FINAL REPORT

Optimal Sensor Management and Signal Processing for New EMI Systems

SERDP Project MR-1708

MAY 2012

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

This document has been cleared for public release

SERDP
DOD ▪ EPA ▪ DOE

This report was prepared under contract to the Department of Defense Strategic Environmental Research and Development Program (SERDP). The publication of this report does not indicate endorsement by the Department of Defense, nor should the contents be construed as reflecting the official policy or position of the Department of Defense. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the Department of Defense.

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)*
27-04-2012

2. REPORT TYPE
Follow-On Report

3. DATES COVERED *(From - To)*
07/2010-05/2012

4. TITLE AND SUBTITLE
2009 ESTCP UXO Discrimination Study San Luis Obispo, CA:

5a. CONTRACT NUMBER
MR-1708

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

Objective: A follow-on study was performed for the Signal Innovations Group, Inc. (SIG) unexploded ordnance (UXO) demonstration at Camp San Luis Obispo, CA. There were four major objectives. 1) Analyze the previous demonstration performance and validate the technical approach (supervised and semi-supervised classifiers) used for discrimination. 2) Explore new active learning procedures for acquiring training labels. 3) Evaluate the adequacy of training data used in the demonstration. And, 4) evaluate new enhancements to the target/sensor model for acquiring object polarizabilities from the sensor data.

Results: The SIG multi-anomaly dipole model is effective at producing features that are easily recognized as being UXO where a single anomaly model would not be effective. The semi-supervised classifier will outperform supervised classifier when a difficult UXO is connected to other UXO by neighboring observations (in feature space). Both the semi-supervised and supervised classifiers are robust to changes in their user-defined parameters. Non-myopic active learning is effective, in general at capturing difficult UXO, and enhancing the quality of training data for feature selection. When active learning is used in conjunction with feature selection, appropriate features can be selected from few training data.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**	SAR	26	Levi Kennedy
U	U	U			**19b. TELEPHONE NUMBER** *(include area code)* 919-323-3456

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

Table of Contents

List of Figures ii
List of Acronyms iii
Keywords iii
Abstract 1
Objective 2
Background 2
Materials and Methods 3
Performance analysis and technical approach validation 3
Active learning / Training Data Adqueacy 3
Sensor/Target Model Enhancements Study 4
Results and Discussion 5
SLO Discrimination Study Performance Analysis and Technical Approach Validation (TEMTADS, MetalMapper, EM61 Array, BUD) 5
Active Learning Follow-on Analysis 9
Training Data Adequacy Study 14
Sensor/Target Model EnhancemeNts Study 15
Conclusions and Implications for Future Research/Implementation 20
Literature Cited 20

List of Figures

Figure 1. The supervised (left) and semi-supervised (right) classifier performances on MetalMapper data for SLO. The features were provided by Snyder Geoscience, which were further pruned by the SIG feature selection algorithm. 6

Figure 2. Sensitivity analysis of the user-defined parameters for SBC (left) and PNBC (right). The colored values are the AUCs of classifiers trained using a single combination of parameters. The parameter settings for the performance evaluation are highlighted..... 7

Figure 3. EM61 responses for 4 anomalies that were challenging to discriminate at SLO. The top two anomalies are shown with their pictures. 8

Figure 4. Feature weights for all the even rounds of active learning. Selected features that were used for classification are highlighed (red circles). Feature names for selected features are also shown. The abbreviations are: t*x* – time gate *x*, a*x* – axis *x*, r01 – exponential model intercept, and r03 – exponential model decay rate. 9

Figure 5. Fisher information gain during NMAL for the MetalMapper sensor data. 10

Figure 6. Histogram of predicted probabilities at the end of training PNBC for the MetalMapper sensor using NMAL. 10

Figure 7. ROC curve for the PNBC classifier at the end of NMAL. Data are from the MetalMapper sensor. The training data are included in the plot (dashed line) along with the last UXO acquired (red). 12

Figure 8. ROC curves during different rounds of NMAL for the PNBC classifier. 12

Figure 9. The number of UXO as a function of their posterior probability for 3 different sampling methods with a sample size of 120. Note: AL in the legend is non-myopic. 13

Figure 10. BENet feature weights used for feature selection during the even rounds of NMAL. The marginal posterior distributions of feature weights are shown with +-1 standard deviation (thick blue segments) and a 95% credible interval (thin blue segments). Selected features are highlighted in red. The bottom right plot show the feature weights for a BENet model where all labels were used for training. 14

Figure 11. KL divergence from the posteriors of BENet feature weights to the BENet model fitted using all labels. Lower values are more similar to the model using all the labels. 15

Figure 12 Polarizabilities of a one anomaly model (left) and two anomaly model (right) as a function of time for a 37mm mortar. 16

Figure 13 Eigenvalues as a function of time for a 37 mm mortar. 17

Figure 14 Eigenvalues as a function of time for two-anomaly data (one is 105mm mortar) 18

Figure 15 Eigenvalues as a function of time for two-anomaly data (one is 37mm mortar) 18

Figure 16 Comparison of extracted parameters for anomaly 847. (a) Skip's result; (b) SIG's result. 19

Figure 17 Comparison of extracted parameters for anomaly 956. (a) Skip's result; (b) SIG's result..... 19
Figure 18 Comparison of extracted parameters for anomaly 895. (a) Skip's result; (b) SIG's result..... 19

List of Acronyms

AUC	Area under the ROC curve
BENet	Bayesian Elastic Net
NMAL	Non-myopic active learning
PNBC	Parametric Neighborhood-Based Classifier
SBC	Sparse Bayesian Classifier
ROC	Receiver Operating Characteristic Curve
SIG	Signal Innovations Group, Inc.
SLO	San Luis Obispo
UXO	Unexploded Ordnance

Keywords

UXO, Bayesian classifier, active learning

Abstract

Objective: A follow-on study was performed for the Signal Innovations Group, Inc. (SIG) unexploded ordnance (UXO) demonstration at Camp San Luis Obispo, CA. There were four major objectives. 1) Analyze the previous demonstration performance and validate the technical approach (supervised and semi-supervised classifiers) used for discrimination. 2) Explore new active learning procedures for acquiring training labels. 3) Evaluate the adequacy of training data used in the demonstration. And, 4) evaluate new enhancements to the target/sensor model for acquiring object polarizabilities from the sensor data.

Technical Approach: Specific core technologies were examined and validated during the follow-on analysis. These technologies fall broadly into the four analysis categories: the sensor/target model, feature selection, classification, and active label selection. SIG developed a sensor/target model using a multi-anomaly dipole inversion. 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 supervised, sparse Bayesian classifier, and the other a semi-supervised, parametric Bayesian classier.

Results: The SIG multi-anomaly dipole model is effective at producing features that are easily recognized as being UXO where a single anomaly model would not be effective. The semi-supervised classifier will outperform supervised classifier when a difficult UXO is connected to other UXO by neighboring observations (in feature space). Both the semi-supervised and supervised classifiers are robust to changes in their user-defined parameters. Non-myopic active learning is effective, in general at capturing difficult UXO, and enhancing the quality of training data for feature selection. When active learning is used in conjunction with feature selection, appropriate features can be selected from few training data.

Benefits: 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.

Objective

Signal Innovations Group, Inc. (SIG) participated in the Sibert and the SLO UXO discrimination studies. SIG has applied and iteratively adapted state-of-the-art UXO discrimination technology developed under funding provided over many years by the SERDP and ESTCP program offices. Following the Sibert discrimination study, and once all the ground truth labels were made available, a follow-on analysis provided the opportunity to gain insights and adapt the applied technology, and greatly aided the performance of the results in the SLO discrimination study.

There were four major objectives specified in the statement of work for the follow-on analysis:

- Performance analysis and validation of the technical approach for the TEMTADS, MetalMapper, EM61 array, and BUD sensors
- Additional exploration of the active learning procedures
- Evaluation of the adequacy of the training data used in the study
- Evaluate enhancements to the target/sensor model

Each of these objectives was split into specific subtasks. The subtasks are treated individually in the Results and Discussion.

Background

The detection and remediation of unexploded ordnance (UXO) remains a high priority for the Department of Defense. According to 2003 Defense Science Board (DSB) report [1], the UXO cleanup problem is a very large-scale undertaking involving multi-million acres of land in more than fourteen hundred test sites scattered over the continental United States. It is estimated that only 10% of this area is occupied by UXO. The report estimates the cost of cleanup in the order of tens of billions in dollars. This cost consists of detection and classification of surface, buried, and partially buried UXO, followed by expensive excavation of these anomalies. The problem is aggravated by the fact that most of the UXO detection systems produce a large number of false alarms. This results in digging many holes where no UXO is present.

