

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) 14-11-2014	2. REPORT TYPE Final Report	3. DATES COVERED (From - To) 1-Sep-2009 - 30-Sep-2013
---	--------------------------------	--

4. TITLE AND SUBTITLE Final Report: Statistical Signal Processing for Remote Sensing of Targets: Proposal for Terrestrial Science Program	5a. CONTRACT NUMBER W911NF-09-1-0487
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER 611102

6. AUTHORS Leslie M. Collins, Peter A. Torrione, Kenneth D. Morton	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAMES AND ADDRESSES Duke University 2200 W Main St Ste 710 Durham, NC 27705 -4677	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) 56551-CS.18

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 LIBS Laser induced breakdown spectroscopy (LIBS) provides a unique technology for rapid elemental analysis of arbitrary objects, with limited or no sample preparation. One area where LIBS has previously shown promise is in the identification and classification of chemical, biological, radiological, and nuclear explosive (CBRNE) threats. Such "white powder" analysis can have applications in airport and port security, as well as in active conflict.

15. SUBJECT TERMS landmine, machine learning, detection, radar, EMI, acoustic, LIBS
--

16. SECURITY CLASSIFICATION OF:	17. LIMITATION OF ABSTRACT	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT UU	UU		Leslie M. Collins
b. ABSTRACT UU			19b. TELEPHONE NUMBER 919-660-5260
c. THIS PAGE UU			

Report Title

Final Report: Statistical Signal Processing for Remote Sensing of Targets: Proposal for Terrestrial Science Program

ABSTRACT

LIBS

Laser induced breakdown spectroscopy (LIBS) provides a unique technology for rapid elemental analysis of arbitrary objects, with limited or no sample preparation. One area where LIBS has previously shown promise is in the identification and classification of chemical, biological, radiological, and nuclear explosive (CBRNE) threats. Such “white-powder” analysis can have applications in airport and postal security, as well as in active conflict zones to scan for trace elements on vehicles or clothing, for example.

In many realistic fielded scenarios, techniques like LIBS are expected to discriminate CBRNE threats (the analytes) in trace amounts or fine powders on a wide varieties of backgrounds (the substrates). These substrates can include a wide array of materials including car door handles, clothing, or any material where traces of CBRNE might be identified. Like many other spectroscopic techniques, LIBS responses can be highly sensitive to the underlying sample substrate, and this substrate response can overwhelm responses from any analyte under consideration. Figure 1 illustrates this behavior.

Figure 1: LIBS spectra from OVA (ricin simulant) on several different substrates: steel, aluminum, and polycarbonate. Although the analyte is the same on each of these substrates, the spectra is dominated by the substrate response.

In this work we explored the application of context-dependent learning to improve the classification of arbitrary analytes on a set of known background substrates. Building upon our previous work in GPR [Ratto], and in contrast to previous work with LIBS spectra [Gottfried], we showed that by separating the classification problem into multiple simpler per-substrate classification problems, overall analyte classification results could be improved from 73% correct to 84% correct.

Although this work illustrated the improvements possible via context-dependent learning for mitigating interference from various substrates, significant additional research has further improved on these results. The context dependent framework described here presumes a closed set of possible substrates, and assumes that the available training data spans the space of possible substrates. In contrast to this approach, our more recent investigations into learning with LIBS data utilizes variational Bayesian techniques to learn the distinct responses from each analyte, even when the training data only provides samples of each analyte in a mixture with a substrate. This ongoing research should help to further improve classification of spectral data in a wide range of applications, not only for LIBS based classification of CBRNE threats.

The majority of algorithms for spectroscopic signal discrimination and classification inherently assume that the data is well modeled as a low dimensional representation projected onto a linear subspace. Detection algorithms based on this assumption include partial least squares discriminant analysis (PLSDA) and principal components analysis (PCA). These algorithms determine a set of factors that can be thought of as base spectra. These factors can then be weighted and combined to approximate all of the observations within the spectra. It is well known LIBS spectra are sparse in nature, meaning that most of the wavelengths in the spectra have no information. It is also well understood a peak in LIBS spectrum at a particular wavelength contributes energy into the neighboring wavelengths with a shape known as a Voigt response. The standard algorithms for estimating the spectral factors for PLSDA and PCA do not exploit the inherent sparsity in the spectra and although there are methods to enforce sparsity in the learned PCA or PLSDA parameters, these methods cannot model the presence of the Voigt response in the spectra. We developed a model for spectroscopic signals that directly models the sparse nature of the peaks as well as the Voigt response. This model has been shown to be more robust than standard PLSDA by exploiting sparse nature of the spectra, as well as more robust than the standard sparse PLSDA by learning an even more sparse representation by exploiting knowledge of the underlying physics.

