

FINAL REPORT

Structural Acoustic UXO Detection and Identification in Marine Environments

SERDP Project MR-2103

MAY 2016

B. H. Houston
M. Saniga
H. Simpson
Z. J. Waters
D. Amon
E. Williams
T. Yoder
Naval Research Laboratory

J. A. Bucaro
Excet, Inc.

A. Sarkissian
Sotera Defense Technology & Systems, Inc.

Distribution Statement A

This document has been cleared for public release



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

REPORT DOCUMENTATION PAGE					Form Approved OMB No. 0704-0188	
<p>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.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p>						
1. REPORT DATE (DD-MM-YYYY) 18-07-2016		2. REPORT TYPE Memorandum Report			3. DATES COVERED (From - To)	
4. TITLE AND SUBTITLE Structural Acoustic UXO Detection and Identification in Marine Environments - Final report for Follow-on Work- MR-2103				5a. CONTRACT NUMBER		
				5b. GRANT NUMBER		
				5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S) J. A. Bucaro*, A. Sarkissian**, B.H. Houston, M. Saniga, H. Simpson, Z.J. Waters, D. Amon, E. Williams, and T. Yoder				5d. PROJECT NUMBER MR-2103		
				5e. TASK NUMBER		
				5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval research Laboratory, Code 7130 4555 Overlook Avenue, SW Washington, DC 20375-5350					8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) SERDP Program Office 4800 Mark Center Drive, Suite 17D08 Alexandria, VA 22350-3605					10. SPONSOR/MONITOR'S ACRONYM(S) SERDP	
					11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited						
13. SUPPLEMENTARY NOTES *On-site contractor with Excet, Inc.Springfield, VA 22151 **On-site contractor with Sotera Defense Technology & Systems, Inc., Columbia, MD 21046						
14. ABSTRACT This work completes the post measurement analysis of the data collected in June 2013 in the Gulf using the BOSS-40. Inclusion of AUV position data has improved image quality and the ability to correctly register buried target depths. Along-range resolution resolved the UXO long axes, and cross-range resolution was within a factor of two of the UXO diameters. In a study that included all targets, results demonstrate that acoustic color can separate the detections from the UXO/non-UXO groups. BOSS source calibration studies demonstrated that the source/AUV system does not radiate spherically causing a significant impact on false target classification performance. "Leave one out" studies indicated that training experimentally can produce classification algorithms that have a near 1 probability of classification for UXO targets of the type used in training and a false positive percentage of about 33% against false targets unlike the UXO from a structural acoustics point of view. While the false alarm rate is higher against targets similar to the UXO, this performance would improve significantly with the use of a more uniform source.						
15. SUBJECT TERMS Underwater UXO Sonar detection of UXO Structural acoustic classification						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON	
a. REPORT	b. ABSTRACT	c. THIS PAGE			B.H. Houston	
Unclassified	Unclassified	Unclassified	UL	95	19b. TELEPHONE NUMBER (include area code) 202-404-3840	

TABLE OF CONTENTS

LIST OF TABLES	ii
LIST OF FIGURES	ii
LIST OF ACRONYMS	viii
KEYWORDS	viii
ACKNOWLEDGEMENTS	viii
ABSTRACT	1
OBJECTIVE	2
TECHNICAL APPROACH	2
RESULTS AND DISCUSSION	4
SUBTASK 2.4. Re-configure imaging algorithm and include AUV position data and north/south AUV path data	4
SUBTASK 2.5. Include North/South AUV multiple path data in target separation study	14
SUBTASK 2.6. Include all additional proud targets (+16) in the RVM target separation study	17
SUBTASK 2.7. Analyze acoustic color from non-epoxy-filled UXO (proud) compared to epoxy-filled (buried) UXO.	18
SUBTASK 2.8. Perform acoustic BOSS source calibration in the NRL near-field holography facility.	20
SUBTASK 2.9 Using new BOSS source calibration data, generate and analyze specular-filtered plan view images of the buried UXO	21
SUBTASK 5.1 Develop methodology for integrating BOSS source calibration data into the data processing and analysis.	23
SUBTASK 5.2 Using BOSS calibration data and optimized features, determine improvement in feature separation results and elastic imagery.	31
SUBTASK 6.1: Develop training methods for the selected identification algorithms	71
SUBTASK 6.2: Demonstrate effectiveness of trained identification methods in separating UXO from clutter.	79
CONCLUSIONS	94
LITERATURE CITED	97
APPENDIX	100

LIST OF TABLES

TABLE 1 – Proud and Buried Target List

TABLE 2 – Number Of Times A Particular Target Is A False Positive On All Paths

TABLE 3 – False Positive Occurrences (Before Source Calibration)

TABLE 4 – False Positive Occurrences (After Combined Ring Calibration)

TABLE 5 – False Positive Calls For The Various Targets In The “Leave One Out” Tests

TABLE 6 – False Positives (FP) and True Negatives (TN) 3-13.5kHz

TABLE 7 – False Positives (FP) and True Negatives (TN) 3-20kHz

LIST OF FIGURES

Figure 1. Drawing illustrating a BOSS flight path along x , the ping sequence, and the interrogated swath on the bottom.

Figure 2. Cartoon illustrating the Δz ambiguity for a sloping sediment when the flight paths are laterally far from the target.

Figure 3. The EW and NS flight paths executed by BOSS over the target field.

Figure 4a. Plan view and two depth images for n1 through n4 targets using EW flight path h .

Figure 4b. Plan view and two depth images for n5 through n8 targets using EW flight path h .

Figure 4c. Plan view and depth images for n9 - n11 targets using EW flight paths g , h , and c , respectively.

Figure 5a. The plan view and two depth images for n1 through n4 targets using NS flight paths e and k .

Figure 5b. The plan view and two depth images for n5 through n8 targets using NS flight paths m , o , and s .

Figure 5c. The plan view and two depth images for n9 through n11 targets using NS flight paths y , Ba , and v .

Figure 6. Assumed rock orientation with respect to x , y scans with its long axis (orange arrow) rotated up out of the horizontal plane $\sim 80^\circ$ about the x axis.

Figure 7. The three 2-D images of the rock (n10) produced from the EW flight path h (left) and from the NS flight path Ba (right).

Figure 8a. Discriminatively-trained RVM classification algorithm trained on even numbered source pings and tested on odd numbered pings. Targets include 9 buried UXO and 9 non-UXO – 2 buried and 7 proud. NS paths (3-20 kHz).

- Figure 8b. Discriminatively-trained RVM classification algorithm trained on even numbered source pings and tested on odd numbered pings. Targets include 9 buried UXO and 9 non-UXO – 2 buried and 7 proud. NS paths (3-13.3 kHz).
- Figure 9. Discriminatively-trained RVM classification algorithm trained on even numbered source pings and tested on odd numbered pings. Targets include 9 buried UXO and 8 non-UXO – 2 buried and 6 proud. EW paths (3 – 13.3 kHz).
- Figure 10. ROC curves for the results in Fig 9.
- Figure 11. Discriminatively-trained RVM classification algorithm trained on even numbered source pings and tested on odd numbered pings. Targets include 9 buried UXO and 9 non-UXO – 2 buried and 7 proud. NS and EW paths (3-13.3 kHz).
- Figure 12. ROC curve associated with the results in Fig 11.
- Figure 13. Discriminatively-trained RVM classification algorithm trained/tested on even/odd numbered source pings, respectively. Targets: 9 buried and 10 proud UXO and 14 non-UXO – 2 buried and 12 proud. NS paths (3-13.3 kHz).
- Figure 14a. RVM discriminative algorithm trained and tested on even and odd numbered source pings, respectively using the acoustic color maps. On average, we used epoxy-filled 155mm shell data from 7 EW paths and empty 155mm shell data from 5 EW paths.
- Figure 14b. RVM discriminative algorithm trained and tested on even and odd numbered source pings, respectively using acoustic color maps. On average, we used epoxy-filled 155mm shell data from 7 EW paths and empty 155mm shell data from 7 EW paths.
- Figure 15. The NRL Structural Acoustic Pool facility.
- Figure 16. The two orientations for hanging BOSS in the NRL Structural Acoustic Pool facility.
- Figure 17. Source pressure levels in dB measured at the 5 receivers 3.4 meters from BOSS as it is rotated in steps of 0.5 degrees. Lower right: what the pattern would be were the BOSS source omnidirectional.
- Figure 18. The measured time waveforms at the R1 through R4 receivers and the prediction from TVR curve.
- Figure 19. The measured signal spectra measured at the R1 through R4 receivers and the prediction from TVR curve.
- Figure 20. Geometry describing the mapping between the measured source radiation pattern and the radiating source strength spatial distribution.
- Figure 21. The supersonic images mapped onto the BOSS AUV upper half and the BOSS AUV lower half. The photo indicates the actual locations on BOSS.
- Figure 22. Geometry used to model a spherical source located within the AUV compartment represented with image sources.
- Figure 23. An example of the radiation pattern versus frequency obtained with the model illustrated in Fig. 22.
- Figure 24. Radiation patterns predicted using a model which includes direct radiation from a vertically displaced spherical source plus scattering from a 12.75inch diameter infinite

cylindrical shell. (Left) calibration ring data; (Middle) model with rigid cylinder; (Right) model with elastic cylinder.

- Figure 25. Diagram for simple calibration ring projection method.
- Figure 26. Geometry for the synthetic holography projection technique. Radial calibration rings in blue, axial calibration ring in red, and spherical hologram surface in green.
- Figure 27. Comparisons of the source amplitude onto the sediment surface for the four projection techniques at low frequency.
- Figure 28. Comparisons of the source amplitude onto the sediment surface for the four projection techniques at mid frequency.
- Figure 29. Comparison of source spatial distribution over target field for a spherical source versus approximately projecting source calibration ring data for three frequencies.
- Figure 30. Cartoon illustrating the BOSS data collection geometry and the relative positions of the various targets and the flight paths. Paths g , h , and i are dashed.
- Figure 31a. False targets considered in studies along paths h and i .
- Figure 31b. False targets considered in studies along path g .
- Figure 32. Target field (left) and spectral/spatial (one dimensional) source pattern: spherical (upper); radial calibration ring (lower).
- Figure 33. Target field (left) and spectral/spatial (one dimensional) source pattern: spherical (upper); axial calibration ring (lower).
- Figure 34. Generative target feature separation using alternate pings for the 12 TS average feature set over path h .
- Figure 35. Generative target feature separation using alternate pings for the 7 TS average feature set over path h .
- Figure 36. Generative target feature separation using alternate pings for the max TS, bin number, standard deviation feature set over path h .
- Figure 37. Generative target feature separation using alternate pings for the max TS, bin number feature set over path h .
- Figure 38. Generative target feature separation using alternate pings for the max TS, bin number feature set over path h .
- Figure 39. Generative target feature separation using alternate pings for the mean TS feature set over path h .
- Figure 40. Generative target feature separation using alternate pings for the 140 frequencies x 22 positions, 20m aperture feature set over path h .
- Figure 41. Generative target feature separation using alternate pings for the 140 frequencies x 22 positions, 2m aperture feature set over path h .
- Figure 42. Generative target feature separation using alternate pings for the 12 TS average feature set over path i .

- Figure 43. Generative target feature separation using alternate pings for the 7 TS average feature set over path *i*.
- Figure 44. Generative target feature separation using alternate pings for the max TS, bin number, standard deviation feature set over path *i*.
- Figure 45. Generative target feature separation using alternate pings for the max TS, bin number feature set over path *i*.
- Figure 46. Generative target feature separation using alternate pings for the max TS feature set over path *i*.
- Figure 47. Generative target feature separation using alternate pings for the mean TS feature set over path *i*.
- Figure 48. Generative target feature separation using alternate pings for the 140 frequencies x 22 positions, 20m feature set over path *i*.
- Figure 49. Generative target feature separation using alternate pings for the 140 frequencies x 22 positions, 2m feature set over path *i*.
- Figure 50. Generative target feature separation using alternate pings for the 12 TS averages feature set over path *g*.
- Figure 51. Generative target feature separation using alternate pings for the 7 TS averages feature set over path *g*.
- Figure 52. Generative target feature separation using alternate pings for the max TS, bin number, standard deviation feature set over path *g*.
- Figure 53. Generative target feature separation using alternate pings for the max TS, bin number feature set over path *g*.
- Figure 54. Generative target feature separation using alternate pings for the max TS feature set over path *g*.
- Figure 55. Generative target feature separation using alternate pings for the max TS feature set over path *g*.
- Figure 56. Generative target feature separation using alternate pings for the 140 frequencies x 22 positions, 20m aperture feature set over path *g*.
- Figure 57. Generative target feature separation using alternate pings for the 140 frequencies x 22 positions, 2m aperture feature set over path *g*.
- Figure 58. Generative target feature separation using alternate pings for Max TS feature set over path *h*: original (upper); eliminating UXO n4 in training (lower).
- Figure 59. Generative target feature separation using alternate pings for 140 frequencies x 22, 2m aperture feature set over path *h*: original (upper); eliminating UXO n4 in training (lower).
- Figure 60. Generative target feature separation using alternate pings for Max TS feature set over path *i*: original (upper); eliminating UXO n4 in training (lower).

- Figure 61. Generative target feature separation using alternate pings for 140 frequencies x 22 positions, 2m aperture feature set over path *i*: original (upper); eliminating UXO n4 in training (lower).
- Figure 62. Generative target feature separation using alternate pings for 140 frequencies x 22 positions, 20m aperture feature set over path *g*: original (upper); eliminating UXOs n1 and n7 in training (lower).
- Figure 63. Acoustic images of a buried cinder block and a buried 5" rocket generated using bi-static acoustic scattering returns in the NRL sediment pool facility.
- Figure 64. Experiment geometry and algorithm used to generate the acoustic images.
- Figure 65a. Elastic highlight images: (upper) bi-static sediment pool scanning; (middle) time gating of Duck, NC rail data; (lower) time gating Boss data in Gulf.
- Figure 65b. The three 2-D images of the horizontally buried five inch rocket (left) and the 155mm shell (right) produced from the E-W flight path *h*.
- Figure 65c. Diagram explaining how creeping or elastic circumferential wave can produce a time-extended extended scattered echo and thus an under resolved extended 2D image at an angle in the *z,y* plane.
- Figure 66. Generative acoustic color 20m aperture feature separation (left) and corresponding ROC curves (right) using alternate pings over path *h*. No added noise (upper) and 0.05% max added noise (lower).
- Figure 67. Generative acoustic color 20m aperture feature separation (left) and corresponding ROC curves (right) using alternate pings over path *h*. 0.1% max added noise (upper) and 0.3% max added noise (lower).
- Figure 68. Generative acoustic color 20m aperture feature separation (left) and corresponding ROC curves (right) using alternate pings over path *h*. 0.5% max added noise (upper) and 1% max added noise (lower).
- Figure 69. Generative acoustic color 20m aperture feature separation (left) and corresponding ROC curves (right) using alternate pings over path *h*. 2% max added noise (upper) and 3% max added noise (lower).
- Figure 70. Generative acoustic color 20m aperture feature separation (left) and corresponding ROC curves (right) using alternate pings over path *h*. 3.5% max added noise (upper) and 1% max added noise (lower).
- Figure 71. 1 – area under ROC curves versus added noise level in % max from Figs. 66 – 70.
- Figure 72. Five inch rocket target strength versus incidence angle as measured in large pool facility.
- Figure 73. The plot of field value for each UXO for the acoustic color value measured at each BOSS receiver divided by the average value in the spectral band in the Gulf measured data base used to generate the results shown in Figs. 66 – 70.
- Figure 74. Generative target feature separation training on path *g* and testing on paths *h* and *i* for the 200 frequencies x 22 positions, 20m aperture feature set over the band 3-13 kHz.

- Figure 75. Generative target feature separation training on path h and testing on paths i and g for the 200 frequencies x 22 positions, 2m aperture feature set over the band 3-13kHz.
- Figure 76. Generative target feature separation training on path h and testing on paths i and g for the 200 frequencies x 22 positions, 20m aperture feature set over the band 3-13kHz.
- Figure 77. Generative target feature separation training on path i and testing on paths h and g for the 200 frequencies x 22 positions, 2m aperture feature set over the band 3-13kHz.
- Figure 78. Generative target feature separation training on path i and testing on paths h and g for the acoustic color, 20m aperture feature set over the band 3-13kHz.
- Figure 79. Generative target feature separation training on path g (without $n1$ and $n7$) and testing on paths h and i for the 200 frequencies x 22 positions, 2m aperture feature set over the band 3-13kHz.
- Figure 80. Generative target feature separation training on path h (without $n4$) and testing on paths i and g for the 200 frequencies x 22 positions, 20m aperture feature set over the band 3-13kHz.
- Figure 81. Generative target feature separation training on path h (without $n4$) and testing on paths i and g for the 200 frequencies x 22 positions, 2m aperture feature set over the band 3-13kHz.
- Figure 82. Generative acoustic color feature separation training on path g and testing on paths h and i over band 3-13kHz after equivalent source calibration.
- Figure 83. Generative acoustic color feature separation training on path g and testing on paths h and i over band 3-20kHz after equivalent source calibration.
- Figure A1. Photograph of BOSS in the Water.
- Figure A2. Plan view and head-on view of BOSS.
- Figure A3. Photograph of BOSS prior to complete assembly showing some of the internal components.
- Figure A4. Preliminary 3-D images obtained from June 23, 2014 BOSS data by Richard Holtzapple and Nick Pineda.

LIST OF ACRONYMS

AUV	Autonomous Underwater Vehicle
BOSS	Buried Object Scanning Sonar
DVL	Doppler Velocity Log
EW	East/West
IMU	Inertial Measurement Unit
NRL	Naval Research Laboratory
NSWC-PCD	Naval Surface Warfare Center – Panama City Division
NS	North/South
PMA	Post Measurement Analysis
RVM	Relevance Vector Machine
ROC	Receiver Operating Characteristic
SA	Structural Acoustic
SAS	Synthetic Aperture Sonar
S/N (SNR)	Signal to Noise
SVD	Singular Value Decomposition
TREX	Transmission and Radiation Experiments
TS	Target Strength
TVR	Transducer Voltage Response
UXO	Unexploded Ordnance

KEYWORDS

Underwater buried UXO, structural acoustic target identification, sonar UXO detection

ACKNOWLEDGEMENTS

This work was performed under direct support from the SERDP Program Office. We would like to acknowledge the very positive and encouraging program management role played by Dr. Herb Nelson, Acting Director of SERDP and ESTCP and SERDP Program Manager for Munitions Response. We would also like to acknowledge the significant impact made in these studies by the leveraging of results from scientifically related efforts funded by ONR.

ABSTRACT

SERDP goals require the development of sonar technologies able to separate UXO and false target detections. The specific goal addressed in the follow-on work reported here is to complete analysis studies using the data collected in June 2013 in the Gulf of Mexico using the BOSS-40 sonar. The measurements and preliminary post measurement analysis (PMA) results have been described in the April 2014 MR2103 Final Report.

In the following, we describe the results of this PMA. In summary, inclusion of the AUV position data in the imaging algorithm has improved the quality of the images and their ability to correctly register buried targets with respect to the location of the sediment interface. Along-range resolution was sufficient to resolve the UXO long axes, and the cross-range resolution typically was within a factor of two of the UXO diameters. We find that the performance of the RVM feature separation algorithm is slightly better for the NS paths than for the EW paths. Further, in a study that included all targets, the results demonstrate that the multi-dimensional feature extracted from acoustic color can separate the detections from the UXO/non-UXO groups and that this feature can be used to separate the epoxy-filled shells from the others. BOSS source calibration studies were completed demonstrating that the coupled source/AUV system does not radiate as a spherically radiating source and at any particular frequency the acoustic pressure fluctuates considerably with direction. Structural acoustic models developed to predict these patterns, while accounting for many of the details, did not exhibit sufficient fidelity. Several methods to project the calibration ring data to the sediment area below were then used to partially correct for the source spatial patterns. The resulting improvement in feature separation between UXO and non-UXO was determined and documented for eight feature sets training generatively and then testing on alternate source pings. The results of these studies suggest that more precise knowledge of the source spatial characteristics are required than is available from the projection techniques we used particularly in view of the limited number of calibration rings we had obtained. A study was carried out to estimate the accuracy or repeatability required in the acoustic color spatial/spectral functions when training on acoustic color features in order to have good classification performance. These results suggested that training experimentally with the actual sonar in an environment similar to the intended test site should be favored over training on numerically generated data. In the final task, studies named “leave one out” were carried out using the 2013 Gulf data base to test the efficacy of the experimental training approach. In these studies, the RVM algorithms were trained generatively using all pings along a single path and then tested using all the pings on two other paths. The results indicate that training experimentally in this way can produce BOSS RVM classification algorithms that can have a near 1 probability of detection and classification for UXO targets of the type used in the training and a fairly low false positive percentage of about 33% when used to discriminate against targets unlike the UXO from a structural acoustics point of view. While the false alarm rate is much higher against targets similar to the UXO, it is argued that this performance would improve significantly with the use of a more uniform source.

OBJECTIVE

Many active and former military installations have ordnance ranges/training areas with adjacent water environments in which unexploded ordnance (UXO) now exists due to wartime activities, dumping, and accidents. SERDP goals require the development of innovative technologies¹⁻⁸ able to separate UXO from false targets and to discriminate amongst UXO targets themselves. The objective of the overall program is to address the scientific and technical issues whose resolution would result in an efficient, high performance structural acoustic (SA) feature-based underwater sonar technology that can detect and localize buried (and proud) targets and separate the detections into UXO vs non-UXO. Our focus is on marine-based sonars that could look both downward (and to some degree sideways) in water depths ranging from several meters to tens of meters. The overall goal here is to develop a sonar approach which results in robust identification algorithms based on structural acoustic features and complementary 3D SAS (synthetic aperture sonar) images and to demonstrate the ability to detect and classify proud and buried UXO in the presence of natural and man-made clutter with actual structural acoustic sonar systems at-sea. The specific goal addressed in the follow-on work reported here is to complete additional analysis studies using the data collected in June 2013 in the Gulf of Mexico using the BOSS-40 sonar⁹. The measurements and preliminary post measurement analysis (PMA) results have been described in Final Report “Structural Acoustic UXO Detection and Identification in Marine Environments, April 2014.”¹⁰

TECHNICAL APPROACH

In June of 2013 NSWC, together with NRL and Bluefin Robotics, successfully flew the BOSS-40⁹ structural acoustic sonar in the Gulf of Mexico over a target field in which were buried nine NRL epoxy-filled UXO’s and 2 false targets and some 20 proud targets. Through this study we established (1) the ability to produce both 3-D images and target strength spatial/spectral maps of the buried targets and (2) that UXO targets and false targets can be separated using a multi-dimensional but otherwise straightforward feature set based on acoustic color. The details of the BOSS system are described in the Appendix.

The general location of the target field was roughly 2 miles off the coast of Panama City, Florida. This location was the sight of a broader set of measurements made about the same time as part of the so-called TREX-13 exercise. For that program, a proud target field was created (see the proud targets listed in Table 1), and scattering measurements were carried out using sources and receivers mounted on towers along the northern perimeter of the target field. These assets were removed just prior to our exercise using the BOSS AUV system.

TABLE 1 - Proud and Buried Target List

Proud Target				NRL Epoxy-Filled Buried Targets	
t1	DEU Trainer	t14	Scuba Tank w/water w/ stem	n1	5inch Rocket nose-up 60°
t2	Rock	t15	2:1 Aspect Phone Pole Section	n2	5inch Rocket nose-up 30°
t3	55 Gallon Filled Drum	t17	2 ft Aluminum cylinder	n3	5inch Rocket horizontal
t5	5:1 Aspect Phone Pole Section	t18	Cement Block	n4	155mm Projectile horizontal
t7	3ft Aluminum Cylinder	t19	Tire	n5	155mm Projectile horizontal 90°
t8	155mm Projectile w/o collar	t20	Aluminum UXO Replica	n6	155mm Projectile horizontal 20cm
t9	155mm Projectile w/ collar	t22	Original Material UXO	n7	155mm Projectile nose-up 30°
t10	Panel Target	t25	Bullet #1	n8	155mm Projectile nose-up 60°
t11	152mm TP-T	t28	155mm Projectile w/collar	n9	120mm Mortar horizontal
t12	81mm Mortar	t29	Bullet #2	n10	Large Rock (no epoxy filler)
t13	Scuba Tank w/water w/o stem	t30	Finned Shell #1	n11	Cinder Block (no epoxy filler)

The nine buried NRL epoxy-filled UXOs include three 5 inch rockets, five 155mm projectiles, and a 120mm mortar. The two buried NRL false targets included a large rock and a conventional cinder block. These eleven NRL targets were buried by Dr. Kevin Williams (Applied Physics Laboratory at University of Washington) and his diving team aided by Dr. Michael Richardson about three weeks before the BOSS exercise. The plan was to bury all but one of the NRL targets in such a way that the highest point of the target would be 10cm below the sediment-water interface. The one exception was to be a horizontally buried 155mm projectile at a 50cm depth. Generally, the divers executed the plan fairly well given the difficulty of burying such relatively large objects in 60 foot waters. Although it was not possible to know the burial condition of each target precisely, the general consensus was that the depth and orientation achieved for each target were roughly according to plan. Two exceptions were that the cinder block burial depth was less than 10cm and the deeply buried 155mm projectile was closer to 20cm from the sediment interface. There were 23 targets placed proud on the sediment. These included some UXOs not filled with epoxy, a rock, and a number of man-made objects. Together with the buried rock and cinderblock, 7 of these 23 proud targets serve as false targets in the target feature separation studies to be discussed later. These seven include the rock (t2), the 55 gallon water filled drum (t3), the 5:1 aspect telephone pole section (t5), the panel (t10 or CP), the 2:1 aspect telephone pole section (t15), the cement block (t18), and the tire (t19). The proud and buried targets in the target field are listed and labelled in Table 1.

In this SERDP program (MR2103), we demonstrated the structural acoustic technology with an autonomous underwater vehicle (AUV) - based SA sonar successfully detecting UXO buried in the sediment in the Gulf of Mexico and showing that their structural acoustic features could be used to separate them from proud and buried false targets¹⁰. In particular, the Buried Object Scanning Sonar (BOSS)⁹ was set up to fly at several meter altitudes in 60 foot water depths off the Coast of Panama, City in the Gulf of Mexico. This AUV-based sonar carried out north-south, east-west, and diagonal paths over a target field comprised of nine epoxy-filled UXOs buried about 10cm below the sediment/water interface, two buried false targets, and 23 proud targets. The scattered acoustic pressure signals from the target field received at each of the 40 wing sensors as the vehicle moved in a straight line were processed in a synthetic aperture manner yielding both 3-D images and several acoustic color constructs for all the buried targets and for seven of the

proud targets considered to be non-UXO. Most of the images gave useful information related to the target size and burial orientation. A multi-dimensional acoustic color feature extracted from the acoustic color maps demonstrated almost perfect separation between the nine UXO and the nine false targets. These results demonstrate that typical buried UXO can be detected, imaged, and in principle classified (UXO versus non-UXO) using a structural acoustic sonar and a relevance vector machine identification algorithm (RVM)¹¹. What remained to be demonstrated was that the RVM algorithm could be successfully trained a-priori. In addition, earlier measurements made with a rail-based structural acoustic sonar in 25 foot waters off the Duck, NC coast validated the new concept of short time (specular) versus long time (elastic highlight)¹² plan view imaging. For reasons associated with the BOSS projector, such images could not be produced from the Gulf BOSS data. Finally, the first accomplishment early in the program involved the sediment pool demonstration of a numerically trained classifier. Here we demonstrated good classifier performance training the relevance vector machine algorithm on a finite element^{13,14} target burial angle simulation data set and testing on UXO and false targets buried in the sediment pool¹⁵.

RESULTS AND DISCUSSION

The 3-D images^{16,17} previously generated and reported in the Final Report¹⁰ for MR2103 were generated in the following way.

Figure 1. Drawing illustrating a BOSS flight path along x, the ping sequence, and the interrogated swath on the bottom.

(See Fig. 1.) For the EW flight paths, the vehicle moved in the x or $-x$ direction. To produce images or carry out SAS processing, the measured signal was initially time windowed to remove the bottom bounce (and everything arriving prior to it) and the surface reflection (with everything arriving after it). Next, using the measured acoustic signal from 40 pings (40 x positions of the vehicle), the image was produced as a function of x, y, z using data over the frequency range of 3 to 20 kHz. Next the images at the same (x, y, z) locations were produced from a different set of 40 pings starting from the 6th ping in the previous set. This was repeated to produce a total of 33 image points at a given (x, y, z) location, each from a different set of 40 pings. Out of the 33 image points at a given (x, y, z) location, the maximum image value was chosen which became the final image point at that (x, y, z) location.

The images were obtained using the following time-delay beam-forming algorithm operating on the signals at the n^{th} receiver to obtain the image strength at r_i .

$$\text{Image Strength at } \vec{r}_i \equiv \sigma_i(\vec{r}_i) = \frac{1}{N} \sum_{n=1}^N 4\pi |\vec{r}_n - \vec{r}_i| d_n \left(\vec{r}_n \left| \vec{r}_i, \frac{|\vec{r}_n - \vec{r}_i|}{c} \right. \right).$$

where d_n = Inverse Fourier transform of D_n and $D_n(\vec{r}_n | \vec{r}_i, \omega) = \frac{P_n(\vec{r}_n | \vec{r}_i, \omega)}{V(\omega)} = \sigma_i(\vec{r}_i, \omega) \frac{e^{ik|\vec{r}_n - \vec{r}_i|}}{4\pi |\vec{r}_n - \vec{r}_i|}$

Since we are interested in forming images over volumes about the size of a target, we can ignore the variation in the range dependent pre-factor in the image strength expression above.

To obtain a 2-d plot as a function of x and y (plan view image), the maximum image value as a function of z is chosen for a given (x, y) coordinate. Similarly, to obtain a plot as a function of y and z (a depth image), the maximum value along x is picked and the same with the plot as a function of x and z (the orthogonal depth image) where the maximum in the y -direction is used.




Figure 2. Cartoon illustrating the Δz ambiguity for a sloping sediment when the flight paths are laterally far from the target.

In all these previous cases, no attempt was made to integrate into the image processing algorithm the AUV position data available from the Doppler Velocity Log (DVL) and Inertial Measurement Unit (IMU) systems. In our imaging algorithm, the 2D depth image of a target, i.e. one mapped over x and z or y and z , presents the image strength at the target's z coordinates where the latter is relative to the sediment interface z coordinate directly below the AUV. As a result, the depth coordinate for a target displaced some in-plane distance from the receiver (at $x_{\text{lat}}, y_{\text{lat}}$ in Fig. 1) will

be the sum of the actual burial depth, the depression or elevation of a sloping interface if the interface is not level, and the uncertainty (unknown fluctuations) in the AUV altitude. (See Fig.2.)

Regarding subtask 2.4, we have generated new images produced after reconfiguring the imaging algorithm to account for AUV position data thus removing the uncertainty due to unknown sensor z coordinates. We considered both AUV altitude and roll but not pitch which we assumed to have a small effect on the receiver height. We can now assume with the actual 3D AUV locations taken into account that the z-locations displayed in the x,z and y,z depth images are accurate (i.e. not an artifact of unaccounted for sensor altitude variations). The remaining

ambiguity is now associated with the fact that both the actual target burial depth and AUV altitude are referenced with respect to the sediment interface. In particular, if the sediment interface is sloping in the x or y directions (with reference to Fig. 2), there is a Δz between the sediment interface z co-ordinate directly below the receiver versus one at the x,y target position. As a result, we can only determine the *sum* of the burial depth plus Δz . When the imaged target is not too far laterally from the AUV receivers, we can ignore Δz so that the observed burial depth can be considered the actual target burial depth. For a linearly sloping sediment, a rough rule of thumb is that one could ignore Δz (only several centimeters) for lateral receiver-to-target separations less than $\sim 3\text{cm} / \theta_s$ where θ_s is the angle in radians between the sloping sediment interface and the horizontal measured in radians. For example, for a 1° sloping sediment (about the maximum seen from the AUV position data), we can ignore Δz if the receiver-to-target lateral separation is less than $\sim 2\text{m}$.