Significant progress has been made in the development of UXO detection and classification technologies over the past several years. However, most of the testing and validation of this technology (both in hardware designs and associated software algorithms) have been primarily limited to test sites with only limited application at live sites. The SLO demonstration was a 'live site' application. For the SLO demonstration, SIG performed discrimination on the TEMTADS, MetalMapper, EM61 array, and BUD sensors.

The SIG classification algorithms performed well in the sense that for the classifiers considered and thresholds considered, most of the UXO items were detected, and generally a substantial number of non-UXO were left unexcavated. This was also true of the active-learning algorithms, which employed no prior labeled data, and ultimately employed substantially less labeled data than initially provided by the ESTCP office. SIG developed an improved forward model for the SLO demonstration which incorporates the actual shape of the transmitter/receiver loops and redesigned the inversion algorithm to make it more robust. The performance of the MetalMapper

data were particularly impressive. In general, multi-axis and multi-coil sensor systems significantly outperformed the single coil electromagnetic and magnetometer sensors.

Materials and Methods

The technical methods used for the follow-on study are summarized by study objective.

Performance analysis and technical approach validation

SIG performed both supervised and semi-supervised classification at SLO. The supervised classifier was a sparse Bayesian classifier (SBC, [2]) and the semi-supervised classifier was a parametric neighborhood-based classifier (PNBC, [3]). These two approaches share some characteristics. First, they both incorporate user-defined prior distributions to prevent the trained model from over-fitting the training data. This purpose of this attribute is to improve prediction on test data. Second, both SBC and PNBC represent predictions as posterior probabilities rather than decisions. This distinguishes them from other common machine learning classifiers such as the support vector machine whose predictions are decisions rather than probabilities.

PNBC is semi-supervised because it attempts to improve prediction by leveraging information about the autocorrelation of unlabeled samples in feature space. This contrasts with supervised classification, where the only information that influences prediction is the responses of labeled data. PNBC and SBC will give identical results when the influence of unlabeled observations is removed from PNBC. Thus, SBC is a subset of PNBC. This highlights one key point of interpretation. When many labeled observations are available for training, the need for information about the configuration of unlabeled samples is reduced. The result is that, given a large training set, PNBC should be tuned to perform more like SBC. Conversely, the benefits of semi-supervised classification are most apparent when sample sizes are small. These qualities prove important when comparing the performance of PNBC and SBC.

Active learning / Training Data Adequacy

During the performance assessment active learning was done only for the TEMTADS and EM61 data. The actively learned training data yielded a better performing classifier with fewer training data than a random sample of observations. This was highlighted in the final report. SIG has abandoned the active learning technique used in those cases, however. This myopic approach required that one label was selected at a time, and then the classifier needed to be retrained. It is not practical to acquire a single label of training data, dig it in the field, and then retrain the model. Instead SIG has adopted a non-myopic active learning approach (NMAL) that allows for a batch of new labels to be requested at once.

NMAL occurs in two phases. The first phase is called basis selection. The algorithm acquires a broad representation of feature space in the absence of a classifier. In this phase, outliers are typically identified depending on how many bases are chosen. SIG developed an algorithm for selecting an optimal number of bases dependent on the natural clustering of observations in feature space and the desired representation of these clusters in the dataset. The features are first transformed into an orthogonal set of principal components. The amount of data variability encompassed by each component determines its importance. For the MetalMapper sensor data

most of the data variability was captured using 4 components. After transforming the data into these 4 component axes, an unsupervised clustering was performed on the data using a Gaussian Mixture Model. The segmentation algorithm for this GMM selected not only the cluster membership but also the number of clusters. For the MetalMapper data, this resulted in 5 clusters. SIG set the number of initial bases to be 20. This yielded one observation from each of these natural clusters across each important dimension of the dataset. The results of such basis selection should not only be representative of the breadth of responses in the dataset, but should also capture outlier responses that do not cluster with other observations in feature space.

The second phase of NMAL is done with a classifier in place. In this phase NMAL seeks to acquire new labels that simultaneous decrease uncertainty in under-sampled regions of the feature space and refine the existing classifier boundary. New labels are acquired until additional labels have little new information for the classifier. SIG performed both phases of NMAL on the MetalMapper sensor data at SLO. A total of 100 training points were acquired.

The features selected for a given classifier can and should change during active learning as the amount of training data increases. There are two reasons for this. First, the quantity of features could change to accommodate increasing model complexity. As the number of training points increase, the number of features in the model may increase as the data suggests with little loss in generality assuming a sparse classification method like SBC or PNBC is used. Second, when training data are few, there may be spurious correlations with features that are not, in general, discriminative features. As the number of training data increases during active learning, these spurious correlations should decrease and better, more general features should be selected. SIG examined these two processes are examined in the context of non-myopic active learning (NMAL).

The technique used for feature selection is the Bayesian elastic net (BENet). BENet employs a sparseness prior equivalent to a convex combination of $L1$-norm and $L2$-norm penalties in a least squares optimization formulation [4], [5]. 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.

Sensor/Target Model Enhancements Study

Spatial overlap between UXO and clutter is not uncommon. For those cases, a single dipole model [6] will fail and the extracted features will ‘UXO-like’. SIG has developed an efficient multi-anomaly dipole model to extract better features. The multi-anomaly model is a straightforward extension of the single anomaly model. It assumes that the received magnetic field at each receiver is simply the sum of scattered magnetic fields from all possible anomalies. The scattered field from each single anomaly is estimated using the well-known dipole model. The received magnetic field at one specific receiver can be expressed as follows

$$\vec{H}(\vec{r}) = \frac{1}{4\pi\mu_0}\sum_{i=1}^{N}\left(3\hat{R}_i\hat{R}_i - \bar{\bar{I}}\right)\cdot\vec{m}_i$$

where $\vec{r}$ is the receiver position vector, $\hat{R}_i$ is the unit vector along $\vec{R}_i = \vec{r} - \vec{r}_i$, and $\vec{r}_i$ is the *i*th anomaly's position vector. $\bar{\bar{I}}$ is the identity dyad, and $\vec{m}_i$ is the dipole moment associated with the *i*th anomaly, which is determined by the incident magnetic field and the polarizability tensor of that anomaly. The coupling among all anomalies is not taken into account. Numerical experiments have shown that ignoring coupling works fine for most cases.

To extract the parameters efficiently, the unknowns are grouped into two parts associated with the anomaly positions and polarizability parameters, respectively. An iterative approach is employed to solve the problem. The anomaly positions are assumed to be known in each iteration. The polarizability parameters of all anomalies are estimated directly according to the raw data without optimization. Once polarizability parameters are known, the anomaly positions are updated via global optimization. The above iterative process repeats until the results converged.

Results and Discussion

SLO Discrimination Study Performance Analysis and Technical Approach Validation (TEMTADS, MetalMapper, EM61 Array, BUD)

Comparison of the results from both the supervised and semi-supervised classifier approaches

There was little difference in the performance of SBC and PNBC for most sensor datasets at SLO. In general, PNBC had dug UXO with more unnecessary digs, initially. The fact that SBC typically dug 60-80% of the UXO with less than 50 unnecessary digs suggests that knowledge about the larger structure of the labeled and unlabeled data available with PNBC was important only in select cases. This was especially true when UXO and clutter were not highly separable. In these cases adding information about the structure of unlabeled data ends up correlating more unlabeled clutter to labeled UXO. As a result the unlabeled clutter is assigned a higher probability of being UXO than if these correlations were not known by the classifier.

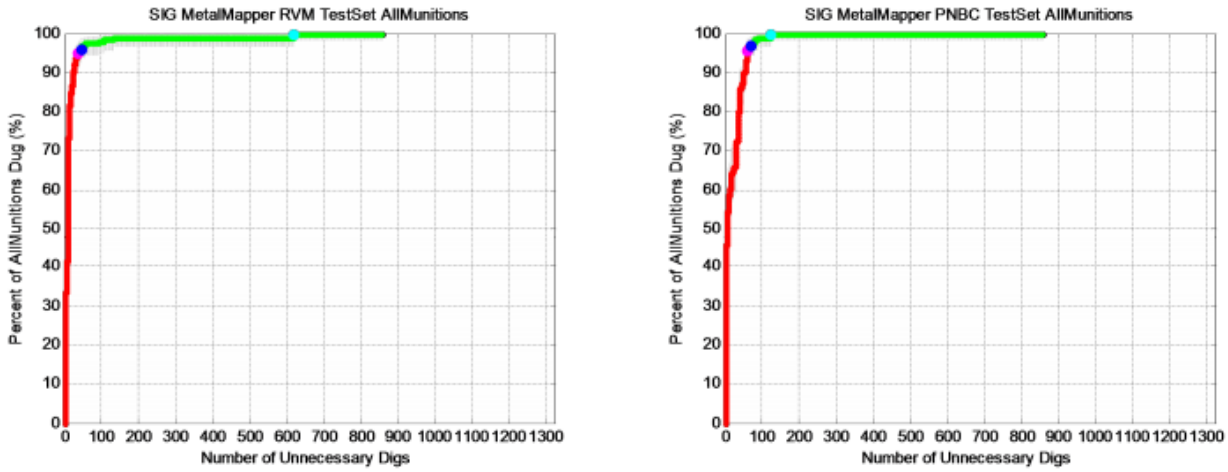

Figure 1. The supervised (left) and semi-supervised (right) classifier performances on MetalMapper data for SLO. The features were provided by Snyder Geoscience, which were further pruned by the SIG feature selection algorithm.