Standoff detection of trace amounts of explosive threats using spectroscopic methods such as LIBS requires an algorithm to detect the presence of the threat (the residue) on an unbounded set of potential surfaces (substrates). Given a collection of training data representing trace amounts of residues of interest on various surfaces it is possible to use the standard spectroscopic approaches of PLSDA and PCA to generate algorithms for residue detection. However giving the limited resources for training data collection, it is unlikely that all potential substrates will be adequately represented in the training data. PLSDA and PCA each learn a set of spectral factors so as to represent the entire collection of training data. Therefore the learned factors are representative of both the residues in the training data as well as the substrates. It has been shown PLSDA performs near chance levels when the model is trained using training data with several substrates but is evaluated using data collected on previously unseen substrates. To enable standoff residue detection that is capable of detecting a residue on substrates not contained in the training data, we developed multi-label PCA (MLPCA). MLPCA exploits the knowledge of the residue and the substrate of each spectrum in the training data, to learn a set of spectral factors to represent each residue and each substrate. A detection algorithm can then be developed using the factors representing only the residue interest. We demonstrated that algorithms based on MLPCA are capable of outperforming standard techniques for residue detection when the underlying substrate is uncertain.

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)

Received Paper

TOTAL:

Number of Papers published in peer-reviewed journals:

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

Received Paper

TOTAL:

Number of Papers published in non peer-reviewed journals:

(c) Presentations

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

<u>Received</u>	<u>Paper</u>
08/21/2013 1.00	Christopher Ratto, Peter Torrione, Kenneth Morton, Leslie Collins. A Hidden Markov Context Model for GPR-Based Landmine Detection Incorporating Dirichlet Process Priors, International Conference on Geoscience and Remote Sensing. 14-JUL-11, . . . ,
08/21/2013 2.00	Kenny Morton, Leslie Collins, Pete Torrione, Stacy Tantum. Investigatyion of the Effects of Operator Technique on Handheld Sensor Data for Landmine Detection, SPIE Defense and Security. 23-APR-12, . . . ,
08/21/2013 3.00	Kenneth Morton, Christopher Ratto, Jordan Malof, Michael Gunter, Leslie Collins, Peter Torrione. Change Based Threat Detection in Urban Environments with a Forward Looking Camera, SPIE Defense and Security. 21-MAY-12, . . . ,
08/21/2013 4.00	Jordan Malof, Kenneth Morton, Leslie Collins, Peter Torrione. Processing forward looking data for anomaly detection: single look, multi look, and spatial classification, SPIE Defense and Security. 21-MAY-12, . . . ,
08/21/2013 5.00	Christopher Ratto, Kenneth Morton, Ian McMichael, Brian Burns, William Clark, Leslie Collins, Peter Torrione. Integration of LIDAR wit the NIITEK GPR for Improved Performance on Rough Terrain, SPIE Defense and Security. 21-MAY-12, . . . ,
08/21/2013 6.00	Rayn Sakaguchi, Kenneth Morton, Leslie Collins, Peter Torrione. Keypoint Based Image Processing for Landmine Detection, SPIE Defense and Security. 21-MAY-12, . . . ,
08/21/2013 7.00	Achut Manandhar, Kenneth Morton, Leslie Collins, Peter Torrione. Multiple Instance Learning for Landmine Detection using GPR, SPIE Defense and Security. 21-MAY-12, . . . ,
08/21/2013 8.00	Christopher Ratto, Kenneth Morton, Leslie Collins, Peter Torrione. A Bayesian Method for Discriminative Context Dependent Fusion of GPR Based Detection Algorithms, SPIE Defense and Security. 21-MAY-12, . . . ,
11/14/2014 9.00	Jordan Malof, Kenneth Morton, Leslie Collins, Peter Torrione. Fusion of forward-looking infrared and ground penetrating radar forimproved stopping distances in landmine detection, SPIE Defense. . . . ,
11/14/2014 10.00	Jordan Malof, Kenneth Morton, Leslie Collins, Peter Torrione. A novel algorithm for buried target detection, evaluated on a collection of seismo-acoustic data, SPIE Defense. . . . ,
11/14/2014 13.00	Daniel Reichmann, Leslie Collins, Peter Torrione, Kenneth Morton. Target Localization and Signature Extraction in GPR DataUsing Expectation-Maximization and Principal ComponentAnalysis, SPIE Defense. . . . ,
11/14/2014 14.00	Mary Knox, Stacy Tantum, Leslie Collins. Target detection and identiφcation using synthetic apertureacoustics, SPIE Defense. . . . ,
TOTAL:	12

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

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

Received Paper

TOTAL:

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

(d) Manuscripts

Received Paper

- 11/14/2014 12.00 Christopher Ratto, Peter Torrione, Kenneth Morton, Leslie Collins. Analysis of Linear Prediction in GPR Data for Soil Characterization in Countermining Applications, Sensing and Imaging: An International Journal (04 2013)
- 11/14/2014 15.00 Achut Manandhar, Peter Torrione, Leslie Collins, Kenneth Morton. Multiple-Instance Hidden Markov Model for GPR-Based Landmine Detection, IEEE TRANSACTIONS ON GEOSCIENCE AND REMOTE SENSING (04 2013)
- 11/14/2014 16.00 Christopher Ratto, Peter Torrione, Kenneth Morton, Leslie Collins. Bayesian Context-Dependent Learning for Anomaly Classification in Hyperspectral Imagery, IEEE TRANSACTIONS ON GEOSCIENCE AND REMOTE SENSING (12 2012)
- 11/14/2014 17.00 Peter Torrione, Rayn Sakaguchi, Leslie Collins, Kenneth Morton. Histograms of Oriented Gradients for Landmine Detection in Ground-Penetrating Radar Data, IEEE TRANSACTIONS ON GEOSCIENCE AND REMOTE SENSING (06 2012)