First, let us consider the new images obtained from the EW flight paths (see Fig. 3) for the buried NRL targets, n1 –n11. These images are shown in Fig. 4a-4c wherein the three images (plan view and two depth views) are shown for each target. All but two of these images were processed from path *h* data except for n9 and n11 which used data from paths *g* and *c*, respectively. With reference to Fig. 3, for targets n1, n2, and n3 the projected plan view distance between the target and the AUV at its closest approach is $\sim 7\text{m}$, 1.5m , and 5m respectively. (Here and in what follows, the AUV target distance that is given refers to the projected distance in the x,y plane.) The depth images of these targets show them to be below the sediment interface by $\sim 10\text{cm}$. Next, these distances for n4, n5, n6, n7, and n8 are 7.5m , 6m , 4.5m , 4m , and 2.5m , respectively. In these cases, the images show them to be below the sediment surface except for n7 and n8 which appear to be at least partially above the interface. n4 and n5 appear $\sim 10\text{cm}$ below the interface and n6 $\sim 20\text{cm}$ which agrees with our burial plan. Finally, for n9, n10, and n11 these distances are 0m , 10m , and

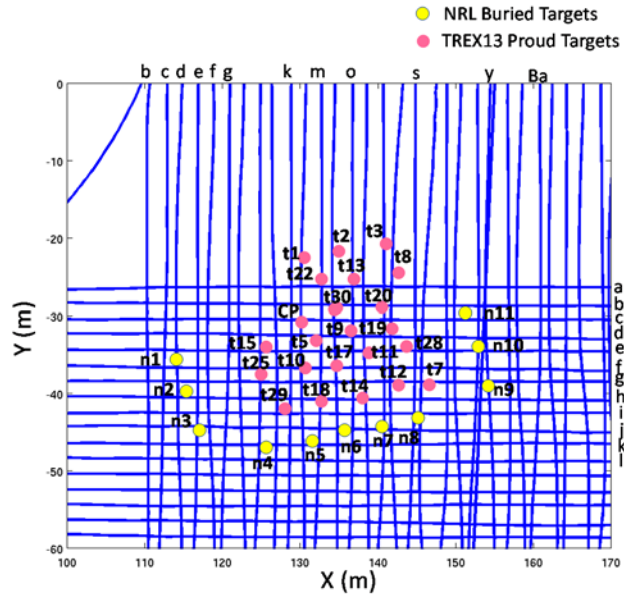


Figure 3. The EW and NS flight paths executed by BOSS over the target field (● NRL buried ● TREX13 proud)

2.5m respectively. These targets appear to be below the sediment surface except for n10 which appears only partially buried.

We make the following specific comments based on the images about each of these targets.

n1h 60° rocket: Plan view image is ~ the same size as the horizontal target which is probably due to a resolution limit; x,z & y,z images now (correctly) below (10cm) the sediment interface; y,z image has a 60° slant suggesting target burial orientation is along y (although large distance from target i.e. ~8m).

n2h 30° rocket: Plan view image indicates ~30° burial along x axis; x,z & y,z images are now (correctly) below (~10cm) the sediment interface.

n3h horizontal rocket: Ignoring double image, size is 50cm by 30 cm versus 45cm by 13cm; target depth images indicate ~10cm burial as planned; x,z & y,z images also seem to suggest the rocket is not quite horizontally buried. We will come back to this image later on page 71.

n4h horizontal 155mm: plan view indicates a 60cm by 30cm size versus 64cm by 16cm. However, depth images indicate near vertical burial and slightly off the x and y directions. n4 is almost 8m away from the *h* flight path. **n5h** 90° rotated horizontal 155mm: these images look similar to those for n4h (as did the old images); target still appears ~ 10cm below the sediment interface as it should. **n6h** horizontal deep burial 155mm: Plan view shows a 70cm by 30cm horizontal target.

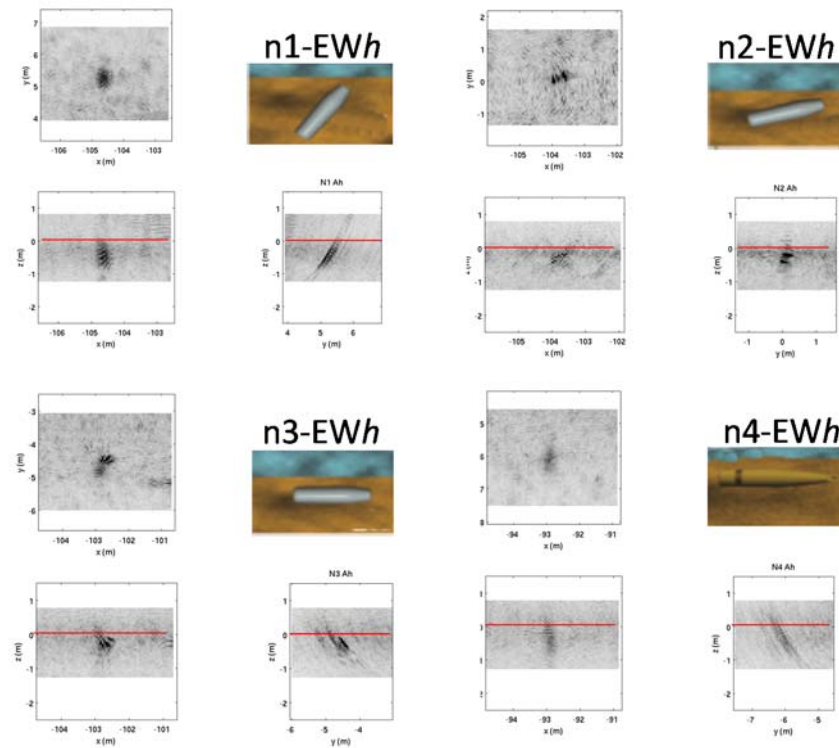


Figure 4a. Plan view and two depth images for n1 through n4 targets using EW flight path *h*.

The x,z depth view also correctly shows a horizontal target and one buried ~20cm. However, the y,z image again shows a slanted target.

n7h 30° 155mm: x,z & y,z images seem to indicate a partially buried target. Even though n7 is closer to path *h*, this partial burial may indicate that the sediment is sloping in the y direction here on the right side of the field.

n8h 60° 155mm: x,z image indicates ~50° burial which would be correct; but both x,z & y,z images show some extension above interface although not as much as in the n7 image. Since this target is closer to path *h*, this may still indicate a sloping sediment.

n9g horizontal 120mm mortar: There is a high background making it hard to interpret all three images. The plan view indicates a horizontal target about 60cm by 40cm; however, this is not at all clear. The z,x and z,y images indicate burial with a large vertical angle. Both of these depth images indicate ~10cm burial as they should since path *g* is directly over the target.

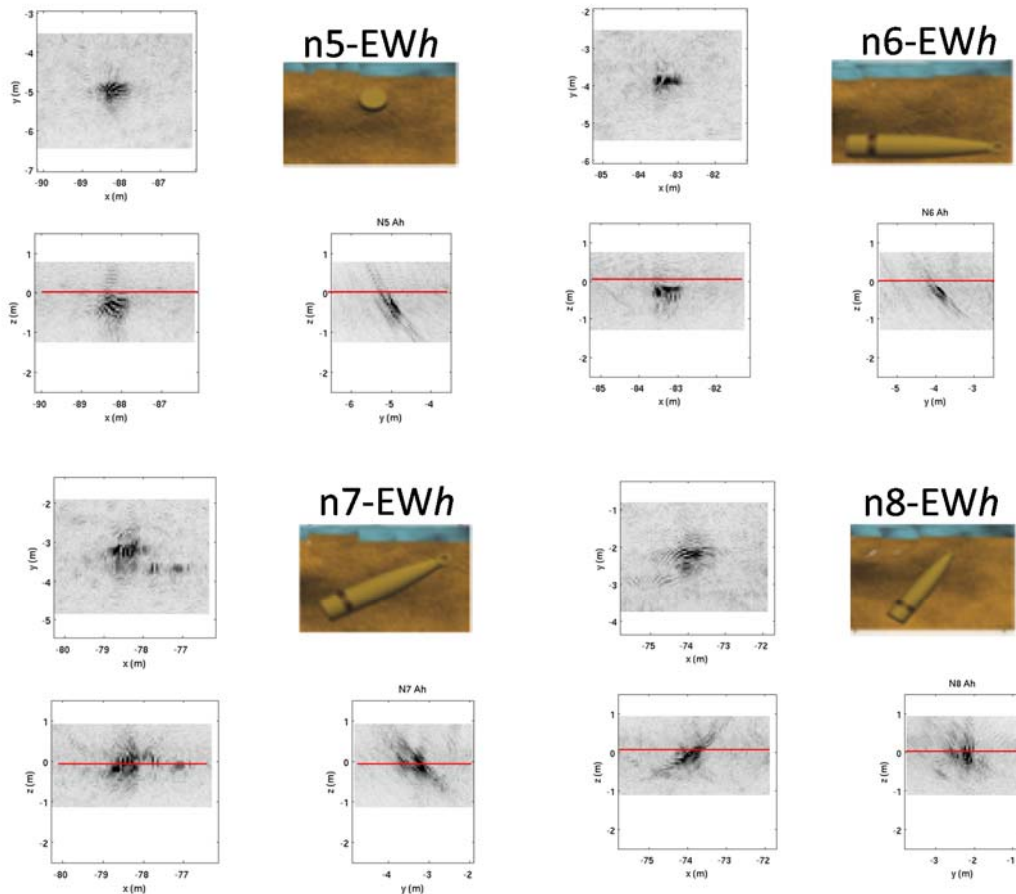


Figure 4b. Plan view and two depth images for n5 through n8 targets using EW flight path *h*.

n10h rock: Unlike the old images, the new depth images indicate a partially buried target as for n7 and n8. However, flight path h is far from the target ($\sim 7\text{m}$) so this may still indicate a sloping interface along y . The x,z and y,z depth images indicate that the plane of the somewhat flat rock is angled near 80° about the x axis, and the x,y plan view image seems to indicate that the somewhat flat rock is buried with its thin axis along the x direction. In point of fact, the three images are very consistent with this, and these images are probably the best to interpret as we discuss later.

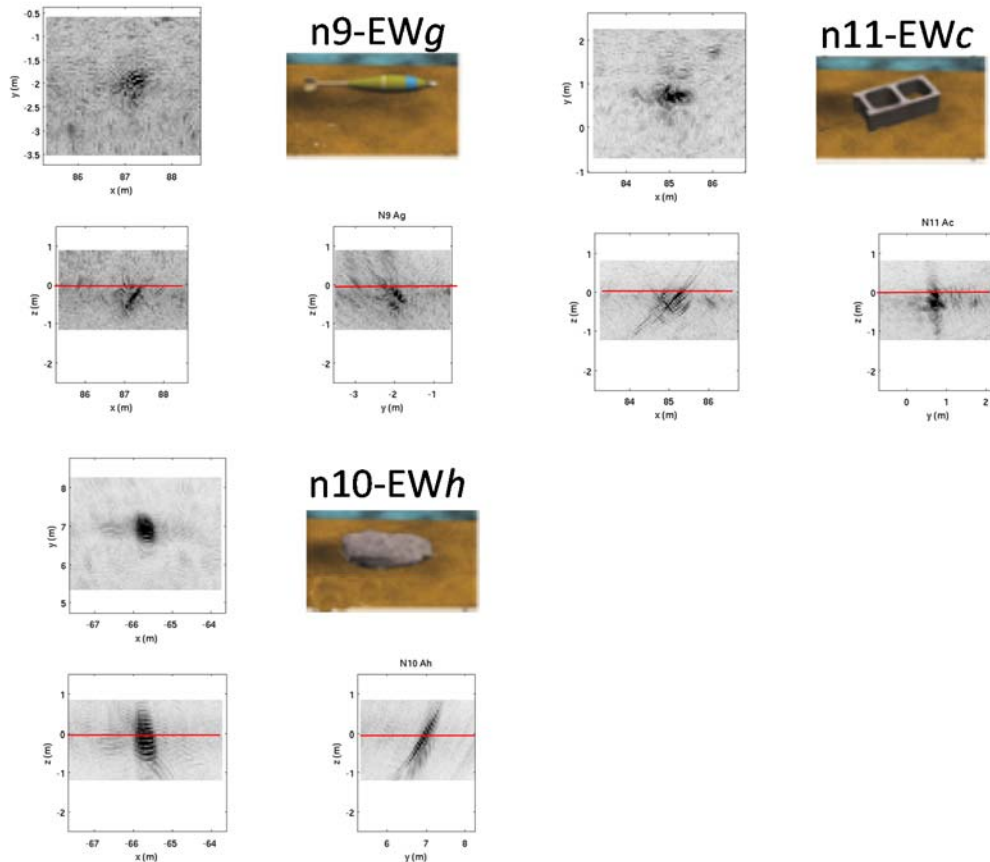


Figure 4c. Plan view and depth images for n9 - n11 targets using EW flight paths g,h , and c , respectively.

n11c cinder block: There are no old images for comparison. In the new images, the depth views indicate a buried target. Taken together, the images suggest that the block is buried horizontally but rolled along its long axis $\sim 80^\circ$ such that the exposed large holes are mostly in the x,z plane and at a steep angle. The plan view image presents one of the block's solid sides but at an angle ($\sim 10^\circ$) sufficient to produce a return only from the long edge. The y,z image shows the other solid side rotated $\sim 10^\circ$ from vertical. This all seems very interpretable as are the rock images.

We now show the images generated from the NS path data in Figs 5a -c.




Figure 5a. The plan view and two depth images for n1 through n4 targets using NS flight paths e and k .

Comparisons of the NS path images with those generated from the EW paths shown earlier seem to indicate that the latter are of higher quality to some degree. This may be due to the fact that in general, the EW flight paths present scattering aspects of the buried targets along their long, cylindrical axes versus circumferentially. Assuming that all the targets are in fact buried as planned, the images processed from the NS flight paths incorrectly show that seven of the eleven are protruding significantly above the sediment surface while only one from the EW paths is incorrect in this regard. Further, the plan view images of the horizontally-buried targets seem to be more reasonable for the EW paths than for the NS paths. In the former case, these plan view images show targets with aspect ratios of about two to one. The actual target aspects are closer to three or four to one, with the smallest dimension we see in these images being about 0.25m. This is consistent with the expected resolution. Given a bandwidth ΔB of 3 kHz to 20 kHz, the limiting (infinite aperture) image resolution ($C/2\Delta B$) would be $\sim 0.04\text{m}$. For a 1m aperture (L_{CT}), the cross-track resolution ($\lambda R/(2L_{CT})$) at the center of our band would be $\sim 0.16\text{m}$ and 0.4m at 2m and 5m ranges, respectively; and at the center of our band, the limiting along-track resolution ($\lambda R/(2L_{SAS})$) would be $\sim 0.08\text{m}$ and 0.18m at 2m and 5m ranges, respectively for a 2m track (L_{SAS}).

Notwithstanding the limited resolution we seem to have achieved, it is interesting to note that the 2-D images for the deepest target burial (20cm for n6 the horizontal 155mm shell) for the EW path h (see Fig. 4b) are perhaps the best of the entire set in that the x-z image correctly shows \sim a 20cm burial and the plan view and x,z images are about the same as they should be.

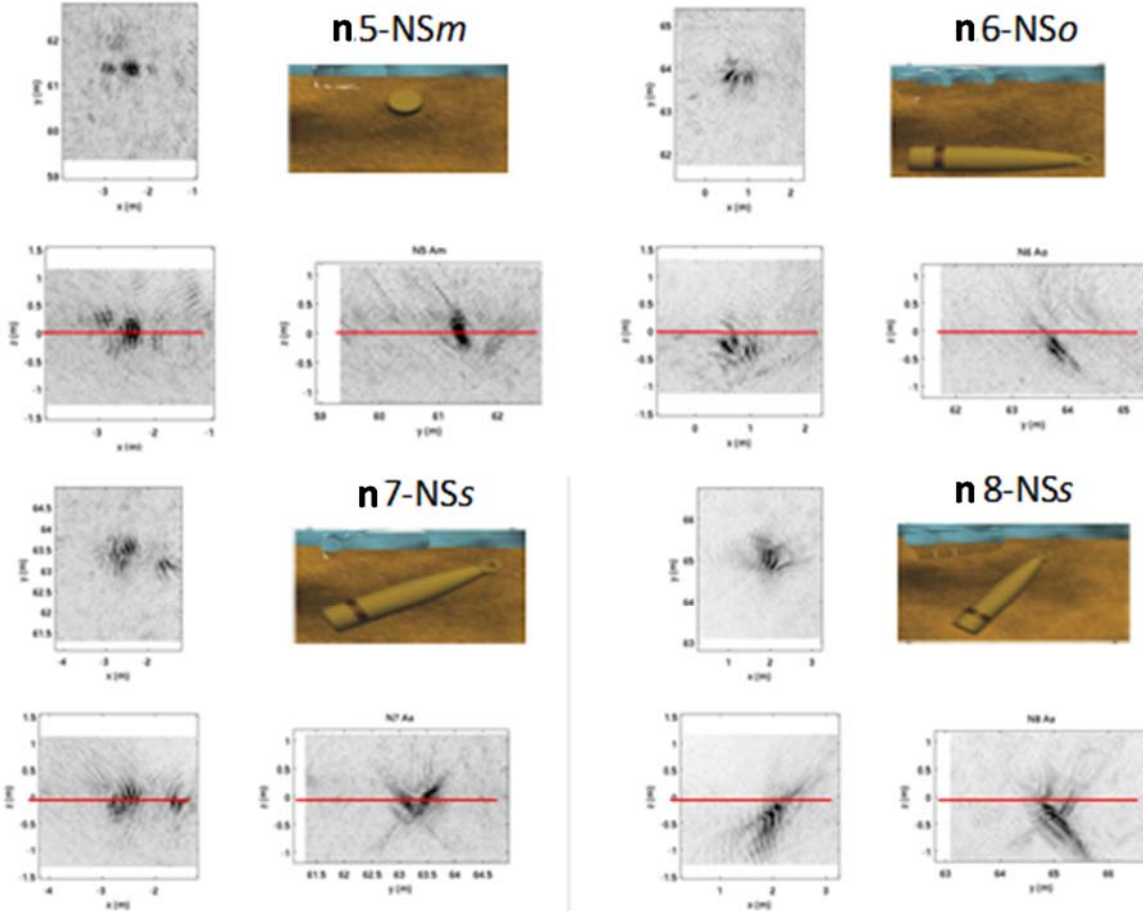


Figure 5b. The plan view and two depth images for n5 through n8 targets using NS flight paths m , o , and s .

Regarding n10, we will consider the associated images in more detail. We show in Fig. 6 a photo of the rock. Although the actual burial orientation is unknown, we assume a specific orientation and then show that the images are consistent with that assumption. As can be seen from Fig. 3, the flight path (h) runs quite a distance from the rock i.e. \sim 10m. As described in the figure, the rock is assumed buried with its long axis (orange arrow) rotated up out of the horizontal plane $\sim 80^\circ$ about the x axis. For both scans, the height above the sediment interface is 3 meters.

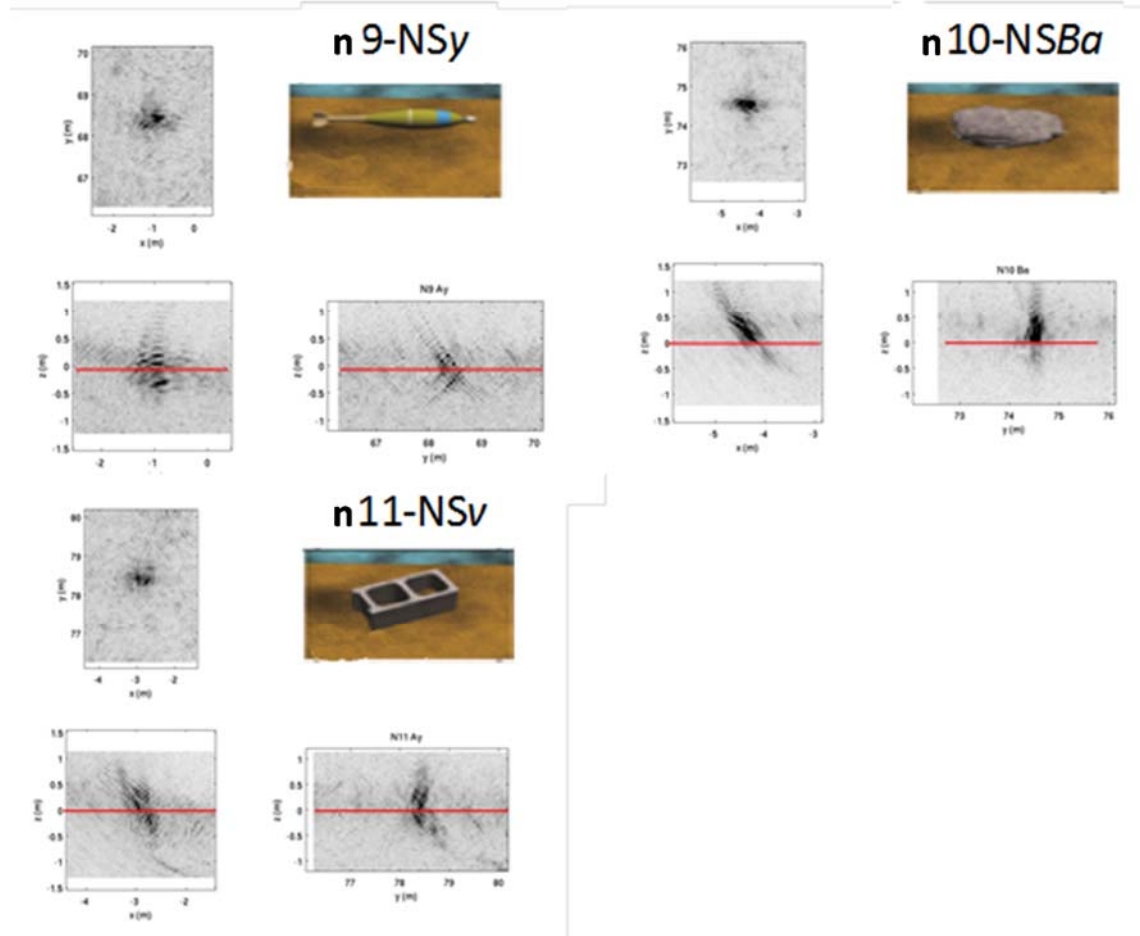


Figure 5c. The plan view and two depth images for n9 through n11 targets using NS flight paths *y*, *Ba*, and *v*.

We show in Fig. 7 the *x* and *y* scan images of the buried rock to facilitate understanding the following discussion. Consider first the *x* scan. The plan view image seems consistent with the rock being buried with its long axis rotated up out of the *x,y* (horizontal) plane by $\sim 80^\circ$ about the *x* axis so that what one sees in the plan view image is roughly the top edge of the rock somewhat smeared and oriented along the *x* axis. Further, with the synthetic scan aperture along the *x* axis, we might see some spatial structure with *x*. This is indeed consistent with the plan view image seen in the *x* scan of Fig. 7. This would then imply that one would see a fairly large area in the *x,z* depth image, and this is indeed the case. Next, considering the detailed shape of the rock (see Fig. 6), in the *x* scan *y,z* image one would again see a thin edge of the rock. And since the smaller *y* coordinates of the rock are deeper, this “thin edge” line would be slanting left to right as indeed it is in the image. Moving to the *y* scan, the plan view image should present roughly the same as in the *x* scan but now with some spatial structure in *y*




Figure 6. Assumed rock orientation with respect to *x*, *y* scans with its long axis (orange arrow) rotated up out of the horizontal plane $\sim 80^\circ$ about the *x* axis.

since the synthetic aperture is along y . This is indeed the case, except that the image is not as long along x as it is in the x scan. In the x,z plane, the smaller rock x coordinates are shallower so that the line should slant from right to left as indeed it does. Finally, in the y,z image, the larger rock y coordinates are shallower so that the line would slant upward from left to right. And this is the case as can be seen in Fig. 7. So although the six images of the rock shown in Fig. 7 at first sight seem inconsistent, we have shown that for one plausible rock burial orientation the images are all consistent.

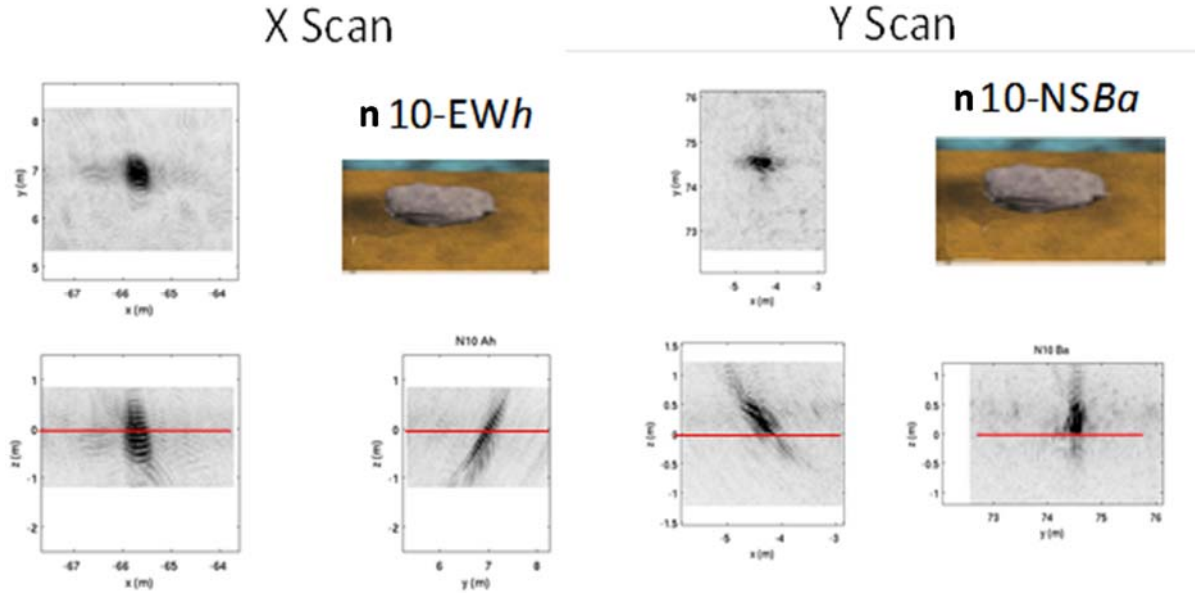


Figure 7. The three 2-D images of the rock (n10) produced from the EW flight path h (left) and from the NS flight path Ba (right).

In summary, inclusion of the AUV position data in the imaging algorithm has improved both the quality of the images and their ability to correctly register the buried targets with respect to the location of the sediment interface. The few cases in which the images indicate a partially buried target are either correct (the target became somewhat unburied over time) or imply that the sediment surface has a small slope in the y direction at the center right of the target field near target n7. The EW path images appear to be better in general than those from the NS paths. This is consistent with the fact that the EW paths sample angles predominately along the cylindrical targets' length versus circumferentially. However, two of the horizontal depth images seem to be at odds with the intended burial angle: n1Hh appears to be oriented along the y versus x axis; and n4Hh appears in the depth images to be vertical rather than horizontal. We do not have an explanation for this. Finally, for the most part the images present targets whose lengths are approximately correct but whose widths are about double what they should be. We take this to be a measure of the resolution limit for our imaging process (data collection, aperture, and imaging algorithm) which would be about 0.25m. At the highest frequency in our band (20 kHz), this is approximately three acoustic wavelengths. From a theoretical point of view, we would expect the following limiting resolutions. Given the 17 kHz bandwidth ΔB , the limiting range resolution ($C/2\Delta B$) would be about 0.04m; and at the center of our band, the limiting along-track resolution ($\lambda R/(2L_{SAS})$) would be about 0.16m and 0.4m at 2meter and 5 meter ranges, respectively. The

across-track resolution ($\lambda R/(2L_{CT})$) is 0.16m and .40m at 2m and 5m ranges respectively. As in previous BOSS target measurements, the multi-aspect imaging procedures used for BOSS cause image distortion for some target aspects that prevent accurate registration of target dimensions and shape. Further, the conventional SAS method of sweeping the synthetic aperture along track frequently does not allow specular illumination of the UXO targets at aspects that generate echo levels with adequate SNR for imaging. Given these facts, the previously quoted average along-track resolution of 0.25m is not surprising.

SUBTASK 2.5. Include North/South AUV multiple path data in target separation study.

As part of this task, we considered whether the feature separation was better for NS or EW paths. We did not include diagonal paths in this study since the latter were only implemented at 5m altitudes. In the following, we trained *discriminatively*¹⁸ using target data from n1 through n9 (the buried, filled NRL UXO) and false target data from n11 and n12 (the buried cinder block and rock) and proud targets t2, t3, t5, CP, t15, t18, and t19 (the

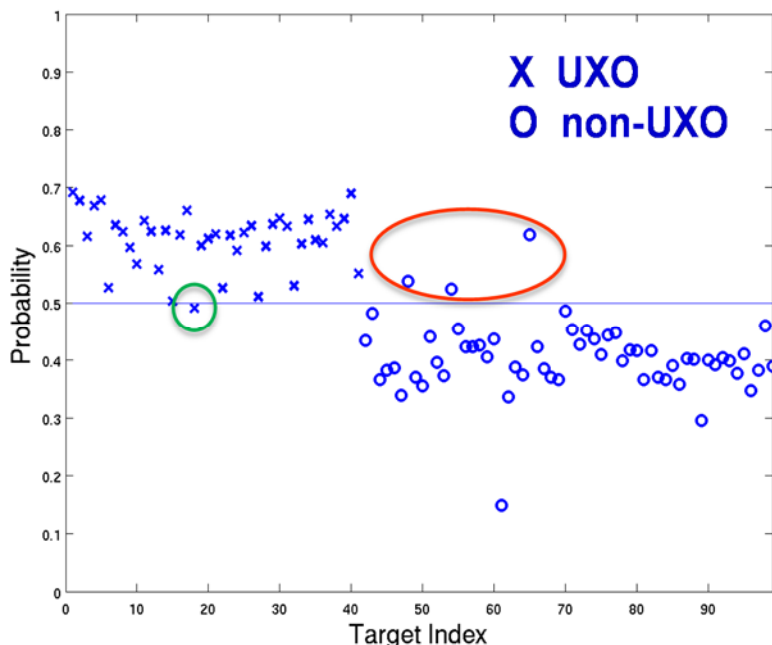


Figure 8a. Discriminatively-trained RVM classification algorithm trained on even numbered source pings and tested on odd numbered pings. Targets include 9 buried UXO and 9 non-UXO – 2 buried and 7 proud. NS paths (3-20 kHz).

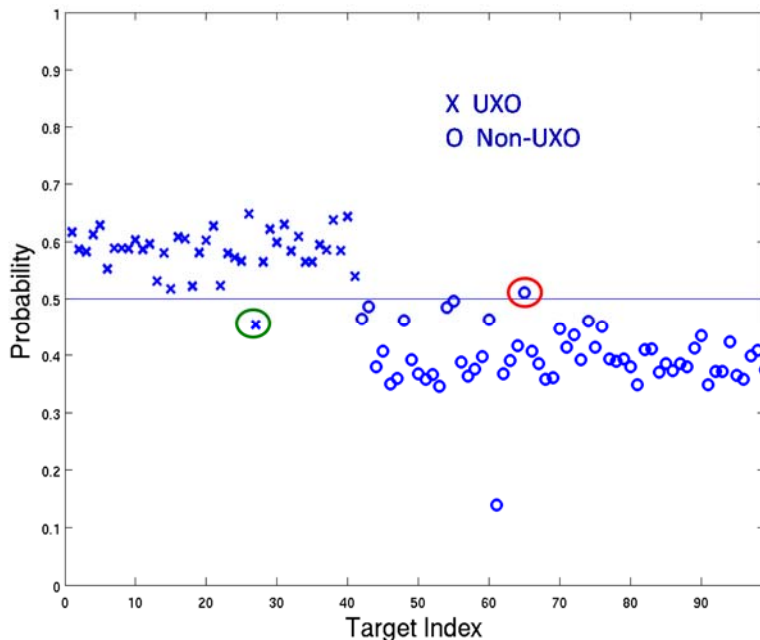


Figure 8b. Discriminatively-trained RVM classification algorithm trained on even numbered source pings and tested on odd numbered pings. Targets include 9 buried UXO and 9 non-UXO – 2 buried and 7 proud. NS paths (3-13.3 kHz).

rock, 55 gallon filled drum, 5:1 aspect pole section, the panel, 2:1 aspect pole section, cement block, and the tire). The training was done using data from odd numbered source pings and tested on data from even numbered pings versus the original procedure which trained on even numbered receivers and tested on odd. Both of these are considered a “sanity check” in the sense that if the chosen features do not separate under these conditions, the features themselves are not useful for target classification.