The situation where PNBC outperformed SBC was in the total number of clutter dug to capture the last UXO. This, of course, is the most important measure of classifier performance, as the cost for leaving UXO in the ground is treated as infinite. There were at least four challenging UXO, and the ability of a classifier to capture these UXO determined the total number of clutter dug. In general, PNBC and SBC required many digs to capture these challenging UXO. For the MetalMapper sensor, however, PNBC captured every UXO with only 120 clutter dug (Figure 1). SBC required 600. It is important to note that the global rate at which UXO were dug relative to clutter were virtually identical between SBC and PBNC in this case. The only difference between the two classifiers was that it took an extra 480 digs to capture the last UXO. The interpretation of this result is that PNBC benefitted from knowing the clustering of unlabeled observations in feature space. The UXO were highly separable from the clutter, in general. So, the supervised classifier boundary was well-defined. However, a single, challenging UXO was present past the boundary and there were many clutter that were slightly closer to labeled UXO in features space. As a result, the supervised algorithm dug these clutter ahead of the final, difficult UXO. However, this UXO was part of a cluster in feature space that was connected to a labeled UXO. Even though there were many observations that were closer in terms of raw feature space, these observations were not as well-connected to the labeled UXO. Thus, the challenging, unlabeled UXO was given higher probability in the classification.

Optimality of the parameters chosen for each classifier in the SLO study

There are two parameters that can be adjusted for both SBC and PNBC. The first parameter, which is shared by both PNBC and SBC, is the correlation width of the radial Gaussian kernel. This parameter determines how similar the labels of nearby observations should be. It is this correlation length that allows the classifier to be nonlinear in feature space. The second parameter of SBC, γ, controls the sparseness of the classifier. To make a discrimination model general, a subset of all possible basis vectors is used in the final classification. The weights of all other basis vectors go to 0 and are effectively removed from the model. If all training basis

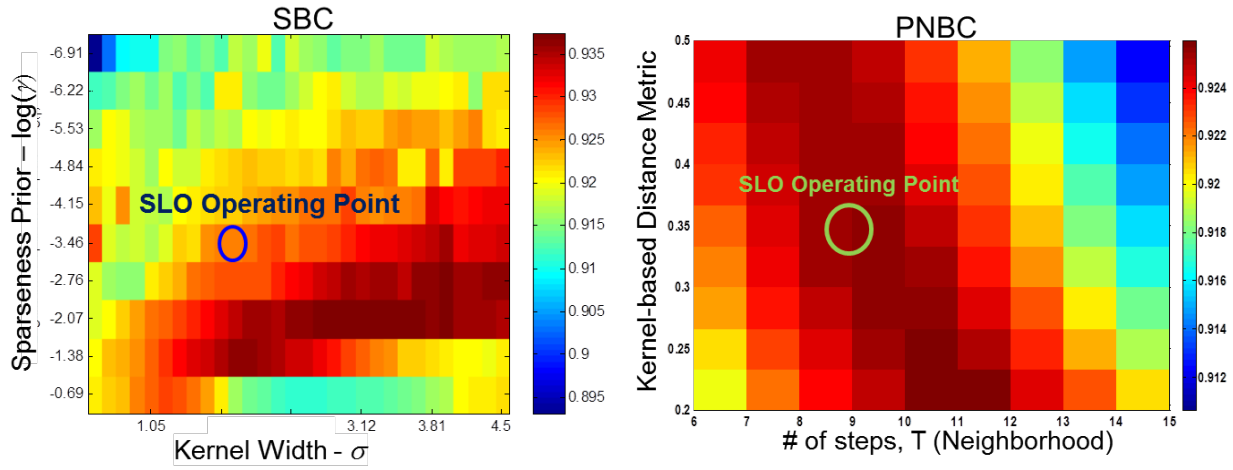

Figure 2. Sensitivity analysis of the user-defined parameters for SBC (left) and PNBC (right). The colored values are the AUCs of classifiers trained using a single combination of parameters. The parameter settings for the performance evaluation are highlighted

vectors were used, then the model would fit perfectly to the training data, but could be a poor predictor of the unlabeled data. Increasing γ causes the weights on many basis vectors to tend toward zero and results in a sparse solution. The second parameter of PNBC is the number of T-steps on the graph of observations. The graph in PNBC is described by the probability of moving from one observation to its nearest neighbor under a Markov random walk. One T-step is equivalent to a single move from a focal point to its nearest neighbor. T-steps greater than one increase the length of the Markov random walk to include additional neighbors. Thus, a T-step of 3 corresponds to a Markov random walk from the focal point to its nearest neighbor, the nearest neighbor to the second point, and the nearest neighbor to the third point. The larger the T-step is the stronger the correlation within clusters of observations. Increasing the T-step is similar in effect to increasing the correlation length of the radial Gaussian kernel except that the correlation between points is constrained by the neighborhood structure of the data.

SIG performed a sensitivity analysis of the SBC and PNBC parameters using the TEMTADS sensor data for SLO. For each parameter combination a classifier was created using the predefined training set and then posterior probabilities were calculated for the unlabeled test data. Then, the area under the ROC curve (AUC) for each classifier was calculated. While this performance metric is not more important than the number of unnecessary digs before the last UXO is captured, it is less sensitive to the particularities of the test dataset. So, it was used as the performance measure for the sensitivity analysis. The results of the sensitivity analysis were compared to the parameter settings during the performance assessment (Figure 2). AUCs were above 0.895 for all parameter combinations, and across a broad range of parameter settings the AUCs were above 0.92. For SBC extremely sparse models (high gamma) and extremely over-fit models (low gamma) tended to perform worse. PNBC performed well over the space of parameters examined and selected operating point was at near peak performance. We conclude that both PNBC and SBC are robust to changes in their user-defined parameters.

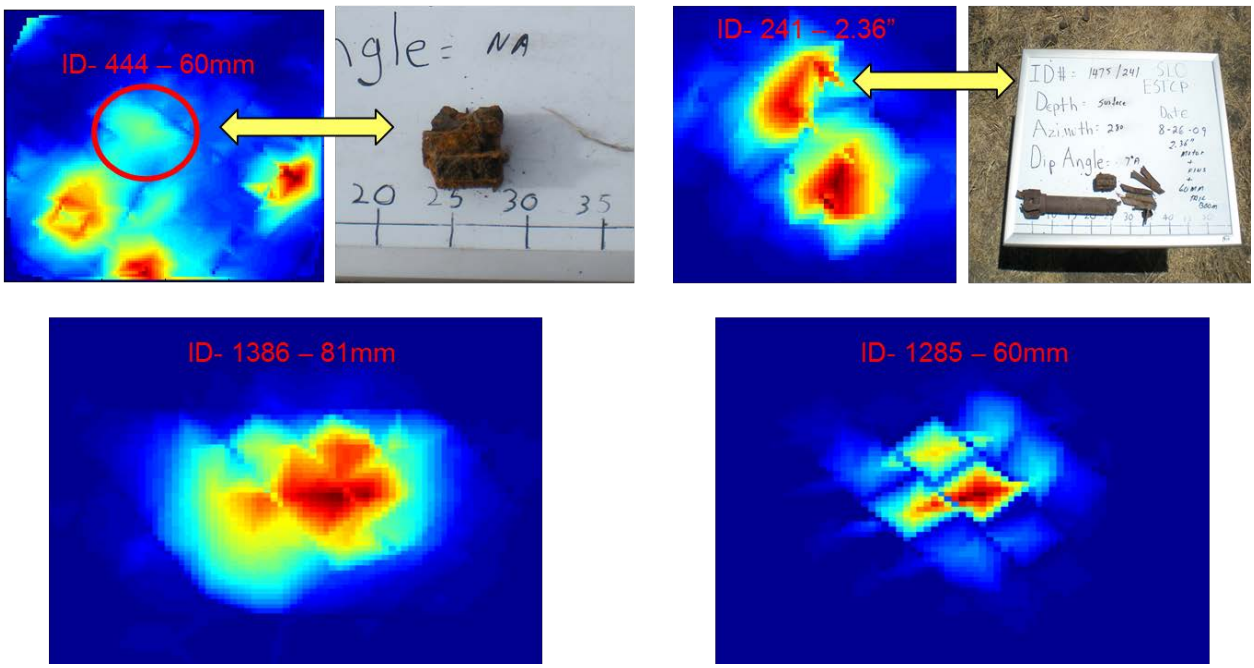

Figure 3. EM61 responses for 4 anomalies that were challenging to discriminate at SLO. The top two anomalies are shown with their pictures.

