SERDP

FINAL REPORT

Multichannel Detection and Acoustic Color-Based Classification of Underwater UXO in Sonar

Mahmood Azimi-Sadjadi
Colorado State University

Steven Kargl
University of Washington

February 2020

SERDP Project MR-2416

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	**2. REPORT TYPE**	**3. DATES COVERED** *(From - To)*
08/31/2015	SERDP Final Report	February 1, 2017 - February 9, 2020

4. TITLE AND SUBTITLE
Multichannel Detection and Acoustic Color-Based Classification of Underwater UXO in Sonar

5a. CONTRACT NUMBER
W912HQ-14-C-0007

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Mahmood Azimi-Sadjadi
Colorado State University

Steven Kargl
University of Washington

5d. PROJECT NUMBER
MR-2416

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Department of Electrical and Computer Engineering Colorado State University Fort Collins, CO 80523

Applied Physics Laboratory
University of Washington
Seattle, WA 98105

8. PERFORMING ORGANIZATION REPORT NUMBER
MR-2416

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Strategic Environmental Research and Development Program (SERDP)
4800 Mark Center Drive
Suite 17D08
Alexandria VA 22350-3600

10. SPONSOR/MONITOR'S ACRONYM(S)
SERDP

11. SPONSOR/MONITOR'S REPORT NUMBER(S) MR-2416

12. DISTRIBUTION / AVAILABILITY STATEMENT
DISTRIBUTION STATEMENT A: This document has been cleared for public release.

13. SUPPLEMENTARY NOTES
N/A

14. ABSTRACT
The Department of Defense (DoD) is currently responsible for clearing many sites that are potentially contaminated with munitions as a result of past training and weapons testing activities. In many cases, these activities occurred near or were performed in shallow water environments where munitions pose threats to public safety and the environment. The objective of this SERDP Exploratory Development (SEED) project is to develop efficient signal processing techniques for the detection and classification of military munitions in shallow underwater environments using data collected from synthetic aperture sonar (SAS) systems.

15. SUBJECT TERMS
Underwater military munitions, detection and classification, signal processing, synthetic aperture sonar (SAS).

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			Dr. M. R. Azimi-Sadjadi
UNCLASS	UNCLASS	UNCLASS	UNCLASS	43	**19b. TELEPHONE NUMBER** *(include area code)* 970-491-7956

Contents

List of Acronyms

Acronym	Meaning
APL-UW	Applied Physics Laboratory, University of Washington
ARL	Applied Research Laboratory
AUV	Autonomous Underwater Vehicle
BB	Broadband
BOSS	Buried Object Scanning Sonar
BPNN	Back-Propagation Neural Network
CMAC	Collaborative Multi-Aspect Classifier
CSU	Colorado State University
DFT	Discrete Fourier Transform
DoD	Department of Defense
FRM	Fast Ray Model
FFT	Fast Fourier Transform
GLRT	Generalized Likelihood Ration Test
HF	High Frequency
MBCD	Multichannel Broadband Coherence Detector
MSC	Matched Subspace Classifier
MSE	Mean Squared Error
MTL	Multi-Task Learning
NRL	Naval Research Laboratory
NSWC-PCD	Naval Surface Warfare Center-Panama City Division
LF	Low Frequency
LFM	Linear Frequency Modulated
LKDL	Linearized Kernel Dictionary Learning
LS	Least Squares
OMP	Orthogonal Matching Pursuit
PDF	Probability Density Function
PSD	Positive Semi-Definite
ROC	Receiver Operating Characteristic
ROI	Region of Interest
RVM	Relevance Vector Machine
SAS	Synthetic Aperture Sonar
SVM	Support Vector Machine
SWAT	Shallow Water Acoustic Toolset
TIER	Target-In-the-Environment Response
TREX13	Target and Reverberation Experiment 2013
UXO	Un-exploded Ordnance
WSS	Wide-Sense Stationary

Acknowledgments

The investigators would like to thank APL-UW and NSWC-Panama City for their support in providing the data used in this study.

Abstract

This work is concerned with the development of robust methods for detection and classification of military munitions in shallow underwater environments using data collected from synthetic aperture sonar (SAS) systems. In this final report we first address the problem of detecting the presence of underwater munitions using the adaptive multichannel coherence analysis framework. Our detection hypothesis is that the presence of munitions in the sonar returns collected from a hydrophone array will lead to higher levels of coherence compared to the retruns from the seafloor alone. This method has been found to produce excellent detection performance on a wide variety of sonar data sets. Here, detection results are presented on two sonar data sets. The first dataset was simulated using a combination of the target-in-environment-response (TIER) model and Personal Computer-Shallow Water Acoustic Toolset (PC-SWAT). Results are presented using standard performance metrics such as probability of detection (P_D), probability of false alarm (P_{FA}), and Receiver Operating Characteristic (ROC) curve of the detector.

The goal of the second part of this work is to develop a robust target classification method that can be applied to the detected contacts to discriminate munitions from non-hostile man made objects and competing clutter. This framework is developed based upon the Matched Subspace Classifier (MSC) using multivariate Acoustic Color data extracted from the raw sonar returns. Scattering models developed by APL-UW were acquired to generate the required training dataset for various UXO and non-UXO objects. This was done owing to the fact that actual sonar data from a wide range of UXO and non-UXO objects is scarce in realistic situations. Although, it may be somewhat ambitious to expect model data capture all the essential features of these objects for target characterization, it will provide us with clues on how to augment the training datasets using perhaps a limited training samples from sonar returns of actual objects to improve the robustness in different environmental conditions. Our classification hypothesis is that spectral content of the sonar backscatter display unique acoustic signatures providing excellent discrimination between different classes of detected objects. Classification results of the MSC with incremental dictionary learning, kernelized MSC, and collaborative multi-aspect classification (CMAC) are provided using two real sonar data sets. The first data set, TREX13, was also collected in the Gulf of Mexico near Panama City Beach, FL, using a rail system. The second data set, BAYEX14, was collected in St. Andrew's Bay (Panama City, FL), also with a rail system. Both of these data sets provide realistic challenges, where factors such as schools of fish, water turbulence, seafloor roughness, and target range were prominent. Results are presented using standard performance metrics such as probability of correct classification (P_{CC}), probability of false alarm (P_{FA}), ROC curve, and area under the curve (AUC).

The detection and classification algorithms developed in this project allow for near real-time assessment of large underwater areas using data collected from low frequency sonar systems. These algorithms not only give the user the ability to assess the degree to which the site is contaminated but also provide capabilities to localize and characterize individual detected objects in varying conditions. The developed methods could be useful in a multitude of sonar applications used to search or survey underwater areas including environmental and oceanographic studies, undersea exploration, the search for wreckage on the sea floor, and mine-hunting. Additionally, the application of the methods can be extended to other multi-sensory systems for remote sensing and surveillance.

MR-2416 Executive Summary

M. R. Azimi-Sadjadi
Colorado State University

E.1 Introduction

The detection and classification of unexploded ordnances (UXO) in sonar data sets is a challenging and unsolved problem that remains an area of active research for government and private researchers, and militaries across the globe. Since the sonar response from a target can vary significantly depending on the object's burial condition and orientation, range, and grazing angles as well as seafloor properties, the identification of hostile objects becomes exceedingly difficult. The signal processing and machine learning communities have made great strides in recent years towards finding classifiers which provide accurate object discrimination in sonar data sets under a variety of conditions. However, several challenges remain before these solutions can be transitioned into autonomous operational system which can continually learn and provide accurate and robust classifications in novel environments, without sacrificing performance on the previous environments.

E.2 Objective

The Department of Defense (DoD) is responsible for clearing many sites which are contaminated with Unexploded Ordnance (UXO) as a result of past training and weapons testing activities. In many cases, these activities occurred near or in shallow water environments where UXOs pose serious threats to public safety and the environment. The main objective of this work is to provide the DoD with robust and reliable detection and classification solutions as it strives to find safer and more cost-effective technologies for underwater munition remediation in varying operating and environmental conditions.

E.3 Technical Approach

The theme of this work is the development of dedicated methods for detection and classification of military munitions that remain robust to operating and environmental changes using data collected from low frequency sonar systems. In Task 1, we developed an environmentally adaptive multi-channel broadband coherence detection (MBCD) method which tests the hypothesis that the presence of a munition leads to coherence patterns among the sonar channels that differ from those of background clutter alone. Owing to the fact that collecting abundant sonar data for underwater munitions in a variety of operating conditions is difficult, if not impossible, empirically validated models of the frequency-dependent acoustic responses from munitions with known geometrical and physical characteristics are generated using a fast ray model (FRM) over a wide range of aspect orientations. These model-generated data sets are then used to construct signal subspaces for a matched subspace classification (MSC) system to classify the detected contacts in real sonar data sets. Our classification hypothesis is that the signal subspaces constructed using acoustic-color templates from model-generated data can effectively serve as a "fingerprint" to classify munitions from actual sonar data sets. To account for variations of the acoustic-color data depending on the operating and environmental conditions, in Task 2 we developed a novel incremental learning algorithm to guarantee the classifier's robustness in such situations. Task 3 aimed at extending MSC classification framework to nonlinear kernel-based systems in order to further improve correct classification rates of UXO and non-UXO objects in difficult conditions. Additionally, multi-aspect classification fusion strategies were studied in Task 4 in order to provide confusion-free decision-making based upon a limited number of aspects/looks. The developed detection and classification

methods are extensively tested and benchmarked against the existing methods on the experimental Target and Reverberation Experiment 2013 (TREX13) experimental sonar data sets.

E.4 Results and Discussion

This two-year research project not only led to the development of new algorithms for robust detection and classification of underwater UXOs in varying operating and environmental conditions but also revealed many important observations and strategies on how to improve the overall performance of such systems. More specifically, these developments and resultant observations drawn from the comprehensive experimental testing and validation are listed below for every task in this two-year project.

TASK 1: Adaptive Multichannel Broadband Coherence Detection

Research efforts in Task 1 of this project led to the development of a multi-channel (multi-sonar) broadband coherence detector (MBCD) and its adaptive version which adjusts the detection threshold automatically based upon the clutter distribution when operating in new environments. The results of this detector on the model-generated TIER-SWAT illuminated many interesting observations. First, multi-sonar (e.g., HF and BB)-based detection using MBCD provides better underwater target detection when compared with the commonly used detectors that rely on some form of matching. Second, the MBCD system can be configured in many variety of ways to perform important trade-off studies. For example, we applied the adaptive version of the MBCD to both the stave and SAS-processed data using two sonar channels namely HF and LF. It can be seen that the adaptive version of the MBCD provided much better detection results when compared with the fixed (non-adaptive) MBCD case and this adaptation can be done easily when the background clutter distribution undergoes significant variation e.g., moving from smooth sandy bottom to rocky conditions. The detection results and the corresponding ROC curve can be seen in Figures 1 and 2, respectively. Many other important trade-off studies in terms of frequency, bandwidth, number of sonar channels, number of array elements, etc. and effects of the detection performance can also be carried out using the MBCD.

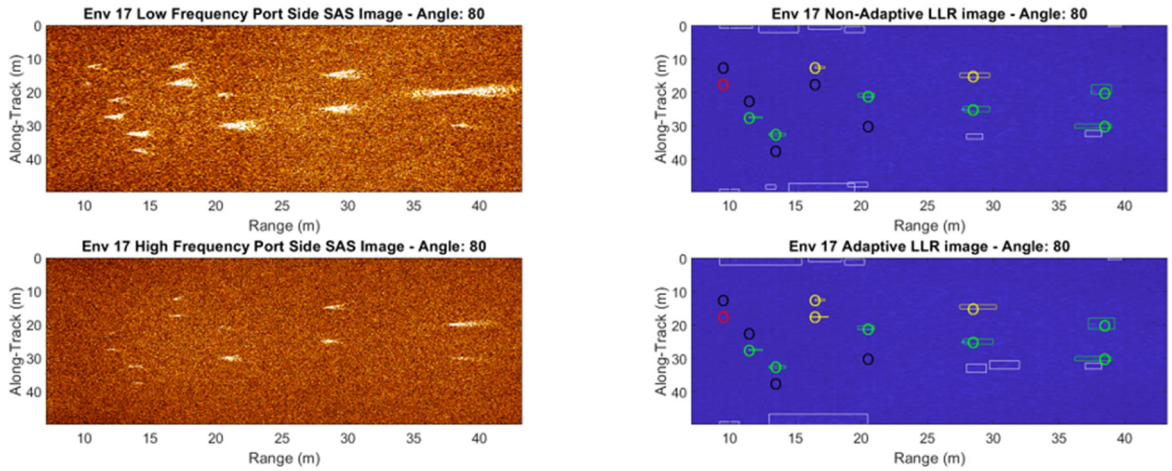

Figure 1: SAS imagery and MBCD detections for 80° object orientation in hard environment. Detection regions and target locations are overlayed on the log-likelihood ratio images.

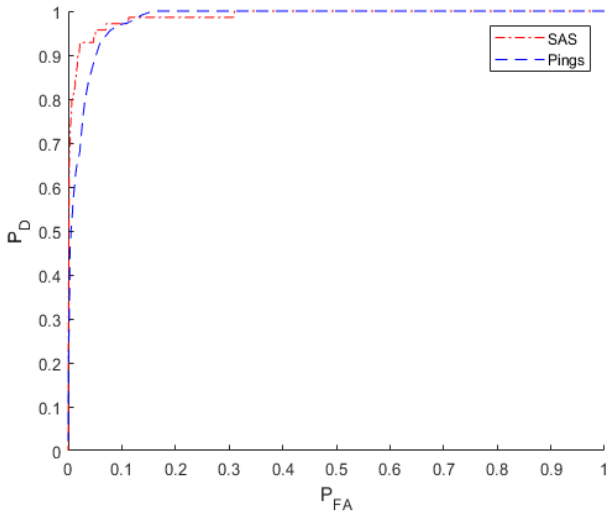

Figure 2: ROC curves of adaptive MBCD on TIER-SWAT data-Rocky bottom.

TASK 2: Incremental Learning of Matched Subspace Classifier

In Task 2, we developed a new incremental learning algorithm for the MSC-based classifier in order to provide the flexibility to learn new UXO and non-UXO objects that can be encountered in new environments without requiring to retrain the entire system. The results of our efforts in this task led to several important observations and conclusions. First, the classification performance of the MSC system when baseline trained using a portion of the TIER model-generated data and then incrementally updated using about 9% of the TREX13 data sets was found to be >10% better than the system without incremental updating. Second, the incrementally updated system possesses excellent generalization ability for classifying unseen (not included in the training data) UXO or non-UXO objects. Third, the proposed incremental learning algorithm does not impact the stability of the previously learnt data, an important property if the system to be used for long-term training in many difficult environments. Fourth, this incremental learning can be accomplished efficiently without the need to carry over the previous training data sets. Finally, the incremental learning algorithm is simple to implement by adding new atoms to the dictionary matrices. Thus, we strongly advocate the use of such incremental learning algorithm for any target classification application in varying conditions. The ROC curves in Figure 3 show the performance of the MSC classifier when incrementally trained using only a small portion of objects found in the TREX13 operating environment with (solid line) and without (dashed-line) two UXO objects left out during the training and then tested on the rest of this data set. Comparing the two ROC curves, the degradation in performance as a result of excluding two UXO objects during the incremental training is found to be less than 2% which is indicative of the great generalization ability of this system.

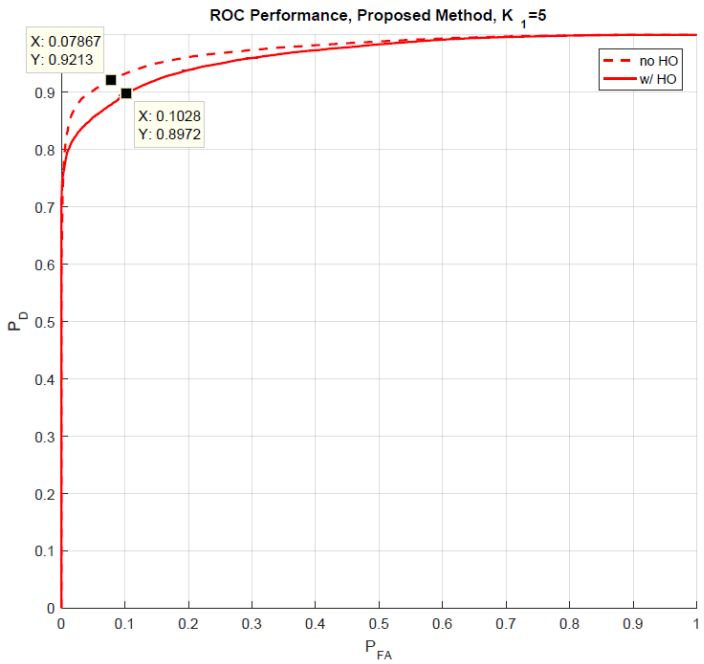

Figure 3: ROC curves for incrementally trained MSC using our new method with and without UXO object holdout.

