

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 the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.
PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY)		2. REPORT TYPE Final Technical Report		3. DATES COVERED (From - To) July 2016 - June 2019	
4. TITLE AND SUBTITLE Advanced Physics & Statistics-Based Algorithms for Standoff IED Detection/ Classification				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER N00014-16-1-2623	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Joseph Burns, Benjamin Hart, Joseph Lindgren, Matthew Masarik, Christopher Rickerd, Christopher Roussi, Brian Thelen, Erick Vega, and Ismael Xique				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Michigan Technological University, Michigan Tech Research Institute 3600 Green Ct., Ste. 100, Ann Arbor, MI 48105				8. PERFORMING ORGANIZATION REPORT NUMBER FR-1	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) US Department of Defense, Office of Naval Research 875 N. Randolph St., Ste. 1425 Arlington, VA 22203-1995				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for Public Release: Distribution is Unlimited.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT Standoff radar systems have the potential to detect buried hazards at safe distances. However, low target-to-clutter ratios and limited spatial resolution resulting from forward-looking measurement geometries limit the performance of automated detection and classification algorithms. The project researched and developed algorithms and measurement approaches to mitigate these limitations. Specifically, the project investigated buried object radar imaging, neural network ATR feature extraction, transfer learning to estimate algorithm extensibility, improvement of buried object imaging using quadratic lifting inversion, and buried object detection and imaging using a ground penetrating radar on a small UAS platform.					
15. SUBJECT TERMS automatic target recognition, buried objection detection, ground penetrating radar, radar imaging					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Joseph Burns, PhD
U	U	U	UU		19b. TELEPHONE NUMBER (Include area code) 734-913-6857

Advanced Physics & Statistics- Based Algorithms for Standoff IED Detection

(Grant N00014-16-1-2623)

Final Report

July 2016 – June 2019

Joseph Burns, Benjamin Hart, Joseph Lindgren, Matthew Masarik,
Christopher Rickerd, Christopher Roussi, Brian Thelen, Erick Vega,
and Ismael Xique

Table of Contents

Table of Figures	iv
1.0 Executive Summary	1
1.1 Research Areas	1
1.2 Summary of Year One Results	3
1.3 Summary of Year Two and Three Results.....	5
1.4 Recommendations for Future Work.....	7
2.0 3D Subsurface Imaging Studies	13
2.1 Problem Description	15
2.1.1 Half-Space Geometry	15
2.1.2 Scattering Equations	16
2.1.3 Spectral Green’s Functions	17
2.1.4 Computational Considerations.....	22
2.2 Approximations	23
2.2.1 Born Approximation	23
2.2.2 Scaled Born Approximation	24
2.2.3 Quadratic Approximation	30
2.2.4 “Ray” Approximation	31
2.2.5 Stationary Phase Approximation	32
2.3 3D Imaging Baseline	34
2.3.1 Ground Penetrating Radar Data Used for Testing	34
2.3.2 Comparisons of Reconstructions to Ground Truth Data.....	36
3.0 ATR Feature Extraction	46
3.1 Deep Neural Networks as Feature Extractors	46
3.1.1 Convolutional Neural Networks	46
3.1.2 Training the Network.....	47

3.2	Convolutional Neural Network Feature Extraction Experiment.....	48
3.2.1	Network Topology Used for Experiment	48
3.2.2	Hyperparameter Selection for Experiment	49
3.2.3	GPR Measurements of Buried Targets Used for Experiment.....	50
3.2.4	Visualization of Feature Maps	51
3.2.5	Detection Performance of Neural Network Features.....	55
4.0	Transfer Learning to Estimate Algorithm Extensibility.....	58
4.1	Bound Formulation	58
4.2	Simulation to Assess Bounds.....	60
5.0	Improvement of Buried Object Imaging Using Lifting Approach to Inversion.....	62
5.1	Lifting Approach to Inversion and Comparison to Alternatives.....	63
5.1.1.	Comparison of Backprojection to Quadratic Lifting Inversion	64
5.1.2	Comparison of Interferometric Imaging to Quadratic Lifting Inversion	65
5.2	Lifting Inversion Algorithm Implementation	67
5.3	Example Lifting Inversion Results	69
5.3.1	Robustness of Lifting to Phase Error and Noise	69
5.3.2	Resolution Comparable to Backprojection on Measured Data.....	72
5.3.3	Application to Buried Target Imaging.....	75
5.4	Observations on the Application of Quadratic Lifting Inversion	79
6.0	sUAS based Radar Measurements of Buried Objects.....	80
6.1	sUAS Measurement Hardware	80
6.2	Field Data Collection	83
6.3	Data Processing.....	86
6.4	Imaging Results	87
	References	91
	Appendix A. Conference Papers	96

Table of Figures

Figure 1. ATR Framework for Detection and Classification with Standoff GPR Data	1
Figure 2: Statistical Inverse-Problem Framework Flowchart.....	14
Figure 3 Half-Space Geometry.....	15
Figure 4 Air-Soil Green's Function Comparison.....	19
Figure 5 Air-Soil Green's Function Error.....	19
Figure 6 Soil-Air Green's Function Comparison.....	20
Figure 7 Soil-Air Green's Function Error.....	20
Figure 8 Soil-Soil Green's Function Comparison.....	21
Figure 9 Soil-Soil Green's Function Error.....	21
Figure 10 Born Approximation gprMax Comparison for 3cm Cube	24
Figure 11 2cm Cube Truth	26
Figure 12 2cm Cube Reconstruction (Born Approximation)	27
Figure 13 2cm Cube Reconstruction (Scaled Born Approximation).....	27
Figure 14 Born Approximation Comparison on 5cm Cube	28
Figure 15 5cm Cube Truth Data.....	29
Figure 16 5 cm Cube Reconstruction (Born Approximation)	29
Figure 17 5cm Cube Reconstruction (Scaled Born Approximation).....	30
Figure 18 GPR Measurement Geometry	31
Figure 19 Extended Ray Approximation gprMax Comparison.....	33
Figure 20 Extended Ray Approximation Error	34
Figure 21 Georgia Tech Dataset Ground Truth.....	35
Figure 22 VS-1.6 and TS-50 Ground Truth	36
Figure 23: Reconstruction of Area of Interest.....	37
Figure 24 VS-1.6 Depth Profile vs. Ground Truth.....	38
Figure 25 TS-50 Depth Profile vs. Ground Truth	39
Figure 26 Mine Simulant Depth Profile vs. Ground Truth.....	40
Figure 27 Reconstruction of VS-1.6 and TS-50 with Spatial Extent Overlays	41
Figure 28 TMA-5 with Crushed Can Ground Truth.....	42
Figure 29 Reconstruction of TMA-5 with Crushed Can	43

Figure 30 Reconstruction of TMA-5 with Crushed Can with Spatial Extent Overlay	44
Figure 31 TMA-5 Depth Profile vs. Ground Truth	45
Figure 32 Hierarchical Feature Extraction with Convolutional Neural Network (Source: Wikipedia)47	
Figure 33 Neural Network Topology Used for Experiment.....	49
Figure 34-B-Scan Containing Target	51
Figure 35 Machine-derived feature maps for the target in Figure 34.....	52
Figure 36 - Model Accuracy and Loss as a Function of Training Epochs	53
Figure 37-Impact of Early Stopping Criteria on Classifier Performance	54
Figure 38 Effects of Bagging and Undersampling on Detection Performance	55
Figure 39 Comparison of Detection Performance Using Boosting with Human-Derived Features (Boosting), Neural Network Features (Neural Network), and the Combination of Human- Derived and Neural Network Features (Boosting NN).....	57
Figure 40 Assumed Distributions for Transfer Learning Experiment.....	61
Figure 41 Error Rates for Perturbed Data.....	61
Figure 42 Windowing of data via selection matrix E: data from frequency bins in a fixed spectral window on the left (of width 11) and a fixed spatial window on the right (red is cross-correlated with green) are all included in the cross-correlations selected via E.....	66
Figure 43 Collection geometry with sinusoidal measurement error with a peak amplitude of 5 cm. The ideal and perturbed measurement locations (used for the forward model and test data synthesis, respectively) are shown, along with the true point target locations and imaging grid.	70
Figure 44 Quadratic lifting improvement (bottom) over backprojected images (top) in the presence of increasing sinusoidal measurement error (left to right from 0cm to 7.5cm) and constant 7dB SNR. Small red circles indicate truth location of the point target in each plot. Coordinates are in pixels. Images are energy normalized, with dynamic range is 30dB and peaks set (per column) by the backprojected image.	71
Figure 45 Values of the half power mainlobe width percentage reduction statistic $1-G$ for a variety of measurement error and SNR levels, showing the improvement of QLI over backprojection.	72
Figure 46 Regularized least squares does not share robustness of QLI. Images were formed from noise-free synthetic data with 5cm amplitude sinusoidal measurement errors. From left to right: backprojection using $F \dagger$, backprojection using $F \ddagger$, least	

squares optimization using $F, F \dagger$, and QLI using $F, F \dagger$. Coordinates are in pixels; dynamic range is 30dB. Each image is peak normalized.....	72
Figure 47 Collection geometry, targets, and buried mine configuration in Georgia Tech study ⁸ (images reproduced and cropped from Counts et al. 2007).	73
Figure 48 “GT Plywood” Target: Comparison of backprojected image vs. QLI reconstruction. We show horizontal slices of interest at heights near expected true height at a dynamic range of 25dB to match that of Counts et al., 2007. No post-processing/normalization has been performed to equalize color scales.....	75
Figure 49 Two-layer forward model for wave scattering from buried target that includes refraction/propagation through ground.....	76
Figure 50 Optimization behavior of QLI objective function $L\mu, E$. In order to surmount the lack of algorithm-specific preconditioning in our implementation, we select as our buried target reconstruction the estimate from the optimizer iteration where $\ X\ _2$ first begins to increase (yellow dashed line).	77
Figure 51 Buried Targets: backprojection vs QLI. We show horizontal slices of interest at nominal depths of 4, 9, and 14cm below ground-air interface. Note that targets appear a few centimeters deeper than predicted by ground truth (see Figure 47 or the Georgia Tech paper ⁸), and that there is no target at a depth of 9cm.	78
Figure 52 Bergen Quad 8 with radar and array antennas attached ready for flight.....	81
Figure 53 The sUAS with four Vivaldi element array flying over beach measurement location (top left), and the Vivaldi array is shown attached to radar (bottom left) along with the Akela antenna array (bottom right).	83
Figure 54 Overhead view of the beach measurement area before the data collection.....	84
Figure 55 Flight path of a typical data collection (blue line). The positions of targets are indicated by the red dots.	85
Figure 56 Data processing pipeline	86
Figure 57 Data processing pipeline	87
Figure 58 (left) Location of targets. (Right) 2D maps of the 3D image energy integrated over height are shown for regions around each target.....	88
Figure 59 Contour plots of the 3D images generated about each target are shown. Coordinates relative to ground-truth positions.	89

1.0 Executive Summary

This report summarizes technical activities performed by the Michigan Tech Research Institute (MTRI) on the Advanced Physics and Statistics-Based Algorithms for Standoff IED Detection program (Grant N00014-16-1-2623) for the period July 2016-June 2019.

1.1 Research Areas

Standoff, low-frequency, wideband ground penetrating radar (GPR) systems have recently been developed and offer the potential to find buried explosive hazards at safe distances. Low target-to-clutter ratios and limited spatial resolution resulting from forward-looking measurement geometries at low grazing angles limit the performance of automated detection and classification algorithms on these data. The objectives of the Advanced Physics and Statistics-Based Algorithms for Standoff IED Detection program are to research and develop algorithms to overcome these measurement limitations and to generalize machine learning algorithms to fundamentally improve automatic target recognition (ATR) performance with standoff detection explosive hazard detection systems.

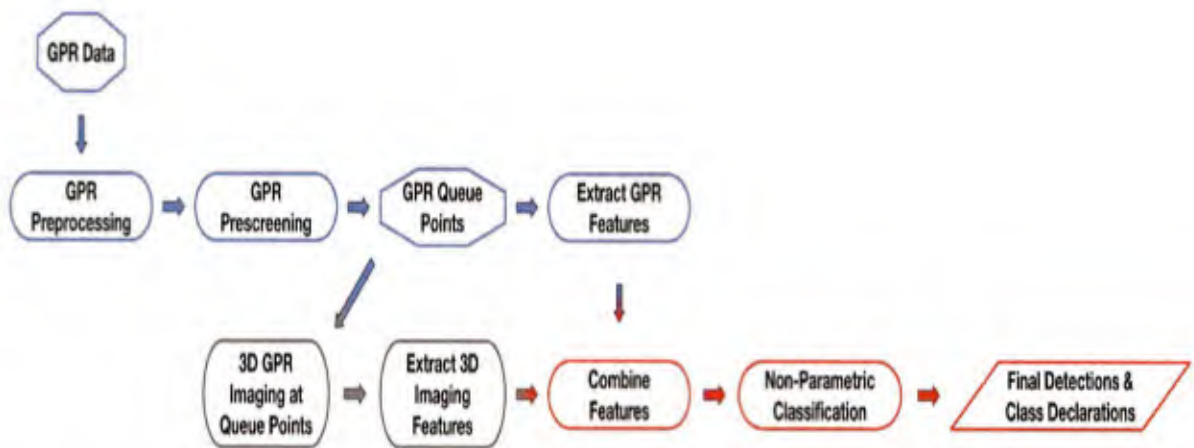


Figure 1. ATR Framework for Detection and Classification with Standoff GPR Data

MTRI has successfully demonstrated algorithms for improvised explosive device (IED) detection and classification on both forward-looking ground penetrating radar (FLGPR) systems and combined handheld GPR (HHGPR) / swept-frequency metal detector systems using the general framework shown in Figure 1 [Masarik et. al., 2016]. However,

previous experience indicates that significant performance advances in IED ATR will require significant advances in several key, basic research topic areas, which are summarized in Table 1.

Table 1. Research Areas

Research Areas	Research Topic	Benefit to Standoff ATR
1	Develop physics-based and statistics-based data filtering schemes to reject clutter and enhance target-to-clutter ratio in data and/or images	<i>Enhancement of target-to-clutter ratio will increase detection</i>
2	Develop high-fidelity 3D subsurface image reconstruction for physical object property estimation (e.g. shape, dielectric contrast).	<i>Target features extracted from 3D subsurface images enhance standoff ATR, specifically estimation of material properties may provide the ability to classify objects and reject confusers</i>
3	Identify methods of extracting more informative features from data that will be used to perform detection and classification	<i>More informative features can improve detection/classification performance</i>
4	Develop algorithms that exploit feature sets to yield optimal detection and classification.	<i>Overcoming issues with machine learning, such as training issues due to imbalanced training datasets, can improve detection/classification performance</i>
5	Develop quantitative bounds on ATR performance in terms of target vs. clutter separability bounds and target class separability bounds.	<i>Provides bounds on potential detection performance as a function of target, clutter and sensor parameters, and thus can guide sensor/collection design</i>
6	Develop quantitative understanding of extensibility of trained classifier to new environments and new targets.	<i>Provides bounds on potential detection performance to estimate the extensibility of a trained algorithm to new data/measurements</i>

Note that several of these topic areas are fundamentally related (for example, some features may generalize better to new environments and targets than others; or an initial clutter rejection scheme can help define a useful feature set to aid in target vs. clutter discrimination). Hence, MTRI believes that revolutionary performance enhancements in standoff IED ATR can only be obtained if these topics are studied jointly to understand and exploit the interplay between each area. The research performed during the reporting period is summarized in the forthcoming sections.

1.2 Summary of Year One Results

As summarized in Table 2, much of the technical activity during reporting period was focused on investigating penetrative three-dimensional (3D) radar imaging with the goal of producing high fidelity representations of buried targets, thus enabling a significant improvement in the downstream ATR processing. Specifically, alternate formulations of the forward operator, specifically spectral representations, which incorporate more physics, were investigated to provide a formulation that supports the estimation of buried object material properties. Features derived from these images, namely location, shape and material properties, can then be exploited in algorithms to produce automated detection and classification.

Table 2. Technical Activity During Reporting Period in the Research Areas

Research Areas	Research Topic	Activity During Period
1	Develop physics-based and statistics-based data filtering schemes to reject clutter and enhance target-to-clutter ratio in data and/or images	Identified data to evaluate RPCA for FLGPR geometries
2	Develop high-fidelity 3D subsurface image reconstruction for physical object property estimation (e.g. shape, dielectric contrast).	<p>Investigated development of a 3D subsurface imaging algorithm incorporating a higher fidelity forward model using a spectral representation to provide a higher fidelity images and material estimates.</p> <ul style="list-style-type: none"> - Spectral representation was too computationally intensive, so have moved to an 3D algorithm using an asymptotic representation - Evaluated its ability to estimate 3D position and shape of buried objects using available laboratory data - Investigated use of regularized imaging with L1 and TV constraints
3	Identify methods of extracting more informative features from data that will be used to perform detection and classification	<p>Examined use of convolutional neural networks to extract features from data</p> <ul style="list-style-type: none"> - Evaluated using available downward-looking GPR measurements of wide range of targets - Compared neural net detection using "machine-extracted" features to boosting using "human-derived" features (see March Quarterly) - Observed improved performance when neural net classification outputs were used as features combined with "human-derived" features

Research Areas	Research Topic	Activity During Period
4	Develop algorithms that exploit feature sets to yield optimal detection and classification.	No activity in period
5	Develop quantitative bounds on ATR performance in terms of target vs. clutter separability bounds and target class separability bounds.	No activity in period
6	Develop quantitative understanding of extensibility of trained classifier to new environments and new targets.	Examined the use of bound metrics from Transfer Learning to estimate the bounds on detection error when a classifier trained using one data set is applied to a new data set <ul style="list-style-type: none"> - Provides a measure of extensibility - Simulation suggested initial metrics considered provided a "loose" bound

As summarized in Section 2, alternate forward operators were formulated, implemented in software, and predictions compared to first principle numerical simulations. The general observation was that the spectral formulations provided accurate predictions of the field propagation, but at the expense of significantly increased computational complexity. The net result is that an inversion algorithm incorporating these models would have difficulty processing large GPR data sets imaging relatively subsurface regions large enough to be of practical use. Consequently, formulations using asymptotic formulations of the spectral operators were pursued. A baseline 3D imaging algorithm, incorporating a ray-based propagation model was exercised on a set of controlled laboratory measurements to assess how well the location and shape of a representative buried targets can be determined. This work provides a baseline to determine what extensions to the baseline algorithm are required to estimate material properties and to provide useful features for automatic target detection and classification.

Once images are formed, features are extracted from the images and used to perform the detection and classification. As discussed in Section 3, we began to evaluate novel methods, based on convolution neural networks (CNN), to extract features from radar images and use them in detection/classification algorithms. A CNN architecture was used to extract features from an available downward-looking GPR B-scans. We then compared buried object detection our CNN architecture using "machine-extracted" features to boosting-based

detector architecture using “human-derived” features developed under a previous program [Masarik, et al. 2016]. The general observation was that the detection performance using “machine-extracted” features was comparable to that obtained using a highly-tuned set of “human-derived” features. Interestingly, improve detection performance was observed when neural net classification outputs were used as features combined with “human-derived” features suggesting that the “machine-derived” features captured some additional useful information over that contained in the “human-derived” features.

Machine learning-based detection and classification algorithms typically train an algorithm on a set of data which is considered to be a representative subset of the potential target and clutter realizations the sensor is intended to detect. To understand the robustness of the algorithm, it would be useful to have a bound on algorithm performance when the algorithm is applied to data outside the original training data set. As described in Section 4, a bound from translation learning was investigated as a possible means of estimating detection performance when a machine learning classified trained on one data set is applied to a new data. As an example, this type of bound would provide guidance on how the detection performance of a machine learning algorithm trained on measured data from attest site would change when the algorithm is applied to data collected at a new location. A simulation was generated to examine the potential performance of the bound before applying it to measured data. The general observation was that the bound was correct but it was not close enough to the true performance to be particularly useful in applications. Thus, methods of improving the tightness of the bound were considered.

1.3 Summary of Year Two and Three Results

Since Michigan Technological University did not recognize that an additional increment of funds had been authorized for Year Two activities until the Spring of 2018, Year Two and Three technical activity occurred primarily in the last quarter of Year Two and first quarter of Year Three.

On the second increment of funding, we addressed two basic challenges to achieving reliable automatic target recognition (ATR) of buried explosive hazards using standoff measurement geometries:

1. How do we improve buried target radar imagery perturbed by propagation through the ground sufficiently to extract useful features for ATR?
2. How do we achieve standoff radar interrogation of the ground at measurement geometries favorable for penetration into the ground to increase detection performance?

As discussed in Section 5, to address the first issue (item 1), we investigated the use of the lifting approach to inverse problems [Demanet and Jugnon, 2017] to generate radar imagery of buried objects that are robust to mismatch or errors in the forward model. Since buried object imaging involves propagation through media that is, at best, partially known, there is mismatch between the forward model used in the inversion and the actual physical propagation mechanisms. The mismatch can cause degradation and/or reduce resolution in the imagery, which limits the features that can be extracted from the imagery. The image degradation limits ATR performance. The lifting approach to inversion can potentially improve imagery, and the resulting image features, which can lead to improve ATR performance. This activity supports Research Area 2 and 3 in Table 1.

During the subject reporting period, we developed software to apply the lifting approach to inverse problems to available ground penetrating radar measurements of buried targets. As discussed in section 5, images of buried targets were generated using the lifting techniques and compared to standard backprojection images. The lifting images had improved localization of target scattering and suppression of surrounding clutter. The variation of image characteristics with algorithm parameters was also examined.

To investigate the second issue (item 2), we investigated the use of a ground penetrating radar on a small unmanned autonomous systems (sUAS) to detect buried targets. A radar mounted on a sUAS can interrogate the ground at near normal angles of incidence, thus providing better propagation of energy into the ground, yet still be at a standoff distance

from friendly dismount operators. A wide bandwidth, low frequency, multi-channel ground penetrating radar with a four element antenna array was flown on the sUAS to provide additional spatial diversity for better spatial resolution. We collected measurements of buried targets along a Lake Michigan beach using a low SWAP radar on one of MTRI's sUAS platforms to demonstrate buried object detection and provide data for algorithm development. The results presented in Section 6, show that useable ground penetrating radar data can be collected from a sUAS platform and that the resulting data can be used to detect and image buried targets.

1.4 Recommendations for Future Work

To further advance the research performed during the subject reporting period, the following activities are recommended.

In Research Area 1, where we are developing image enhancement and initial clutter rejection algorithms through physics-based and statistics-based filtering schemes, we recommend investigating whether robust principal components analysis (RPCA) can reduce clutter to increase target-to-clutter ratio, which improves target detection performance. In previous work [Masarik et al., 2015], clutter reduction using RPCA has been demonstrated for downward-looking GPR. For standoff forward-looking GPR measurement geometries, target signatures are obscured by distributed surface scattering. To advance clutter reduction in forward-looking GPR, we recommend evaluating whether RPCA can reduce distributed clutter using available forward-looking GPR measurements of buried targets, and revise the algorithm as required. We also recommend evaluating the applicability of hyperbolic-Radon transform filtering for the forward-looking geometry [Masarik et al. 2016], which was shown to be effective for downward-looking GPR. The hyperbolic-Radon transform filtering looked for the characteristic hyperbolic signature of a buried object in a B-scan to separate target returns from clutter. The work would determine if a similar geometric signature can be exploited in forward looking geometries.

The third approach to clutter reduction that should be considered is statistical combining of sequences of images. Sequences of images can be generated from the data collected by an advancing standoff GPR sensor. Under the assumption that image pixels

containing distributed clutter will randomly change from image to image, while a localized buried target signature will be statistically more stable, the sequence of images can be combined to enhance the target response relative to the distributed clutter. The study should define the optimal statistical method of combining the sequence of images. The analysis would use available forward looking GPR measurements of buried targets. The development of signal processing algorithms to improve target-to-clutter ratio advances standoff ATR by improving target detection performance.

In Research Area 2, to continue to develop high-fidelity 3D subsurface image reconstruction algorithms for physical object property estimation, we suggest developing and validating a stationary phase approximation of the spectral Green's function based forward model to improve computational efficiency with minimal accuracy loss. This incorporates the radiometric quantities necessary to quantitatively estimate material properties. The ability of the algorithm to estimate material properties using numerical simulations and available measured data could then be evaluated. The study should also consider the impact of limited measurement diversity, which may be associated with certain standoff measurement geometries, on image reconstruction.

All subsurface imaging algorithms must reconstruct buried objects using GPR measurements involving propagation through a medium with properties that are assumed or have limited characterization. These unknown medium properties produce phase and amplitude errors in the GPR measurements, which perturb the reconstructed buried object image. These errors due to unknown information are similar to the uncompensated motion errors that impact conventional synthetic aperture radar imaging algorithms.

In the current effort, we quadratic lifting inversion (QLI), proposed by Demanet and Jugnon, [2017], as a robust approach to inversion in the presence of mismatch or errors in the forward model, which can produce reconstructions with fidelity comparable to direct inversion. On the basis of theory and simulations, Demanet and Jugnon claim that QLI has to potential to be robust to model mismatch, provide comparable resolution with classical techniques (backprojection, least squares auto-focus, etc.), and statistically spatially stable recovery (with respect to randomly heterogeneous media). We applied this technique to radar

imaging of buried targets to determine if it can produce enhanced imagery in the presence of limited knowledge of the surrounding ground geometry and/or material properties. The approach can potentially improve imagery, and the resulting image features, which would lead to improve ATR performance. To further investigate the application and utility of the lifting approach to inverse problems to generation of enhanced radar imagery of buried objects and its impact on ATR feature extraction, we suggest exercising the approach on a wider array of buried objects using both simulated and measured data, and determining whether the enhanced imagery impacts the extraction of ATR features. This approach has the potential to improve the fidelity of subsurface imagery, and thus improve the shape information that is available to an ATR.

Finally, we suggest evaluating the utility of features extracted from 3D images for standoff ATR. 3D imaging requires additional data collection and data processing. Specifically, the future work should investigate whether the additional feature information obtained from 3D imaging increases the detection and classification performance of an ATR system. This analysis is relevant to standoff ATR since it would quantitatively determine if the potential gain in performance is worth the additional data collection and processing.