TOTAL: 4

Number of Manuscripts:

Books

Received Book

TOTAL:

Received Book Chapter

11/14/2014 11.00 Peter Torrione, Leslie Collins, Kenneth Morton. Multivariate analysis, chemometrics, and machine learning in laser spectroscopy, United States of America: Woodhead Publishing Limited, (10 2014)

TOTAL: 1

Patents Submitted

Patents Awarded

Awards

None

Graduate Students

<u>NAME</u>	<u>PERCENT SUPPORTED</u>	<u>Discipline</u>
Christopher Ratto	1.00	
Rayn Sakaguchi	1.00	
Kenneth Morton	1.00	
Jordan Malof	1.00	
FTE Equivalent:	4.00	
Total Number:	4	

Names of Post Doctorates

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

Names of Faculty Supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>	National Academy Member
Leslie Collins	0.10	
Peter Torrione	1.00	
Kenneth Morton	0.75	
FTE Equivalent:	1.85	
Total Number:	3	

Names of Under Graduate students supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>	Discipline
Patrick Wang	0.50	BME
FTE Equivalent:	0.50	
Total Number:	1	

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: 1.00

The number of undergraduates funded by this agreement who graduated during this period with a degree in science, mathematics, engineering, or technology fields:..... 1.00

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:..... 1.00

Number of graduating undergraduates who achieved a 3.5 GPA to 4.0 (4.0 max scale):..... 1.00

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>	
Chris Ratto	
Rayn Sakaguchi	
Jordan Malof	
Total Number:	3

Names of personnel receiving PHDs

<u>NAME</u>	
Chris Ratto	
Kenneth Morton	
Peter Torrione	
Total Number:	3

Names of other research staff

NAME

PERCENT SUPPORTED

FTE Equivalent:

Total Number:

Sub Contractors (DD882)

Inventions (DD882)

Scientific Progress

See attached

Technology Transfer

HoG code given to NVESD and Chemring/NIITEK

2013 Final Project Report

Statistical Signal Processing for Remote Sensing of Targets: Proposal for Terrestrial Science Program

Grant Number 56551CS

Leslie M. Collins, Peter A. Torriane, Kenneth D. Morton
Electrical and Computer Engineering Department
Duke University

Abstract

The fundamental objectives of this work were to (1) develop new algorithms for both prescreening and discrimination to support the HMDS; (2) assess utility of information-theoretic approaches we have been developing under the MURI and under follow on ARO support for consideration by the various NVESD programs, particularly handheld systems; (3) develop algorithms for forward looking radar systems and assess their performance; (4) develop algorithms for emerging sensors, such as side aperture acoustic and laser induced breakdown spectroscopy (LIBS). In particular, for each system of interest we are carefully considering robustness issues, particularly with respect to terrain, speed, antenna motion, etc. We work closely with other groups supported on these efforts to insure both collaborative algorithm development (algorithm fusion) and technology transfer. In particular, code for several feature-based classification algorithms and a terrain tracking algorithm was transferred to Chemring, manufacturers of the HMDS system.

Technical Achievements

All of the details and results of the various projects are provided in the SPIE publications accompanying this submission.

LIBS

Laser induced breakdown spectroscopy (LIBS) provides a unique technology for rapid elemental analysis of arbitrary objects, with limited or no sample preparation. One area where LIBS has previously shown promise is in the identification and classification of chemical, biological, radiological, and nuclear explosive (CBRNE) threats. Such "white-powder" analysis can have applications in airport and postal security, as well as in active conflict zones to scan for trace elements on vehicles or clothing, for example.

In many realistic fielded scenarios, techniques like LIBS are expected to discriminate CBRNE threats (the analytes) in trace amounts or fine powders on a wide varieties of backgrounds (the substrates). These substrates can include a wide array of materials including car door handles, clothing, or any material where traces of CBRNE might be identified. Like many other spectroscopic techniques, LIBS responses can be highly sensitive to the underlying sample substrate, and this substrate response can overwhelm responses from any analyte under consideration. Figure 1 illustrates this behavior.

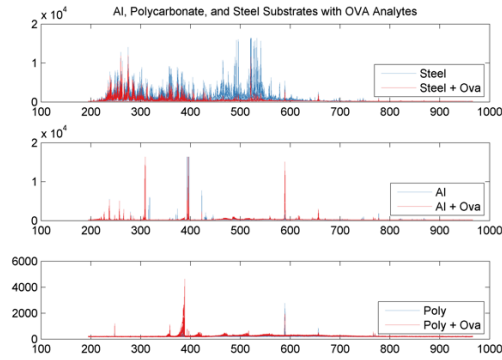


Figure 1: LIBS spectra from OVA (ricin simulant) on several different substrates: steel, aluminum, and polycarbonate. Although the analyte is the same on each of these substrates, the spectra is dominated by the substrate response.