TASK 3: Non-linear Extensions of Matched Subspace Classifier

The question that we attempted to answer in Task 3 is: "Can a kernel extension of the MSC classification framework be developed (much like the support vector machine (SVM) or relevance vector machine (RVM)) in order to provide substantially better UXO versus non-UXO discrimination performance"? This theoretical extension requires not only extending the theory of the MSC to nonlinear (kernel) based classification but also kernelizing the dictionary learning method used to build the classifier using the subspace matrices. Our investigations, however, indicated that this approach would lead to a very complicated process unsuitable for incremental learning settings. Therefore, we decided to resort to a different technique that takes advantage of the improved discrimination in nonlinearly mapped feature space and at the same time can easily be amenable to incremental leaning in new environments. Our method relies on mapping the training and testing data sets to the corresponding *virtual sample space* which would allow the use of any linear dictionary learning algorithm while providing the benefits of a kernel-based solution. Our recent results presented in Figures 4 and 5 indicated several important key observations. First, the importance of the nonlinear feature mapping on the overall discrimination of the system is clearly evident by the increase in the overall classification performance by more than 7%. Second, this increase in the performance can be gained even when two UXO objects in the TREX13 data set were completely held out during the training. Third, the proposed method relies on using linear dictionary learning while providing the benefits of kernel mapping via the virtual sample space. Fourth, the sampling in this method retains only the essential samples needed for generating the virtual mapping hence it offers a useful mechanism for avoiding overfitting in life-long learning applications. Finally, the proposed method can still take advantage of the incremental learning based on some limited data samples drawn from the new environment without sacrificing the previous learning.

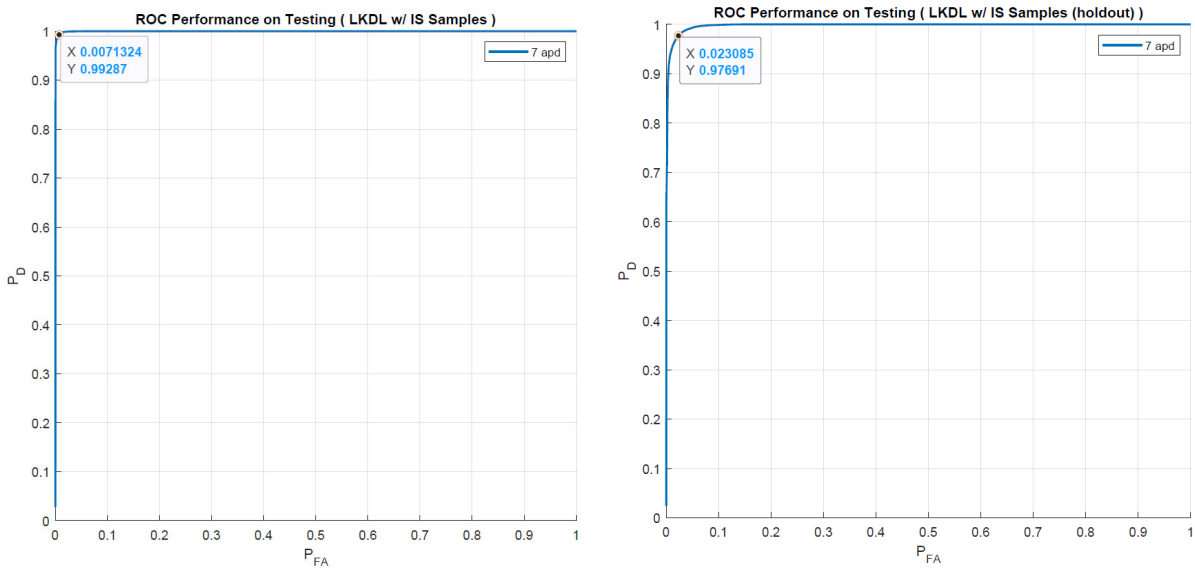

(a) without hold out (b) with hold out

Figure 4: ROC curves of MSC with Virtual Features: (a) without hold out; (b) two UXO objects held out.

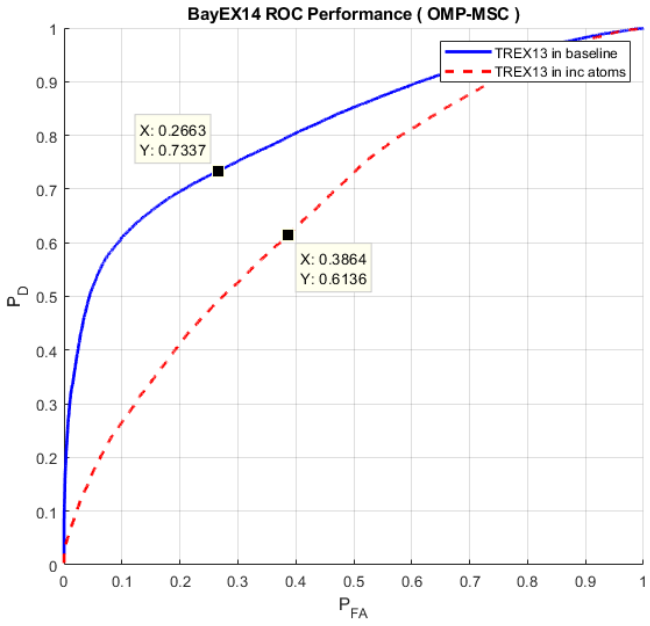

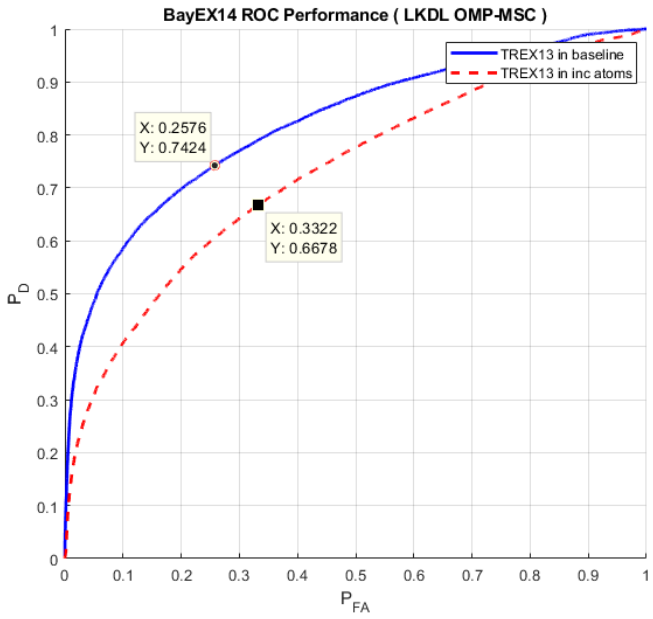

(a) Without Embedding (b) With Embedding

(b)

Figure 5: ROC curves of MSC in two training configurations on unseen BAYEX14 test samples.

TASK 4: Multi-Aspect Classification Fusion

Task 4 research involved the development of an optimum multi-aspect UXO classification fusion system that can provide high confidence decision-making even in situations where only a few useful aspects/looks on the object are available, e.g., targets in vicinity of clutter or partially obscured targets. A new version of the classification fusion system known as the Collaborative Multi-Aspect Classifier (CMAC) was developed utilizing a MSC and a neural network subsystems together with a fusion center. The system makes decisions based upon the acoustic color features at a particular aspect/ping as well as prior decisions made at M previous aspects/pings (decision feedback). The results presented in Figures 6 (a) and (b) led to many interesting and useful conclusions and observations. First, the results on the same TREX13 test data sets showed over 10% improvement in correct classification rate when compared to the MSC-based multi-aspect classification. More specifically, the knee-point of the ROC curve gave $P_{CC} = 98\%$ and $P_{FA} = 2\%$ which are indeed great results given that a linear MSC classifier was used in this system. Second, the system offers the best generalization capability as the degradation in performance due to two UXO object hold out during the training is less than 0.1%. Third, the decision feedback mechanism exploited in this system increases the separation between the distributions of the test statistic (see Figures 7 and 8) under both target and non-target hypothesis hence leading to significantly improved discrimination. Finally, the proposed system is simple to implement and can easily incorporate a variable number of aspects/pings to determine the final class label of an object, and hence provides a very versatile and modular means of performing multiple-ping UXO classification fusion.

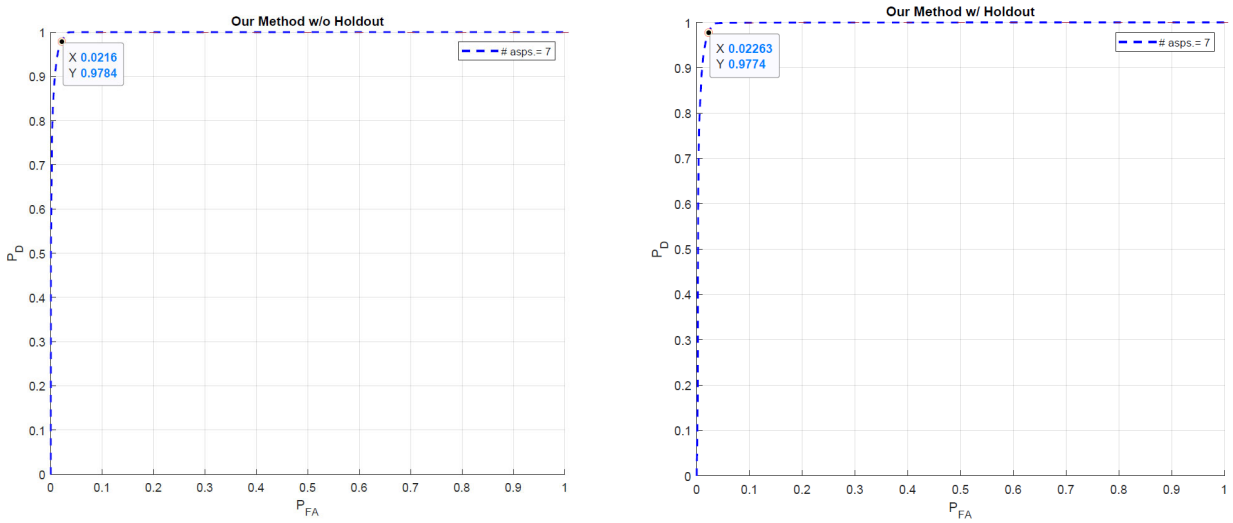

(a) no object held out (b) with objects held out

Figure 6: ROC curves of CMAC with decision-feedback and incrementally trained MSC-without and with object holdout.

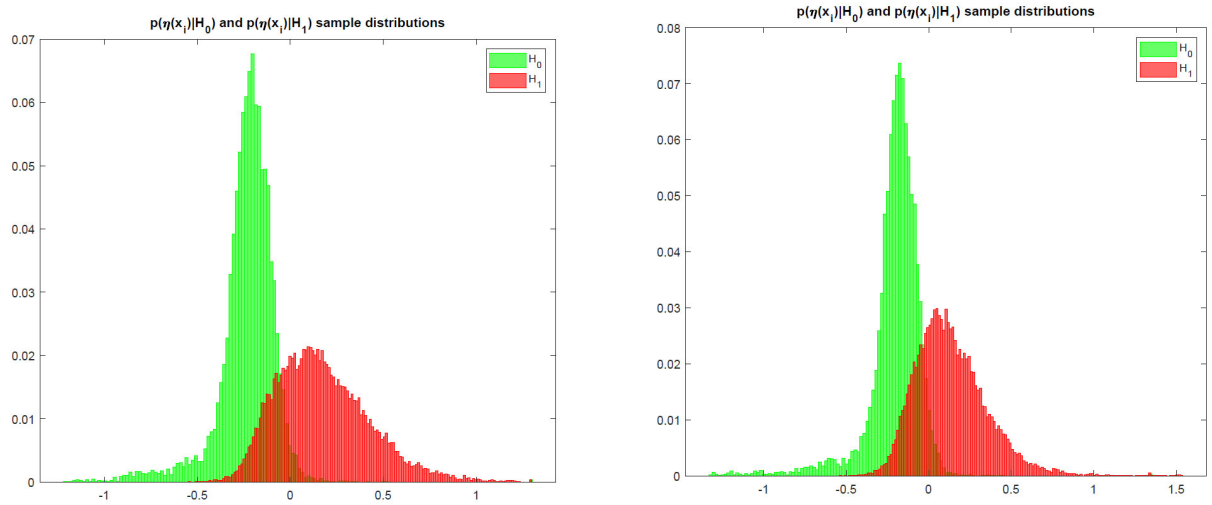

(a) no object held out (b) with objects held out

Figure 7: Class conditional distribution of MSC test statistics- without and with object holdout.

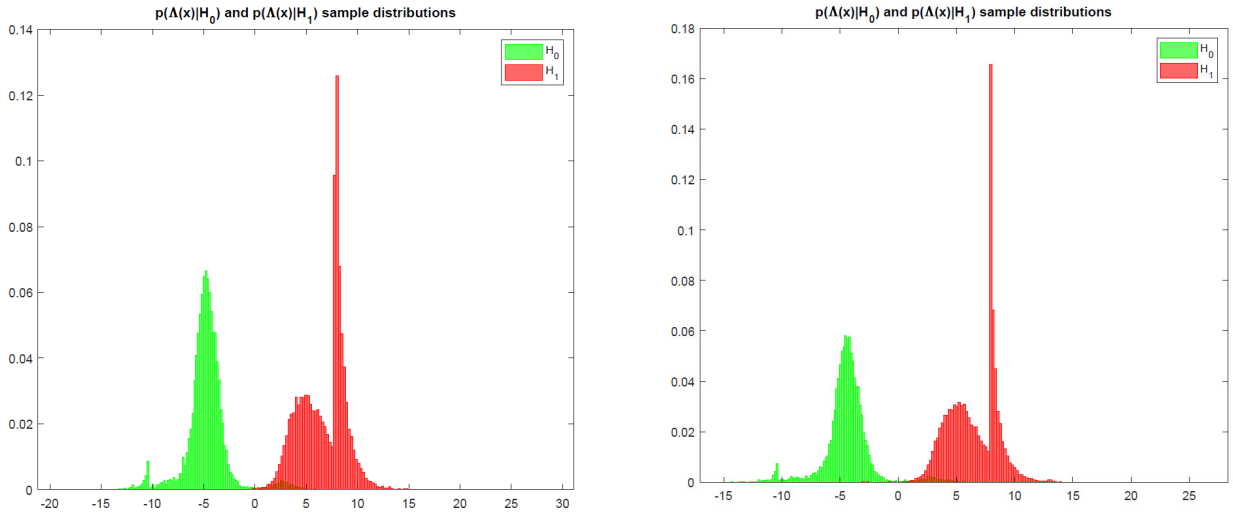

(a) no object held out (b) with objects held out

Figure 8: Class conditional distribution of final test statistics- without and with object holdout

E.5 Implications for Future Research and Benefits

We believe this two-year research project made significant strides toward developing underwater UXO detection and classification systems that remain robust to variations in operating and environmental conditions and hence addressed many critical questions on how to construct such systems for successful deployment in different shallow water settings. In particular, our research studies culminated to conclusive and thorough experimental results that highlight the importance and merit of all the above-mentioned algorithms as critical ingredients of any successful UXO detection and classification system. Consequently, we propose to continue our current research efforts by investigating the following important areas:

1. Comprehensive testing, evaluation, and fine-tuning of our developed systems on more realistic data sets e.g., CLUTTEREX17 or BOSS data sets.

2. Conduct tradeoff studies of the MBCD to determine impact of different sonar configurations and choice of parameters on the ability to provide unambiguous UXO detection in different settings.

3. Study the effects of life-long incremental learning and how to overcome performance drift and structural complexity of the linear MSC or kernel-based MSC classifiers.

4. Develop a new incremental updating mechanism for the mapping of the original data to the virtual feature space. This, will allow the full implementation of the kernel-based incremental learning when new data samples become available in small batches.

5. Develop a new mechanism for the CMAC-based multi-aspect classification fusion system in order to automatically decide the number of aspects/pings needed to perform confusion-free decision-making in different conditions.

Transition the code and documentations to our collaborators at APL-UW and NSWC-PCD and help with testing and validation

1 Objective

The Department of Defense (DoD) is responsible for clearing many sites which are contaminated with Unexploded Ordnance (UXO) as a result of past training and weapons testing activities. In many cases, these activities occurred near or in shallow water environments where UXOs pose serious threats to public safety and the environment. The main objective of this work is to provide the DoD with robust and reliable detection and classification solutions as it strives to find safer and more cost-effective technologies for underwater munition remediation in varying operating and environmental conditions. This research project responds to SERDP Statement of Need (SON) MR-16-04 in *Wide Area and Detailed Surveys* for rapid and highly efficient detection and classification of underwater UXOs found in contaminated sites.