Overall discrimination failure analysis and performance assessment

The classifier performances for both PNBC and SBC across all sensors were excellent in terms of AUC. The major difference between approaches was in the number of anomalies dug to reveal the last few UXO. Some classifier and sensor combinations performed well on these difficult anomalies (e.g. the PNBC classifier on the MetalMapper sensor data captured all UXO with only 120 unnecessary digs). For most classifiers, however, there was at least one anomaly that proved difficult to discriminate. Four of the challenging anomalies from SLO are highlighted here and an explanation is given for why they were difficult (Figure 3). Two of the anomalies (ID s 444 and 241) were difficult because there were multiple objects nearby (the top two anomalies of Figure 3). The poor features of these anomalies are partially due to the fact that the two munitions were among the smallest in size. No 81mm or larger munitions were challenging due to nearby clutter, though these munitions existed in the dataset. The features of these, small, challenging anomalies could be improved by using a multi-anomaly dipole inversion. The other two difficult anomalies (IDs 1386 and 1285) were difficult because signal quality was poor. Unlike the other challenging anomalies, these responses could not be improved by using multi-anomaly dipole inversions. They could be discriminated by using active learning if their responses were significantly different from other clutter, even if they did not exhibit the 'standard' UXO response. Their responses are similar to other clutter, though. No clear analysis method exists for distinguishing them. Better field protocols are the only obvious solution to such difficult anomalies. Inferior signal quality can make even large munitions (e.g. ID1386: 81mm) difficult to discriminate.

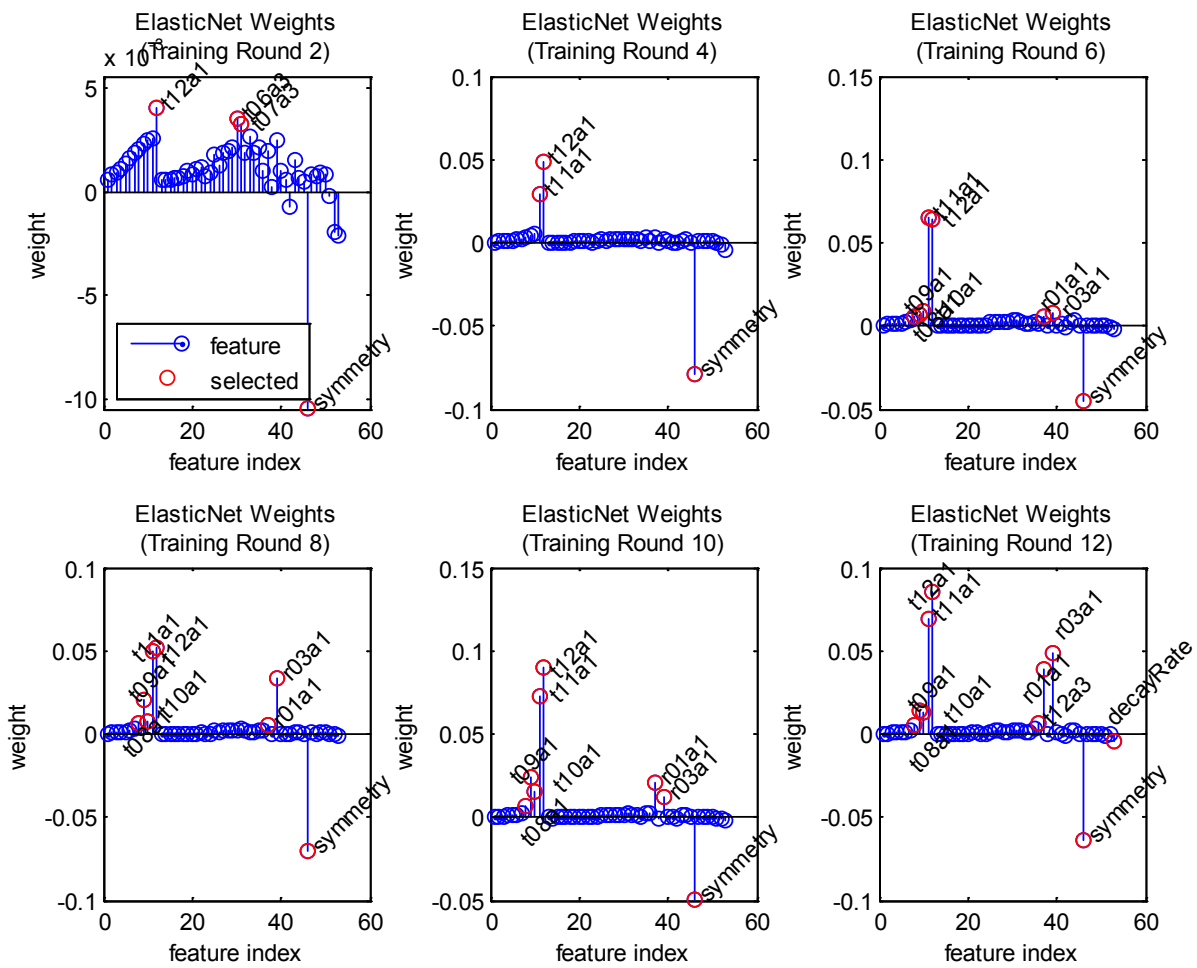

Figure 4. Feature weights for all the even rounds of active learning. Selected features that were used for classification are highlighed (red circles). Feature names for selected features are also shown. The abbreviations are: t*x* – time gate *x*, a*x* – axis *x*, r01 – exponential model intercept, and r03 – exponential model decay rate.

Active Learning Follow-on Analysis

Investigation of an integrated method of active learning with feature selection

For this analysis SIG used the MetalMapper sensor at SLO. There were a total of 53 features: the log polarizabilities for each of the three axes at 12 time gates, 3 parameters for each axis fitting a negative exponential model to the polarizabilities, and 9 shape features describing nonlinear combinations of polarizabilities (e.g. symmetry, eccentricity, aspect ratio). SIG performed feature selection using BENet and then a non-linear classification based on the selected features using PNBC. 20 labels were acquired initially using NMAL. A classifier was then created, and 10 rounds of NMAL were performed. The total number of actively acquired labels was 120. For each round of active learning feature selection was performed using BENet. The selected features evolved as more observations were collected during batch active learning, but eventually came to a relatively stable distribution (Figure 4). In every round the polarizabilities of axis one at the last few time-gates were important, and SIG's measure of

symmetry (i.e. the mean squared differences between the log polarizabilities of axis 2 and 3 across all time gates) was important. These features highlight two important physical features of UXO. First, UXO are typically larger than the surrounding clutter and thus have responses that remain strong at later time gates. Second, they are cylindrical: the two transvers axes are roughly the same size. There was a general trend of selecting more features as more training data acquired. This is expected, since more training data can justify a more complex model. But, these additional features were typically correlated with previously selected features. For instance at later rounds of active learning, more time gates from axis one were used. This highlights the strength of BENet for feature selection as opposed to SBC. SBC would not select such correlated features because it tries to select an orthogonal set. BENet keeps correlated features, and it may be that one of these correlated features has some additional information for discrimination that might predict those few anomalies that are difficult to discriminate.

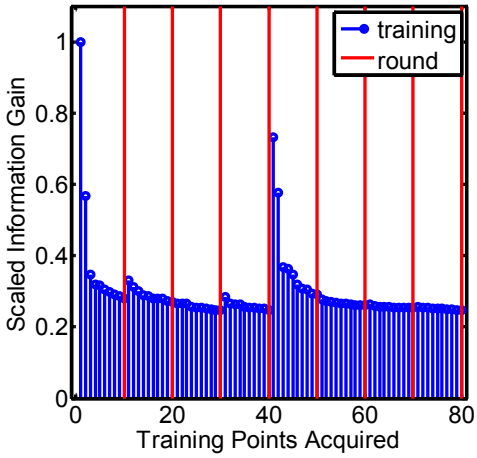

Figure 5. Fisher information gain during NMAL for the MetalMapper sensor data.

Perform active learning with SIG MetalMapper sensor model:

SIG developed a classification model for the MetalMapper sensor that leverages NMAL in a PNBC classifier. The focus in this section is on the information gained from active learning and the decision for stopping the active learning process rather than the performance of the classifier itself. Figure 5 shows the Fisher information gain during the second phase of active learning.