In Research Area 3, we suggest continuing to investigate methods of identifying and extracting features from standoff measurements that will be used to perform detection and classification. Most machine learning-based approach to ATR operate on features extracted from the raw sensor measurements to detect and classify buried objects. Traditionally, these features have been derived by either analysis of characteristics of the data by researchers or the application of features sets from related technical areas, such as computer vision. Recently, feature extraction using deep neural network architectures has been advanced in the machine learning literature, although the applications have generally been web-based applications with very large training sets. Machine extraction of features has advantages in terms of efficiency, but has limitations in terms of interpretability and extensibility. We specifically suggest evaluating the applicability of residual networks, which are a novel neural network architecture that propagates residuals. The technique can be applied to simulated data and available forward-looking GPR measurements to observe what features

are extracted and to compare detection algorithm performance with machine-extracted features to that obtained with human-engineered features.

We suggest applying the Gaussian copula framework [Thelen, et al., 2014] to give a quantitative assessment of the most useful features for separating targets from clutter and for discriminating between targets and confusers. This approach has been successfully used to determine the most important features to separate targets in ATR systems for HRR and SAR data, and can potential provide similar insight for buried object detection and discrimination. Developing efficient methods of extracting informative features and assessing their impact on ATR performance provides a useful tool for the development of improved standoff ATR systems.

In Research Area 4, which is investigating methods of exploiting feature sets to obtain optimal detection and classification, we suggest investigating solutions to the algorithm training issues that arise from imbalanced training sets. A common characteristic of measured data sets used to train ATRs for buried object detection is that there are usually many more samples of clutter than samples of targets. This imbalance can lead to biased training results, which reduce performance. Somewhat ad hoc methods of compensating for this imbalance, such as subsampling or augmenting with simulated samples, have been tried, but we plan to consider a more structured approach. Specifically, we suggest evaluating the Naïve Bayes SVM (NBSVM) algorithm to overcome fundamental issues of imbalanced training datasets and unequal misclassification costs. We also suggest extending this framework to work with a boosting-based classifier, as well as the SVM. This study will improve standoff ATR performance by overcoming machine learning training issues due to imbalanced training datasets.

In Research Area 5, which aims to develop quantitative bounds on ATR performance in terms of target vs. clutter separability bounds and target class separability bounds, we suggest quantifying the fundamental separability of object classes using the Gaussian copula framework [Thelen, et al. 2016] to inform sensor/scenario design and yield fundamental performance bounds for algorithm performance. This research work will advance standoff

ATR by providing bounds on potential performance as a function of target, clutter and sensor parameters, which can guide the development of sensors and data collection methodologies.

Finally, in Research Area 6, which aims to develop quantitative understanding of extensibility of trained classifier to new environments and new targets, we recommend investigating extensions to current transfer learning bounds to provide an approach for estimating the extensibility of a trained algorithm to new data. This research work will advance standoff ATR by providing a method of estimating the detection performance of a trained ATR algorithm when applied to new data and/or measurement scenarios.

Most of the limitations of current standoff IED detection systems arise from the fact that they view the ground at oblique angles of incidence, which reduces the propagation of interrogating radar energy into the ground and produces much more distributed clutter. Collecting downward-looking GPR data from a small unmanned autonomous system provides a method of interrogating the ground at near normal incidence while keeping the operator at a safe standoff distance. In the current program, the detection and imaging of buried targets using a wideband, low frequency, ground penetrating radar on an sUAS system was demonstrated. To further investigate the use of a ground penetrating radar on a small unmanned autonomous systems (sUAS) to detect buried targets, we suggest additional collects of data from sUAS systems to understand the measurement and data processing issues that affect buried object detection and imaging performance. Many of these issues relate to sUAS control and motion compensation. Once this behavior is understood and compensated, detection and classification algorithms previously developed for downward-looking sensors can be applied. We feel that fielding these systems can significantly advance the standoff buried target detection and classification performance.

1.5 Report Overview

The results of the study are summarized in the following sections. Three dimensional subsurface imaging studies are discussed in Section 2. ATR feature extraction using neural networks is described in Section 3. Transfer learning to estimate algorithm extensibility is discussed in Section 4. The application of quadratic lifting inversion to buried object radar

imaging is described in Section 5, and buried object detection and imaging using a GPR on a small UAS is summarized in Section 6.

2.0 3D Subsurface Imaging Studies

Our research effort into 3D imaging is based on a statistical inverse problem approach to making inferences from measured data. The underlying principle of this methodology is that by understanding the physical relationships between the environment, the imaging system, and the collected data, one can more effectively use available measurements to make inferences about the scene being imaged. The estimation problem at hand is said to be part of an inverse-problem approach because information regarding the unknown parameters is only observed through its influence on the output of a physical process. A statistical inverse-problem approach to parameter estimation is one which properly accounts for all of the available sources of information in the generation of the estimates. This includes knowledge of the forward processes, the types and sources of statistical uncertainty, available prior information, and the data itself.

In order to accomplish this, we create a model of the physical process to be inverted. This model is known as the forward model, and in most modern inverse-problem approaches, is actually composed of two models. Using the terminology of [Rosenblueth and Wiener, 1945], the forward model always includes a “symbolic description in logical terms” which acts as an abstract description of the forward process. With regard to 3D imaging using ground penetrating radar, the physical process is governed by Maxwell’s equations, the scattering geometry, and other parameters related to the sensors and materials involved.

In general, the inverse problems framework would correspond to a forward model of the type

$$d = \Lambda(\theta, \eta) + n$$

where θ is the parameter of interest, η represents the “nuisance” parameters that are not of interest, but affect the data and so need to be estimated, Λ is a forward model operator, n is noise, and d is collected data. In the inverse problems framework, we begin with collected data d and try to derive an estimate of θ and η . To do this we apply regularizations on θ and η corresponding to a priori information, and we make use of a metric Γ measuring closeness between the modeled data and the collected measurements which takes into account a

statistical model of the noise. These in conjunction comprise the objective function. The final equation takes on the general form below.

$$(\hat{\theta}, \hat{\eta}) \approx \arg \min_{\theta, \eta} \{ \Gamma(\Lambda(\theta, \eta) - d) - R_1(\theta) - R_2(\theta) \}$$

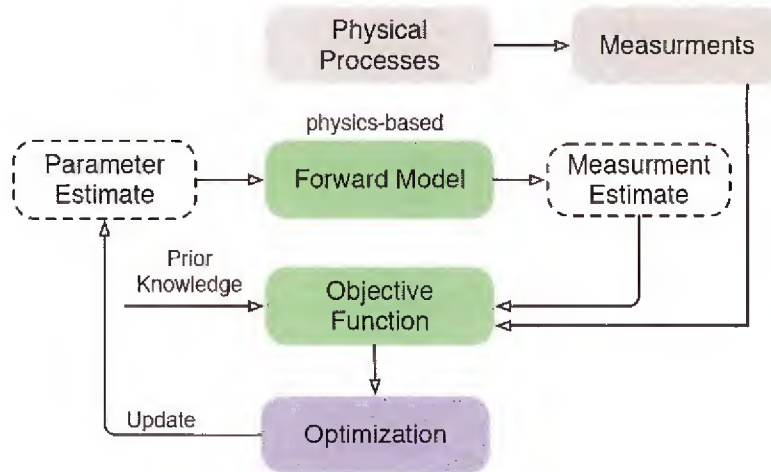


Figure 2: Statistical Inverse-Problem Framework Flowchart

Much of the technical activity during reporting period was focused on investigating penetrative three-dimensional (3D) radar imaging with the goal of producing high fidelity representations of buried targets. Specifically, alternate formulations of the forward operator, specifically spectral representations, which incorporate more physics, were investigated to provide a formulation that supports the estimation of buried object material properties. The objective was primarily to develop a practical inversion formulation that could estimate material properties of buried objects. Features derived from these images, namely location, shape and material properties, can then be exploited in algorithms to produce automated detection and classification.

The following subsections will provide a discussion of alternate forward operators will be discussed. A baseline 3D imaging algorithm was evaluated on a set of controlled laboratory measurements to assess how well we can locate and determine the shape of a representative buried targets. This work provides a baseline we can use to determine what

extensions to the baseline algorithm are required to provide useful features for automatic target detection and classification.

2.1 Problem Description

2.1.1 Half-Space Geometry

In order to model a ground penetrating radar system, we describe our system within the context of a half-space geometry. The space above the ground we refer to as the “collection domain,” or “ Σ ”, and the lower space we call the “investigation domain,” or “ D ”. The air-soil interface is defined to be the plane $z = 0$. In general, we have an antenna that generates an electromagnetic field in the collection domain, reacts to the targets in the investigation domain, and then is measured again in the collection domain. From these measurements, we wish to gain as accurate an understanding as possible of exactly what materials were buried underground.

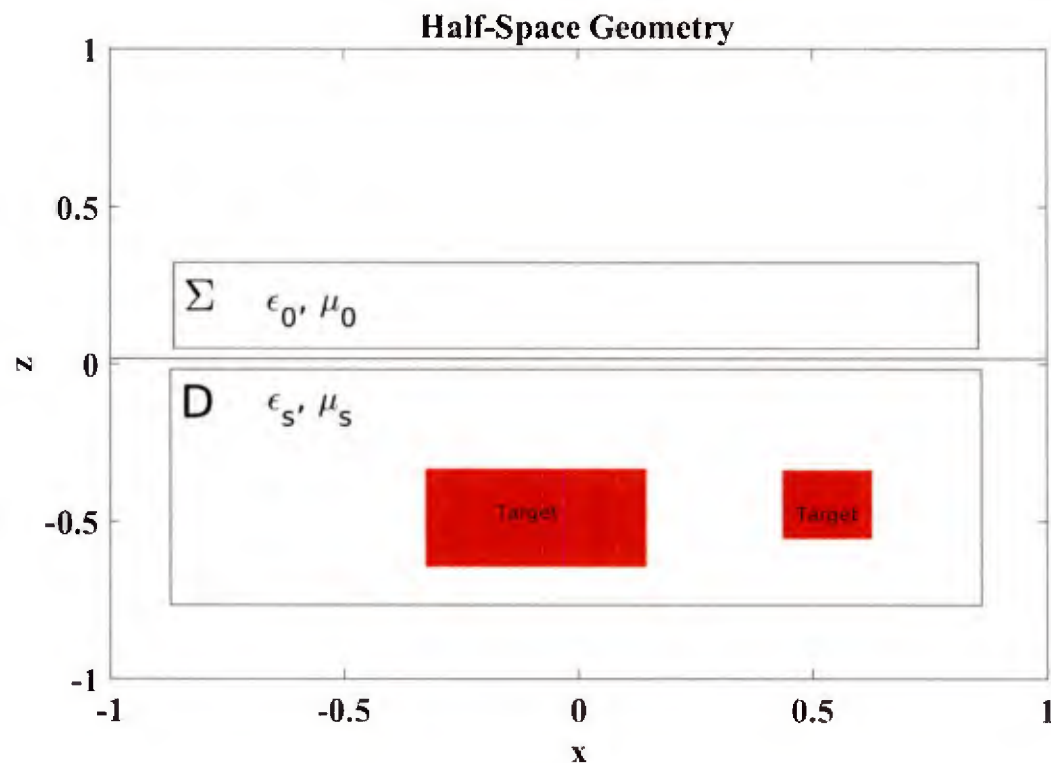


Figure 3 Half-Space Geometry

In this problem, a reconstruction of a dielectric contrast function is the goal. The dielectric contrast function represents the ratio of the dielectric constant of whatever happens to be in the investigation domain to the background dielectric permittivity of the soil, ϵ_s . The dielectric contrast function that we used in this work is [Persico 2005]:

$$\chi_e(r') = \frac{\epsilon(r') - \epsilon_s}{\epsilon_s}.$$

The dielectric contrast is equal to zero when the dielectric in the investigation domain matches the dielectric of the soil.

2.1.2 Scattering Equations

The way in which the electric field responds to the half-space geometry and the buried targets is described completely and exactly by the scattering equations given below. [Persico 2005] Dependence on frequency has been made explicit to emphasize the spectral nature of the approach we have chosen to use.

$$e_{inc}(r', \omega) = -i\omega\mu_0 \int_{\Sigma} G_{as}(r', r_s, \omega) J(r_s, \omega) d^3r_s \quad (r_s \in \Sigma, r' \in D)$$

$$e(r', \omega) = e_{inc}(r', \omega) + \omega^2\mu_0\epsilon_s \int_D \chi_e(r'') G_{ss}(r', r'', \omega) e(r'', \omega) d^3r'' \quad (r'' \in D)$$

$$e_{sc}(r, \omega) = \omega^2\mu_0\epsilon_s \int_D \chi_e(r') G_{sa}(r, r', \omega) e(r', \omega) d^3r' \quad (r \in \Sigma)$$

The first equation describes the “incident field” $e_{inc}(r', \omega)$ in the soil due to the current source $J(r_s, \omega)$ above the ground by invoking an “air-soil” green’s function, $G_{as}(r', r_s, \omega)$. The second equation, the “consistency equation,” describes how the dielectric contrast function $\chi_e(r'')$ and the incident electric field $e_{inc}(r', \omega)$ interact to produce the “total electric field” $e(r', \omega)$ in the soil by invoking a “soil-soil” green’s function, $G_{ss}(r', r'', \omega)$. The third equation, the “data equation,” describes how the total electric field $e(r', \omega)$ interacts with the dielectric contrast function $\chi_e(r')$ to produce the “scattered field,” $e_{sc}(r, \omega)$ which is what one would extract from their GPR measurements after removing the “self-response” and the “ground response.” In a way, these equations are a bit tautological.

The green's functions used above are *defined* such that these equations are true. They are the functions such that multiplying them by a source term and integrating them over the problem domain provides the solution to Maxwell's differential equations describing the half-space geometry.

It is worth noting that each of the Green's functions above is represented by a 3x3 matrix of components, and each electric field is represented by a 3x1 vector of components. Wherever relevant, standard matrix multiplication rules apply.

2.1.3 Spectral Green's Functions

In order to relate the quantities $e_{sc}(r, \omega)$ and $J(r_s, \omega)$ in the above scattering equations, we need to be able to compute the green's functions mentioned above. There are various ways of deriving and expressing these Green's functions. We have chosen the spectral form given below. [Lo Monte, et al., 2012]

$$\begin{aligned}
G_{as}(r, r', \omega) &= \frac{i}{4\pi^2 \omega^2 \mu_0 \epsilon_0} \iint_{-\infty}^{\infty} \frac{e^{iu(x-x')} e^{iv(y-y')} e^{iz' \sqrt{\omega^2 \mu_0 \epsilon_0 - u^2 - v^2}} e^{-iz \sqrt{\omega^2 \mu_0 \epsilon_s - u^2 - v^2}}}{\sqrt{\omega^2 \mu_0 \epsilon_s - u^2 - v^2} + \frac{\epsilon_s}{\epsilon_0} \sqrt{\omega^2 \mu_0 \epsilon_0 - u^2 - v^2}} M(u, v, \epsilon_s, \epsilon_0) du dv \\
G_{sa}(r, r', \omega) &= \frac{-i}{4\pi^2 \omega^2 \mu_0 \epsilon_0} \iint_{-\infty}^{\infty} \frac{e^{iu(x-x')} e^{iv(y-y')} e^{iz \sqrt{\omega^2 \mu_0 \epsilon_0 - u^2 - v^2}} e^{-iz' \sqrt{\omega^2 \mu_0 \epsilon_s - u^2 - v^2}}}{\sqrt{\omega^2 \mu_0 \epsilon_s - u^2 - v^2} + \frac{\epsilon_s}{\epsilon_0} \sqrt{\omega^2 \mu_0 \epsilon_0 - u^2 - v^2}} M^T(u, v, \epsilon_s, \epsilon_0) du dv \\
G_{ss}(r, r', \omega) &= \left(\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} + \frac{1}{\omega^2 \mu_0 \epsilon_s} \begin{pmatrix} \frac{\partial^2}{\partial x^2} & \frac{\partial}{\partial x} \frac{\partial}{\partial y} & \frac{\partial}{\partial x} \frac{\partial}{\partial z} \\ \frac{\partial}{\partial x} \frac{\partial}{\partial y} & \frac{\partial^2}{\partial y^2} & \frac{\partial}{\partial y} \frac{\partial}{\partial z} \\ \frac{\partial}{\partial x} \frac{\partial}{\partial z} & \frac{\partial}{\partial y} \frac{\partial}{\partial z} & \frac{\partial^2}{\partial z^2} \end{pmatrix} \right) \frac{e^{i\omega \sqrt{\mu_0 \epsilon_s} |r-r'|}}{4\pi |r-r'|} \\
&+ \frac{i}{4\pi^2 \omega^2 \mu_0 \epsilon_0} \iint_{-\infty}^{\infty} \frac{e^{iu(x-x')} e^{iv(y-y')} e^{-i(z+z') \sqrt{\omega^2 \mu_0 \epsilon_s - u^2 - v^2}}}{\sqrt{\omega^2 \mu_0 \epsilon_s - u^2 - v^2} + \frac{\epsilon_s}{\epsilon_0} \sqrt{\omega^2 \mu_0 \epsilon_0 - u^2 - v^2}} M(u, v, \epsilon_s, \epsilon_0) \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & \frac{\epsilon_s}{\epsilon_0} \end{pmatrix} du dv
\end{aligned}$$

$$M(u, v, \epsilon_s, \epsilon_0) = \begin{pmatrix} \sqrt{\omega^2 \mu_0 \epsilon_0 - u^2 - v^2} \sqrt{\omega^2 \mu_0 \epsilon_s - u^2 - v^2} + v^2 & -uv & u \sqrt{\omega^2 \mu_0 \epsilon_s - u^2 - v^2} \\ -uv & \sqrt{\omega^2 \mu_0 \epsilon_0 - u^2 - v^2} \sqrt{\omega^2 \mu_0 \epsilon_s - u^2 - v^2} & v \sqrt{\omega^2 \mu_0 \epsilon_s - u^2 - v^2} \\ u \sqrt{\omega^2 \mu_0 \epsilon_0 - u^2 - v^2} & v \sqrt{\omega^2 \mu_0 \epsilon_0 - u^2 - v^2} & u^2 + v^2 \end{pmatrix}$$

There are three functions listed above: “air-soil,” “soil-air,” and “soil-soil.” They have been named intuitively, such that the “air-soil” green’s function describes the response in the soil due to a source in the air, etc. They are represented using 3x3 matrices because there are three possible transmission polarizations and three possible measurement polarizations, one for each principal axis in 3 dimensions. Each of the nine combinations of transmitting and receiving polarizations of our implementations have been tested and verified against the software package gprMax. Figures 4, 6, and 8 show a direct comparison of the output of our green’s function implementations to the output of a gprMax simulation of the same geometry. gprMax is an open-source electromagnetic simulation code designed to simulate ground-penetrating radar. gprMax employs Yee’s algorithm to solve Maxwell’s equations in using the finite difference time-domain method. [Giannopoulos 2005]. Figures 5, 7 and 9 show the fraction of the total error (F_i), defined below, between x_i , our implementation, and y_i , the gprMax output, at each time step.

$$F_i = \frac{|x_i - y_i|}{\sum_{j=1}^N |x_j - y_j|^2}$$

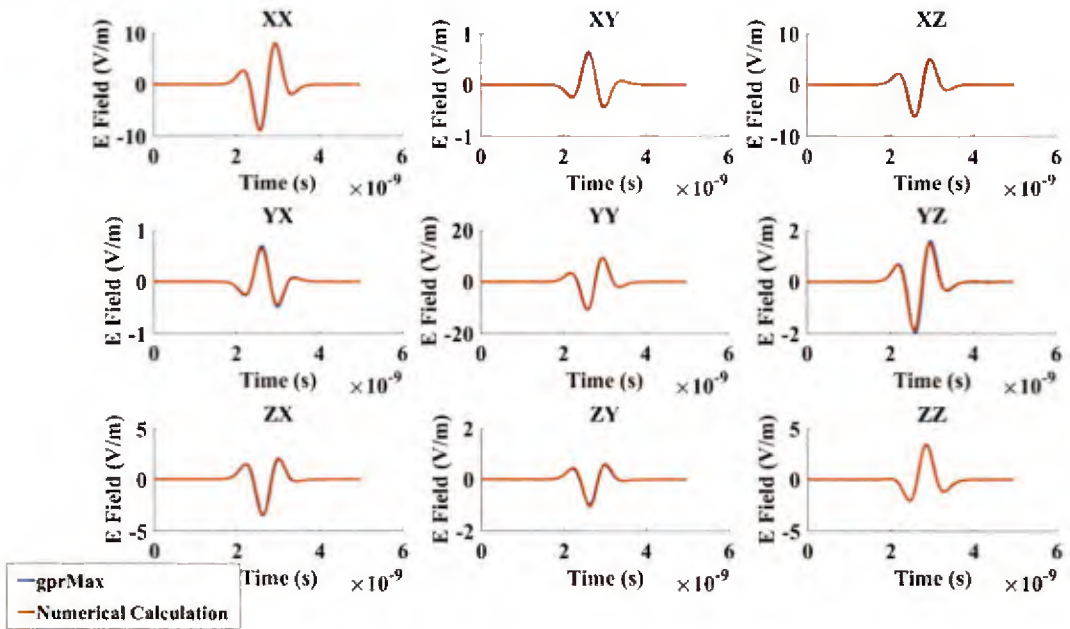


Figure 4 Air-Soil Green's Function Comparison




Figure 5 Air-Soil Green's Function Error

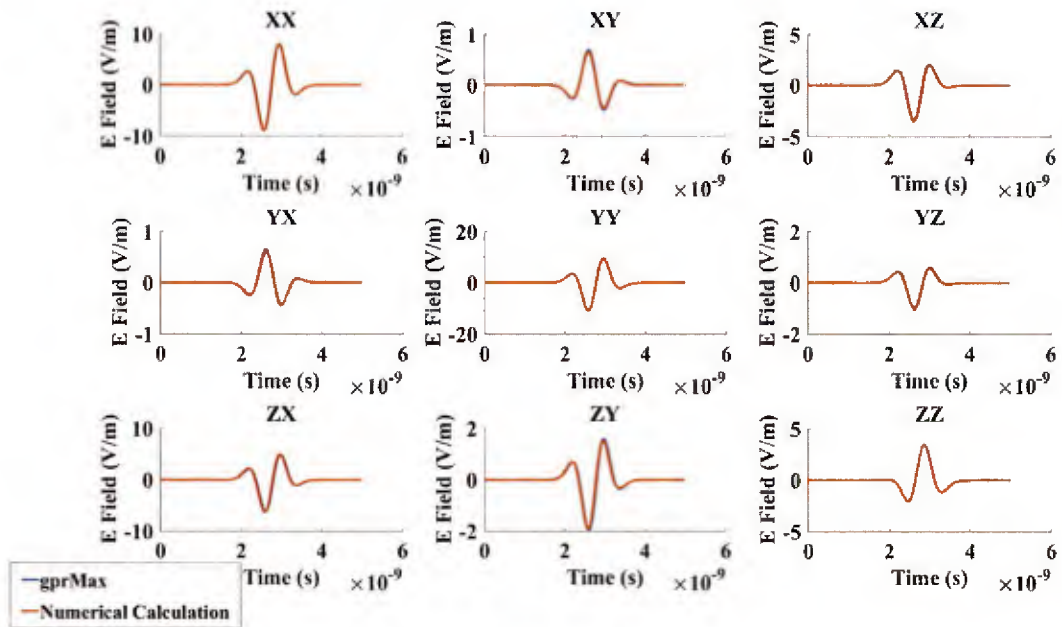


Figure 6 Soil-Air Green's Function Comparison

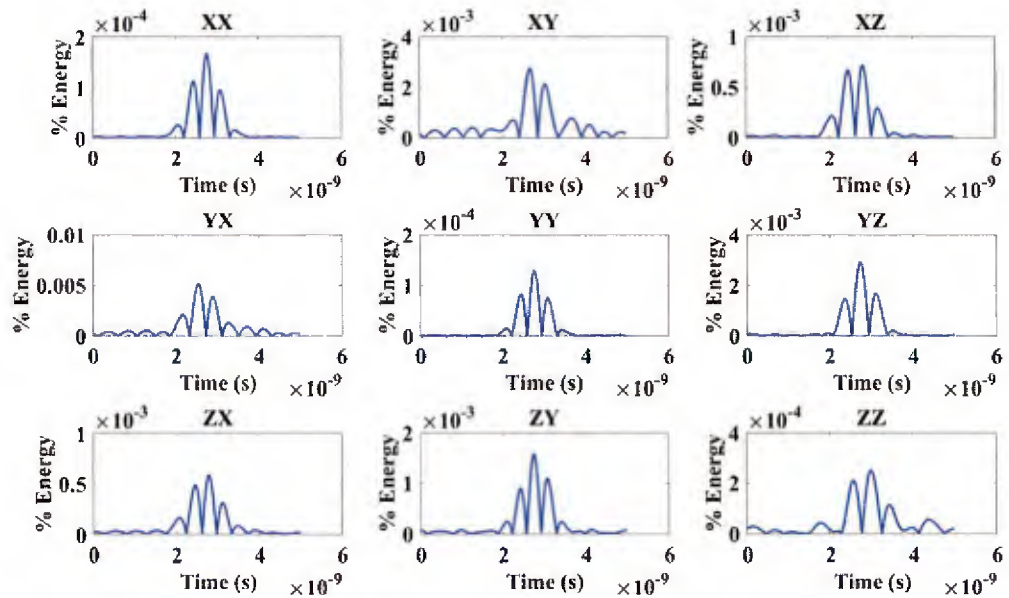


Figure 7 Soil-Air Green's Function Error

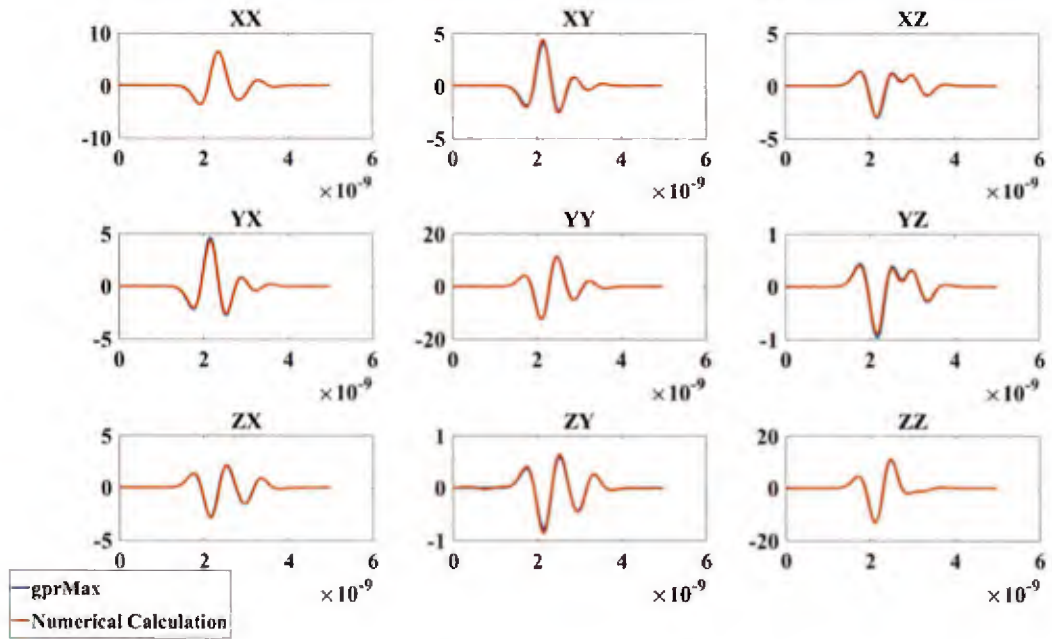


Figure 8 Soil-Soil Green's Function Comparison

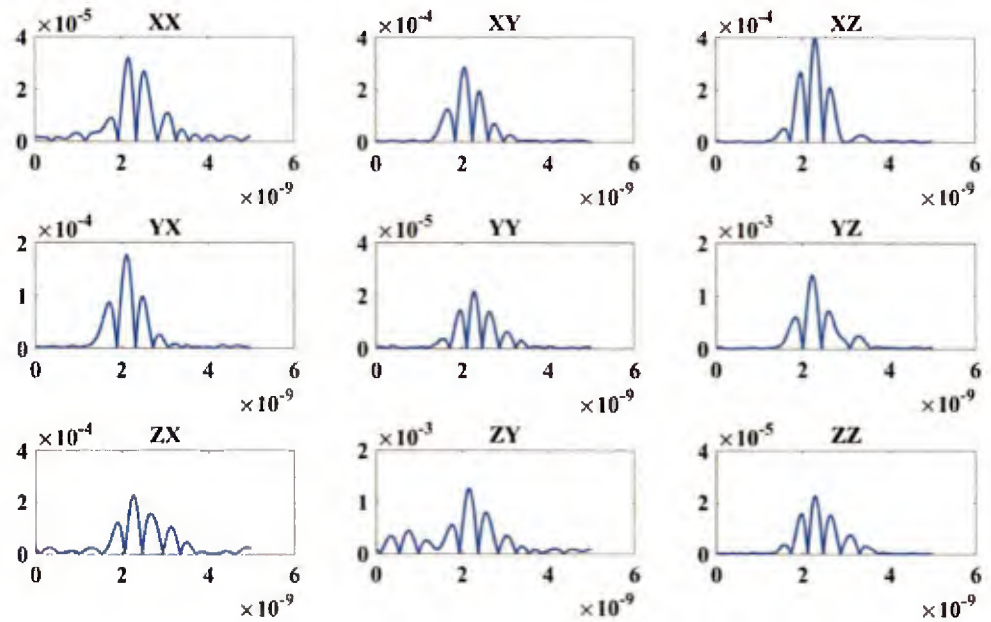


Figure 9 Soil-Soil Green's Function Error

We can conclude that our implementations of these green's functions are verified against gprMax to be an accurate numerical approximation to the spectral green's functions defined above.