2 Background

Over the past few years, many researchers have developed different classes of approaches to robust decision-making in the presence of changes in target features and environmental conditions. Multi-channel detection methods [1]-[3] that use multiple sonar frequency bands, e.g., high frequency (HF) and broadband (BB), have recently attracted some attentions due to their ability to provide substantially better detection performance comparing to single-channel systems in varying conditions. In [4], a non-parametric approach to multi-channel detection is proposed by defining the generalized coherence measure among multiple channels. In [3], the problem was posed as a test of independence among the complex beamformed data of multiple sensor channels. More specifically, this test for independence becomes a test of block-diagonal structure in a composite covariance matrix of the channels. It was shown [3] that under the 2-D wide-sense stationary (WSS) assumption (i.e. both along-track and range dimensions) of each channel, the Generalized Likelihood Ratio Test (GLRT) leads to a *broadband coherence* test which can be implemented very efficiently. In our final report of the SEED proposal [5], we presented the results of this multichannel detector when applied to the PondEX10 data set. These results showed that the broadband coherence statistic is indeed capable of detecting munitions lying on the seafloor from background with probability of detection of $P_D = 100\%$ and an average of 1 false alarm per image.

Various methods have also been developed for modeling the acoustic response of underwater objects with geometries typically observed in mine and UXO-hunting applications and using this information for the purposes of classification. In [6], the authors considered SAS imaging of simple targets by combining models for reverberation, acoustic penetration, and target scattering into a unified model. This is then used to generate pings suitable for SAS simulations over a range of environmental and experimental conditions. Experimentally measured target scattering from proud and buried targets are then used to validate the model through several simulations. In [7], the authors analyze experimental results from a SAS data set collected in a fresh water pond. These measurements were conducted to investigate discrimination capabilities based on the acoustic response of targets for underwater UXO applications. Results from this study showed that it is possible to use the acoustic template as a fingerprint to uniquely identify a given target. In [8], it was further shown that these acoustic-color features are useful for discriminating similarly shaped targets.

Considerable efforts have recently been made to develop different feature extraction and classification methods [9]-[14] for underwater UXO discrimination. In [9], Bucaro et. al. used a finite-element (FE)-based structural acoustics model to simulate sonar returns from underwater UXOs at various orientations. A Relevance Vector Machine (RVM) classifier [15] was trained using these model-generated data. The trained RVM classifier was shown to be successful in classifying different objects using their sonar

returns. Fischell et. al. [10] used 3-D bistatic scattered field data from spherical and cylindrical targets that were collected by an autonomous underwater vehicle (AUV) to perform UXO versus non-UXO classification using a Support Vector Machine (SVM) [15]. They used both experimental and simulated data to evaluate classifier's performance. In [11], the authors developed new coherent-based feature extraction and synthetic aperture sonar (SAS)-like acoustic color for detection and classification of underwater objects. New multi-aspect classification algorithms were also developed using hidden Markov models (HMMs) [12],[13] and a Collaborative Multi-Aspect Classifier (CMAC) [14] to improve classification accuracy while reducing the false alarms when multi-pings/aspects sonar data are available. By using physical models [6],[7], [16],[17] to construct a signal subspace spanned by the acoustic response of a particular target over a range of aspect orientations, we showed [18] that the matched subspace classifier (MSC) in [19] can provide an effective classification method that remains robust to changes in target orientation and range. Moreover, by keying in on specific target responses, the classifier was shown to exhibit low false alarm rates.

3 Technical Approach

The theme of this work is the development of dedicated methods for detection and classification of military munitions that remain robust to operating and environmental changes using data collected from low frequency sonar systems. In Task 1, we developed an environmentally adaptive multi-channel broadband coherence detection (MBCD) method [1] which tests the hypothesis that the presence of a munition leads to coherence patterns among the sonar channels that differ from those of background clutter alone. Owing to the fact that collecting abundant sonar data for underwater munitions in a variety of operating conditions is difficult, if not impossible, empirically validated models of the frequency-dependent acoustic responses from munitions with known geometrical and physical characteristics are generated using a fast ray model (FRM) [16], [17] over a wide range of aspect orientations. These model-generated data sets are then used to construct signal subspaces for a matched subspace classification (MSC) [18] system to classify the detected contacts in real sonar data sets. Our classification hypothesis is that the signal subspaces constructed using acoustic-color templates from model-generated data can effectively serve as a *fingerprint* to classify munitions from actual sonar data sets. To account for variations of the acoustic-color data depending on the operating and environmental conditions, in Task 2 we developed a novel incremental learning algorithm [20] to guarantee the classifier's robustness in such situations. Task 3 aimed at extending MSC classification framework to nonlinear kernel-based systems in order to further improve correct classification rates of UXO and non-UXO objects in difficult conditions. Additionally, multi-aspect classification fusion [14] strategies are studied in Task 4 in order to provide confusion-free decision-making based upon a limited number of aspects/looks. The developed detection and classification methods are extensively tested and benchmarked against the existing methods on the Target and Reverberation Experiment 2013 (TREX13) experimental sonar data sets.

4 Materials and Methods

This final report gives a detailed review of our accomplishments and results for the work performed on all tasks since the start of the project. Along with a brief overview of the developed theory for each task, recent results for all tasks are provided. More specifically, we report Task 1 detection results of the adaptive MBCD Test on the TIER-SWAT data set. Results of Task 2 on our incremental learning algorithm [18], [20] for updating the MSC system are presented on TREX13 data sets. New results for Task 3, which involved the development of a nonlinear extension of the MSC classification framework are also presented which attest to the benefits of this extension for improving the overall classification accuracy. Finally, theory and results of our multi-aspect decision fusion technique (Task 4) developed and tested during this effort are presented.

The organization of this final report is as follows. In Section 5, we present an overview of the adaptive MBCD and its results on raw sonar ping data from the TIER-SWAT data sets. In Section 6, we briefly describe the theory behind the MSC system and its incremental learning when baseline trained using the model-generated TIER data sets and then incrementally trained and tested on the TREX13 data sets. In Section 7, we present the theory and extension of the kernel MSC system together with the results and comparison with those of the original MSC system. This section also includes results on the BAYEX14 data set using both the incremental and kernel MSC systems. Section 8 presents a brief review of the collaborative multi-aspect classification (CMAC) fusion system with decision feedback and its results on the TREX13 data sets. Finally, conclusions on the results of this report and ideas for future work are discussed in Section 9.

5 Task 1: Adaptive Multichannel Broadband Coherence Detection

In this section, we provide an overview of the MBCD [3] and how it can be adapted for operations in different environments. The adaptability allows the MBCD to maintain a predetermined probability of false alarm (*PFA*) when faced with operating environments that provide new and potentially changing statistical conditions.

5.1 Multichannel Broadband Coherence Detection (MBCD)

Consider a set of L random data matrices$\{X_i\}_{i=1}^{L}$ with each matrix

$$X_i = \begin{bmatrix} x_i[0,0] & x_i[0,1] & \cdots & x_i[0,N-1] \\ x_i[1,0] & x_i[1,1] & & x_i[1,N-1] \\ \vdots & \vdots & \ddots & \vdots \\ x_i[M-1,0] & x_i[M-1,1] & \cdots & x_i[M-1,N-1] \end{bmatrix} \in \mathbb{C}^{M\times N} \tag{1}$$

representing a 2-dimensional, zero-mean random process captured at sonar platform i. In this particular application, the random variable $x_i[m,n]$ represents the n^{th} temporal sample collected from the m^{th} hydrophone element in an array from the i^{th} sonar channel. Stacking the columns of X_i to form the vector $\mathbf{x}_i = \text{vec}(X_i)$, the composite vector $\mathbf{z} = [\mathbf{x}_1^T \ \cdots \ \mathbf{x}_L^T]^T$ has covariance matrix

$$R = E[\mathbf{z}\mathbf{z}^H] = \begin{bmatrix} R_{11} & R_{12} & \cdots & R_{1L} \\ R_{12}^H & R_{22} & & R_{2L} \\ \vdots & \vdots & \ddots & \vdots \\ R_{1L}^H & R_{2L}^H & \cdots & R_{LL} \end{bmatrix} \in \mathbb{C}^{LMN\times LMN}$$

with $R_{ik} = R_{ki}^H = E[\mathbf{x}_i\mathbf{x}_k^H] \in \mathbb{C}^{MN\times MN}$. This matrix captures all space-time second-order information within and between the random vectors $\{\mathbf{x}_i\}_{i=1}^{L}$.

We now assume we are given an experiment producing P *iid* realizations $\{\mathbf{x}_i[p]\}_{p=1}^{P}$ of the random vector associated with the i^{th} channel. In the contexts of this application, these P independent copies could, for example, represent multiple pings collected from each frequency band. All P realizations may then be used to form the data matrix

$$\mathcal{Z} = \left[\mathbf{z}[i] \cdots \mathbf{z}[P]\right] = \begin{bmatrix} \mathbf{x}_1[1] & \cdots & \mathbf{x}_1[P] \\ \vdots & \ddots & \vdots \\ \mathbf{x}_L[1] & \cdots & \mathbf{x}_L[P] \end{bmatrix} \in \mathbb{C}^{LNM \times P} \tag{2}$$

where $\mathbf{z}[p] = [\mathbf{x}_1^T[p] \cdots \mathbf{x}_L^T[p]]$. Using the Generalized Likelihood Ratio Test (GLRT) [21] to test the null hypothesis that matrix R is block diagonal $R = \text{blkdiag}\{R_{11}, \cdots, R_{LL}\}$ involves forming an estimate of the composite covariance matrix

$$\hat{R} = \frac{1}{P}\mathcal{Z}\mathcal{Z}^H = \frac{1}{P}\sum_{p=1}^{P}\mathbf{z}[p]\mathbf{z}^H[p] = \begin{bmatrix} \hat{R}_{11} & \hat{R}_{12} & \cdots & \hat{R}_{1L} \\ \hat{R}_{12}^H & \hat{R}_{22} & & \hat{R}_{2L} \\ \vdots & \vdots & \ddots & \vdots \\ \hat{R}_{1L}^H & \hat{R}_{2L}^H & \cdots & \hat{R}_{LL} \end{bmatrix}$$

And computing the likelihood ratio [1]

$$\Lambda = \frac{\det\hat{R}}{\prod_{i=1}^{L}\det\hat{R}_{ii}} \tag{3}$$

The likelihood ratio given in (3) is referred to as a *generalized* Hadamard ratio [22]. A measure of *bulk coherence* among all the channels may be written as

$$\gamma^2 = 1 - \frac{\det\widehat{\mathbf{R}}}{\prod_{i=1}^{L}\det\widehat{\mathbf{R}}_{ii}} \tag{4}$$

which takes a value of $\gamma^2 = 0$ when $\widehat{\mathbf{R}}$ is block-diagonal (i.e. null hypothesis). An important and useful property of the test statistic in (3) or (4) is its invariance to channel-wise nonsingular linear transformation, i.e. $\mathbf{T}_i\mathbf{x}_i$, of the data including scaling, filtering, unitary transformation (e.g., DFT), and also permutation or ordering of the channel index [3]. *Invariance to scaling is, in particular, very useful property for sonar target detection applications.*

The likelihood ratio in (3) is general in that no specific structure is imposed on the composite covariance matrix other than it is block-diagonal under the null hypothesis. However, doing so requires a data rich environment in which $P \gg LMN$. In cases of data poverty, one can impose additional structure on the covariance matrix, generate the maximum likelihood estimates under that model, and construct the likelihood ratio. Such is the case for the MBCD where we assume that the random vectors from each channel are realizations from 2-D wide-sense stationary (WSS) processes. Results on asymptotically large block-Toeplitz matrices [22] show that, as $M, N \to \infty$, the likelihood ratio in (3) converges to the *broadband coherence test* [3],[22],

$$\Lambda^{\frac{1}{MN}} \to \exp\left\{\int_{-\pi}^{\pi}\int_{-\pi}^{\pi}\ln\frac{\det\hat{S}(e^{j\theta}, e^{j\phi})}{\prod_{i=1}^{L}\hat{S}_{ii}(e^{j\theta}, e^{j\phi})}\frac{d\theta d\phi}{4\pi^2}\right\} \tag{5}$$

with

$$\hat{S}(e^{j\theta}, e^{j\phi}) = \begin{bmatrix} \hat{S}_{11}(e^{j\theta}, e^{j\phi}) & \hat{S}_{12}(e^{j\theta}, e^{j\phi}) & \cdots & \hat{S}_{1L}(e^{j\theta}, e^{j\phi}) \\ \hat{S}_{12}^*(e^{j\theta}, e^{j\phi}) & \hat{S}_{22}(e^{j\theta}, e^{j\phi}) & & \hat{S}_{2L}(e^{j\theta}, e^{j\phi}) \\ \vdots & \vdots & \ddots & \vdots \\ \hat{S}_{1L}^*(e^{j\theta}, e^{j\phi}) & \hat{S}_{2L}^*(e^{j\theta}, e^{j\phi}) & \cdots & \hat{S}_{LL}(e^{j\theta}, e^{j\phi}) \end{bmatrix} \in \mathbb{C}^{L\times L} \tag{6}$$

being an estimated composite power spectral density matrix. Here, $\hat{S}_{ik}\left(e^{j\theta}, e^{j\phi}\right)$ represents a quadratic estimate of the cross-power spectrum between channels i and k at frequency $-\pi < \theta \leq \pi$ and

wavenumber $-\pi < \phi \leq \pi$. Thus, the likelihood ratio becomes a frequency/wavenumber-dependent Hadamard ratio integrated over the Nyquist band. This likelihood ratio can then be compared to a threshold (determined experimentally based upon some training data) to decide whether or not UXO's are present. Moreover, comparing to the likelihood test in (3), the test statistic given in (5) is computationally more efficient as the estimated cross-spectral matrix in (6) can be computed efficiently using a Fast Fourier Transform (FFT) of the waveforms received by each channel. More specifically, if we let $X_i^{(p)}(e^{j\theta}, e^{j\phi}) \in \mathbb{C}$ represent the two dimensional FFT of the p^{th} realization of the data matrix given in (1) at frequency θ and wavenumber ϕ, then the cross-power spectrum may be computed as

$$\hat{S}_{ik}\left(e^{j\theta}, e^{j\phi}\right) = \frac{1}{P}\sum_{p=1}^{P} X_i^{(p)}\left(e^{j\theta}, e^{j\phi}\right)\left(X_k^{(p)}\left(e^{j\theta}, e^{j\phi}\right)\right)^* \tag{7}$$

where * is complex-conjugate operation.

5.2 Adaptive MBCD for UXO Detection

When using the detector in (5) in a new munition site with different sediment type and bottom conditions, the detection threshold η must be automatically adjusted in order to maintain a desired detection performance (P_D vs P_{FA}). To achieve this, we adapt this parameter as data is being collected in the new environment by fitting a parametric distribution to the null distribution of the broadband coherence statistics produced from that data. That is, we assume that the likelihood ratio given in (5) is distributed with some density function $f_{\boldsymbol{\theta}}(\lambda)$ where the function f is known but the parameter vector $\boldsymbol{\theta}$ that represents it is unknown. Given a sequence of likelihood ratio measurements $\{\lambda_i\}_{i=1}^n$ collected by applying (5) to some data from the new environment, one can replace these unknown parameters with their maximum likelihood estimates

$$\widehat{\boldsymbol{\theta}} = \arg\max_{\boldsymbol{\theta}} \prod_{i=1}^{n} f_{\boldsymbol{\theta}}(\lambda_i)$$

Once this parametric fit is formed, the threshold η may then be adjusted such that the detector achieves the desired probability of false alarm $0 < P_{FA} = \alpha < 1$,

$$\int_{\eta}^{\infty} f_{\widehat{\boldsymbol{\theta}}}(\lambda) d\lambda = \alpha \tag{8}$$

The estimated parameter vector $\widehat{\boldsymbol{\theta}}$ may then be recursively updated as new data is being collected to give the detector the ability to adapt to changing statistical behavior in the likelihood ratio. This adaptive process allows the MBCD to be used in environments for which the background composition produces measurements that are drawn from a different distribution than the environment the system was originally trained on.