In this work we explored the application of context-dependent learning to improve the classification of arbitrary analytes on a set of known background substrates. Building upon our previous work in GPR [Ratto], and in contrast to previous work with LIBS spectra [Gottfried], we showed that by separating the classification problem into multiple simpler per-substrate classification problems, overall analyte classification results could be improved from 73% correct to 84% correct.

Although this work illustrated the improvements possible via context-dependent learning for mitigating interference from various substrates, significant additional research has further improved on these results. The context dependent framework described here presumes a closed set of possible substrates, and assumes that the available training data spans the space of possible substrates. In contrast to this approach, our more recent investigations into learning with LIBS data utilizes variational Bayesian techniques to learn the distinct responses from each analyte, even when the training data only provides samples of each analyte in a mixture with a substrate. This ongoing research should help to further improve classification of spectral data in a wide range of applications, not only for LIBS based classification of CBRNE threats.

The majority of algorithms for spectroscopic signal discrimination and classification inherently assume that the data is well modeled as a low dimensional representation projected onto a linear subspace. Detection algorithms based on this assumption include partial least squares discriminant analysis (PLSDA) and principal components analysis (PCA). These algorithms determine a set of factors that can be thought of as base spectra. These factors can then be weighted and combined to approximate all of the observations within the spectra. It is well known LIBS spectra are sparse in nature, meaning that most of the wavelengths in the spectra have no information. It is also well understood a peak in LIBS spectrum at a particular wavelength contributes energy into the neighboring wavelengths with a shape known as a Voigt response. The standard algorithms for estimating the spectral factors for PLSDA and PCA do not exploit the inherent sparsity in the spectra and although there are methods to enforce sparsity in the learned PCA or PLSDA parameters, these methods cannot model the presence of the Voigt response in the spectra. We developed a model for spectroscopic signals that directly models the sparse nature of the peaks as well as the Voigt response. This model has been shown to be more robust than standard PLSDA by exploiting sparse nature of the spectra, as well as more robust than the standard sparse PLSDA by learning an even more sparse representation by exploiting knowledge of the underlying physics.

Standoff detection of trace amounts of explosive threats using spectroscopic methods such as LIBS requires an algorithm to detect the presence of the threat (the residue) on an unbounded set of potential surfaces (substrates). Given a collection of training data representing trace amounts of residues of interest on various surfaces it is possible to use the standard spectroscopic approaches of PLSDA and PCA to generate algorithms for residue detection. However giving the limited resources for training data collection, it is unlikely that all potential substrates will be adequately represented in the training data. PLSDA and PCA each learn a set of spectral factors so as to represent the entire collection of training data. Therefore the learned factors are representative of both the residues in the training data as well as the substrates. It has been shown PLSDA performs near chance levels when the model is trained using training data with several substrates but is evaluated using data collected on previously unseen substrates. To enable standoff residue detection that is capable of detecting a residue on substrates not contained in the training data, we developed multi-label PCA (MLPCA). MLPCA exploits the knowledge of the residue and the substrate of each spectrum in the training data, to learn a set of spectral factors to represent each residue and each substrate. A detection algorithm can then be developed using the factors representing only the residue interest. We demonstrated that algorithms based on MLPCA are capable of outperforming standard techniques for residue detection when the underlying substrate is uncertain.

Handheld Landmine Detection Systems

Ground penetrating radar (GPR) is a commonly employed sensing modality for landmine detection. It has been successfully deployed in vehicular systems, and is also being integrated into handheld systems. Handheld mine detection systems are typically deployed in situations where either the terrain or mission renders a vehicular-based system less effective. Handheld systems are often more compact and maneuverable, but quality of the sensor data may also be more dependent on the operators experience with and technique in using the system. In particular, the sensor height with respect to the air-ground interface may be more variable than with a vehicular-based system. This variation in sensor height above the air-ground interface may have the potential to adversely affect mine detection performance with the GPR sensing modality. In this work, the effects of operator technique on handheld sensor data quality was investigated, and ground alignment was explored as a potential approach to reducing variability in the sensor data quality due to operator technique.

The handheld sensor used in this analysis is a two-channel frequency-domain ground-penetrating radar (GPR) sensor. The two channels perform similarly, so without loss of generality the focus here is on the front channel. The frequency domain data for each scan (the response due to a single electromagnetic pulse) is inverse transformed to create time-domain range data, termed an A-scan. Although each A-scan is a function of time, time roughly corresponds to depth under the assumption of a constant wave speed. Thus, the time at which each reflection arrives at the sensor roughly corresponds to the depth of the electrical discontinuity from which it was generated. The A-scans are measured at regular time intervals resulting in a variable spatial resolution of the measured data that depends on the speed at which the sensor is swept across the ground. The series of A-scans, termed a B-scan, forms an image which is utilized for automated landmine detection. In the case of handheld data, the portion of the B-scan that constitutes a single pass over a target is termed a sweep. The operator uses a marking button to denote the beginning and end of a sweep over a target. Assuming a constant sweep speed (regular spatial sampling), a point (or sufficiently small) dielectric discontinuity produces a hyperbolic signature in the B-scan image. This hyperbolic signature forms the basis for many GPR landmine detection algorithms, including hidden Markov models, edge histogram descriptor, and texture feature coding method.