2.1.4 Computational Considerations

The green's functions above are implemented in MATLAB by numerically computing the integrand over a range of spatial frequencies and then taking the 2D fast Fourier transform to estimate the integrals over a large range of x and y values. In order to make the green's function as accurate as in Figures 4 through 9, the necessary rate at which the spatial frequencies must be sampled exceeds the number of points in the X and Y directions by a factor of 6. That is, if we wish to compute the electric field on a 50 x 50 grid of x and y locations, we numerically computed the integrand over a 300 x 300 grid of spatial frequencies. Then we compute the fast Fourier transform and discard the extra 87,500 points. This process is repeated for each x-y planar "slice" in our investigation domain because performing the operation all at once would require too much memory.

The total electric field in the "consistency equation" in Section 2.1.2 depends on the dielectric contrast in a nonlinear fashion. It is possible to express a discrete version of this equation as

$$e(\omega) = A(\chi_e, \omega)e_{inc}(\omega)$$

Note that the dependence on position has been dropped because the dimensions of the matrices represented by A , e , e_{inc} account for the functional dependence on r, r' . In the above, e is a $3N_D \times N_C$ matrix, A is a $3N_D \times 3N_D$ matrix, and e_{inc} is a $3N_D \times N_C$ matrix, where N_D and N_C are the number of discrete sample points in the investigation and collection domains, respectively. The incident field and the total field are each comprised of 3 components, and are each sampled at N_D points in the investigation domain, and dependent on N_C different collection geometries. If we were to consider a polarimetric model, N_C would become $3N_C$.

Generally, the forward operator will be too large to store in memory. This will be true regardless of whether or not a linear approximation is made. For example, to store the

forward operator which maps a 25×25 grid of collection points, each collected over 100 frequencies, onto a 3D space of $50 \text{ cm} \times 50 \text{ cm} \times 50 \text{ cm}$ at 1 cm grid spacing would require $25 \times 25 \times 100 \times 51 \times 51 \times 51 = 8.3 \text{ Billion}$ complex, double-precision floating point values, or approximately 132.8 Gigabytes of memory. Additionally, storing the matrix representing the full nonlinear relationship between the incident field and the total field in the same scenario would require $(3 \times 51^3)^2 = 158 \text{ Billion}$ complex, double-precision floating point values, or approximately 2.5 Terabytes of memory. Therefore, it becomes necessary to implement a highly optimized subroutine which emulates the *action* of the forward operator. From that point, we can use this operator in conjunction with a gradient-descent based algorithm to find a sufficient minimum for some objective function.

2.2 Approximations

Given the computational burdens outlined above, it would be difficult to directly solve the full nonlinear spectral inverse scattering problem given reasonable computational resources for a typical standoff GPR data collection. Therefore, approximations need to be made and tradeoffs need to be carefully considered.

2.2.1 Born Approximation

In terms of the scattering equations given in Section 2.1.2, the “Born approximation” is equivalent to approximating the total electric field with the incident electric field. This approximation makes the scattered field linearly dependent on the dielectric contrast function and easy to compute; however, it is well-known that the Born approximation is not sufficient for quantitative reconstructions of strongly scattering objects. [Lo Monte, et al. 2012] The Born approximation is generally only valid for weak scatterers, since it does not take into account interactions between multiple scatterers, which can also lead to artifacts and “ghost” targets in reconstructions [Picco, 2015].

$$e(r', \omega) \approx e_{inc}(r', \omega)$$

$$e_{sc}(r, \omega) \approx \omega^2 \mu_0 \epsilon_s \int_D \chi_e(r') G_{sa}(r, r', \omega) e_{inc}(r', \omega) d^3 r'$$

The Born approximation is “most valid,” when scatterers are small and weak. [Li 2010] In the following example, we compare a gprMax simulation of a half-space with a buried 3x3x3 cm dielectric cube, and compare the output to what we would predict using the Born approximation and our numerical implementations of the Green’s function operators. The relative permittivity of the buried cube in the example below is 16, and the relative permittivity of the soil is 4.




Figure 10 Born Approximation gprMax Comparison for 3cm Cube

Even for a fairly small target, the Born approximation prediction departs from the output of the gprMax simulation. We have developed a slight improvement over the Born approximation which more accurately predicts the scaling of the electric field. The yellow line in Figure 10 is our scaled version of the Born approximation, described in detail in Section 2. 2.2.

2.2.2 Scaled Born Approximation

Assuming a single point-scatterer in the investigation domain leads to the following expression for the dielectric contrast.

$$\chi_e(r) \approx \begin{cases} \chi_e & r = r_t \\ 0 & r \neq r_t \end{cases}$$

We can then take this into the consistency equation given in Section 2.1.2 to retrieve the following expression for the total electric field.

$$e(r', \omega) \approx \begin{cases} e_{inc}(r_t, \omega) + \omega^2 \mu_0 \epsilon_s \chi_e G_{ss}(r_t, r_t, \omega) e(r_t, \omega) & r = r_t \\ e_{inc}(r', \omega) & r \neq r_t \end{cases}$$

The soil-soil green's function consists of two terms. The first term represents a direct path through the soil from point r to point r' . The second term represents a “bounce” off of the air-soil interface, beginning at point r and terminating at point r' . In this approximation we assume that the first term dominates the second term in the soil-soil green's function and we neglect the second term. According to [Chew 1990] the expression for the first term given in Section 2.1.2 is not valid when $r = r'$. Instead, there should be an additional Dirac delta function in order to uniquely and correctly define the green's function.

$$G_{ss}(r, r', \omega) \approx \left(I + \frac{1}{\omega^2 \mu_0 \epsilon_s} \nabla \nabla \right) \frac{e^{i\omega^2 \mu_0 \epsilon_s |r-r'|}}{4\pi |r-r'|} - \frac{1}{3\omega^2 \mu_0 \epsilon_s} \delta(r-r')$$

Integrating over this approximate green's function while also assuming a single point-scatterer leads to what we call the “scaled Born approximation.”

$$e(r_t, \omega) \approx e_{inc}(r_t, \omega) - \frac{1}{3} \chi_e e(r_t, \omega)$$

$$e(r_t, \omega) \approx \frac{3 e_{inc}(r_t, \omega)}{1 + 3 \chi_e}$$

$$e_{sc}(r, \omega) \approx \omega^2 \mu_0 \epsilon_s G_{sa}(r, r_t, \omega) \frac{3 \chi_e}{3 - \chi_e} e_{inc}(r_t, \omega)$$

This approximation adds negligible overhead compared to the Born approximation, but its benefits are seen in Figure 8. For small enough scatterers such that the Born approximation is relatively “valid,” this scaling produces results that allow one to more

accurately estimate the scattered electric field, which then leads to more accurate reconstructions of dielectric contrast.

In Figures 11 through 13 below is a reconstruction of a buried 2x2x2 cm cube using both the Born approximation and the scaled Born approximation. The Born approximation reconstruction estimates a contrast of approximately 0.45 and the scaled Born reconstruction estimates a contrast of approximately 0.55. The true dielectric contrast was 3. This is a marginal improvement of 3% of the true contrast.




Figure 11 2cm Cube Truth




Figure 12 2cm Cube Reconstruction (Born Approximation)




Figure 13 2cm Cube Reconstruction (Scaled Born Approximation)

As shown above, the Born approximation does provide somewhat accurate reconstructions of very small targets, up to a scaling factor. However, as shown in Figure 14, when we increase the size of the target from a 2 cm cube to a 5 cm cube, the Born approximation can no longer recover the correct shape of the pulse.

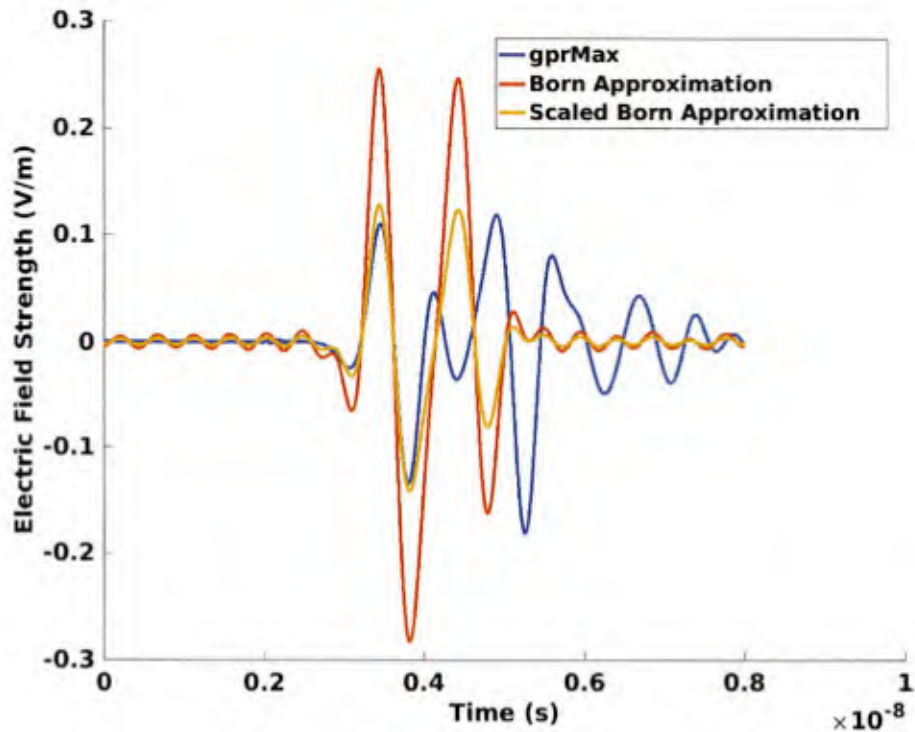


Figure 14 Born Approximation Comparison on 5cm Cube

The forward model no longer matches the gprMax simulation very well when we expand the target to a 5 cm cube. Attempting to image this model produces “ghost” targets above and below the true target, as shown in Figures 15-17.

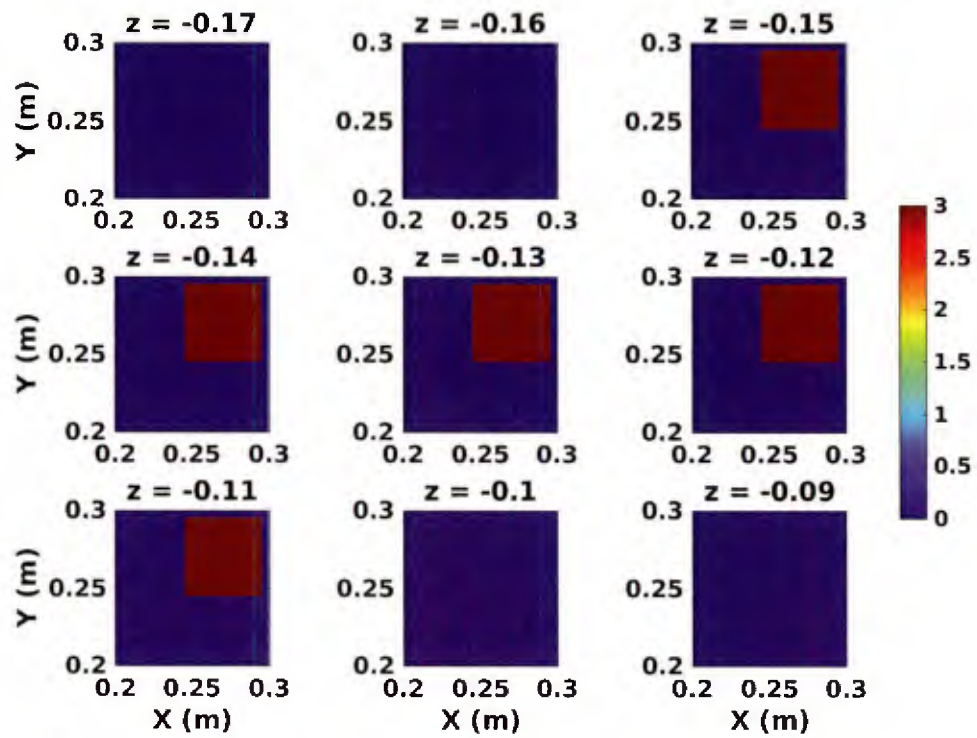


Figure 15 5cm Cube Truth Data




Figure 16 5 cm Cube Reconstruction (Born Approximation)




Figure 17 5cm Cube Reconstruction (Scaled Born Approximation)

2.2.3 Quadratic Approximation

The consistency equation in Section 2.1.2 can be written out as a Neumann series (commonly called the “Born series”) which will converge if the operator which represents the integral over the “soil-soil” green’s function and the total electric field has a small enough norm. Taking into account only the first term in this series leads to the Born approximation. Taking into account the first two terms leads to a “second order” Born approximation, or a Quadratic approximation. This has recently been attempted by [Vittorio, 2015] and shown to be an improvement over the Born model in terms of shape, but not in terms of scale of dielectric contrast. The same paper claims inversion attempts using the Quadratic model outside the range of validity of the Born model still underestimate the true dielectric contrast function, and that the largest source of error is modeling error. We have chosen not to pursue this option as an advanced imaging technique for these reasons.

2.2.4 “Ray” Approximation

Intuitively, this model simply computes the time (τ) it would take for a “ray” of light to travel from a transmit point to a point r' through the air-soil interface, then from that point back to a collection point, once again through the air-soil interface.

The i th data sweep contains complex radar data that is received at location x_i^R , was transmitted at location x_i^T , and is collected over the set of discrete frequencies f_k , $k = 1, \dots, N_f$. The imaging region (Figure 1) is sampled at the grid points x_j^g and these are indexed linearly over $j = 1, \dots, N_{grid}$. If we denote the two-way lag time from x_i^T to x_j^g and from x_j^g to x_i^R as τ_{ij} , and the reflectivity associated with point x_j^g as σ_j , then the data model can be written as

$$\tilde{r}_i(f) = \sum_{j=1}^{N_{grid}} \sigma_j e^{-2\pi i f \tau_{ij}} + \tilde{n}_i(f).$$

Here we have denoted the i th data sweep as $\tilde{r}_i(f)$ and the system noise as $\tilde{n}_i(f)$, which is modeled as circularly complex additive white Gaussian noise (AWGN). The two-way lag depends on the relative location of the radar transmitter/receiver and scattering location relative to the ground-air interface and the material properties of the ground, due to the refraction at the interface. See Figure 1 for an illustration of the problem geometry. We have implemented a fast algorithm for computing the lags perturbed by the air-ground interface.

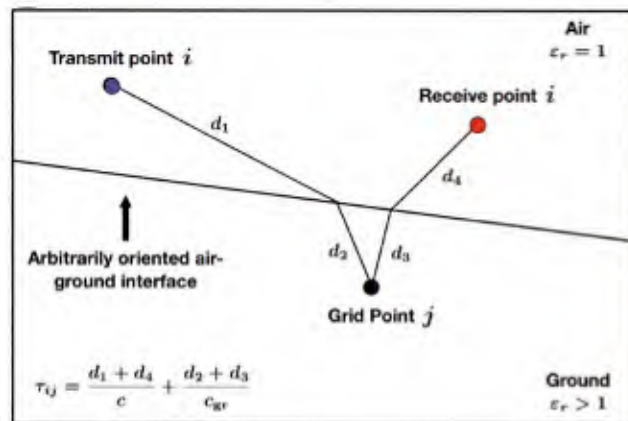


Figure 18 GPR Measurement Geometry

The data equation can then be written as a matrix-vector equation of the form:

$$d_i = S_i \sigma + n_i.$$

Then, upon stacking the d_i 's, the S_i 's, and the n_i 's vertically for $i = 1, \dots, N_{coll}$, we obtain a model for the entire phase history:

$$d = S\sigma + n.$$

The backprojection solution to which we will refer is defined as

$$\sigma_{BP} = S^H d.$$

Note also that we have only used the matrix S to facilitate writing down our solutions. Forming the full matrix S is extremely inefficient and can generally be avoided by computing the action of S and S^H by exploiting their resemblance to a Fourier operator.

2.2.5 Stationary Phase Approximation

One more reasonable approach to compute the electric field with a feasible set of resources is to make an approximation of the green's functions which eliminates the need for Fourier transforms or integral estimation. Ideally, we would be able to do this without introducing significant error in the approximation. One common approach in the estimation of oscillatory integrals is an asymptotic stationary phase approximation. [Hansen, 2000; Persico, 2005] We currently have an ongoing effort to develop a stationary phase approximation of the air-soil and soil-air green's functions.

To help motivate the stationary phase approach, a model has been developed which uses the interpretation in the "ray approximation" to produce an approximation which requires much less computational resources than the computation of the spectral green's functions but is still able to predict the electric field with errors no greater than 4 parts in 1000.

This approximation requires computation of the classical “diffraction points” as mentioned in the ray model given above and the computation of the computationally inexpensive free-space green’s function given below.

$$G_{fs}(r, r', \omega) = \left(I + \frac{1}{\omega^2 \mu \epsilon} \nabla \nabla \right) \frac{e^{-i\omega^2 \mu \epsilon |r - r'|}}{4\pi |r - r'|}$$

Then, we use this green’s function to propagate from the source point r' to the diffraction point, apply the interface conditions for a dielectric interface, then propagate from the diffraction point to the measurement point in the ground. The goal is to derive a reasonable approximation of the computationally expensive “air-soil” green’s function given above. Then, we can use this approximation and spend more of our computational budget dealing with the nonlinearities of the scattering problem rather than computing large Fourier transforms to numerically estimate the green’s functions.

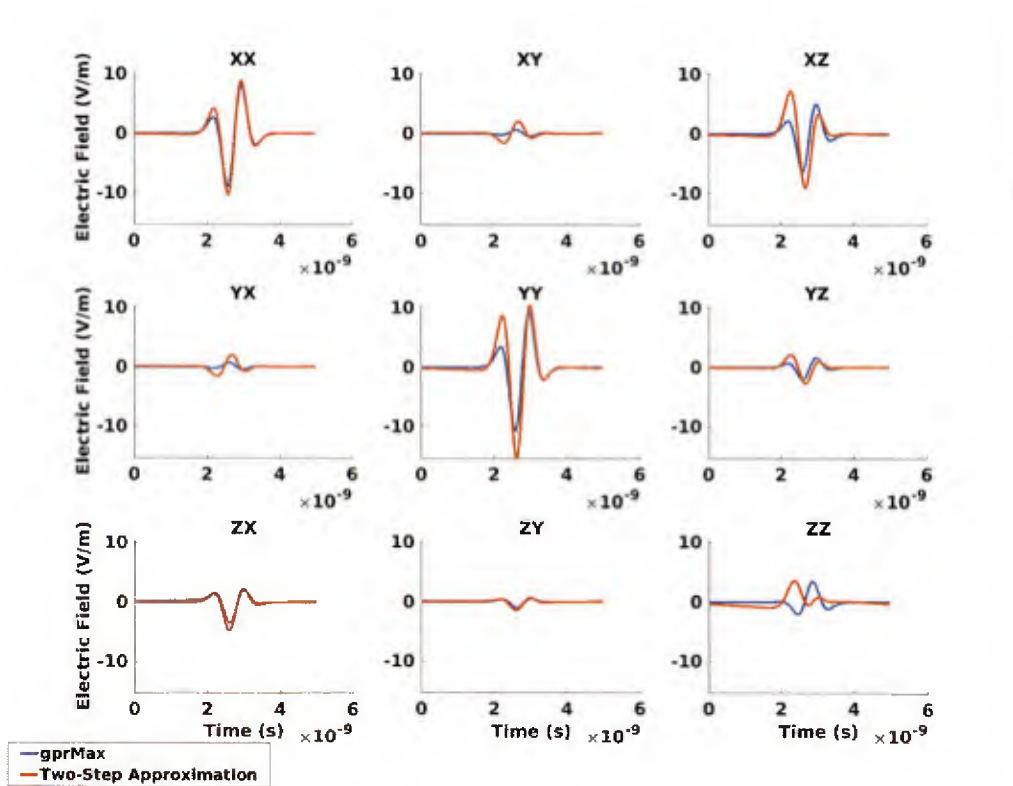


Figure 19 Extended Ray Approximation gprMax Comparison

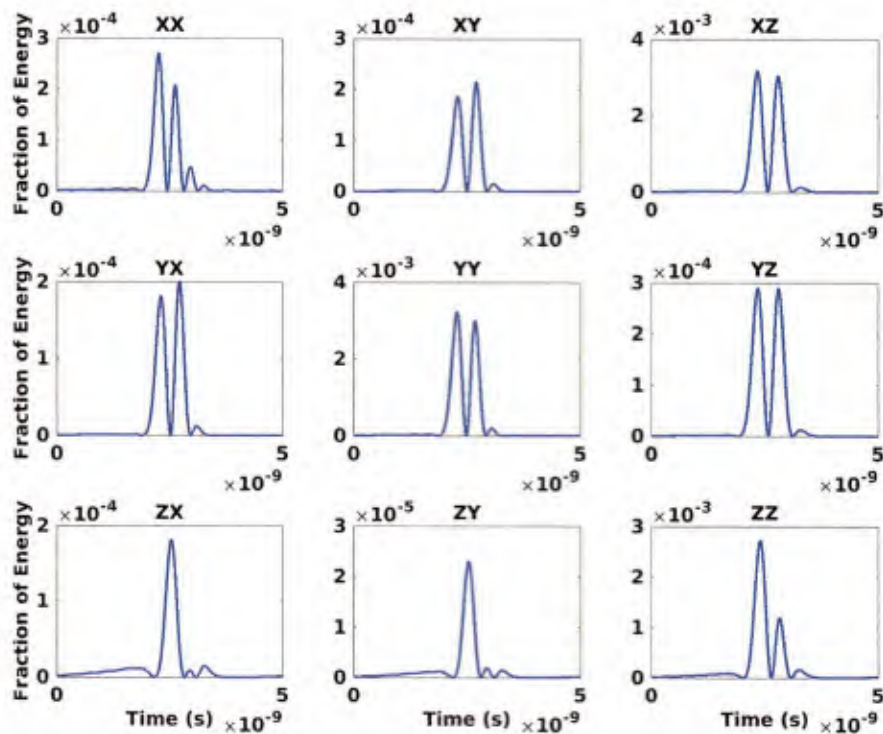


Figure 20 Extended Ray Approximation Error

2.3 3D Imaging Baseline

We quantitatively evaluated a baseline 3D imaging algorithm on a set of controlled laboratory measurements to assess how well we can locate and determine the shape of a representative buried targets. This work provides a baseline we can use to determine what extensions to the baseline algorithm are required to provide useful features for automatic target detection and classification. In this subsection, we will first describe the measurement data model and backprojection-based imaging model, and then show representative imaging performance on available controlled laboratory measurements.

2.3.1 Ground Penetrating Radar Data Used for Testing

A highly-controlled GPR dataset was made available by Professor Waymond Scott's group at Georgia Tech [Counts, et al. 2007]. Measurements were made of known targets at known positions buried in a container of sand in a laboratory setting. The transmitter and

receiver were moved to known positions using translation stages. Thus, these measurements provide a highly controlled data set to evaluate the ability of the imaging algorithm to locate and determine the shape of the buried targets.

A subset of the available measurement data was used to form 3D images of the buried targets. Specifically, a grid of $45 \times 45 = 2025$ collection points spanning an effective aperture area of 0.81 m^2 and centered on the target region were used to form the images. At each collection point, the scattering at 401 frequencies spanning a bandwidth of 8 GHz (60 MHz to 8.06 GHz), were collected, and the transmitter and receiver were separated by 12 cm. This data collection has a large amount of angle and frequency diversity, and can be considered an upper bound on the amount of data that may be available from field-portable equipment. The wide large amount of angular measurement diversity could potentially be achieved with a standoff system by mounting a GPR on a small UAS system.