5.3 Test Results on TIER-SWAT Data Sets

The adaptive MBCD method was used to process both the raw stave data as well as Synthetic Aperture Sonar (SAS) imagery from the TIER-SWAT data set. This data set uses target scattering signals generated using APL-UW's target-in-environment-response (TIER) model [23] that has been embedded in environments generated using NSWC-PCD's PC-SWAT tool [24]. Two sonar frequency band were used

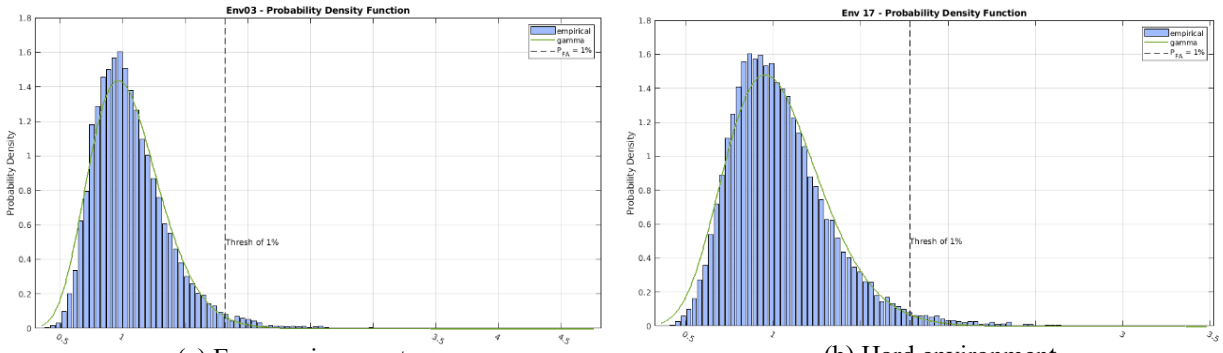

(a) Easy environment (b) Hard environment

Figure 1: Empirical distributions for test statistic in (5).

with the low frequency (LF) band covering 2-10 kHz and the high frequency (HF) band covering 10-30 kHz. The simulated runs were generated for a linear path with vehicle altitude at 5m off the bottom and for 13 different objects (7 UXOs and 6 non-UXOs) which were laid out proud on the bottom at various ranges and orientation angles. A list of the target types can be seen in Table 1. Data were generated for 17 different simulated environments ranging from Clay to Rock. The data of only two environments were used corresponding to the sandy slit (easy) and rocky (relatively difficult) bottom conditions

Table 1: Object types available in the TIER-SWAT data set.

Object Types	**Target (T) or Non-Target (NT)**
Artificial Clutter Blob (Point Scatterers)	NT
55 Gallon Drum	NT
Scuba Tank, no stem	NT
2:1 Alum Pipe	NT
2:1 Alum Cylinder	NT
3:1 Alum Cylinder	NT
Alum UXO replica	T
Steel UXO replica	T
105 mm Bullet, air-filled	T
105 mm Bullet, water-filled	T
155 mm Howitzer, no endcap	T
155 mm Howitzer, w/ endcap, air-filled	T
155 mm Howitizer, w/ endcap, water-filled	T

The detection process was performed using both the original and adaptive MBCD. In the case of the original MBCD, the detection process was performed on both the easy and hard environments with the same $P_{FA} = \alpha$ determined by the easy environment. For the adaptive MBCD, the probability density function $f_{\boldsymbol{\theta}}(\lambda)$ of the test statistic in (5) under the null hypothesis was estimated to be from the exponential family, and specifically gamma distributed. The empirical PDF of each environment for the SAS images can be seen in Figure 1. The threshold was then chosen to provide $P_{FA} = 0.01$. The threshold value can be seen as the vertical dashed line in Figure 1.

As mentioned before, two versions of the detector were implemented- one using the actual stave data and the other using the SAS images (omega-k beamformer) of the two sonar channels. For the former case,

the complex-valued data of each ping captured by a single hydrophone of the LF and HF sonar **channels** were segmented into several range gates each composed of N samples (range cells). That is, each channel **vector** consists of N elements for every range gate of a sonar return at a particular ping. To compute the cross-power spectrum in (7) **averaging** is performed over P contiguous pings. Thus, MBCD test statistic is generated for every pair of range gates of the LF and HF sonar data and at every ping. For the latter case, SAS images of LF and HF sonar **channels** were partitioned into regions of interest (ROI) which are in turn partitioned into blocks (in cross-track and along-track). Thus, each channel **vector** is the vector-arrangement of each block. For this case, averaging in (7) takes place over all P blocks within every pair of ROIs. For this case, MBCD test statistic is generated for every pair of ROIs in the LF and HF sonar images.

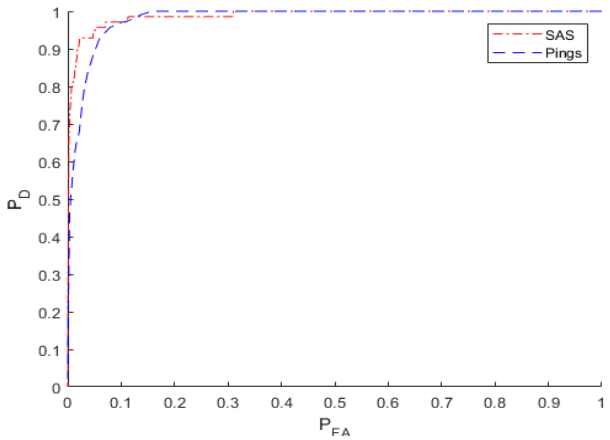

Figure 2: ROC curves of adaptive MBCD on TIER-SWAT data-Rocky bottom.

Figure 2 gives the receiver operating characteristic (ROC) curves of the two adaptive detectors when tested on all the 9 runs of rocky environments and for both methods of implementation. The SAS-based MBCD detector provided probability of detection P_D= 91% at an average of 8.8 false detections per image. Some examples of the LF and HF SAS images and their associated detections superimposed on the likelihood maps can be seen in Figure 3 for 80º orientation (near end-on insonification). These results are shown for both the non-adaptive (top right) and adaptive (bottom right) MBCD methods. Note that 0 degree orientation corresponds to broadside view of the object while 90 degree corresponds to an end-on view of objects. The rectangular boxes are the detection regions, and the circles represent the placement of the different objects. Each white rectangular box in the image is a false alarm, while green and yellow colored boxes and circles are correct detections. Red and black circles signify missed detected objects. It was noted that missed detections were primarily caused by lack of any useful target signature in one or both sonar channels at certain orientation and particularly at short ranges. In general, the adaptive version of the MBCD performed better (both P_D and P_{FA}) than the non-adaptive version with the former improving the probability of detection while reducing the number of false alarms per image. Additionally, it was found that the detection performance of the MBCD and the adaptive MBCD on the SAS images is better than that of the raw stave data.

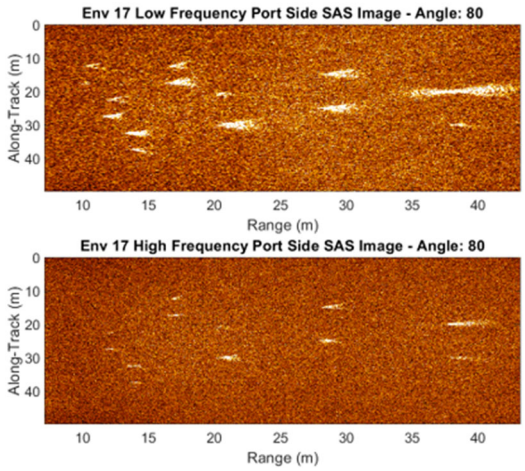

(a) SAS image for 80° object orientation

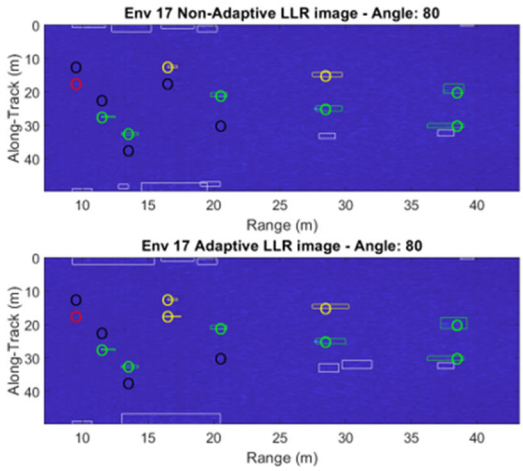

(b) SAS based detections for 80° object orientation

Figure 3: SAS imagery and MBCD detections for 80object orientation in hard environment. Detection regions and target locations are overlaid on the LLR images.

6 Task 2: Incremental Learning of Matched Subspace Classifier

6.1 Review of Matched Subspace Classifier

In this section, we first give a brief review of the theory behind the Matched Subspace Classifier (MSC) and a simple modification of this classifier which allows for utilization of a dictionary learned via the K-SVD method [25]. This classifier operates on the aspect-dependent spectral feature vectors with the assumption that such vectors, each belonging to a given class of objects, can be accurately represented using a linear combination of a small subset of basis vectors associated with that class.

We consider a general M-ary classification problem in which the observations (feature vectors) come from $m = 0, \dots, M-1$ possible classes satisfying the signal model,

$$\mathbf{Z} = \mathbf{H}_m\mathbf{X} + \mathbf{N} \quad m \in [0, M-1] \tag{9}$$

where $\mathbf{Z} = [\mathbf{z_1}\ \mathbf{z_2} \dots \mathbf{z}_K]$ is the observation matrix of size $N \times K$ containing spectral feature vectors $\mathbf{z}_i \in \mathbb{R}^N$, $\mathbf{H}_m \in \mathbb{R}^{N \times L}$ *is* the dictionary (or subspace) matrix whose columns are the basis vectors, $\mathbf{h}_j s$ that span the subspace associated with the m^{th} object class, $\mathbf{X} = [\mathbf{x_1}\ \mathbf{x_2} \dots \mathbf{x}_K]$ is an unknown parameter matrix with $\mathbf{x}_i \in \mathbb{R}^L$ being the parameter vector associated with $\boldsymbol{z}_i$ and $\mathbf{N}$ denotes an additive zero-mean noise matrix that represents the inaccuracy in this representation.

The subspace matrix $\mathbf{H}m$ can be constructed utilizing the training data representing the objects belonging to class m. Once these matrices are formed, the decision about the class membership of a particular unknown observation vector (e.g., aspect-dependent spectral feature vector) $\boldsymbol{z}_k$ can be made using,

$$m^* = \arg \min_{m \in [0, M-1]} \{\|\mathbf{z}_k - \mathbf{H}_m \hat{\mathbf{x}}_k\|^2\} \tag{10}$$

where $\hat{\mathbf{x}}_k$ represents an estimate of the actual parameter vector $\mathbf{x}_k$. This decision rule assumes that the reconstruction error takes its smallest magnitude for the correct class m^*. Thus, $J_m = \|\mathbf{z}_k - \mathbf{H}_m \hat{\mathbf{x}}_k\|^2$ serves as the *discriminant function* for this classifier.

For making decision on several (P) aspect-dependent spectral feature vectors forming data matrix $\mathbf{Z} = [\mathbf{z}_1\ \mathbf{z}_2\ ...\ \mathbf{z}_P]$, the expression in (10) becomes,

$$m^* = \arg \min_{m\in[0,M-1]} \{\left\|\mathbf{Z} - \mathbf{H}_m\widehat{\mathbf{X}}\right\|_F^2\}, \tag{11}$$

where $\|\mathbf{A}\|_F^2 = \mathrm{tr}\{\mathbf{A}\mathbf{A}^{\mathrm{T}}\}$ represents the squared Frobenius norm of matrix $\mathbf{A}$.

As mentioned before, training of this classifier amounts to constructing class-dependent dictionary matrices $\mathbf{H}_m$ from representative training data sets in each class m. In this work, we applied the K-SVD [25] signal-specific dictionary learning method to build $\mathbf{H}_m$s using the TIER model-generated data. The details are discussed in [18]. This dictionary learning algorithm uses sparse coding methods such as Orthogonal Matching Pursuit (OMP) or Basis Pursuit [26] to reduce the data into sparse vectors as columns of $\widehat{\mathbf{X}}$ that contain only a few nonzero elements. Since the estimates generated using these sparse coding methods differ from those of the standard Least Squares (LS) solution used to develop MSC decision rule in (11), we use the modified the decision rule,

$$m^* = \arg \min_{m\in[0,M-1]} \{\left\|\mathbf{Z} - \mathbf{H}_m\widehat{\mathbf{X}}_{\mathrm{OMP}}^{(m)}\right\|_F^2\}, \tag{12}$$

where $\widehat{\mathbf{X}}_{\mathrm{OMP}}^{(\boldsymbol{m})}$ is the estimate of $\mathbf{X}$ generated using the OMP algorithm when dictionary $\mathbf{H}_m$ is used. For a two-class discrimination (M = 2) based upon P aspects/pings, (12) can be done using,

$$\eta(\mathbf{Z}) = \frac{\left\|\mathbf{Z} - \mathbf{H}_0\widehat{\mathbf{X}}_{\mathrm{OMP}}^{(0)}\right\|_F^2}{\left\|\mathbf{Z} - \mathbf{H}_1\widehat{\mathbf{X}}_{\mathrm{OMP}}^{(1)}\right\|_F^2} \begin{matrix} \mathrm{UXO} \\ \gtrless \\ \mathrm{Non\text{-}UXO} \end{matrix} \gamma. \tag{13}$$

Aside from its simplicity in structure, training, and decision-making, the proposed classification framework offers many other benefits including: ease in performing multi-class classification, certain invariance properties, and incremental in-situ training in new environments [18],[20] without jeopardizing the prior training, i.e., offers flexibility without sacrificing the stability. The latter is discussed next.

6.2 A New Incremental Learning for UXO Classification in Changing Environments

Acoustic color features for a specific target vary significantly depending on the object's burial condition and orientation, range, and sonar grazing angle as well as seafloor properties and roughness. Hence, it is unrealistic to expect that training on model-generated or even real data will produce a classifier that can capture all such variations when operating conditions differ from those of the prior training. To address this important issue, we developed [18], [20] a new incremental learning algorithm that updates the previously trained classifier using a limited set of sonar returns from the detected objects in the new environment. *The system offers flexibility to learn incrementally using small batches of data in the new environment without impairing the prior learning even in the presence of new munition types.* Here, we briefly review this novel dictionary learning algorithm that allows for incremental updating of the dictionary matrices when new unseen data samples are encountered without requiring to retrain the entire system.

Let us assume that the set of dictionary atoms has already been constructed using an *old* training data set $\{\boldsymbol{z}_i\}_{i=1}^{K}$ and for each $\boldsymbol{z}_i$ we have generated, using fast OMP [16], a compact dictionary matrix $\overline{H}_i^{old}$ based upon these dictionary atoms. Upon arriving *new* observations $\{z_i\}_{i=K+1}^{K+\bar{K}}$, we first examine if they can be

successfully (wrt sparsity and minimum error) represented using a small subset the old dictionary atoms $\mathbf{h}_l, l \in [1, L]$. Otherwise, new dictionary atom, $\mathbf{h}_{L+1}$, needs to be created and added one-by-one to the column space of $\bar{H}_i^{old}$ to generate new dictionary matrix $\bar{H}_i^{new} = [\bar{H}_i^{\boldsymbol{old}} \mathbf{h}_{L+1}]$ without retraining on the entire augmented data set that includes the new sample. Also, we would like to do this by interfacing this recursive updating with the fast OMP algorithm in [27] for computational efficiency. This problem can then be posed as a constrained minimization,

$$\min_{\mathbf{h}_{L+1}} \left\{ \sum_{i=K+1}^{K+\bar{K}} \left\| \mathbf{z}_i - \sum_{j \in S_i^{new}} \mathbf{h}_j x_{i,j} \right\|^2 \right\} \text{ s.t. } \|\mathbf{h}_{L+1}\|^2 = 1 \tag{14}$$

where the set of indices $S_i^{new} \subset \{1, \ldots, L+1\}$ and the constraint is added to impose the unit norm requirement for $\mathbf{h}_{L+1}$ [12]. Here, we assume that all the previously learned atoms $\mathbf{h}_l, l \in [1, L]$ are unchanged from their previously estimated values while the new atom $\mathbf{h}_{L+1}$ and associated coefficients $x_{i,j}, i \in [K+1;\ K+\bar{K}],\ j = L+1$ need to be recursively determined in conjunction with the fast OMP. Considering these assumptions, the Lagrangian associated with optimization problem in (14) is,

$$\mathcal{L}(\mathbf{h}_{L+1}, \lambda) = \sum_{i}^{K} \left\| \mathbf{r}_{i,L}^{new} - \mathbf{h}_{L+1} x_{i,L+1} \right\|^2 + \lambda(\mathbf{h}_{L+1}^{\top} \mathbf{h}_{L+1} - 1) \tag{15}$$

where $\mathbf{r}_{i,L}^{new} = \mathbf{z}_i - \sum_{\boldsymbol{j} \in S_i^{old}} \mathbf{h}_j x_{i,j}$ is the residual error component corresponding to the old compact dictionary matrix $\bar{H}_i^{\boldsymbol{old}}$ but new coefficients $x_{i,j}, j \in S_i^{old}$ obtained using fast OMP and S_i^{old} contains the same indices as in S_i^{new} except $L+1$. Taking the partial derivatives of (15) wrt $x_{i,L+1}$ and $\mathbf{h}_{L+1}$, setting the results to zero and solving for all unknowns give the following solutions,

$$\boldsymbol{x}_{L+1} = \mathbf{R}^{\top} \mathbf{h}_{L+1} \tag{16}$$

and

$$\mathbf{h}_{L+1} = \frac{\mathbf{R}\mathbf{x}_{L+1}}{\sqrt{\mathbf{x}_{L+1}^{\top} \mathbf{R}^{\top} \mathbf{R} \mathbf{x}_{L+1}}} \tag{17}$$

where $\mathbf{R} = [\mathbf{r}_{K+1,L}^{new} \ldots \mathbf{r}_{K+\bar{K},L}^{new}]$ and $\mathbf{x}_{L+1} = \left[x_{K+1,L+1} \ldots\ x_{K+\bar{K},L+1}\right]^{\top}$.

