

REPORT DOCUMENTATION PAGE			Form Approved OMB NO. 0704-0188		
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 22-12-2014		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 1-Oct-2008 - 30-Sep-2013	
4. TITLE AND SUBTITLE Final Technical Report: Multi-Sensor Detection of Obscured and Buried Objects			5a. CONTRACT NUMBER W911NF-08-1-0410		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER 622712		
6. AUTHORS Jeremy Bolton, Brandon Smock			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAMES AND ADDRESSES University of Florida Office of Engineering 339 Weil Hall Gainesville, FL 32611 -6550			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211			10. SPONSOR/MONITOR'S ACRONYM(S) ARO		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S) 55033-CS.9		
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited					
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.					
14. ABSTRACT The subject research was performed at the University of Florida between the primary contract dates of Fall 2008 and Summer 2013, with an extension through Spring 2014. The research was performed to support the ability to detect landmines and such subsurface objects by a variety of sensors and platforms employed in systems being studied by NVESD. The work in this period of research was concerned with discovering and/or evaluating: Year 1 • Detection and discrimination of land mines in ground penetrating radar based on edge histogram descriptors					
15. SUBJECT TERMS Landmine detection, University of Florida, pattern recognition, image processing, multi-sensor fusion, classifier development, multiple-instance learning (MIL), joint orthogonal matching pursuit (JOMP), sweep detection					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Paul Gader
a. REPORT UU	b. ABSTRACT UU	c. THIS PAGE UU			19b. TELEPHONE NUMBER 352-505-2551

Report Title

Final Technical Report: Multi-Sensor Detection of Obscured and Buried Objects

ABSTRACT

The subject research was performed at the University of Florida between the primary contract dates of Fall 2008 and Summer 2013, with an extension through Spring 2014. The research was performed to support the ability to detect landmines and such subsurface objects by a variety of sensors and platforms employed in systems being studied by NVESD. The work in this period of research was concerned with discovering and/or evaluating:

Year 1

- Detection and discrimination of land mines in ground-penetrating radar based on edge histogram descriptors and a Possibilistic K-Nearest Neighbor Classifier,

Year 2

- Gradient Angle Model Algorithm on Wideband EMI data Classifier,
- Context Dependent Multi-Sensor Fusion and its Application to Land Mine Detection,
- Airborne and Ground Sensor Fusion for Target Detection, and
- Variational Mixture of Experts for Classification.

Year 3

- A Large Scale Evaluation of Several Fusion Algorithms for Anti-tank Landmine Detection and Discrimination. This evaluation including the investigation and analysis of several preprocessing techniques, features, detectors, and fusion approaches for landmine detection, including the following:

- i) HMM detector
- ii) EHD detector
- iii) SPECT detector
- iv) GEOM detector
- v) TFCM detector
- vi) GMRF detector
- vii) GFIT detector
- viii) Bayesian-based fusion
- ix) Mahalanobis distance-based fusion
- x) Dempster-Shafer based Fusion
- xi) Decision template fusion
- xii) Discrete Choquet integral
- xiii) Context-dependent fusion

Year 4

- Mixture of HMM Experts with Applications to Landmines Detection
- Landmine Detection Using Two-Tapped Joint Orthogonal Matching Pursuits

Year 5 (plus extension)

- Multiple-instance learning (MIL) for landmine detection
- Sweep detection in hand-held ground penetrating radar data

Enter List of papers submitted or published that acknowledge ARO support from the start of the project to the date of this printing. List the papers, including journal references, in the following categories:

(a) Papers published in peer-reviewed journals (N/A for none)