In this study we explored some of the factors related to operator technique with a handheld GPR sensor that may be related to variations in performance often observed across operators. Operator technique can differ in many ways, including sweep speed, rate of forward progress, and the consistency with which the sensor head remains parallel to the ground at a constant height above the ground. Quantitative measures suggest that the three operators who collected the data utilized in this study employed three distinctive techniques: faster sweep speed and faster rate of forward progress, slower sweep speed and slower rate of forward progress, and slower swing speed and faster rate of forward progress. Since sensor height variations have been shown to adversely impact vehicular-based GPR systems,⁶ data alignment was considered as an approach to mitigate variations in the location of the air-ground interface that may result from how the sensor head is positioned above the ground. Although no clear trends emerged, this initial exploration into some of the factors related to operator technique that may explain some of the variations in performance across operators suggests some avenues for future research which may result not only in algorithm improvements, but may also in suggestions for best practices in data collection with handheld sensors.

Change Detection in Video for IED Detection

As the nature of hidden explosive threats continues to evolve, direct detection of many threat emplacements is becoming more difficult. On commonly traversed routes, change detection may provide an efficient technique for sharing information amongst and between convoys and detecting threats using a combination of automated processing and operator situational awareness. This work has expanded upon previous approaches to automated change detection by considering more challenging urban environments and comparing and contrasting different distance metrics for temporal alignment and change confirmation. This work indicates that for the data under consideration, pre-screener performance is relatively invariant to the specific inter-frame distance used to determine optimal temporal warping. This work also showed that although SAD distances are simple to compute, they are not suitable for change confirmation for even relatively slight changes in ambient lighting. Other distance metrics, such as normalized cross-correlation and BRIEF features provide much more robust metrics for change confirmation, and fusion of these distance metrics using standard machine learning techniques can further improve performance. Furthermore, optimization of the effective field of view and conservative thresholding can be applied to system outputs to efficiently flag significant changes at very low false alarm rates, as shown in <http://youtu.be/X24AVXtitLM>. Several modifications to the existing processing streams can be undertaken to further improve performance. First, our future work for change detection in forward looking video will attempt to mitigate the causal requirements in the temporal-warping stages of this processing. Also, throughout this work, no attempts were made to take advantage of the camera parameters or any possible inference regarding object distances. Incorporating this information into the processing could help with scale selection for change detection and automatic determination of local neighborhood sizes for template matching.

Downward Looking Ground Penetrating Radar

There were several projects pursued (by different students) in support of the HMDS downward looking GPR System. This included context dependent processing, prescreening efforts, investigation of new features and algorithms, and incorporating LIDAR into the sensing system to improve ground tracking. Vehicle-mounted ground-penetrating radar (GPR) has proved to be a valuable technology for buried threat detection, especially in the area of military route clearance. However, detection performance may be degraded in very rough terrain or off road conditions. This is because the signal processing approaches for target detection in GPR first identify the ground reflection in the data, and then align the

data in order to remove the ground Reflection. Under extremely rough terrain, antenna bounce and multipath effects render finding the ground reflection a difficult task, and errors in ground localization can lead to data alignment that distorts potential target signatures and/or creates false alarms. In this work, commercial-off-the-shelf light detection and ranging (LIDAR), global positioning system (GPS), and inertial measurement unit (IMU) were integrated with a GPR into a prototype route clearance system. The LIDAR provided high-resolution measurements of the ground surface profile, and the GPS/IMU recorded the vehicle's position and orientation. Experiments investigated the applicability of the integrated system for finding the ground reflection in GPR data and decoupling vehicle motion from the rough surface response. Assessment of ground-tracking performance was based on an experiment involving three prepared test lanes, each with different configurations of buried targets and terrain obstacles. Several algorithms for target detection in GPR were applied to the data, both with traditional preprocessing and incorporating the LIDAR and IMU. Experimental results suggest that the LIDAR and IMU may be valuable components for ground tracking in next-generation GPR systems.

The success of techniques such as the HMM and EHD illustrate the potential for statistical modeling and machine learning to provide robust results for object detection in GPR data.

In a parallel vein, modern advances in the field of computer vision also make use of statistical object descriptors coupled with machine learning to develop accurate algorithms for both instance and category recognition. Notable examples include: 1) the scale invariant feature transform (SIFT) and a related method; 2) speeded up robust features (SURF) which provide a low-dimensional representation of visual images for instance matching with images; 3) Viola–Jones' cascade of simple features used for robust face detection; and 4) histogram of oriented gradients (HOG) which has been successfully used for pedestrian detection. In the context of subsurface threat detection, HOG features are particularly interesting since analysis of the literature illustrates that effective feature extraction techniques from two disparate image processing fields (MPEG encoding and pedestrian detection in images) share very similar underlying mathematical and theoretical structure. These features are also widely utilized in the field of computer vision and video processing, so are necessarily computationally inexpensive and efficient. This observation naturally leads to the question: can HOG and other tools from the computer vision literature be successfully brought to bear on the problem of landmine detection in GPR data?