The resulting combinatorial probability using the north/south paths is shown in Fig. 8a for the bandwidth 3 – 20 kHz. This is obtained by combining the probabilities over the 40 y positions (receivers) i.e. taking the product of the probabilities at each y raised to the $1/40$ power. Here the target index is sequenced each time a target is seen in a path. (We will see later that the maximum target index for UXO was smaller for NS paths than for EW paths.) The one false negative corresponds to target n5 and path n which is the 155mm shell buried horizontally perpendicular to the x direction. Path n is about 3m to the right of this target. The three false positives correspond to the following targets: n11 the buried cinder block for path w which goes about 1m to its right; t5 the proud 5:1 aspect pole section for path m which goes right over the target; and t15 the proud 2:1 aspect pole section for path m which also goes nearly directly over the target. Considering that there were on average 5 realizations for each target (five different paths), the four incorrect calls above represent only about 4% of the 96 realizations.

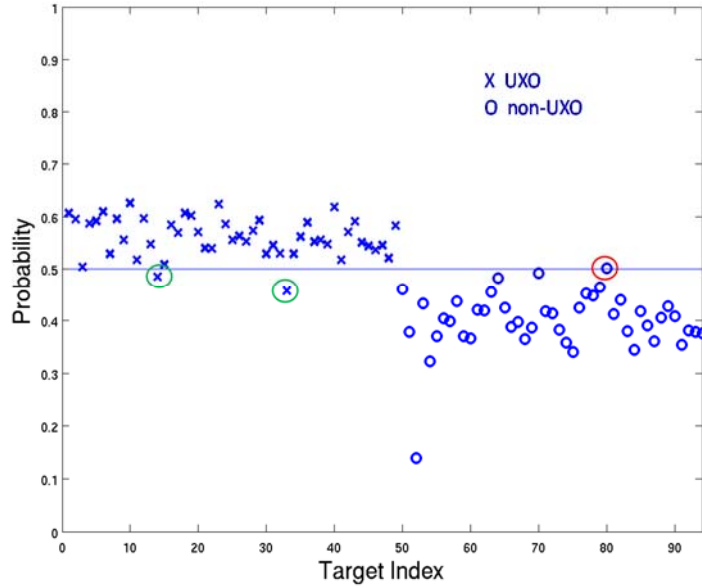


Figure 9. Discriminatively-trained RVM classification algorithm trained on even numbered source pings and tested on odd numbered pings. Targets include 9 buried UXO and 8 non-UXO – 2 buried and 6 proud. EW paths (3 – 13.3 kHz).

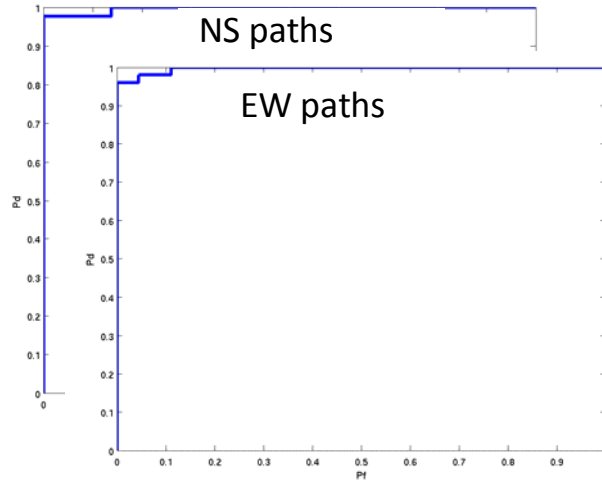


Figure 10. ROC curves for the results in Fig 9.

Next, we repeated this identical study but now with the bandwidth consistent with the Nyquist sampling rate, i.e. 3 -13.3 kHz. For this next figure as well as all those following, the bandwidth will be 3 – 13.3 kHz. The new results for the NS paths are shown in Fig 8b. As can be seen, while there is still one false negative, there is only one vice three false positives. We conclude that there is some merit in limiting the bandwidth, although the impact of not doing so is minimal. The one false negative corresponds to target n7 and path n which is the 155mm shell buried nose-up 30 degrees. Path n is about 7m to the right of this target. The false positive corresponds to t15 and path m which is the 2:1 aspect pole section. Path m goes about 8m to the right of the target.

Next we show the results for the east/west paths. In this case, we had no unambiguous image of t3, so we could not include this target. Again we use the combinatorial probability as defined above, and the results are shown in Fig. 9. Overall, the results appear comparable in effectiveness to those shown above for the NS paths. In the horizontal case, the two false negatives are as follows: n4 which is the 155mm shell buried horizontally along the x direction and path h which is about 7m above the target; and n7 which is the 155mm shell buried nose-up 30 degrees and path g which is about 6m above the target. The one false target is t18 which is the proud cement block and path m which is 3m above the target.

While the results appear comparable in general to those for the NS paths, there are differences including the following. As can be seen in Fig. 8b, there are many UXO target entries which have probabilities above 0.6 whereas in Fig. 9 there are much fewer. However, the distribution of the false target probabilities is about the same. The ROC curves for each case are shown in Fig. 10. These also indicate a slightly better performance (25% area above the curves) for the NS paths. To the extent that this is meaningful, this improved performance may be related to the fact that for these paths the receivers are moving perpendicular to the cylinder axis for most of the UXO targets. We note that this is




Figure 11. Discriminatively-trained RVM classification algorithm trained on even numbered source pings and tested on odd numbered pings. Targets include 9 buried UXO and 9 non-UXO – 2 buried and 7 proud. NS and EW paths (3-13.3 kHz).




Figure 12. ROC curve associated with the results in Fig 11.

opposite to the case for the images where moving along the target's long axis produced better images.

Next we combined the data from the EW and NS paths. These results are shown for the combinatorial probability in Fig. 11. As can be seen, rather than improving the separation, combining the data from the two paths actually resulted in a small degradation in performance. The three false negatives are as follows: n1 which is the five inch rocket buried nose up at sixty degrees and EW path *h* which is about seven meters below the target; n7 which is the 155mm shell buried nose-up thirty degrees and EW path *g* which about seven meters above the target; and n7 and NS path *n* which is about eight meters to the left of the target. The five false positives are as follows: t18 and EW path *i* which is about 2 meters below the target; n10 and NS path *w* which is right on the target; n11 and NS path *w* which is about 1.5 meters to the right of the target; t5 and NS path *m* which is about 0.5 meters to the right of the target; and t15 and NS path *m* which is almost 9 meters to the right of the target. The corresponding ROC curve is shown in Fig. 12 where it can be seen that the separation performance has indeed been degraded.

In summary, we find that the performance of the RVM target feature separation algorithm is slightly better for the NS paths than for the EW paths. For the most part, the NS paths sample the cylinder scattering circumferentially rather than along its length. This is opposite to what we found for the imaging algorithm where images appeared better for the EW paths. In addition, limiting the bandwidth according to the Nyquist sampling criteria improved the performance. Further, using both the NS and EW path data in the training and testing actually degraded the performance. With regard to the percentage of incorrect calls, we found the following: NS paths – 2%; EW paths – 3%; combining paths – 4%. Why this is the case is not currently understood.

SUBTASK 2.6. Include all additional proud targets (+16) in the RVM target separation study

Next, we included all 33 targets both proud and buried. In the list of UXOs, we included the 9 buried, epoxy-filled UXOs (n1 – n9) and the 10 proud UXOs (t8, t9, t11, t12, t20, t22, t25, t28, t29, and t30). In the list of non-UXOs, we included the 2 buried targets (n10 and n11) and the




Figure 13. Discriminatively-trained RVM classification algorithm trained/tested on even/odd numbered source pings, respectively. Targets: 9 buried and 10 proud UXO and 14 non-UXO – 2 buried and 12 proud. NS paths (3-13.3 kHz).

12 proud non UXOs (t1, t2, t3, t5, t7, CP, t13, t14, t15, t17, t18, and t19). Again, we trained on even numbered source pings and tested on odd. The results for the combinatorial probability including the ROC curve are shown in Fig. 13. Given the large number of targets and their many realizations (~ 5 paths for each) and the mixture of both proud and buried targets, we consider this separation result to be very encouraging.

The two false negatives are n7 which is the 155mm shell buried nose-up thirty degrees and NS path n and t29 which is the small bullet #2 proud on the sediment surface and NS path n . The three false positives are t18 which is the cement block proud on the sediment surface and NS path n , t19 which is the tire proud on the sediment surface and NS path n , and t14 which is the scuba tank with water and a stem and NS path s .

In summary, in this study the chosen UXO targets included those that were both buried and proud and those that were epoxy-filled, air-filled, and one that was water-filled. Except for their internal states, they were all UXOs. The false targets included many non-UXO targets that were both proud and buried. The results above demonstrate that the multi-dimensional feature extracted from the acoustic color maps can indeed separate the detections from these two groups. In fact, the five miss-calls represent less than 3% of the total calls. We would expect that separation of these targets into two slightly different groups than the one described above would still offer reasonable target separation.

SUBTASK 2.7. Analyze acoustic color from non-epoxy-filled UXO (proud) compared to epoxy-filled (buried) UXO.

The set of acoustic color maps generated from the EW paths for eight of the 155mm targets were analyzed. The targets included the epoxy-filled buried targets n4 – n8 and the proud targets t8, t9, and t28. We assume that the latter three targets are empty, but we have not been able to clarify the details of their internal state at this time. For this study in particular, we inputted the correct range factor in constructing the acoustic color plots in order to have range-corrected target strength numbers for the comparisons. In particular, the target strength formed from the pressure measured at the receiver at r_n from the scattered echo from the target centered at r_c is given by⁴:

$$\begin{aligned} \text{Target Strength at } \vec{r}_n &= 10 \log |X_n(\omega)|^2 \\ X_n(\omega) &= D_n(\omega) \frac{|\vec{r}_n - \vec{r}_c|}{e^{ik|\vec{r}_n - \vec{r}_c|}} \\ \text{where } D_n(\vec{r}_n | \vec{r}_i, \omega) &= \frac{P_n(\vec{r}_n | \vec{r}_i, \omega)}{P_i(\omega)} \end{aligned}$$

We examined the acoustic color maps visually for the presence of a frequency/x position (aspect angle) feature that we had seen earlier in laboratory measurements in our free-field⁴ and sediment pool facilities⁵ and which we had subsequently associated with elastic waves in the epoxy filler material. We considered fifty color maps associated with the eight targets. Thirty of these acoustic color maps were deemed “good” in the sense that presence or non-presence of the specific frequency/angle feature seemed definitive. Using these thirty “good” color maps, 90% were

correct calls and only 10% incorrect assuming that the proud targets were indeed empty. Of the twelve flight paths used in this study, paths *h* and *i* had no associated “wrong” calls while path *l* had the most, i.e. five. Regarding the twenty “ambivalent” color maps, not surprisingly the calls with them split almost equally i.e. 45% correct and 55% incorrect. We believe that these results would improve if we were able to de-convolve the correct incident waveform from the scattered signals (see Subtask 2.8).

Next we trained the RVM algorithm *discriminatively*¹⁸ on even numbered source pings and tested on odd numbered pings using the acoustic color maps. On average, we used the epoxy-filled 155mm shell data collected on seven paths and the empty 155mm shell data collected on five paths.




Figure 14a. RVM discriminative algorithm trained and tested on even and odd numbered source pings, respectively using the acoustic color maps. On average, we used epoxy-filled 155mm shell data from 7 EW paths and empty 155mm shell data from 5 EW paths.




Figure 14b. RVM discriminative algorithm trained and tested on even and odd numbered source pings, respectively using acoustic color maps. On average, we used epoxy-filled 155mm shell data from 7 EW paths and empty 155mm shell data from 7 EW paths.

The results for the combinatorial probability of being an epoxy-filled 155mm shell are shown in Fig. 14a. As can be seen, there is perfect separation between the two target sets. After we had carried out this target separation study, we learned that the two targets called “bullet #1” (t25) and “bullet #2” (t29) are also 155mm shells. Further, t25 is air filled while t29 is water filled. We added the data from these two targets to the false target data and retrained and retested. The results are shown in Fig. 14b. As can be seen, the epoxy-filled 155mm shells still separate from the empty and water-filled shells.

In summary, the fact that the multi-dimensional feature extracted from the acoustic color maps can be used in principle to separate the epoxy-filled shells from the others is encouraging. However, we point out that these two target sets differ not only because of the epoxy filler material or the lack thereof, but also by the fact that one target set is buried while the other is proud. In this regard, future measurements ought to include epoxy-filled and empty UXO that are all buried.

SUBTASK 2.8. Perform acoustic BOSS source calibration in the NRL near-field holography facility.

The necessary facility preparations were made for executing the source calibration measurements for BOSS in the NRL Laboratory for Structural Acoustics. We took delivery of BOSS June 9, 2014 and successfully mounted the AUV-based system in our Building 5 structural acoustic pool facility (see Fig. 15) without any problems. Measurements began on June 12, 2014. Given the time allotted for BOSS at NRL (two weeks), we did not attempt a full holographic scan of the source/vehicle system. Rather, we designed a more limited data collection plan which gave us radiated pressure (phase and amplitude) at 3.5 meter distances over five complete rings in steps of one/half degree. For four of the rings, the vehicle was hung vertically in a manner permitting 360 degree vehicle rotation in steps of 1/2 degree. As shown on the left of Fig. 16, the monitoring hydrophone was positioned 3.4 meters from the vehicle at four different vertical locations (near the nose, adjacent to the projector, near the wings, and near the tail section). Data was collected at each of these positions over the full 360 degrees. A final fifth ring configuration was executed in which the vehicle was hung horizontally and the receiver rotated 360 degrees in the horizontal plane. (See the right side of Fig. 16.) Measurements ended on June 18, 2014, and BOSS was shipped back to Panama City, FL on June 23, 2014. We show in Fig. 17 the frequency-angle plots for the response measured on each of the five hydrophones in Fig. 17. We also show in Fig. 17 (bottom right) the response that would be seen if the source were indeed a monopole (omnidirectional) and it were excited by the chirp electrical waveform used in our experiments. To obtain this response map, we used the TVR curve provided by the source manufacturer for the actual BOSS source, a ITC-1007



Figure 15. The NRL Structural Acoustic Pool facility.

omnidirectional 6 inch diameter zirconate titanate transducer.




Figure 16. The two orientations for hanging BOSS in the NRL Structural Acoustic Pool facility.

Clearly, the coupled source/AUV system does not radiate as a monopole. Unlike the isolated spherically radiating source, at any particular frequency the acoustic pressure fluctuates considerably with direction. Most post mission processing analyses (PMA) of BOSS data to date assume an incident signal independent of angle, and this assumption impacts the various constructs obtained from the BOSS data collection flights. This would include imaging (mainly due to source drop-outs) but more seriously acoustic color since this produces modulation in frequency and space, although what one does with each of these constructs in terms of target classification will determine the ultimate impact.

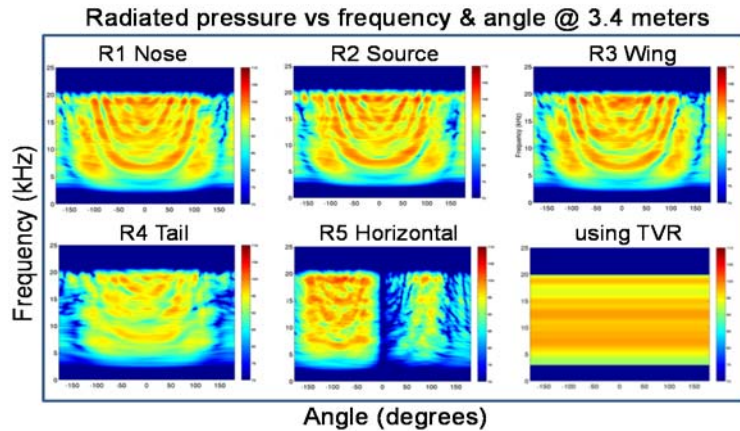


Figure 17. Source pressure levels in dB measured at the 5 receivers 3.4 meters from BOSS as it is rotated in steps of 0.5 degrees. Lower right: what the pattern would be were the BOSS source omnidirectional.

characteristics at the scattering target, we identified two overall approaches. One approach involves the development of a structural acoustic model describing the angle-frequency characteristics that is validated by the five ring data we have collected. A second approach involves a number of ring data projection algorithms. The models and/or the projection algorithms would then be used to obtain a close approximation to the actual source characteristics that had existed on the sediment. These source characteristics would then be integrated into the algorithms we use to obtain the various constructs from the BOSS receiver data.

SUBTASK 2.9 Using new BOSS source calibration data, generate and analyze specular-filtered plan view images of the buried UXO.

The data obtained from the BOSS calibration measurements (Subtask 2.8) was processed and analyzed. The measurements had been made in the NRL Building 5 structural acoustic pool facility and involved the following five configurations. The vehicle was hung vertically in a manner permitting 360 degree vehicle rotation in steps of 1 degree. The monitoring hydrophone was positioned 3.4 meters from the vehicle at four different vertical locations (near the nose, adjacent to the projector, near the wings, and near the tail section). Data was collected at each of these

We show in Figs 18 and 19 the time waveforms and frequency response of each of the four rings at 0 degrees. For comparison, we show the same signal and spectrum assuming the source is omnidirectional and excited by our particular chirped waveform in the lower left hand side in each figure.

In considering how one might use the above limited BOSS calibration data in order to intuit more accurate incident sound spatial and temporal

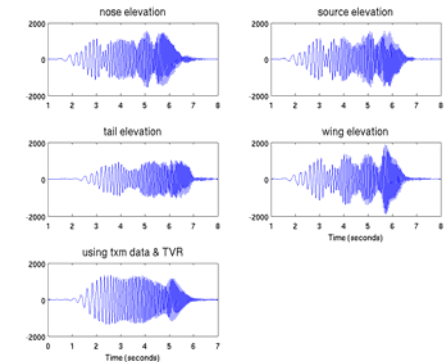


Figure 18. The measured time waveforms at the R1 through R4 receivers and the prediction from TVR curve.

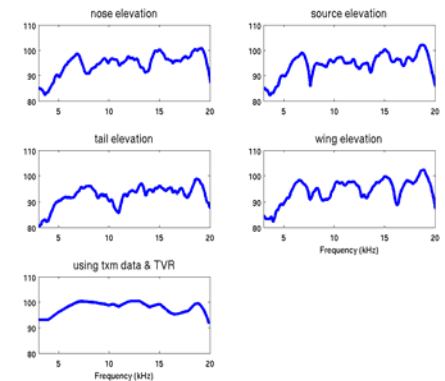


Figure 19. The measured signal spectra measured at the R1 through R4 receivers and the prediction from TVR curve.

positions over the full 360 degrees (the circumferential data). A final fifth configuration was executed in which the vehicle was hung horizontally and the receiver rotated 360 degrees in the horizontal plane. The first four rings of data viz. the circumferential data spaced along the AUV length clearly show time waveform variations from what is expected from the isolated source confirming our earlier suspicion. Further, there is strong azimuthal pseudo-periodic structure in the source radiation pattern demonstrating that the source as mounted on the AUV is far from an omni-directional radiator. This means that assuming an incident pressure waveform on the sediment interface below

that is (1) a simple convolution of the electrical waveform with the isolated source transfer function and (2) spatially uniform is not justified. As far as we know, these assumptions have been made by all previous BOSS users including ourselves up to now. This complication impacts all the post measurement analysis (PMA) results, particularly those involving large angular apertures and/or range variations in the data collection. This task focused on image improvement was originally planned recognizing problem (1) above but not problem (2). Redoing the images in a meaningful manner will now first require the resolution of both issues (1) and (2). Addressing the latter issue will require a significant study to develop at least an approximate approach for obtaining the more complicated spatially dependent pressure waveforms incident on the sediment interface as a function of range and angle. We devoted the remaining resources allocated to Subtask 2.9 to making progress on this issue.

To this end, we completed the generation of supersonic intensity images¹⁹ to identify the actual radiating source spatial distribution as it exists over the surface of BOSS for the acoustically coupled source/vehicle over the length of the AUV. With reference to Fig. 20, the supersonic imaging algorithm constructs the radiating source strength distribution as a function of z , ϕ , and frequency given the measured source radiation pattern given in Fig. 17 of Subtask 2.8 and marked R5 Horizontal.

The resulting supersonic images are shown in Fig. 21 for the two halves of the cylindrical surface, one labeled bottom side (the half containing the source) and the other labeled top side (the half area opposite the source). Immediately we see two dominant characteristics: (1) the radiation is coming from the 25% or less of the surface that includes the source; and (2) the bottom side source distribution is two to three times wider than that for the bottom side.

Examination of the data shown in Fig. 17 of the Subtask 2.8 section for the four circumferential




Figure 20. Geometry describing the mapping between the measured source radiation pattern and the radiating source strength spatial distribution.

rings indicates that the basic frequency/angle structure is about the same for all four rings even though the levels are much lower for the tail ring. Further, the frequency structure shows the frequency modulation for zero angle has a periodicity Δf of about 3 kHz. If this were associated with the resonances of a rectangular cavity, one would estimate the width of the cavity to be $T = C/(2\Delta f)$ or about 10 inches if the cavity were water filled. Here C is the speed of sound. This is close to the diameter of the BOSS AUV which is 12.8 inches.

These three findings – (1) the supersonic image somewhat localized to the cavity containing the source; (2) the frequency/angle structure being the same along the length of the AUV; and (3) the periodicity in frequency at zero angle being consistent with resonances of a cavity about the dimension of the AUV diameter - led us to conclude that the important structural acoustic interaction to consider is that between the spherical source, its mounting structure, and the associated *cylindrical* section of the AUV. The next step is to develop a relatively simple model to predict our measured results and then to integrate this into our PMA of current and future BOSS data.

SUBTASK 5.1 Develop methodology for integrating BOSS source calibration data into the data processing and analysis.

The development of methods to incorporate the limited BOSS source calibration data²⁰ obtained as described in Subtasks 2.8 and 2.9 focused on four approaches. These are:

- (1) Structural Acoustics Model. Develop a structural acoustic model which could be used to determine the source pressure phase and amplitude at any point on the sediment surface.
- (2) Spatial Pattern Projection. Use simple projection of the amplitude and phase values measured on the calibration source ring at each angle along a line that meets the sediment surface at a point a distance r away and adjusting the phase according to $\exp(ikr)$ and the amplitude by $1/r$.
- (3) Synthetic Hologram. Create a spherical harmonic-based hologram generated from the ring calibration data that can be used to project the field outward to any point.
- (4) Equivalent Source Method. Generate an equivalent source distribution determined by matching the ring responses using SVD which can be easily projected onto the plane below.

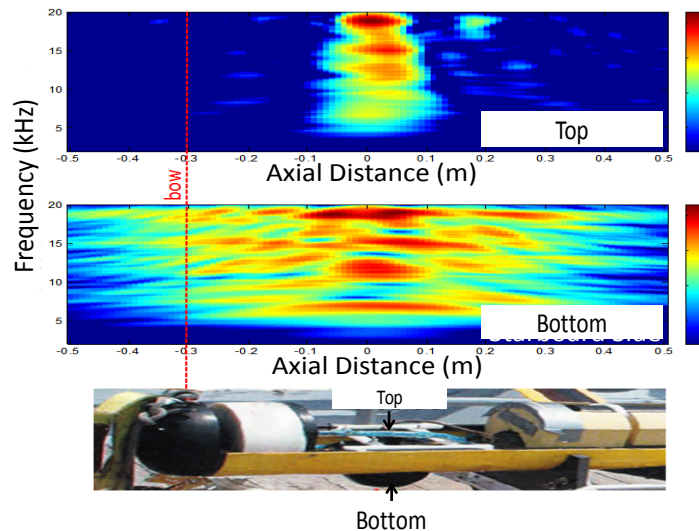


Figure 21. The supersonic images mapped onto the BOSS AUV upper half and the BOSS AUV lower half. The photo indicates the actual locations on BOSS.

With respect to (1), we attempted to model the structural acoustic mechanisms responsible for the radiation patterns as measured at the calibration rings. Basically, this implies developing a reasonable description of the nearfield scattering from the AUV and adding this coherently to the direct radiation from the spherical source. This approach was pursued taking into account two parts of the AUV structure which were considered to be dominant: (1) the two circular walls defining the compartment which housed the source; and (2) the cylindrical AUV body itself.

Shown in Fig.22 is the geometry used in tackling the first structure. One can see the position of the real, spherical source in green and the two image sources in red used to model the effect of the two compartment discontinuities. The resulting patterns from this and several similar geometries




Figure 22. Geometry used to model a spherical source located within the AUV compartment represented with image sources.

were computed. None of them were able to produce anything similar to the complex patterns associated with the calibration ring data. Basically, there was not enough interference to produce the necessary phase variations. An example of the patterns obtained using this model is shown in Fig. 23 where one can see very little resemblance between this pattern and those from the calibration rings.




Figure 23. An example of the radiation pattern versus frequency obtained with the model illustrated in Fig. 22.

As an extension of this model, we considered the radiation produced with the source between two rigid side disks using a finite sphere source equation, and this failed to produce interesting results. A model with the source between two rigid side disks *and* two image sources laterally on the other side of each of the two disks failed to produce interesting results as well.

Next, we considered the source plus an image source centered on the AUV circumference immediately opposite (above) the real source. Although we found that the radiation pattern has some of the calibration ring data characteristics, unlike the latter it is perfectly symmetric top-to-bottom. Obviously, the scattering for the AUV shell must be included.

We then used an equivalent source method in which a spatial distribution of source strengths within the actual spherical source were specified in order to produce a uniform velocity at the spherical surface of the source. This source distribution was then used to excite a finite rigid cylinder the size of the BOSS AUV together with a rigid plate at the position of the source mounting plate. The radiation pattern was then calculated at positions along an AUV radial ring used in the BOSS source calibration exercises. Very little of the characteristic frequency-angle structure was seen.

In a more sophisticated model we included scattering from the cylindrical AUV body itself. In particular, we added the direct source radiation to the scattering from an infinite cylindrical shell with a 12.75 inch diameter using a formal model which could apply to both a rigid AUV shell as well as one with the elastic parameters of the actual Bluefin shell material. We point out that although we are using an infinite shell model for the AUV, scattering contributions from shell sections far from the source are very small given that the source is localized and close to the shell. In addition, the model allowed placing the source at different positions along a vertical line through the center of the source compartment. The results are shown in Fig. 24. The best result so far comes from a model in which the source is artificially displaced 0.16m below its actual position in the

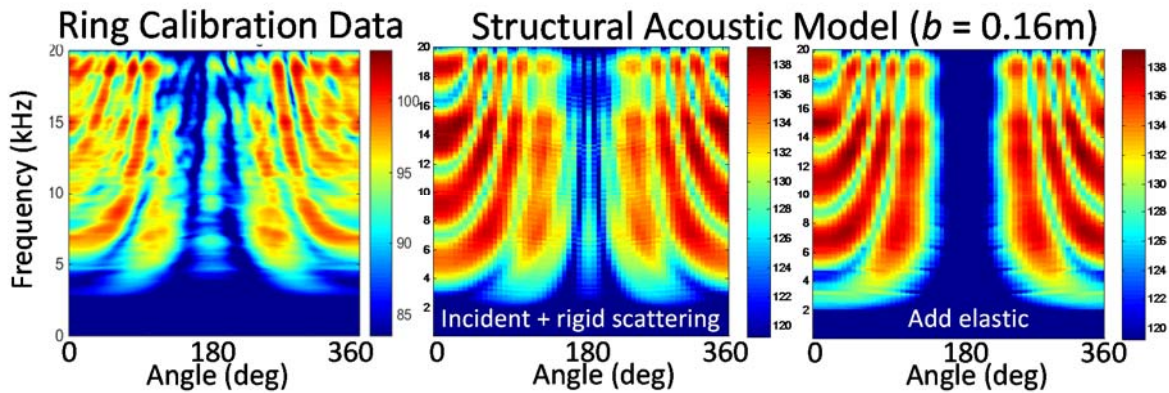


Figure 24. Radiation patterns predicted using a model which includes direct radiation from a vertically displaced spherical source plus scattering from a 12.75inch diameter infinite cylindrical shell. (Left) calibration ring data; (Middle) model with rigid cylinder; (Right) model with elastic cylinder.

real system. When the incident field is added to that scattered from a *rigid* version of the AUV, the frequency-angle contours are close to those seen in the calibration ring measurements although they are somewhat thicker in the model. The dynamic range of the highlights is about the same for both the model and the calibration ring data. While the patterns do indeed resemble those from the ring calibration data, we do not understand the rationale behind the artificial source offset which we have introduced. Perhaps the offset compensates for the fact that in this model we are ignoring the two axial discontinuities which define the source compartment. Further, we expected the results to improve with the addition of *elastic* effects associated with the AUV shell, and this was not the case as seen in Fig. 23. Most importantly, the degree to which the model predicts the phase and amplitude of the measured ring data was not deemed sufficient for us to move forward with this model as a predictor of the actual source pattern on the sediment surface for the BOSS Gulf data. However, should it become important to have a good BOSS source model at a later time, we believe this model could be improved.

Consider next (2) Spatial Pattern Projection. Here we use the amplitude/phase values measured on the calibration source ring at each point and simply project each one along a line from the ring center through that point to a point a distance r away (at an angle θ) adjusting the phase according to $\exp(ikr)$ and the amplitude by $1/r$. The geometry is illustrated in Fig. 25.




Figure 25. Diagram for simple calibration ring projection method.

projection using a multipole expansion of the pressure with terms containing Hankel functions h_n . We assume that a finite number of Fourier coefficients are needed to project the data from the 5 rings to the plane below the BOSS vehicle. That this is reasonable is suggested by the fact that the circumferential transforms of the ring calibration data cut off at about $n=15$ at 20 kHz; and even though the data is severely aliased with respect to the acoustic wavelength at 20 kHz ($ka = 255$) requiring harmonic orders of $n = 255$, the observed maximum of $n = 15$ based on examination of the ring data should remove any potential aliasing issue. The multi-pole expansion is of the form:




Figure 26. Geometry for the synthetic holography projection technique. Radial calibration rings in blue, axial calibration ring in red, and spherical hologram surface in green.

We carried out this simple projection for radial, axial, and combined radial/axial ring calibration (where in the combined calibration we averaged the results at each point from the radial and axial rings). These projected results looked reasonable from the perspective of our understanding of the structural acoustics of the situation, although we were unable to determine the accuracy achieved nor which ring calibration produced the most accurate projection. This projection method is one of the techniques we subsequently used to re-process the gulf scattering data.

Next, we developed formalism based on what we call (3) Synthetic Holography. In this approach, we formulate a spherical harmonic

$$p(r, \theta, \varphi) = \sum_{n,m} [h_n(kr) / h_n(kr_h)] P_{nm}(r_h) Y_{nm}(\theta, \varphi)$$

where the Y_{nm} are spherical harmonics. In the multipole expansion, the $P_{nm}(r_h)$ coefficients represent the unknown Fourier coefficients in the sum describing the pressure measured on a hologram surface located on the sphere at $r = r_h$. They are determined by the pressure p provided on the 5 rings of data, 4 at fixed values of θ with $-\pi \leq \varphi \leq \pi$ and one ring with $0 \leq \theta \leq \pi$ at $\varphi = 0$ and $\varphi = \pi$. (see Fig. 26.) If we can get a stable solution for P_{nm} using this data, then we can use the same expansion to determine p anywhere in the outward direction, e.g. on the sediment surface below. A greatly simplifying approach which avoids the very complicated analysis needed to connect the ring geometries to the spherical harmonic coordinates (see Fig. 26) is to rotate the spherical harmonic coordinate system so that the FT of say the 5th ring is done only on the θ coordinate, not the φ coordinate. The rotation of spherical harmonics is provided by the Wigner D matrix based on the Euler angle rotations which is well described in Wikipedia and where the Wigner D matrix elements are the Fourier coefficients of the rotated spherical harmonic given in terms of the Euler angles.

We successfully developed an algorithm based on this method and then successfully tested it using synthesized data for the radiation from a point source. We then applied the algorithm to the BOSS ring data in order to obtain the BOSS source spherical hologram from which we would project the BOSS source down onto a plane below (the sediment surface).