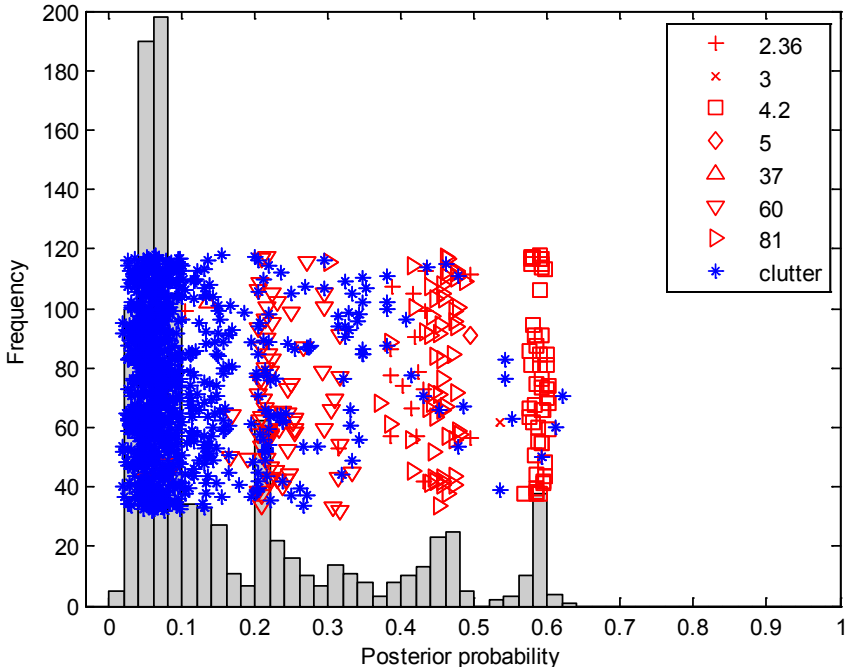

Figure 6. Histogram of predicted probabilities at the end of training PNBC for the MetalMapper sensor using NMAL.

Each red line in Figure 5 separates a single round of NMAL. There were 10 new labels acquired during each round of active learning. Acquiring multiple labels whose joint information gain is maximal is what distinguishes non-myopic active learning from the myopic active learning used previously. As more labels are acquired through NMAL the relative contribution of each additional point decreases. Occasionally, a new UXO is found by NMAL that changes the structure of the model either by revealing new features that are important or by revealing a new mode of UXO in feature space. When such a label is acquired, the information gain increases suddenly. This effect can be seen in Figure 5 at rounds 2, 4, and particularly, 5. The decision to stop digging occurs when it is obvious that future rounds of NMAL will not reveal significant jumps in information gain. This occurred after 60 labels had been acquired. But 2 more rounds of training labels were acquired to be sure the information gain would not change because of some hidden mode of UXO.

Semi-supervised active learning investigation

SIG selected an initial data basis using NMAL of 20 training points for the MetalMapper sensor data at SLO. Feature selection was performed. A PNBC classifier was trained, and then NMAL was used to acquire 10 addition labels. This process was repeated 8 more times until a total of 100 training labels were collected. The T-step for the PNBC classifier was set to 5. The predicted probabilities of each anomaly reveal a good separation between UXO and clutter (Figure 6). The different types of UXO varied in their predicted probabilities. The 4.2 in mortars, being the largest UXO in the dataset, were among the easiest to discriminate. Small munitions such as the 2.36in and 37mm mortars were more difficult. All but 4 of the UXO were dug with only 150 unnecessary digs (Figure 7). Two of these UXO could have been acquired by using a multi-anomaly model.

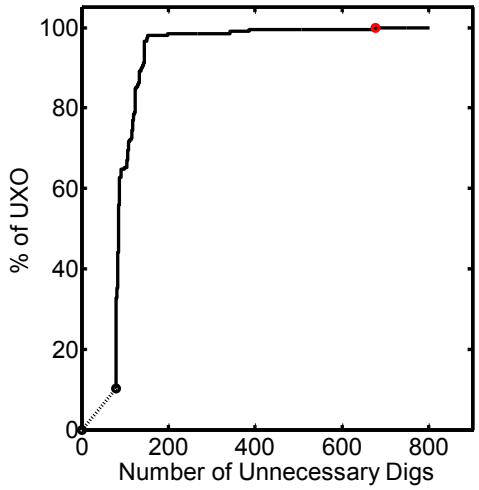

Figure 7. ROC curve for the PNBC classifier at the end of NMAL. Data are from the MetalMapper sensor. The training data are included in the plot (dashed line) along with the last UXO acquired (red).

Active learning did increase the rate at which UXO were acquired by the classifier resulting in slightly larger AUCs for later rounds of NMAL (Figure 8). But, active learning did not have a significant effect on the total number of unnecessary digs to acquire the last UXO. Given that NMAL is effective

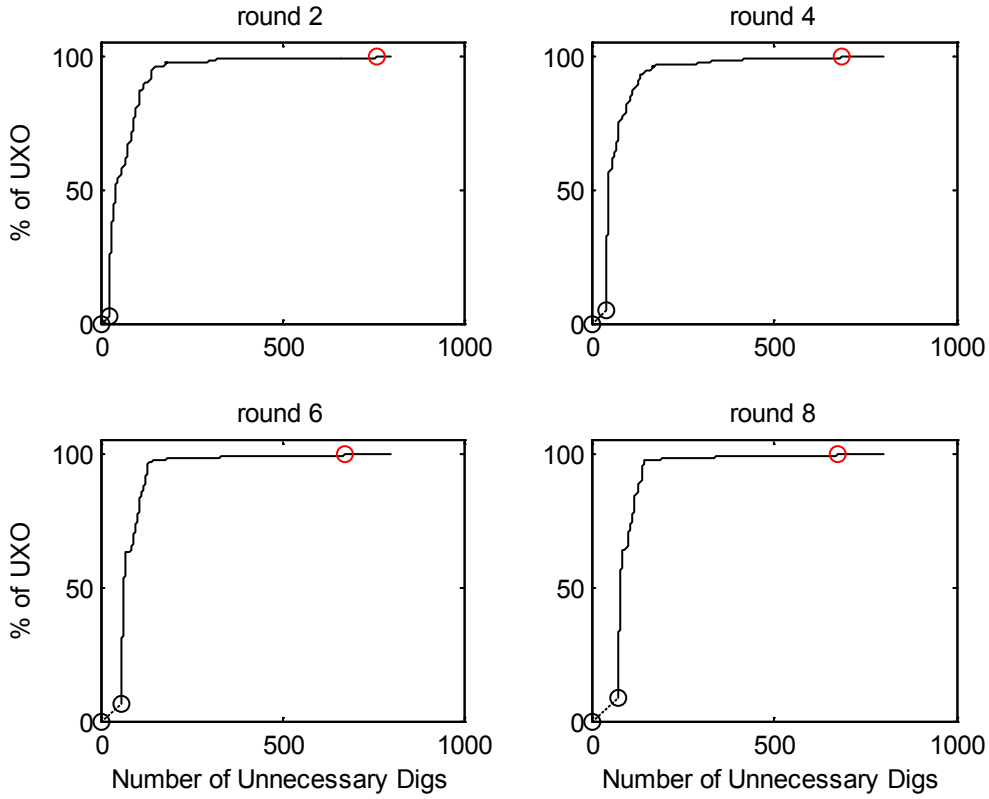

Figure 8. ROC curves during different rounds of NMAL for the PNBC

at finding outlier anomalies the fact that NMAL did not select these challenging observations for training suggests that they were similar to the average clutter response. Again, this highlights the possibility that multiple anomalies were present. Including a multi-anomaly dipole model in the classification would increase separation between the UXO and clutter responses.

Potential of active learning and other methods for finding outliers in feature distributions

Active learning seeks two different objectives simultaneously. The first objective is to reduce uncertainty around the classifier boundary by requesting labels where posterior probabilities are near 0.5. The second objective is to explore areas of the feature space that are poorly represented in the training data. These are not necessarily areas on the boundary between clutter and UXO, but are areas of feature space that have not yet been sampled. Meeting the first objective will not find outlier UXO, but meeting the second will. And, while the both objectives are present in the second phase of active learning, where a classifier is already trained, the second objective is not present in the first phase of active learning. In this section the two different phases of NMAL are examined for their efficacy at finding UXO outliers.