The size and scope of modern computer vision techniques is enormous, we initially focused only on the specific case of HOG features. The rationale for an initial focus on HOG features is motivated by the similarity of HOG to preexisting feature extraction techniques [e.g., edge histogram descriptor (EHD)], and also to several aspects of HOG features which make them particularly well suited to landmine detection in GPR. For example, unlike SIFT, SURF, and others, HOG features are typically applied as a dense feature extraction technique. This enables HOG features to be extracted without previous interest point identification, which is an unsolved though interesting problem in the GPR literature. Furthermore, unlike many modern feature descriptors, HOG features are fairly robust to moderate changes in object location within a HOG window. This enables the use of preexisting techniques for depth-binning in GPR classification. Our results indicated that HOG features, coupled with off-the-shelf classification techniques, can enable significant improvements in target classification.

Historically, for vehicle-mounted GPR processing to maintain speeds commensurate with operationally relevant rates of advance, computationally simple algorithms were utilized to rapidly process large volumes of GPR data and down-select locations of interest, alarms, which were then utilized in more advanced feature-based processing. This separation between prescreening and classification has always seemed artificial, as it could be advantageous to incorporate advanced statistical models into prescreening if feature extraction and classification could both be accomplished in a computationally

tractable manner. This paper is the first to develop a technique for threat identification in vehicle-mounted GPR data that is fast enough and robust enough to operate as a prescreener (on every pixel) versus only at previously flagged locations of interest. The primary novel advances contained in this paper include the first application of techniques explicitly developed for feature extraction and object classification techniques from the the development of a series of preprocessing steps suitable for simple target extraction and classifier training across multiple feature extraction approaches (e.g., HOG, and EHD), and the development of a novel joint prescreening/classification algorithm for rapid pre-screening.

Our results indicated that the performance of the HOG algorithms derived in this work enable significant performance improvements over existing methods, and that HOG might provide a useful technique to help break down the artificial divide between prescreening and target classification used in GPR systems due to computational bottlenecks. This code was provided to the government as well as NIITEK/Chemring. We also developed a number of preprocessing and classification steps that enabled modifications to the EHD algorithm to overcome some previous limitations (e.g., manual target localization, special-purpose classification algorithms) and enabled simpler algorithm retraining. A comparison of EHD and HOG performance illustrated the importance that preprocessing steps have on algorithm performance and the difficulty in providing accurate feature-to-feature performance comparisons. From a computational perspective, the HOG feature extraction procedure is quite tractable; for example, HOG features are widely used in image and video processing, where algorithms using these features are capable of achieving realtime performance. Furthermore, several open-source implementations of the HOG feature extraction code are available.

Based on our work, interesting avenues for further research and development are possible. First, throughout the training process, a rather ad hoc energy-based technique to localize actual responses of interest was utilized. Optimal extraction of locations of interest to identify and classify target versus nontarget responses is an unsolved problem in GPR processing. In the computer vision literature, corner points are often used to find regions of interest for additional processing—could a similar motivation be used to identify interesting locations in GPR data? Also, during classifier development and execution, other ad hoc techniques to aggregate decision metrics from sets of images are required (e.g., average the top three confidences). This process is necessary because the underlying problem being solved is not a simple pattern recognition problem but can instead be well represented as a multiple-instance learning (MIL) problem (see below). Treating the learning problem in GPR as an MIL problem is a young but active area of research that is currently being explored. At a bigger picture level, HOG is certainly not the only technique available from the computer vision literature that might have implications for target classification in GPR data. It is possible that the work described here will lead to significant leveraging of other alternative approaches for class identification from the computer vision literature being applied to solve problems in GPR data. For example, numerous powerful and computationally simple feature extraction techniques such as BRIEF, random features, combinations of features, and others may enable faster and more robust target classification. Other approaches to prescreening are also possible—the Viola–Jones algorithm is of particular interest due to simple computational requirements and the inclusion of cascade detectors to help maintain rapid rates of advance. As mentioned previously, HOG is related to (but distinct from) previous approaches to target classification in GPR data, such as the EHD algorithm. As more and more image processing techniques are brought to bear on target classification, an interesting question is what properties of different techniques make them successful for processing GPR data. Our results make it clear that direct head-to-head comparison of feature extraction approaches with different preprocessing steps is quite difficult, so care should be taken to ensure that

the same preprocessing steps are utilized when attempting to make direct comparisons of feature performance.

In terms of classical machine vision properties and their relevance toward landmine detection in GPR, we hypothesize that invariance to exact target (or interest point) location and invariance to received energy will be shown to be useful properties, while invariances to rotations or affine transformations can be safely ignored or discarded. Also, novel techniques in image processing and understanding, including part-based models, may enable much more robust modeling of target response variations due to changing soil properties, target sizes, and target emplacements.

