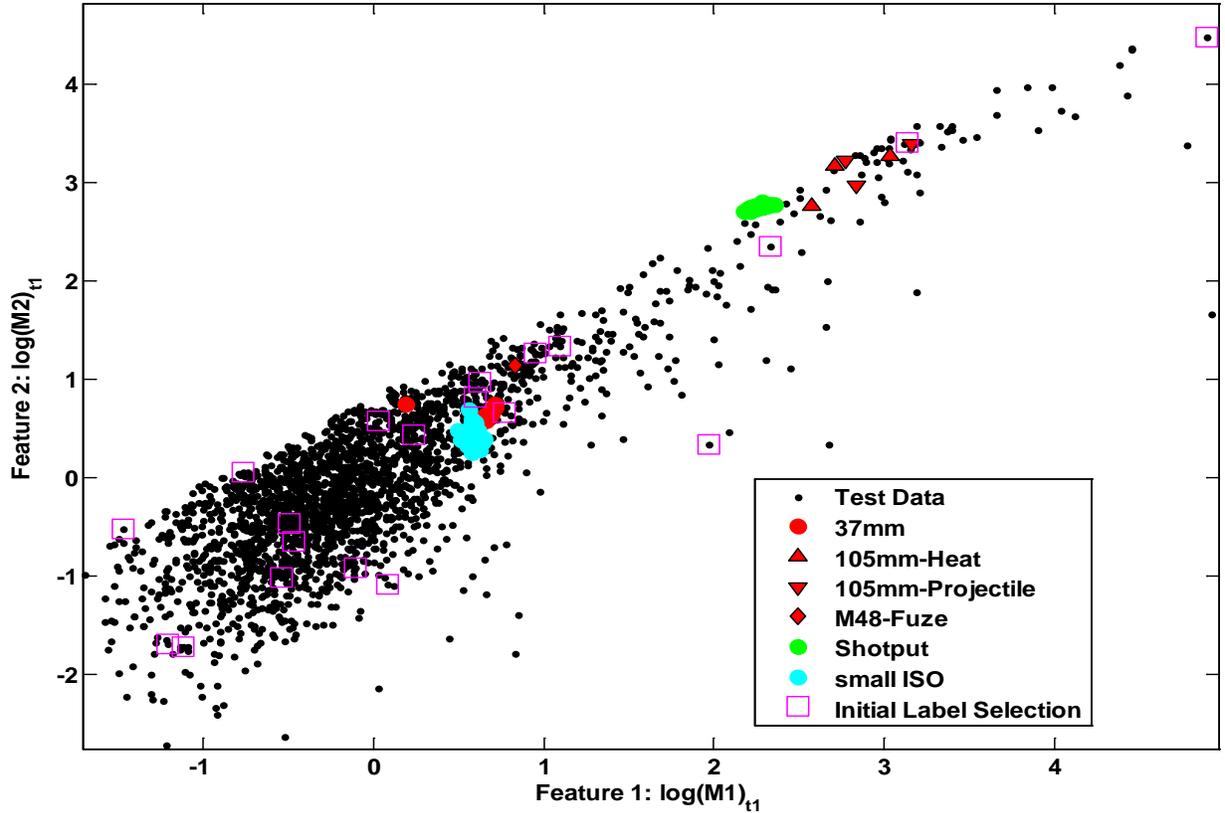


Figure 6. Initial basis selection of Butner TMTADS data along two primary features. The two features are the log polarizability of the primary secondary axes at the first time-gate. Test pit UXO are also shown. M1 is the primary axis. M2 is the secondary axis. T1 is the first time gate.

dig/no-dig decision. The goal of feature extraction is to invert the responses at the receivers and estimate the polarizabilities of the measured anomaly along its three axes. A non-linear least square approach was used to find the solutions to this dipole inversion. The fit-error of the non-linear least square model gives information about the degree to which the dipole model is appropriate for the anomaly (Figure 4). Generally speaking, the single anomaly model is appropriate when fit errors are less than 0.05. If the one anomaly model fit errors were large, then a two-anomaly model was created. Each anomaly is a triaxial dipole. Fit errors for the two-anomaly model are always less than the one anomaly model because the number of parameters is larger. Of the 2291 observations, SIG created a two-dipole inversion for 1464. If the fit errors for the two-anomaly model were also large, then a three-anomaly model was created. SIG created three-anomaly dipole inversions for 320 of the observations.