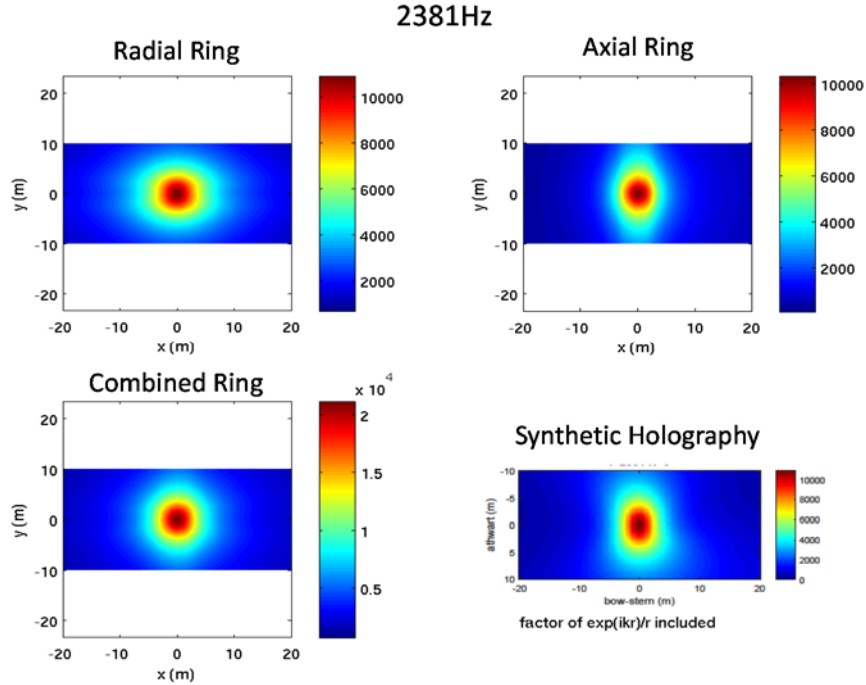


Figure 27. Comparisons of the source amplitude onto the sediment surface for the four projection techniques at low frequency.

We show projections of the BOSS calibration data onto the surface of the sediment (3m depths) in Figs. 27 and 28 using the various projection methods i.e. simple spatial pattern projection for the radial ring (at the source position), axial ring, and combined ring cases and the spherical holography projection. Fig 27 shows projections using the 2.3 kHz calibration data and Fig 28 using the 5.4 kHz data.

At the lower frequency, all projections look reasonable, although there are some differences among them. We are not able to determine the accuracy of each nor which is closest to the actual source pattern. Also, as shown in Fig. 28, the differences in the projection results increase noticeably at the higher frequency. Further, the forward beaming displayed in the spherical holography is not seen in the other projections and is almost certainly spurious indicating a problem with this projection algorithm. We believe this is related to a regularization issue, but we have not addressed this problem at this point.

Finally, we discuss (4) the Equivalent Source Method in which we approximate the BOSS source radiation in terms of imaginary point sources on the surface of an imaginary sphere of radius a centered on the BOSS source. We undertook this development sometime after development of the approaches just presented and after having applied those to the BOSS scattering data and the target

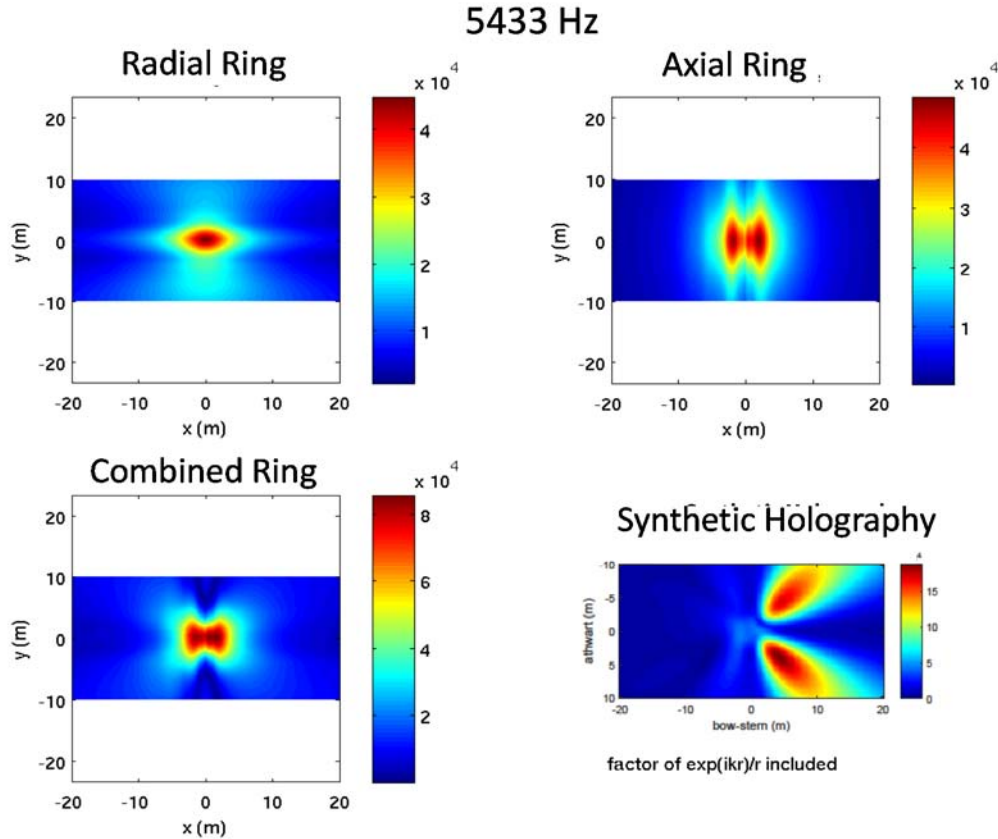


Figure 28. Comparisons of the source amplitude onto the sediment surface for the four projection techniques at mid frequency.

classification results. As a consequence, we do not show results for this method until near the end of the report; but we present the technique itself here. In the Equivalent Source Method, we generate an equivalent source distribution as determined by matching the predicted response of the sum of sources to the ring responses using SVD.

Specifically, the locations of the imaginary sources are

$$\vec{r}_j \text{ for } j=1, \dots, N$$

and the source strength of the imaginary source j is

$$q_j \text{ for } j=1, \dots, N$$

The field produced by the N imaginary sources at location \vec{r} is

$$p(\vec{r}) = \sum_{j=1}^N q_j \frac{e^{ik|\vec{r}-\vec{r}_j|}}{|\vec{r}-\vec{r}_j|}.$$

We evaluate the pressure in the equation above at the ring measurement locations

$$P_i = \sum_{j=1}^N q_j \frac{e^{ik|\vec{r}_i^m-\vec{r}_j|}}{|\vec{r}_i^m-\vec{r}_j|}.$$

We write P in terms of the Greens function

$$P_i = \sum_{j=1}^N G_{ij} q_j,$$

where

$$G_{ij} = \frac{e^{ik|\vec{r}_i^m-\vec{r}_j|}}{|\vec{r}_i^m-\vec{r}_j|}.$$

Finally, the source strengths are given by the matrix inversion

$$Q = G^{-1}P.$$

Simple checks on this approach as well as success we have had applying it to related problems indicated that it should be included as one of the approaches we should examine for removing the source pattern effect from our processed BOSS Gulf data, and as just mentioned we will present these results later.

SUBTASK 5.2 Using BOSS calibration data and optimized features, determine improvement in feature separation results and elastic imagery.

In what follows, we will first present results regarding feature separation based on using the simple projection approach to insert source calibration information into the scattering data. Shown in Fig. 29 is an approximate description of the spatial source mapping onto the target field for different frequencies using the radial ring projection compared to that for a spherical source. The results which follow were obtained prior to our development of the equivalent source method which was applied to the data only after the simple projection results were compiled. The results for applying the equivalent source method to several important cases will be presented at the end of this section.

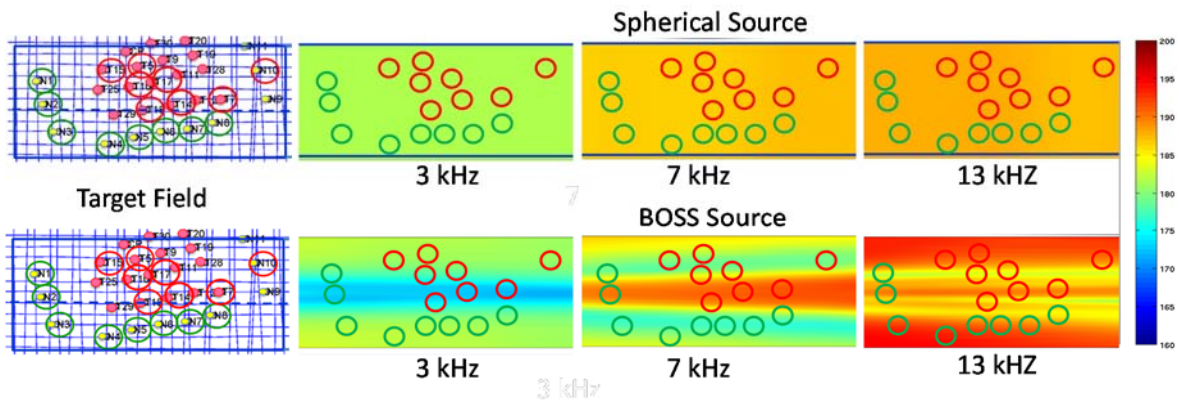


Figure 29. Comparison of source spatial distribution over target field for a spherical source versus approximately projecting source calibration ring data for three frequencies.

In order to lay the proper framework for the studies to insert BOSS source calibration data into the feature separation studies, we first carried out generative training and testing studies with alternate pings using single paths (h, i, g) for eight different feature sets. (Note that *generative* training uses only the UXO targets themselves.) The three paths as well as the general geometry and target layout are shown in Fig. 30. The feature sets were: (1) 12 TS averages; (2) 7 TS averages; (3) Max TS, bin #, and standard deviation; (4) Max TS and bin #; (5) Max TS; (6) Mean TS; (7) Acoustic color (phase and amp) 20m aperture; and (8) Acoustic color (phase and amp) 2m aperture. In general, we trained generatively on buried UXOs n1 – n8 using even indexed pings and tested on the buried UXO using odd indexed pings and on the non-UXOs (buried and proud) that were within 10m of the flight path using *both* odd and even indexed pings. More specifically, over paths h and i we trained with the 8 UXOs n1 – n8 and over path g with the 5 UXOs n1 – n3 and n7 and n8. To understand these results, it is important to realize that we are not presenting them as tests of a *trained* classification algorithm. Rather the results are used to determine whether the particular features being considered are capable of being used to separate the UXO targets from the non-UXO targets.

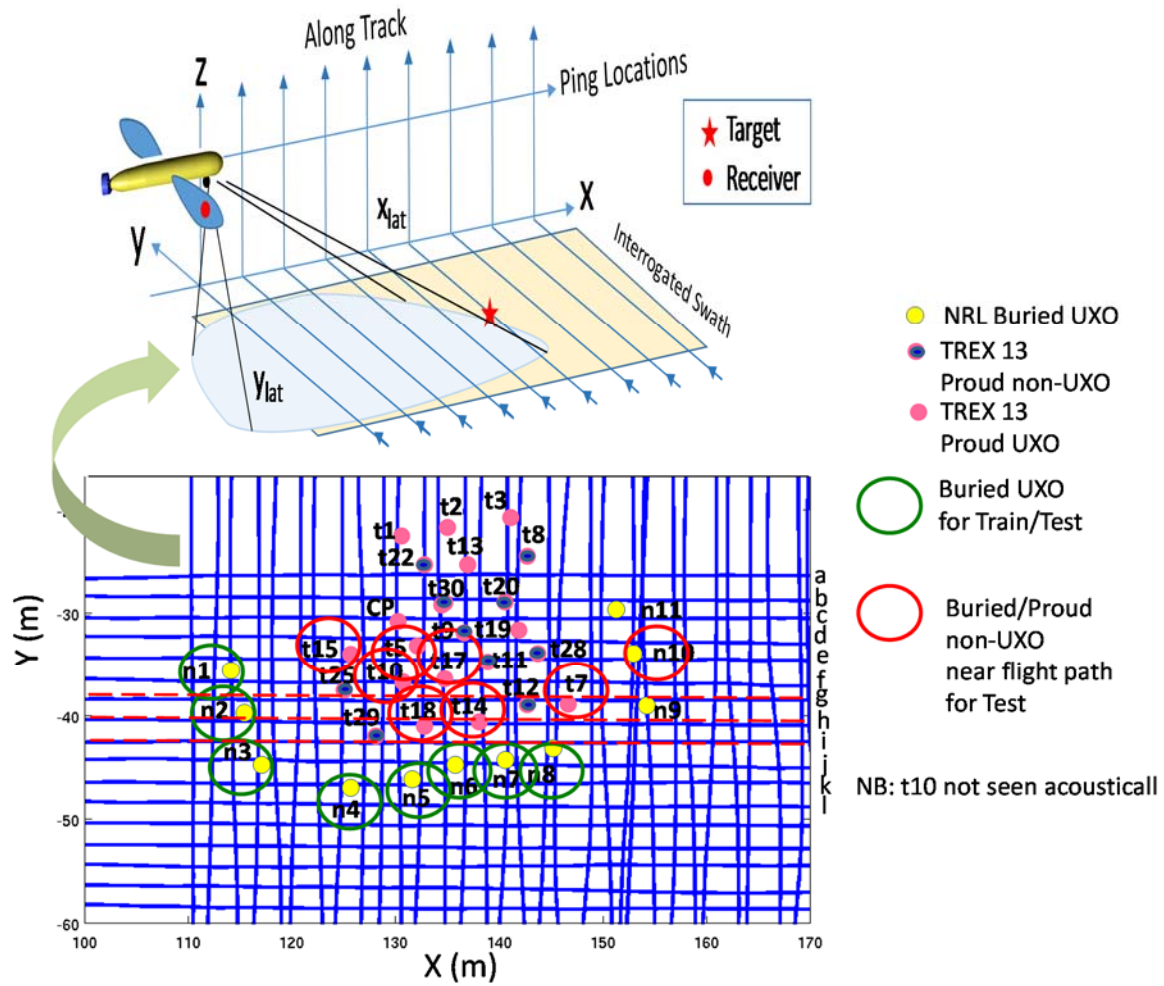


Figure 30. Cartoon illustrating the BOSS data collection geometry and the relative positions of the various targets and the flight paths. Paths *g*, *h*, and *i* are dashed.

The false targets used with paths *h* and *i* are shown in Fig. 31a and those with path *g* in Fig. 31b.

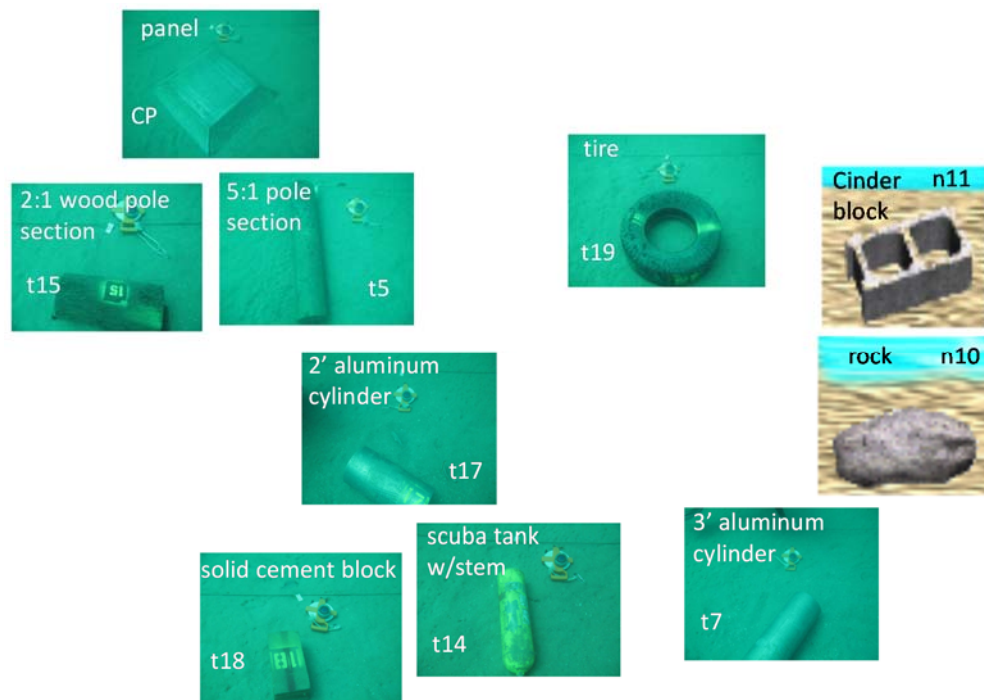


Figure 31b. False targets considered in studies along path *g*.




Figure 31a. False targets considered in studies along paths *h* and *i*.




Figure 32. Target field (left) and spectral/spatial (1 dimensional) source pattern: spherical (upper); radial calibration ring (lower).




Figure 33. Target field (lower) and spectral/spatial (1 dimensional) source pattern: spherical (right); axial calibration ring (left).

We first compute the target strengths versus frequency and angle (or x) using projections of the incident pressure field. The first is the assumption of a spherical source with frequency dependence given by the manufacturer's transfer voltage response curve. This frequency-position pattern is shown in the upper right hand corner of Figs. 32 and 33. The second is a spatial pattern determined by projecting the pressure/phase values measured on the calibration source ring at each angle along a line that meets the sediment surface at a point a distance r away

and adjusting the phase according to $\exp(ikr)$ and the amplitude by $1/r$. This frequency-position pattern is shown in the lower right hand corner of Fig. 32. The third uses a similar procedure with the axial ring. This frequency pattern is shown in the lower right hand corner of Fig. 33. The fourth uses a combined approach in which the radial and axial ring data are projected and then averaged at each point on the sediment.

Having obtained these four different source descriptions, we carried out the alternate ping generative feature separation studies along the three paths described above. The results are shown in the following figures for the eight feature sets and for the three paths g , h , and i .

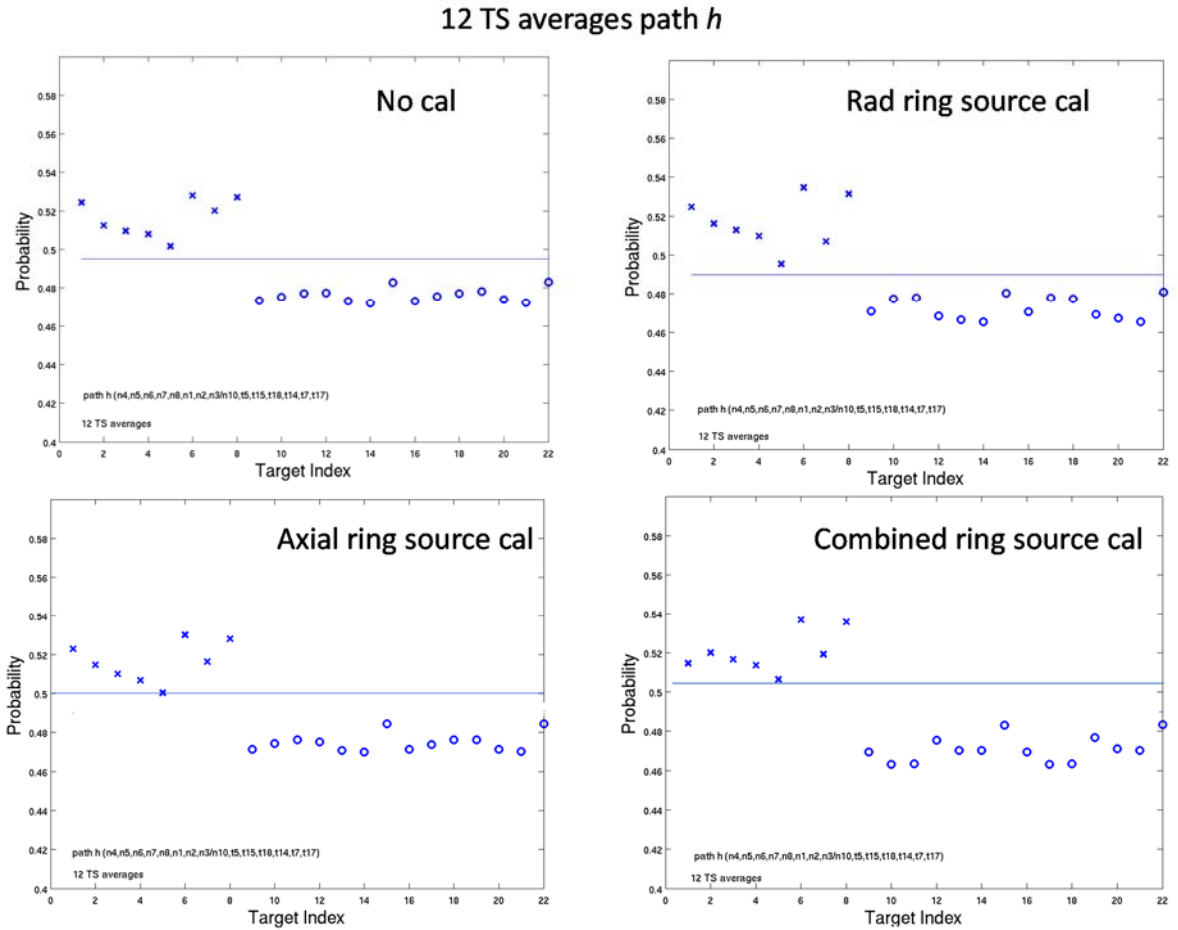


Figure 34. Generative target feature separation using alternate pings for the 12 TS average feature set over path h .

As can be seen in Fig. 34, the no calibration case for the 12 TS averages feature set demonstrates complete UXO/non-UXO target separation as do the three calibration cases. The combined calibration case does show a slightly improved average separation between the true and false targets.

The results shown in Fig. 35 for the 7 TS average feature set are similar. In this case, the combined calibration case shows even better separation between the true and false targets.




Figure 35. Generative target feature separation using alternate pings for the 7 TS average feature set over path h .

The results for the max TS, bin number, standard deviation feature set over path h are shown in Fig. 36. In this case, the radial ring calibration case introduces a false positive, and as indicated it is associated with the two foot aluminum cylinder false target. Actually, one can see two false positive entries in Fig 36 for the radial ring calibration case. Both of these are associated with the same target, and this is a consequence of the training/testing protocol in that the generative *training* (done on alternate pings) by definition is done only on true targets so that both odd and even source

pings are available for *testing* the false targets. This results in two data points for each false target. For this reason, in the accounting that is done later, we will count each false target data point as one half a data point.

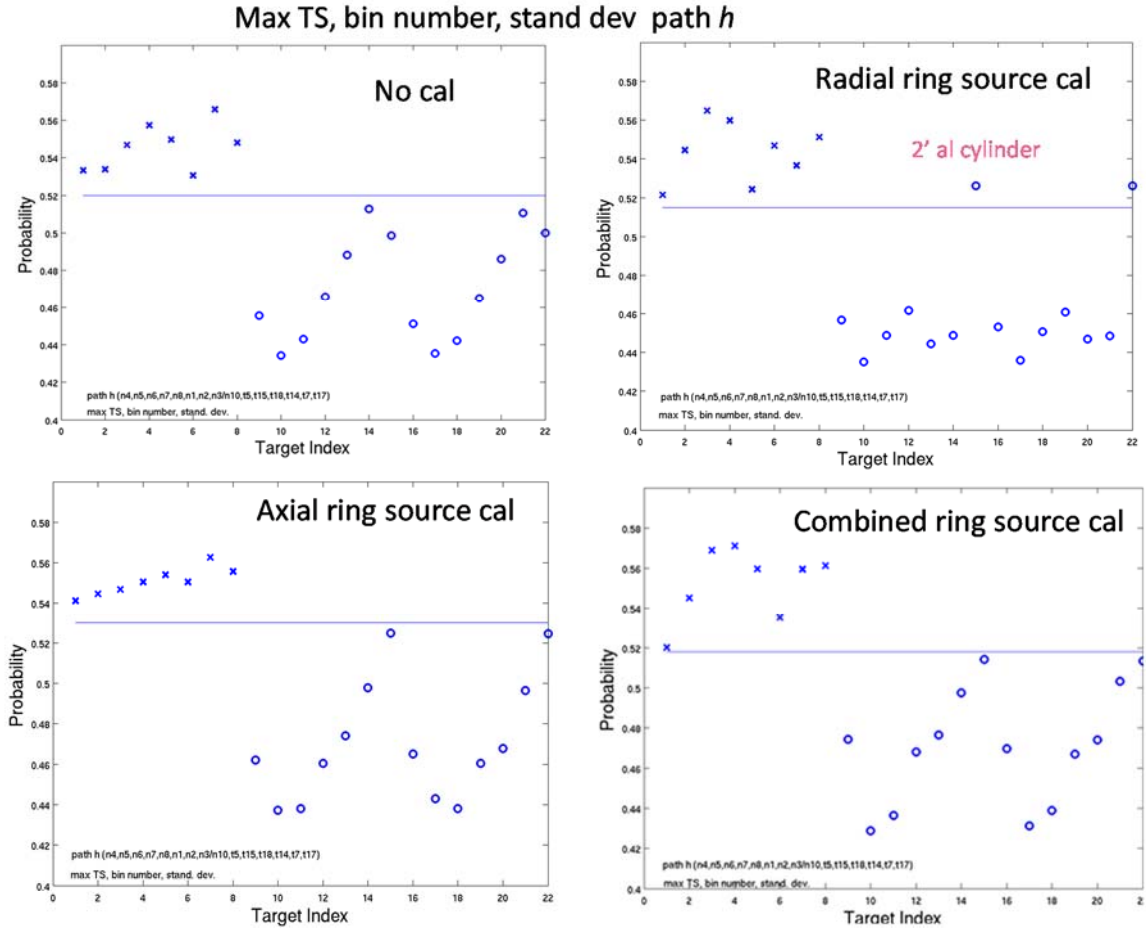


Figure 36. Generative target feature separation using alternate pings for the max TS, bin number, standard deviation feature set over path h .

In the next figure (Fig. 37) are shown the results for the max TS, bin number feature set over path h . Similar to the previous feature set, only the radial ring calibration case has a false positive; and again it is the two foot aluminum cylinder. Certainly, this false target is in some respects (metal, solid, size about the same) one of the false targets most similar to the UXO targets. On the other hand, the other two calibration cases (axial ring and combined ring) not only have no false positives but they also show an improved target separation.




Figure 37. Generative target feature separation using alternate pings for the max TS, bin number feature set over path h .

Next, we show in Fig. 38 the results for the max TS feature set. As can be seen, the non-calibration case contains no false positives; and while this is true also of the radial and axial calibration cases, the target separation is near zero, especially for the radial ring case. We also note that for the first time the n4 UXO target falls a considerable distance from the other UXOs, and this is seen for other feature and path cases as we will see. While we do not yet understand why n4 may be a problematic target, after we complete showing all the feature sets and flight paths with n4 included we will repeat the studies without using the n4 UXO target.




Figure 38. Generative target feature separation using alternate pings for the max TS, bin number feature set over path h .

Next, we show in Fig. 39 the results for the mean TS feature set over path h . Here the radial ring calibration data is worse, and the axial and combined ring calibration data better, than the non-calibration case with regard to number of false positives. As in several of the previous cases, we see that in general the false targets that are almost always called correctly are the rock and the two wooden pole sections. From the perspective of shape (the rock) and material (wooden pole sections), that this is the case is not surprising given the UXO material (steel) and shape ($\sim 4:1$ aspect ratio cylindrical shape).




Figure 39. Generative target feature separation using alternate pings for the mean TS feature set over path h .

Next, consider the 140 frequencies x 22 positions, 20m aperture acoustic color feature set on path h shown in Fig.40. Unlike the other two calibration cases, the combined calibration case eliminates one false positive from the original result. We also point out once again that the false targets that seem to be called correctly consistently are the proud rock, 5:1 wooden pole section, and 2:1 wooden pole section.

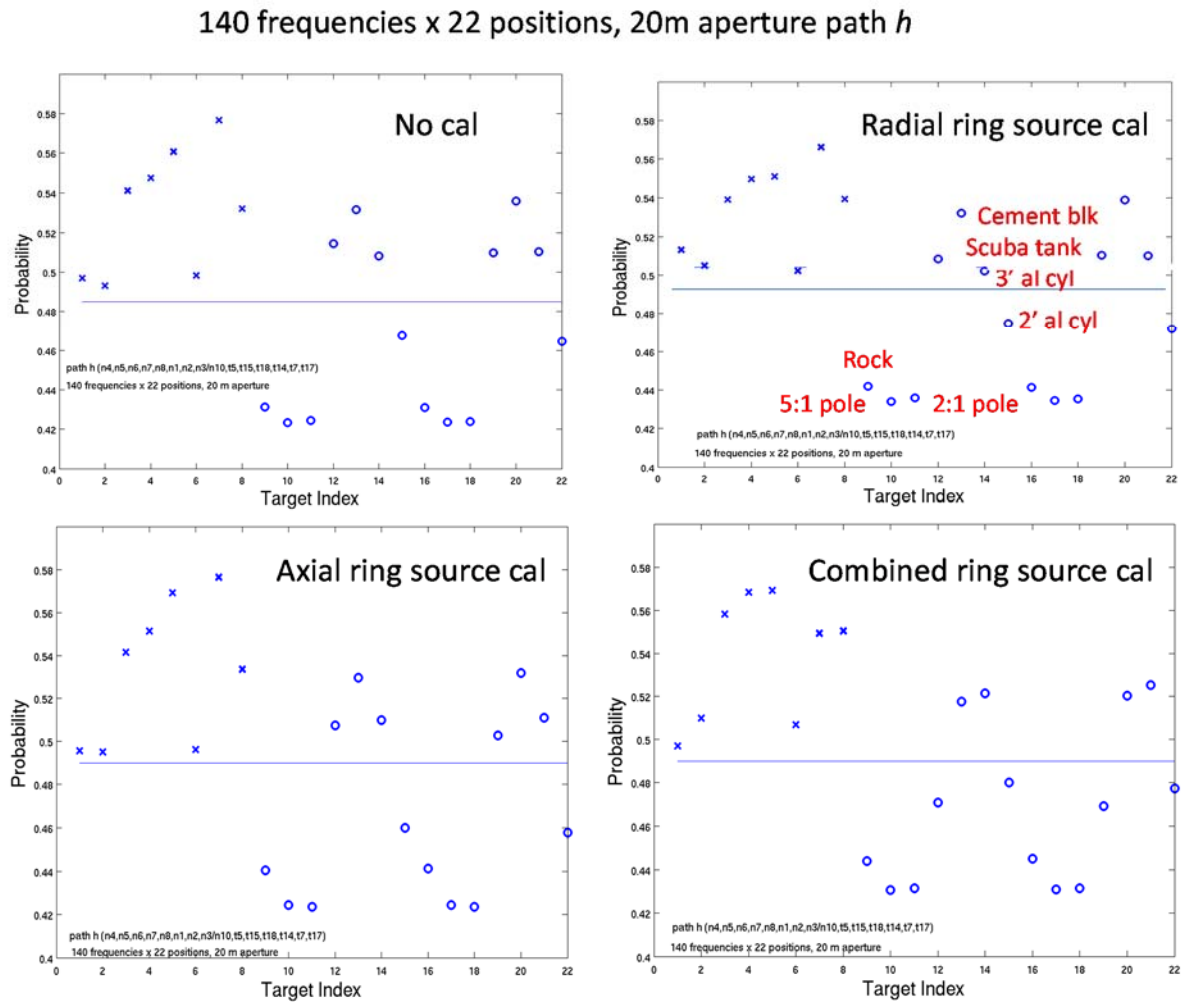


Figure 40. Generative target feature separation using alternate pings for the 140 frequencies x 22 positions, 20m aperture feature set over path h .

The final case for path h is the 140 frequencies x 22 positions, 2m aperture acoustic color feature set shown in Fig. 41. As can be seen, the radial ring and axial ring calibration cases improve the performance only marginally, if at all. The combined calibration case shows again the anomalous behavior of the n4 UXO as can be seen by how far it falls from the other UXO training targets. As said before, we will come back to this after running through the cases for the other two paths.