The coupled equations in (16) and (17) must be iterated until convergence. To start the process, we initialize $\mathbf{h}_{L+1}$ and use fast OMP [27] for every $\mathbf{z}_i, i \in [K+1, K+\bar{K}]$ to compute the entire matrix $\boldsymbol{X}$ after adding $\mathbf{h}_{L+1}$. After this initialization, $\mathbf{h}_{L+1}$ and $\mathbf{x}_{L+1}$ can be updated for the first time using (16) and (17). The solutions of these equations provide a rank-1 approximation of the error matrix $\mathbf{R}$. That is, if we let $\mathbf{R} = \mathbf{u}_1 \lambda_1 \mathbf{v}_1^{\top}$, then its is easy to verify using (16) and (17) that at convergence $\mathbf{h}_{L+1} \rightarrow \mathbf{u}_1$ and $\mathbf{x}_{L+1} \rightarrow \lambda_1 \mathbf{v}_1$.

Now, the remaining issue is to devise a criterion for creating new atom $\mathbf{h}_{L+1}$ based upon its importance to the data representation. More specifically, we choose to create new $\mathbf{h}_{L+1}$ if its addition reduces the residual error. That is, $\mathbf{h}_{L+1}$ is useful if the residual errors before and after updating satisfy,

$$E_i[\left\| \mathbf{r}_i^{old} - \mathbf{r}_i^{new} \right\|^2] > \delta \tag{18}$$

where δ is a positive tolerance quantity and $E_i[\cdot]$ is the *time-average* over all the data samples $\boldsymbol{z}_i, i \in [K+1, K+\bar{K}]$. It can easily be shown [20], that (18) becomes,

$$E_i\left[x_{i,L+1}^2\right] > \delta / \sin^2 \psi_{L+1} \tag{19}$$

where ψ_{L+1} is the principal angle between subspaces $\boldsymbol{h}_{L+1}$ and $\langle \bar{H}_i^{old} \rangle$. That is, when principal angle between these two subspaces is large (i.e. less mutual coherence) the lower bound on the Mean Squared Error (MSE) of the associated coefficients $E_i[x_{i,L+1}^2]$ reduces comparing to the case when ψ_{L+1} is small. Thus, the dictionary atom $\boldsymbol{h}_{L+1}$ will be added only if the MSE of the associated with coefficients $x_{i,L+1}$ satisfies (19).

6.3 Training and Testing Data Sets for UXO Classification

6.3.1 Baseline Training Data Set and Procedure

In this study, we used the model-generated aspect-dependent spectral feature vectors (acoustic color) for several UXO and non-UXO objects with known geometrical and physical characteristics at different ranges and orientations to construct dictionary matrices suitable for baseline training the MSC-based classification system. The TIER model [16] was utilized to create acoustic color matrices for different environments and simulated runs. Using the procedures described in [18] simulated sonar runs for 8 different objects (six UXO and two non-UXO) were generated. These simulated runs were designed to replicate the conditions of the experimental TREX13 data collection. In particular, these runs were generated for a 40 m linear SAS path length at ranges of 10, 15, 20, 25, 30, 35, 40 m to the target and with a sonar interface elevation of 3.8 m. Additionally, object orientations ranged from -80° to +80° in 20° increments, with 0° orientation corresponding to broadside view of the object and 90° corresponding to an end-on view of objects. These synthetic sonar data sets and their corresponding acoustic color matrices were generated for two different environments with sound speed matching those conditions in TREX13 experiments. In particular, one environment used the sediment sound speed of sand and the other used that of a slightly denser material simulating a mixture of sand and silt or clay.

For each object, the aspect-dependent spectral features in the associated acoustic color matrix were decimated along the frequency dimension to have N = 271 frequency bins spanning the 3-30 kHz frequency range (i.e. 100 Hz separation between frequency bins). However, along the aspect dimension the original ping separation was used resulting in aspect resolution that changed depending on target range. The training data matrix, $\mathbf{Z}_{m,n}$, for object n of class m will then contain all the model-generated aspect-dependent spectral features as its columns for all the above-mentioned ranges and orientations of the UXO and non-UXO objects.

6.3.2 TREX13 Testing Data Set and Procedure

The TREX13 target field contained 8 different UXO and non-UXO objects (same as TIER data) with varying shapes, sizes, and compositions, all of which were located between 10 m to 40 m horizontally from the rail system and are proud on the bottom. The sonar system was mounted on a mobile tower to minimize platform motion as the sonar tower traversed along the rail. The length of the rail was approximately 40 m. The sonar transmitted a 6 ms Linear Frequency Modulated (LFM) signal over 3-30 kHz with a 10% taper between the leading and trailing edges to minimize ringing effects. Sonar backscatter was received by a 6-element linear array. For the formation of acoustic color matrices, data from only the third hydrophone element of this array was used.

There were a total of nine runs through the target field in which the physical orientation of all the objects in the scene differed with each run, with each object having the same orientation for a given run. The

object orientations varied from -80° to $+80^\circ$ in 20° increments. Each run consisted of 1600 pings in which the sonar platform moved along the fixed rail in increments of 0.025 m, transmitting and receiving once for each sonar position. The data was sampled at 100 kHz and the sonar platform was tilted at either a 10° or 20° grazing angle depending on range to the targets being observed (angle of the sonar main response axis with respect to the horizontal plane).

6.4 Results of Incremental Learning

As mentioned before, one of the important benefits of the MSC-based classifier is that the dictionary matrices, $\mathbf{H}_m$s, can be updated incrementally to incorporate some limited labeled data drawn from the new environment in which the sonar is operating. This incremental learning can be accomplished without the need to carry over the previous training data sets. In this subsection, we adopted our newly developed algorithm [20] that keeps the old dictionaries trained using the TIER model-generated data for UXO and non-UXO classes intact while augmenting them with the incremental dictionary matrices that are learned using less than 7% of the randomly drawn TREX13 samples. We allowed adding an incremental dictionary with only five columns (i.e. $\tilde{L} = 5$) while keeping the sparsity factor $\tau = 15$ as before. It was experimentally determined that adding five dictionary atoms is adequate for this particular application.

Figure 4 shows the ROC curve (dashed-line) for the incrementally trained system using the proposed algorithm in Section 6.2 when baseline trained on all the model-generated data, incrementally updated on less than 10% of the randomly selected TREX13 data sets, and tested on the remaining TREX13 data set i.e. for all target ranges and orientations. For these results, we used $P = 7$ aspects/pings with some separation [18]. Comparing this ROC curve with that of the MSC-based system without incremental learning that was presented in [18] reveals significant improvement (> 10%) in the correct classification performance as a consequence of this incremental learning.

To further examine the generalization ability of the MSC-based classifier after incremental training on unseen objects, a separate experiment was conducted where all samples from the two actual UXO objects were left out of the incremental training set. However, all samples (at all ranges and orientations) of these unseen objects were included in the testing data set to evaluate the classification performance using the incrementally updated dictionaries. Figure 4 also illustrates the ROC curve (solid-line) of the MSC classifier for this experiment. Comparing to the ROC curve associated with the no holdout case (dashed-line), the degradation in performance as a result of excluding two UXO objects during the incremental training is found to be less than 2%. These results demonstrate the generalization ability of the MSC classifier when incrementally trained using only a subset of UXO objects found in the TREX13 operating environment. Moreover, they show the flexibility and ease of the MSC-based system for adaptive UXO classification applications.

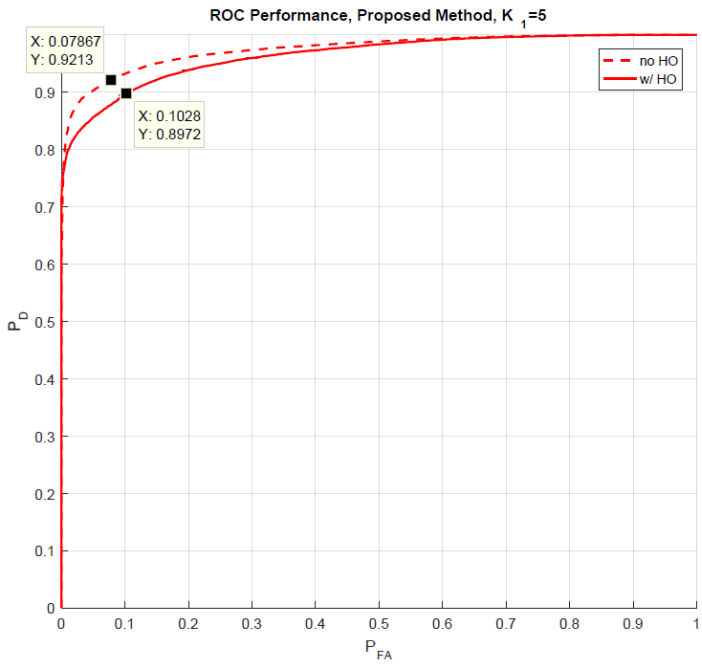

Figure 4: ROC curves for incrementally trained MSC using our new method with and without UXO object holdout.

7 Task 3: Non-linear Extensions of Matched Subspace Classifier

Kernel machines have extensively been studied for many complex pattern recognition and machine learning applications [28]. This is due to the fact that certain non-linearly separable classification problems can easily be converted to linearly separable ones by mapping the data to a higher dimensional feature space using kernel-producing nonlinear mapping functions. However, most kernel machines, e.g., support vector machines (SVM) or relevance vector machines (RVM) [15] don't offer the flexibility to learn incrementally using a small batch of data encountered during the actual operation. As illustrated in Section 6.4 results, incremental learning of the MSC system not only provides substantial improvements in classification performance as compared to the fixed classifier but also offers much more efficient learning since retraining the entire system requires carrying a growingly large data set which would in turn require substantial training time and resources especially when novel threats are becoming available in small batches. In our investigations we began laying the groundwork for an expandable kernel MSC-type classifier which can be incrementally trained in a manner similar to its linear counterpart in Section 6.

In an attempt to design a kernel-based classification system capable of providing higher probability of correct classification while remaining amenable to incremental learning strategies, we have investigated several possible extensions of our MSC classifier [18]. Such extensions would rely on kernelizing not only the MSC formulations but also the dictionary learning and sparse coding [29] used in the training phase of this method. A version of such kernel-based MSC system was then developed and tested on the TREX13 data sets. Owing to the fact that the classification results of this kernel MSC classifier were found to be slightly worse than those of the original MSC counterpart, we decided to look for an alternative approach that takes advantage of the full benefits of nonlinear kernel machines while allowing any linear dictionary learning algorithm to be used. To this end, the method in [30] which is referred to as linearized kernel dictionary learning (LKDL) was adopted. The LKDL method allows for linear dictionary learning methods to be applied to the non-linearly mapped virtual representatives of the data

samples without increasing the dimensionality of the data representation. That is, discriminative benefits of kernel machines can be gained without the need for the full kernel matrix in representing the dictionary. Additionally, the inherent sampling in this method significantly reduces the number of useful samples required to generate highly discriminative kernelized features. Owing to all these benefits, the results in this section are concentrated on the application and testing of this method.

7.1 Review of Linearized Kernel Dictionary Learning Method

The LKDL method introduces a pre-processing technique [30] to transform the training and testing samples into "virtual samples" which not only takes benefit from the inherent non-linear kernel mapping of the data but also allows for the use of the standard linear dictionary learning methods. This implies that unlike the other kernel-based subspace classifiers [29] that rely on nonlinear extension of the dictionary learning, this method works with the original dictionary learning methods hence can easily be incrementally trained. A brief overview of the LKDL technique is given below.

The basic idea in any kernel-based classification system is that if data is not linearly separable in the original data domain $\mathcal{Z}$, there could be a nonlinear mapping $\Phi(\cdot) : \mathcal{Z} \rightarrow \mathcal{F}$ which maps data in $\mathcal{Z}$ to points in $\mathcal{F}$, which are of much higher dimensionality when compared to the original signal domain, and in this mapped space, the mapped samples become linearly separable.

Let $\mathbf{Z} \in \mathbb{R}^{N \times K}$ be the data matrix defined as before and $\mathbf{K} = \Phi(\mathbf{Z})^{\top}\Phi(\mathbf{Z})$ be the associated kernel Gram matrix of the nonlinearly mapped data with elements $\mathbf{K}_{i,j} = \kappa(\mathbf{z}_i, \mathbf{z}_j) = \Phi(\mathbf{z}_i)^{\top}\Phi(\mathbf{z}_j)$ where $\kappa(\mathbf{z}_i, \mathbf{z}_j)$ is an appropriate kernel function [28]. Since $\Phi(\mathbf{z}_i)$ is of high or potentially infinite dimension (Gaussian kernel), the goal here is to come up with equivalent low dimensional inner product factorization of the kernel matrix i.e. $\mathbf{K} = \mathbf{F}^{\top}\mathbf{F}$ that possesses the same properties as the nonlinearly mapped version using $\Phi(\cdot)$. The columns, $\mathbf{f}_i$s, of matrix $\mathbf{F}$ then represent the *virtual data samples* which could be used in conjunction with any linear dictionary learning method e.g. K-SVD and hence in the MSC framework of Section 6. However, the problem that arises is that if the training data set contains a large number of samples (such as in the TIER model-generated dataset used in this study), the dimension of the virtual samples $\mathbf{f}_i$s will also be large. To overcome this problem, the authors in [30] used the *Nyström Method* in conjunction with an appropriate sampling method [31] to reduce the number of samples used to approximate matrix $\mathbf{K}$. This process goes like this.

Let $\mathbf{K} \in \mathbb{R}^{K \times K}$ be a symmetric positive semi-definite (PSD) kernel matrix and assume $c \leq K$ columns of $\mathbf{K}$ are sampled according to some sampling method (e.g., uniform) to form matrix $\mathbf{C} \in \mathbb{R}^{K \times c}$. Then we can write, $\mathbf{C} = \begin{bmatrix} \mathbf{W} \\ \mathbf{S} \end{bmatrix}$ and $\mathbf{K} = \begin{bmatrix} \mathbf{W} & \mathbf{S}^{\top} \\ \mathbf{S} & \mathbf{B} \end{bmatrix}$, where $\mathbf{W} \in \mathbb{R}^{c \times c}$ is the kernel matrix associated with the retained samples while $\mathbf{B} \in \mathbb{R}^{(K-c)\times(K-c)}$ is the one associated with discarded (as a result of sampling) samples. An approximated $\mathbf{K}$ can then be constructed using,

$$\mathbf{K} \approx \mathbf{C}\mathbf{W}^{+}\mathbf{C}^{\top} \tag{20}$$

where $\mathbf{W}^{+}$ is the pseudo-inverse of matrix $\mathbf{W}$ which can be decomposed using $\mathbf{W} = \mathbf{V}\mathbf{\Sigma}\mathbf{V}^{\top}$ where is a diagonal matrix containing all its eigenvalues and $\mathbf{V}$ contains the associated eigenvectors as its columns. Using this eigenvalue decomposition, (20) can be expressed as,

$$\mathbf{K} \approx \mathbf{C}\mathbf{W}^{+}\mathbf{C}^{\top} = \mathbf{C}\mathbf{V}\mathbf{\Sigma}^{+}\mathbf{V}^{\top}\mathbf{C}^{\top} = \left((\mathbf{\Sigma}^{+})^{1/2}\mathbf{V}^{\top}\mathbf{C}^{\top}\right)^{\top}\left((\mathbf{\Sigma}^{+})^{1/2}\,\mathbf{V}^{\top}\mathbf{C}^{\top}\right) = \mathbf{F}^{\top}\mathbf{F}$$

Therefore, matrix $\mathbf{F}$ takes the form,

$$\mathbf{F} = (\mathbf{\Sigma}^{+})^{1/2}\mathbf{V}^{\top}\mathbf{C}^{\top}$$

Choosing the k largest eigenvalues $\mathbf{\Sigma}_k = diag[\sigma_1, \sigma_2, \dots, \sigma_k]$ of $\mathbf{W}$ and the associated eigenvectors $\mathbf{V}_{\mathrm{k}} = [\mathbf{v}_1, \mathbf{v}_1 \dots \mathbf{v}_{\mathrm{k}}]$, the reduced dimensional *virtual sample matrix* can be derived in a similar manner using,

$$\mathbf{F}_{\mathrm{k}} = (\mathbf{\Sigma}_{\mathrm{k}}^{+})^{1/2}\mathbf{V}_{\mathrm{k}}^{\top}\mathbf{C}^{\top}$$

Given the matrices $\mathbf{\Sigma}_k^{+}$ and $\mathbf{V}_{\mathrm{k}}$, any test sample $\mathbf{z}$ can be mapped to the corresponding virtual sample using,

$$\mathbf{f} = (\mathbf{\Sigma}_k^{+})^{1/2}\mathbf{V}_k^{\top}\,[\kappa(\mathbf{z}, \mathbf{z}_1), \dots, \kappa(\mathbf{z}, \mathbf{z}_c)]^{\top}$$

where $\mathbf{z}_1 \dots \mathbf{z}_c$ are the chosen (via sampling) c training samples and $\kappa(\boldsymbol{x}, \boldsymbol{y})$ is the chosen kernel function [28]. The steps in the entire algorithm are given in the table below.