6.2. Feature Selection

Polarizabilities were estimated for 5 time gates along each of the three object axes for a total of 15 features. An initial subset of these features was chosen based on prior knowledge obtained from the SLO and Sibert demonstrations. This set included polarizabilities from early and late time gates for axis 1 and 3. This set of features was changed after the initial labels were received. This second round of feature selection was performed using BENet. The selected features were *similar* to those shown in Figure 5. This figure highlights the fact that large responses along the smallest axis were indicative of UXO, as was the late time response of the

primary axis. Having selected this subset of features, additional discrimination and active learning were performed.

6.3. Training and Classification

The initial set of 20 requested observations were selected to cover the breadth of feature responses in the training dataset (Figure 6). After initial basis selection, a series of additional training data points were acquired via NMAL. A total of 65 additional training labels were acquired (Figure 7). Among those 16 were UXO (6 – 37mm, 5 – 105mm, 5 – M48). After active learning was complete, two different classifiers were trained. The first was a linear PNBC classifier where the input features were the original features selected by BENet. The second was a nonlinear PNBC classifier. For this classifier a radial basis kernel function was applied to the original features. The input features to the classifier was, then, a $N_1 \times N_2$ matrix where N_1 was all of the data and N_2 were the labeled training points. The values in each row of the kernel were weights to all the labeled training points. So, for a given observation, a high weight would be given to a labeled training point that was close (in feature space) to the focal point, and a low weight would be given to a labeled point that was far. The nonlinear PNBC classifier identified two clusters of high probability UXO (Figure 8). The first was associated with the 105mm munitions; the second was associated with the smaller projectiles (M48-Fuzes and 37mm projectiles). 37mm projectiles were the most difficult to discriminate and there were 4, in particular, that proved particularly difficult.