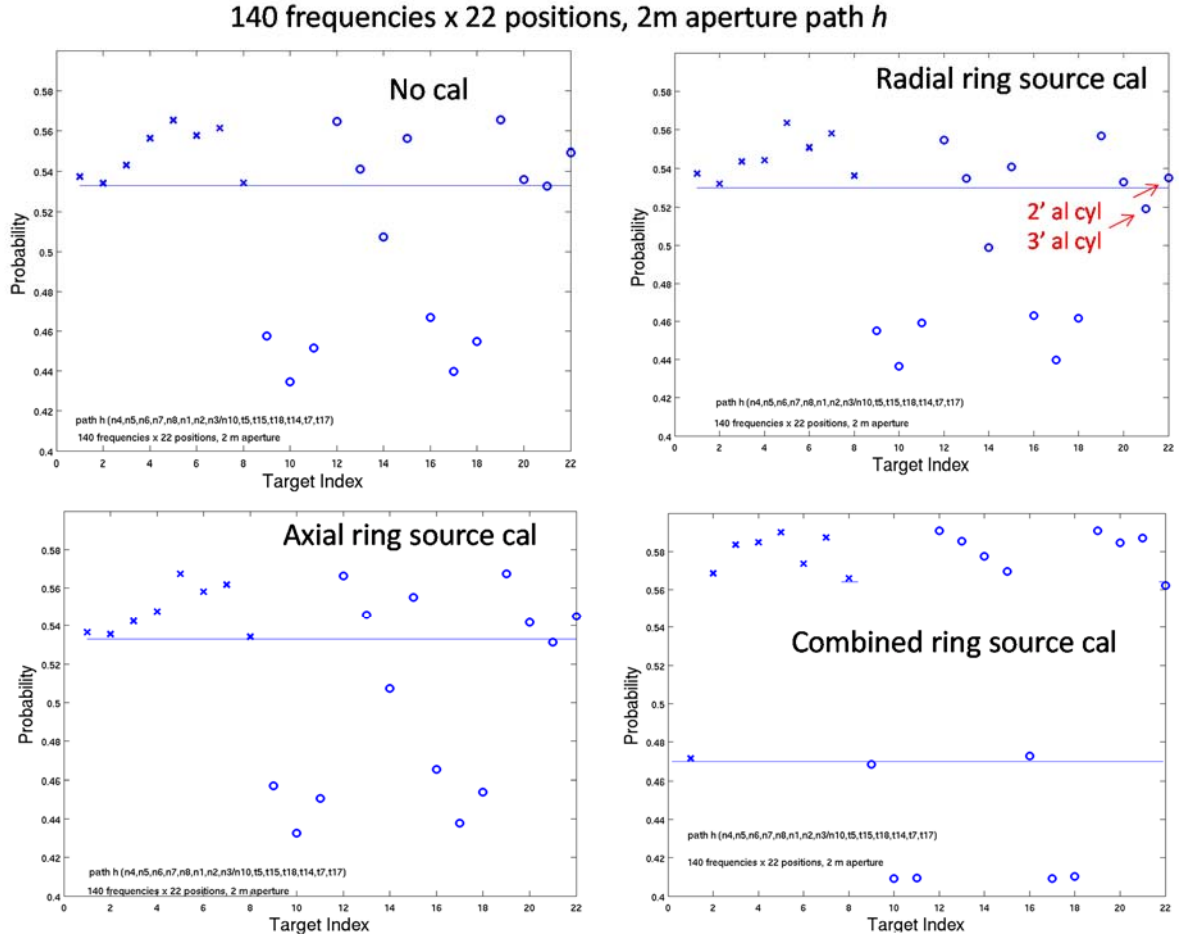


Figure 41. Generative target feature separation using alternate pings for the 140 frequencies x 22 positions, 2m aperture feature set over path h .

Next, we present the same four generative training and testing results for path i . As mentioned earlier, the same targets are used for training and for testing that were used in the path h case. In Fig. 42 is shown the 12 TS averages case where we see that there are no false positives for the no calibration case and for all three source calibration cases similar to what was found for path h . However, here the separation between true and false positives is much smaller for path i . In the latter case we note that on path i the training targets are somewhat closer while the false targets are somewhat farther than for the path h case. On the other hand, the combined calibration case for path i still shows an increase in the separation over that for the no calibration case.

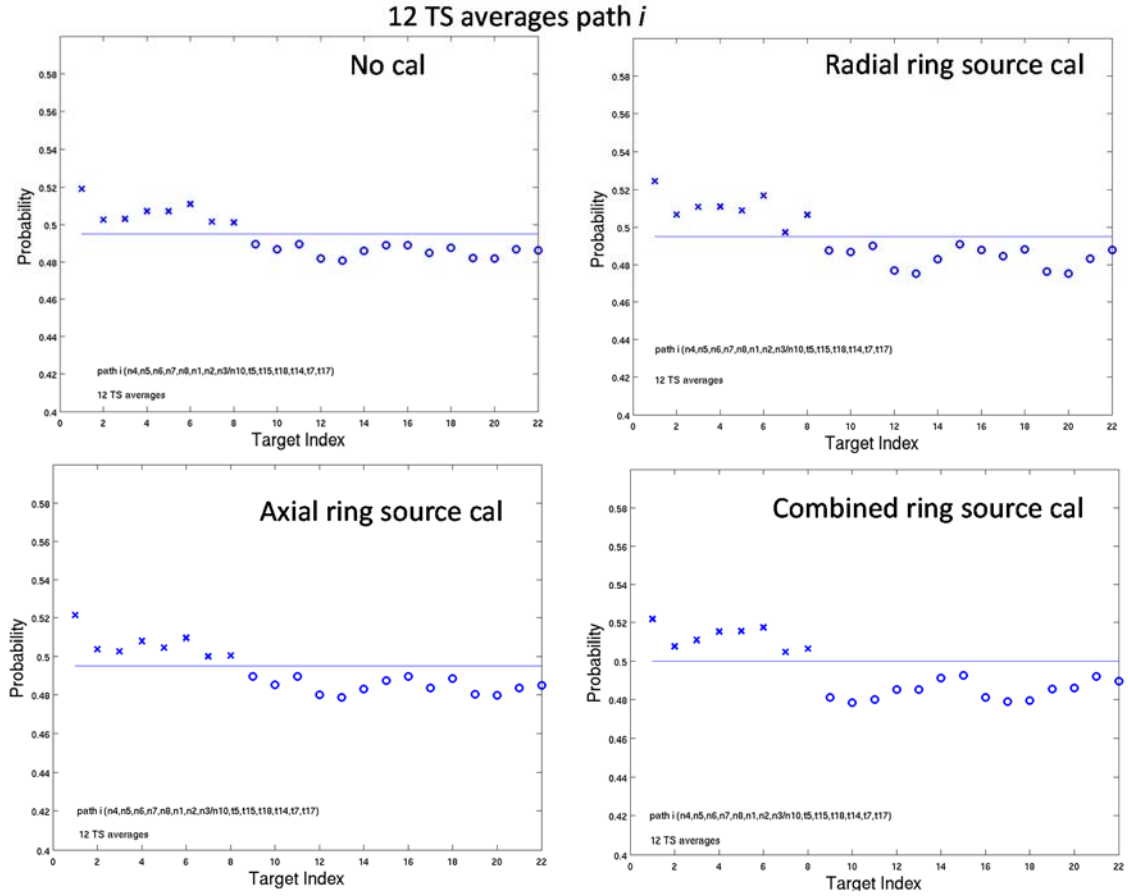


Figure 42. Generative target feature separation using alternate pings for the 12 TS average feature set over path i .

Next in Fig. 43 is shown the 7 TS average feature set case. Here the results are seen to be very similar to that found for the 12 average TS feature set case. In particular, we see that there are no false positives for the no calibration case and for all three source calibration cases similar to what was found for path h . Again, here the separation between true and false positives is much smaller for path i . In the latter case we note again that the training targets are somewhat closer while the false targets are somewhat farther than they were for the path h case. On the other hand, the combined calibration case for path i still shows an increase in the separation over that for the no calibration case as was the case for the previous feature set.

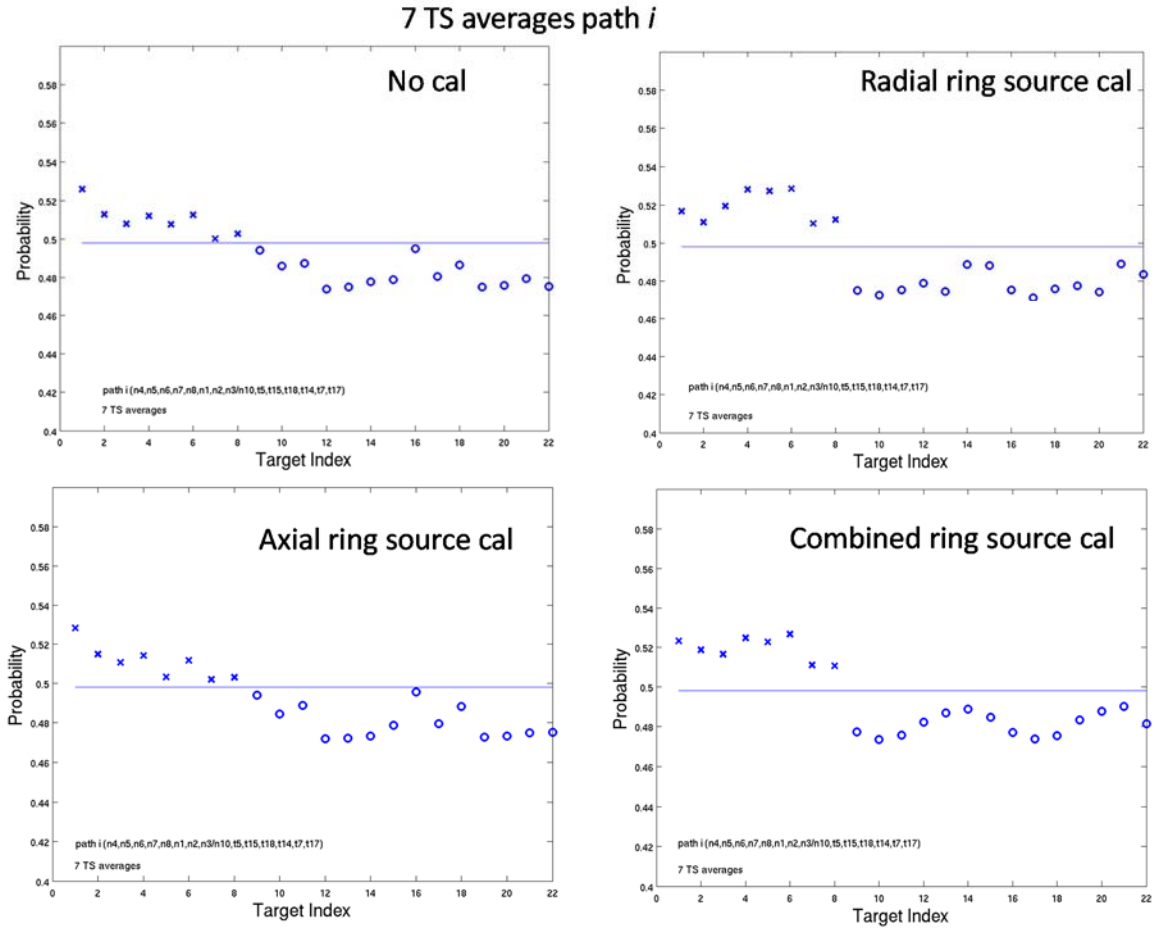


Figure 43. Generative target feature separation using alternate pings for the 7 TS average feature set over path i .

Next, we show the results for the max TS, bin number, standard deviation feature set in Fig. 44. Here again, the results for path i are very similar to what was found for path h . As in the previous case, except for the radial ring calibration case, the other three cases (no calibration, axial ring calibration, and combined calibration) demonstrate no false positives. One interesting difference is the fact that for path g the one false positive in the radial ring calibration case is the 3 foot aluminum cylinder whereas for path h the false positive is the 2 foot aluminum cylinder. As mentioned previously, both these aluminum targets are expected to be somewhat similar to the UXO targets themselves both with respect to size and shape and to their metal composition.

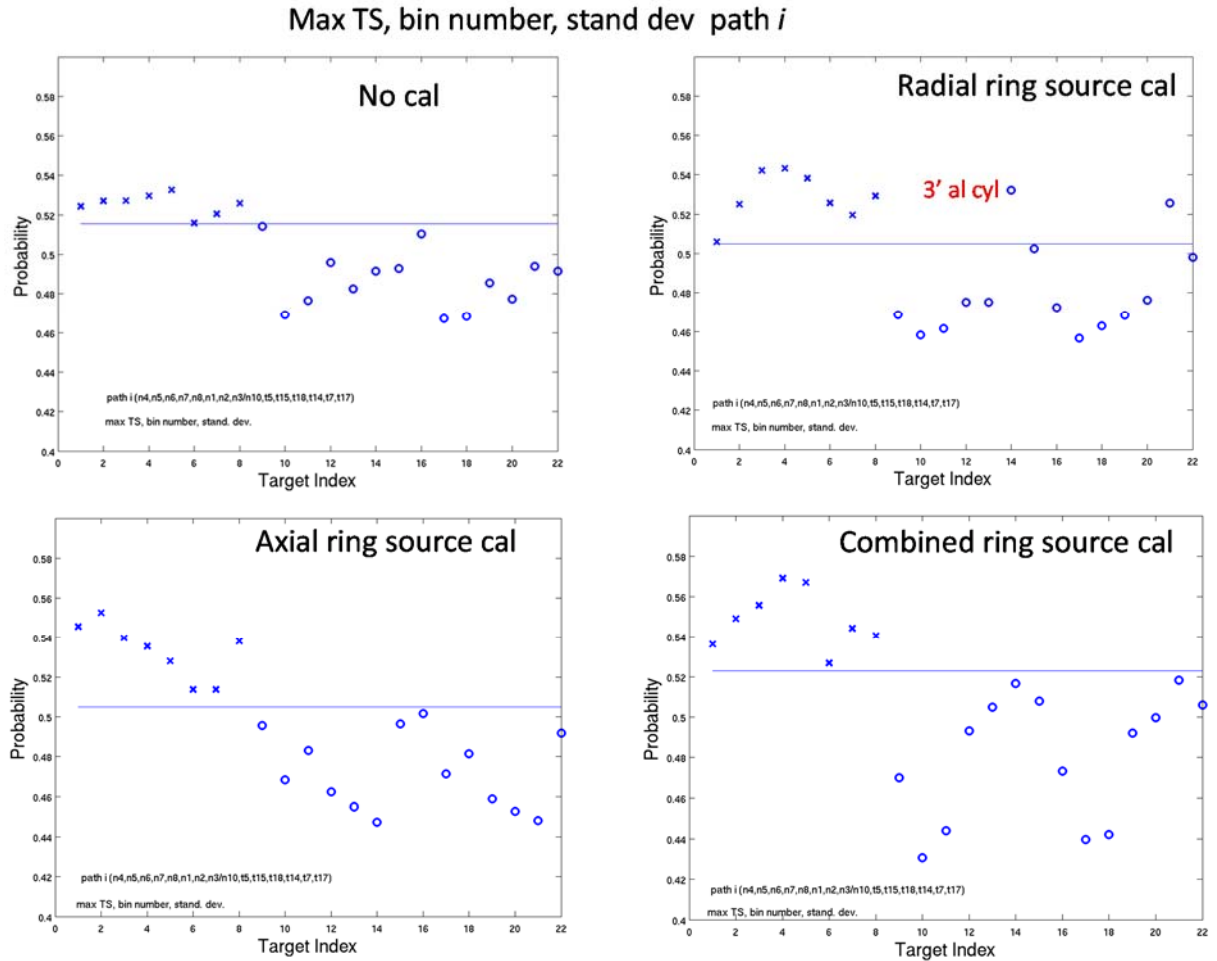


Figure 44. Generative target feature separation using alternate pings for the max TS, bin number, standard deviation feature set over path i .

Next, we show the results for the max TS, bin number feature set in Fig. 45. Here, the results for path i are somewhat worse than what was found for path h . In particular, in the radial ring calibration case, the 3 inch aluminum cylinder has now been added to the 3 inch aluminum cylinder as a false positive. This is not surprising since as mentioned previously, both these aluminum targets are expected to be somewhat similar to the UXO targets themselves both with respect to size and shape and to their metal composition. Further, in the axial ring calibration case, unlike for path h , there is now a false positive (the proud rock). On the other hand, the combined calibration case has no false positives and has a better target separation than the no calibration case.




Figure 45. Generative target feature separation using alternate pings for the max TS, bin number feature set over path i .

Next, we show the results for the max TS feature set in Fig. 46. Here, the results for the no calibration and radial ring calibration cases for path i are worse than what was found for path h . In particular, compared to no false positives there are now three. Further, the axial ring case which had no false positives in path h now has two. As was seen in path h for the combined calibration case, the n4 UXO target has become an outlier; and in this case the false positives grow from two to four. Again, we point out that later we will redo these cases without including n4 in the training.

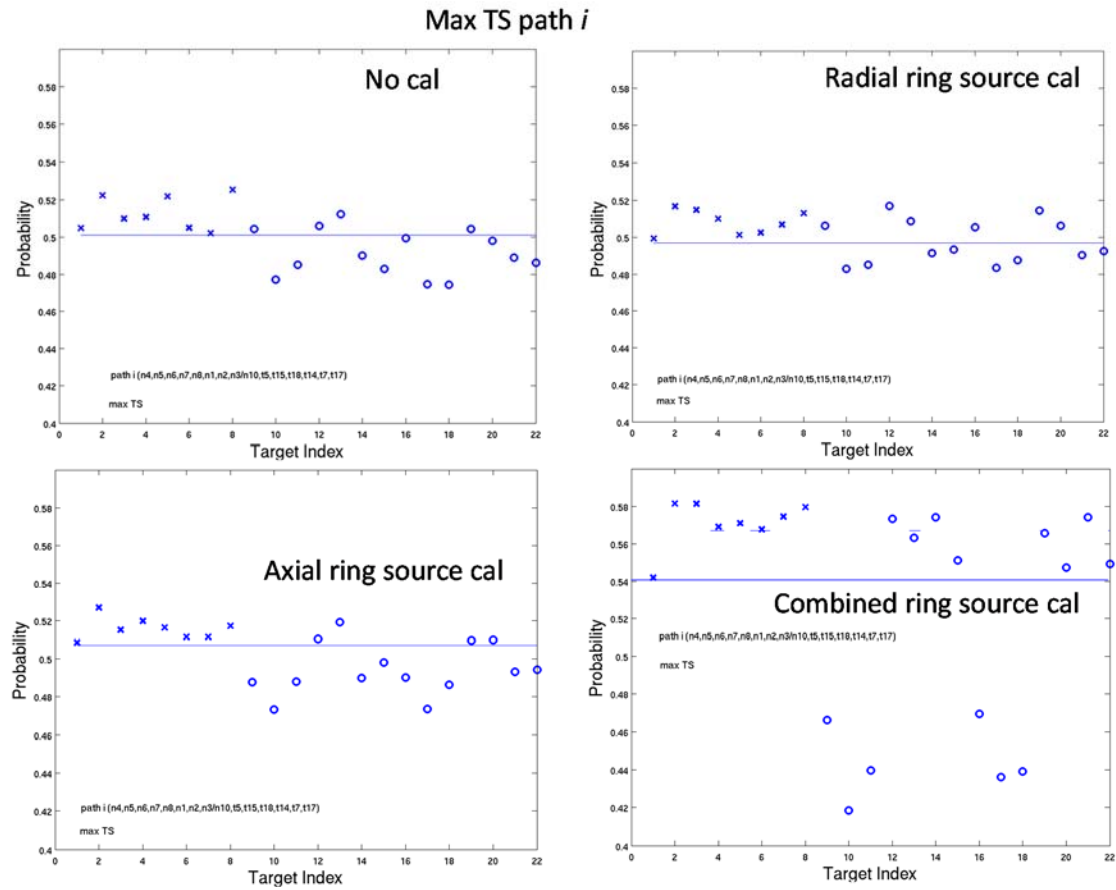


Figure 46. Generative target feature separation using alternate pings for the max TS feature set over path i .

Next, we show the results for the mean TS feature set in Fig. 47. Here, the results for path i no calibration and all three calibration cases are worse than what was found for path h . Regarding the effect of calibration over path i , both the axial ring calibration and the combined calibration show some improvement in the number of false positives over the no calibration case. In particular, compared to two false positives with no calibration, there is now only one in the axial and combined cases. However, the number of false positives has grown from two to four in the radial calibration case.

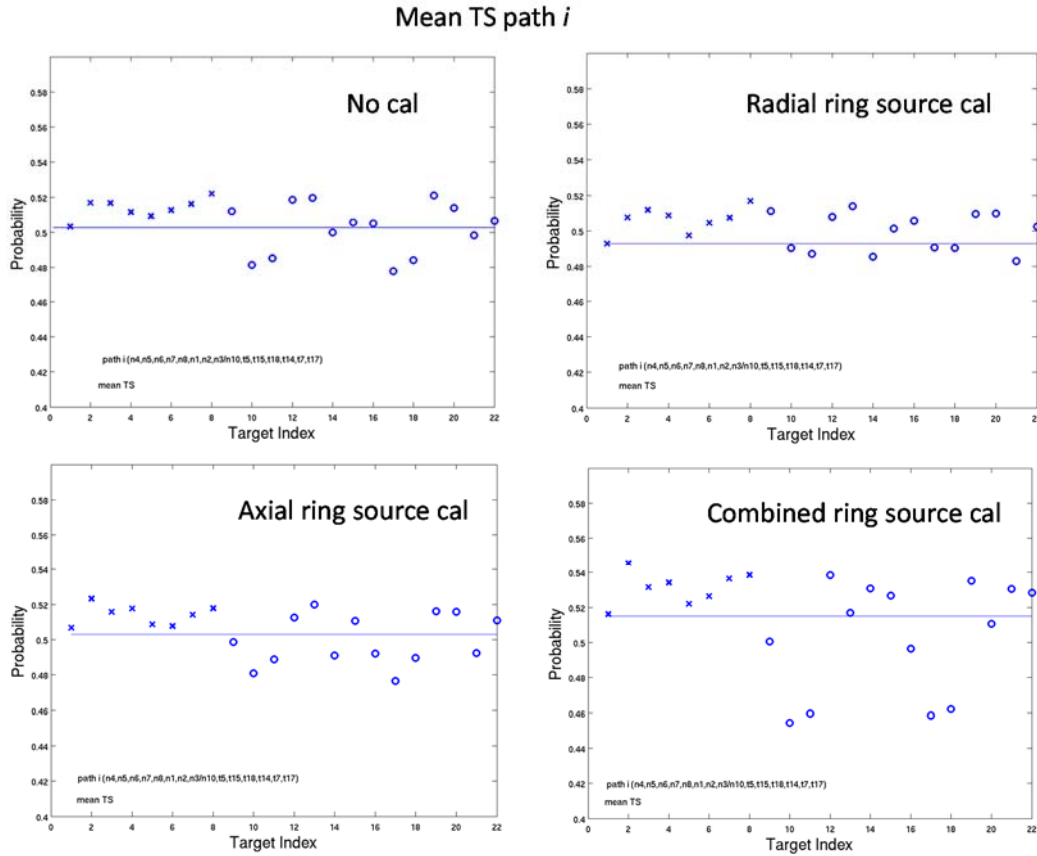


Figure 47. Generative target feature separation using alternate pings for the mean TS feature set over path i .

Next, we show the results for the 140 frequencies x 22 positions, 20m aperture acoustic color feature set. As can be seen in Fig. 48, all four cases (the no calibration case, the radial calibration case, the axial calibration case, and the combined calibration case) exhibit four false positives. As displayed in the upper left corner, the four false positives are the cement block, the scuba tank, the three foot aluminum cylinder and the two foot aluminum cylinder. These targets show up as false positives more often than any others, and as mentioned previously these targets are expected to be somewhat similar to the UXO targets themselves either with respect to size and shape or to their metal/hard material composition. This is similar to the results for the 140 frequencies x 22 positions, 2m aperture feature set shown in Fig. 49. As can be seen there, all four cases exhibit three or four false positive calls. In particular, in the no calibration, the radial ring calibration, and the axial ring calibration cases the three false positives are the cement block, the scuba tank, and the 3 foot aluminum cylinder while the 2 foot aluminum cylinder is added to the other three as false positives. In contrast to the previous feature set, the combined calibration case is worse than the no calibration case. However, as we have seen before, in the combined calibration case, the n_4

UXO target has become an outlier in the group of UXO targets used in the training. We point out again that later we will consider training without the n4 target.




Figure 48. Generative target feature separation using alternate pings for the 140 frequencies x 22 positions, 20m feature set over path *i*.

Next, we present the same four generative training and testing results for path *g*. For path *g* a somewhat different set of targets is used for training and for testing than were used in the path *h* and *i* cases. As mentioned earlier, this is in response to, and to compensate for, the movement of the flight path further from the lower line of UXOs and closer to the field of proud false targets. We now use the n1, n2, n3, n7, and n8 UXOs for training and the n10, n11, t5, cp, t15, t18, t19, t14, t7, and t17 targets as false targets. In Fig. 50 is shown the 12 TS averages case where we see that there are no false positives for the no calibration case and for all three source calibration cases

similar to what was found for paths *h* and *i*. However, here the separation between true and false positives is much smaller for the other two paths especially for path *h*.

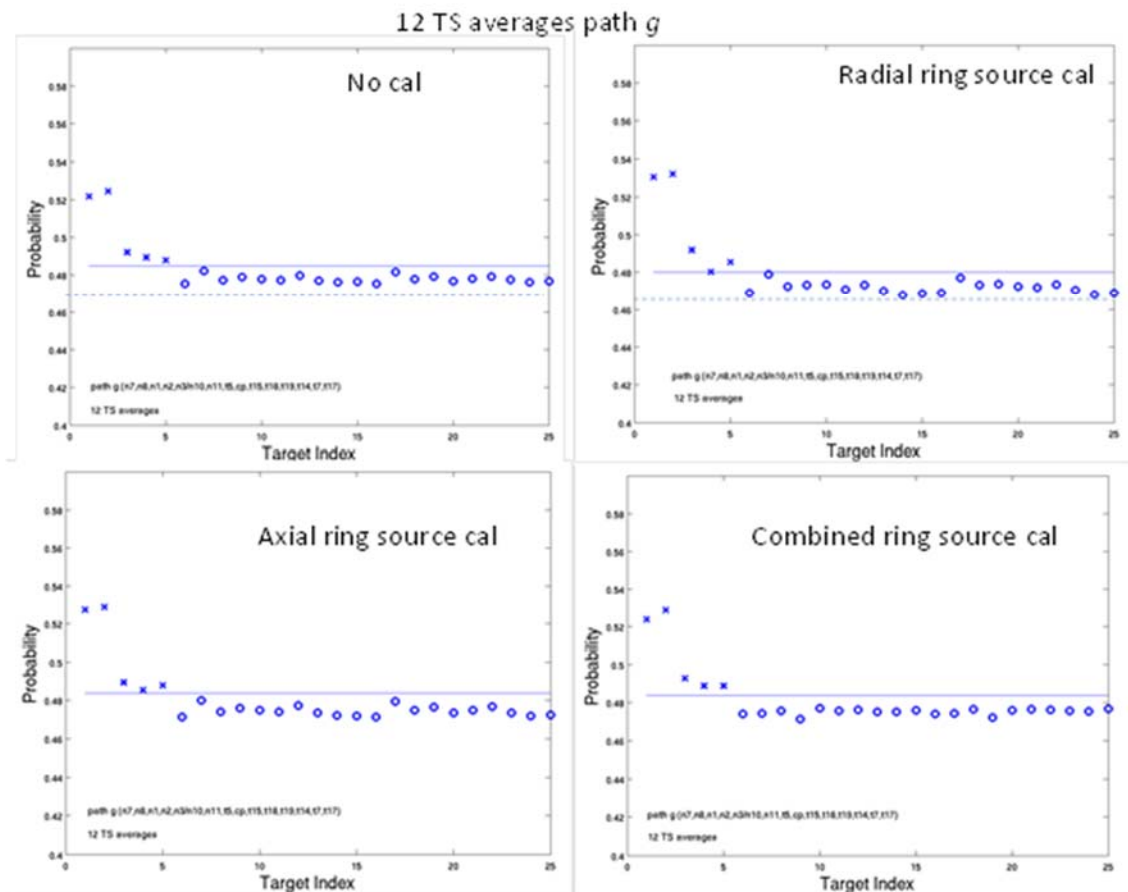


Figure 50. Generative target feature separation using alternate pings for the 12 TS averages feature set over path *g*.

Next in Fig. 51 is shown the results for the 7 TS averages feature set over path *g*. Again, all four cases (the no calibration and the three calibration cases) demonstrate no false positives. Further, the target separation is larger than for the previous feature space, and the combined calibration case shows an increased target separation over the no calibration case.




Figure 51. Generative target feature separation using alternate pings for the 7 TS averages feature set over path g.

Next is shown the results for the maximum TS, bin number, standard deviation feature set over path g. As can be seen in Fig. 52, the no calibration case has one false positive, the tire. The axial ring calibration removes this false positive, and the radial ring calibration removes this as well but adds another false positive, the solid cement block. However, the combined calibration process not only removes all false positives completely but also increases the target separation.




Figure 52. Generative target feature separation using alternate pings for the max TS, bin number, standard deviation feature set over path *g*.

Next is shown the results for the max TS, bin number feature set over path *g*. As can be seen in Fig. 53, the results are very similar to the results for the preceding feature set. One difference is that the axial calibration case does not now remove the tire false positive. However, the combined calibration case still removes all false positives.




Figure 53. Generative target feature separation using alternate pings for the max TS, bin number feature set over path g.

Next we show the results for the max TS feature set over path g. As can be seen in Fig. 54, unlike the previous feature set case (max TS, bin number), now all three calibration cases exhibit no false positives. This was the case for path *h* as well but not for path *i*. More importantly, the calibration cases have removed the false positive (the tire) seen in the no calibration case. Further, the combined calibration case in addition to removing the false positive has produced a somewhat improved overall target separation compared to the no calibration case.

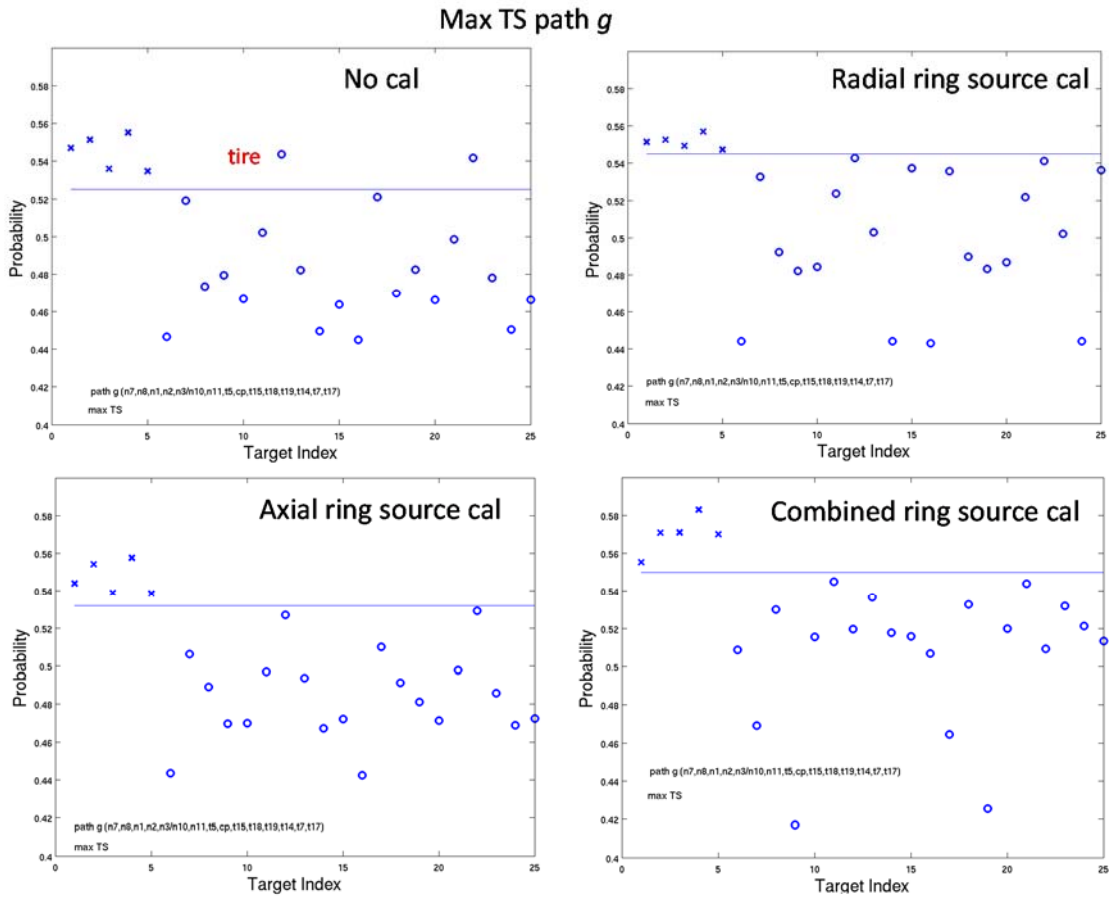


Figure 54. Generative target feature separation using alternate pings for the max TS feature set over path g .