Algorithm 1 LKDL Preprocessing Algorithm Steps.

Input: $\mathbf{Z}_{train} = [\mathbf{z}_1, \dots, \mathbf{z}_K] \in \mathbb{R}^{N\times K}$; $\mathbf{Z}_{test} \in \mathbb{R}^{N\times L}$ where L is the number of testing data samples; kernel function $\kappa(.,.)$; and sampling method [21] to find c.

1. $\mathbf{Z}_R = sub\_sample\,(\mathbf{Z}_{train}, \text{sampling method}, c) \in \mathbb{R}^{N\times c}$
2. Compute $\mathbf{C}_{train} = \mathbf{K}\,(\mathbf{Z}_{train}, \mathbf{Z}_R) \in \mathbb{R}^{K\times c}$
3. Compute $\mathbf{W} = \mathbf{K}\,(\mathbf{Z}_R, \mathbf{Z}_R) \in \mathbb{R}^{c\times c}$
4. Choose k and approximate $\mathbf{W}$ through a thin SVD to get $\mathbf{W}_k = \mathbf{V}_k\mathbf{\Sigma}_k\mathbf{V}_k^{\top} \in \mathbb{R}^{c\times c}$
5. Compute virtual training samples $\mathbf{F}_{train} = \left(\mathbf{\Sigma}_k^{\dagger}\right)^{1/2}\mathbf{V}_k^{\top}\mathbf{C}_{train}^{\top} \in \mathbb{R}^{k\times K}$
6. Compute $\mathbf{C}_{test} = \mathbf{K}\,(\mathbf{Z}_{test}, \mathbf{Z}_R) \in \mathbb{R}^{L\times c}$
7. Compute virtual testing samples $\mathbf{F}_{test} = \left(\mathbf{\Sigma}_k^{\dagger}\right)^{1/2}\mathbf{V}_k^{\top}\mathbf{C}_{test}^{\top} \in \mathbb{R}^{k\times L}$.

7.2 Results of MSC Classification with LKDL Virtual Features

An experiment was conducted using the method described above to test the feasibility of coupling this nonlinear pre-processing with the MSC classification framework in Section 6 used in previous testing of sonar samples. In this experiment, 5440 random samples (i.e. about 10% of total samples) were chosen from the synthesized TIER model-generated training data set and 2161 *in-situ* samples (i.e. about 2.5 % of the total test set) were selected from the TREX13 testing set to augment the training data. A total of c = 1084 samples were then selected (using column-norm sampling) from this augmented data set in the pre-processing stage in order to approximate the kernel matrix via the Nyström method. Thus, the size of the data set in this case is almost 1/8 of that used for the incremental learning in Section 6.4. The virtual sample space was chosen to be of the same dimension as the original samples (i.e. k = 271-dimensional) and a Gaussian kernel with smoothing parameter 0.75 was utilized in the inherent nonlinear mapping. After all training and testing samples were mapped to their virtual sample representatives using the LKDL algorithm outlined previously. The K-SVD dictionary learning was run on the UXO and non-UXO training virtual samples to learn the associated dictionary matrices $\mathbf{H}_0$ and $\mathbf{H}_1$ with 542 and 1626 columns, respectively. The sparsity factor was chosen to be $\tau = 15$ as before. The MSC classification was then performed on the virtual testing samples for $P = 7$ aspects/pings per decision as shown in the

ROC curve of Figure 5 (a). As evident from this result the pairing of the LKDL embedded virtual features with our MSC classification system indeed provided excellent performance on the TREX13 testing set. The system attained a knee-point performance of $P_{CC} = 99.3\%$, $P_{FA} = 0.7\%$ which is the highest performance achieved by the authors on this data set to date. Figure 5 (b), on the other hand, presents the ROC curve of the same system but when two UXO objects were held out from the in-situ data set i.e. they were neither used in the training of the LKDL process nor in the incremental dictionary learning. Comparing to the ROC curve in Figure 5 (a) associated with the no holdout case, the degradation in performance as a result of excluding two UXO objects is found to be less than 2%. Again, these results show the generalization ability of the kernel-based MSC classifier when incrementally trained using only a subset of UXO objects found in the TREX13 operating environment.

Comparison with the ROC curves in Figure 4 using the incremental learning of the dictionary matrices and the original MSC (i.e. without kernel mapping) shows over 7% improvements in correct classification rate in both cases. These results clearly illustrate the importance of nonlinear mapping of the data prior to classification using the MSC framework. However, it should be pointed out that here the augmented (with in-situ data) data set was used both for the batch training of the LKDL embedding and also for the K-SVD dictionary training. That is, no incremental learning based on only the limited in-situ data was done owing to the fact that currently no such mechanism exists for incremental updating of the LKDL embedding. Thus, a natural question that arises is that how the LKDL embedding paired with MSC system would perform when it only allowed to see synthesized TIER model training samples but then updated incrementally using a very limited data captured in the new operating environment? This calls for extending the LKDL framework to perform incremental learning without requiring the old training data sets. We propose to address this issue in future research.

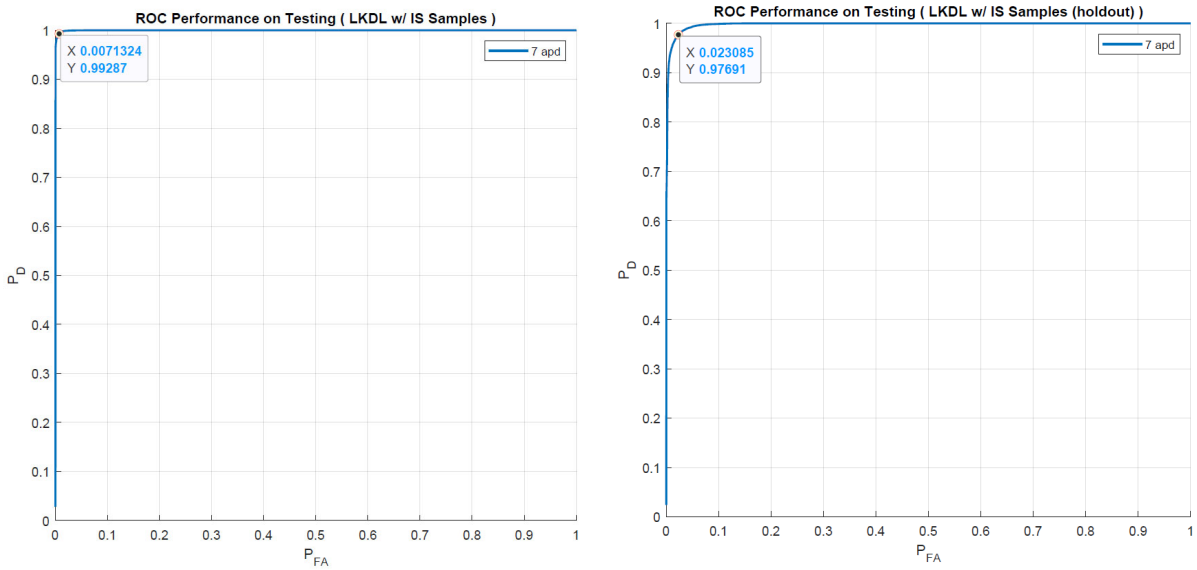

(a) without holdout (b) with holdout

Figure 5: ROC Performance of MSC with LKDL Virtual Features: (a) without hold out; (b) two UXO objects held out.

7.3 Results of MSC Classification with LKDL Virtual Features on Unseen BAYEX14 Data

In order to demonstrate the robustness of systems trained incrementally using both the linear form of the MSC and the LKDL enhanced version, a series of experiments were performed using the BAYEX14 dataset. BAYEX14 was collected in St. Andrew's Bay (Panama City, FL) in 2014, using a rail system similar to the system used in the TREX13 collection. This sonar data set features a wide variety of UXO and non-UXO objects, many of which were also present during the TREX13 data collection. Targets in the BAYEX14 experiment were located between 10 m and 40 m from the rail collection system and all of the objects in this dataset were partially buried, as opposed to the TREX13 set, where all testing and training objects were proud on the seafloor.

An experiment was run using two training schemes of the linear MSC and the LKDL-MSC. First, both systems were configured to use dictionaries that were trained in their baseline configuration on purely TIER-model generated samples. That is, for the linear MSC, a K-SVD dictionary with L=271 columns was trained for each of the 8 objects using only TIER model data. Each of these object-dependent dictionaries were then augmented with 5 more atoms trained using a small portion (~2.5%) of the TREX13 testing set. For this first set of experiments, the LKDL-MSC was trained almost identically however the LKDL embedding was first learned, using the TIER training samples. This embedding was then applied to all training and testing samples before training baseline dictionaries or adding incremental atoms. The second set of experiments was similar to the previous result given for the LKDL-MSC where in addition to the TIER generated data, we also used the incremental TREX13 samples for baseline K-SVD training. This implies that when learning the LKDL embedding, TREX13 samples could be chosen as critical samples (i.e. those samples in $\mathbf{Z}_R$ from Algorithm 1) for the LKDL embedding. For both sets of experiments, the classifiers were then tested on the BAYEX14 data set consisting of a total of 109782 samples using corresponding to the same 8 objects used in training.

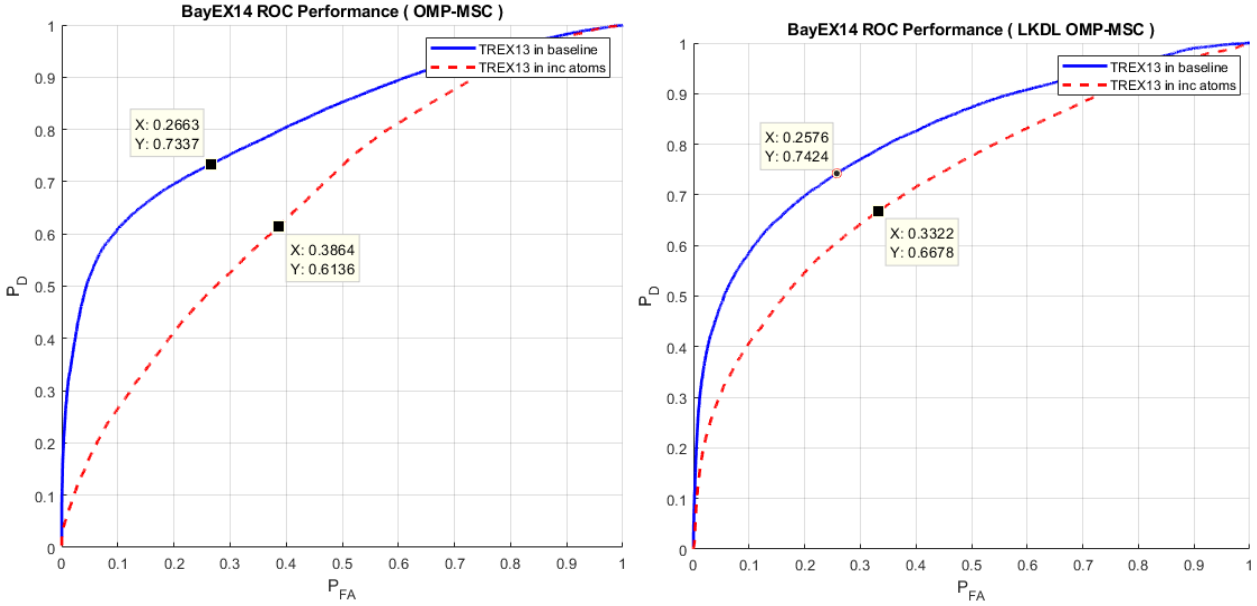

(a) without LKDL embedding (b) with LKDL embedding

Figure 6: ROC Performance of MSC in two training configurations on BAYEX14 test samples.

Comparison between ROC curves in Figures 6 (a) and (b) reveal that under both training schemes, the LKDL enhanced classifier manages to achieve superior performance on the unseen BAYEX14 testing samples. From these results it is clear that the systems have an advantage when they are able to utilize in-situ samples from TREX13 data set during the baseline training. When atoms are incrementally learned after baseline training is performed solely on model data, the linear classifier struggled to generate discriminative statistics while the LKDL enhanced system managed to correctly discriminate approximately 66% of the unseen testing samples on the BAYEX14 set. When TREX13 samples were available for use in baseline training (i.e. solid blue curves in Figures 6 (a) and (b)), both systems greatly improved in performance. These results clearly illustrate the importance of nonlinear mapping of the data prior to classification using the MSC framework.

Training Scheme	**AUC**	P_{CC}	P_{FA}
TREX13 Incremental	0.6708	61.36%	38.64%
TREX13 In Baseline	0.8160	73.37%	26.63%

Table 2: AUC and Knee-Point Performance, No Embedding

Training Scheme	**AUC**	P_{CC}	P_{FA}
TREX13 Incremental	0.7291	66.78%	33.22%
TREX13 In Baseline	0.822	74.24%	25.76%

Table 3: AUC and Knee-Point Performance, With Embedding

The classification performance on the BAYEX14 attained by the two classifiers in two training schemes is presented in Figures 6 (a) and (b) and Tables 2 and 3. Statistics from (13) were used to generate the ROC curves for the modified MSC systems with P=7 aspects per decision. The resulting ROCs for the first training scheme (incrementally training atoms) had knee-point performances of $P_{CC} = 61.36\%, P_{FA} = 38.64\%$ and $P_{CC} = 66.78\%, P_{FA} = 33.22\%$ for real and virtual samples respectively. While using the second training scheme with limited TREX13 samples added to the baseline training, the modified MSC systems achieved knee-point performance of $P_{CC} = 73.37\%,\ P_{FA} = 26.63\%$ and $P_{CC} = 74.24\%,\ P_{FA} = 25.76\%$ for real and virtual samples, respectively. The knee-point performance and the area under the curve (AUC) for each dataset was evaluated using the averaged ROCs of 30 Monte Carlo trials. Tables 2 and 3 list the knee-point probability of correct classification and probability of false alarm along with AUC without and with the embedding, respectively.

Although the training data featured no samples from the BAYEX14 dataset, the classifiers in both scenarios still managed to correctly classify a great number of the samples. This is particularly impressive because all of the objects that were present in the training data sets corresponded to objects proud on the seafloor. However, all of the objects in the BAYEX14 testing dataset were partially buried. Although the performance was comparable, these results indicate that the MSC does indeed benefit from the non-linear embedding of all training and testing samples and that these benefits can extend even into completely unseen samples.

To further demonstrate the utility of the non-linear LKDL mapping, an additional test was run using the BAYEX14 testing set. In this experiment, the LKDL enhanced MSC was trained using ONLY TIER model generated samples and was subsequently tested on the full BAYEX14 testing set. As with previous experiments, the same rule was used for decision making and P=7 aspects were used for every decision. Figure 7 demonstrates the ROC performance of this classifier that was trained on purely TIER-generated samples using identical LKDL embedding parameters to previous experiments. When testing on the full

BAYEX14 set, this purely-model trained modified MSC achieved a knee point performance of $P_{CC} = 67\%$, $P_{FA} = 33\%$ with an AUC of 0.7294.

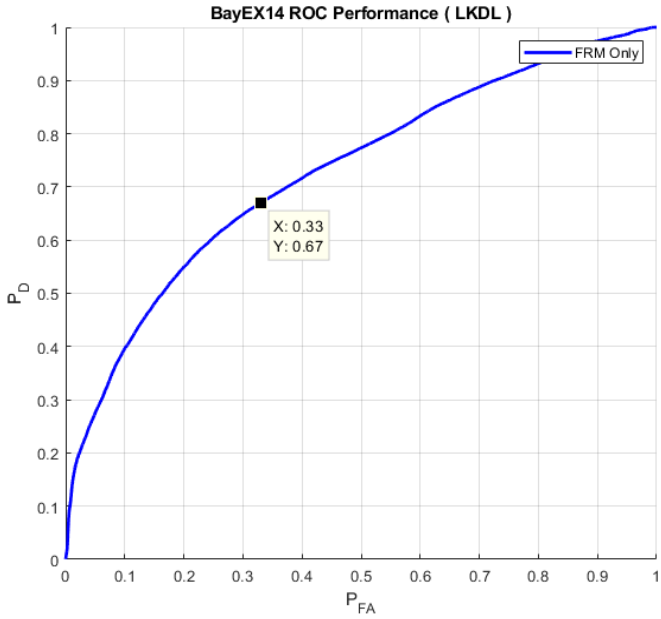

Figure 7: ROC Performance of TIER-only trained LKDL Enhanced MSC on BAYEX14 test samples.