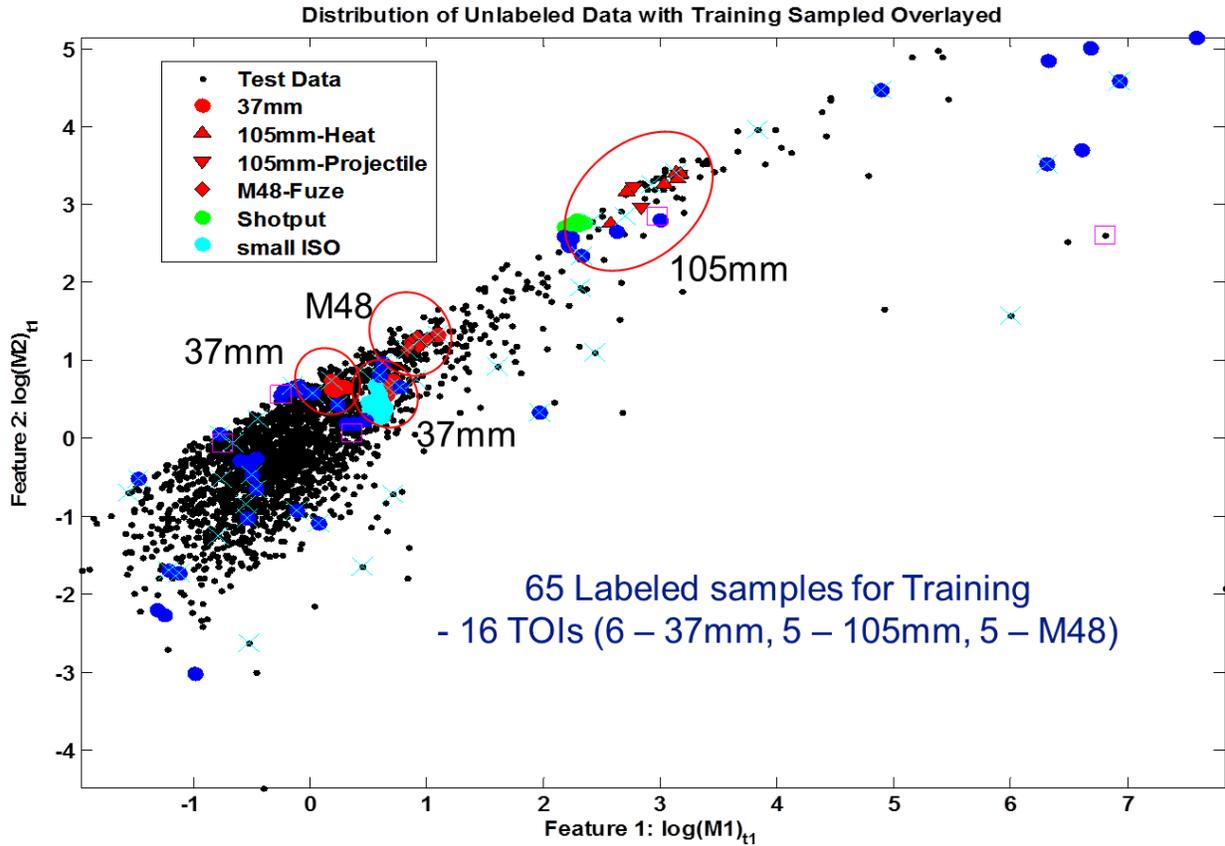


Figure 7. Labeled data at the end of training. UXO (red) and clutter (blue) training labels are shown along with the last round of actively learned labels (pink squares). Obvious clusters of munitions are highlighted.