In order to evaluate the performance of an imaging method, we need to compare to ground truth data. Waymond Scott has provided along with the dataset the ground truth diagram shown in Figure 21.




Figure 21 Georgia Tech Dataset Ground Truth

In the sections that follow we examine various different areas of interest in this dataset and benchmark the ability of the “ray” approximation described above to image the buried targets.

2.3.2 Comparisons of Reconstructions to Ground Truth Data

Centered around the point $(0, -40)$ and rotating the image 90 degrees to the left gives us the following area of interest:




Figure 22 VS-1.6 and TS-50 Ground Truth

The VS-1.6 in the center is a plastic-cased anti-tank mine with very few metal components. It is approximately 9.2 cm tall and 22.2 cm in diameter. To the left and right of the VS-1.6 are two TS-50's. Each TS-50 is approximately 4.5 cm tall and 9 cm in diameter. Below in Figure 23 is a reconstruction of the scene in Figure 22 using our simplified “ray model”.




Figure 23: Reconstruction of Area of Interest

We can clearly see the outline of all the mines and even the mine simulant. Looking at a look at a 1-dimensional “slice,” shown in Figure 24, where the VS-1.6 is supposed to be and compare it to the known dimensions of the buried target, we see that the target has been placed exactly where it belongs, according to the ground truth.




Figure 24 VS-1.6 Depth Profile vs. Ground Truth

Due to the linearity of the scattering model, instead of seeing one big “blob,” we see one localized excitation where the top of the target is and one small excitation where the bottom is. Note that the vertical axis on the “expected image” in Figure 24 is completely arbitrary and is only intended to show the accuracy of the imaging model at showing the depth of the target. We see two spikes clearly corresponding to the top and the bottom of the mine. Similarly, if we look at a 1 dimensional slice of one of the TS-50’s, we see in Figure 25 that the TS-50 mines appear just under the surface, matching the ground truth data.




Figure 25 TS-50 Depth Profile vs. Ground Truth

Once again, the vertical axis is arbitrary and the red line is used as a tool to show that the spikes in our reconstructed image correspond to the correct location and size in the ground truth data. Assuming the mine simulant used is the same size as the TS-50, we can also compare our image of the mine simulant to what we would expect.




Figure 26 Mine Simulant Depth Profile vs. Ground Truth

We can also evaluate the localization ability of the imaging algorithm by overlaying the expected spatial extent of each target on the reconstructed image, as seen in Figure 27 below.




Figure 27 Reconstruction of VS-1.6 and TS-50 with Spatial Extent Overlays

For this dataset we can reasonably conclude that the baseline imaging model is able to achieve an acceptable level of accuracy in terms of depth and location.

The following dataset contains a TMA-5 anti-tank mine with a crushed can on top of it. Our area of interest is centered around the point $(-0.45, 0.50)$.




Figure 28 TMA-5 with Crushed Can Ground Truth

The TMA-5 is a rectangular, plastic-cased mine enclosure with dimensions 27.5×30 cm and a height of 11 cm. Below in Figure 29 is a reconstructed image of the scene.




Figure 29 Reconstruction of TMA-5 with Crushed Can

The TMA-5 is rectangular, and we can almost start to make out the corners of the rectangle in the above image. As Figure 30 shows, the “blob” produced by our baseline imaging algorithm stays roughly confined to the expected extent of the rectangular target. One could argue we can even begin to see a slight peak where we would expect to see the crushed can.




Figure 30 Reconstruction of TMA-5 with Crushed Can with Spatial Extent Overlay

Looking at a one-dimensional slice through the middle of the TMA-5 reveals the accuracy of our depth estimation.




Figure 31 TMA-5 Depth Profile vs. Ground Truth

The general conclusion is that the baseline imaging model is able to achieve a reasonable level of accuracy in terms of depth and location for these targets in sandy soil given measurements with ample diversity in angle and frequency. In future work we intend to upgrade the imaging algorithm to estimate target material properties, to enhance shape estimation for more complicated geometries, and to preserve performance with data collections with more restricted angular and frequency diversity.

3.0 ATR Feature Extraction

Once images are formed, features are extracted from the images and used to perform the detection and classification. In the third quarter, we began to evaluate novel methods, based on convolution neural networks, to extract features from radar images and use them in detection/classification algorithms. In addition, a fundamental issue with machine learning-based algorithms was addressed. Namely, we began to consider methods of estimating the performance of a machine learning algorithm when it is applied to data not used to train the algorithm.

3.1 *Deep Neural Networks as Feature Extractors*

Machine learning detection and classification algorithms typically operate on feature vectors derived from raw sensor measurements and/or images generated from sensor measurements to perform detection and classification. Typically, features have been defined as part of the algorithm development process and are usually based on physical attributes, such as shape and dielectric contrast function values obtain from images, statistical attributes, such as local region statistics and contextual statistics, or image processing-based attributes, such as computer vision features (HoG, SIFT, etc.) and curvelets. Recently, deep neural networks have been advanced as a method of developing “machine-derived” features, which emerge from processing the training data. To begin to examine this concept, we applied some state-of-the-art deep learning methodologies to a set of available downward looking GPR measurements collected at a government test range to see how the machine-derived features compared to a previously developed feature set developed manually.

3.1.1 Convolutional Neural Networks

Convolutional neural networks are a class of machine learning algorithms that can be used to perform automatic feature extraction from rasterized images. Though their representational capacity as powerful function approximators has been well known for some time [Cybenko, 1989], they have enjoyed recent success in tackling complex problems in the fields of computer vision [Szegedy, 2015] and natural language processing [LeCun, 2015]. As shown notionally in Figure 32 convolutional neural networks perform automatic feature

extraction in a hierarchical fashion -- layers near the beginning extract low order features such as edges, and layers near the end extract high order features through nonlinear combinations of the low order features.

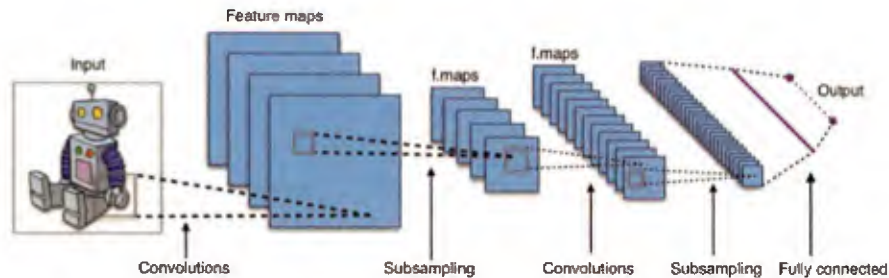


Figure 32 Hierarchical Feature Extraction with Convolutional Neural Network (Source: Wikipedia)

Similar to the multi-layer perceptron model, the output of each layer becomes the input to the subsequent layer. The convolution layers consist of multiple filtering operations followed by the pixel-wise application of a nonlinear function called the activation function. Mathematically, this is written as

$$x_i^k = f(w_i^k * o^{k-1})$$

where o^{k-1} denotes the output of layer $k - 1$, w_i^k is the kernel of the i 'th convolution filter in the k 'th layer and f is the activation function. The output of the k 'th layer consists of the "stack" of filtered images $o^k = [x_1^k, \dots, x_{n_k}^k]$ where n_k denotes the number of filters in the k 'th layer.

3.1.2 Training the Network

The final layer in the network is *fully-connected* in that each node is directly connected to every node in the previous layer. The convolutional layers, by contrast, are only connected to local patches of the previous layer (the size of these patches is determined by the filter size). Training is performed by backpropagating the error between the fully-connected layer output and the true class label. The weights are then updated according to the gradient of the backpropagated error.

Training is generally performed via variants of stochastic-gradient descent (SGD). Rather than computing the loss for all the training examples simultaneously as in classical gradient descent algorithms, SGC approximates the loss by randomly partitioning the training data into small subsets called batches or mini-batches. This training technique addresses the issue of memory constraints arising from massive datasets and has the added benefit of acting as a passive regularizer [Zhang 2016].

3.2 Convolutional Neural Network Feature Extraction Experiment

To investigate the ability of a convolutional neural network architecture to extract ATR features from GPR data, a deep convolution neural network was implemented in software and applied to an available downward-looking GPR data set. The features and detection performance produced by the convolutional neural network was compared to the detection performance produced by a boosting-based classifier using human-derived ATR features.

3.2.1 Network Topology Used for Experiment

Figure 33 shows one network architecture that was used for our experiments. For this application, we developed a deep convolutional neural network, with rectified linear activation functions and a single dropout layer for regularization. The network was created using the Python library Keras which is a wrapper for both Theano and Tensorflow. Theano is a deep learning library developed by the University of Montreal [Bergstra, et al., 2010], whereas Tensorflow was developed by Google [Abadi, et al, 2016]. Although our network was designed to classify downward looking GPR B-scans, the architecture is sufficiently flexible to perform classification on data acquired under different geometries.

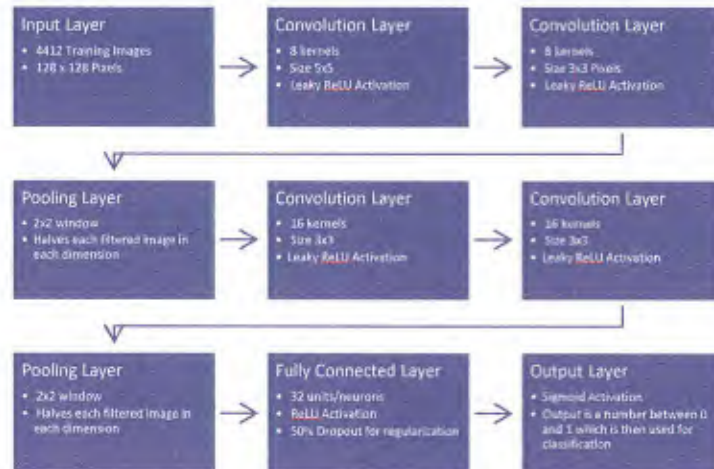


Figure 33 Neural Network Topology Used for Experiment

3.2.2 Hyperparameter Selection for Experiment

Hyperparameters are variables that are typically set before training starts and which determine various aspects of the network architecture. Hyperparameters include network depth/width (i.e., the number of layers and filters/nodes per layer), the size of minibatch used for backpropagation, the choice of optimization algorithm, the learning rate for such an algorithm, the choice of activation function, and regularization parameters. We summarize our choices for these in the table below.

Hyperparameter	Choice
Depth (Number of Layers)	4 conv. layers, 2 pooling layers, 2 dense
Width (Number of conv. filters/layer)	8-8-16-16 + 32 units in fully conn. layer
Filter Size	5x5, 3x3, 3x3, 3x3
Regularization	Single dropout layer with parameter 0.5
Optimizer	SGD with learning rate 0.01
Activation Function	Leaky ReLu with parameter 0.3
Kernel Initialization	He uniform

The simplest way to mitigate overfitting is by reducing the number of parameters in the network. It was found that a depth of 4 layers with the given number of filters and filter size was sufficient for the network to fit the data while maintaining the ability to generalize. Though we ran experiments with much larger networks, it was found that these would quickly overfit the data unless aggressive weight regularization was applied. Rather than going deeper and introducing further hyperparameters (in the form of weight decay coefficients), we opted for a smaller network with a single regularization parameter: the dropout rate. A dropout layer sets the output of each node in the preceding layer to zero with a certain fixed probability (the dropout parameter). Recent work has suggested the underlying reason for the effectiveness of dropout is that it approximates Bayesian inference [Gal 2015].

While classically the logistic function $f(x) = \frac{1}{1+e^{-x}}$ has been used as the activation in neural networks, in very deep neural networks this activation leads to the problem of vanishing gradients – the gradient of the backpropagated errors approach zero too quickly for effective learning. Recently, rectified linear units or ReLU's defined by $f(x) = \max(0, x)$ have been developed to address this issue. Leaky ReLU's defined by $L(x) = \begin{cases} x, & x \geq 0 \\ cx, & x < 0 \end{cases}$ where c is a small positive number have been shown to have similarly desirable properties, while not having the sparsifying effects of standard ReLU's [Xu 2015].

3.2.3 GPR Measurements of Buried Targets Used for Experiment

The GPR data used for the experiment were collected for the US Army NVESD using a prototype handheld GPR at a government test facility in the eastern US, where a wide variety of ground-truthed targets are buried in realistic configurations over multiple lanes. The measurements used to train and test the algorithm were B-scans over the target and clutter regions. The convolutional neural network was trained with multiple realizations appropriately labeled as target or clutter. As part of the training process, the system derived a set of machine-derived features from the raw B-scans, which were then used to perform target detection. Cross-validation was used to quantify detection performance.

The data used for our experiments were collected at a US government test facility, where a wide variety of ground-truthed targets are buried in realistic configurations over multiple lanes. Measurements were collected August 2015 using a Niitek demonstrator sensor which simultaneously collected impulse GPR and frequency domain EMI data. The data used for training our algorithms consisted of downward-looking GPR B-scans; a total of 4412 training examples, with 569 of these containing targets. Data includes metal and non-metal targets ranging in size and buried depth.

The B-scans were initially preprocessed using Robust Principal Component Analysis (RPCA) to perform ground-removal [Masarik et al, 2015]. The target signature was then centered in each scan using the Apex-Shifted Hyperbolic Radon Transform. As the data was complex-valued, only the magnitude of each B-scan was used. Finally, the B-scans were resized to 128 x 128 pixels before being fed into the neural network. The figure below shows a sample B-scan containing a target.




Figure 34-B-Scan Containing Target

3.2.4 Visualization of Feature Maps

As the input data moves through the network, multiple nonlinear filters and downsampling operations are applied. At the end of the network, we obtain a stack of filtered images which are then used for classification. The final convolution layer in our network has 16 filters. Figure 35 shows the effect of the first 8 of these on a B-scan containing a target shown in Figure 34.




Figure 35 Machine-derived feature maps for the target in Figure 34.

In effect, these filtered images become the new features that are used for classification. The final fully-connected layer of the neural network acts as the classifier using these features as inputs. It is worth noting that deeper networks typically include more downsampling operations in order to reduce the number of parameters in the network. Thus, the machine-derived features of a deeper network will not resemble the original input image due to the multiple downsampling operations. Previous work has suggested that the application of statistical classifiers to these machine-derived features can outperform the performance of the fully connected layers [Sakaguchi 2015].

3.2.5 Effects of Class Imbalance

Early-stopping is a technique commonly employed to prevent overfitting. A small portion of the training data, 10% in our experiments, is reserved for testing. Training of the network ceases when performance on the test set stops increasing. As shown in Figure 36, after approximately 8 epochs (i.e. times the algorithm has seen the entire dataset), the test/train curves begin to diverge, indicating that the algorithm has begun to overfit the data.




Figure 36 - Model Accuracy and Loss as a Function of Training Epochs

Though total loss on the test set is the most common metric used to track performance, the imbalanced nature of our data means that monitoring this quantity alone can bias an algorithm to correctly classify only those examples from the class for which more data is available. Furthermore, the fact that accuracy on the test set remains at around 87% suggests that the algorithm is simply classifying all B-scans as non-targets as 87% of the training data was from the non-target class. Two early stopping criteria were then tested and compared. The first criteria we tested was total loss, and the second was the area under the test set ROC curve (AUC).




Figure 37-Impact of Early Stopping Criteria on Classifier Performance

As shown in Figure 37, performance is significantly improved when using the AUC score on the test set as the stopping criteria. For the remainder of our experiments, AUC on the test set was used to monitor training progress and to stop training once performance was no longer increasing.

Previous work [Wallace 2011] has also suggested bagging with undersampling as a means of dealing with imbalanced classes. For our experiments a convolutional neural network using the architecture outlined in Figure 33 was trained using this technique. Undersampling was performed by randomly choosing examples from the non-target class without replacement to match the number of target-class examples. Performance was again compared against the boosting-based classifier that used human-engineered features as inputs as well as a CFAR prescreeener. Early stopping by monitoring test-set ROC AUC was used to determine when to stop training. The results of these experiments are shown in Figure 38.

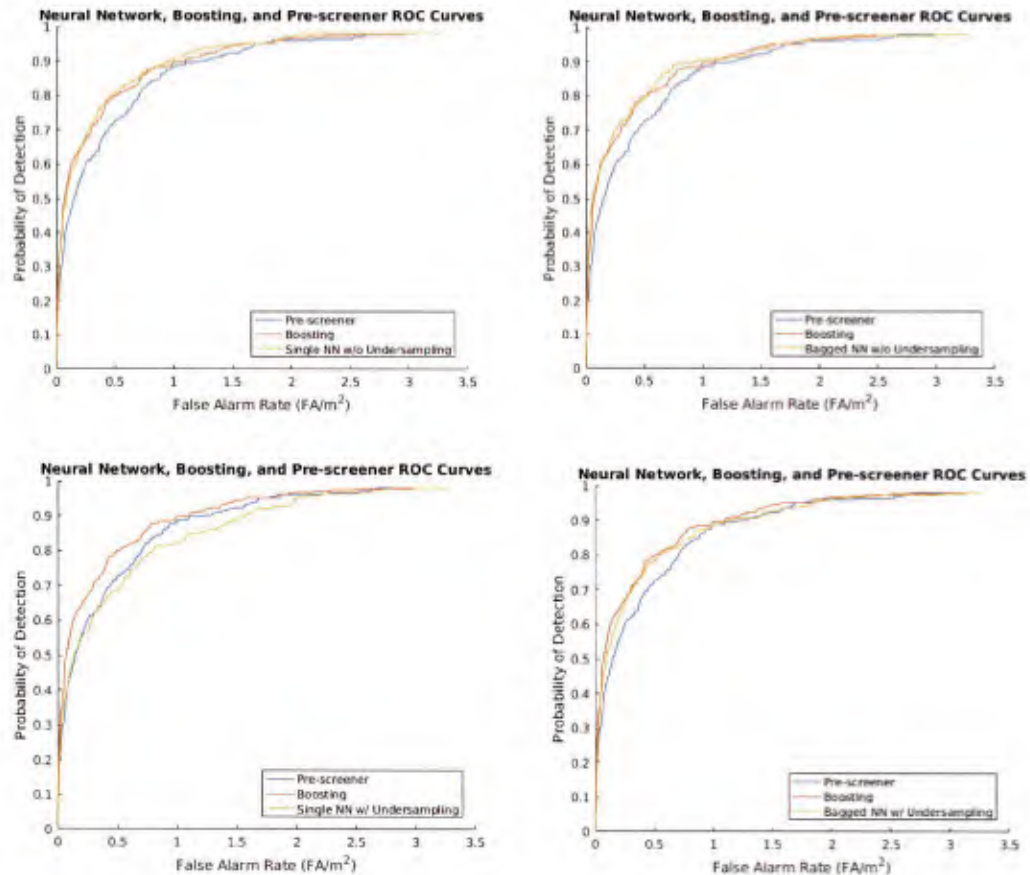


Figure 38 Effects of Bagging and Undersampling on Detection Performance

It was found that the performance of a single neural network without undersampling and without bagging was comparable to the performance when using bagging and undersampling. It was also noted, however, that bagging yielded more stable results as the performance varied significantly when using different random seeds to initialize the weights of the neural network.

3.2.5 Detection Performance of Neural Network Features

A boosting-based classifier using a human-derived feature set was previously developed for the data set [Masarik, et al., 2016], and its performance was used as the baseline for comparing the performance of the boosting-based classifier with machine-derived features.

The features that were used as inputs to the boosting classifier were based on absolute energy values and local energy contrast values. Specifically, the features were:

1. Energy at the detection location (1 feature)
2. Energy in a disk of radius 20 cm in the integrated energy ground map (1 features)
3. Ratio of energy in disks of radius 10, 20, 30, and 40 cm to the energy in a disk of radius 50 cm in the integrated energy ground map (4 features)
4. Ratios of energy in circles of radius 10, 20, and 30% of image size and centered about the image energy maximum to the total energy in the B-scan image (3 features)

Note that this feature extraction algorithm yields a total of 10 features from the GPR sensor. Moreover, several different feature extraction algorithms have been examined (e.g. using the discrete curvelet transform, using downsampled imagery, using standard computer vision features such as HoG and SIFT), and none of the additional feature sets yielded significant performance enhancements relative to the set of features described above. It is likely that for a diverse target set, generic features (like those described above) are most applicable.

Although we were able to perform the final classification directly using the neural network directly, performance was on-par with our boosting-based classifier using physics-based features. The ROC curves given in Figure 36 illustrate the performance of the different approaches. The blue pre-screener curve corresponds to a commonly employed CFAR detector. The yellow boosting curve corresponds to the boosting-based classifier using only the human-derived physics-based feature set. The red neural network curve corresponds to the convolutional neural network described above using the B-scans as inputs, and the purple Boosting NN curve corresponds to the boosting based-classifier using both the neural network and human-derived features. More concretely, the best performance was obtained by taking the classification scores output by the neural network and appending these as additional features for the boosting-based classifier. This approach of combining human and

machine-derived features leverages the power of neural networks as feature extractors while maintaining partial interpretability of the results via physics-based features.

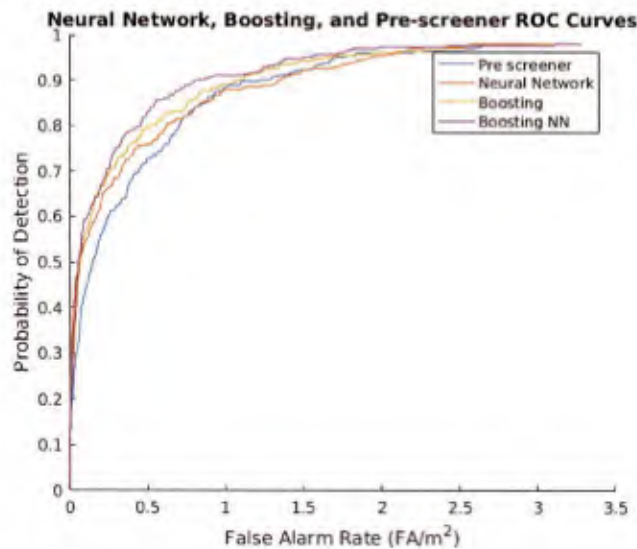


Figure 39 Comparison of Detection Performance Using Boosting with Human-Derived Features (Boosting), Neural Network Features (Neural Network), and the Combination of Human-Derived and Neural Network Features (Boosting NN)

The general conclusion from this initial experiment is that deep neural network architectures may provide an approach to feature extraction for standoff automatic target detection that performs as well or better than human-derived features sets, although the physical interpretability of machine-derived features may not be obvious.

4.0 Transfer Learning to Estimate Algorithm Extensibility

Machine learning-based detection and classification algorithms typically train an algorithm on a set of data which is considered to be a representative subset of the potential target and clutter realizations the sensor is intended to detect. To understand the robustness of the algorithm, it would be useful to have a bound on algorithm performance when the algorithm is applied to data outside the original training data set. Concepts from transfer learning were used to evaluate recent bounds on domain adaptation error.

4.1 Bound Formulation

In the context of this project, transfer learning refers to the application of machine learning algorithms to data drawn from a different distribution from what was used to train the algorithm. Recent bounds on transfer learning error using the triangle divergence [Berisha, 2016] were investigated to assess their viability in predicting field performance of existing algorithms.

In a best-case scenario, data that is used to train a machine-learning classifier is considered to be representative subset of the potential target and clutter realizations the sensor is used to detect. In practice, however, due to the high variability of targets as well as operating environments, this is seldom the case. Transfer learning refers to the application of machine learning algorithms to data drawn from a different distribution from what was used to train the algorithm. Previous works has established bounds on the Bayes error associated with a classifier when there are equal costs associated with false alarms and missed detections. In this section we detail our extension of transfer learning bounds to the case of unequal costs.

A hypothesis $h: \mathbb{R}^d \rightarrow \{0,1\}$ learned on some dataset S is a function that assigns to each feature vector its class membership. Given such a hypothesis we have the following bound on the expected risk ϵ_T on a separate target dataset T [Berisha 2016]:

$$\epsilon_T(h, y_T) \leq \epsilon_S(h, y_S) + E_{f_S(x)}[|y_S - y_T|] + 2 \sqrt{\hat{D}_1(f_S, f_T)}$$

Here, y_S and y_T are the true labeling functions in the source and target domains and f_S, f_T are the density functions of the feature vectors in each domain. The function \widehat{D} is closely related to the triangle divergence and is defined by

$$\widehat{D}_p(f, g) = \int \frac{(pf(x) - (1-p)g(x))^2}{p \cdot f(x) + (1-p)g(x)} dx, \text{ for } p \in (0,1).$$

Though the labeling function y_T is not known ahead of time, a reasonable assumption is that it is similar to y_S . The *covariate shift assumption*, commonly used in the transfer learning literature posits that $y_S = y_T$. If a hypothesis h achieves the Bayes error rate in the source domain, under the covariate shift assumption, the best case performance on the target domain is bounded above by

$$\epsilon_T(h, y_T) \leq \frac{1}{2} - \frac{1}{2} \widehat{D}_{\frac{1}{2}}(f_{S,0}, f_{S,1}) + 2 \sqrt{\widehat{D}_{\frac{1}{2}}(f_S, f_T)}.$$

The functions $f_{S,0}$ and $f_{S,1}$ are the densities of the features vectors in the source domain coming from class 0 and class 1 respectively. This bound was based on assigning equal weight to false alarms and missed detections. In practice, however, it is often the case that we wish to assign more weight to missed detections and in the most general setting we have a cost matrix C that associates a cost to each type of decision.

$$C = \begin{bmatrix} c_{00} & c_{01} \\ c_{10} & c_{11} \end{bmatrix}$$

We write c_{ij} for the cost associated with classifying a feature vector as coming from class i when it really came from class j . We may then define the risk associated with a particular hypothesis h as follows:

$$\epsilon^C(h, y) = c_{01} \mathbb{E}[(1-h)|h-y|] + c_{10} \mathbb{E}[h|h-y|] + c_{00} \mathbb{E}[(1-h)(1-y)] + c_{11} \mathbb{E}[hy]$$

Where the expectation is taken with respect to the underlying distribution of the data. Our first result gives an upper bound on the Bayes risk in terms of the triangle divergence.

$$\epsilon_{Bayes}^C \leq \frac{1}{2} + (p_0 c_{00} + q_0 c_{11}) - \widehat{D}_{p_0}((c_{10} - c_{00})f_0, (c_{01} - c_{11})f_1)$$

Here p_0, q_0 denote the prior probabilities of drawing a feature vector from class 0 and class 1 respectively. Our result is valid whenever $c_{10} - c_{00} > 0$ and $c_{01} - c_{11} > 0$ – this is a common situation wherein we associate a higher cost with incorrect decisions as compared to correct ones. We note that in the case that $c_{00} = c_{11} = 0$ and $c_{01} = c_{10} = 1$, the equation above reduces to the familiar expression

$$\epsilon_{Bayes} \leq \frac{1}{2} - \frac{1}{2} \widehat{D}_{\frac{1}{2}}(f_0, f_1).$$

given in [Berisha 2016]. The analysis was further extended to give a bound on the expected transfer risk in the covariate shift scenario

$$\begin{aligned} \epsilon_T^c(h) \leq & \frac{1}{2} + (p_0 c_{00} + q_0 c_{11}) - \widehat{D}_{p_0}((c_{10} - c_{00})f_0, (c_{01} - c_{11})f_1) + 2c^{10} \sqrt{\widehat{D}_{\frac{1}{2}}(hf_T, hf_S)} + \\ & 2c^{01} \sqrt{\widehat{D}_{\frac{1}{2}}((1-h)f_T, (1-h)f_S)} + 2c^{00} \sqrt{\widehat{D}_{\frac{1}{2}}((1-h)f_T, (1-h)f_S)} + 2c^{11} \sqrt{\widehat{D}_{\frac{1}{2}}(hf_T, hf_S)}. \end{aligned}$$

4.2 Simulation to Assess Bounds

To investigate the utility of the bound given in [Berisha 2016], a simple numerical experiment was generated. In the example shown in Figure 39, we consider simple 1-D classification problem. The per-class distribution for the source data are Skewed-t distributions centered at $x = 4$ and $x = 10$. In other words, the blue and red curves correspond to $f_{S,0}$ and $f_{S,1}$ respectively and f_S is the mixture distribution $f_S = \frac{1}{2}f_{S,0} + \frac{1}{2}f_{S,1}$. Nominally, this corresponds to data available at training time.