Next in Fig. 55, we show the results for the mean TS over path g . In this case, the no calibration case exhibits two false positives, the solid cinder block and the tire. The combined calibration case is successful in removing both these false positives but just barely so for the tire. The radial ring calibration case and the axial ring calibration case remove the solid cement block as a false positive but not the tire.

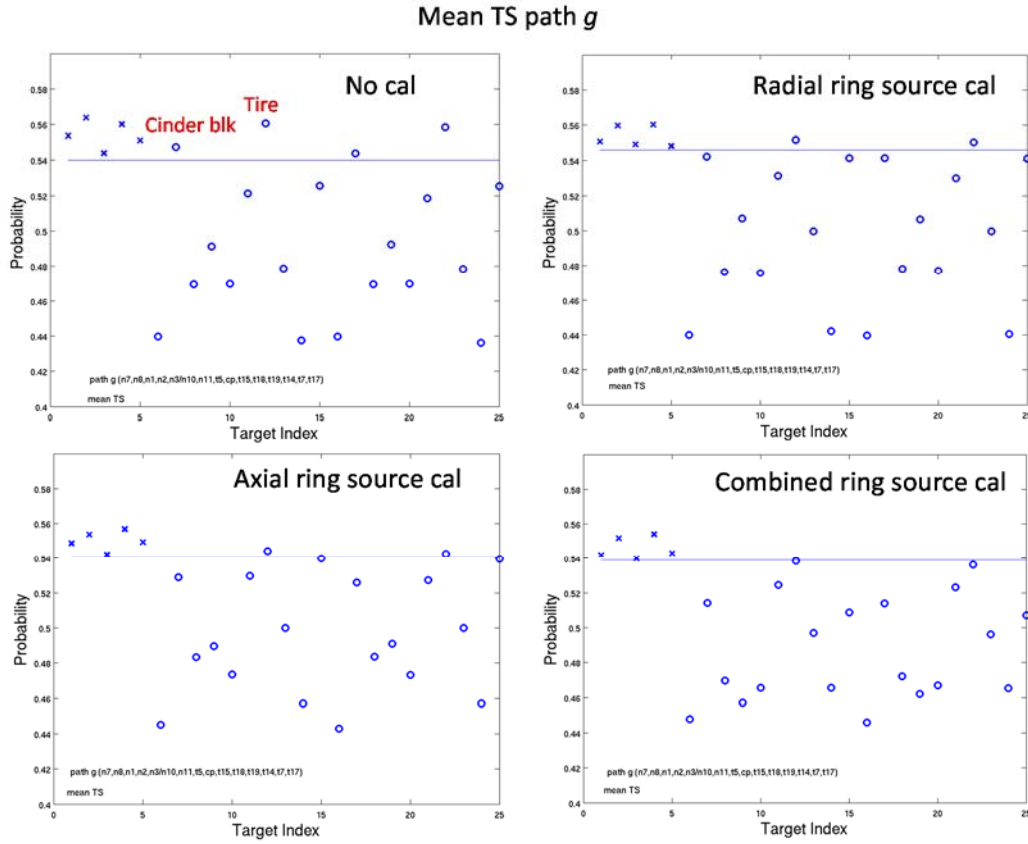


Figure 55. Generative target feature separation using alternate pings for the max TS feature set over path g .

Next we show the 140 frequencies, $\times 22$ positions, 20m aperture acoustic color feature set over path g in Fig. 56. For this case the worst performance is exhibited in the no calibration case which resulted in eight false positives. The axial ring calibration case is only slightly better having six false positives followed by the radial calibration case which has five false positives. The combined calibration case itself has six false positives. We point out that we see in the axial and combined calibration cases that there are two UXO outliers (n7 and n1) that have dragged the performance down considerably. As in the case of n4 in the previous paths, we will redo these training cases after eliminating n7 and n1 from the training.




Figure 56. Generative target feature separation using alternate pings for the 140 frequencies x 22 positions, 20m aperture feature set over path g.

Finally, we show in Fig. 57 the results for the 140 frequencies, x 22 positions, 2m aperture acoustic color feature set over path g. Here the no calibration case has only one false positive the 2 foot aluminum cylinder. This same false positive shows up for the radial calibration case and the axial calibration case. For the combined calibration case, in addition to the 2 foot aluminum cylinder now another false positive is added – the 3foot aluminum cylinder.

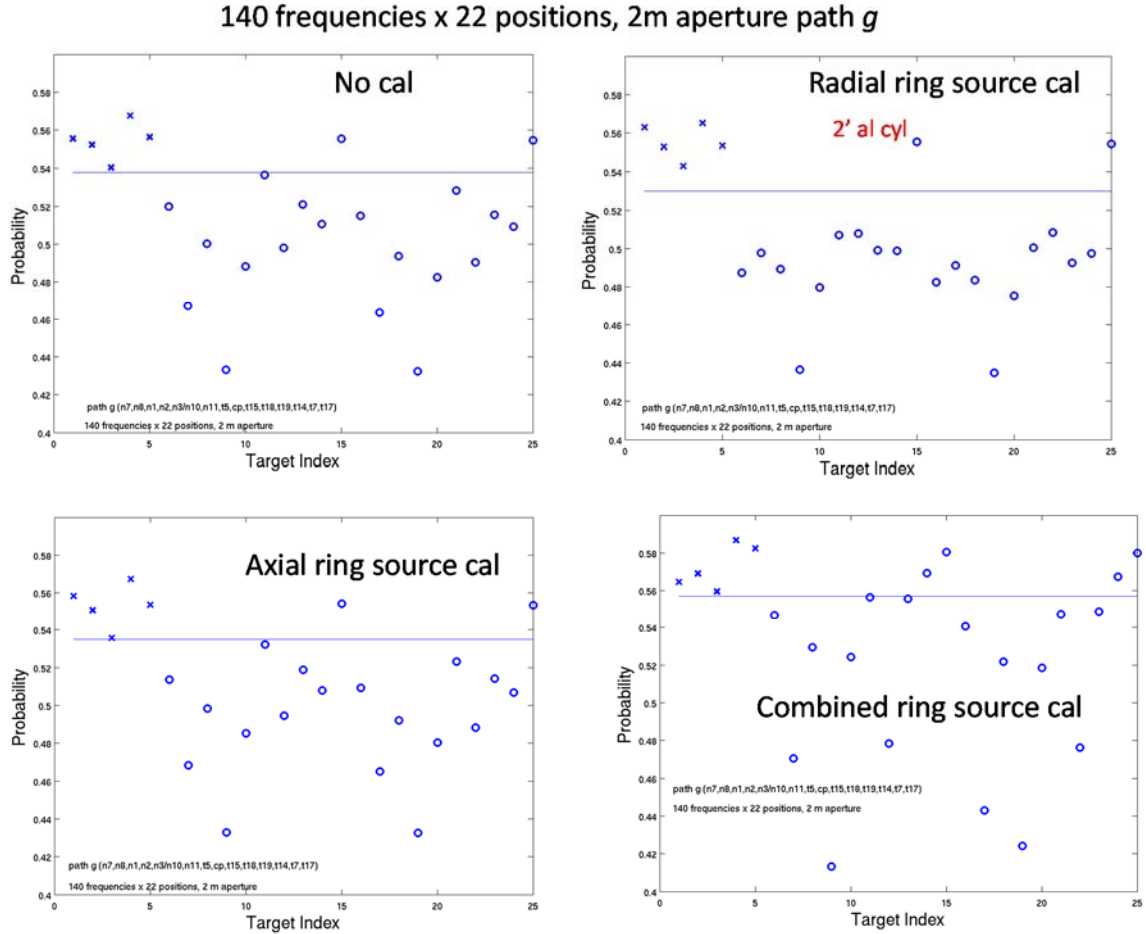


Figure 57. Generative target feature separation using alternate pings for the 140 frequencies x 22 positions, 2m aperture feature set over path g.

We had mentioned earlier that there appeared to be several problematic UXO targets used in the generative training. In these cases, the tested UXO targets exhibited an outlier whose probability number fell well below several of the false targets causing them to be called as false positives.

Of course, we could raise the threshold level and decrease the number of false positives at the expense of a less than perfect classification probability (i.e. $P_c < 1$). Nevertheless, we will retrain the classifier without using these two problematic targets holding $P_c = 1$.

One example of this can be seen in Fig. 38 for the Max TS feature set which is associated with UXO target n4 for the combined ring source calibration case over path *h*. The second example can be seen in Fig. 41 for the 140 frequencies x 22 positions, 2m aperture feature set which is again associated with UXO target n4 for the combined ring source calibration case over path *h*. The third example can be seen in Fig. 46 for the Max TS feature set which is again associated with UXO

target n4 for the combined ring source calibration case over path *i*. The fourth example can be seen in Fig. 49 for the 140 frequencies x 22 positions, 2m aperture feature set which is again associated with UXO target n4 for the combined ring source calibration case over path *i*. The fifth and final example which can be seen in Fig. 56 for the 40 frequencies x 22 positions, 20m aperture acoustic color feature set which is associated with UXO targets n1 and n7 for the combined ring source calibration case over path *g*. We show in Fig. 58 the result of eliminating UXO target n4 from the training for the combined ring source calibration case for the Max TS feature set over path *h*. As can be seen, with the elimination of this apparently problematic UXO training target, there are no longer any false positives. We also show the impact of eliminating n4 from the training on the performance of the axial ring source calibration case, and as can be seen there is an improvement of the target separation. Next in Fig. 59 is shown the result for eliminating UXO n4 from training for the 140 frequencies x 22 positions, 2m aperture acoustic color case for the combined ring calibration case over path *h*. Also shown for comparison is the original case after training with n4 included. As can be seen, this process has eliminated two but not all of the false positives.




Figure 58. Generative target feature separation using alternate pings for Max TS feature set over path *h*: original (upper); eliminating UXO n4 in training (lower).




Figure 59. Generative target feature separation using alternate pings for 140 frequencies x 22, 2m aperture feature set over path h : original (upper); eliminating UXO n4 in training (lower).

Next we consider the impact of eliminating n4 from the training over path i for the Max TS feature set for the combined ring source calibration case. This is shown in Fig. 60 along with the result for eliminating n4 from the training for the axial ring source calibration course. As in the previous cases, we also show the original results for both cases keeping UXO target n4 in the training. As can be seen, we now have reduced the false positives from four to two in the combined case and from two to none in the axial calibration case.




Figure 60. Generative target feature separation using alternate pings for Max TS feature set over path i : original (upper); eliminating UXO n4 in training (lower).

Next we consider the impact of eliminating UXO target n4 from the training over path i for the 140 frequencies x 22 positions, 2m aperture acoustic color feature set. This is shown in Fig. 61 along with the original case with n4 included. As can be seen, the elimination of n4 from the training set has reduced the false positives from four to three.





Figure 61. Generative target feature separation using alternate pings for 140 frequencies x 22 positions, 2m aperture feature set over path *i*: original (upper); eliminating UXO n4 in training (lower).

The last case to consider is that associated with the 140 frequencies x 22 positions, 20m aperture acoustic color feature set where we had seen UXO target n1 and n7 as UXO outliers over path *g*. The results after eliminating UXO targets n1 and n7 from the training are shown in Fig. 62 along with the original results which included n1 and n7. As can be seen, the elimination of n1 and n7 from the training set has resulted in a decrease in false positives from six to three in the combined ring source calibration case and from three to two in the axial ring source calibration case.



We now summarize these overall results especially regarding the impact on target separation of inserting source calibration data into the training and testing in the approximate manner described above. The overall impact is illustrated below in Table 2. Here the number of times a particular non-UXO target is called a UXO (a false positive) is shown both before source calibration is included as well as after combined ring source calibration is implemented. These number counts include the results from all three paths and for all eight feature sets. Also they apply to the case in which UXO training was carried out without including UXO target *n4* for feature sets 5 and 8 over paths *h* and *i* and excluding UXO training targets *n1* and *n7* for feature set 7 over path *g*. With regard to the statistical sample represented in Table 2, the total number of testing events over the three paths on non-UXO targets is 192. This number is arrived at in the following manner. Recall that in training with the UXO targets, we train on even pings and test on odd pings; but because the training is generative, in the case of the non-UXO targets we can test on both even and odd

pings. However, we choose to count tests on two successive pings as one test for a particular non-UXO target. For paths *h* and *i* we have used seven false targets for the eight feature sets; and for path *g* we have used ten false targets for the eight feature sets. This gives us 192 non-UXO target tests.

TABLE 2 – Number of times a particular target is a false positive on *all paths*, eliminating training targets n4 for feature sets 5 and 8 over paths *h, j* and n1, n7 for feature set 7 over path *g*.

False Positive Occurrences before Source Cal (with elim targets)											
Target	t18	t14	t17	t7	t19	n10	t5	n11	t15	CP	14%
Occurrence	6	6	6	3	4	1	0	1	0	0	
27 False positives											
False Positive Occurrences after Combined Ring Source Cal (with elim targets)											
Target	t18	t14	t17	t7	t19	n10	t5	n11	t15	CP	11%
Occurrence	4.5	4.5	4	8	0	0	0	0	0	0	
21 False positives											
False Positive Occurrences> (with elim targets)											
Target	t18	t14	t17	t7	t19	n10	t5	n11	t15	CP	13%
Occurrence	5.3	5.3	5	5.5	2	0.5	0	0.5	0	0	
24.1 false positives											
The worst offenders											
t18 – solid cement block			t14 – water-filled scuba tank			t17 – 2’ al cyl			t7 – 3’ al cyl		

Examination of Table 2 shows that before inserting source calibration into the training/testing, the false positive percentage is 14%. We consider this to be good performance since the discrimination thresholds were set to suffer no false negatives. Having to deal with only 14% false positives while guaranteeing no false negatives ($P_C = 1$) is considered excellent in most decision making applications where a missed true positive can lead to disastrous results (e.g. explosions, cancer fatalities, etc.).

As seen in Table 2, after implementing source calibration (in this case using the combined axial/radial ring calibration information), the false positive percentage drops to 11% which is about a 25% reduction. This we believe is not an insignificant drop, and it indicates the merit of applying source calibration. We point out that the calibration approach we implemented is based on only an approximate representation of the true incident source spatial distribution, and we continue to believe that development of a more accurate source projection technique in the case of the existing BOSS data or the use of an improved source in the collection of new data would result in significant

target separation improvement. Further, implementing a new, improved source would show further improvements since it could eliminate the source level dropouts.

Before moving on to more detailed examination of the performance versus path and versus feature set, we consider which false targets seem to be the worst offenders, i.e. those most confused with the set of UXOs we studied. The bottom row in Table 2 presents the false positive occurrence for each target generated by the tests averaged over the two cases, one with and one without source calibration. As can be seen, by far the worst offenders are the solid concrete block, the water-filled scuba tank, the two foot aluminum cylinder and the 3 foot aluminum cylinder. The average false positive percentage for these four targets is up to 22%. The least offenders are the 2:1 aspect wood pole section, the 5:1 aspect wood pole section, the solid rock, the flat panel, the rubber tire, and the hollowed-out cinder block. The average false positive percentage for these six targets is down to 2%.

The great performance difference between these two false target sets is roughly understandable in terms of the degree of structural acoustic similarities between these targets and the UXO. We see that in general the false targets that are almost always called correctly are the rock, the flat panel, the cinder block, and the two wooden pole sections. From the perspectives of shape (the rock and the panel) and of material (wooden pole sections), that this is the case is not surprising. Regarding shape, the UXO targets we have used for training have cylindrical shapes with aspect ratios around 4:1 which are quite different from the irregular shape of the rock with an average aspect ratio close to 1.5:1 and the flat shape of the panel. With respect to material, while the UXO targets have thick steel walls and a hard polymer filling, the wooden pole sections are much softer having sound speeds near 3500 m/s and densities near 500kg/m³ giving acoustic impedances very close to that of water. Steel materials on the other hand have speeds near 6000m/s and densities near 8000 kg/m³ with acoustic impedances 30 times that of water.

On the other hand, the false targets most often called incorrectly (i.e. as false positives) are the solid concrete block, the water-filled scuba tank, the two foot aluminum cylinder and the three foot aluminum cylinder. The last two are made of metal and have cylindrical shapes with large aspect ratios. Further, the solid concrete block material has a sound speed of ~ 3400m/s and a density ~ 2400 kg/m³ giving it an acoustic impedance ~ 8×10^6 which is significantly higher than that of water (1.5×10^6) as is the acoustic impedance of the UXO. Finally, the water-filled scuba tank like the UXO targets has a thick steel wall, is cylindrical, and has an aspect ratio of about 5:1. Further, the interior water can support acoustic waves as does the epoxy filler material, although ones with lower speeds. Given that their geometric and/or material parameters are similar to those of the UXO, that these become confused with the UXO targets is not surprising.

One target that at first does not seem to “fit the mold” described above is the cinderblock. It is made of material similar to that of the solid concrete block; and the shapes and sizes of the two targets are about the same. However, unlike the solid block (which is often registered as a false positive) the cinderblock is identified correctly as a false target. We believe that this behavior might be dependent to some degree upon the roll aspect of the cinderblock. Specifically, if the opened - cell side of the cinderblock is oriented toward the direction of the incoming acoustic wave, one would expect a small reflection component since that face is predominately water with only a skeletal concrete structure to scatter the sound. This is borne out in the bi-static acoustic

images we have generated for the cinderblock shown in Fig. 63. Another difference which can also be seen in Fig. 63 is the relatively low level of interior waves seen in the cinderblock compared to that for the rocket. If rolled 90 degrees, however, a uniform concrete slab with a 2:1 aspect ratio is now presented to the incident wave with a specular reflection similar to that from the solid block of the same dimensions. In this orientation it might then be more often called as a false positive for the reasons that pertain to the solid block.

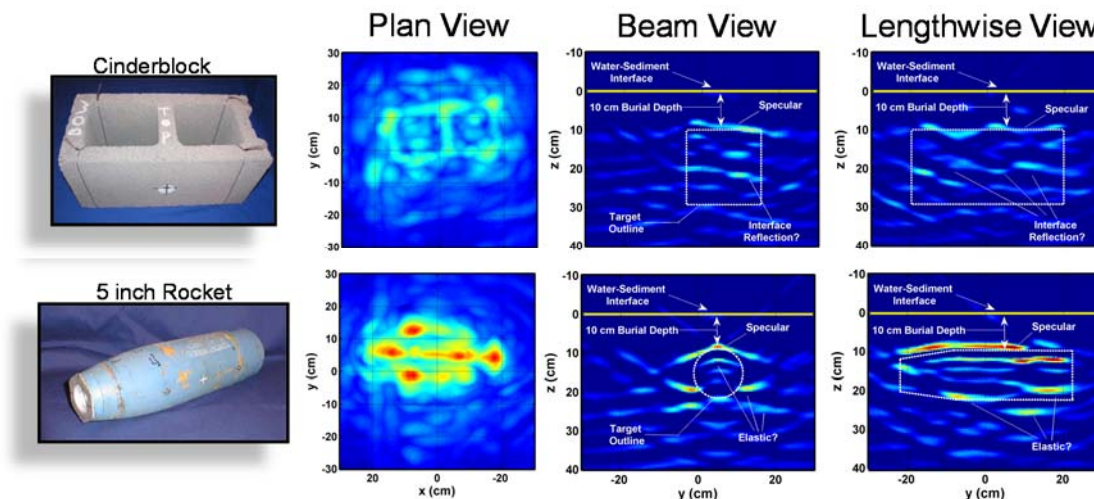


Figure 63. Acoustic images of a buried cinder block and a buried 5" rocket generated using bi-static acoustic scattering returns in the NRL sediment pool facility.

We show in Table 3 and Table 4 the performance over the individual paths for both the before-source-calibration case and the after-combined-ring-calibration case. An interesting trend is the following: for both the no-calibration and after-calibration cases, the performance for path *g* is the best followed by that for path *h*. For path *g*, the original percentage of false positives is 10% which falls to 5% after calibration. For path *h*, the original percentage is 14% which falls to 11%. And for path *i*, the original percentage of 23% falls to 20% after calibration.

TABLE 3 – False positive Occurrences (before source calibration)
Eliminating training targets n4 for feature sets 5 and 8 over paths h, j

Number of times a particular target is a false positive on path h											
Target	t18	t14	t17	t7	n10	t5	t15	14% FP			
Occurrence	3	2	2	1	0	0	0				
Number of times a particular target is a false positive on path i											
Target	t18	t14	t17	t7	n10	t5	t15	23% FP			
Occurrence	4	3.5	2	2	1.5	0	0				
Number of times a particular target is a false positive on path g											
Target	t18	t14	t17	t7	t19	n10	t5	n11	t15	CP	10%
Occurrence	0	1	2	0	4	0	0	1	0	0	
Number of times a particular target is a false positive on all paths											
Target	t18	t14	t17	t7	t19	n10	t5	n11	t15	CP	15 %
Occurrence	7	6.5	6	3	4	1.5	0	1	0	0	

Examination of Table 3 indicates that the superior performance along path g is associated with its ability to identify as non-UXOs the false targets most confused along the other two paths. Further, the relative performance as measured by *percentage* of false positives is helped on path g by the addition of three new non-UXO targets on that path which for the most part are not very similar to the UXO (a flat panel, a rubber tire, and a buried cinder block) and in fact are identified correctly. We have not determined how well these targets would be identified if added to the non-UXO list for the other two paths. Finally, we point out that on path g we train with only 5 UXO targets versus 8 on paths h and i . This may also be a factor in the improved performance along path g .

TABLE 4 – False positive Occurrences (after combined ring calibration)
Eliminating training targets n4 for feature sets 5 and 8 over paths h, j
And n1, n7 for feature set 7 over path g


Number of times a particular target is a false positive on path h											
Target	t18	t14	t17	t7	n10	t5	t15	11% FP			
Occurrence	1	2	1	2	0	0	0				
Number of times a particular target is a false positive on path i											
Target	t18	t14	t17	t7	n10	t5	t15	20% FP			
Occurrence	3.5	2.5	1	4	0	0	0				
Number of times a particular target is a false positive on path g											
Target	t18	t14	t17	t7	t19	n10	t5	n11	t15	CP	5%
Occurrence	0	0	2	2	0	0	0	0	0	0	
Number of times a particular target is a false positive on all paths											
Target	t18	t14	t17	t7	t19	n10	t5	n11	t15	CP	11 %
Occurrence	4.5	4.5	4	8	0	0	0	0	0	0	

Before moving on to imaging, we identify the following outstanding issue from the alternating ping studies: Why does elimination of certain targets from training improve performance? We had mentioned earlier that there appeared to be several problematic UXO targets used in the generative training for a few feature sets. In these cases, the tested UXO targets exhibited an outlier whose probability number fell well below those of the other UXO and several of the false targets as well causing these UXO to be called as false positives. Most often these have been UXO target n4. The cases include the Max TS feature set for combined ring source calibration over path h , the 140 frequencies x 22 positions, 2m aperture acoustic color feature set for combined ring source calibration over path h , the Max TS feature set for combined ring source calibration over path i , and the 140 frequencies x 22 positions, 2m aperture acoustic color feature set for combined ring source calibration over path i . Target n4 is the horizontally-buried 155mm projectile. It is not clear why these cases are problematic.

Turning now to the impact of lack of source calibration on acoustic images, although our image formation process begins with the de-convolution of the source function from the received scattered pressure, we do not believe that using an incorrect source function plays a significant role in the quality of the traditional acoustic images. This can be understood by the following argument. In forming an image (see Fig. 64), we use scattered data measured over a large angular aperture. Using the signals on the 40 receivers from 40 pings (40 vehicle x positions), the image (x,y,z) is produced over the band 3-20 kHz using time-delay beam-forming (algorithm on the bottom). Images at the same (x,y,z) locations (image points within a particular target) are produced from the next set of $i=6$ to $i=45$ pings. This is repeated to produce 33 realizations of each of the image points at each (x,y,z) location. The maximum image value in the 33 images at each (x,y,z) image point becomes the final image. (Resulting multi-static aperture is 10m or 90° .) Although the source function from target to target can vary in phase and amplitude in a significant way (and thus has a significant effect on the differences registered in their acoustic color constructs), given the small relative size of the targets of interest, this source pattern is approximately the same for each image point on a particular target. The thirty three source positions do produce different incident pressures, but this set of thirty three pressures is about the same for each image point within the target. The scattering strength at each target focal point is obtained by applying the appropriate time delay (determined by the acoustic speed and distance from the receiver to the focal point) and summing this result from each receiver. Thus the relative scattering strength determined for each point in the target is approximately independent of the source spatial function.

This may not be the case for the formation of the new type of image we have introduced in this program, namely the long-time (or elastic highlight) image. By this we mean a target image formed by first removing the short-time portion of the echo, which component is predominately due to specular scattering. The elastic image is then produced by predominately elastic scattering mechanisms which by their nature produce echoes spread out in time. We have successfully obtained these elastic highlight images in laboratory experiments in which the scattering geometry was predominately bi-static, a geometry which emphasized non-specular scattering. (See Fig. 65a.) We have also obtained much lower quality elastic highlight images in the Duck, NC ocean measurements where we attempted to time-gate out the specular returns. However, in the current

BOSS data case, we have been unsuccessful using this time-filtering approach to generate any reasonable elastic highlight images. This may in fact be due to failing to account for the proper source calibration information. In particular, the short-time specular component is very broadband in nature, and we have noticed that the spatial deviation of the BOSS source from one that is spherical is most significant at the higher frequencies in our band. These images will be revisited when we develop a more accurate source projection capability.



$$\text{Image Strength at } \vec{r}_i \equiv \sigma_i(\vec{r}_i)$$

$$= \frac{1}{N} \sum_{n=1}^N 4\pi |\vec{r}_n - \vec{r}_i| d_n \left(\vec{r}_n \left| \vec{r}_i, \frac{|\vec{r}_n - \vec{r}_i|}{c} \right. \right)$$

* Focal points extend 1.5m above
and 2.5m below sediment surface

Figure 64. Experiment geometry and algorithm used to generate the acoustic images.

Before leaving the subject of imaging, we report on a recent finding related to the acoustic images and a new classification feature we have discovered connected to the 2-D depth views. As indicated in the images shown in the Fig. 65b for the buried five inch rocket buried at 10cm and for the 155mm shell buried at 20cm, the targets have been buried horizontally along the x direction. The plan view and x,depth 2-D acoustic images are consistent with this. The y,depth 2-D image, however, shows a target slanted at about -45 degrees from the vertical.




Figure 65a. Elastic highlight images: (upper) bi-static sediment pool scanning; (middle) time gating of Duck, NC rail data; (lower) time gating Boss data in Gulf.

This we believe is an exploitable artifact which exists for cylindrical targets and as shown in Fig. 65c is related to the generation of circumferential creeping waves or circumferential elastic waves which are generated near the point on the cylindrical target of specular reflection and travel around the circumference of the cylinder with a delay time equal to the travel time around the circumference, or $\Delta t = \pi D/C$ where D is the cylinder diameter and C is the velocity of the circumferential wave. The wave then reradiates back to the receiver from the point it was first launched but with the delay time $\sim \Delta t$. With much better resolution than we achieve here, this would produce two imaged highlight points separated by a distance $\sim D$. However, for the limited cross-range resolution we obtain, this creates a smeared image in the y,depth plane which appears as a strip oriented in the direction of the incident wave and having a length (given for the case of a creeping wave) by approximately $2.5D$. In the case of a circumferential elastic wave, the length of the artifact image would be somewhat less due to the faster propagation speed compared to the creeping wave.





Figure 65b. The three 2-D images of the horizontally buried five inch rocket (left) and the 155mm shell (right) produced from the E-W flight path h .



Path difference between creeping wave return and specular echo

$$\Delta = D/2 + \pi D/2 + D/2 = (1 + \pi/2)D$$

Figure 65c. Diagram explaining how creeping or elastic circumferential wave can produce a time-extended scattered echo and thus an under-resolved extended 2D image at an angle in the z,y plane (lower). Y, Z and X,Z images (upper).

SUBTASK 6.1: Develop training methods for the selected identification algorithms.

We considered two general approaches to training the RVM classification algorithms prior to deploying them in an actual test or clean-up operation. (1) The first is based on training with numerically generated UXO scattering data using existing, validated finite element structural acoustic models available to us. (2) The second involves training with direct measurements made by flying the actual sonar system on linear flight paths over a field of surrogate UXO targets of the type expected to be encountered at the test site. The targets would be buried in a training site having similar sediment and acoustic propagation characteristics to the test site.

In order to evaluate whether either of these two approaches represents a viable method for performing the RVM classification algorithm training, we first carried out a simple study that would allow us to estimate the accuracy (probably more correctly, repeatability) required for the acoustic color training data.

We centered this study on the 20m aperture acoustic color classification results for the alternating ping, single path case using path h and combined ring calibration. As a rough simulation of error, we added random noise to the measured acoustic color plots as a function of position of each receiver, with the average magnitude of the added noise given as a percentage of the maximum color value along each receiver path for each target.

In the following figures (Figs. 66 -70) we show the results of these studies both in terms of the actual probability calls (on the left) and the associated ROC curves (on the right). Each figure shows the results for two cases having sequential values for the added noise given as a percentage of the maximum color value. The first case shown in Fig. 66 (upper plots) is for the no-added noise case originally shown in Fig. 40. One can examine how the correctly called non-UXO points move closer to the discrimination threshold as the added noise level is increased and the cross over creates more false positive calls. The ROC curves also show the drop in performance as the added noise is increased.

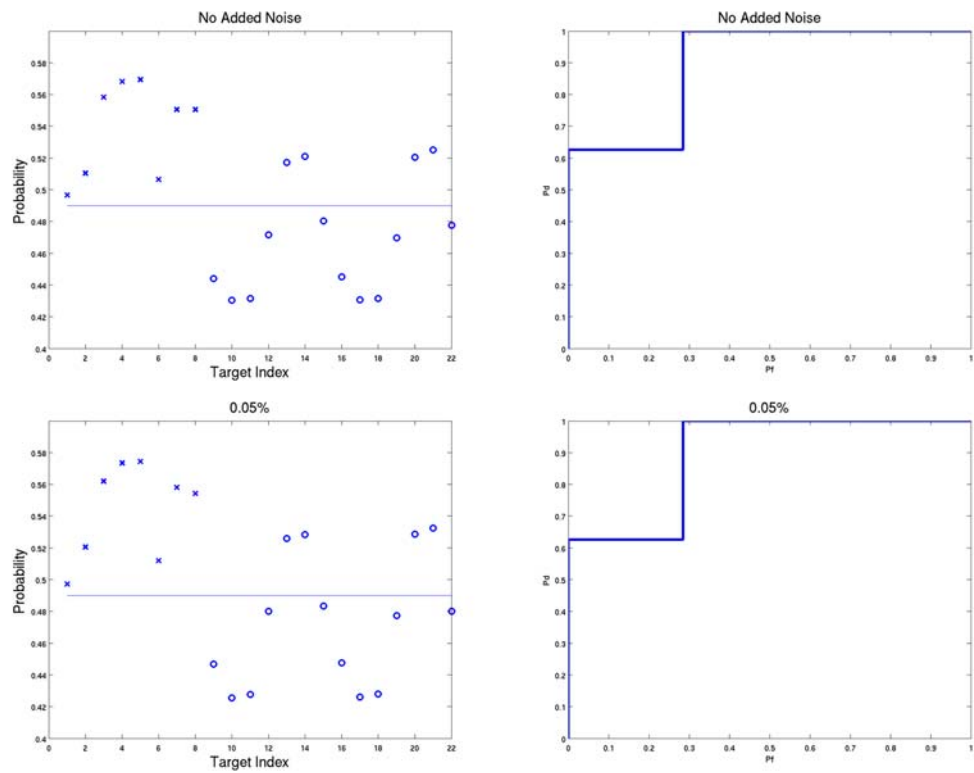


Figure 66. Generative acoustic color 20m aperture feature separation (left) and corresponding ROC curves (right) using alternate pings over path h for combined ring source calibration. No added noise (upper) and 0.05% max added noise (lower).