To assess the ability of NMAL to reveal outliers SIG compared the number of UXO captured by three sampling methods as a function of posterior probability of being UXO in the final PNBC classifier for the MetalMapper sensor data. The three sampling methods were random sampling, NMAL without a classifier and the NMAL on the PNBC classifier. The posterior probability associated with the sampled observations can be interpreted as the degree to which a UXO is similar to the average UXO or similar to the average clutter. Thus, UXO with posteriors close to zero are 'outliers' in that they are very similar to other clutter. The results are intuitive given the objectives of the two different types of active learning (Figure 9). NMAL associated with the PNBC classifier tries to refine the classifier boundary and explore the feature space. Many of the UXO captured in this sample come from the boundary (probability approximately 0.5). But, it also samples low probability UXO at a rate greater than a random sample presumably because these outliers occupy unique regions of feature space. NMAL without the classifier does not focus the samples along a classifier boundary, and, thus is more effective at capturing UXO that are most similar to clutter responses. Both types of active learning capture more UXO in their sample than a random sample does. The conclusion of this examination is that NMAL should be very efficient at capturing outlier UXO compared to random sampling. The type of NMAL used for outlier detection is dependent upon whether training the classifier is a simultaneous objective or not. NMAL without the classifier could be performed after all

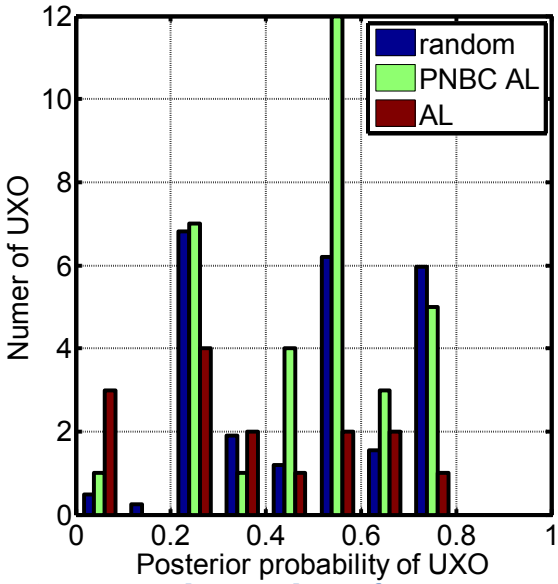

Figure 9. The number of UXO as a function of their posterior probability for 3 different sampling methods with a sample size of 120. Note: AL in the legend is non-myopic.

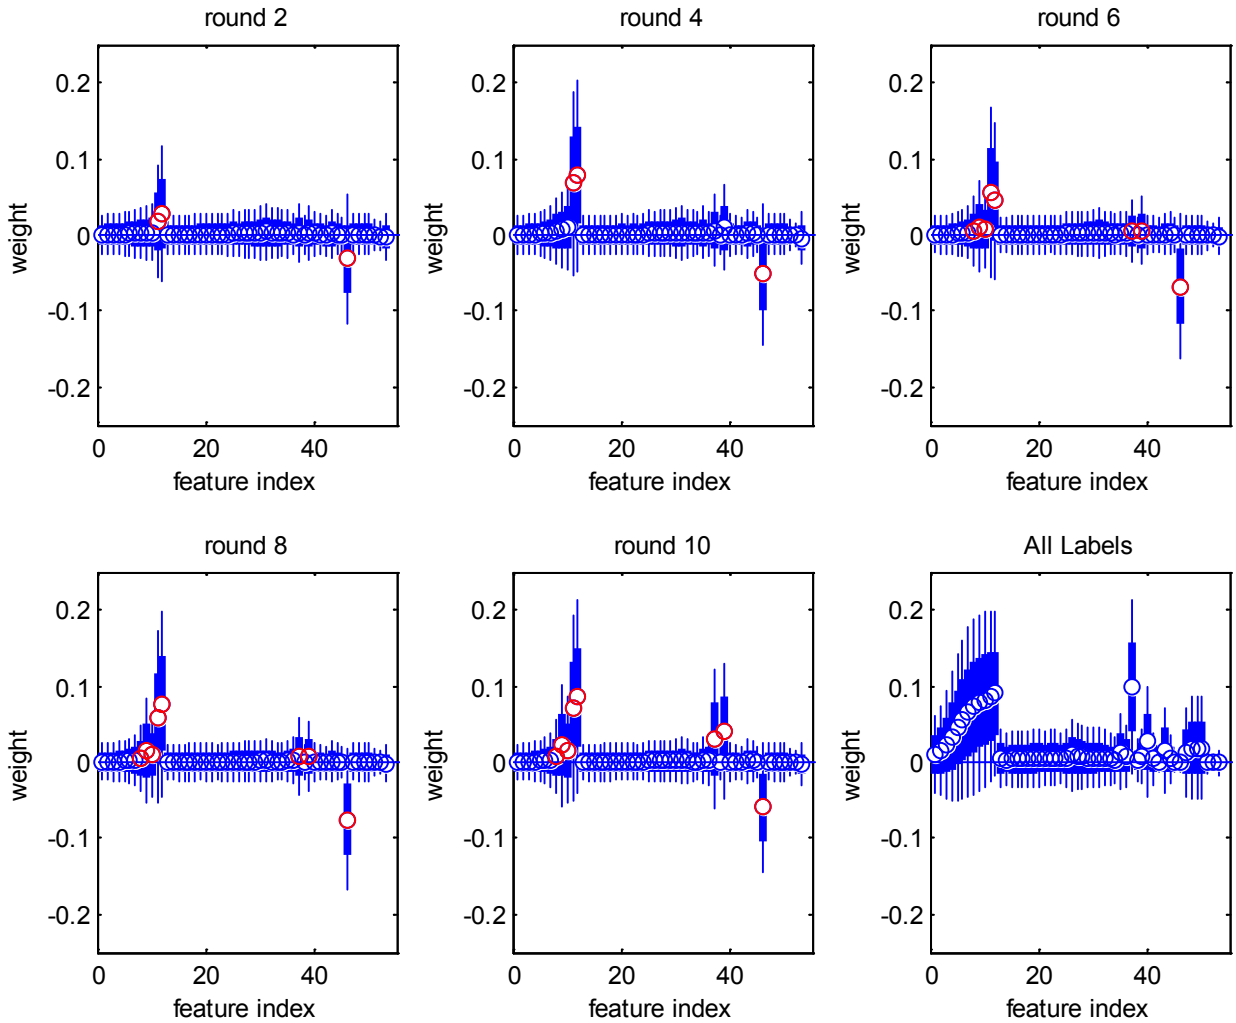

Figure 10. BENet feature weights used for feature selection during the even rounds of NMAL. The marginal posterior distributions of feature weights are shown with +-1 standard deviation (thick blue segments) and a 95% credible interval (thin blue segments). Selected features are highlighted in red. The bottom right plot show the feature weights for a BENet model where all labels were used for training.

the expected UXO are dug as a means of quantifying for stakeholders the degree to which the feature space has been covered by the sampling.

Training Data Adequacy Study

SIG used BENet for feature selection. Since this method is Bayesian, the feature weights have a full posterior distributions estimated through Variational Bayes. Typically, however, only the maximum *a posteriori* estimates for the weights are used. The uncertainty in these posterior distributions quantifies the degree to which the data constrain the feature weights around a particular value. Thus, information from the posterior distributions describes the degree to which feature selection is appropriate given the sample.

Figure 10 shows the distribution of BENet feature weights during a few rounds of training. Few of the features that were selected actually have high confidence of being different than zero.

This is true even for a model where all the labels were used for training, the best possible model. The message from this pattern is that a lot of overlap exists between clutter and UXO in a linear feature space. In spite of this, SIG's feature selection approach tended to select features during active learning that were important in the best model. SIG is confident that feature selection is appropriate with few samples.

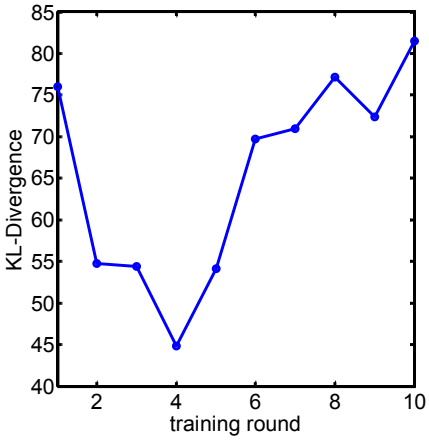

Figure 11. KL divergence from the posteriors of BENet feature weights to the BENet model fitted using all labels. Lower values are more similar to the model using all the labels.

The above comparison is qualitative, comparing the features selected during active learning to the features that would be important the best model. A more quantitative comparison can be achieved by measure the difference between the distributions of posteriors of the feature weights during active learning and the distribution of posterior weights for the best model. The features weights are assumed to be drawn from a multivariate normal distribution. The difference is between the feature weights during NMAL to the feature weights of the best model using Kullback-Leibler (KL) divergence (Figure 11). If the amount of training data were a limiting factor in selecting the best features, then the KL divergence would decrease monotonically as the number of training data increased. This was not the case, however, and suggests that feature selection can be done effectively with few training points. At the very least, the type and quality of the training points is more important for feature selection than quantity of training points.