Figure 40 Assumed Distributions for Transfer Learning Experiment

It is reasonable to expect that the distribution of data whereupon the algorithm will be deployed will “drift” slightly from the data that was used to train the model. This is partially due to the fundamental problem of having insufficient data to create an exhaustive model of target/non target distributions under all possible environmental conditions (sometimes called dataset bias). In our experiment, this drift was manifested as a shifted mean in the class-1 distribution. Our target distribution was thus given by $f_T(x) = f_{S,0}(x) + f_{(S,1)}(x - \delta)$. By allowing the mean to vary over a small range, we were able to quantify the best-case performance, the real-world performance, and the theoretical upper bound on the classification error.




Figure 41 Error Rates for Perturbed Data

After perturbing our data, the new Bayes error rate, true error rate and upper bounds were calculated using Monte Carlo methods. As shown in Figure 41, while the upper bound given in [Berisha, 2016] is indeed valid, it is not particularly useful. While the empirical error rate remains close to the Bayes error rate for small perturbations of the Class-1 distribution (as to be expected), the upper bound quickly grows as we move away from the source distribution. This suggests a tighter bound is necessary in order to quantify transfer error bounds than what is provided by the triangle divergence. Deriving a tighter bound will be a subject for future research.

5.0 Improvement of Buried Object Imaging Using Lifting Approach to Inversion

Since buried object imaging involves propagation through media that is, at best, partially known, there is mismatch between the forward model used in the inversion and the propagation behavior actually observed in the measured data. The mismatch can cause degradation and/or reduce resolution in the imagery, which limits the features that can be extracted from the imagery, which, in turn, limits ATR performance.

In recent years, several groups have advocated that backpropagation of interferometric measurements can produce more statistically stable estimates of targets in the presence of forward model errors and in the presence of clutter [Borcea et al, 2005; Borcea et al 2011]. MTRI has applied variations of these techniques to field measurements of buried targets collected by handheld ground penetrating radar systems [Burns et al. 2016; Masarik et al., 2016] and vehicle-based forward looking ground penetrating radar (FLGPR) systems [Webb et al., 2017; Xique et al. 2018]. While these techniques do produce more statistically stable estimates, which enable weak targets to be observed, the inherent averaging in the algorithms tends to reduce the spatial localization of targets, which reduces discriminability.

More recently, Demanet and Jugnon, [2017] proposed quadratic lifting inversion (QLI) as a robust approach to inversion in the presence of mismatch or errors in the forward model, which can produce reconstructions with fidelity comparable to direct inversion. On

the basis of theory and simulations, Demanet and Jugnon claim that QLI has to potential to accomplish all three of the following goals:

- robustness to model mismatch,
- comparable resolution with classical techniques (backprojection, least squares auto-focus, etc.),
- and statistically spatially stable recovery (with respect to randomly heterogeneous media).

We applied this technique to radar imaging of buried targets to determine if it can produce enhanced imagery in the presence of limited knowledge of the surrounding ground geometry and/or material properties. The approach can potentially improve imagery, and the resulting image features, which would lead to improve ATR performance.

In the following, we evaluate these claims on measured data. Sections 5.1 and 5.2 introduces our implementation of QLI. Section 5.3 is divided into three subsections, each of which details an experiment which evaluates one of the above claims made about QLI. In the first (Subsection 5.3.1), we show by simulation that QLI is robust to model mismatch. In the second (Subsection 5.3.2), we show that QLI can produce imagery in a controlled experiment, with measured data, that has resolution comparable to backprojection. Subsection 5.3.3 completes the survey by testing the stability of the method against data collected in an experimental but randomly heterogeneous (buried target) setting. Section 5.4 summarizes our findings and comments on future directions for research

5.1 Lifting Approach to Inversion and Comparison to Alternatives

Consider a basic forward problem of the form

$$d = FX + n$$

where d is the observed data, F is a forward model, X is the desired image and n is noise. Assume measurements are made of various cross-products, i.e., the actual data for inversion is of the form

$$B_{ij} = d_i d_j^*$$

From this, the lifting approach carries out an estimation of x via solving the following problem

$$\tilde{X} := \underset{X}{\operatorname{argmin}} \left\{ \left\| E \circ (\tilde{F} X X^\dagger \tilde{F}^\dagger - d d^\dagger) \right\|_2^2 + \mu \rho(X) \right\}$$

where E is an operator that only keeps some of the interactions, i.e., those that are deemed to be valuable and potentially informative and $\mu \rho(X)$ is a regularization term. Typically, E would correspond to measurements at locations that are in the same spatial neighborhood. Demanet and Jugnion [2017] have proposed lifting as a robust approach to inversion in the presence of mismatch or errors in the forward model, which can produce reconstructions with fidelity comparable to direct inversion. They have applied the algorithm to seismic and medical imaging applications. We examined the application to buried object imaging using ground penetrating radar data.

5.1.1. Comparison of Backprojection to Quadratic Lifting Inversion

One of the most straightforward ways to form an image of a target from synthetic aperture radar (SAR) data is via backprojection of the recorded phase delay measurements. Backprojection approximates inversion of the true scattering map

$$F: \text{Scene} \rightarrow \text{Data}, X \mapsto d$$

using the adjoint of a model (scattering) operator \tilde{F} , which is believed to be a reasonable approximation of F . The resulting image is thus

$$\tilde{X}_{BP} = \tilde{F}^\dagger d.$$

An alternative technique is to compute the least squares fit

$$\tilde{X}_{LS} := \operatorname{argmin}_X \left\{ \left\| \tilde{F}X - d \right\|_2^2 + \mu \rho(X) \right\},$$

where the regularization term $\rho(X)$ with weight μ imposes certain constraints on the reconstructed scene \tilde{X}_{LS} . The limitation of the above formulations is that they rely directly on the fidelity of the model \tilde{F} .

5.1.2 Comparison of Interferometric Imaging to Quadratic Lifting Inversion

As mentioned above, interferometric imaging forms cross-correlations dd^\dagger , obtaining a matrix of products $\{m_{ij}\} = d_i \bar{d}_j$ from indexed measurements $\{d\}$. In radar imaging / phase retrieval problems, the heuristic motivation is that this matrix of quadratic measurements can be thought of as recording travel-time differences. As such, products of spatially and spectrally local pairs d_i, d_j are likely to be less sensitive to typical model errors such as sensor imperfections and imprecise sensor locations. This is explained more fully in section III.D of Demanet and Jugnon's paper.[2017] For this reason, interferometric algorithms like coherent interferometric imaging (CINT) [Borcea, et al. 2005] and QLI employ a matrix E of weights/masks to choose only those correlations that occur on short spatial and frequency scales. (Choosing E to be the identity matrix yields phaseless imaging as a special case.)

The CINT imaging process can be expressed as

$$X_{CINT} := \operatorname{diag}\{\tilde{F}^\dagger (E \circ [dd^\dagger]) \tilde{F}\}$$

while QLI relates the image \tilde{X}_{QLI} to the solution of the optimization problem

$$\tilde{R} := \operatorname{argmin}_R \left\{ \left\| E \circ (\tilde{F}RR^\dagger \tilde{F}^\dagger - dd^\dagger) \right\|_2^2 + \mu \rho(X) \right\}$$

A strength of QLI is that the optimization can be reformulated as a higher dimensional problem, in which $R \in \mathcal{C}^{N \times K}$, for some $K > 1$. Then we still have $RR^\dagger \in \mathcal{C}^{N \times N}$, but the problem is less susceptible to trapping by local minima due to this rank K relaxation.[Demanet and Jugnon, 2017]. The image \tilde{x}_{QLI} is then constructed by computing the principal eigenvector of $\tilde{R}\tilde{R}^\dagger$. Unfortunately, this relaxation step is not computationally

feasible for $K \geq 2$ in most of our experiments due to the dimensionality ($d = 3$ in Sections 3.2 and 3.3) and large pixel counts (often more than 10^4) of our images. For this reason, we set $K = 1$ in the relaxation, resulting in an estimation of the scene $\tilde{X}_{QLI} := \tilde{R}$ directly as a vector. For the rest of this paper, therefore, we will identify the two and use X to represent the variable of optimization. We choose $\rho(X) = \|X\|_2^2$ as our regularization term, and we refer to the objective function as $L_{\mu,E}(X)$.

Spectral gap analysis of the graph Laplacian associated to the selection matrix E shows that auto-correlations ($E = I$) and even purely spatial correlations are -- in general -- insufficient to achieve model robustness [Demagnet and Jugnon, 2017]. However, experiments by Mason and Yazici [2018] suggest that such kinds of phaseless or incoherent interferometry of this flavor can succeed if additional information is available. Figure 1 shows the spatial ordering and correlation windows for our experiments on measured data. Here we are including cross-correlations between local frequency/spatial locations pairs as depicted in Figure 1.

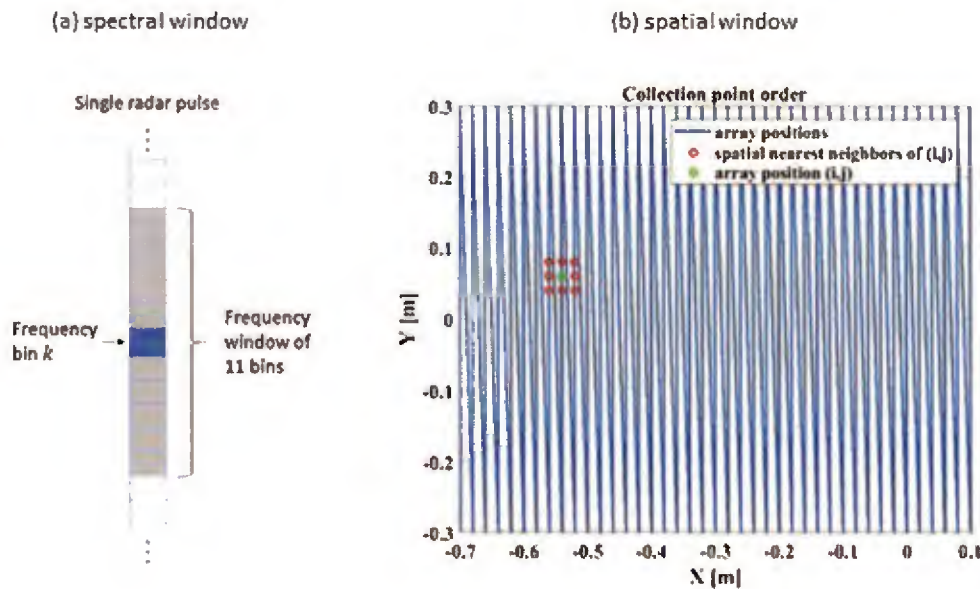


Figure 42 Windowing of data via selection matrix E : data from frequency bins in a fixed spectral window on the left (of width 11) and a fixed spatial window on the right (red is cross-correlated with green) are all included in the cross-correlations selected via E .

We note that due to the diagonal nature of our choice of E and the linear ordering of the data vector d , we end up performing some correlations across ends of frequency sweeps and edges of 2-D array apertures that exceed any reasonable coherence length. Elimination of these non-local correlations could be accomplished by generating a dense mask that directly computes the proximity of data samples in frequency and spatial dimensions, but such an adjustment would give up at least part of the computational efficiency of our vectorized implementation for computing $L_{\mu,E}$ and its directional derivatives. Because the windows we have chosen are short relative to our data dimensions, these non-local correlations are rare enough that we expect them to have little effect on the QLI reconstruction.

5.2 Lifting Inversion Algorithm Implementation

Minimization of the loss functional $L_{\mu,E}(X)$ is achieved via a limited memory Broyden-Fletcher-Goldfarb-Shanno (hereafter “LBFGS”) quasi-Newtonian optimization algorithm [Liu and Nocedal, 1989] written in MATLAB. We provide this scheme with a MATLAB function which computes the value of this functional and its gradient with respect to X . The latter quantity is as follows:

$$\nabla_X L_{\mu,E} = 4F^{\wedge \dagger} (E \circ A) F^{\wedge} X + 2\mu X$$

where we have abbreviated $A := (\tilde{F} R R^{\dagger} \tilde{F}^{\dagger} - d d^{\dagger})$. The LBFGS optimizer computes iterative estimates of the target scene, terminating when the L_2 norm of this gradient drops below a prescribed threshold (our default value was 2×10^6 times the L_2 norm of the initial gradient). For each of the three experiments described below, reasonable regularization parameters have been empirically chosen from a range of candidate values.

Note, however, that the structure of the regularized functional is not invariant to scaling transformations, nor to variations in the density of matrix E . Perhaps for this reason, we have so far been unable to identify a heuristic analogous to that of Borcea et al. [2006] for tuning the matrix E to improve image quality. Continuing this search is a likely direction for future research.

Moreover, the regularization parameter μ is an absolute weight. For our experiments, this means that different preprocessing or sampling of the SAR data (or even different image cube pixel counts/spacings) can lead to different effective regularization weights, thus motivating a more automatic selection of regularization as a topic of future research.

For the same imaging regime as mentioned earlier in this section, it is also memory-prohibitive to form XX^\dagger explicitly. Instead, by choosing E to have constant-diagonal structure, we can compute only those entries we consider spectrally and spatially relevant (those for which $e_{ij} = 1$) in a manner reminiscent of discrete convolutions of $\tilde{F}X$ with itself. This diagonal structure of E corresponds to fixed windows with equal weight in frequency and space, and is (mostly) consistent with the model depicted in Figure 1 (see earlier discussion at the end of Section 5.1.2. More precisely, we compute

$$\left[E \circ (\tilde{F}XX^\dagger\tilde{F}^\dagger) \right]_{ij} = \left[E \circ (\tilde{F}X(\tilde{F}X)^\dagger) \right]_{ij} = e_{ij}(\tilde{F}X)_i \overline{(\tilde{F}X)_j}$$

for each nonzero entry e_{ij} in E . For each cross-correlation index pair (i, j) , we form a copy of $(\tilde{F}X)$ shifted by $i - j$ positions. By zero-padding and stacking these vectors into a matrix D_E , we can calculate all the entries (in diagonal format) using a single vectorized, element-wise product $D_E \circ (\tilde{F}X)^\dagger$. Since E is a sparse diagonal matrix (with density less than 0.1% in our experiments), the reduction of the storage and computational burdens of this process from $O(N^2)$ to $O(N \cdot |e_{ij} \neq 0|)$ is substantial. Due to this reduction, the main algorithmic cost is computation of the linear operation $\tilde{F}(X)$ (and its adjoint \tilde{F}^\dagger , in calculation of the gradient $\nabla_X L_{\mu,E}$). For the problem we consider, these operators employ non-uniform fast Fourier transforms (NUFFT) on data sets of order $O(N_{pulses} \cdot N_f)$.⁷ The operators thus have complexity of $O(N_{pulses} \cdot mN_f \log_2 N_f)$, where m is a number related to the oversampling and interpolation performed inside the NUFFTs.

5.3 Example Lifting Inversion Results

In this section we show the results of utilizing QLI on simulated and measured data. In particular, in the first (Subsection 5.3.1), we show by simulation that QLI is robust to a periodic model mismatch. In the second (Subsection 5.3.2), we show that QLI can produce imagery in a controlled experiment with measured data that has resolution comparable to backprojection. Finally, subsection 5.3.3 completes the survey by testing the stability of the method against data collected in an experimental but randomly heterogeneous (buried target) setting.

5.3.1 Robustness of Lifting to Phase Error and Noise

In this section, we demonstrate robustness of QLI in a radar imaging application to a fundamental model mismatch caused by unknown sensor position errors. The effect of this, to first order, is to induce phase errors. Our experiment shows that successful application of QLI reduces the need for exquisite autofocus/motion compensation/estimation. We also show that QLI is stable against additive noise up to very low (even negative) signal to noise ratio (SNR). We generate data from a simulated point scatterer at $(0, 0, 0)$ using 100 measurements over a wide frequency bandwidth of 0.5-2.0GHz, and sampling returns at 201 positions on a 6m synthetic aperture at a 2m standoff from the scene center. The aperture sample spacing was chosen to achieve better than Nyquist rates ($\Delta x < \lambda_{min}/4$, where $\lambda_{min} = 15\text{cm}$). According to our calculations, this geometry admits a resolution of 6.1cm in azimuth (X) and 10cm in range (Z). We choose a far smaller pixel spacing of 1.67cm in each dimension to expose the point spread function and facilitate comparison between the resolution and performance of QLI imaging and backprojection.

In these experiments, we simulate an ideal (linear) array geometry and a perturbed array that differs from it by an additive sinusoidal perturbation in the Z direction (see Figure 43). We synthesize clean data using the true distance measurements ($d_{clean} = FX_{target}$). Complex-valued additive Gaussian white noise (AWGN) with variance $\sigma^2 = 2\eta^2$ is then

added across frequencies and pulses to simulate instrumental/environmental effects so that the synthesized data has the form

$$d := d_{clean} + \eta N_{complex} = FX_{target} + \eta[N(0,1) + iN(0,1)]$$

The optimization utilizes the ideal array geometry for both imaging and model operators ($\sim Fy, \sim F$), which means that the algorithm fits the noisy data using a flawed model. Despite this disadvantage, QLI is able to improve location and resolution of the point target over that of backprojection and regularized least squares. For this experiment, the selection matrix E used only the two nearest-in-memory neighboring array positions instead of the nine pictured in Figure 42. The regularization used in the optimization was $\mu = 0.1$. The resulting images from QLI and backprojection are compared in Figure 44.




Figure 43 Collection geometry with sinusoidal measurement error with a peak amplitude of 5 cm. The ideal and perturbed measurement locations (used for the forward model and test data synthesis, respectively) are shown, along with the true point target locations and imaging grid.

To measure the impact of QLI over backprojection, we compute a ratio of geometric means over range and cross-range half power mainlobe widths $G = \sqrt{\frac{W_{lift,x}W_{lift,y}}{W_{BP,x}W_{BP,y}}}$, computed for parametrically varying measurement errors and noise realizations, with $W_{m,d}$ representing this width across axis d for an image formation method m . These widths are computed from

data slices through the center of the scene (ground truth for the point scatterer). The statistic $1 - G$ presented in Figure 45 is thus measure of the overall reduction in width by QLI reconstruction over backprojection. Due to the cyclic ambiguity of phase information, it is natural to expect phases with errors of $\frac{1}{2}\lambda_{min} = 7.5\text{cm}$ (or more) to be difficult or impossible to recover without an error model. We therefore call these errors “(super-)critical”, and do not show results for this regime in Figures 44 or Figure 45. We found in our simulations that the QLI algorithm is largely stable up to low (even negative) SNRs, and that it only fails to form a unified return as we approach the super-critical error regime.

This robustness to phase error is not shared by linear least squares optimization. For a wide range of regularization parameters, minimization of the least squares (LS) loss/objective function even accentuates split returns. (See the third plot in Figure 46 for one such result at regularization weight $\mu = 500$.) This is the case even when the initialization image is backprojected using the true adjoint operator F^\dagger (see the leftmost plot of Figure 46. Thus in simulation, QLI has showed robustness to mild model mismatch.




Figure 44 Quadratic lifting improvement (bottom) over backprojected images (top) in the presence of increasing sinusoidal measurement error (left to right from 0cm to 7.5cm) and constant 7dB SNR. Small red circles indicate truth location of the point target in each plot. Coordinates are in pixels. Images are energy normalized, with dynamic range is 30dB and peaks set (per column) by the backprojected image.

		SNR		
		∞ dB	15dB	6dB
error	0 cm	11%	11%	11%
	2.5 cm	14%	14%	15%
	5 cm	24%	24%	25%

Figure 45 Values of the half power mainlobe width percentage reduction statistic $I-G$ for a variety of measurement error and SNR levels, showing the improvement of QLI over backprojection.




Figure 46 Regularized least squares does not share robustness of QLI. Images were formed from noise-free synthetic data with 5cm amplitude sinusoidal measurement errors. From left to right: backprojection using F^\dagger , backprojection using \tilde{F}^\dagger , least squares optimization using $\tilde{F}, \tilde{F}^\dagger$, and QLI using $\tilde{F}, \tilde{F}^\dagger$. Coordinates are in pixels; dynamic range is 30dB. Each image is peak normalized.

5.3.2 Resolution Comparable to Backprojection on Measured Data

Next, we evaluated the ability of QLI to produce imagery with resolution comparable to backprojection from actual free space measurements when there is limited model

mismatch. On images formed from measured data, we have observed resolutions and clutter levels that compare favorably to fast backprojection. Our experiments were performed on one of the free space targets from a data set recorded at Georgia Tech University under carefully controlled conditions and made freely available at

<https://waymondscott.ece.gatech.edu/multistaticbeamformingdata/multistaticbeamformingdata/>

[Counts et al, 2007]. The right half (sub-figures “a” and “b”) of Figure 47 depicts the ground truth and collection geometry for this data. Although multiple receivers were used in the collection, in our tests we used only the first transmit/receive antennae pair (“T1-R1”).

For this experiment, the target was a $38.5 \times 46.5 \times 1.8$ cm plywood Georgia Tech logo supported (above ground) on a Styrofoam base. The data consist of 8281 radar pulses, collected over a 2D square grid of positions at a constant height above the target. Each pulse contains 401 frequencies on a wideband of 60MHz-8.06GHz. We used the 1681 pulses collected in an 80×80 cm region directly above the target. In order to obtain an even number of frequencies (400), we ignored the highest frequency bin.




Figure 47 Collection geometry, targets, and buried mine configuration in Georgia Tech study⁸ (images reproduced and cropped from Counts et al. 2007).

Before image formation, we de-meant the frequency data across pulses to remove instrumental noise or bias and then additionally range-gate the data to remove the instrumental noise, feedthrough, and clutter from outside the target volume. Having thus preprocessed the data, we backproject an initial $31 \times 31 \times 11$ pixel scene estimate (select slices are displayed in the first row of Figure 48) and take $\mu = 0.3$ as our L_2 (absolute) regularization parameter for the QLI optimization. For this experiment, the selection matrix E was chosen so as to use cross-correlations across 9 spatial positions and 11 frequency bins as depicted in Figure 42. After 75 iterations, our LBFGS optimization scheme achieved gradient convergence (see discussion in Section 2.2) and we obtained the 3D image shown as slices in the second row of Figure 48. The optimization took less than 15 minutes of execution on a 3.9 GHz quad-core i7-3940XM cpu in a DELL Precision M4700 laptop with 32 GB of DDR3 RAM. The QLI reconstruction shows resolution comparable to backprojection, and contrast has been improved.

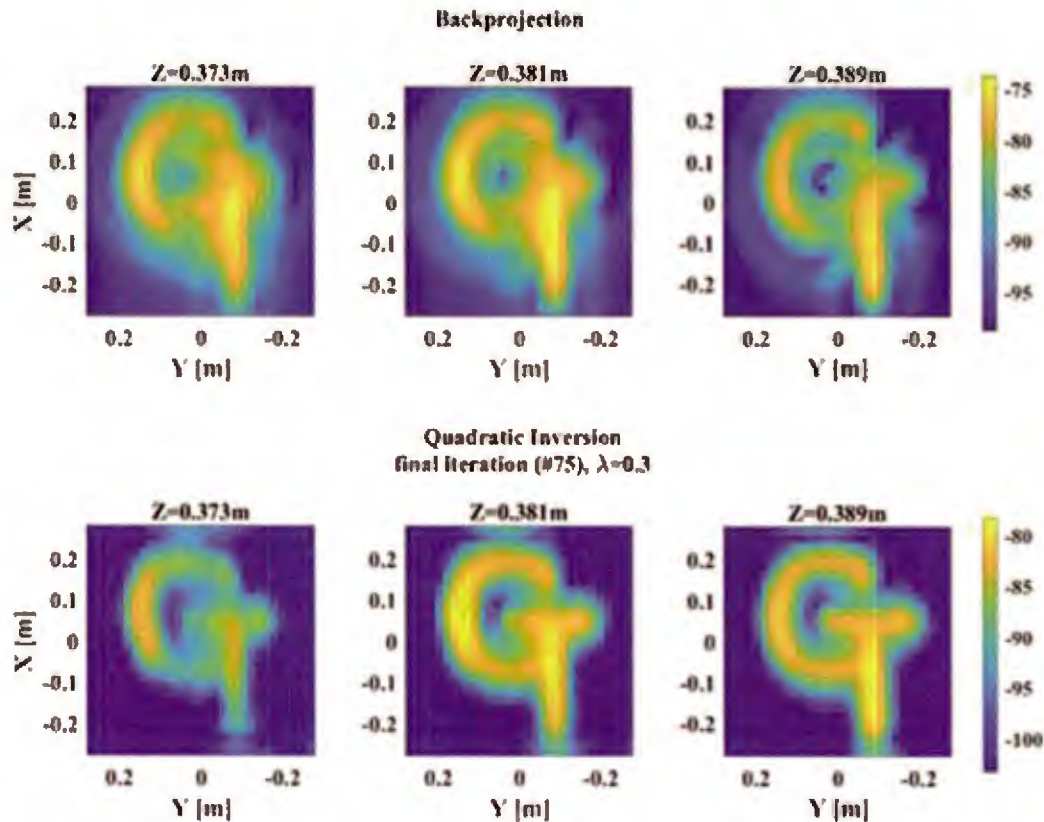


Figure 48 “GT Plywood” Target: Comparison of backprojected image vs. QLI reconstruction. We show horizontal slices of interest at heights near expected true height at a dynamic range of 25dB to match that of Counts et al., 2007. No post-processing/normalization has been performed to equalize color scales.