6.4. Excavation

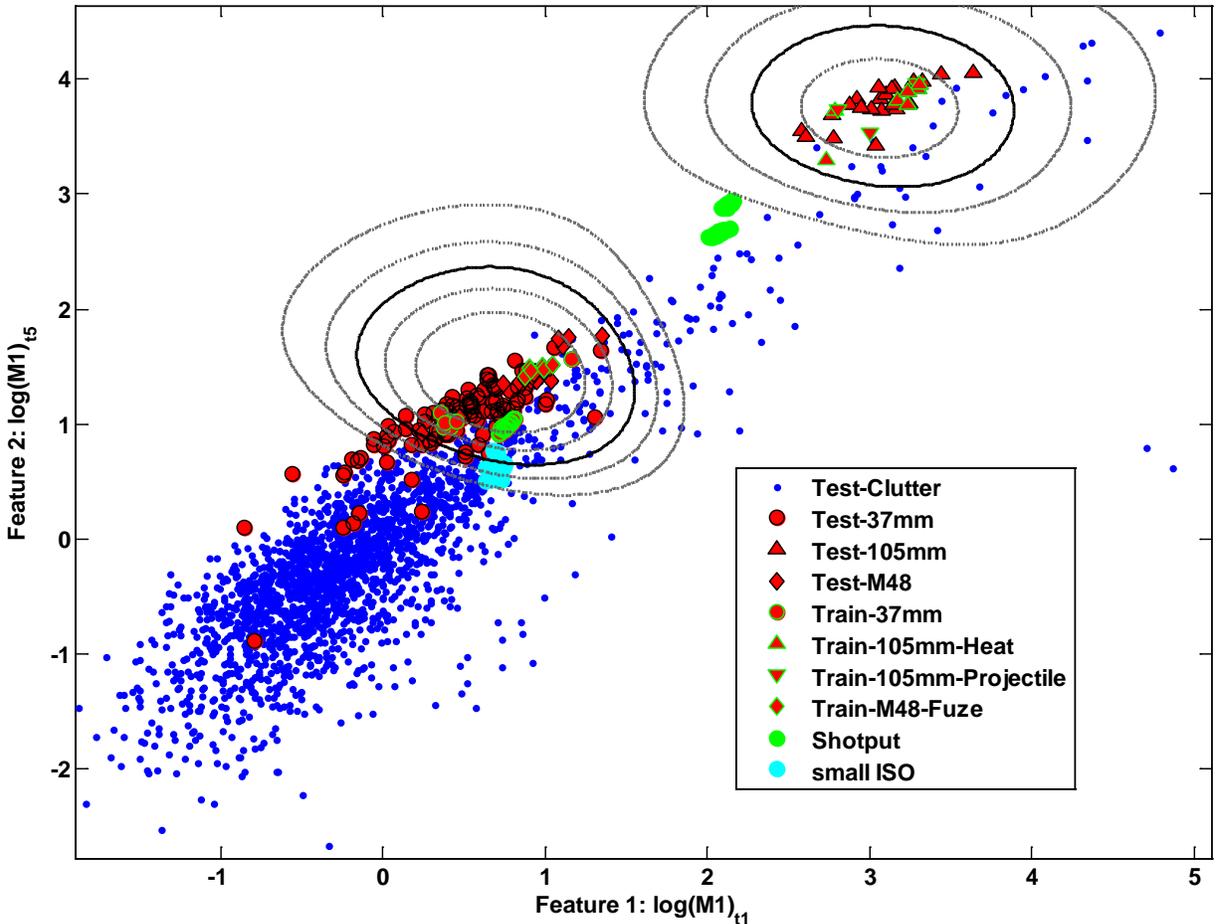


Figure 8. Non-linear PNBC classification boundary along with test and training data. Contour intervals are 0.1 posterior predicted probability of being UXO.

Two dig lists were submitted to the program office, one corresponding to the linear classifier and one corresponding to the non-linear classifier. Both used the same training data, and both included approximately 30 anomalies whose features were too difficult to extract. These were labeled, ‘can’t analyze’, and were marked for digging. Initial dig lists were submitted and the program office returned partial receiver operating characteristic (ROC) curves for the linear and non-linear classifiers. Both methods revealed approximately 130 UXO with less than 10 unnecessary digs, and another 30 UXO with approximately 60 extra digs. We then retrained the models with the additional labels of the dug anomalies from the first list. Then, a second set of lists were submitted that requested a few more digs for each model. The final ROC curves for these classifiers can be seen in Figure 9 and Figure 10.

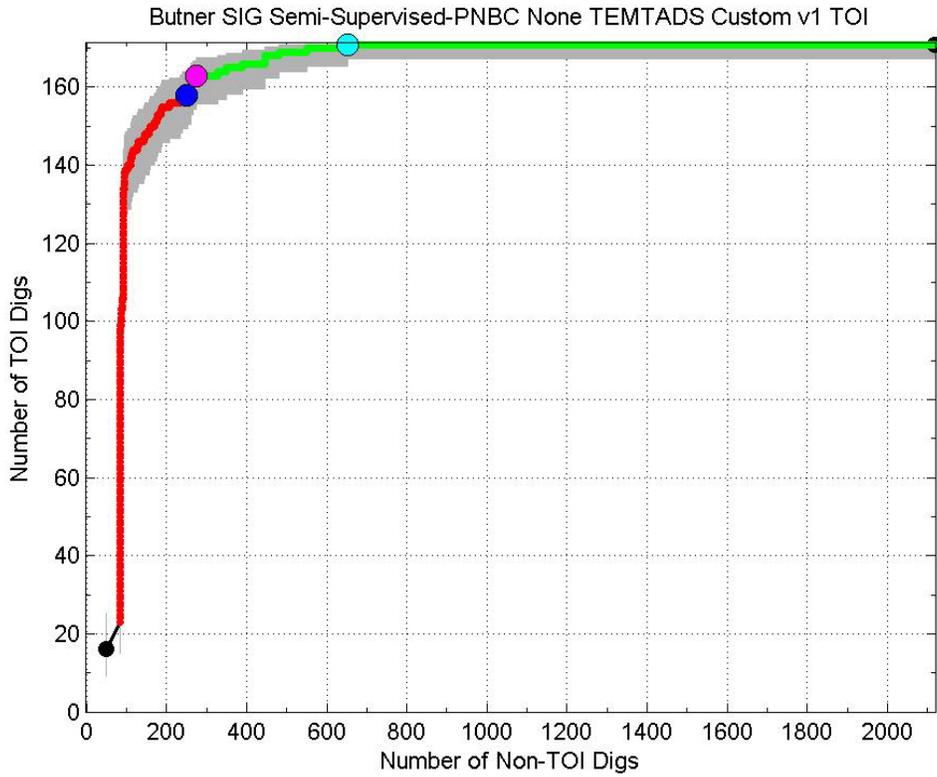


Figure 9. ROC curve for the linear classification of the Camp Butner TEMTADS data