<u>Received</u>	<u>Paper</u>
08/08/2012 1.00	H. Frigui, P. Gader. Detection and Discrimination of Land Mines in Ground-Penetrating Radar Based on Edge Histogram Descriptors and a Possibilistic χ^2 -Nearest Neighbor Classifier, IEEE Transactions on Fuzzy Systems, (02 2009): 0. doi: 10.1109/TFUZZ.2008.2005249
08/22/2012 2.00	Hichem Frigui, Lijun Zhang, Paul D Gader. Context-Dependent Multisensor Fusion and Its Application to Land Mine Detection, IEEE Transactions on Geoscience and Remote Sensing, (06 2010): 0. doi: 10.1109/TGRS.2009.2039936
08/22/2012 3.00	Ganesan Ramachandran, Paul D. Gader, Joseph N. Wilson. GRANMA: Gradient Angle Model Algorithm on Wideband EMI Data for Land-Mine Detection, IEEE GEOSCIENCE AND REMOTE SENSING LETTERS, (07 2010): 0. doi: 10.1109/LGRS.2010.2041184
08/22/2012 5.00	Hichem Frigui, Lijun Zhang, Paul Gader, Joseph N. Wilson, K.C. Ho, Andres Mendez-Vazquez. An evaluation of several fusion algorithms for anti-tank landmine detection and discrimination, Information Fusion, (4 2012): 0. doi: 10.1016/j.inffus.2009.10.001
09/09/2012 7.00	Seniha Esen Yuksel, Joseph N. Wilson, Paul D. Gader. Twenty Years of Mixture of Experts, IEEE Neural Networks, (08 2012): 0. doi: 10.1109/TNNLS.2012.2200299
12/04/2014 13.00	Oualid Missaoui, Hichem Frigui, Paul Gader. Multi-stream continuous hidden Markov models with application to landmine detection, EURASIP Journal on Advances in Signal Processing, (03 2013): 0. doi: 10.1186/1687-6180-2013-40
TOTAL:	6

Number of Papers published in peer-reviewed journals:

(b) Papers published in non-peer-reviewed journals (N/A for none)

<u>Received</u>	<u>Paper</u>
-----------------	--------------

TOTAL:

Number of Papers published in non peer-reviewed journals:

(c) Presentations

Number of Presentations: 0.00

Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

<u>Received</u>	<u>Paper</u>
09/09/2012 6.00	Seniha Esen Yuksel, Paul D. Gader. MIXTURE OF HMM EXPERTS WITH APPLICATIONS TO LANDMINE DETECTION, IEEE IGARSS. 28-JUL-12, . : ,
12/04/2014 12.00	Jeremy Bolton, Paul Gader, Hichem Frigui. Embedding the multiple instance problem: applications to landmine detection with ground penetrating radar, SPIE Defense, Security, and Sensing. 29-APR-13, Baltimore, Maryland, USA. : ,
12/04/2014 10.00	Joseph N. Wilson, Jeremy Bolton, Peter J. Dobbins. Sweep detection and alignment in handheld GPR detection devices, SPIE Defense, Security, and Sensing. 29-APR-13, Baltimore, Maryland, USA. : ,
12/12/2014 14.00	Alina Zare, Miranda Silvius, Ryan Close, Paul Gader. Quantifying the benefit of airborne and ground sensor fusion for target detection, SPIE Defense, Security, and Sensing. 05-APR-10, Orlando, Florida. : ,
TOTAL:	4

Number of Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Peer-Reviewed Conference Proceeding publications (other than abstracts):

Received

Paper

08/22/2012 4.00 Paul Gader, Seniha Esen Yuksel. Variational Mixture of Experts for Classification with Applications to Landmine Detection, 2010 20th International Conference on Pattern Recognition (ICPR). 23-AUG-10, Istanbul, Turkey. : ,

09/09/2012 8.00 Sean Goldberg, Taylor Glenn, Joseph N. Wilson, Paul D. Gader. Landmine detection using two-tapped joint orthogonal matching pursuits, Detection and Sensing of Mines, Explosive Objects, and Obscured Targets XVII. 27-APR-12, Baltimore, Maryland, USA. : ,

12/04/2014 11.00 Seniha Esen Yuksel, Jeremy Bolton, Paul D. Gader. Landmine detection with Multiple Instance Hidden Markov Models, 2012 IEEE International Workshop on Machine Learning for Signal Processing (MLSP). 23-SEP-12, Santander, Spain. : ,

TOTAL: 3

Number of Peer-Reviewed Conference Proceeding publications (other than abstracts):

(d) Manuscripts

Received

Paper

TOTAL:

Number of Manuscripts:

Books

Received

Book

TOTAL:

Received

Book Chapter

TOTAL:

Patents Submitted

Patents Awarded

Awards

Graduate Students

<u>NAME</u>	<u>PERCENT SUPPORTED</u>	<u>Discipline</u>
Pete Dobbins	0.06	
Taylor Glenn	0.17	
Brandon Smock	0.12	
Ken Watford	0.14	
Ryan Close	0.03	
Dmitri Dranishnikov	0.03	
Joshua Wood	0.03	
Sean Goldberg	0.08	
FTE Equivalent:	0.66	
Total Number:	8	