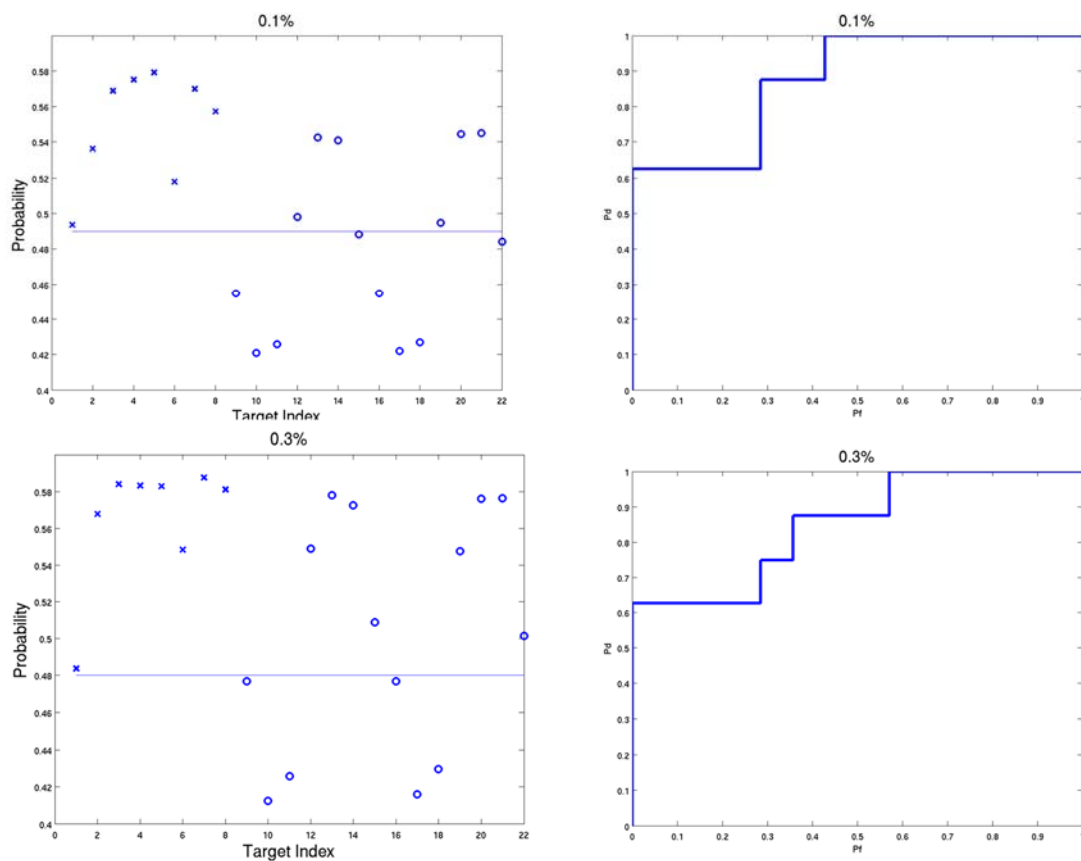


Figure 67. Generative acoustic color 20m aperture feature separation (left) and corresponding ROC curves (right) using alternate pings over path h for combined ring source calibration. 0.1% max added noise (upper) and 0.3% max added noise (lower).




Figure 68. Generative acoustic color 20m aperture feature separation (left) and corresponding ROC curves (right) using alternate pings over path h for combined ring source calibration. 0.5% max added noise (upper) and 1% max added noise (lower).

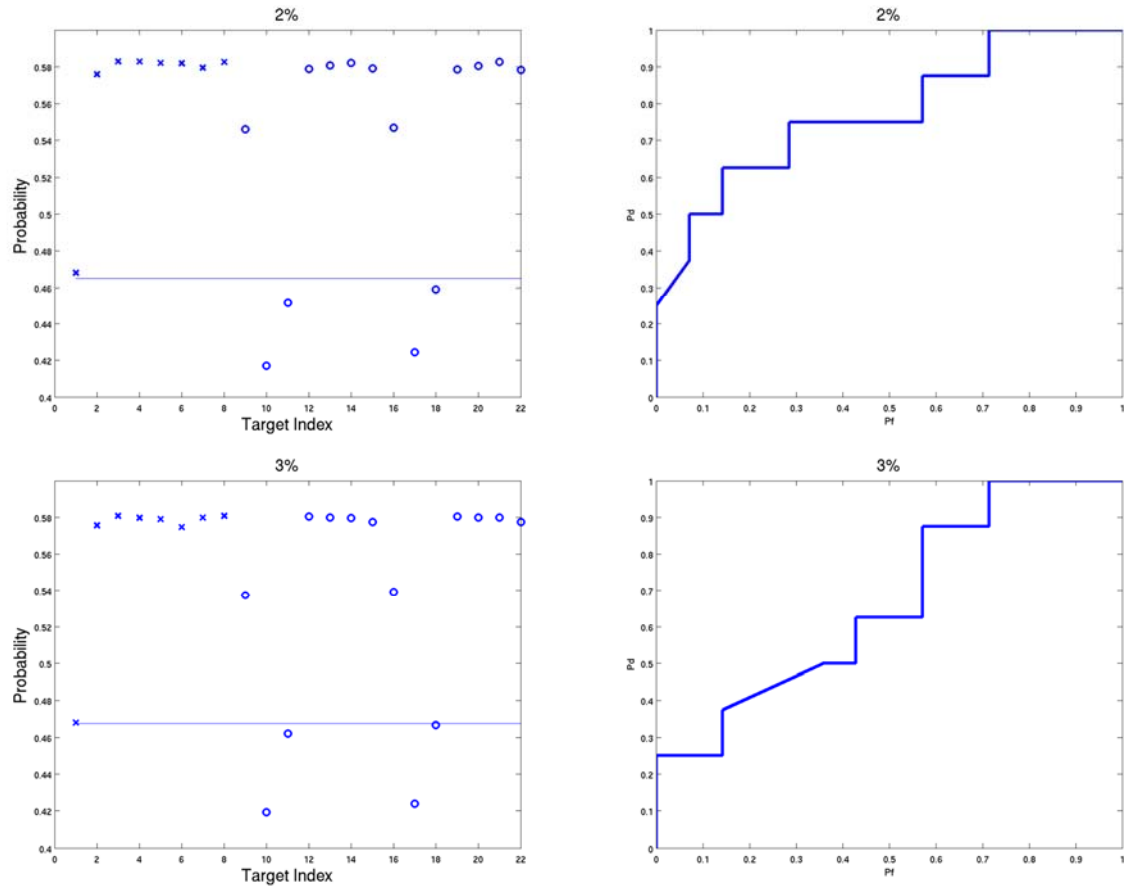


Figure 69. Generative acoustic color 20m aperture feature separation (left) and corresponding ROC curves (right) using alternate pings over path h for combined ring source calibration.. 2% max added noise (upper) and 3% max added noise (lower).




Figure 70. Generative acoustic color 20m aperture feature separation (left) and corresponding ROC curves (right) using alternate pings over path h for combined ring source calibration. 3.5% max added noise (upper) and 1% max added noise (lower).

We summarize the results from the above figures by plotting the function $\tilde{A} = 1 - \text{area under the ROC curves}$ versus the increased noise percentages, and this is shown in Figure 71.

Based on the results in Fig. 71, and somewhat arbitrarily taking the maximum acceptable percentage decrease in the area under the ROC curve as 5%, we have that the maximum tolerable percentage error in acoustic color would be 0.3% of the maximum acoustic color value in the band. To help put this and the curve in Fig. 71 into perspective, we will illustrate the impact of these results by considering the experimentally determined acoustic color spectrum measured for the five inch rocket as shown in Fig. 72.

In the Gulf measured data base used to generate the results shown in Figs. 66 - 70, the plot of field value for each UXO for the acoustic color value measured at each BOSS receiver divided by the

1 - area under ROC curve




Figure 71. $1 - \text{area under ROC curves}$ versus added noise level in % max from Figs. 66 - 70.

average value in the spectral band varies from about 20 to 300 as shown in Fig. 73.

If we add random levels of 0.3% of the *maximum* field value to the acoustic color map then we are adding values as large as $0.9 \times$ the *average* values. (See plot in Fig. 73.) With reference to Fig. 72, we determine that the angular average TS @ 5kHz is about -23 dB. Using the 0.9 factor above we would then have $-23\text{dB} - 1\text{dB} = -24\text{dB}$ for the maximum added TS fluctuations we could tolerate. This implies that the background deep blue levels (-50dB) could at most become yellow (-24dB) while the peak orange-red levels (-18dB) could at most become red (-14.5dB).

Considering these results and trying to generalize, we concluded that in general the accuracy and efficiency expected from *simulating* acoustic color data for representative UXO targets using current state-of-the-art finite element-based structural acoustics models and propagation codes is probably not yet sufficient for planning to train the classification algorithms using simulated, numerical data. This limitation stems from several issues. The first is the basic fidelity of the finite element model especially at the higher frequencies given that we know the required UXO material, geometric, and boundary condition parameters in the first place. The second and more important consideration is the likelihood that we will not know these UXO details sufficiently well. Further, there will be additional errors related to the sonar system itself together with its deployment conditions. The former include the actual source spatial characteristics as discussed here previously while the latter include the degree to which one can control the particular vehicle flight paths used or realize the flight paths intended. Finally, in general modeling of acoustic propagation in the water column and sediment could introduce additional errors as well.




Figure 72. Five inch rocket target strength versus incidence angle as measured in large pool facility.




Figure 73. The plot of field value for each UXO for the acoustic color value measured at each BOSS receiver divided by the average value in the spectral band in the Gulf measured data base used to generate the results shown in Figs. 66 – 70.

Based on the above, we believe that at this time the better way to train a Boss-like system is to carry out preliminary measurements with the actual sonar system on buried surrogate targets of the type expected to be encountered in the actual clean-up operation. These measurements would use linear tracks as would the test measurements, and they would be carried out in a similar environment or if possible in a smaller region of the actual test site.

SUBTASK 6.2: Demonstrate effectiveness of trained identification methods in separating UXO from clutter.

In order to access the second training approach, i.e. one based on surrogate measurements, we carried out test studies using the available data from the BOSS Gulf exercise. In particular, we carried out what are called “Leave One Out” studies to simulate a realistic train and test procedure at a particular site using the existing BOSS Gulf data set. The argument is that using this data from one path to train and then testing with data collected on a different path (even though carried out at the same site) replicates approximately the uncontrollable variations (apart from different targets) such as look angle, range, position and propagation fluctuations, etc. that we would experience in using any pre-trained algorithm in an actual exercise.

In carrying out the “Leave One Out” procedure using the BOSS Gulf data, we train using UXOs sampled over one path (say path *g*) and then test on two other nearby paths (*h* and *i*). Unlike our previous studies using alternate pings for training and testing, in this procedure we obtain generative classification results when training and testing on *all* the pings on a particular path. Since the distance traveled between pings is now reduced by a factor of 2 (compared to odd or

even pings), the maximum frequency can be increased by a factor of 2 and still fall within the Nyquist sampling rate limit. Accordingly, the frequency range for these runs is now 3 to 13 kHz.

We carried out a number of “Leave One Out” studies on each of the eight feature sets using paths g , h , and i where we trained and tested using each of the three path combinations. In all cases, the performance is very poor for all the feature sets except the 200 frequencies x 22 positions 20m and 2m aperture acoustic color feature sets. Accordingly, in the following we will present only the acoustic color feature results.

We show in Fig. 74 one result for the case in which we train over path g and test over paths h and i for the 200 frequencies x 22 positions 20m aperture acoustic color feature set. Here the no calibration case is shown in the upper left, the radial ring calibration case in the upper right, the axial ring calibration case in the lower left, and the combined ring calibration case in the lower right. The data display in Fig. 74 is arranged and labeled in the following way. The x’s relate to UXO targets and the o’s to non UXO targets. The training which is carried out over path g uses UXO targets $n7$, $n8$, $n1$, $n2$, and $n3$. The order of x’s in Fig. 37 is as follows: $n4$, $n5$, $n6$, $n7$, $n8$, $n1$, $n2$, and $n3$ tested over path h then the same targets tested over path i . The order of o’s in Fig. 74 is as follows: $n10$, $n11$, $t5$, CP , $t15$, $t18$, $t19$, $t14$, $t7$, $t17$ tested over path g (for completeness), $n10$, $t5$, $t15$, $t18$, $t14$, $t7$, and $t17$ tested over path h , and then the same targets tested over path i . One interesting difference seen here compared to the results for training and testing over the same path is that the worst performance (zero) is found for the combined ring calibration case while a modest performance is seen for the no calibration and radial ring calibration cases. Further, while the 200 frequencies x 22 positions acoustic color feature sets perform the worst for the test and train on the same path studies, they perform best here in the “leave one out” scenario. In part, this may be related to the use of twice the bandwidth used previously since this acoustic color feature is expected to perform better when the band is opened to include additional structural acoustic echo mechanisms.

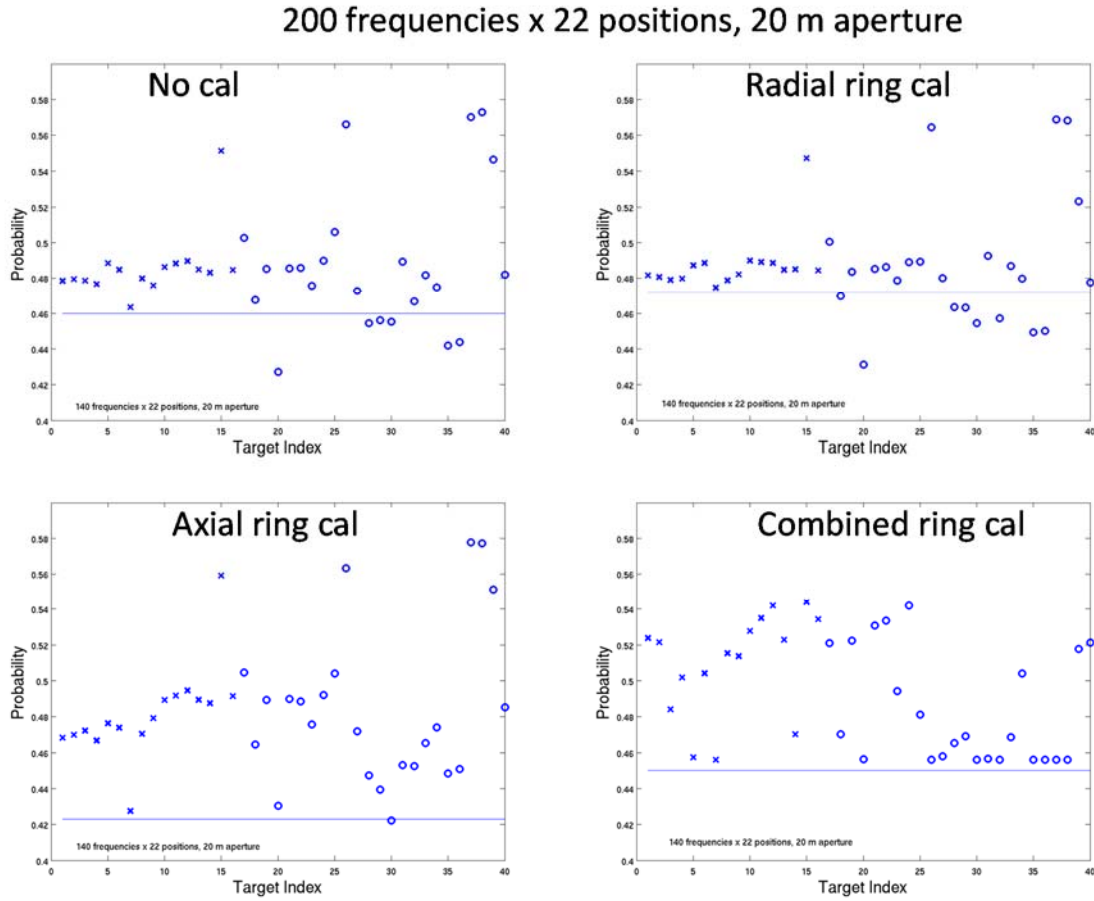


Figure 74. Generative target feature separation training on path *g* and testing on paths *h* and *i* for the 200 frequencies x 22 positions, 20m aperture feature set over the band 3-13kHz.

The best overall performance is seen for the no calibration and radial ring calibration cases. A more detailed examination of Fig. 74 for these cases shows the correctly called non-UXO targets are CP, n11, and t5 over path *g* (the training path), t5, t7, t15, and t18 over path *h*, and t5 and t15 over path *i*. We point out that in the previous single path studies, CP, n11, t5, and t15 are the non-UXO targets most often correctly called. In Fig. 74, the non-UXO targets that perform the worst (highest probability of being miss-called as a UXO) are targets t7, t14, t17, and t18. It is interesting to note that these are the very same four non-UXO targets that were most often miss-called as UXOs in the previous single path studies. Although not shown here, the results for the 200 frequencies x 22 positions, 2m aperture feature set although somewhat worse are similar to the 20m aperture case. In particular, the correctly called non-UXO targets are CP and n11 on path *g* and t5 and t15 on path *i*. Again, these are the non-UXO targets most often called correctly in the previous single path studies.

Next we discuss the “Leave One Out” case for training on path *h*. We show this result in Fig. 75 for the 200 frequencies x 22 positions 2m aperture acoustic color feature set where again the no calibration case is shown in the upper left, the radial ring calibration case in the upper right, the axial ring calibration case in the lower left, and the combined ring calibration case in the lower right. For this path, the 2m aperture result is somewhat better than those for the 20m aperture (and for the other feature sets as well). The data display in Fig. 75 is arranged and labeled in the following way. As before, the x’s relate to UXO targets and the o’s to non UXO targets. The training which is carried out over path *h* uses UXOs targets n4, n5, n6, n7, n8, n1, n2, and n3. The order of x’s in Fig. 75 is as follows: n7, n8, n1, n2, and n3 tested over path *g* and then n4, n5, n6, n7, n8, n1, n2, and n3 tested over path *i*. The order of o’s in Fig. 38 is as follows: n10, t5, t15, t18, t14, t7, t17 tested over path *h* (for completeness), n10, n11, t5, CP, t15, t18, t19, t14, t7, t17 tested over path *g*, and then n10, t5, t15, t18, t14, t7, and t17 tested over path *i*. Examination of Fig. 75 shows that for the no calibration, radial ring calibration, and axial ring calibration cases there are three non-UXOs called correctly, viz. targets CP, t5 and t15. These targets were always called correctly in the single path studies as well. Further, unlike the results for training on path *g*, the combined ring calibration case does the best and correctly calls six targets, t5 and t15 on path *i*, CP and n11 on path *g*, and t5 and t15 on path *h* (the training path). Except for n11, all these targets were called correctly in the previous single path studies while n11 was missed less than 2% of the time.

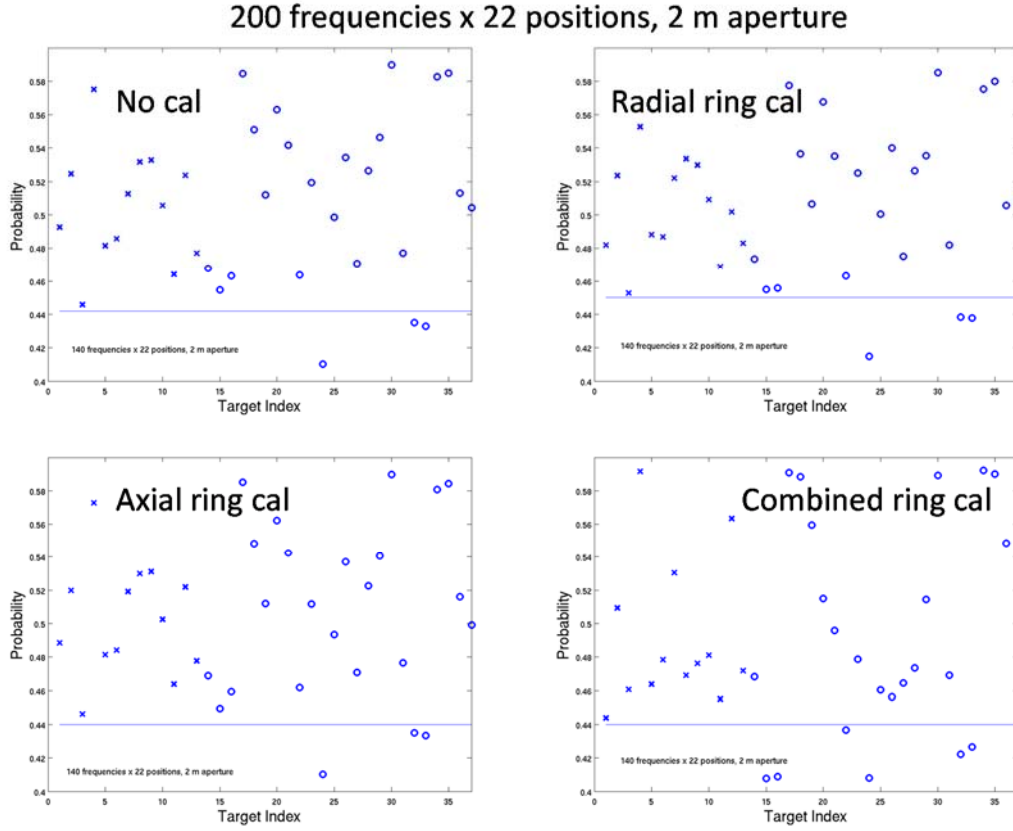


Figure 75. Generative target feature separation training on path h and testing on paths i and g for the 200 frequencies x 22 positions, 2m aperture feature set over the band 3-13kHz.

For comparison purposes, in Fig. 76 we show the results for the 200 frequencies x 22 positions, 20m aperture acoustic color feature set for training on path h . In contrast to the 2m aperture case, the combined ring calibration case has virtually zero correct non-UXO calls. However, similar to the path g training case for this feature set, there are three correct non-UXO target calls viz. CP, t5, and t15 for the no calibration, radial ring calibration, and axial calibration cases. Again, these targets were called correctly 100% of the time in the single path studies.

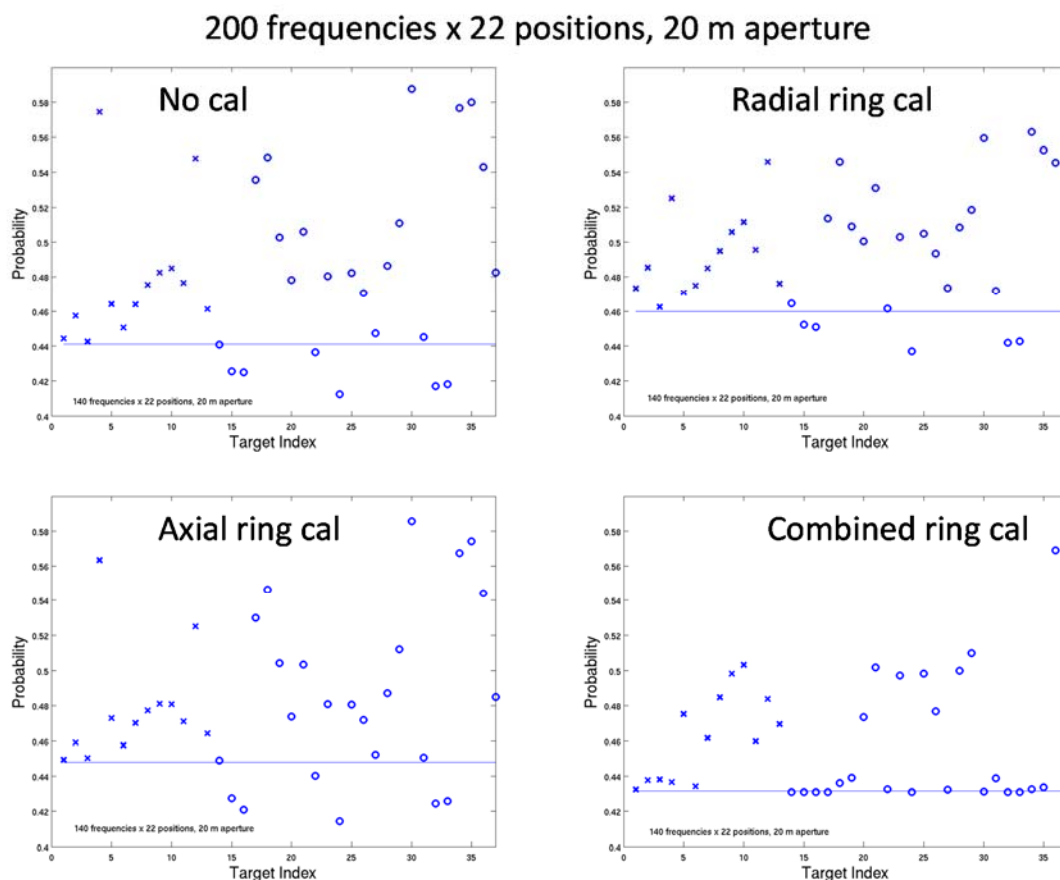


Figure 76. Generative target feature separation training on path h and testing on paths i and g for the 200 frequencies x 22 positions, 20m aperture feature set over the band 3-13kHz.

Next we discuss the “Leave One Out” case for training on path i . We show this result in Fig. 77 for the 200 frequencies x 22 positions 2m aperture acoustic color feature set where again the no calibration case is shown in the upper left, the radial ring calibration case in the upper right, the axial ring calibration case in the lower left, and the combined ring calibration case in the lower right. For path i as in path h , the 2m aperture result is somewhat better than those for the 20m aperture (and for the other feature sets). The data display in Fig. 77 is arranged and labeled in the following way. As before, the x’s relate to UXO targets and the o’s to non UXO targets. The training which is carried out over path i uses UXOs targets n4, n5, n6, n7, n8, n1, n2, and n3. The order of x’s in Fig. 40 is as follows: n7, n8, n1, n2, and n3 tested over path g and then n4, n5, n6, n7, n8, n1, n2, and n3 tested over path h . The order of o’s in Fig. 77 is as follows: n10, t5, t15, t18, t14, t7, t17 tested over path i (for completeness), n10, n11, t5, CP, t15, t18, t19, t14, t7, t17 tested over path g , and then n10, t5, t15, t18, t14, t7, and t17 tested over path h . Examination of Fig. 77 shows that for the no calibration, axial ring calibration, and combined ring calibration cases there are three non-UXOs called correctly, viz. targets CP, t5 and t15. These targets were always

called correctly in the single path studies as well. Further, unlike the results for training on paths *g* and *h*, the radial ring calibration case does the best and correctly calls six targets, t5 and t15 on paths *i* (the training path) and *h* and n11 and CP on path *g*. As before, except for n11, all these targets were called correctly in the previous single path studies while n11 was missed less than 2% of the time.




Figure 77. Generative target feature separation training on path *i* and testing on paths *h* and *g* for the 200 frequencies x 22 positions, 2m aperture feature set over the band 3-13kHz.

The results for training on path *i* for the 200 frequencies x 22 positions, 20m aperture acoustic color feature set is shown in Fig. 78. As can be seen, the axial and combined calibration cases perform very poorly (only one correct false target call between them). The no calibration case has four correct false target calls (t5 and t15 on path *i*), CP on path *g*, and t18 on path *h*. The radial calibration case also has four correct false target calls (t5 on path *i*, CP on path *g*, and t15 and t18 on path *h*). Except for t18, these are the same targets that are called correctly in both the single path and previous “leave one out” cases.

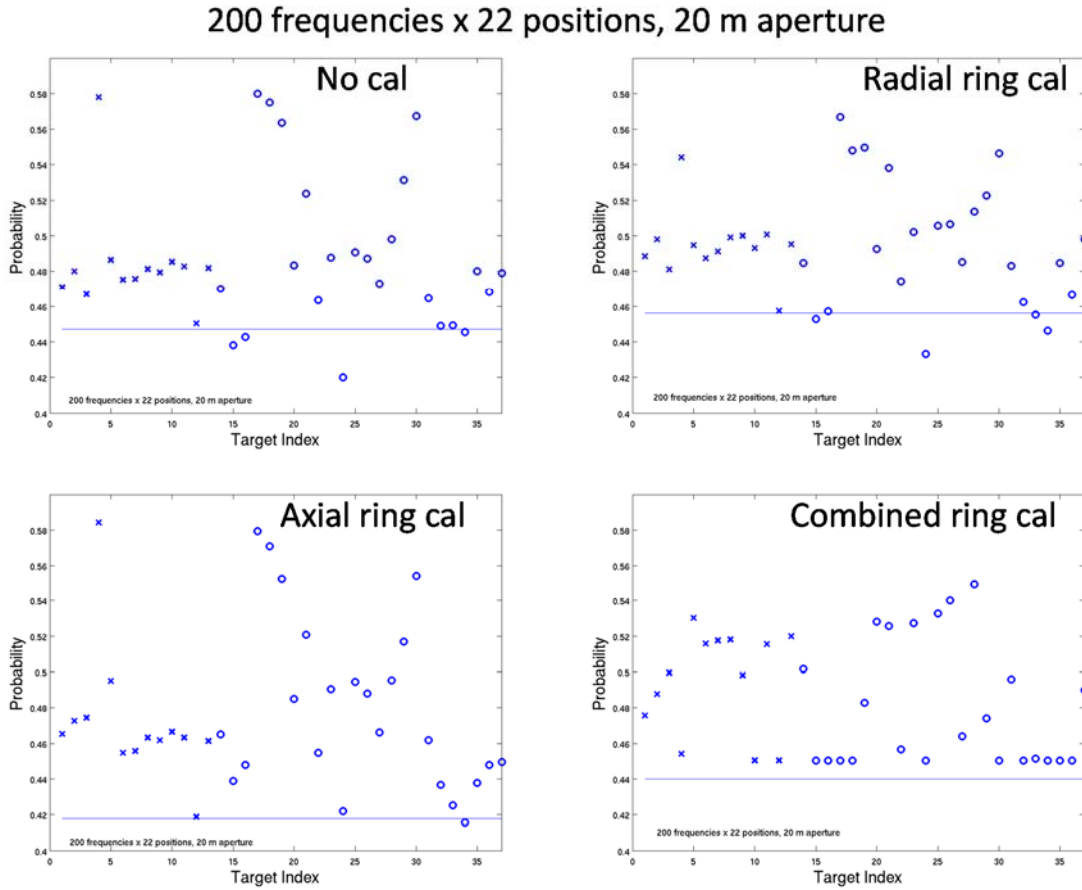


Figure 78. Generative target feature separation training on path i and testing on paths h and g for the acoustic color, 20m aperture feature set over the band 3-13kHz.

In the earlier single path, alternating ping studies, we considered cases in which we eliminated some training targets that seemed problematic. In the same vein, we look at eliminating targets n1 and n7 from the training over path g . Except for the 200 frequencies x 22 positions 2m acoustic color feature set, there was not much difference. We show the results for that case in Fig. 79 where we see very consistent performance for all but the combined ring calibration case. The correctly called non-UXOs are again n11, CP, t5, and t15.