Sensor/Target Model EnhancemeNts Study

Investigation of the multi-anomaly model for SLO cued sensor data

An example is shown in Figure 12. There are two anomalies under the ground. One is a 37mm mortar, and the other is clutter. The polarizability curves based on the one anomaly mode are not indicative of a UXO. The fit error is quite large (31% relative to the measured data), which implies that the extracted features are not usable. With the two-anomaly model, the fit error decrease to 8.9%, and the polarizability curves associated with the second and third axes of the anomaly reveal a more symmetric anomaly. That is, the extracted features based on two-anomaly model look more like a UXO.

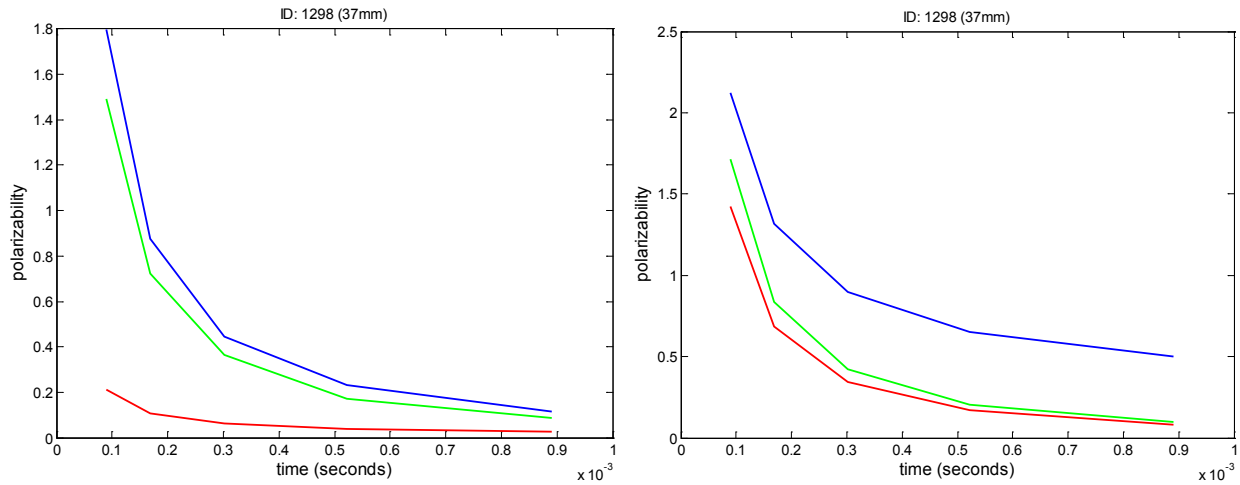

Figure 12 Polarizabilities of a one anomaly model (left) and two anomaly model (right) as a function of time for a 37mm mortar.

When overlapping signatures are present and detectable, investigate reliability of multi-anomaly model inversions for UXO discrimination in SLO data

The success of UXO discrimination depends on the quality of the classifier and the accuracy of the anomalies' features. The latter is strongly associated with the data quality and the assumptions of the multi-anomaly dipole model. One of these assumptions is the number of anomalies. Selecting the wrong number of anomalies could result in worse features.

Determining the number of anomalies is a challenging topic in the UXO discrimination community. However, multi-sensor detecting system like TEMTADS makes this problem solvable. TEMTADS consists of 25 matched transmitter/receiver pairs. Since the number of transmitters and receivers are equal, it is possible to construct a full incident field matrix for each time gate and then find the number of anomalies by performing principal component analysis.

For TEMTADS, the matrix with one specific time gate can be formed as follows

$$M = \begin{bmatrix} H_{1,1} & H_{1,2} & \cdots & H_{1,25} \\ H_{2,1} & H_{2,2} & \cdots & H_{2,25} \\ \vdots & \vdots & \ddots & \vdots \\ H_{25,1} & H_{25,2} & \cdots & H_{25,25} \end{bmatrix}$$

where each column represents the received magnetic fields at different receivers due to the incident field of one transmitter. Either SVD (singular value decomposition) or eigenvalue decomposition can be applied to the above matrix to find the principle components. The principal components of all time gates are then aligned. Joint diagonalization can be used for this purpose.

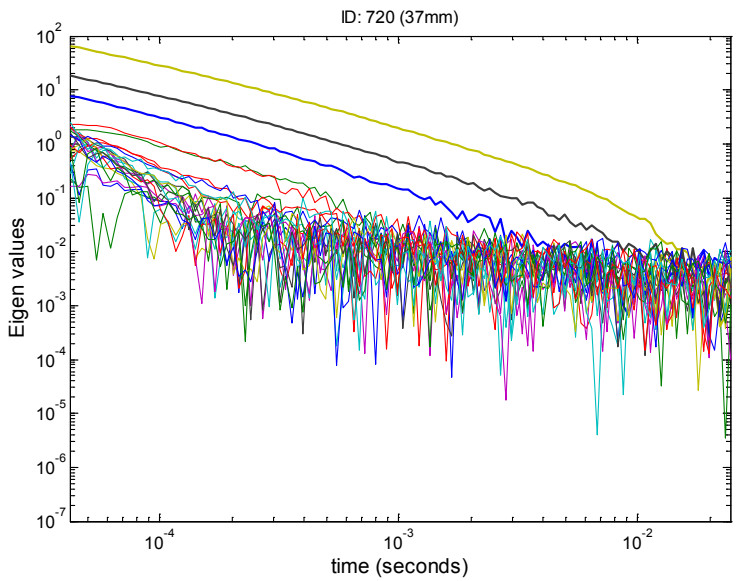

Figure 13 Eigenvalues as a function of time for a 37 mm mortar.

Figure 13 shows an example of the eigenvalues through time for a single anomaly. Most eigenvalues are very small (less than 2) and there are 3 larger eigenvalues. Based on the dipole model, one anomaly is represented by 3 principal axes. Therefore, the 3 eigenvalues in this case are associated with one anomaly. All the other non-significant eigenvalues are due to noises. One-anomaly model should be used in this case.

Two examples are shown in Figures 14 and 15. They each have two anomalies, and they are apparent in the eigenvalues. There are six eigenvalues that are separated from the noise and have large magnitudes. These 6 eigenvalues correspond to 2 anomalies, 3 axes for each. Note that this technique is on applicable for sensors with an equal number of transmitters and receivers, because an equal number is required to generate the full matrix from which principle components are derived.

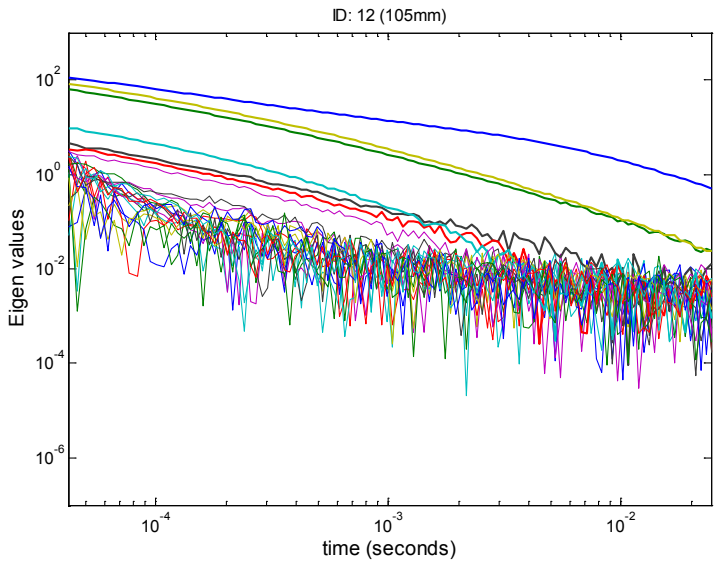

Figure 14 Eigenvalues as a function of time for two-anomaly data (one is 105mm mortar)

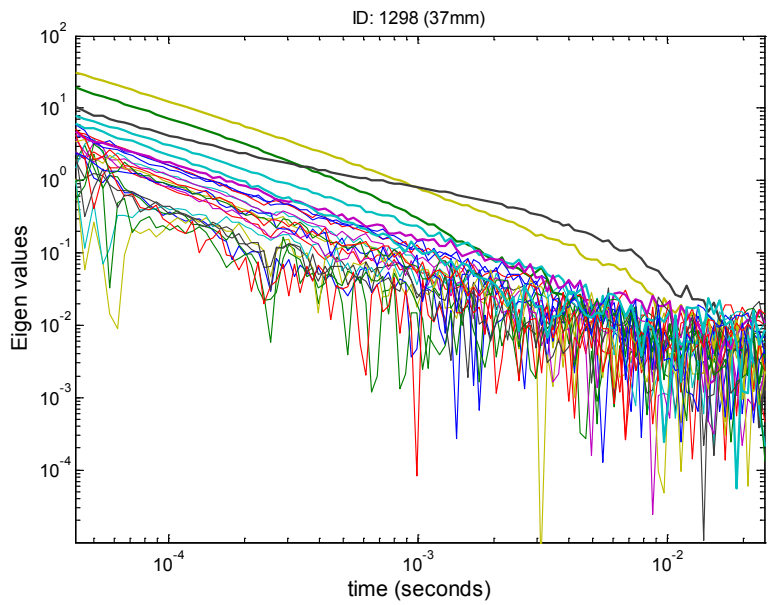

Figure 15 Eigenvalues as a function of time for two-anomaly data (one is 37mm mortar)

Performance evaluation of SIG developed MetalMapper sensor model for SLO discrimination

SIG has developed a MetalMapper model for discrimination. The unknowns are grouped into two parts: positions of the anomalies and the dipole parameters, which are nonlinear and linear, respectively. An efficient, iterative approach is employed to obtain all the parameters. The extracted polarizability parameters were compared with Skip Snyder's results. Some example results are shown in Figures 16-18. They are identical for all 3 cases.