As noted above, utilizing methods from the image processing and computer vision fields has led to advances in high resolution Ground Penetrating Radar (GPR) based threat detection. By analyzing 2-D slices of GPR data and applying various image processing algorithms, it is possible to discriminate between threat and non-threat objects. In initial attempts to utilize such approaches, object instance-matching algorithms were applied to GPR images, but only limited success was obtained when utilizing feature point methods to identify patches of data that displayed landmine-like characteristics. While the approach worked well under some conditions, the instance-matching method of classification was not designed to identify a type of class, only reproductions of a specific instance. In contrast, our current approach is focused on identifying methods that can account for within-class variations that result from changing target types and varying operating conditions that a GPR system regularly encounters. Image category recognition is an area of research that attempts to account for within class variation of objects within visual images. Instead of finding a reproduction of a particular known object within an image, algorithms for image categorization are designed to learn the qualities of images that contain an instance belonging to a known class. The results of this research show that while the bag-of-words method has promise in providing additional information for the discrimination of landmines, this method currently does not adequately take advantage of the spatial layout of the characteristic shapes in a landmine response. Research in the field of image categorization is currently evolving at a very fast rate. Progress on difficult categorization datasets is growing quickly as new methods are developed. Many of these make use of the relative spatial location of the parts which make up a given object. These methods are expected to benefit this research in future studies. These methods presented investigated here provide a promising direction in understanding and processing landmine data and serve as a basis for which methods can be utilized for further research in this area.

In a separate effort, Image keypoints, which are widely used in computer vision for object matching and recognition, where they provide the best solution for matching and instance recognition of complex objects within cluttered images were considered for the landmine detection problem. Most matching algorithms operate by first finding interest points, or keypoints, that are expected to be common across multiple views of the same object. A small area, or patch, around each keypoint can be represented by a numerical descriptor that describes the structure of the patch. By matching descriptors from keypoints found in 2-D data to keypoints of known origin, matching algorithms can determine the likelihood that any particular patch matches a pre-existing template. The objective in this research is to apply these methods to two-dimensional slices of Ground Penetrating Radar (GPR) data in order to distinguish between landmine and non-landmine responses. In this work, a variety of established object matching algorithms have been tested and evaluated to examine their application to GPR data. In addition, GPR specific keypoint and descriptor methods have been developed which better suit the landmine detection task within GPR data. These methods improve on the performance of standard image processing techniques, and show promise for future work involving translations of technologies from the computer vision yield to landmine detection in GPR data.

Another effort involved investigation of novel Multiple Instance Learning approaches to GPR data processing. Although GPR data contains a representation of 3D space, during training, target and false alarm locations are usually only provided in 2D space along the surface of the earth. To overcome uncertainty in target depth location, many algorithms simply extract features from multiple depth regions that are then independently used to make mine/non-mine decisions. A similar technique is employed in hidden Markov models (HMM) based landmine detection. In this approach, sequences of downtrack GPR responses over multiple depth regions are utilized to train an HMM, which learns the probability of a particular sequence of GPR responses being generated by a buried target. However, the uncertainty in object depth complicates learning for discriminating targets/non-targets since features at the (unknown) target depth can be significantly different from features at other depths but in the same volume. To mitigate the negative impact of the uncertainty in object depth, mixture models based on Multiple Instance Learning (MIL) have previously been developed. MIL is also applicable in the landmine detection problem using HMMs because features that are extracted independently from sequences of GPR signals over several depth bins can be viewed as a set of unlabeled time series, where the entire set either corresponds to a buried threat or a false alarm. In this work, a novel framework termed as multiple instance hidden Markov model (MIHMM) was developed. This framework is suitable for GPR-based landmine detection because features generated from depth regions can be considered as sets of samples, with the entire set corresponding to a target or a false alarm. In this work, we extended the basic HMM approach to an HMM in MIL framework that inherently incorporates the ambiguity in the individual sample labels. We have evaluated the proposed MIHMM method against the HMM method on synthetic and landmine datasets. Our results indicate that the proposed approach performs better than the HMM technique. In addition, the MIHMM model parameters can be learned using Variational Bayesian methods, yielding fast and computationally tractable inference. For future work, it may be beneficial to explore the possibility of performing Markov Chain Monte Carlo (MCMC) inference to obtain better parameter estimates at the expense of additional computation. We also intend to modify our current model in two different ways. The first modification is to use a single observation model but different transition matrices and different initial state probabilities for the H1 and H0 HMMs. This adaptation will reduce the number of parameters and potentially improve parameter inference. The second modification is to use a single observation model and a mixture of transition matrices and a mixture of initial state probabilities for each of the HMMs. This alteration will potentially improve modeling different types of targets and false alarms.