5.3.3 Application to Buried Target Imaging

For a more difficult problem on which to test the lifting formulation, we turned to the buried target data publicly available from the same Georgia Tech research team.[Counts et al., 2007] The subset of this data that we chose included an assortment of targets buried up to 12cm deep in clean, packed sand. See the right half of Figure 47 for ground truth target locations (depths are given in parentheses beneath the label for each target), and the paper for more details of the data collection.

Before lifting, we preprocessed the data, de-meaning across pulses in an attempt to remove the mostly uniform surface returns (<2cm height variation). We used a two-layer scattering model with calculation of Snell refraction points (Figure 49) when computing time lags for our imaging routines (using code obtained from [Masarik et al., 2015]). Determination of these refraction points is accomplished by solving a fourth degree polynomial equation for each transmit/receive pair and pixel location, using the Georgia Tech team's observation of wave propagation speed through the sand. These solutions for the refraction points were obtained by means of Newton's method for finding zeros of a function.




Figure 49 Two-layer forward model for wave scattering from buried target that includes refraction/propagation through ground

After searching the parameter space (and based on results from previous experiments), we settled on an absolute L_2 regularization parameter of $\mu = 0.05$. We selected 11 inter-frequency and 9 inter-pulse cross correlations depicted in Figure 1. The LBFGS optimization scheme described in Section 5.2 was run for 66 iterations (less than 8 minutes on the same laptop mentioned in Section 5.3.2). We chose as our solution an intermediate estimate of the target scene that was selected by identifying the first iteration on which the norm of the

solution began to increase (see Figure 50). This is our proxy termination condition in the absence of sufficient gradient decay, since consistent achievement of the latter has so far proved elusive. Figure 51 depicts the solution thus chosen. It is important to note that in the QLI reconstruction, the impulse response has been narrowed. This is especially apparent on the deep target. Clutter has also been significantly reduced.




Figure 50 Optimization behavior of QLI objective function $L_{\mu,E}$. In order to surmount the lack of algorithm-specific preconditioning in our implementation, we select as our buried target reconstruction the estimate from the optimizer iteration where $\|X\|_2$ first begins to increase (yellow dashed line).




Figure 51 Buried Targets: backprojection vs QLI. We show horizontal slices of interest at nominal depths of 4, 9, and 14cm below ground-air interface. Note that targets appear a few centimeters deeper than predicted by ground truth (see Figure 47 or the Georgia Tech paper⁸), and that there is no target at a depth of 9cm.

We did explore finding an optimal combination of optimization and regularization for carrying out QLI on these data, but we have not been successful to date. However, we did find some significant patterns in this investigation. From a plot of the objective function values $L_{\mu,E}(X_i)$ (a.k.a., loss) alongside the L_2 norm of the estimated scene X_i (a.k.a., solution) at each iteration, it appears that the target to clutter ratio was improving principally when the norm of the scene was decreasing. This plot is shown in Figure 50 and iterations of interest are labeled. Our first attempted solution to this problem was to increase the regularization parameter, but an optimal value has proved exceedingly elusive. The transition between over-regularization and the behavior exhibited in the results shown here seems to be rather sharp. To surmount this difficulty, we eventually introduced the stopping criterion described above and illustrated in Figure 50. With this criterion, we were able to reconstruct a sharper image of buried targets than backprojection.

5.4 Observations on the Application of Quadratic Lifting Inversion

In this section we have described our rank 1 implementations of quadratic lifting inversion and evaluated three claims about the method by conducting imaging experiments using simulated and measured data. We describe how to efficiently compute data-space cross-correlations by designing the selection matrix E to be constant along diagonals and avoiding the full formation of the product XX^\dagger , at the cost of including relatively few unwanted data correlations. Via a synthetic data experiment, we have showed that QLI has robustness to model errors not shared by backprojection or least squares inversion. In the presence of a defective array geometry, these classical methods produced split returns from a single point target, while QLI was able to recover a unified return for array model inaccuracies of up to 7.5cm. On measured data, we have demonstrated that QLI attains image resolution comparable to that of backprojection, improves clutter suppression, and increases contrast of the targets. Finally, we were able to image buried targets using a two-layer model without carefully tuning the propagation speed.

Along the way, we encountered several promising directions for future work. Prominent among these are (i) the search for a heuristic to tune density of the selection matrix E and the regularization constant μ to characteristics of the measured data; (ii)

exploration of a more principled automatic truncation rule or preconditioning of the problem to replace the one outlined in Section 5.3.3 and illustrated in Figure 50; and (iii) the testing of QLI for higher rank ($K > 1$) relaxations on 3D/GPR imaging problems, which currently appear out of reach of our current implementation and hardware. A less pressing direction might be the design of a scaling invariant QLI formulation.

6.0 sUAS based Radar Measurements of Buried Objects

Current standoff radar data collection systems, such as vehicle-based FLGPR systems, view the ground region of interest (ROI) at shallow grazing angles. Thus, only a small amount of the transmitted radar energy propagates into the ground, and an even smaller portion of the target backscattered energy is returned to the sensor. Mounting a low size, weight and power (SWAP) radar on a small unmanned aerial system (sUAS) platform provides a method of interrogating the ground ROI at more favorable angles of incidence with the radar operator being at a safe standoff distance from the buried explosive threat. Reliable detection and/or imaging of a buried object from a sUAS is influenced by several issues related to sensor power, measurement of sensor position and control of the sUAS in the proximity to the ground. In principle, these issues can potentially be mitigated by hardware and/or software solutions. To begin investigation of some of the basic issues associated with this type of system, we collected ground penetrating radar measurements of buried objects from using a low SWAP radar on one of MTRI's sUAS and processed the resulting data to determine the ability to detect and image buried objects.

6.1 sUAS Measurement Hardware

MTRI mounted a four port, wideband, low frequency stepped frequency continuous wave (SFCW) radar built by AKELA on a Bergen octocopter to collect downward looking ground penetrating radar measurements. The aircraft used was a Bergen Quad 8, which is shown ready for flight in Figure 47. This "octocopter" has eight propellers in coaxial pairs which aid in stability and in increasing the maximum payload capacity. It is manufactured in Michigan and is capable of flying with an 8kg payload for up to 20 minutes per charge. The four arms also fold for easier transport. MTRI has successfully retrofitted similar aircraft for high-resolution 3D mapping and thermal assessments. The aircraft was outfitted with a DJI

A3 redundant avionics system which provides for continued flight in the face of GPS failure, as well as increased flight stability and path accuracy. Without this system, maintaining a level, straight flightpath (either manually or on autopilot) is challenging, especially in turbulence near the ground. By increasing the GPS accuracy and update rate, this is much less of an issue.



Figure 52 Bergen Quad 8 with radar and array antennas attached ready for flight

Accurate measurement of the radar position during data collection is required to form high quality images of the buried targets. To further increase accuracy, a Blade RTK GPS/IMU unit was used to monitor the sensor position during the flight. Data from this unit was post-processed to provide relative position estimates with differential GPS level accuracy.

A four port, wideband, low frequency stepped frequency continuous wave (SFCW) radar was assembled using an AKELA VNA and was mounted on the Bergen X8 to collect downward looking ground penetrating radar measurements. During the data collection, 256 frequencies between 500MHz and 2000MHz were transmitted with a notch-filter at 1575MHz (100MHz 3dB width) to suppress frequencies in the GPS band (to avoid self-jamming of the navigation system). The frequency hopping rate was 15300 Hz, which resulted in a sweep repetition frequency of 59.8 Hz. A summary of the radar system is provided in Table 3.

Table 3. Radar System Parameters


System Parameters	Specification
Transmit Frequency	500 to 2000 MHz
Step Frequency	15300 Hz
Number of Samples Per Sweep	256
Weight	1 kg
Size	193 x 107 x 39 mm
Transmit Power	50 mW

The radar was connected to a four-antenna array mounted under the sUAS. Connecting the multiport radar to an antenna array allows each antenna to serve as either a transmitter or receiver. Thus, the measurement system has multiple phase centers that provide additional measurement diversity to produce spatially resolved images of the ground as the sUAS advances. In the nominal configuration for a data collection, one antenna in the array was used as a transmitter and the three other antennas were used as receivers. This produced data with three effective phase centers transverse to the sUAS direction of movement. Two antenna arrays have been used for measurements; a linear array of four Vivaldi antennas and an array of four offset wideband antenna elements developed by AKELA. The sUAS with radar antenna arrays and an image showing the sUAS flying during measurement are shown in Figure 48. For our initial measurements, the radar was controlled by a laptop computer via an Ethernet tether.

sUAS Flying Above Beach



Four Element Vivaldi Array



AKELA Array



Figure 53 The sUAS with four Vivaldi element array flying over beach measurement location (top left), and the Vivaldi array is shown attached to radar (bottom left) along with the Akela antenna array (bottom right).

6.2 Field Data Collection

To evaluate the ability of the system to detect and image buried targets, ground penetrating radar measurements of buried targets were collected at a sandy private beach in South Haven, MI on Lake Michigan. Measurements of the beach were collected with and without the targets present. The sandy soil had a maximum volumetric water content measured as 1-2% to a depth of 1m. Flight conditions were good; light and variable winds from the Northwest, gusting to 8kts, with partly cloudy skies. A view of the measurement area before the data collections is shown in Figure 49.



Figure 54 Overhead view of the beach measurement area before the data collection.

Six targets were chosen with diverse dielectrics, sizes, and geometries for use in this study. Short descriptions, sizes, and burial depths of all these targets measured are listed in Table 4. We note that all targets were buried at depths above the underlying water table.

Table 4. Buried Targets Used For Data Collection

TARGET	SIZE	DEPTH
Disk, foam w/ metal face	30cm	5cm
Disk, foam w/ metal face	30cm	15cm
Disk, metal	12.5cm diameter, 5cm thick	5cm
Jug	1 gal, 14cm x 14cm x 26cm	5cm - 31cm
Wooden Beam	4" x 4" x 36"	5cm
Leaf Rake	triangular, 22" on a side	0cm

The positions of the targets are indicated as red points in in Figure 50. They were spaced between 2m and 3m apart with the final target about 13m from the take-off point (marked by an s). Flights were made at an average altitude of 1.5m AGL. A typical flight is path is shown as the blue line in Figure 50. Due to limitations in the flight-control software,

automatic flight was not possible at this low altitude; all flights were controlled manually, which resulted in some irregularities in the path (correcting for wind drift, etc.).

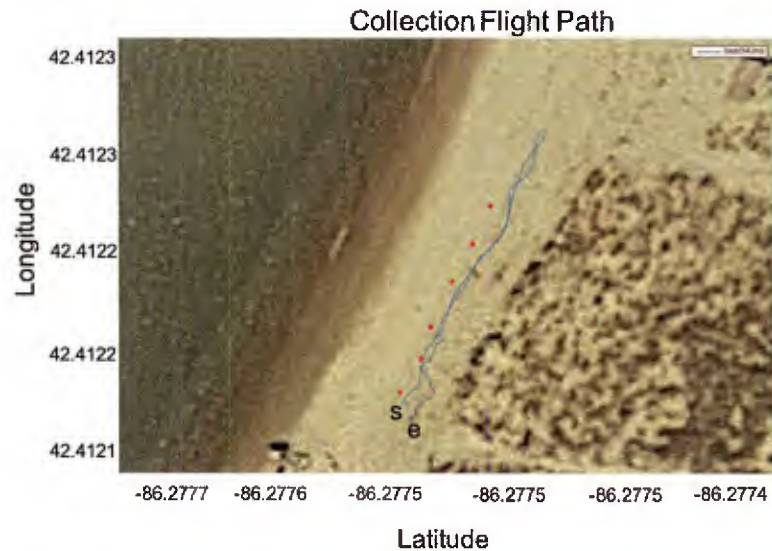


Figure 55 Flight path of a typical data collection (blue line). The positions of targets are indicated by the red dots.

A data collection consisted of taking off from a point just south of the first target (point s in Figure 50), flying north at nominally 1 m/s until past the last target, then returning along the same path, landing near the original take-off point (point e in Figure 50). Several flights were made from the take-off point, flying west until over the water, then returning. The data collection flights are listed in Table 5.

Table 5. Data Collection Flights

Data File Name	Time (EDT)	Antenna Array	Description
beach1.imb	11:38	Akela	No targets present
beach2.imb	11:41	Akela	No targets present
beach3.imb	12:11	Akela	All targets
beach4.imb	12:25	Akela	All targets
beach5.imb	12:32	Akela	To the water
beach6.imb	12:53	Vivaldi	Range test (varying altitudes)

Data File Name	Time (EDT)	Antenna Array	Description
beach7.imb	12:55	Vivaldi	All targets
beach8.imb	12:57	Vivaldi	All targets
beach9.imb	13:03	Vivaldi	No targets present
beach10.imb	13:05	Vivaldi	No targets present
beach11.img	13:07	Vivaldi	To the water

6.3 Data Processing

The collected data consists of measurements of the radar backscattering as a function of frequency at specific measurement positions for each Tx-Rx pair of antennas. The data were detrended by subtracting the mean over measurement positions at each frequency to remove constant returns, such as the direct feed through from the transmitter to receiver. A complete description of the data processing pipeline is shown in Figure 51.




Figure 56 Data processing pipeline

The resulting data were processed into three-dimensional (3D) radar images using standard backprojection. The 3D images were examined by looking at contour plots of the 3D images, or by integrating the energy over height at each transverse spatial position to produce a 2D map of integrated backscattered power. In each case, the resulting displays were examined to look for evidence of the buried target. A complete imaging pipeline description is shown in Figure 52.

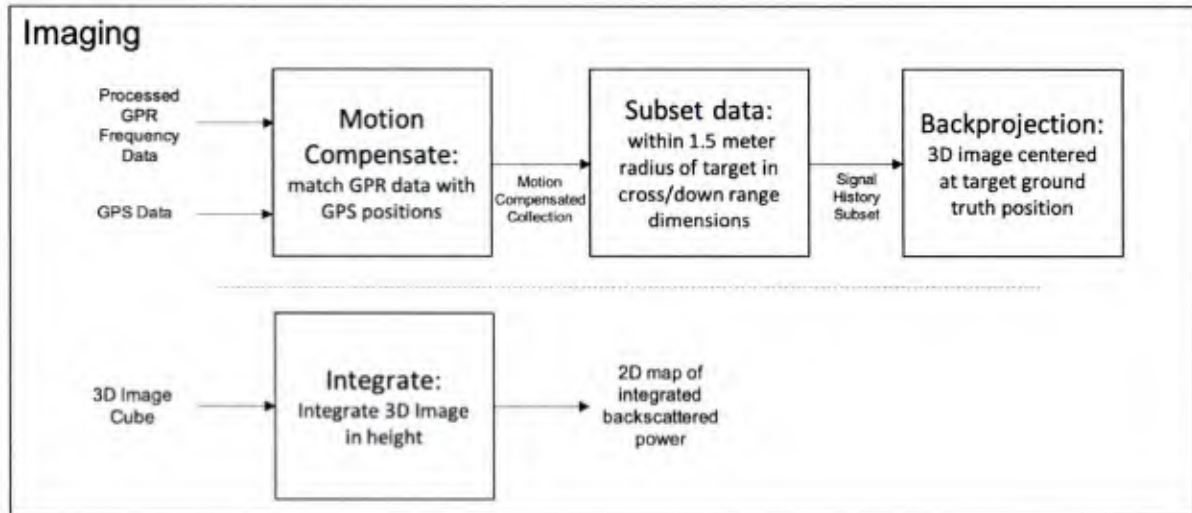


Figure 57 Data processing pipeline

6.4 Imaging Results

We discuss results from run 7 collected with the Vivaldi Antenna and with all targets present. Figures 53 and 54 show the 3D image results. In Figure 53, The image on the left of Figure 53 shows the locations of targets. On the right of Figure 53, integrated images of the 3D images are shown. In this image, the 30 cm diameter disks can be easily distinguished. The wood and jug of water are not as easily distinguished. In Figure 54, contour plots of the 3D images generated about each target are shown. The 30 cm disks are localized in ground truth positions, while the wood and jug of water are not as distinguishable. In general, the targets can be distinguished and detected, although future work will examine aspects of the data collection and processing to further improve the reliability of the technique.





Figure 58 (left) Location of targets. (Right) 2D maps of the 3D image energy integrated over height are shown for regions around each target.

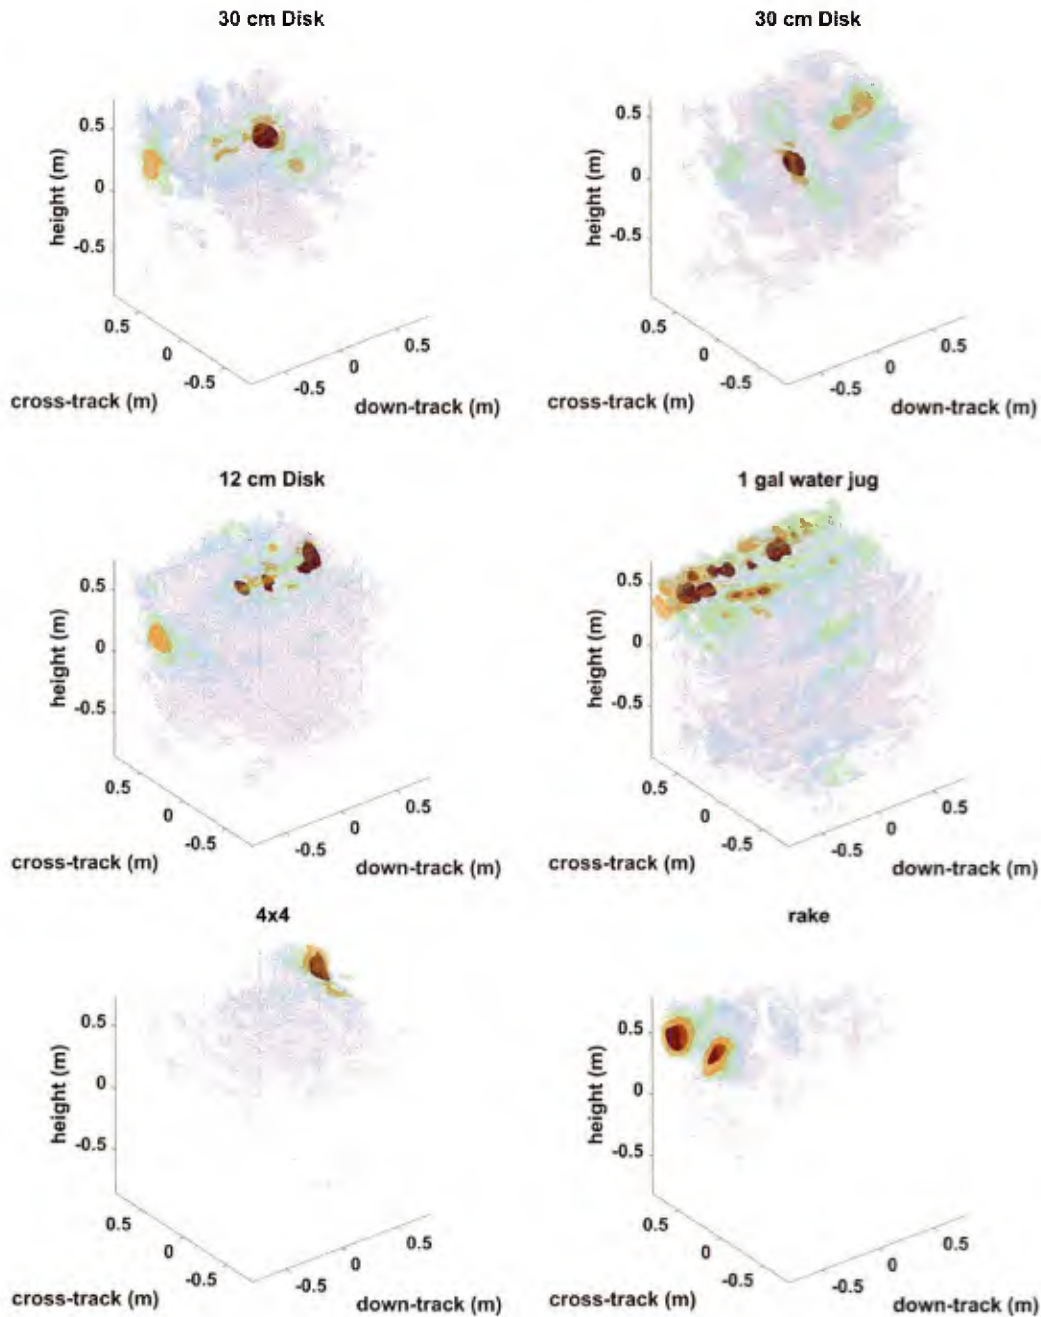


Figure 59 Contour plots of the 3D images generated about each target are shown. Coordinates relative to ground-truth positions.

The results presented show that useable ground penetrating radar data can be collected from a sUAS platform and that the resulting data can be used to detect and image buried targets. Future work will look at implementing different imaging approaches that account for ground-plane diffraction, incorporating additional motion compensation, and making improvements

to the sUAS platform through the incorporation of a sonar/laser range finder to track elevation.

References

Berisha, V., Wisler, A., Hero, A. O., & Spanias, A. (2016). Empirically estimable classification bounds based on a nonparametric divergence measure. *IEEE Transactions on Signal Processing*, 64(3), 580-591.

Borcea L., G. Papanicolaou, and C. Tsogka (2005), Interferometric array imaging in clutter, *Inverse Probl.*, vol. 21, pp. 1419–1460

Borcea, L., Papanicolaou, G., and Tsogka, C. (2006), Adaptive interferometric imaging in clutter and optimal illumination, *Inverse problems*, vol. 22(4), 1405–1436.

Borcea L., J. Garnier, G. Papanicolaou, and C. Tsogka (2011), Coherent interferometric imaging, time gating and beamforming. *Inverse Problems*, Vol. 27, Issue 6.

Burns J., M. Masarik, B. Thelen, J. Kelly, R. Lipa, R. LaRose, A. Pinar A. and T. Havens (2016), *Advanced Signal Processing and Detection Algorithms for Handheld Explosive Hazard Detection*, Final Technical Report, Contract #: W909MY-13-C-0029

Byrd R., P. Lu and J. Nocedal. A Limited Memory Algorithm for Bound Constrained Optimization, (1995), *SIAM Journal on Scientific and Statistical Computing* , 16, 5, pp. 1190-1208.

Chew, W.C. (1990). *Waves and Fields in Inhomogeneous Media*. Van Nostrand Reinhold

Colwell, K. A., & Collins, L. M. (2016, May). Attribute-driven transfer learning for detecting novel buried threats with ground-penetrating radar. In *SPIE Defense+ Security* (pp. 982319-982319). International Society for Optics and Photonics

Counts, T., Gurbuz, A. C., Scott, W. R., McClellan, J. H., & Kim, K. (2007). Multistatic ground-penetrating radar experiments. *IEEE Transactions on Geoscience and Remote Sensing*, 45(8), 2544–2553. <https://doi.org/10.1109/TGRS.2007.900677>

Craig Warren, Antonios Giannopoulos, and Iraklis Giannakis. gprMax: Open source software to simulate electromagnetic wave propagation for ground penetrating radar. *Computer Physics Communications*, 2016.

Dahl, G. E., Sainath, T. N., & Hinton, G. E. (2013, May). Improving deep neural networks for LVCSR using rectified linear units and dropout. In *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on* (pp. 8609-8613). IEEE.

Demanet, Laurent, and Vincent Jugnon (2017), Convex recovery from interferometric measurements. *IEEE Transactions on Computational Imaging* 3.2: 282-295.

Giannopoulos, A., (2005) GPRMAX 2D/3D, User's Manual, Version 2.0

Gorham, L. and L. Moore (2010), SAR image formation toolbox for MATLAB, Proc. SPIE 7699, Algorithms for Synthetic Aperture Radar Imagery XVII, 769906 (18 April 2010); doi: 10.1117/12.855375; <https://doi.org/10.1117/12.855375>

Gal, Y., & Ghahramani, Z. (2015). Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. arXiv preprint arXiv:1506.02142, 2.

Greengard, L. and Lee, J.-Y. (2004), Accelerating the nonuniform fast fourier transform, *SIAM Review*, vol. 46, 443– 454 (2004).

Hansen, T. B., & Johansen, P. M. (2000). Inversion scheme for ground penetrating radar that takes into account the planar air-soil interface. *IEEE Transactions on Geoscience and Remote Sensing*, 38(1 II), 496–506. <https://doi.org/10.1109/36.823944>

He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 770-778).

Li, J., Wang, X., & Wang, T. (2010). On the Validity of Born Approximation. *Progress In Electromagnetics Research*, 107(July), 219–237. <https://doi.org/10.2528/PIER10070504>

Liu, D. C. and Nocedal, J. (1989), On the limited memory bfgs method for large scale optimization, *Mathematical Programming*, vol. 45 (1-3), 503–528.

Lo Monte, L., Soldovieri, F., Erricolo, D., Giannopoulos, A., & Wicks, M. C. (2012). A Comprehensive Forward Model for Imaging under Irregular Terrain Using RF Tomography. *International Journal of Antennas and Propagation*, 2012, 1–15.

<https://doi.org/10.1155/2012/735414>

Long, J., Shelhamer, E., & Darrell, T. (2015). Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 3431-3440).

Masarik, M. P., Burns, J., Thelen, B. T., and Sutter, L. (2015), Fast 3D subsurface imaging with stepped-frequency GPR, in [*Detection and Sensing of Mines, Explosive Objects, and Obscured Targets XX*], Proc. SPIE 9454, 945417.