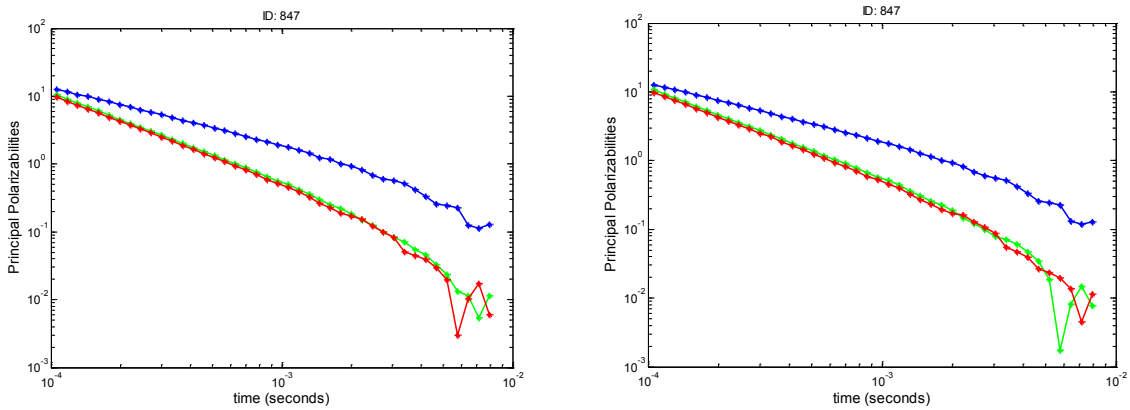

(a) (b)

Figure 16 Comparison of extracted parameters for anomaly 847. (a) Skip's result; (b) SIG's result.

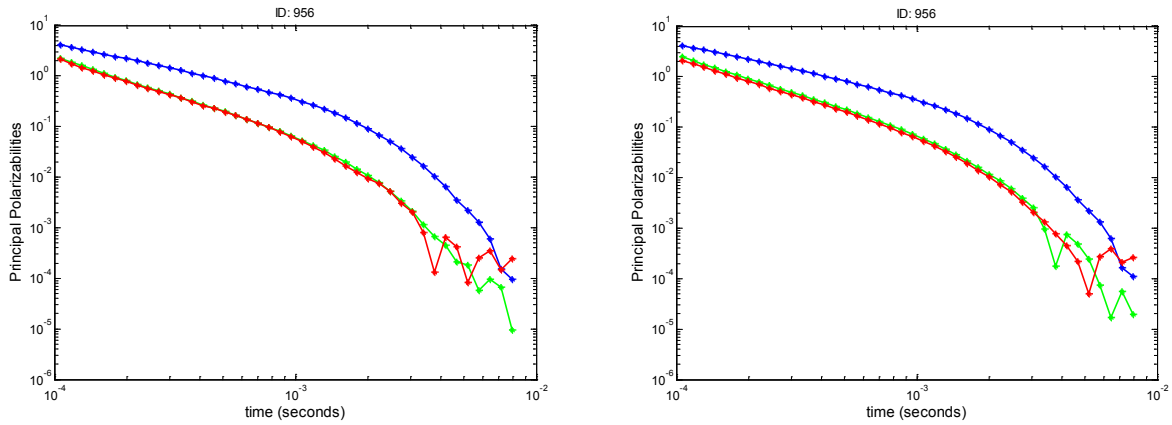

(a) (b)

Figure 17 Comparison of extracted parameters for anomaly 956. (a) Skip's result; (b) SIG's result.

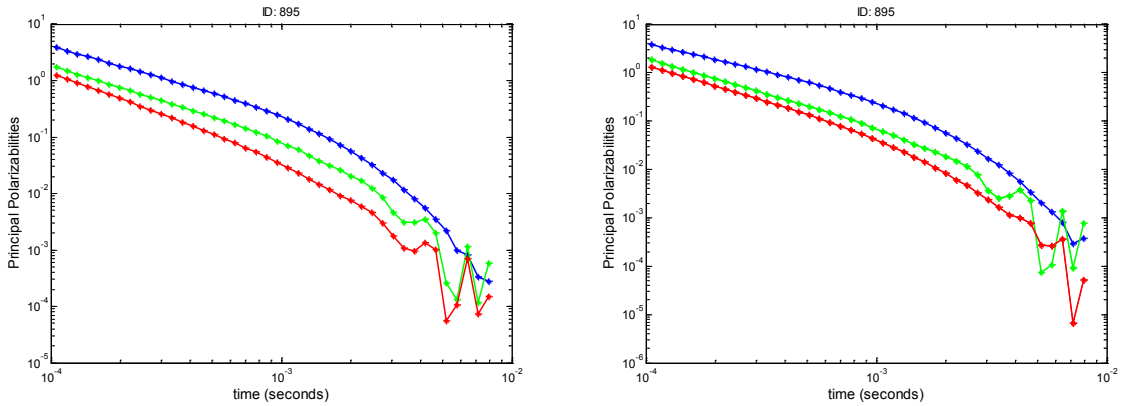

(a) (b)

Figure 18 Comparison of extracted parameters for anomaly 895. (a) Skip's result; (b) SIG's result.

Conclusions and Implications for Future Research/Implementation

The conclusions drawn from the follow-on analysis are summarized with respect to the objectives of the study

Performance Analysis and Technical Approach Validation

While PNBC had fewer unnecessary digs to capture the last UXO in select situations, the rate at which UXO were captured relative to the number of unnecessary digs was similar between SBC and PNBC. This suggests that in select cases, where a difficult UXO is connected to other UXO by neighboring observations (in feature space), PNBC will outperform SBC. Both PNBC and SBC are robust to changes in their user-defined parameters. Finally, active learning can capture UXO that are difficult to discriminate in cases where their responses are different from clutter, even if the response is not similar to the UXO training data.

Active Learning Follow-on Analysis

In general, NMAL should be very efficient at capturing outlier UXO compared to random sampling. For the specific difficult anomalies at SLO, however, active learning did not decrease the total number of unnecessary digs to acquire the last UXO. Including a multi-anomaly dipole model in the classification would increase separation between the UXO and clutter responses. The type of NMAL used for outlier detection is dependent upon whether training the classifier is a simultaneous objective or not. NMAL without the classifier could be performed after all the expected UXO are dug as a means of quantifying for stakeholders the degree to which the feature space has been covered by the sampling.

Training Data Adequacy

The type and quality of the training points is more important for feature selection than quantity of training points. Using NMAL to acquire training points increases their adequacy for feature selection.

Sensor/Target Model Enhancements

The features extracted using SIGs inversion model are not significantly different from Skip Snyder's features.

Literature Cited

[1] B. Delaney and D. Etter, Report of the Defense Science Board Task Force on Unexploded Ordnance, http://www.acq.osd.mil/dsb/reports/uxo.pdf, December 2003.

[2] M.E. Tipping, "Sparse Bayesian learning and the relevance vector machine," The Journal of Machine Learning Research, vol. 1, pp. 211-244, 2001.

[3] L. Quihua, X. Liao, and L. Carin, "Detection of Unexploded Ordnance via Efficient Semisupervised and Active Learning," IEEE Transactions on Geoscience and Remote Sensing, vol. 46, pp. 2558-2567, September 2008.

[4] 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.

[5] 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.

[6] L. Carin, H. Yu, Y. Dalichaouch, A. Perry, and P. Czipott, “On the wideband EMI response of a rotationally symmetric permeable and conducting target,” IEEE Transactions on Geoscience and Remote Sensing, 2001.