Yet another effort involved evaluating the impact of incorporating contextual information into the processing architecture. Several pattern classification algorithms have been proposed and evaluated which enable GPR systems to achieve robust performance. However, comparisons of these algorithms have shown that their relative performance varies with respect to the environmental context under which the GPR is operating. Context-dependent fusion has been proposed as a technique for algorithm fusion and has been shown to improve performance by exploiting the differences in algorithm performance under different environmental and operating conditions. Early approaches to context-dependent fusion clustered observations in the joint confidence space of all algorithms and applied fusion rules within each cluster (i.e., discriminative learning). Later approaches exploited physics-based features extracted from the background data to leverage more environmental information, but decoupled context learning from algorithm fusion (i.e., generative learning). In this work, a Bayesian inference technique which combines the generative and discriminative approaches is proposed for physics-based context-dependent fusion of detection algorithms for GPR. The method uses a Dirichlet process (DP) mixture as a model for context, and relevance vector machines (RVMs) as models for algorithm fusion. Variational Bayes is used as an approximate inference technique for joint learning of

the context and fusion models. Experimental results compared the novel Bayesian discriminative technique to generative techniques developed in past work by investigating the similarities and differences in the contexts learned as well as overall detection performance.

The proposed technique drew from previous work by Frigui's group by learning contexts discriminatively, i.e. finding the contexts that allow for best overall classification. It also drew from previous work by our group by adopting a Bayesian approach and utilizing separate features for learning context and target discrimination. The proposed model for joint context learning and target classification was based on the mixture-of-experts model, but utilized a DPGMM gating network and sparse linear classifiers based on the RVM. An example was first shown to illustrate the performance of discriminative context learning on synthetic data with known context and target labels. Discriminative context learning was able to determine an effective number of contexts while jointly clustering the context features and classifying the target features. An experiment was performed on a large set of GPR data collected over a diverse target population in varying environments. Context features were extracted from prescener alarms using techniques proposed in previous work and fusion was performed on the confidence values yielded by the prescener and three feature-based algorithms for each alarm. Although a large number of contexts were identified by the proposed technique, each context yielded a unique weighting of the four algorithms. Overall discrimination performance was compared to previous context-dependent fusion techniques, as well as to conventional fusion with a linear RVM. Results indicated that discriminative context learning may be biased by easy-to-detect targets and/or clutter. Generative context learning, which is unbiased to any particular target type, may yield better improvements in discriminating more difficult observations.

In future work, first principles must be revisited to determine the conditions in which it would be appropriate to learn context generatively or discriminatively. Although previous work¹⁴ illustrated that GPR context features are indicative of a variety of quantitative soil properties, the benefit of performing classification and clustering in separate feature spaces should also be explored. It is indeed possible that good performance could be obtained by using a highly-nonlinear model on the combined target and context features. However, good performance comes at the expense of model simplicity, and future work should determine the conditions of this trade-off for context-dependent learning.

Forward Looking Infrared (FLIR)

This work investigated the value of multi-look information in FLIR-based landmine detection. In order to do this a planview framework was developed where FLIR video is decomposed based on scene-to-camera distance. Experiments were run on a large FLIR dataset using the RX detectors (RX) and the RX detector with mean-shift (MSRX) on raw FLIR data and within the plan-view framework. The following conclusions were drawn from the experimental results:

- Multi-look information can improve detection performance.
- Multi-look information from some image regions can be detrimental to detection performance.
- The MSRX detector performs better in the plan-view framework than in the original FLIR video.
- The plan-view makes more efficient use of multi-look information for detection.

Further study could investigate whether other algorithms besides RX show improvement when operated in the planview. Also, this work utilized a relatively inaccurate camera transform and further experiments with improved transforms may yield even stronger performance improvements. Finally, future studies may also address the computational costs associated with the plan-view framework. A more computationally efficient algorithm is an important step in order to exploit the benefits of the framework in real-time applications.

Electromagnetic Induction

Frequency-domain electromagnetic induction (EMI) sensors have been shown to provide target signatures which enable discrimination of landmines from harmless clutter. In particular, frequency-domain EMI sensors are well-suited for target characterization by inverting a physics-based signal model. In many model-based signal processing paradigms, the target signatures can be decomposed into a weighted sum of parameterized basis functions, where the basis functions are intrinsic to the target under consideration and the associated weights are a function of the target sensor orientation. When sensor array data is available, the spatial diversity of the measured signals may provide more information for estimating the basis function parameters. After model inversion, the basis function parameters can form the foundation of model-based classification of the target as landmine or clutter. In this work, sparse model inversion of spatial frequency-domain EMI sensor array data followed by target classification using a statistical model is investigated. Results for data measured with a prototype frequency-domain EMI sensor at a standardized test site were determined. Results obtained on data measured at a standardized test facility in a temperate climate indicate that structured RVM regression model inversion to find DSRF parameter estimates from multiple spatial measurements appears to provide sparse, stable representations of landmine targets using frequency-domain EMI sensor data. The stability is further improved when the measurements are aggregated over multiple targets of the same target type to generate a single DSRF model for a given target type. This approach does not require specifying the model order a priori, and simultaneously enforces commonality of selected relaxation frequencies across spatial measurements, whether those measurement come from a single target or multiple targets of the same type. In addition, it is a computationally efficient inversion technique that provides sparse solutions well-suited for subsequent detection and/or classification algorithms. Classification with distance-based classifiers demonstrates that training with multiple passes across a target or with data from multiple target examples and testing with data from a single target pass provides performance that is comparable to training and testing with the single maximum energy measurement in each square. It is anticipated that training with multiple target will result in target models that are more robust to variations in the target burial depth and/or orientation.