Masarik M., J. Burns; B. Thelen; J. Kelly; T. Havens (2016), Enhanced buried UXO detection via GPR/EMI data fusion, Proc. SPIE 9823, *Detection and Sensing of Mines, Explosive Objects, and Obscured Targets XXI*, 98230R (3 May 2016);

doi: 10.1117/12.2223009

Mason E., I. Son and B. Yazıcı (2015), Passive Synthetic Aperture Radar Imaging Using Low-Rank Matrix Recovery Methods, in *IEEE Journal of Selected Topics in Signal Processing*, vol. 9, no. 8, pp. 1570-1582.

Persico, R., Bernini, R., & Soldovieri, F. (2005). The role of the measurement configuration in inverse scattering from buried objects under the born approximation. *IEEE Transactions on Antennas and Propagation*, 53(6), 1875–1887. <https://doi.org/10.1109/TAP.2005.848468>

Persico, R., Negri, S., Soldovieri, F., & Pettinelli, E. (2012). Pseudo 3D imaging of dielectric and magnetic anomalies from GPR data. *International Journal of Geophysics*, 2012.

<https://doi.org/10.1155/2012/512789>

Picco, V., Gennarelli, G., Negishi, T., Soldovieri, F., & Erricolo, D. (2015). Experimental Validation of the Quadratic Forward Model for RF Tomography. *IEEE Geoscience and Remote Sensing Letters*, 12(7), 1461–1465. <https://doi.org/10.1109/LGRS.2015.2409025>

Rosenblueth A. and N. Wiener (1945). The role of models in science. *Philosophy of science*, 12(4):316–321.

Sakaguchi, R. T., Morton, K. D., Collins, L. M., & Torriano, P. A. (2015). Recognizing subsurface target responses in ground penetrating radar data using convolutional neural networks, 9454, 94541A. <https://doi.org/10.1117/12.2177747>

Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. (2014). Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1), 1929–1958.

van den Berg, E. and M. P. Friedlander. SPGL1: A solver for large-scale sparse reconstruction. [<http://www.cs.ubc.ca/labs/scl/spgl1>]. June 2007.

van den Berg, E. and M. P. Friedlander. Probing the Pareto frontier for basis pursuit solutions. *SIAM Journal on Scientific Computing* Vol. 31 No. 2 pp 890-912. 2008.

Wallace, B. C., Small, K., Brodley, C. E., & Trikalinos, T. A. (2011, December). Class imbalance, redux. In *Data Mining (ICDM), 2011 IEEE 11th International Conference on* (pp. 754-763). IEEE.

Webb A., I. Xique, B. Thelen, and J. Burns (2017), Forward Looking Ground Penetrating Radar (FLGPR) Processing, Final Technical Report, CACI PO# P000028947

Xique I., J. Burns; B. Thelen and R. LaRose (2018), Novel application of windowed beamforming function imaging for FLGPR, *ProcSPIE 10628, Detection and Sensing of Mines, Explosive Objects, and Obscured Targets XXIII*, 106280H (30 April 2018); doi: 10.1117/12.2303724

Xu, B., Wang, N., Chen, T., & Li, M. (2015). Empirical evaluation of rectified activations in convolutional network. arXiv preprint arXiv:1505.00853.

Zhang, C., Bengio, S., Hardt, M., Recht, B., & Vinyals, O. (2016). Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*.

Zhang, Y., Huston, D., & Xia, T. (2016). Underground object characterization based on neural networks for ground penetrating radar data, 9804(April), 980403.

<https://doi.org/10.1117/12.2219345>

Appendix A. Conference Papers

The following two papers were presented at the 2019 SPIE Defense and Security Conference in Baltimore, MD.

Buried Object Imaging Using a Small UAS-based GPR

Christopher Roussi, Ismael Nique, Joseph Burns, Benjamin Hart

Michigan Tech Research Institute, Michigan Technological University, Ann Arbor, MI USA

ABSTRACT

Ground-based systems for imaging objects at standoff distances usually interrogate the ground region of interest at shallow grazing angles. Thus, only a small amount of the transmitted radar energy propagates into the ground, and an even smaller portion of the target backscattered energy is returned to the sensor, thus limiting detection performance. Mounting a small size, weight and power radar on a small unmanned aerial system (sUAS) platform provides a method of interrogating the ground at more favorable angles of incidence with the radar operator being at a safe standoff distance. Reliable detection and/or imaging of a buried object from a sUAS is influenced by several issues related to sensor power, measurement of sensor position and control of the sUAS in the proximity to the ground. We performed a series of measurements of buried objects using a small radar on a tethered sUAS to explore these issues. The experiments and issues encountered will be discussed, and three-dimensional radar imagery of buried objects generated from the measurements will be presented.

Keywords: Ground Penetrating Radar, Imaging, UAS, Buried Object

INTRODUCTION

The problem of generating images of subsurface objects using ground penetrating radar (GPR) is of general interest to many groups but of particular interest to military and humanitarian organizations looking to safely conduct mine and explosive hazard clearance. Demining efforts can be dangerous because they can require operators to be within proximity of the explosive hazard to accurately detect and identify it. Over the years a variety of GPR platforms have been proposed such as forward looking GPR (FLGPR), but these systems have had to operate at a grazing angle of incidence to provide operator standoff at the expense of reduced detection performance compared to a downward looking system¹⁻⁵. In this paper, we investigate the use of a ground penetrating radar on a small unmanned aerial system (sUAS) to detect buried targets. A radar mounted on a sUAS can interrogate the ground at near normal angles of incidence, thus providing better propagation of energy into the ground, yet still be at a standoff distance from friendly dismount operators. In this study, we collected measurements of buried targets along a Lake Michigan beach using a low size, weight, and power (SWAP) radar on one of MTRI's sUAS platforms to demonstrate buried object imaging and provide data for algorithm development. A four-element antenna array was flown on a sUAS to provide additional spatial diversity for better spatial resolution. The results show that useable ground penetrating radar data can be collected from a sUAS platform and that the resulting data can be used to detect and image buried targets.

This paper is organized as follows. In section 2, we describe the UAS platform and GPR system used in this study. We give specifics such as the GPR system parameters, antenna configurations and schematics, and UAS platform capabilities. In section 3, we give an overview of the field data collect. We specify the buried targets used, details on the test site, and an overview of the data collections. In section 4, we assess the recovered images of buried targets from the collected GPR data. We also detail the preprocessing steps and imaging steps in this section. Lastly, in section 5, we provide a summary of results and detail steps that can be taken to further improve imaging when using a radar mounted on a sUAS.

UAV PLATFORM AND GPR SYSTEM DESCRIPTION

Recent advances in remote controlled unmanned aerial vehicles have made them much cheaper and capable. These vehicles have been used by MTRI and others in conjunction with lightweight and advanced sensors, such as optical, thermal, and LIDAR sensors, to create effective systems for the collection of remote sensing data. Recently, small UAS systems have been applied successfully to tasks such as synthetic aperture radar (SAR) imaging⁶. In this paper, we adopt a similar multi-rotor platform equipped with a GPR system for the measurement of buried targets. The components of the sUAS, which is shown ready for flight in Figure 1, are detailed in this section. We discuss the Bergen Quad 8 sUAS

platform, the DJI A3 redundant avionics system, the Blade Real-Time Kinematic (RTK) GPS/IMU, and the ground penetrating radar system based on an AKELA vector network analyzer (VNA).

1.1 UAS Platform

The aircraft used was a Bergen Quad 8. This "octocopter" has eight propellers in coaxial pairs which aid in stability and in increasing the maximum payload capacity. It is manufactured in Michigan and is capable of flying with an 8kg payload for up to 20 minutes per charge. The four arms also fold for easier transport. MTRI has successfully retrofitted similar aircraft for high-resolution 3D mapping and thermal assessments. The aircraft was outfitted with a DJI A3 redundant avionics system which provides for continued flight in the face of GPS failure, as well as increased flight stability and path accuracy. Without this system, maintaining a level, straight flightpath (either manually or on autopilot) is challenging, especially in turbulence near the ground. By increasing the GPS accuracy and update rate, this is much less of an issue.



Figure 1. Bergen Quad 8 with radar and array antennas attached ready for flight.

Accurate measurement of the radar position during data collection is required to form high quality images of the buried targets. To further increase accuracy, a Blade RTK GPS/IMU unit was used to monitor the sensor position during the flight. Data from this unit was post-processed to provide relative position estimates with differential GPS level accuracy.

1.2 GPR Sensor and Antenna Configuration

A four port, wideband, low frequency stepped frequency continuous wave (SFCW) radar was assembled using an AKELA VNA and was mounted on the Bergen X8 to collect downward looking ground penetrating radar measurements. During the data collections, 256 frequencies between 500MHz and 2000MHz were transmitted with a notch-filter at 1575MHz (100MHz 3dB width) to suppress frequencies in the GPS band (to avoid self-jamming of the navigation system). The frequency hopping rate was 15300 Hz, which resulted in a sweep repetition frequency of 59.8 Hz. A summary of the radar system is provided in Table 1.

Table 1. Radar System Parameters.

System Parameters	Specification
Transmit Frequency	500 to 2000 MHz
Step Frequency	15300 Hz
Number of Samples Per Sweep	256
Weight	1 kg
Size	193 x 107 x 39 mm
Transmit Power	50 mW

The RF sensing unit was connected to a four-antenna array mounted under the sUAS. Connecting the multiport unit to an antenna array allows each antenna to serve as either a transmitter or receiver. Thus, the measurement system has multiple phase centers that provide additional measurement diversity to produce spatially resolved images of the ground as the sUAS advances. For initial measurements, the radar was controlled by a laptop computer via an Ethernet tether. This configuration simplified equipment mounting, and provided longer flight times. It is possible to mount a small, dedicated control system on the aircraft if needed.

Two antenna arrays were available for use during collection but we discuss the linear array of four Vivaldi antennas (Figure 2) in this paper. In the nominal configuration for a data collection, one antenna in the array was used as a transmitter and the three other antennas were used as receivers. This produced data with three effective phase centers transverse to the sUAS direction of movement. Note, that in principle, each antenna could be used as a transmitter, with the other three antennas as receivers, to produce a total of twelve phase centers.

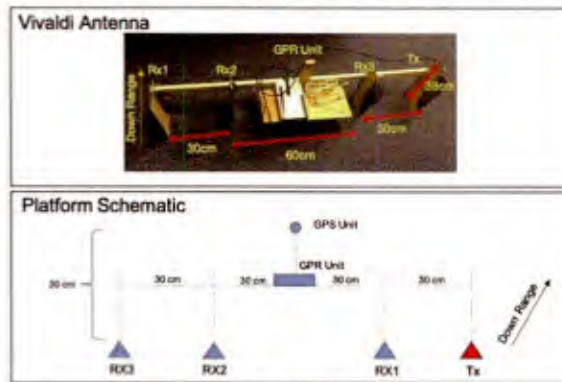


Figure 2. Vivaldi antenna array and platform schematic.

FIELD DATA COLLECTION

1.3 Test Site Description

To evaluate the ability of the system to detect and image buried targets, ground penetrating measurements of buried targets were collected at a sandy private beach in South Haven, MI on the shore of Lake Michigan. The lake water is west of the beach. Measurements of the beach were collected with and without the targets present. A view of the measurement area before the data collections is shown in Figure 3. The sandy soil had a maximum volumetric water content measured as 1-2% to a depth of 1m. Flight conditions were good during the data collection with light and variable winds, from the Northwest, that were gusting to 8kts and partly cloudy skies.



Figure 3. Overhead view of the beach measurement area before the data collection.

1.4 Target Description

Six targets were chosen with diverse dielectrics, sizes, and geometries for use in this study. Short descriptions, sizes, and burial depths of all these targets measured are listed in Table 2. We note that all targets were buried at depths above the underlying water table.

Table 2. Buried Targets Used for Data Collection

Target	Size	Depth
Disk, foam w/ metal face	30cm	5 cm
Disk, foam w/ metal face	30cm	15 cm
Disk, metal	12.5cm diameter, 5cm thick	5 cm
Jug	1gal, 14cm x 14cm x 26cm	5 cm – 31 cm
Wooden Beam	4" x 4" x 36"	5 cm
Leaf Rake	triangular, 22" on a side	0 cm

The positions of the targets are indicated in the image in Figure 4 as red points. They were spaced between 2 meters and 3 meters apart. The final target was 13 meters from the take-off point (marked by S). Flights were made at an average altitude of 1.5 meters above ground level (AGL). A typical flight path over the targets is shown as a blue line in Figure 4. Due to limitations in the DJI flight-control software, automatic flight was not possible at this low altitude; all flights were controlled manually, which resulted in some irregularities in the path (correcting for wind drift, etc.).

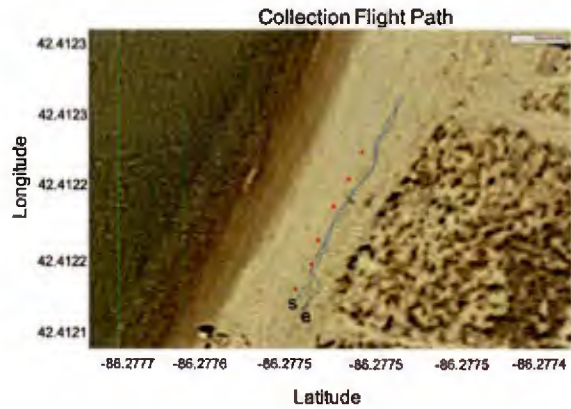


Figure 4. Flight path of a typical data collection (blue) line. The positions of targets are indicated by the red dots.

1.5 Data Collection Description

A data collection consisted of taking off from a point just south of the first target (point s in Figure 4), flying north at nominally 1 m/s until past the last target, then returning along the same path, landing near the original take-off point (point e in Figure 4). Note that the sUAS flies over the targets in reverse on the return path. In this paper, we process and discuss results from one of multiple data collections.

GPS data is collected at 10 Hz as the sUAS platform flies over the targets. It is post-processed to provide differential GPS level accuracy. Figure 5, shows the positions of the effective phase centers along the flight path relative to the targets.




Figure 5. GPS positions compensated to account for platform geometry.

Figure 6 shows the magnitude and phase components of the frequency domain data collected using Vivaldi antennas. Note that the Channel 3 receive antenna is physically closest to the transmit antenna and thus has a visibly stronger direct response. Additionally, a visible drop off at 1.575 GHz was due to the GPS notch filter.

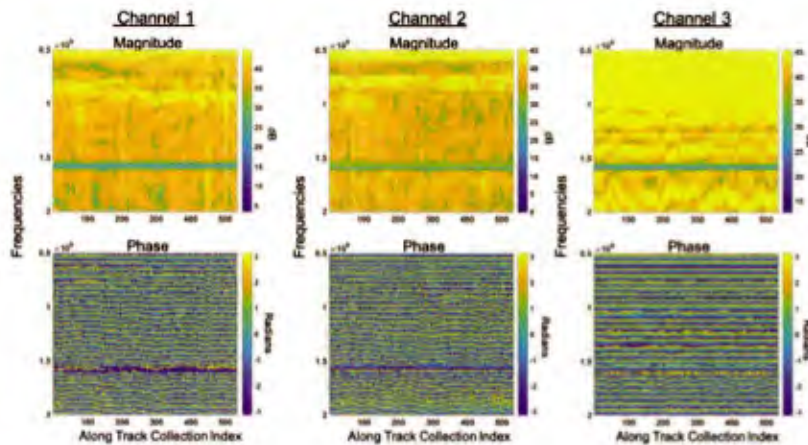


Figure 6. Magnitude (top row) and Phase (bottom row) components of a typical frequency domain data collection using the Vivaldi antenna array. Data from a fixed transmitter/receiver (channel) pair are shown per column.

BURIED OBJECT IMAGING

The collected data consists of measurements of the radar backscattering as a function of frequency at specific measurement positions for each Tx-Rx pair of antennas. The data undergoes some preprocessing which we will discuss in the next section and is then synthesized to recover a three-dimensional (3D) representation of the buried target via the backprojection algorithm as well as a 2D map of integrated backscattered power. 2D and 3D representations of buried targets can be used potentially for detection and localization of buried objects and to provide features for object classification.

1.6 Data Processing

Each measurement taken by the radar system underwent three step processing before imaging. The raw data of a single collection consisted of a set of three-channel frequency sweeps with each sweep associated with a unique spatial position (x, y, z). We calibrated the data to account for antenna frequency response. The calibration data consisted of the response of a reflective plate. Additionally, the collected data were phase shifted to remove a system cable delay of about 2.15 (m). Lastly, the data were detrended by subtracting the mean over all measurement positions in a fixed channel at each frequency to remove constant returns, such as the direct feedthrough from the transmitter to receiver. The processing pipeline is summarized in Figure 7.




Figure 7. Data processing pipeline.

To make a 3D image of a buried target, the processed lane data underwent a subsequent three step process. The data points were first motion compensated to match the GPR data to the GPS positions. Then, using the target ground truth coordinates, we extracted from the data, collection points within a 1.5 meter radius of the selected target. This data was then backprojected⁷ into a 3D image with the following dimensions: 1.2 (m) x 1.2 (m) x 1.2 (m) and 14 cm resolution. A

potential extension of our current processing scheme would be to account for ground-plane diffraction as was done in Masarik's fast 3D imaging approach.⁸ The resulting 3D image, from backprojection, was centered around the ground-truth position of the target. Additionally, we observed the variation in energy over targets along the whole lane by integrating the energy of each 3D image in the height dimension. Figure 8 summarizes the processing pipeline.

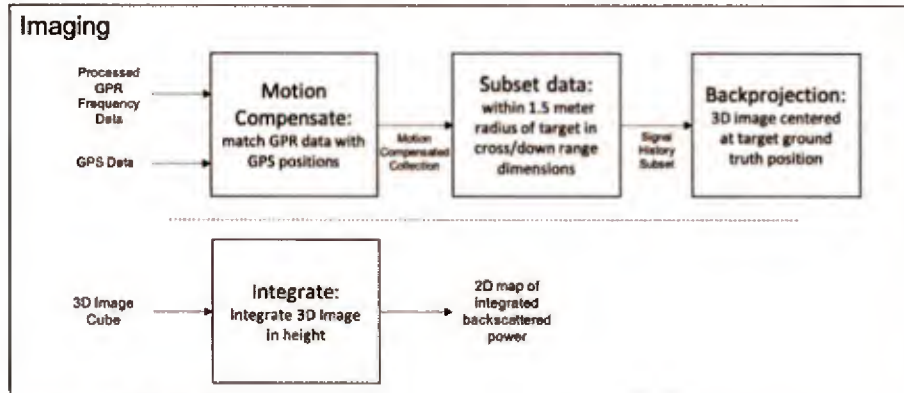


Figure 8. Data imaging pipeline.

1.7 Imaging Results

The 3D images were examined by integrating the energy over height at each transverse spatial position to produce a 2D map of integrated backscattered power and by looking at contour plots of the 3D images. In each case, the resulting displays were examined to look for evidence of the buried target.




Figure 9. (Left) Location of targets. (Right) 2D maps of the 3D image energy integrated over height are shown for regions around each target

Figures 9 and 10 show the imaging results. In Figure 9, 2D maps of the image energy integrated over height are shown for regions around each target. The diagram in the left of Figure 9 shows the locations of targets. The image in the right of Figure 9 shows that the 30 cm diameter disks can be distinguished easily. They have a strong and localized returns in

ground truth positions. The rake is also quite visible which is expected given its size and burial depth. The wood and jug of water are not as distinguishable, but it is possible that the returns from the rake are also saturating those of the weaker targets. The small 12 cm disk is also not as distinguishable but we hypothesize that further processing that incorporates ground-plane diffraction and soil propagation speeds could help in further localizing and strengthen its return. In Figure 10, contour plots of the 3D images generated about each target are shown. The 30 cm disks are localized in ground truth positions, while the wood and jug of water are not as distinguishable. In general, the targets can be detected, although future work will examine aspects of the data collection and processing to further improve the reliability of the technique.

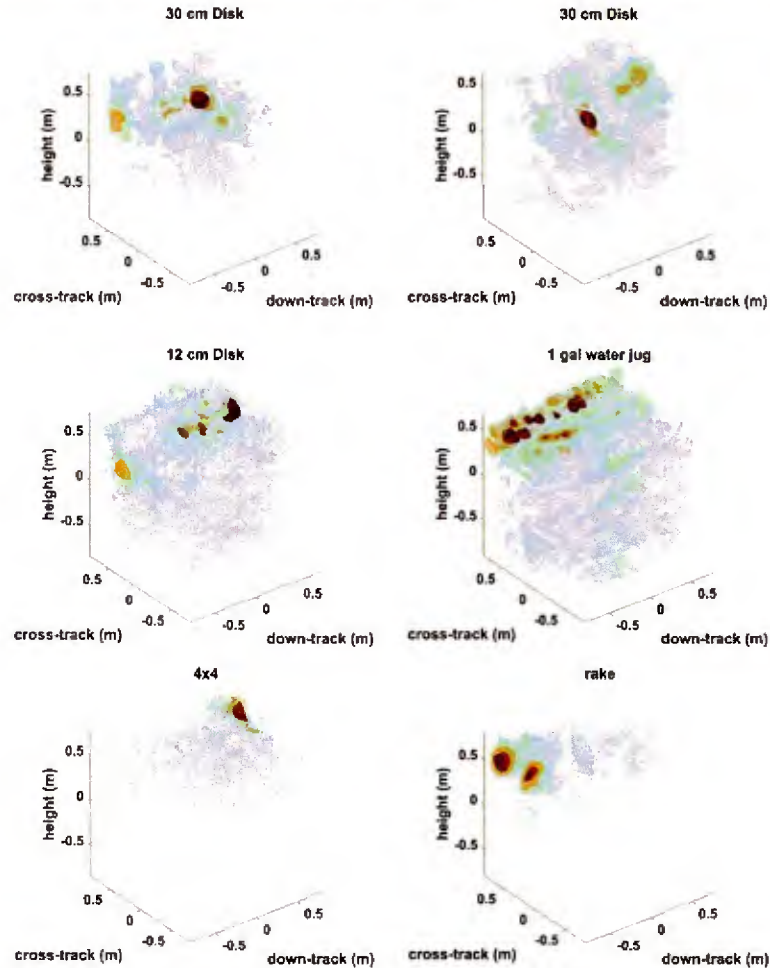


Figure 10. Contour plots of the 3D images generated about each target. Coordinates relative to target ground truth position.

CONCLUSIONS

We investigated the use of a ground penetrating radar on a small unmanned autonomous systems (sUAS) to image buried targets. A radar mounted on a sUAS can interrogate the ground at near normal angles of incidence, thus providing better propagation of energy into the ground, yet still be at a standoff distance from friendly dismount operators. Four element antenna arrays were flown on an sUAS to provide additional spatial diversity for better spatial resolution. We collected measurements of buried targets along a Lake Michigan beach using a low SWAP radar on one of MTRI's sUAS platforms to demonstrate buried object imaging and detection and provide data for algorithm development. The results presented in Section 4, show that useable ground penetrating radar data can be collected from a sUAS platform and that the resulting data can be used to detect and image buried targets. Future work will look at implementing different imaging approaches that account for ground-plane diffraction, incorporating additional motion compensation, and making improvements to the sUAS platform through the incorporation of a sonar/laser range finder to track elevation.

ACKNOWLEDGEMENTS

This work was funded by the Office of Naval Research under Grant N00014-16-1-2623.

REFERENCES

- [1] Burns, J., Masarik, M. P., Nique, I. J., Thelen, B. and Webb, A., "Comparative analysis of image formation techniques for FLGPR," *Proc.SPIE* **10182** (2017).
- [2] Soldovieri, F., Gennarelli, G., Catapano, I., Liao, D. and Dogaru, T., "Forward-Looking Radar Imaging: A Comparison of Two Data Processing Strategies," *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* **10**(2), 562–571 (2017).
- [3] Kositsky, J., "Forward-looking high-resolution GPR system," *Proc. SPIE* **3710**(Part 1&2), 1052–1062 (1999).
- [4] Kositsky, J., Cosgrove, R., Amazeen, C. and Milanfar, P., "Results from a forward-looking GPR mine detection system," *Proc. SPIE - Int. Soc. Opt. Eng.* **4742**, 206–217 (2002).
- [5] Catapano, I., Affinito, A., Del Moro, A., Alli, G. and Soldovieri, F., "Forward-Looking Ground-Penetrating Radar via a Linear Inverse Scattering Approach," *IEEE Trans. Geosci. Remote Sens.* **53**(10), 5624–5633 (2015).
- [6] Lort, M., Aguasca, A. and Carlos, L., "Initial Evaluation of SAR Capabilities in UAV Multicopter Platforms," 127–140 (2018).
- [7] Gorham, L. A. and Moore, L. J., "SAR image formation toolbox for MATLAB," *Proc.SPIE* **7699** (2010).
- [8] Masarik, M. P., Burns, J., Thelen, B. T. and Sutter, L., "Fast 3D subsurface imaging with stepped-frequency GPR," *Proc.SPIE* **9454** (2015).

Quadratic lifting inversion applied to buried object radar imaging.

Joseph B. Lindgren^a, Joseph Burns^a, Brian Thelen^a, and Ismael Xique^a

^aMichigan Tech Research Institute, 3600 Green Ct., Ste. 100, Ann Arbor, MI 48105, U.S.A.;

ABSTRACT

Since radar imaging of buried objects involves propagation through media that are at best partially known, there is mismatch between the forward model used in the inversion and the propagation behavior actually observed in the measured data. The mismatch can cause degradation and/or reduce resolution in the imagery, which limit automatic target recognition features that can be extracted from the imagery. Recently, several research groups have advocated backpropagation of interferometric measurements as a more statistically stable estimator of targets in the presence of forward model errors and in the presence of clutter. Specifically, the lifting approach to inverse problems [Demanet and Jugnon, 2017]¹ has been proposed as a robust approach to inversion in the presence of forward model mismatch that can produce reconstructions with fidelity comparable to direct inversion with the matched model. We apply this technique to radar imaging of buried targets to determine if it can produce enhanced imagery in the presence of limited knowledge of the surrounding ground geometry and/or material properties. In this paper we describe the algorithm implementation and present results for both simulated and measured data. The results show that the approach has significant potential for enhancing images of buried objects from scenarios with realistic forward model mismatch. However, we have observed significant sensitivity to surrounding clutter and to the choice of regularization. Mitigating these sensitivities is a topic of ongoing research.