The results presented on unseen real sonar data set suggest that there is consistent improvement that can be gained without the need for a full kernel solution that no longer requires the storage of all previous training samples, a critical constraint of lifelong learning systems. Furthermore, these benefits can extend into new data sets that have never been encountered in training. In future work, we hope to explore the sensitivity of this embedding technique to parameters that exist such as feature dimension and choice of kernel function and kernel function parameters. Furthermore, it should be pointed out that in the second training scheme, the augmented (with in-situ data) data set was used both for the batch training of the LKDL embedding and also for the K-SVD dictionary training of both methods. That is, no incremental learning based on only the limited in-situ data was done owing to the fact that currently no such mechanism exists for incremental updating of the LKDL embedding. Thus, a natural question that arises is how the LKDL embedding paired with MSC system would perform when it only allowed to see synthesized TIER model training samples but then updated incrementally using a very limited data captured in the new operating environment? This calls for extending the LKDL framework, in a natural way, to perform incremental learning without requiring the old training data sets. We propose to address this issue in future research.

8 Task 4: Multi-Aspect Classification Fusion Using CMAC

Multi-aspect processing yields substantial improvements in classification performance, resolution, and sensing of the physical properties of an underwater object in a non-isotropic environment. However, depending on such factors as clutter or natural features in vicinity of a target, target size and position, poor grazing angle, platform instability, etc. it is highly likely that only a few useful pings/aspects may be available for decision-making. Thus, the idea of this particular task is to devise an optimal multi-aspect classification method to generate confusion free UXO classification decisions using a very small number of aspect/look angles.

There are many different mechanisms for performing multi-aspect/look decision-making for underwater

target classification. These mechanisms are based upon either: (a) a decision-level multi-aspect fusion [32] which linearly or non-linearly combines the individual classification decisions made at several aspects, or (b) a feature-level multi-aspect fusion [11], [33] to generate a decision based upon observing a sequence of feature vectors at various aspects with certain separations, or (c) a collaborative decision-making process [14], which can be viewed as a combination of the feature-level and decision-level fusion methods with collaboration among multiple decision making subsystems. The multi-aspect fusion method in [14] offers an important advantage over those in [11],[32],[33] in that the number of aspects used to generate multi-aspect decisions may be variable after the initial training owing to the modular nature of this system. This is indeed an important benefit of this system since depending on the target type, size, and orientation, nearby clutter, etc. the number of aspects/pings needed for confusion-free decision-making can vary. Due to this benefit, here we adopted this framework to develop and test a collaborative multi-aspect classifier (CMAC) using the decision-feedback principle. The use of decision-feedback has widely been exploited in high-performance digital communication and equalization systems to reduce the number of channel states and improve the overall decision-making performance.

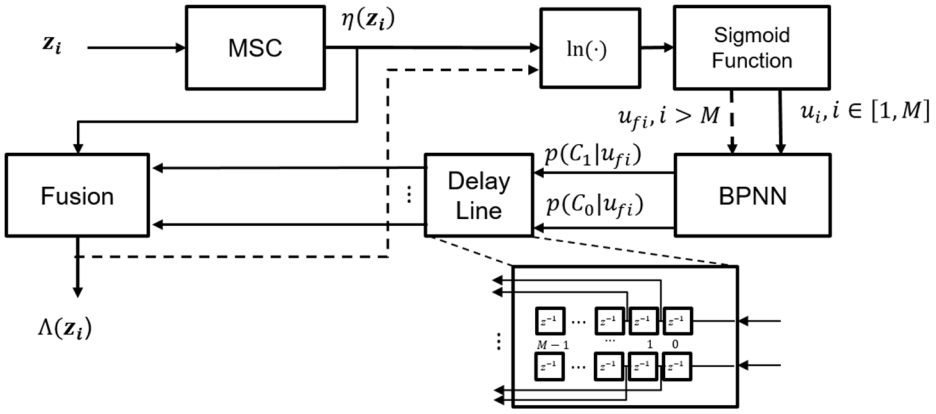

Figure 8: CMAC with decision feedback.

8.1 Multi-Aspect Classification Fusion Using CMAC with Decision Feedback

CMAC with decision-feedback is an optimal multi-aspect fusion system that minimizes the cumulative cost incurred in decision-making. Figure 8 shows the block diagram of a CMAC with decision-feedback system which uses: (a) a trained MSC subsystem to make feature-level decisions based upon the aspect-dependent spectral (acoustic color) feature vectors, $\boldsymbol{z}_i$; (b) a two-layer back-propagation neural network (BPNN) which makes decision-level based classification; and (c) a decision-feedback fusion mechanism which combines the decisions of these two subsystems for several aspects to yield a final decision. The MSC was used here in order to allow a more direct comparison between CMAC decision fusion and the standard MSC-based multi-aspect classification fusion in Section 6.1. A tapped-delay line (or accumulator) is used to keep the previous decisions for the decision-feedback process which is implemented sequentially given a sequence of aspects or sonar pings.

The process goes like this. The MSC subsystem first makes a preliminary decision on the class of a detected object at the current ping/aspect i i.e. it produces test statistic $\eta(\boldsymbol{z}_i)$ using (13). The BPNN subsystem, on the other hand, generates posterior probability $p(C_k|u_{f,i-j})$ of class $C_k, k = 0,1$ given the previous decisions $u_{f,i-j},\ j \in [1, M]$. That is, the BPNN is trained to find the confidence in the final decisions. These decisions are accumulated in the tapped-delay line for M previous aspects/pings for the

fusion system. A final decision, u_{fi} is then made in the fusion center based upon the MSC decision at aspect i and all M prior final decisions. The decision fusion implements the following decision rule,

$$\Lambda(\boldsymbol{z}_i) = \frac{p(C_1|\boldsymbol{z}_i, \boldsymbol{u}_{if})}{p(C_0|\boldsymbol{z}_i, \boldsymbol{u}_{if})} = \frac{p(C_1|\boldsymbol{z}_i)\prod_{j=1}^{M} p(C_1|u_{f,i-j})}{p(C_0|\boldsymbol{z}_i)\prod_{j=1}^{M} p(C_0|u_{f,i-j})} = \eta(\boldsymbol{z}_i)\frac{\prod_{j=1}^{M} p(C_1|u_{f,i-j})}{\prod_{j=1}^{M} p(C_0|u_{f,i-j})} = u_{f,i} \qquad (21)$$

where $\boldsymbol{u}_{if} = \{u_{f,i-1}, u_{f,i-2}, \dots, u_{f,i-M}\}$ represents a sequence of M previous final decisions $u_{f,i-j},\ j \in [1, M]$ and u_{fi} (output of the fusion center) is the final decision about the class membership of $\boldsymbol{z}_i$.

Thus, this multi-aspect classification fusion system generates the likelihood ratio $\Lambda(\boldsymbol{z}_i)$ which tests the class membership (C_0 versus C_1 or the non-UXO class versus the UXO class) of an unknown object given the feature vector $\boldsymbol{z}_i$ and a set of M previous final decisions $\boldsymbol{u}_{if} = \{u_{f,i-1}, u_{f,i-2}, \dots, u_{f,i-M}\}$. Here, it is assumed that the previous decisions and $\boldsymbol{z}_i$ are conditionally independent, i.e. $p(C_k|\boldsymbol{z}_i, \boldsymbol{u}_{if}) = p(C_k|\boldsymbol{z}_i)\prod_{j=1}^{M} p(C_k|u_{f,i-j})$. Also, the MSC's preliminary decision $u_i \in [0, 1]$ is generated using the mapping function

$$u_i = \sigma(\ln(\eta(\boldsymbol{z}_i))) = \frac{1}{1 + \exp(-\ln(\eta(\boldsymbol{z}_i)))} \qquad (22)$$

which combines the log operation and logistic squashing function $\sigma(\cdot)$ that maps MSC's output $\eta(\boldsymbol{z}_i)$.

As stated before a favorable property of this system is that a variable number of previous decisions can be used to form $\Lambda(\boldsymbol{z}_i)$. To train the BPNN, the actual decisions in (21) are used. Therefore, the same BPNN can be used to generate $p(C_k|u_{f,i-j}), k = 0,1$ regardless of how many previous final decisions are used to form $u_{f,i}$.

8.2 Multi-Aspect Classification Fusion Results

In this section, we provide the results of the CMAC classification fusion system with decision-feedback when trained using the same training and testing data sets discussed in Sections 6.3.1 and 6.3.2, respectively. The performance of this multi-aspect classification fusion system is compared against that of the MSC system for multiple aspects using the test statistic in (13). More specifically, as stated in [18] multi-aspect classification results can be generated using the MSC by forming columns of data matrix $\boldsymbol{Z}$ using the acoustic color feature vectors extracted from TREX13 sonar data at several ($P = 7$) pings/aspects with certain separation. However, it is easy to show that the resultant test statistic in (13) is just the sum of the individual test statistic associated with each acoustic color observation vector in $\boldsymbol{z}_i$ in $\boldsymbol{Z}$. Although, this is an easy way of making decisions on a sequence of multiple sonar aspects, it is not necessarily an optimum way of performing multi aspect decision-making. This is due to the fact that the individual contribution and importance of each individual decision is lost in the summation operation. As will be shown shortly the proposed system in Section 8.1 circumvents all these issues and provides a very powerful approach to multi-aspect/ping classification especially using a small number of aspects/pings.

Figures 9 (a) and (b) display the ROC curves of the proposed multi-aspect ($P = 7$ aspects) classification fusion system using the incremental learning procedure discussed in Section 6.4 for the MSC subsystem without and with two UXO objects held out, respectively. Several interesting observations can be drawn from these results. First, comparing to the ROC curves in Figure 4 for the incrementally trained MSC without multi-aspect classification fusion, over 10% improvement in the correct classification rate can be

observed for this CMAC with decision-feedback. This is indeed a substantial improvement in the performance especially considering the fact that both systems used exactly the same number of aspects for decision-making and the same linear classifier namely the MSC. Second, comparing the two ROC curves in Figures 9 (a) and (b) reveals the great generalization ability of this system as the degradation in performance due to holding out (i.e. not included in the incremental learning) two UXO objects is almost negligible (0.1%). This robustness in performance is indeed another important benefit of this system for UXO classification. Finally, the modularity property of this system in terms of being able to change the number of aspects/pings used in decision feedback is a unique advantage not shared by other fusion mechanisms.

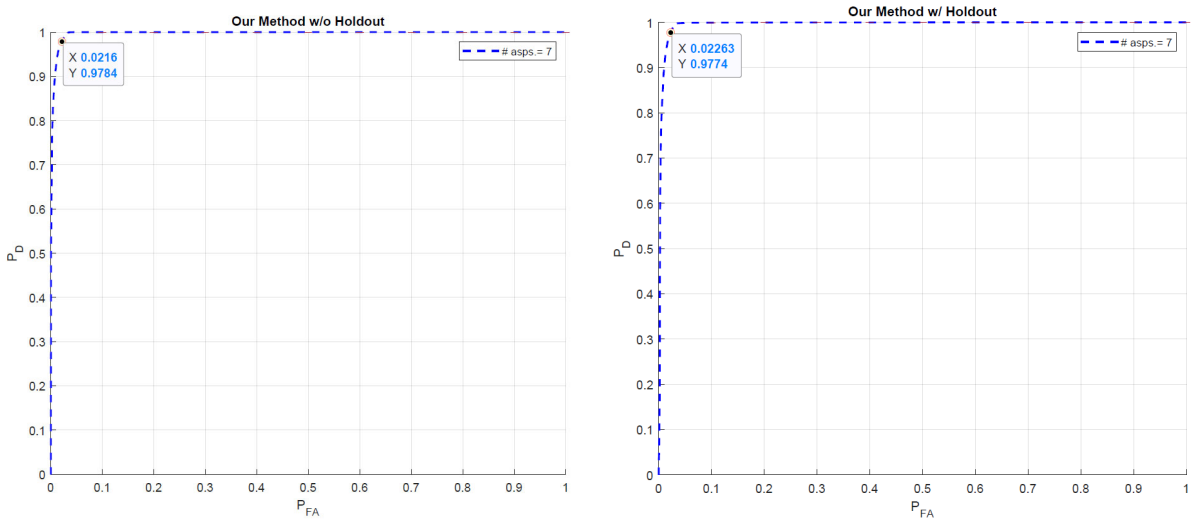

(a) no object held out (b) with objects held out

Figure 9: ROC curves of CMAC with decision-feedback and incrementally trained MSC-without and with object holdout.

To shed some light on the reasons why this system provides such superior results, we plotted the class conditional distributions of the MSC test statistics, $\eta(\boldsymbol{z}_i)$, for the incrementally trained system without and with two UXO objects holdout in Figures 10 (a) and 10 (b), respectively. These figures show that these distributions overlap and this overlap is somewhat more prominent for the case where two objects were held out as expected. Figures 11 (a) and 11 (b), on the other hand, display the distribution of the test statistics, $\Lambda(\boldsymbol{z}_i)$, when using the CMAC classification fusion with decision-feedback under the same conditions. Comparing the distributions of the final test statistics, $\Lambda(\boldsymbol{z}_i)$, with those in Figures 10 (a) and 10 (b) clearly reveals why the latter has done a superb job in UXO versus non-UXO classification. Note that the peak in the test statistic distribution for the UXO class is due to the thresholding operation that occurs at the output of the BPNN.

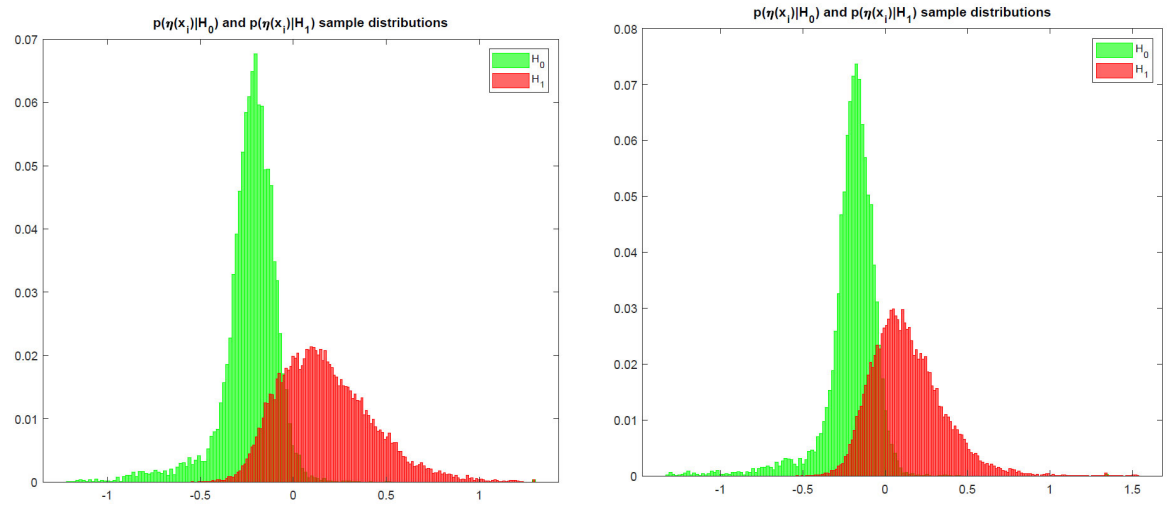

(a) no object held out (b) with objects held out

Figure 10: Class conditional distribution of MSC test statistics- without and with object holdout.