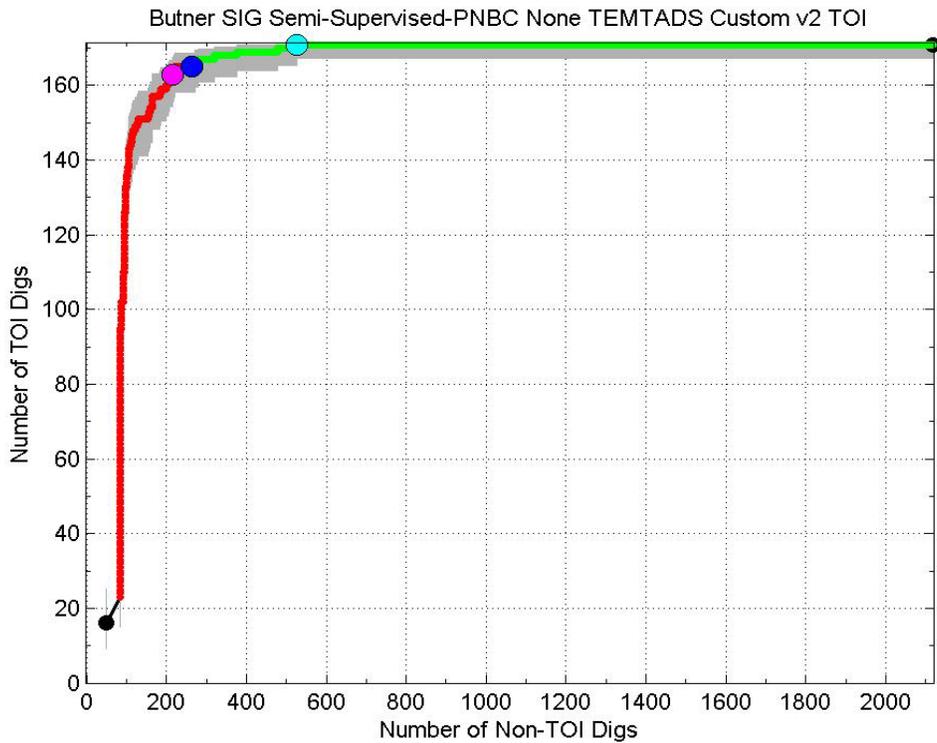


Figure 10. ROC curve for the non-linear classification of the Camp Butner TEMTADS data