Names of Post Doctorates

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
Xuping Zhang	0.06
Alina Zare	0.05
Jeremy Bolton	0.15
FTE Equivalent:	0.26
Total Number:	3

Names of Faculty Supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>	National Academy Member
Paul Gader	0.18	
Joseph Wilson	0.15	
FTE Equivalent:	0.33	
Total Number:	2	

Names of Under Graduate students supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Student Metrics

This section only applies to graduating undergraduates supported by this agreement in this reporting period

The number of undergraduates funded by this agreement who graduated during this period: 0.00

The number of undergraduates funded by this agreement who graduated during this period with a degree in science, mathematics, engineering, or technology fields:..... 0.00

The number of undergraduates funded by your agreement who graduated during this period and will continue to pursue a graduate or Ph.D. degree in science, mathematics, engineering, or technology fields:..... 0.00

Number of graduating undergraduates who achieved a 3.5 GPA to 4.0 (4.0 max scale):..... 0.00

Number of graduating undergraduates funded by a DoD funded Center of Excellence grant for Education, Research and Engineering:..... 0.00

The number of undergraduates funded by your agreement who graduated during this period and intend to work for the Department of Defense 0.00

The number of undergraduates funded by your agreement who graduated during this period and will receive scholarships or fellowships for further studies in science, mathematics, engineering or technology fields: 0.00

Names of Personnel receiving masters degrees

<u>NAME</u>
Ken Watford
Joshua Wood
Total Number:

Names of personnel receiving PHDs

<u>NAME</u>
Xuping Zhang
Ryan Close
Total Number:

Names of other research staff

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Sub Contractors (DD882)

Inventions (DD882)

Scientific Progress

Technology Transfer

See attached.

This document contains an overview of research and work performed and published at the University of Florida from October 1, 2008 to March 31, 2014 pertaining to proposal 55033CS: *Multi-Sensor Detection of Obscured and Buried Objects*.

Overview

Topics By Year


1. 2008-2009
 - EHD and a possibilistic k-nearest neighbor classifier
2. 2009-2010
 - Gradient angle model algorithm on wideband EMI data classifier
 - Context dependent multi-sensor fusion and its application to land mine detection
 - Airborne and ground sensor fusion for target detection
 - Variational mixture of experts for classification
3. 2010-2011
 - A large scale evaluation of several fusion algorithms for anti-tank landmine detection and discrimination
 - o This evaluation including the investigation and analysis of several preprocessing techniques, features, detectors, and fusion approaches for landmine detection, including the following:
 - o HMM detector
 - o EHD detector
 - o SPECT detector
 - o GEOM detector
 - o TFCM detector
 - o GMRF detector
 - o GFIT detector
 - o Bayesian-based fusion
 - o Mahalanobis distance-based fusion
 - o Dempster-Shafer based Fusion
 - o Decision template fusion
 - o Discrete Choquet integral
 - o Context-dependent fusion
4. 2011-2012
 - Landmine detection using two-tapped joint orthogonal matching pursuits
 - Application of a mixture of hidden Markov experts to WEMI data
5. 2012-2013 + Extension
 - Multiple-instance learning (MIL) for landmine detection
 - Sweep detection in hand-held ground penetrating radar data

2008-2009


EHD and a possibilistic k-nearest neighbor classifier

This research investigated an algorithm for land mine detection using sensor data generated by a ground- penetrating radar (GPR) system that uses edge histogram descriptors (EHD) for feature extraction and a possibilistic K-nearest neighbors (K-NNs) rule for confidence assignment (Frigui et al, 2009). The proposed algorithm demonstrated the best performance among several high-performance algorithms in extensive testing on a large real-world datasets associated with the difficult problem of land mine detection. The superior performance of the algorithm is attributed to the use of the possibilistic K-NN algorithm, thereby providing important evidence supporting the use of possibilistic methods in real-world applications.