Keywords: interferometry, lifting, non-linear optimization, ground penetrating radar, model mismatch

1. INTRODUCTION

Since radar imaging of buried objects involves propagation through media that are at best partially known, there is mismatch between the forward model used in the inversion and the propagation behavior actually observed in the measured data. The mismatch can cause degradation and/or reduce resolution in the imagery, which limit automatic target recognition features that can be extracted from the imagery. Instead of attempting to construct detailed physical models with more parameters, interferometric methods focus on the differences between the spatially/spectrally local cross-correlations of the forward model and the spatially/spectrally local cross-correlations of the data as the basis for inversion. However, interferometric techniques such as Coherent INTERferometric (CINT)² or Windowed Beamforming Energy (WBE) trade image resolution for this model robustness and statistical stability.³ Recently, quadratic lifting inversion (QLI) has been proposed as an alternative imaging problem formulation to surmount this limitation. On the basis of theory and simulations,¹ Demanet and Jugnon claim that QLI has to potential to accomplish all three of the following goals:

- robustness to model mismatch,
- comparable resolution with classical techniques (backprojection, least squares auto-focus, etc.),
- and statistically spatially stable recovery (with respect to randomly heterogeneous media).

Further author information: (Send correspondence to J. L.)

J. L.: E-mail: jblindgr@mtu.edu, Telephone: 1 734 994 7231

In the present paper we evaluate these claims on measured data. Section 2 introduces our implementation of QLI. Section 3 is divided into three subsections, each of which details an experiment which evaluates one of the above claims made about QLI. In the first (Subsection 3.1), we show by simulation that QLI is robust to model mismatch. In the second (Subsection 3.2), we show that QLI can produce imagery in a controlled experiment, with measured data, that has resolution comparable to backprojection. Subsection 3.3 completes the survey by testing the stability of the method against data collected in an experimental but randomly heterogeneous (buried target) setting. Section 4 summarizes our findings and comments on future directions for research.

2. PROBLEM OVERVIEW

One of the most straightforward ways to form an image of a target from synthetic aperture radar (SAR) data is via backprojection of the recorded phase delay measurements. Backprojection approximates inversion of the true scattering map

$$F : \text{Scene} \rightarrow \text{Data}, X \mapsto d$$

using the adjoint of a model (scattering) operator \tilde{F} , which is believed to be a reasonable approximation of F . The resulting image is thus

$$\tilde{X}_{BP} := \tilde{F}^\dagger X, \quad (1)$$

An alternative technique is to compute the least squares fit

$$\tilde{X}_{LS} := \underset{X}{\operatorname{argmin}} \{ \|\tilde{F}X - d\|_2^2 + \mu\rho(X) \}, \quad (2)$$

where the regularization term $R(X)$ with weight μ imposes certain constraints on the reconstructed scene \tilde{X}_{LS} . The limitation of Formulations (1) and (2) is that they rely directly on the fidelity of the model \tilde{F} .

2.1 Interferometric Imaging

As mentioned in the introduction, interferometric imaging forms cross-correlations dd^\dagger , obtaining a matrix of products $\{m_{ij}\} := d_i d_j$ from indexed measurements $\{d_i\}$. In radar imaging / phase retrieval problems, the heuristic motivation is that this matrix of quadratic measurements can be thought of as recording travel-time differences. As such, products of spatially and spectrally local pairs d_i, d_j are likely to be less sensitive to typical model errors such as sensor imperfections and imprecise sensor locations. This is explained more fully in section III.D of Demanet and Jugnon's paper.¹

For this reason, interferometric algorithms like CINT and QLI employ a matrix E of weights/masks to choose only those correlations that occur on short spatial and frequency scales. (Choosing E to be the identity matrix yields phaseless imaging as a special case.)

The CINT imaging process can be expressed as

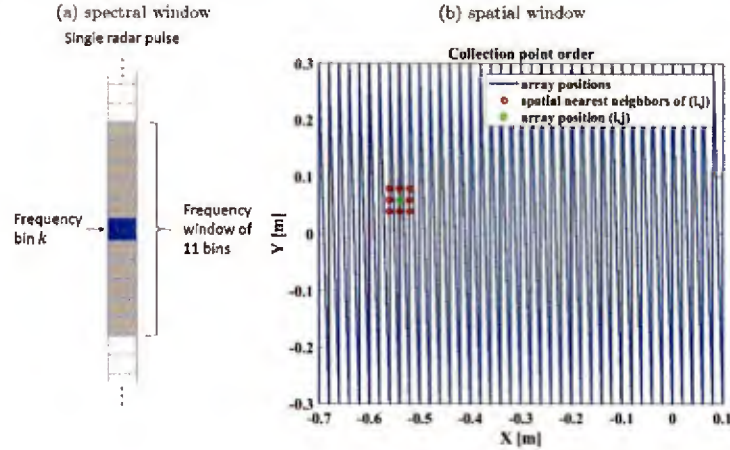
$$\tilde{X}_{CINT} := \operatorname{diag} \left\{ \tilde{F}^\dagger (E \circ [dd^\dagger]) \tilde{F} \right\}, \quad (3)$$

while QLI relates the image \tilde{X}_{QLI} to the solution of the optimization problem

$$\tilde{R} := \underset{R}{\operatorname{argmin}} \{ \underbrace{\|E \circ (\tilde{F}RR^\dagger \tilde{F}^\dagger - dd^\dagger)\|_2^2}_{L_{\mu,E}} + \mu\rho(X) \}. \quad (4)$$

A strength of QLI is that the optimization can be reformulated as a higher dimensional problem, in which $R \in \mathbb{C}^{N \times K}$, for some $K > 1$. Then we still have $RR^\dagger \in \mathbb{C}^{N \times N}$, but the problem (4) is less susceptible to trapping by local minima due to this rank K "relaxation".¹ The image \tilde{X}_{QLI} is then constructed by computing the principal eigenvector of $\tilde{R}\tilde{R}^\dagger$. Unfortunately, this relaxation step is not computationally feasible for $K \geq 2$ in most of our experiments due to the dimensionality ($d = 3$ in Sections 3.2 and 3.3) and large pixel counts (often more than 10^4) of our images. For this reason, we set $K = 1$ in the relaxation, resulting in an estimation of the scene $\tilde{X}_{QLI} := \tilde{R}$ directly as a vector. For the rest of this paper, therefore, we will identify the two and use X

Figure 1: Windowing of data via selection matrix E : data from frequency bins in a fixed spectral window on the left (of width 11) and a fixed spatial window on the right (red is cross-correlated with green) are all included in the cross-correlations selected via E .



to represent the variable of optimization. We choose $\rho(X) = \|X\|_2^2$ as our regularization term, and we refer to the objective function in Equation 4 as $L_{\mu,E}(X)$.

Spectral gap analysis of the graph Laplacian associated to the selection matrix E shows that auto-correlations ($E = I$) and even purely spatial correlations are in general insufficient to achieve model robustness.¹ However, experiments by Eric Mason et al. suggest that such kinds of phaseless or incoherent interferometry of this flavor can succeed if additional information is available.⁴ Figure 1 shows the spatial ordering and correlation windows for our experiments on measured data. Here we are including cross-correlations between local frequency/spatial locations pairs as depicted in the Figure 1.

We note that due to the diagonal nature of our choice of E and the linear ordering of the data vector d , we end up performing some correlations across ends of frequency sweeps and edges of 2-D array apertures that exceed any reasonable coherence length. Elimination of these non-local correlations could be accomplished by generating a dense mask that directly computes the proximity of data samples in frequency and spatial dimensions, but such an adjustment would give up at least part of the computational efficiency of our vectorized implementation for computing $L_{\mu,E}$ and its directional derivatives. Because the windows we have chosen are short relative to our data dimensions, these non-local correlations are rare enough that we expect them to have little effect on the QM reconstruction.

2.2 Implementation

Minimization of the loss functional $L_{\mu,E}(X)$ is achieved via a limited memory Broyden-Fletcher-Goldfarb-Shanno (hereafter "LBFGS") quasi-Newtonian optimization algorithm⁵ written in MATLAB. We provide this scheme with a MATLAB function which computes the value of this functional and its gradient with respect to X . The latter quantity is as follows:

$$\nabla_X L_{\mu,E} = 4\hat{F}^H (E \circ A) \hat{F} X + 2\mu X, \quad (5)$$

where we have abbreviated $A := (\hat{F} X X^H \hat{F}^H - dd^H)$.

The LBFGS optimizer computes iterative estimates of the target scene, terminating when the L_2 norm of this gradient drops below a prescribed threshold (our default value was 2×10^{-6} times the L_2 norm of the initial

gradient). For each of our three experiments, reasonable regularization parameters have been empirically chosen from a range of candidate values.

Note, however, that the structure of the regularized functional is not invariant to scaling transformations nor to variations in the density of matrix E . Perhaps for this reason, we have so far been unable to identify a heuristic analogous to that of Borcea et al.⁶ for tuning the matrix E to improve image quality. Continuing this search is a likely direction for future research.

Moreover, the regularization parameter μ is an absolute weight. For our experiments, this means that different preprocessing or sampling of the SAR data (or even different image cube pixel counts/spacings) can lead to different effective regularization weights, thus motivating a more automatic selection of regularization as a topic of future research.

For the same imaging regime as mentioned earlier in this section, it is also memory-prohibitive to form XX^T explicitly. Instead, by choosing E to have constant-diagonal structure, we can compute only those entries we consider spectrally and spatially relevant (those for which $e_{ij} = 1$) in a manner reminiscent of a discrete convolutions of $\tilde{F}X$ with itself. This diagonal structure of E corresponds to fixed windows with equal weight in frequency and space, and is (mostly) consistent with the model depicted in Figure 1 (see earlier discussion at the end of Section 2.1. More precisely, we compute

$$[E \circ (\tilde{F}X X^T \tilde{F}^T)]_{ij} = [E \circ (\tilde{F}X(\tilde{F}X)^T)]_{ij} = e_{ij}(\tilde{F}X)_i(\tilde{F}X)_j^* \quad (6)$$

for each nonzero entry e_{ij} in E . For each cross-correlation index pair (i, j) , we form a copy of $(\tilde{F}X)$ shifted by $i - j$ positions. By zero-padding and stacking these vectors into a matrix D_E , we can calculate all the entries (in diagonal format) using a single vectorized, element-wise product $D_E \circ (\tilde{F}X)$. Since E is a sparse diagonal matrix (with density less than 0.1% in our experiments), the reduction of the storage and computational burdens of this process from $O(N^2)$ to $O(N \times |e_{ij} \neq 0|)$ is substantial. Due to this reduction, the main algorithmic cost is computation of the linear operation $\tilde{F}(X)$ (and its adjoint \tilde{F}^H , in calculation of the gradient $\nabla_X L$). For the problem we consider, these operators employ non-uniform fast fourier transforms (NUFFT) on data sets of order $O(N_{pulses} \cdot N_f)$.⁷ The operators thus have complexity of $O(N_{pulses} \cdot (mN_f \log_2 N_f))$, where m is a number related to the oversampling and interpolation performed inside the NUFFTs.

3. RESULTS

In this section we show the results of utilizing QLI on simulated and measured data. In particular, in the first (Subsection 3.1), we show by simulation that QLI is robust to a periodic model mismatch. In the second (Subsection 3.2), we show that QLI can produce imagery in a controlled experiment with measured data that has resolution comparable to backprojection. Finally, subsection 3.3 completes the survey by testing the stability of the method against data collected in an experimental but randomly heterogeneous (buried target) setting.

3.1 Synthetic Data: Robustness of Lifting to Phase Error and Noise

In this section we demonstrate robustness of QLI to a fundamental model mismatch caused by unknown sensor position errors. The effect of this, to first order, is to induce phase errors. Our experiment shows that successful application of QLI reduces the need for exquisite autofocusing/motion compensation/estimation. We also show that QLI is stable against additive noise up to very low (even negative) signal to noise ratio (SNR).

We generate data from a simulated point scatterer at $(0, 0, 0)$ using 100 wideband frequencies from 0.5-2.0GHz, sampling returns at the frequencies from 201 positions on a 6m synthetic aperture at a 2m standoff from the scene center. The aperture sample spacing was chosen to achieve better than Nyquist rates ($\Delta x < \lambda_{min}/4$, where $\lambda_{min} = 15$ cm). According to our calculations, this geometry admits a resolution of 6.1cm in azimuth (X) and 10cm in range (Z). We choose a far smaller pixel spacing of 1.67cm in each dimension to expose the point spread function and facilitate comparison between the resolution and performance of QLI imaging and backprojection.

In these experiments, we define an ideal (linear) array geometry and a true array that differs from it by an additive sinusoidal perturbation in the Z direction (see Figure 2). We synthesize clean data using the true distance measurements ($d_{clean} = FX_{target}$). Complex-valued additive gaussian white noise (AWGN) with

Figure 2: Collection geometry with 1-D sinusoidal measurement error, amplitude=5cm. Shown are true and ideal arrays (used for data synthesis and imaging, respectively), image extent, and true point target location. The operators F and \tilde{F} utilize the true and ideal arrays, respectively.





Figure 3: Quadratic lifting improvement (bottom) over backprojected images (top) in the presence of increasing sinusoidal measurement error (left to right from 0cm to 7.5cm) and constant 7dB SNR. Small red circles indicate truth location of the point target in each plot. Coordinates are in pixels. Images are energy normalized, with dynamic range is 30dB and peaks set (per column) by the backprojected image.



variance $\sigma = \sqrt{2}\eta$ is then added across frequencies and pulses to simulate instrumental/environmental effects so that the synthesized data has the form


$$d := d_{clean} + \eta N = FX_{target} + \underbrace{\eta[\mathcal{N}(0, 1) + i\mathcal{N}(0, 1)]}_{\mathcal{C}\text{-valued AWGN}}. \quad (7)$$

The optimization utilizes the ideal array geometry for both imaging and model operators ($\tilde{F}^\dagger, \tilde{F}$), which means that the algorithm fits the noisy data using a flawed model. Despite this disadvantage, QLI is able to improve location and resolution of the point target over that of backprojection and regularized least squares. For this experiment, the selection matrix E used only the two nearest-in-memory neighboring array positions instead of the nine pictured in Figure 1. Regularization was $\mu = 0.1$.

Table 1: Values of the half power mainlobe width percentage reduction statistic $1-G$ for a variety of measurement error and SNR levels, showing the improvement of QLI over backprojection. Here $G = \sqrt{\frac{W_{LS}^{x,z} W_{LS}^{y,z}}{W_{BP}^{x,z} W_{BP}^{y,z}}}$ is a geometric mean of the individual range and cross-range improvement ratios $\frac{W_{QLI}}{W_{BP}}$. SNR is computed as $SNR = \frac{\text{var}(d_{\text{signal}})}{\text{var}(nN)}$.

		SNR		
		∞ dB	15dB	6dB
error	0 cm	11%	11%	11%
	2.5 cm	14%	14%	15%
	5 cm	24%	24%	25%

Figure 4: Regularized least squares does not share robustness of QLI. Images were formed from noise-free synthetic data with 5cm amplitude sinusoidal measurement errors. From left to right: backprojection using F^\dagger , backprojection using \tilde{F}^\dagger , least squares optimization using \tilde{F} , \tilde{F}^\dagger , and QLI using \tilde{F} , \tilde{F}^\dagger . Coordinates are in pixels; dynamic range is 30dB. Each image is peak normalized.




To measure the impact of QLI over backprojection, we compute a ratio of geometric means over range and cross-range half power mainlobe widths $G = \sqrt{\frac{W_{LS}^{x,z} W_{LS}^{y,z}}{W_{BP}^{x,z} W_{BP}^{y,z}}}$, computed for parametrically varying measurement errors and noise realizations, with $W_{m,d}$ representing this width across axis d for a image formation method m . These widths are computed from data slices through the center of the scene (ground truth for the point scatterer). The statistic $1-G$ presented in Table 1 is thus measure of the overall reduction in width by QLI reconstruction over backprojection. Due to the cyclic ambiguity of phase information, it is natural to expect phases with errors of $\frac{1}{2}\lambda_{\text{min}} = 7.5\text{cm}$ (or more) to be difficult or impossible to recover without an error model. We therefore call these errors "(super-)critical", and do not show results for this regime in Table 1 or Figure 3. We found in our simulations that the QLI algorithm is largely stable up to low (even negative) SNRs, and that it only fails to form a unified return as we approach the super-critical error regime.

This robustness to phase error is not shared by linear least squares optimization. For a wide range of regularization parameters, minimization of the least squares (LS) loss/objective function even accentuates split returns. (See the third plot in Figure 4 for one such result at regularization weight $\mu = 500$.) This is the case even when the initialization image is backprojected using the true adjoint operator F^\dagger (see the leftmost plot of Figure 4). Thus in simulation, QLI has showed robustness to mild model mismatch.

3.2 Resolution Comparable to Backprojection on Measured Data

Next, we evaluated the ability of QLI to produce imagery with resolution comparable to backprojection from actual free space measurements when there is limited model mismatch. On images formed from measured data, we have observed resolutions and clutter levels that compare favorably to fast backprojection. Our experiments were performed on one of the free space targets from a data set recorded at Georgia Tech University

Figure 5: Collection geometry, targets, and buried mine configuration in Georgia Tech study⁸ (images reproduced and cropped).



under carefully controlled conditions and made freely available at <https://waymond-scott.ece.gatech.edu/multistaticbeamformingdata/multistaticbeamformingdata/>.⁸ The right half (sub-figures “a” and “b”) of Figure 5 depicts the ground truth and collection geometry for this data. Although multiple receivers were used in the collection, in our tests we used only the first transmit/receive antennae pair (“T1-R1”).

For this experiment, the target was a $38.5 \times 46.5 \times 1.8$ cm plywood Georgia Tech logo supported (above ground) on a Styrofoam base. The data consist of 8281 radar pulses, collected over a 2D square grid of positions at a constant height above the target. Each pulse contains 401 frequencies on a wideband of 60MHz-8.06GHz. We used the 1681 pulses collected in an 80×80 cm region directly above the target. In order to obtain an even number of frequencies (400), we ignored the highest frequency bin.

Before image formation, we de-meant the frequency data across pulses to remove instrumental noise or bias and then additionally range-gate the data to remove the instrumental noise, feedthrough, and clutter from outside the target volume. Having thus preprocessed the data, we backproject an initial $31 \times 31 \times 11$ pixel scene estimate (select slices are displayed in the first row of Figure 6) and take $\mu = 0.3$ as our L_2 (absolute) regularization parameter for the QLI optimization. For this experiment, the selection matrix E was chosen so as to use cross-correlations across 9 spatial positions and 11 frequency bins as depicted in Figure 1. After 75 iterations, our LBFGS optimization scheme achieved gradient convergence (see discussion in Section 2.2) and we obtained the 3D image shown as slices in the second row of Figure 6. The optimization took less than 15 minutes of execution on a 3.9 GHz quad-core i7-3940XM cpu in a DELL Precision M4700 laptop with 32 GB of DDR3 RAM. The QLI reconstruction shows resolution comparable to backprojection, and contrast has been improved.

3.3 Controlled Experiment with Buried Targets

For a more difficult problem on which to test the lifting formulation, we turned to the buried target data publicly available from the same Georgia Tech research team.⁸ The subset of this data that we chose included an assortment of targets buried up to 12cm deep in clean, packed sand. See the right half of Figure 5 for ground truth target locations (depths are given in parentheses beneath the label for each target), and the paper for more details of the data collection.

Before lifting, we preprocessed the data, de-meaning across pulses in an attempt to remove the mostly uniform surface returns (< 2 cm height variation). We used a two-layer scattering model with calculation of Snell refraction points (Figure 7) when computing time lags for our imaging routines (using code obtained from⁹).

Figure 6: "GT Plywood" Target: Comparison of backprojected image vs. QLI reconstruction. We show horizontal slices of interest at heights near expected true height at a dynamic range of 25dB to match that of ⁶. No post-processing/normalization has been performed to equalize color scales.

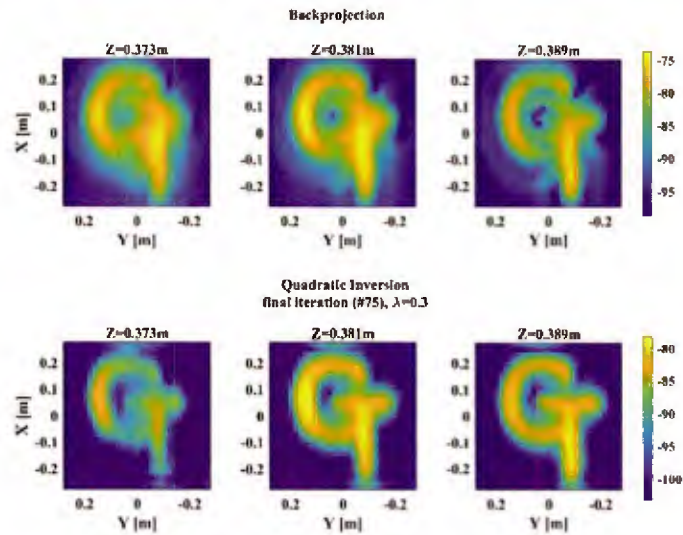




Figure 7: Two-layer forward model for wave scattering from buried target that includes refraction/propagation through ground



Determination of these refraction points is accomplished by solving a fourth degree polynomial equation for each transmit/receive pair and pixel location, using the Georgia Tech team's observation of wave propagation speed through the sand. These solutions for the refraction points were obtained by means of Newton's method for finding zeros of a function.

After searching the parameter space (and based on results from previous experiments), we settled on an absolute L_2 regularization parameter of $\mu = 0.05$. We selected 11 inter-frequency and 9 inter-pulse cross correlations

Figure 8: Optimization behavior of QLI objective function $L_{\mu,E}$. In order to surmount the lack of algorithm-specific preconditioning in our implementation of Equation (4), we select as our buried target reconstruction the estimate from the optimizer iteration where $\|X\|_2$ first begins to increase (dotted line).



depicted in Figure 1. The LBFGS optimization scheme described in Section 2 was run for 66 iterations (less than 8 minutes on the same laptop mentioned in Section 3.2). We chose as our solution an intermediate estimate of the target scene that was selected by identifying the first iteration on which the norm of the solution began to increase (see Figure 8). This is our proxy termination condition in the absence of sufficient gradient decay, since consistent achievement of the latter has so far proved elusive. Figure 9 depicts the solution thus chosen. It is important to note that in the QLI reconstruction, the impulse response has been narrowed. This is especially apparent on the deep target. Clutter has also been significantly reduced.

We did explore finding an optimal combination of optimization and regularization for carrying out QLI on these data, but we have not been successful to date. However, we did find some significant patterns in this investigation. From a plot of the objective function values $L_{\mu,E}(X_i)$ (a.k.a., loss) alongside the L_2 norm of the estimated scene X_i (a.k.a., solution) at each iteration, it appears that the target to clutter ratio was improving principally when the norm of the scene was decreasing. This plot is shown in Figure 8 and iterations of interest are labeled. Our first attempted solution to this problem was to increase the regularization parameter, but an optimal value has proved exceedingly elusive. The transition between over-regularization and the behavior exhibited in the results shown here seems to be rather sharp. To surmount this difficulty, we eventually introduced the stopping criterion described above, and illustrated in Figure 8. With this criterion, we were able to reconstruct a sharper image of buried targets than backprojection.

4. CONCLUSIONS

In this paper we have described our rank 1 implementations of quadratic lifting inversion and evaluated three claims about the method by conducting imaging experiments using simulated and measured data. We describe how to efficiently compute data-space cross-correlations by designing the selection matrix E to be constant along diagonals and avoiding the full formation of the product (XX^T) , at the cost of including relatively few unwanted data correlations. Via a synthetic data experiment, we have showed that QLI has robustness to model errors not shared by backprojection or least squares inversion. In the presence of a defective array geometry, these classical methods produced split returns from a single point target, while QLI was able to recover a unified return for array model inaccuracies of up to 7.5cm. On measured data, we have demonstrated that QLI attains

Figure 9: Buried Targets: backprojection vs QLI. We show horizontal slices of interest at nominal depths of 4, 9, and 14cm below ground-air interface. Note that targets appear a few centimeters deeper than predicted by ground truth (see Figure 5 or the Georgia Tech paper⁸), and that there is no target at a depth of 9cm.




image resolution comparable to that of backprojection, improves clutter suppression, and increases contrast of the targets. Finally, we were able to image buried targets using a two-layer model without carefully tuning the propagation speed.

Along the way, we encountered several promising directions for future work. Prominent among these are (i) the search for a heuristic to tune density of the selection matrix E and the regularization constant μ to characteristics of the measured data; (ii) exploration of a more principled automatic truncation rule or preconditioning of the problem to replace the one outlined in Section 3.3 and illustrated in Figure 8; and (iii) the testing of QLI for higher rank ($K > 1$) relaxations on 3D/GPR imaging problems, which currently appear out of reach of our current implementation and hardware. A less pressing direction might be the design of a scaling invariant QLI formulation.

ACKNOWLEDGMENTS

We would like to acknowledge the support of ONR grant N00014-16-1-2623. We would also like to thank Professor Waymond Scott of Georgia Tech University for making his ground penetrating radar data sets freely available.

REFERENCES

- [1] Demanet, L. and Jugnon, V., "Convex recovery from interferometric measurements," *IEEE Transactions on Computational Imaging* **3**, 252–295 (2017).
- [2] Borcea, L., Papanicolaou, G., and Tsogka, C., "Interferometric array imaging in clutter," *Inverse Problems* **21**(4), 1419 (2005).
- [3] Borcea, L., Garnier, J., Papanicolaou, G., and Tsogka, C., "Coherent interferometric imaging, time gating and beamforming," *Inverse Problems* **27**(6), 065008 (2011).
- [4] Mason, E. and Yazici, B., "Passive phaseless sar imaging," in [2018 IEEE Radar Conference (RadarConf18)], 0292–0297, IEEE (2018).
- [5] Liu, D. C. and Nocedal, J., "On the limited memory bfgs method for large scale optimization," *Mathematical programming* **45**(1-3), 503–528 (1989).
- [6] Borcea, L., Papanicolaou, G., and Tsogka, C., "Adaptive interferometric imaging in clutter and optimal illumination," *Inverse problems* **22**(4), 1405–1436 (2006).
- [7] Greengard, L. and Lee, J.-Y., "Accelerating the nonuniform fast fourier transform," *SIAM Review* **46**, 443–454 (2004).
- [8] Counts, T., Gurbuz, A. C., W. R. Scott, Jr., McClellan, J. H., and Kim, K., "Multistatic ground-penetrating radar experiments," *IEEE Transactions on Geoscience and Remote Sensing* **45** (2007).
- [9] Masarik, M. P., Burns, J., Thelen, B. T., and Sutter, L., "Fast 3D subsurface imaging with stepped-frequency GPR," in [Detection and Sensing of Mines, Explosive Objects, and Obscured Targets XX], *Proc. SPIE* **9454**, 945417 (May 2015).