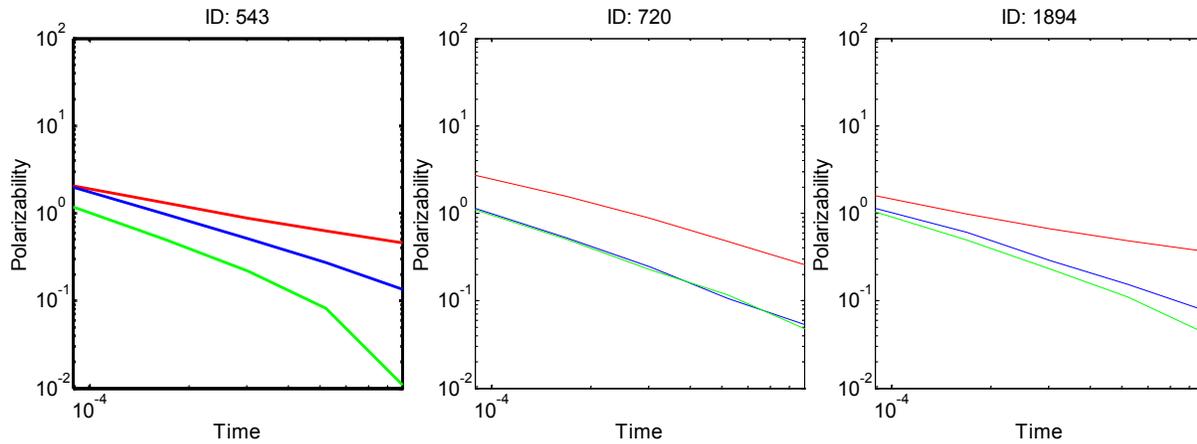


Figure 11. Polarizabilities for anomalies that were difficult to discriminate for both the linear and non-linear models. IDs for each anomaly are shown in the title of the plots.

7. Performance Assessment

Based on the stopping point (blue dots in Figure 9 and Figure 10), the number of UXO revealed by the linear and non-linear approaches were roughly the same: 163 for the linear approach and 164 for the non-linear approach. Also, the number of unnecessary digs was roughly the same, 230. Both the linear and non-linear approaches left UXO in the ground. The non-linear model left fewer UXO than the linear model. The bulk of the missed anomalies were 37mm mortars. This is not surprising given the clustering of 37mm mortars with clutter in feature space (Figure 8). If digging continued until the final UXO was dug, then the non-linear classifier would have outperformed the linear classifier. The non-linear classifier would have had 560 clutter, while the linear classifier would have dug 625.

Three of the missed anomalies were shared between the linear and non-linear classifier: IDs 543, 720, and 1894. Their polarizabilities are plotted in Figure 11. All three anomalies share a distinguishing characteristic. Their overall response is lower than a typically UXO. The maximum polarizabilities for most munitions are on the order 10-100 in our inversion model. These anomalies all have maximum responses slightly higher than 1. Anomaly 543 would be difficult to discriminate no matter what features or model was used because it does not exhibit a standard UXO response, namely, it has low overall magnitude and the transverse axes are not symmetric. Anomalies 720 and 1894 could be discriminated on the basis of symmetry. The symmetry feature was not included in our set of discriminating features.

7.1. Maximize correct classification of targets of interest

The linear and non-linear classifications retained 163 and 164 UXO, respectively. This was the only performance object that was missed. It was missed due to a poorly chosen no-dig threshold. Were the stopping point moved to 625 false alarms, both methods would have met all of the performance objectives.

7.2. Maximize correct classification of non-UXO

If the dig-threshold were chosen correctly, then the reduction of false alarms would have been 75% for the linear classification and 85% for the non-linear classification. The no-dig threshold set too early, however. So, both classifications reduced the number of false alarms by 90%, but left UXO in the ground.

7.3. Specification of no-dig threshold

The operating point for the no-dig threshold was set at approximately 230 false alarms for both the linear and non-linear classifiers. The decision to stop digging was based on the separation between the posteriors predicted probabilities of the anomalies not used for training.

7.4. Minimize the number of anomalies that cannot be analyzed

98% of the anomalies had target parameters extracted effectively. 2% had large fit errors for the non-linear least squares model used for dipole inversion, were labeled “can’t analyze”, and marked for digging.

8. Cost Assessment

This section should provide sufficient cost information such that a professional involved in the field could reasonably estimate costs for implementation at a given site. In addition, this section should provide a discussion of the cost benefit of the technology. The following subsections with detailed discussions and examples should be provided.

8.1. Cost Model

The cost model is summarized in Table 2. The total cost per anomaly is \$21.0. Each cost element is described in subsections below.