The GPR produces a 3-D array of intensity values, representing a volume below the surface of the ground. First, a computationally inexpensive prescreening algorithm for anomaly detection is used to focus attention and identify candidate signatures that resemble mines. The identified regions of interest are processed further by a feature extraction algorithm to capture their salient features. We use translation invariant features that are based on the local edge distribution of the 3-D GPR signatures. Specifically, each 3-D signature is divided into subsignatures, and the local edge distribution for each subsignature is represented by a histogram -- as shown below. Next, the training signatures are clustered to identify prototypes. The main idea is to identify few prototypes that can capture the variations of the signatures within each class. These variations could be due to different mine types, different soil conditions, different weather conditions, etc. Fuzzy memberships are assigned to these representatives to capture their degree of sharing among the mines and false alarm classes. Finally, a possibilistic K-NN- based rule is used to assign a confidence value to distinguish true detections from false alarms.



The researched algorithm was implemented and integrated within a complete land mine prototype system. It is trained, field-tested, evaluated, and compared using a large-scale cross-validation experiment that uses a diverse dataset acquired from four outdoor test sites at different geographic locations. This collection covers over 41 807m² of ground and includes 1593 mine encounters. ROC results are presented below.




2009-2010

Gradient angle model algorithm on wideband EMI data classifier

This research investigates a simple and fast algorithm to analyze wideband electromagnetic induction (EMI) data for subsurface targets (Ramachandran et al, 2010). A well-known four-parameter model, Cole- Cole, is differentiated, resulting in a two-parameter model. A fast lookup table is used to find parameters as opposed to nonlinear optimization. The researched approach provides a computationally faster way to reproduce the results of state-of-the-art methods on landmine EMI data. A detailed mathematical analysis of the model is given that describes the advantages and limitations of the researched method.


The researched method show to improve landmine classification results on two datasets. The experiments were performed on two data sets. The first contained 62 different types of objects, including 26 different types of mines collected over 11 adjoining lanes divided into 220 grid cells. The second contained 24 different types

of objects, including 12 different types of mines collected over 12 adjoining lanes divided into 225 cells.




Context dependent multi-sensor fusion and its application to land mine detection


This research investigates a novel method for fusing the results of multiple land mine detection algorithms which use different sensors (GPR and EMI), features, and different classification methods (Frigui et al, 2010). The researched multisensor/multialgorithm fusion method, which is called context-dependent fusion (CDF), is motivated by the fact that the relative performance of different sensors and algorithms can vary significantly depending on the mine type, geographical site, soil and weather conditions, and burial depth.



CDF is a local approach that adapts the fusion method to different regions of the feature space. The training part of CDF has two components: context extraction and algorithm fusion. In context extraction, the features used by the different algorithms are combined and used to partition the feature space into groups of similar signatures, or contexts. The algorithm fusion component assigns a degree of worthiness to each detector in each context based on its relative performance within the context. To test a new alarm using CDF, each detection algorithm extracts its set of features and assigns a confidence value. Then, the features are used to identify the best context, and the degrees of worthiness of this context are used to fuse the individual confidence values.



Results on large and diverse ground-penetrating radar and wideband electromagnetic data collections show that the researched method can identify meaningful and coherent clusters and that different expert algorithms can be identified for the different contexts. Typically, the contexts correspond to groups of alarm signatures that share a subset of common features. Our extensive experiments have also indicated that CDF outperforms all individual detectors and the global fusion that uses the same method to assign aggregation weights.




Airborne and ground sensor fusion for target detection


This research investigated the detection of buried objects by fusing airborne Multi-Spectral Imagery (MSI) and ground-based Ground Penetrating Radar (GPR) data is investigated (Zare et al, 2010). The benefit of using the airborne sensor to cue the GPR, which will then search the area indicated by the MSI, is investigated and compared to results obtained via a purely ground-based system. State-of-the-art existing algorithms, such as hidden Markov models and possibilistic classifiers will be applied to the GPR data both in queued and non-queued modes. In addition, the ability to measure disturbed earth with the GPR sensor will be investigated. Furthermore, state-of-the-art algorithms for the MSI system will be described. These algorithms require very high detection rates with acceptable false alarm rates in order to serve as an acceptable system.

Variational mixture of experts for classification

This research provided a complete framework for classification using Variational Mixture of Experts (VME), discovered the conceptual variational lower bound; and successfully applied the method to landmine detection (Yuksel et al, 2010). The results include a comparison to Mixtures of Experts trained with Expectation Maximization (EMME). VME has previously been used for regression and Waterhouse explained how to apply VME to classification (which we will call as VMEC). However, the steps to train the model were not made clear since the equations were applicable to vector valued parameters as opposed to matrices for each expert. Our research solidifies this gap and permits a principled implementation for real-world classification as shown below.