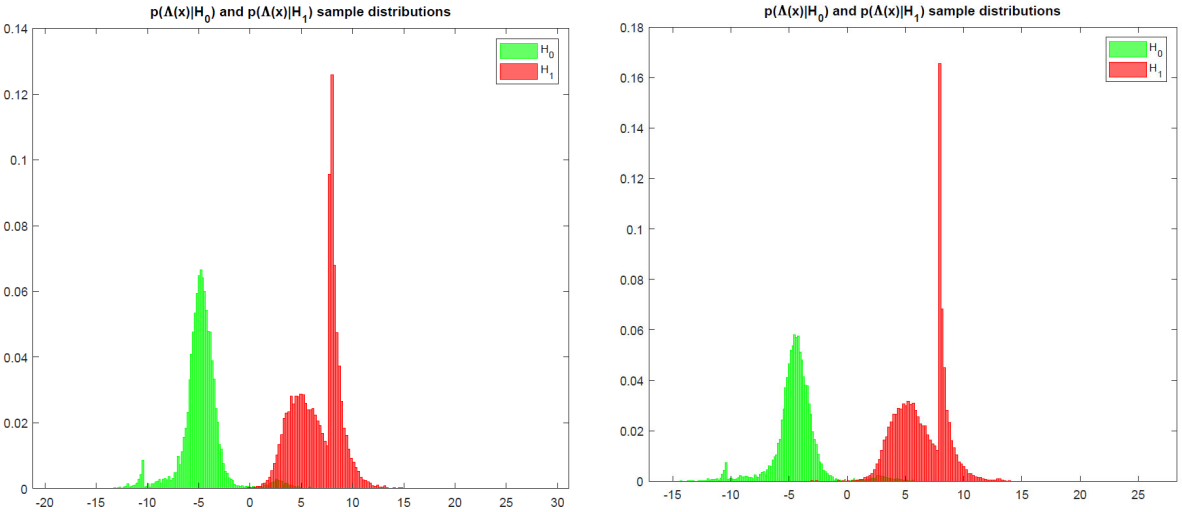

(a) no object held out (b) with objects held out

Figure 11: Class conditional distribution of final test statistics - without and with object holdout

9 Conclusions & Implications for Future Research/Implementation

9.1 Observations & Conclusions

This two-year research project not only led to the development of new algorithms for robust detection and classification of underwater UXOs in varying operating and environmental conditions but also revealed many important observations and strategies on how to improve the overall performance of such systems. More specifically, these developments and resultant observations drawn from the comprehensive experimental testing and validation are listed below for every task in this two-year project.

1. Research efforts in Task 1 of this project led to the development of a multi-channel (multi-sonar) broadband coherence detector (MBCD) and its adaptive version which adjusts the detection threshold automatically based upon the clutter distribution when operating in new environments. The results of this detector the model-generated TIER-SWAT data sets illuminated many interesting observations. First, multi-sonar (e.g., HF and BB)-based detection using MBCD provides substantially better detection of underwater targets when compared with the commonly used detectors that rely on some form of matching. Second, the MBCD system can be configured in many variety of ways to perform important trade-off studies. For example, we applied the MBCD to both the stave and SAS-processed data using two sonar channels namely HF and LF. Third, the adaptive version of the MBCD, in general, provides better detection results when compared with the fixed MBCD case and this adaptation can be done easily when the background clutter distribution undergoes significant variation e.g., moving from smooth sandy bottom to rocky conditions.

2. In Task 2, we developed a new incremental learning algorithm for the MSC-based classifier in order to provide the flexibility to learn new UXO and non-UXO objects that can been countered in new environments without requiring to retrain the entire system. The results of our efforts in this task, i.e. those in Section 6.4 and in [18], led to several important observations and conclusions. First, the classification performance of the MSC system when baseline trained using a portion of the TIER model-generated data and then incrementally updated using about 9% of the TREX13 data sets was found to be >10%better than the system without incremental updating. Second, the incrementally updated system possesses excellent generalization ability for classifying unseen (not included in the training data) UXO or non-UXO objects. Third, the proposed incremental learning algorithm does not impact the stability of the previously learnt data, an important property if the system to be used for long-term training in many difficult environments. Fourth, this incremental learning can be accomplished efficiently without the need to carry over the previous training data sets. Finally, the incremental learning algorithm is simple to implement by adding new atoms to the dictionary matrices. Thus, we strongly advocate the use of such incremental learning algorithm for any UXO classification application in varying conditions.

3. The question that we attempted to answer in Task 3 is: "Can a kernel extension of the MSC classification framework be developed (much like the support vector machine (SVM) or relevance vector machine (RVM) [9],[15]) in order to provide substantially better UXO versus non-UXO discrimination performance"? This theoretical extension requires not only extending the theory of the MSC to nonlinear (kernel) based classification but also kernelizing the K-SVD dictionary learning method used to build the classifier using the subspace matrices. Our investigations, however, indicated that this approach would lead to a very complicated process unsuitable for incremental learning settings. Therefore, we decided to resort to a different technique that takes advantage of the improved discrimination in nonlinearly mapped feature space and at the same time can easily be amenable to incremental leaning in new environments. Our method is inspired by the work in [30] and relies on mapping the training and testing data sets to the corresponding *virtual sample space* which would allow the use of any linear dictionary learning algorithm while providing the benefits of a kernel-based solution. Our recent results presented in Sections 7.2 and 7.3 indicate several important key observations. First, the importance of the nonlinear feature mapping on the overall discrimination of the system is clearly evident by the rise in the overall classification performance by more than 7%. Second, this increase in the performance can be gained even when two UXO objects in the TREX13 data set were completely held out during the training. Third, the proposed method relies on using linear

(e.g., K-SVD) dictionary learning while providing the benefits of kernel mapping via the virtual sample space. Fourth, the sampling in the Nyström method retains only the essential samples needed for generating the virtual mapping hence it offers a useful mechanism for avoiding overfitting in life-long learning applications. Finally, the proposed method can still take advantage of the incremental learning if the LKDL mapping can be updated based on some limited data samples drawn from the new environment without sacrificing the previous learning.

4. Task 4 research involved the development of an optimum multi-aspect UXO classification fusion system that can provide high confidence decision-making even in situations where only a few useful aspects/looks on the object are available. A new version of the classification fusion system known as the CMAC was developed utilizing a MSC and a neural network subsystems together with a fusion center. The system makes decisions based upon the acoustic color features at a particular aspect/ping as well as prior decisions made at M previous aspects/pings (decision feedback).The results presented in Section 8.2 led to many interesting and useful conclusions/observations. First, the results on the same TREX13 test data sets showed over 10% improvement in correct classification rate when compared to the MSC-based multi-aspect classification. More specifically, the knee-point of the ROC curve gave P_{CC}= 98% and P_{FA}= 2% which are indeed great results given that a linear MSC classifier was used in this system. Second, the system offers the best generalization capability as the degradation in performance due to two UXO object hold out during the training is less than 0.1%. Third, the decision feedback mechanism exploited in this system increases the separation between the distributions of the test statistic under both UXO and non-UXO hypothesis hence leading to significantly improved discrimination. Finally, the proposed system is simple to implement and can easily incorporate a variable number of aspects/pings to determine the final class label of an object, and hence provides a very versatile and modular means of performing multiple-ping UXO classification fusion.

9.2 Future Research/Implementation

We believe this two-year research project made significant strides toward developing underwater UXO detection and classification systems that remain robust to variations in operating and environmental conditions and hence addressed many critical questions on how to construct such systems for successful deployment in different shallow water settings. In particular, our research studies culminated to conclusive and thorough experimental results that highlight the importance and merit of all the above-mentioned algorithms as critical ingredients of any successful UXO detection and classification system. Consequently, we propose to continue our current research efforts by investigating the following important areas:

6. Comprehensive testing, evaluation, and fine-tuning of our developed systems on more realistic data sets e.g., CLUTTEREX17 or BOSS data sets.

7. Conduct tradeoff studies of the MBCD to determine impact of different sonar configurations and choice of parameters on the ability to provide unambiguous UXO detection in different settings. This task will be done in conjunction with our collaborators at APL-UW.

8. Study the effects of life-long incremental learning and how to overcome performance drift and structural complexity of the linear MSC or kernel-based MSC classifiers.

9. Develop a new incremental updating mechanism for the LKDL mapping of the original data to

the virtual feature space. This, will allow the full implementation of the kernel-based incremental learning when new data samples become available in small batches.

10. Develop a new mechanism for the CMAC-based multi-aspect classification fusion system in order to automatically decide the number of aspects/pings needed to perform confusion-free decision-making in different conditions.

11. Transition the code and documentations to our collaborators at APL-UW and NSWC-PCD as well as NRL counterparts and help with testing and validation.

Literature Cited

[1] J.D. Tucker and M.R. Azimi-Sadjadi, "Coherence-Based Underwater Target Detection from Multiple Disparate Sonar Platforms", *IEEE Journal of Oceanic Engineering*, vol. 34, pp. 37-51, January 2011.

[2] N. H. Klausner and M.R. Azimi-Sadjadi, "Detection in Multiple Disparate Systems using Multi-Channel Coherence Analysis", *IEEE Trans on Aerospace and Electronic Systems*, vol. 48, pp.3554-3566, October 2012.

[3] N. Klausner, M. Azimi-Sadjadi, and L. Scharf, "Detection of spatially correlated time series from a network of sensor arrays," *IEEE Transactions on Signal Processing*, vol. 62, no. 6, pp. 1396-1407, March 2014.

[4] D. Cochran, H. Gish, and D. Sinno, "A geometric approach to multiple-channel signal detection," *IEEE Transactions on Signal Processing*, vol. 43, no. 9, pp. 2049-2057, September 1995.

[5] M. Azimi-Sadjadi and S. Kargl, "Multichannel detection and acoustic color-based classification of underwater UXO in sonar," *First Year SERDP Report for MR-2416*, March 2018.

[6] S. Kargl, K. Williams and E. Thorsos, "Synthetic aperture sonar imaging of simple finite targets," IEEE Journal of Oceanic Engineering, vol. 37, no. 3, pp. 516-532, July 2012.

[7] S. Kargl, K. Williams, T. Marston, and J. Kennedy, "Acoustic response of unexploded ordnance (UXO) and cylindrical targets," *Proceedings of MTS/IEEE Oceans Conference* 2010, pp. 1-5, September 2010.

[8] R. Lim, "Sonar detection and classification of underwater UXO and environmental parameters," Interim Report, SERDP Project MR-1666, pp. 1-19, October 2010.

[9] J. Bucaro, Z. Waters, B. Houston, H. Simpson, A. Sarkissian, S. Dey, and T. Yoder, "Acoustic identification of buried underwater exploded ordnance using a numerically trained classifier," *Journal of the Acoustical Society of America*, vol. 132, pp. 3614-3617, December 2012.

[10] E. Fischell and H. Schmidt, "Classification of underwater targets from autonomous underwater vehicle sampled bistatic acoustic scattered fields," *Journal of the Acoustical Society of America*, vol. 138, pp. 3773-3784, December 2015.

[11] N. Wachowski and M. R. Azimi-Sadjadi, “A new synthetic aperture sonar processing method using coherence analysis," *IEEE Journal of Oceanic Engineering*, vol. 34, pp. 665-678, October 2011.

[12] P. Runkle, L. Carin, L. Couchman, J. A. Bucaro, and T. J. Yoder, “Multi-aspect identification of submerged elastic targets via wave-based matching pursuits and hidden Markov models," *Journal of Acoustical Society of America*, vol. 106, pp. 605-616, August 1999.

[13] M. Azimi-Sadjadi. M. Robinson and J. Salazar, “Multi-aspect classification using hidden Markov models and neural networks," *IEEE Transactions on Neural Networks*, vol. 16, pp. 447-459, March 2005.

[14] J. Cartmill, M. Azimi-Sadjadi, and N. Wachowski, “Buried underwater object classification using a collaborative multi-aspect classifier," *IEEE Journal of Oceanic Engineering*, vol. 34, pp. 32-44, January 2009.

[15] C. M. Bishop, *Pattern recognition*. Springer, 2006, vol. 128, ch. 7, pp. 345-356.

[16] S. Kargl, A. España, K. Williams, J. Kennedy, and J. Lopes, “Scattering from objects at a water-sediment interface: Experiment, high-speed and high-fidelity models, and physical insight," *IEEE Journal of Oceanic Engineering*, vol. 40, no. 3, pp. 632-642, 2015.

[17] S. Kargl and K. Williams, “Full scale measurement and modeling of the acoustic response of proud and buried munitions at frequencies from 1-30kHz,"*Final Report, SERDP Project MR-1665*, pp. 18-19, May 2012.

[18] J. Hall, M. R. Azimi-Sadjadi, S. G. Kargl, K. L. Williams, and Y. Zhao, “Underwater UXO classification using a matched subspace classifier with adaptive dictionaries," *Oceanic Engr.*, vol. 44, pp. 739-751, July 2019.

[19] A. Salberg, A. Hanssen, and L. Scharf, “Robust multidimensional matched subspace classifiers based on weighted least-squares," *IEEE Transactions Signal Processing*, vol. 55, pp. 873-880, March 2007.

[20] M. R. Azimi-Sadjadi, Y. Zhao, and S. Sheedvash, “Incremental dictionary learning with sparsity," in *Proc. of 2018 Inter. Joint Conf. on Neural Networks (IJCNN)*, July 2018.

[21] L. Scharf, *Statistical Signal Processing: Detection, Estimation, and Time Series Analysis*. Addison-Wesley, 1991.

[22] D. Ramirez, J. Via, I. Santamaria, and L. Scharf, “Detection of spatially correlated Gaussian time series," *IEEE Transactions on Signal Processing*, vol. 58, no. 10, pp. 5006-5015, 2010.

[23] K. L. Williams, “Reverberation, sediment acoustics, and targets-in-the-environment," WASHINGTON UNIV SEATTLE APPLIED PHYSICS LAB, Tech. Rep., 2012.

[24] G. S. Sammelmann, “Propagation and scattering in very shallow water," in *MTS/IEEE Oceans 2001. An Ocean Odyssey. Conference Proceedings (IEEE Cat. No. 01CH37295)*, vol. 1. IEEE, 2001, pp. 337-344.

[25] M. Aharon, M. Elad, and A. Bruckstein, “K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation," *IEEE Transactions on Signal Processing*, vol. 54, no. 11, pp. 4311-4322, November 2006.

[26] S. Chen, D. Donoho, and M. Saunders, “Atomic decomposition by basis pursuit," *SIAM journal on scientific computing*, vol. 20, no. 1, pp. 33-61, 1998.

[27] M. Azimi-Sadjadi, J. Kopacz, and N. Klausner, “K-SVD dictionary learning using a fast OMP with applications," *Proc. of IEEE Intern. Conference on Image Processing (ICIP)*, pp. 1599-1603, October 2014.

[28] B. Scholkopf and A. Smola, *Learning with Kernels*. The MIT press, 2002.

[29] H. Van Nguyen, V. M. Patel, N. M. Nasrabadi, and R. Chellappa, “Design of non-linear kernel dictionaries for object recognition," *IEEE Transactions on Image Processing*, vol. 22, no. 12, pp. 5123-5135, 2013.

[30] A. Golts and M. Elad, “Linearized kernel dictionary learning," *IEEE Journal of Selected Topics in Signal Processing*, vol. 10, no. 4, pp. 726-739, 2016.

[31] S. Kumar, M. Mohri, and A. Talwalkar, “Sampling methods for the Nyström method," *Journal of Machine Learning Research 13 (2012)*, vol. 13, pp. 981-1006, 2012.

[32] M. Azimi-Sadjadi, Q. Yao, D. Huang, and G. Dobeck, “Underwater target classification using wavelet packets and neural networks," *IEEE Transactions on Neural Networks*, vol. 11, no. 3, pp. 784-794, May 2000.

[33] M. Azimi-Sadjadi. A. Pezeshki and L. Scharf, “Undersea target classification using canonical correlation analysis," *IEEE Journal of Oceanic Engineering*, vol. 32, pp. 948-955, October 2007.

10 Appendix:

Publications Resulted from this Research

Journal Publications:

[1] J. Hall, M. R. Azimi-Sadjadi, S. Kargl, Y. Zhao, and K. Williams “Underwater Unexploded Ordnance (UXO) Classification Using a Matched Subspace Classifier with Adaptive Dictionaries," *IEEE Journal of Oceanic Engineering*, vol. 44, pp: 739-751, July 2019.

[2] A. Pezeshki, M. R. Azimi-Sadjadi, and C. Robbiano, “A Multiple Kernel Machine with In-Situ Learning using Sparse Representation," *submitted to IEEE Journal on Selected Topics in Signal Processing*.

Conference Publications/Presentations:

[1] N. Larson, J. Hall, and M. R. Azimi-Sadjadi, “Analyzing Transfer Learning Methods for UXO Classification in Shallow Water Environments ,"*Proc. of 2019 IEEE Intern. Workshop on Machine Learning for Signal Processing*, Pittsburgh, PA, October 2019.

[2] A. Pezeshki, M. R. Azimi-Sadjadi, and C. Robbiano, “A Multiple Kernel Machine with In-Situ Learning using Sparse Representation," *Proc. of IEEE Intern. Joint Conference on Neural Networks (IJCNN 2019)*, Budapest, Hungary, July 2019.

[3] M. R. Azimi-Sadjadi and Y. Zhao, “Incremental Dictionary Learning With Sparsity", *Proc. of IEEE Intern. Joint Conference on Neural Networks (IJCNN 2018)*, Rio de Janeiro, Brazil, July 2018.

[4] J. Hall, M. R. Azimi-Sadjadi, and S. Kargl, “Underwater UXO classification using Matched Subspace Classifier with Synthetic Sparse Dictionaries," *Proc. of OCEANS 2016 MTS/IEEE* Monterey, CA, 2016.

.