Cost Element	Data Tracked During Demonstration	Estimated Costs
Feature Inversion	Unit: \$ per anomaly <ul style="list-style-type: none"> • Time required • Personnel required • Number of sensors • Number of classifier techniques 	11.8
Classifier Training/Testing	Unit: \$ per anomaly <ul style="list-style-type: none"> • Time required • Personnel required • Number of sensors • Number of classifier techniques 	5.5
Reporting	Unit: \$ per anomaly <ul style="list-style-type: none"> • Time required • Personnel required • Number of sensors • Number of classifier techniques 	3.7

Features Inversion

Feature inversion includes any denoising and data preprocessing. The input data product here are the raw sensor data. The output are the polarizabilities from the dipole model. Additional quality checks are performed at this stage. Costs would scale less than linearly with number of anomalies, because the time required for quality control is roughly the same regardless of the number of anomalies.

Classifier Training/Testing

Classifier training and testing encompasses all the data analysis required to move from anomaly polarizabilities to a final dig list. This includes requesting training data from the program office, feature selection, active learning, and quality assurance. Costs scale less than linearly with number of anomalies, because the percentage of training data required should decrease as the total number of anomalies increases.

Reporting

This includes documentation of all feature inversion, classifier training/testing, and classifier performances. The cost should scale linearly with the sensors and classification techniques used.

8.2. Cost Drivers

The purpose of the SIG Isolate discrimination process is to decrease the cost per anomaly and to do so in a manner that scales well with production level discrimination. As the requirement for expert intervention and interpretation decreases, the scaling of the cost per anomaly should improve.

8.3. Cost Benefit

While the SIG Isolate process is not completely automated at this point, increasing automation drives the cost per anomaly toward becoming simply a function of computing time required and quality assurance checks. Since analyst time is the greatest cost in the discrimination process, automation provides excellent cost benefit for discrimination.

9. Implementation Issues

The software for the current SIG Isolate technology is based on MATLAB® and is not freely available. While the software is currently used by the experts who wrote the system, transitioning to minimally trained users is a goal of the software development. Future demonstrations will be used to mature this software.

10. References

[1] Y. Zhang, X. Liao, and L. Carin, "Detection of buried targets via active selection of labeled data: application to sensing subsurface UXO," IEEE Transactions on Geoscience and Remote Sensing, vol. 42, pp. 2535–2543, November 2004.

[2] Y. Zhang, H. Y. L. Collins, C. Baum, and L. Carin, "Sensing of unexploded ordnance with magnetometer and induction data: Theory and signal processing," IEEE Transactions on Geoscience and Remote Sensing, vol. 41, pp. 1005–1015, May 2003.

[3] Y. Zhang, X. Liao, and L. Carin, "Detection of buried targets via active selection of labeled data: application to sensing subsurface UXO," IEEE Transactions on Geoscience and Remote Sensing, vol. 42, pp. 2535–2543, November 2004.

[4] N. Cristianini and J. Shawe-Taylor, An Introduction to Support Vector Machines and other kernel-based learning methods. Cambridge University Press, 2000.

[5] X. Liao and L. Carin, "Migratory Logistic Regression for Learning Concept Drift Between Two Data Sets With Application to UXO Sensing," IEEE Transactions on Geoscience and Remote Sensing, vol. 46, no. 12, December 2008.

[6] D. MacKay, "Information-based objective functions for active data selection," Neural Computation, vol. 4, pp. 589–603, 1992.

11. Appendices

11.1. Appendix A: Points of Contact

POINT OF CONTACT Name	ORGANIZATION Name Address	Phone Fax E-mail	Role in Project
Lawrence Carin	Signal Innovations Group, Inc. 4721 Emperor Blvd., Suite 330 Durham, NC 27703	919 660-5270 919-323-4811 lcarin@ece.duke.edu	Principal Investigator
Levi Kennedy	Signal Innovations Group, Inc. 4721 Emperor Blvd., Suite 330 Durham, NC 27703	919-323-3456 919-287-2578 lkennedy@siginnovations.com	Project Management

Todd Jobe	Signal Innovations Group, Inc. 4721 Emperor Blvd., Suite 330 Durham, NC 27703	919-323-4811 919-287-2578 tjobe@siginnovations.com	Engineer
Xianyang Zhu	Signal Innovations Group, Inc. 4721 Emperor Blvd., Suite 330 Durham, NC 27703	919-323-4811 919-287-2578 xianyang@siginnovations.com	Engineer