The discussed variational lower bound provides an excellent stopping criterion that resists over-training. This provides for a practical, real-world implementation. We demonstrate the efficacy of the method on real-world mine classification; in which, training robust mine classification algorithms is difficult because of the small number of samples per class. In our experiments VMEC consistently improved performance over EMME as shown below.



2010-2011

A large scale evaluation of several fusion algorithms for anti-tank landmine detection and discrimination

Many algorithms have been researched for detecting anti-tank landmines and discriminating between mines and clutter objects using data generated by a ground penetrating radar (GPR) sensor. Our extensive testing of some of these algorithms has indicated that their performances are strongly dependent upon a variety of factors that are correlated with geographical and environmental conditions. It is typically the case that one algorithm may perform well in one setting and not so well in another. Thus, fusion methods that take advantage of the stronger algorithms for a given setting without suffering from the effects of weaker algorithms in the same setting are needed to improve the robustness of the detection system. In this research effort, we investigate, test, and compare seven different fusion methods: Bayesian, distance-based, Dempster-Shafer, Borda count, decision template, Choquet integral, and context-dependent fusion (Frigui et al, 2011). We present the results of a cross validation experiment that uses a diverse data set together with results of eight detection and discrimination algorithms. These algorithms are the top ranked algorithms after extensive testing. The data set was acquired from multiple collections from four outdoor sites at different locations using the NIITEK GPR system. This collection covers over 41,807 m² of ground and includes 1593 anti-tank mine encounters.

The discrimination algorithms and the different fusion methods were implemented and tested with data collected using the NIITEK vehicle mounted GPR system. The data were collected between November 2002 and July 2006 from four geographically distinct test sites. Sites A, B, and D are temperate climate test facilities with prepared soil and gravel lanes. Site C is an arid climate test facility with prepared soil lanes. The four sites have a total of 17 different lanes with known mine locations. All mines are anti-tank (AT) mines. In all, there are 19 distinct mine types that can be classified into three categories: anti-tank metal (ATM), anti-tank with low metal content (ATLM), and simulated mines (SIM). The targets were buried up to 6 in. deep. Multiple data collections were performed at each site at different dates, covering a ground area of 41; 807:57 m, resulting in a large and diverse collection of mine and false alarm signatures. False alarms arise as a result of radar signals that present a mine-like character. Such signals are generally said to be a result of clutter. In this experiment, clutter arises from two different processes. One type of clutter is emplaced and surveyed in an effort to test the robustness of the algorithms. Other clutter result from human activity unrelated to the data collection or as a result of natural processes. We refer to this second kind of clutter as non-emplaced. Non-emplaced clutter includes objects discarded or lost by humans, soil inconsistencies and voids, stones, roots and other vegetation, as well as remnants of animal activity.

The data collected from Sites B and D have emplaced buried clutter. Although the lanes at Sites A and C are prepared, they still contain non-emplaced clutter objects. Both metal and non-metal non-emplaced clutter objects such as ploughshares, shell casings, and large rocks have been excavated from these sites. The emplaced clutter objects include steel scraps, bolts, soft-drink cans, concrete blocks, plastic bottles, wood blocks, and rocks. In all, there are 12 collections having 19 distinct mine types. Many of these mine types are present at several sites. The prescreener detected 1560 of the 1593 mines encountered in the data, yielding a 97.9% probability of detection. It rejected 161 of 211 emplaced clutter objects encountered, and yielded a total of 3435 false alarms associated with non-emplaced clutter objects. As it can be seen, the mines buried at 1 inch through 6 inches occupy 87.5% of the total targets encountered vs. 12.5% surface-laid or flush-buried mines.

To provide an objective and consistent evaluation of all algorithms, we use the TUF system with lane- based cross validation. The results of this process are scored using the Mine Detection Assessment and Scoring (MIDAS) system developed by the Institute for Defense Analysis. The scoring is performed in terms of probability of detection (PD) vs. false alarm rate (FAR). Confidence values are thresholded at different levels to produce Receiver Operating Characteristic (ROC) curve. For a given threshold, a mine is detected if there is an alarm within 0.25 m from the edge of the mine with confidence value above the threshold. Given a threshold, the PD is defined to be the number of mines detected divided by the number of mines. The FAR is defined as the number of false alarms per square meter.




Figure above: Eight Detector Results





Figure Above: Comparison of the EHD and HMM outputs for several mine (green dots) and clutter (red stars) signatures extracted from: (a) a subset of Site A; and (b) a subset of Site B. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




2011-2012

Landmine detection using two-tapped joint orthogonal matching pursuits

Joint Orthogonal Matching Pursuits (JOMP) is used in our research in the context of landmine detection using data obtained from an electromagnetic induction (EMI) sensor. The response from an object containing metal can be decomposed into a discrete spectrum of relaxation frequencies (DSRF) from which we construct a dictionary. A greedy iterative algorithm was developed for computing successive residuals of a signal by subtracting away the highest matching dictionary element at each step. The final confidence of a particular signal is a combination of the reciprocal of this residual and the mean of the complex component. Our two-tap approach comparing signals on opposite sides of the geometric location of the sensor is examined and found to produce better classification (Goldberg et al, 2012).

It is found that using only a single pursuit does a comparable job, reducing complexity and allowing for real-time implementation in automated target recognition systems. JOMP is particularly highlighted in comparison with a previous EMI detection algorithm known as String Match.

Experiments were performed on data collected by personnel from Niitek at a test facility with a desert climate in the western United States. The data included two different types of low metal antipersonnel mines and nine different types of antitank mines, containing both low and high metal. In total there were 88 objects containing metal. The goal is for JOMP to be able to designate each target object as an alarm with a low rate of false alarms. Results are presented below as a ROC curve.




A computationally efficient prescreener for detecting metal-bearing subsurface objects has been created as a result of our research that uses data collected by a wideband electromagnetic induction sensor. Using a physics-based dictionary based on relaxation frequencies, joint orthogonal matching pursuits yield a residual that provides a better confidence value on data collected than previous EMI algorithms. Performance would see an even greater increase in effectiveness when used in conjunction with a GPR prescreener.

Application of a mixture of hidden Markov experts to WEMI data

In many applications including object recognition, data classification may be hindered by the existence of multiple contexts that produce an input sample. To alleviate the problems associated with multiple contexts, context-based classification is a process that uses different classifiers depending on a measure of the context. Context-based classifiers offer the promise of increasing performance by allowing classifiers to become experts at classifying input samples of certain types, rather than trying to force single classifiers to perform well on all possible inputs.

Our research led to a novel mixture of experts model, the Mixture of Hidden Markov Model Experts (MHMME) which is shown below (Yuksel et al, 2012). This model is designed to perform context-based classification of samples that are variable length sequences. The model has a similar high-level structure to previous mixture of experts models but has the novelty that the gates and the experts are HMMs and the input data are sequences. The contexts are determined by the gates and the classifiers are determined by the experts. The gates and the experts are learned simultaneously using a single probabilistic model.



Experiments were done on a data set consisting of time samples of wide-band electro-magnetic data collected for landmine detection. An analysis of the functioning of the internal structure of the model is provided as well as classification and reliability rates. Summary results are shown below.

CLASSIFICATION RATES ON THE LANDMINE DATA FOR 10-FOLD TRAINING

Model	Mean	Standard Deviation
MHMME + SVM	0.83	0.04
MHMME	0.80	0.05
PCA + SVM	0.78	0.04
MCE-HMM	0.75	0.05
PCA + ME	0.73	0.05
Gate	0.71	0.05
CI-HMM	0.70	0.02
Experts	0.61	0.02

2012-2013 + Extension

Multiple-instance learning (MIL) for landmine detection


One of the problems central to landmine detection with ground-penetrating radar (GPR) is precisely labeling data—or, conversely, learning from data that is difficult to label. For each target, a significant amount of data is collected about the target and its surroundings. Much of the data collected is not indicative of target presence whatsoever, but instead corresponds to other portions of the scene, such as the surrounding soil and clutter. However, all of this data is typically labeled as corresponding to a target. Including all of the data under one label inflates the dimensionality of the data, which can severely inhibit learning. It would be preferable to instead label the small fraction of the data that actually indicates target presence as such, and label the rest as indicating non-target. But prior to the development of a learning algorithm capable of doing such a labeling automatically, this process is extremely time-consuming.

A better approach might be to frame the learning problem to handle the uncertainty surrounding the labeling. Multiple-instance learning (MIL) is a framework that specifically tries to learn to label the smaller, unlabeled individual elements that make up a set of larger, labeled elements.

Imagine a collection of data samples that are observed simultaneously, such as a frame of GPR data. If one element in the collection indicates target presence, like a small patch of pixels within the frame, the entire collection is labeled as a target. If


no elements in the collection indicate target presence, the entire collection is labeled as a non-target.

Within the training data, the collection labels are given, and the goal is to learn the smaller sample labels. In the terminology of MIL, the collections are *bags* and the samples are *instances*.




MIL is in some sense a feature-learning algorithm. It learns what it is about the data that indicates target presence. But what it learns is actually what *subset* of a large collection of data is useful or correlates with target presence. This framework captures the nature of GPR data for landmine detection very well.

In our setup, we let a frame of GPR be a bag, and small overlapping windows of data be the instances to be labeled.



During the reporting period, we developed an extension to MIL that incorporates it into a hidden Markov model (HMM) framework, which we refer to as the MI HMM model. HMMs learn to characterize sequences of data. When originally proposed, HMMs for landmine detection with GPR became the state-of-the-art GPR-based

detector. However, the same problems exist for the HMM learning algorithm as for all learning algorithms with GPR, which can only be automatically labeled at the frame-level. Thus, training an HMM within an MIL framework promises to learn an effective HMM model without the need for time-consuming and possibly error-prone manual labeling of sequence data.



In our experiments, we found that not only did the MI-HMM facilitate HMM model training with labels applied only to frames and not individual sequences within the frame, its performance was significantly better than the standard HMM.


We also investigated creating an improved Multiple Instance Relevance Vector Machine (MI-RVM) by combining the fastest known MI-RVMs with a faster sequential learning algorithm developed for standard RVMs. This is important because MI-RVMs typically suffer from a very high computational burden, particularly with high-dimensional GPR data. We found that we were able to develop a much faster MI-RVM while maintaining the same performance level as the less-optimized version. This will help to further the MI-RVM as a feasible option for landmine detection into the future.

Sweep detection and alignment in hand-held ground penetrating radar data

Landmine detection with hand-held sensors predates vehicle-based systems but offers distinct challenges with respect to the use of ground-penetrating radar (GPR). A vehicle-based system is steady and rigid, collecting data at regularly-spaced intervals, which allows for clear 2D snapshots of the scene, as well as the ability to aggregate the 2D snapshots into 3D volumes of data. But in a hand-held system that moves freely, producing even 2D snapshots can be challenging. Consequently, algorithm development with vehicle-based GPR sensors has achieved a higher degree of success.


In this work, we aim to adapt hand-held systems to utilize the advancements made in vehicle-based systems. The goal is to develop a system to automatically convert a constant, unorganized 2D stream of data collected by the hand-held sensor into distinct frames that roughly correspond to successive slices of a 3D volume of data.

This relies on a hand-held system human operator collecting data as a series of *sweeps*. As the operator walks slowly and steadily down a lane, the hand-held system is swung from left-to-right in a single sweep, then right-to-left in another sweep, with each sweep progressively further down the lane. The sweeps are depicted in the figure below.



Altogether, the data is collected as a 2D stream. The first goal is to segment the 2D stream into sweeps, representing 2D snapshots of the scene. This allows for a number of vehicle-based GPR algorithms to be utilized for detection with the hand-held system. This is the first stage of our procedure—the *sweep detection* stage. These snapshots can also be collected and organized into 3D volumes if desired, which corresponds to the *sweep alignment* stage that we explain later on.

Below we depict three 2D streams, corresponding to three channels of data collected simultaneously by a hand-held system. These streams are segmented into frames by detecting the beginning and end of each sweep, and the resulting segment boundaries are shown.



Once the frames are created, they must be aligned so that they can be organized into a 3D volume, which is a non-trivial task. First, every even-numbered frame needs to be reversed, since the successive sweeps are collected in opposite directions. Then, the frames need to be adjusted to account for the different speeds and slight positional offsets at which the sweeps can occur, which create frames that are not the same length. To do this, we use dynamic time warping (DWT) and interpolate to fill in the gaps. This creates frames that are all the same length, which can then be organized into a 3D volume of data