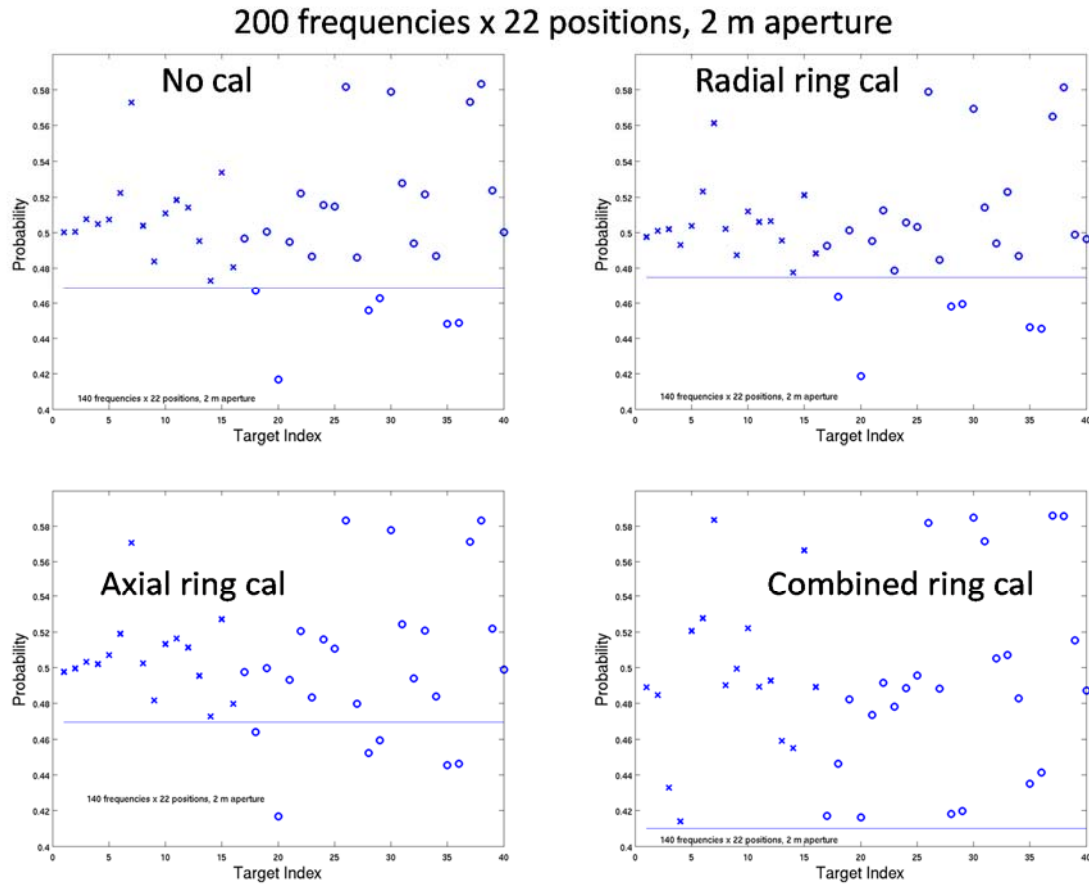


Figure 79. Generative target feature separation training on path *g* (without n1 and n7) and testing on paths *h* and *i* for the 200 frequencies x 22 positions, 2m aperture feature set over the band 3-13kHz.

Along the same lines, we look at the results for eliminating n4 from the training over path *h* as was done in the single path, alternating ping studies. Again, there is not much effect except for the 200 frequencies x 22 positions acoustic color feature sets. The one for the 20m aperture is shown in Fig.80. The combined calibration case has no performance, but the other three consistently have correct calls on non-UXO targets n10, n11, CP, T5, and t15. In the single path studies, all these targets also exhibited near zero false positives.

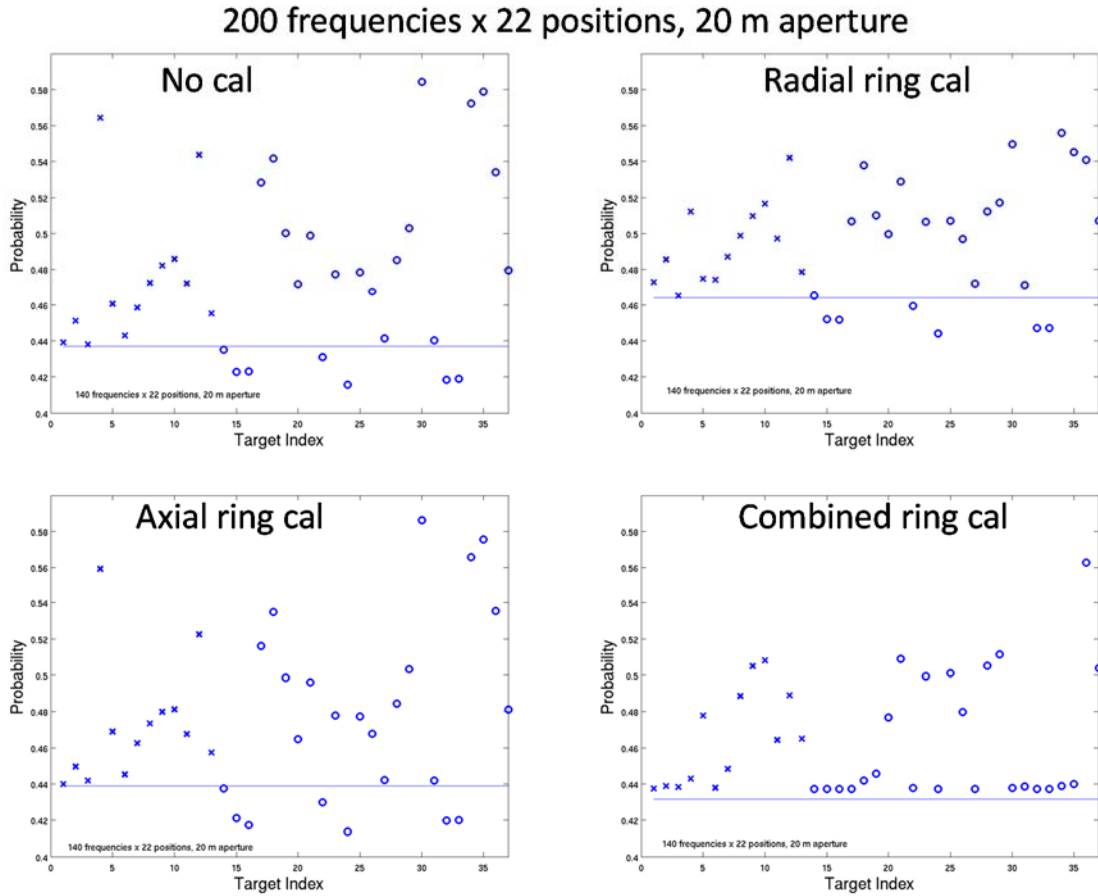


Figure 80. Generative target feature separation training on path h (without n_4) and testing on paths i and g for the 200 frequencies x 22 positions, 20m aperture feature set over the band 3-13kHz.

In Fig. 81 is shown the similar result for the 2m aperture case after eliminating n_4 from the training. As can be seen, there is some improvement in the radial calibration case. Here the correctly called non-UXO targets include n_{10} , n_{11} , t_5 , t_{15} , the most often correctly called non-UXO targets; but they also now include t_{14} and t_{18} which are the most often incorrectly called targets.

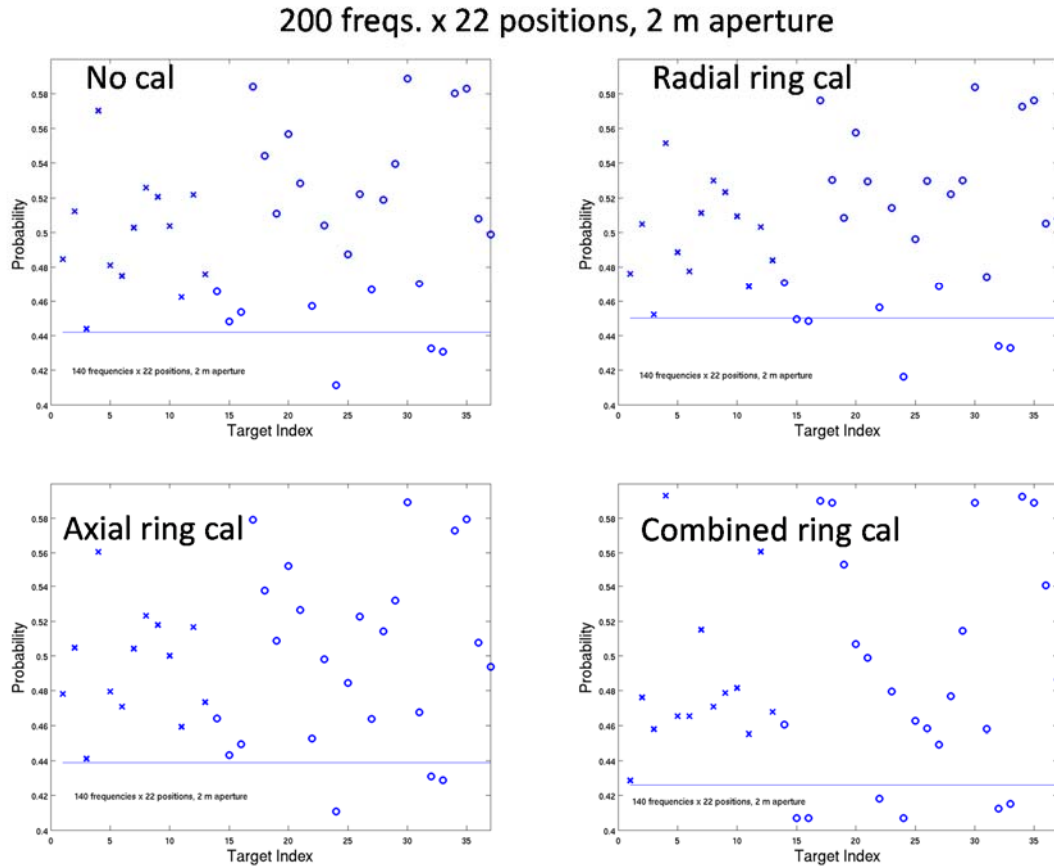


Figure 81. Generative target feature separation training on path h (without n_4) and testing on paths i and g for the 200 frequencies x 22 positions, 2m aperture feature set over the band 3-13kHz.

Some of the findings in these “Leave One Out” studies (together with the previous single path studies) are as follows. (1) For all eight feature sets, but especially for the first six, the good performance seen in the single path, alternate ping procedure is not maintained in the multiple path “Leave One Out” scenario. (2) For the most part, in the latter there is no systematic improvement seen by applying source calibration using an approximate description of the source spatial pattern as there was in the single path studies. (3) One consistently supported result is that of the ten non-UXO targets in the data, those most often called correctly are: n_{10} , n_{11} , CP, t_5 , and t_{15} . These are the buried rock, the buried cinder block, the flat panel, the 5:1 wooden pole section, and the 2:1 wooden pole section. (4) Further, in the “Leave One Out” studies unlike in the single path studies, the best performance is associated with one or the other of the two 200 frequencies x 22 positions acoustic color feature sets. As mentioned earlier, this may be in part associated with the fact that the bandwidth in the “Leave One Out” studies is double that used in the single path studies. (5) Another point is that in the “Leave One Out” studies, some of the more difficult non-UXO targets, i.e. those that are called as false positives in the 200 frequencies x 22 positions acoustic

color feature sets, fall near the discrimination threshold in many of the other feature sets. These include t14, t18, and to a lesser extent t7. These are the solid concrete block, the scuba tank, and the 3 foot aluminum cylinder. This is somewhat encouraging and suggests that it might be profitable to explore the benefit of focusing the identification of these type targets on the first six feature sets, i.e. (1) 12 TS averages, (2) 7 TS averages, (3) Max TS, bin #, and standard deviation, (4) Max TS and bin #, (5) Max TS, and (6) Mean TS and identification of the other non-UXO targets on the 200 frequencies x 22 positions, 20m and 2m aperture feature sets. However, we have not yet explored this idea.

In the following, we continue to focus specifically on the two 200 frequencies x 22 positions feature sets in the “Leave One Out” scenarios. In particular, we generate the statistics on the various non-UXO targets over the three paths. This is shown in Table 5 where we list the number of calls made on each target for each of the three paths, the number of false positives (FP), and the % of false positives. For paths *g* and *h*, we use the 200 frequencies x 22 positions, 20m aperture feature set and for path *i* the 200 frequencies x 22 positions, 2m aperture feature set since these present the best performance of all the feature sets. In all three cases, the numbers in Table 4 refer to the radial ring calibration case, which also demonstrates the best performance compared to the no calibration, the axial ring calibration, and combined ring calibration cases.

As can be seen, for this set of parameters targets n11 and CP are called correctly each time (two and three, respectively) they are presented. Targets t15 and t5 are called correctly 33% of the eight times they are presented. These results are consistent for all three paths. With one exception, all other targets are miss-called 100% of the time. The exception is target t18 which is miss-called only 33% of the time over path *g*.

TABLE 5 – False positive calls for various targets in the “Leave One Out” tests.
Eliminating training targets n4 for feature sets 5 and 8 over paths *h*, *j*
and n1, n7 for feature set 7 over path *g*

Target		n11	CP	t18	t15	t5	t7	n10	t19	t17	t14
<i>g</i>	FP#	0	1	2	2	2	1	3	1	3	3
	FP%	0%	0%	33%	33%	33%	50%	100%	100%	100%	100%
<i>h</i>	FP#	1	1	2	3	3	3	3	1	2	3
	FP%	0%	0%	100%	33%	33%	100%	100%	100%	100%	100%
<i>i</i>	FP#	1	1	3	3	3	3	3	1	3	3
	FP%	0%	0%	100%	33%	33%	100%	100%	100%	100%	100%

What we can conclude with some caution about experimentally-based training from this limited study is that the approach may be effective for certain targets (the buried cinder block, the flat plate, and the two wooden pole sections) but not for others (solid cement block, the aluminum cylinders, the tire, and the buried rock). Further consideration shows that three of the problematic targets (the tire, the buried rock, and the solid block) are correctly called (at least by a small margin)

in one of the other feature sets. As mentioned previously, this suggests that the simultaneous use of *all* the feature sets in implementing target identification could have some merit.

We considered these results with respect to establishing more meaningful threshold levels for the runs. Up until now, we have established the threshold level by simply marking it at the UXO target with the lowest probability which ensures that there are no false negative calls. This is done no matter on which path the UXO target tests are carried out and no matter which UXO targets were used in the generative training; and we have not included “tests” using the training targets on the actual training path. So for example, when we trained on path *g* using *n7*, *n8*, *n1*, *n2*, and *n3* we tested on paths *i* and *h* including these targets as well as the other UXOs. But we did not include tests of *n7*, *n8*, *n1*, *n2*, and *n3* on the training path (path *g*). One way to establish a more meaningful threshold might be to include results for the UXOs used in the training on say path *g* in the tests on path *g*. The lowest probability of these would be used to establish the thresholds for the other paths. Perhaps a better approach would be to set the threshold using only targets that were used in training but as they are tested on the other two paths. Switching to this latter approach could impact the results we have presented for training on path *g* (but not on the others) since not all UXO were used in training on path *g*. What needs to be done in this case is to eliminate the *x*’s not associated with the training targets when establishing the thresholds with the tested data. When we in fact eliminated these targets we found that there was no change in the threshold level.

The fact that the performance results for the “Leave One Out” studies are worse than those obtained training and testing on alternating pings is to be expected since the former represents a much more challenging scenario regarding changing conditions between training and testing. Alternating pings are only separated by 5cm whereas the paths are separated by as much as 5 meters. Among the other reasons for this difference in performance is an inaccurate description of the incident sound. This must surely play an important role in training and testing on different paths given the significant source pattern fluctuations with position we have documented for the BOSS source. And notwithstanding our efforts to remove these source variations, we do not believe any of our approaches do a good job in mitigating this effect particularly at the higher frequencies. To compound this issue, the doubling of the bandwidth used in the “Leave One Out” scenario would accentuate this problem.

After we had completed the studies on inputting source calibration information described above, we were able to complete development of the equivalent source method for projecting the ring source calibration data onto the sediment plane. While we did not have the resources to go back and use this new approach alongside the others for the many cases examined, we were able to




Figure 82. Generative acoustic color feature separation training on path *g* and testing on paths *h* and *i* over band 3-13kHz after equivalent source calibration.

apply this technique to the following “Leave One Out” case for training on path *g* using *n2*, *n3*, and *n8* as before. For the acoustic color feature with 2m aperture we get the results shown in Fig. 82 for the original band of 3 – 13kHz. The data display in Fig. 82 is arranged and labeled in the following way. The *x*’s relate to UXO targets and the *o*’s to non UXO targets. The order of *x*’s in Fig. 82. is as follows: *n4*, *n5*, *n6*, *n7*, *n8*, *n1*, *n2*, and *n3* tested over path *h* then the same targets tested over path *i*. The order of *o*’s in Fig. 82 is as follows: *n10*, *n11*, *t5*, *CP*, *t15*, *t18*, *t19*, *t14*, *t7*, *t17* tested over path *g* (for completeness), *n10*, *t5*, *t15*, *t18*, *t14*, *t7*, and *t17* tested over path *h*, and then the same targets tested over path *i*. Compared to the original result for the no calibration case, there is a slight improvement in that the correct false negative, *n10* (the buried rock), now falls well below the threshold. As can be seen earlier in Fig.79

which used the 20m aperture acoustic color feature, for the no calibration case *n10* falls right on the threshold line. The results in Fig. 82 are summarized in Table 6.

TABLE 6 - False Positives (FP) and True Negatives (TN) 3-13.5kHz

	n10 buried rock	n11 buried cinder block	t5 5:1 pole	t15 2:1 pole	cp panel	t18 solid block	t19 tire	t14 scuba tank	t7 3' al cyl	t17 2' al cyl
# FP	3	0	1	1	0	3	1	3	3	3
# TN	0	1	2	2	1	0	0	0	0	0

Next we show in Fig. 83 the results using the identical parameters as for Fig. 82 except that we increase the bandwidth, now covering 3 – 20kHz. We are able to do this because the method of superposition actually suppresses the aliasing effect which limited us in the previous cases to staying below 13.5 kHz. As can be seen in Fig. 83, there is a definite performance improvement in that there are now eight correctly called false targets. The results are summarized in Table 7.

We point out that while the results summarized in Table 7 indicate a P_{FA} as high as ~60%, they are based on a P_C of 100%. (See Fig. 82.) In this regard, it is also interesting to point out that this perfect P_C performance includes correct UXO calls on UXOs buried at angles not included in the

training set. It is also important to realize that the high P_{FA} above includes false targets that are similar in material and/or shape to the two UXO used for training. Given the aforementioned source problems, we expect to perform poorly against these UXO-like false targets. In fact, when we eliminate these targets (scuba tank, aluminum cylinders, and solid concrete block), the P_{FA} drops to 33% which is close to our overall performance goal.

It is also important to realize that the high P_{FA} above includes false targets that are similar in material and/or shape to the two UXO used for training. Given the aforementioned source problems, we expect to perform poorly against these UXO-like false targets. In fact, when we eliminate these targets (scuba tank, aluminum cylinders, and solid concrete block), the P_{FA} drops to 33% which is close to our overall performance goal.




Figure 83. Generative acoustic color feature separation training on path g and testing on paths h and i over band 3-20kHz after equivalent source calibration.

TABLE 7 - False Positives (FP) and True Negatives (TN) 3-20kHz

	n10 buried rock	n11 buried cinder block	t5 5:1 pole	t15 2:1 pole	cp panel	t18 solid block	t19 tire	t14 scuba tank	t7 3' al cyl	t17 2' al cyl
# FP	2	0	1	1	0	3	0	3	3	3
# TN	1	1	2	2	1	0	1	0	0	0

The major motivation for carrying out the “Leave One Out” studies is to evaluate whether this scenario might form the basis for a practical training approach in the field. In this regard, we believe the results shown in Table 7 are very encouraging in that they indicate that we can expect good performance using an RVM identification algorithm experimentally trained beforehand using known proud and buried UXO targets in an underwater site having similar sediment properties as the actual test site. While the results above validate this conclusion regarding non-UXO targets that are structural acoustically dissimilar to the UXO, we believe this would also be the case for the broader set of non-UXO were one to use a better designed sonar source. This follows in part from the fact that we are using only an approximate source calibration technique, one we believe tends to “wash out” the acoustic color differences between UXO and non-UXO. Further, the serious source drop outs with angle and frequency, which cannot be corrected even with a good

calibration technique, lead to S/N issues. These two problems will be eliminated once a spatially and spectrally uniform source is used; and under these conditions we expect significantly improved performance for all non-UXO targets.

CONCLUSIONS

The following conclusions are based on the results reported both here and in the MR2103 Final Report¹⁰ dated April 2014.

Major Conclusions

(1) Image formation of buried targets using BOSS SAS-processed scattering data can be used to garner approximate dimensions, burial orientation, and burial depth of typical UXO targets.

(2) Overall, we have studied eight different features sets for classification, and all of these were shown in principle able to separate the UXO targets from the non-UXO targets. This was determined by generatively training and testing an RVM algorithm on either alternate receivers (5 cm apart) or alternate source pings (5 cm apart). The robustness of the trained RVM algorithms was determined by training and testing generatively on data from different flight paths (2.5 to 5 meter apart). Only the acoustic color feature sets passed this test. For the set of targets most dissimilar in shape and/or material to the UXO (e.g. panel or wooden pole section), the acoustic color feature-based RVM demonstrated a probability of detection of 1 and a false positive probability of ~ 0.3 . The poorer false positive performance against UXO-like targets (eg. metal cylinder) is a consequence of the spatially varying source pattern produced by BOSS.

The insertion of already available improved source technology is expected to provide significantly improved false positive performance for these targets as well.

Detailed Findings

(1) Image formation of buried targets using SAS-processed scattering data and AUV positional data collected using BOSS executing linear tracks over the buried targets can be used to garner approximate dimensions, burial orientation, and burial depth of typical UXO targets. With the AUV track length on the order of 10m, rough spatial resolutions of 0.25m in both the along-track and cross-track directions were achieved. From a theoretical point of view, we would expect the following limiting resolutions. Given a bandwidth of 3kHz to 20 kHz, the limiting (infinite aperture) image resolution ($C/2\Delta B$) would be ~ 0.04 m. For a 1m aperture, at the top of our band the cross-track resolution ($\lambda R/(2L_{CT})$) would be ~ 0.08 m and 0.2m at 2m and 5m ranges, respectively; and at the top of our band, the limiting along-track resolution ($\lambda R/(2L_{SAS})$) would be ~ 0.04 m and 0.09m at 2m and 5m ranges, respectively for a 2m track. In BOSS target measurements, the multi-aspect imaging procedures used in BOSS post-processing cause image distortion for some target aspects that limits the accurate registration of target dimensions and shape. Further, the conventional SAS method of sweeping the synthetic aperture along a linear track frequently does not allow specular illumination of the UXO targets at aspects that generate echo levels with adequate SNR for imaging. Imaging with data from tracks along the length of the cylindrical targets appears somewhat better than from tracks along the circumference. It

remains to be seen whether combining orthogonal path data in the imaging process improves the images.

(2) We have discovered a new classification feature connected to the 2-D depth view images related to the generation of circumferential creeping waves (or circumferential elastic waves) which are generated at the point on a cylindrical target of specular reflection and travel around the circumference of the cylinder with a delay time equal to the travel time around the circumference. For the limited cross-range resolution, this creates an image in the y,depth plane for flight paths along the x direction which appears as a narrow strip oriented in the direction of the incident wave and having a length (given for the case of a creeping wave) by approximately twice the cylinder diameter. This we believe is an artifact which exists for all cylindrical targets and is exploitable for classification.

(3) The high-dimensioned feature extracted from acoustic color (target strength versus aspect or position and frequency) has been shown to have sufficient information to be able to separate the buried UXO from the buried and proud non-UXO using an RVM classification algorithm trained and tested on either alternate receivers or on alternate pings. Further, this same RVM-processed multi-dimensional feature separated the epoxy-filled 155mm shells from the empty ones. For these targets, visual observation of the acoustic color maps for the presence of the elastic wave feature at quartering aspects also statistically separated the empty and filled shells. In addition, six more conventional feature sets also demonstrated the ability to separate the UXO and the non-UXO. It is important to note that these alternate receiver or ping results demonstrate only that the features separate; they do not in themselves demonstrate that practical training algorithms can be implemented to exploit this feature separation. (For this see (6) and (7) below.)

(4) Detailed BOSS projector/AUV radiation measurements over frequency and angle made in the NRL structural acoustic facility clearly demonstrate that the assumption made in this and all previous BOSS PMA exercises, i.e. that the projected sound is spherically symmetric, is grossly incorrect. The resulting inadequate knowledge of the frequency/angle source characteristics (i.e. the incident pressure waveform amplitude and phase) over the interrogated sediment surface clearly has a significant impact on our ability to separate UXO and non-UXO targets based on their measured acoustic color spatial and spectral structure. This is true to some extent for the feature separation results; but the far larger impact is in regard to the studies we carried out to assess experimental training of the RVM algorithm. These latter results which were obtained by training and testing on *different* flight paths must be considered “worst case” given the imprecise knowledge of the spatial source patterns across flight path separations and the source drop outs.

(5) The methods we developed to correct the acoustic color maps obtained with the existing BOSS system for these source-related artifacts included both source-AUV model development and ring calibration projection techniques. The former, while showing some promise, did not produce sufficient accuracy to warrant application to the BOSS Gulf data. Regarding the latter, several approximate ring calibration projection techniques were used successfully to improve the observed target separation results confirming that these source artifacts are indeed limiting our performance. What is required now is a more accurate method to incorporate these calibration source characteristics into the acoustic color PMA procedures on the existing Gulf data set or the use of

an improved source in future measurements. We would expect that success here would lead to significant improvements to the target classification results.

(6) Progress has been made regarding practical approaches for training the RVM algorithm. One is based on training with numerically generated UXO scattering data using existing, validated finite element structural acoustic models available to us. A second involves training with direct measurements made by flying the actual sonar system over linear flight paths of surrogate UXO targets of the type expected to be encountered buried in a site having similar sediment and acoustic propagation characteristics. We first carried out a simple study that would allow us to estimate the accuracy (probably more correctly, repeatability) required for the acoustic color training data. We concluded that in general the accuracy and efficiency expected from simulating acoustic color data for representative UXO targets using current state-of-the-art finite element-based structural acoustics models is not yet sufficient for planning to train the classification algorithms using simulated, numerical data. This limitation stems from several issues. The first is the basic fidelity of the finite element model especially at the higher frequencies given that we know the required UXO material, geometric, and boundary condition parameters in the first place. The second and more important consideration is the likelihood that we will not know these UXO details sufficiently well. Further, there will be additional errors related to the sonar system itself together with modeling its deployment and performance conditions. The former include the actual source spatial characteristics as discussed here previously while the latter include the degree to which one can control the particular vehicle flight paths used or realize the flight paths intended. Finally, in general modeling of acoustic propagation in the water column and sediment could introduce additional errors as well.

Based on the above, we believe that at this time the better way to train a Boss-like system is to carry out preliminary measurements with the actual sonar system on buried surrogate targets of the type expected to be encountered in the actual clean-up operation. These measurements would use linear tracks as would the test measurements, and they would be carried out in a similar environment or if possible in a smaller region of the actual test site.

We carried out “Leave One Out” studies to evaluate whether this scenario might form the basis for a practical training approach in the field. In these studies, we train on one path with all pings and test on two other paths using all pings. We believe the results are very encouraging in that they indicate that we can expect good performance using an RVM identification algorithm experimentally trained beforehand using known proud and buried UXO targets in an underwater site having similar sediment properties as the actual test site. While the results validate this conclusion regarding non-UXO targets that are structural acoustically dissimilar (quite different material and/or shape) to the UXO, this was not the case for non-UXO similar to the UXO. However, we believe the good performance would also be the case for the broader set of non-UXO were one to use a better designed sonar source. This follows in part from the fact that we are using only an approximate source calibration technique, one we believe tends to “wash out” the acoustic color differences between UXO and non-UXO. Further, the serious source drop outs with angle and frequency, which cannot be corrected even with a good calibration technique, lead to S/N issues. These two problems will be eliminated once a spatially and spectrally uniform source is used; and under these conditions we expect significantly improved performance for all non-UXO targets. It is also important to realize that the high P_{FA} 's for the problematic targets above (scuba

tank, aluminum cylinders, and solid concrete block) are based on a P_C of 100%. In this regard, it is also interesting to point out that this perfect P_C performance includes correct UXO calls on UXOs buried at angles not included in the training set.

(7) Some of the findings in these “Leave One Out” studies (together with the previous single path studies) are as follows. For all eight feature sets, but especially for the first six non-acoustic color sets, the good performance seen in the single path, alternate ping procedure is not maintained in the multiple path “Leave One Out” scenario. For the most part, in the latter there is no systematic improvement seen by applying source calibration using an approximate description of the source spatial pattern as there was in the single path studies. One consistently supported result is that of the ten non-UXO targets in the data, those most often called correctly are: n10, n11, CP, t5, and t15. These are the buried rock, the buried cinder block, the flat panel, the 5:1 wooden pole section, and the 2:1 wooden pole section. Further, in the “Leave One Out” studies unlike in the single path studies, the best performance is associated with one or the other of the two 200 frequencies x 22 positions acoustic color feature sets. Another observation is that in the “Leave One Out” studies, some of the more difficult non-UXO targets, i.e. those that are called as false positives in the 200 frequencies x 22 positions feature sets, fall near the discrimination threshold for many of the other feature sets. These include t14, t18, and to a lesser extent t7. These are the solid concrete block, the scuba tank, and the 3 foot aluminum cylinder. This suggests that it might be profitable to explore the benefit of including in the identification of these type targets one or more of the first six feature sets (i.e. 12 TS averages; 7 TS averages; Max TS, bin #, and standard deviation; Max TS and bin #; Max TS; and Mean TS) with the 200 frequencies x 22 positions, 20m and 2m aperture feature sets. However, we have not yet explored this idea.

LITERATURE CITED

1. B. H. Houston, J.A. Bucaro, T. Yoder, L. Kraus, and J. Tressler, J. Fernandez, T. Montgomery, T. Howarth, “Broadband Low Frequency Sonar for Non-Imaging Based Identification,” IEEE Oceans 2002.
2. J. A. Bucaro, B.H. Houston, M. Saniga, H. Nelson, T. Yoder and L. Kraus, and L. Carin, “Wide Area Detection and Identification of Underwater UXO Using Structural Acoustic Sensors,” 1st Annual SERDP Report, NRL/MR/7130—07-9014, January 12, 2007.
3. J.A. Bucaro, B.H. Houston, M. Saniga, A. Sarkissian, H. Nelson, T. Yoder, L. Kraus, and L. Carin, “Wide Area Detection and Identification of Underwater UXO Using Structural Acoustic Sensors – 2nd Annual SERDP Report,” NRL/MR/7130—08-9103, August 12, 2008.
4. J.A. Bucaro, B.H. Houston, M. Saniga, L.R. Dragonette, T. Yoder, S. Dey, L. Kraus, and L. Carin, “Broadband Acoustic Scattering Measurements of Underwater Unexploded Ordnance (UXO),” J. Acous. Soc. Am. 123, 738-746 (2008).

5. J. A. Bucaro, H. Simpson, L. Kraus, L. R. Dragonette, T. Yoder and B. H. Houston, "Bistatic scattering from submerged unexploded ordnance lying on a sediment," *The Journal of the Acoustical Society of America* 126 (5), 2315-2323 (2009).
6. J. A. Bucaro, B.H. Houston, H. Simpson, L. Kraus, T. Yoder, M. Saniga, A. Sarkissian, and L. Carin, "Wide Area Detection and Identification of Underwater UXO Using Structural Acoustic Sensors," 3rd Annual SERDP Report, January 2009.
7. J. A. Bucaro, B.H. Houston, H. Simpson, M. Saniga, A. Sarkissian, D. Calvo, L. Kraus, T. Yoder and L. Carin, "Wide Area Detection and Identification of Underwater UXO Using Structural Acoustic Sensors," 4th Annual SERDP Report, July 21, 2010.
8. J. A. Bucaro, B.H. Houston, H. Simpson, Z. Waters, M. Saniga, S. Dey, A. Sarkissian, D. Calvo, L. Kraus, and T. Yoder, "Wide Area Detection and Identification of Underwater UXO Using Structural Acoustic Sensors," Final Report to SERDP MR-1513, July 8, 2011.
9. S. G. Schock and J. Wulf, "Buried Object Scanning Sonar for AUV's," *Oceans* 2003, 494-499.
10. J.A. Bucaro, A. Sarkissian, B.H. Houston, H. Simpson, Z.J. Waters, D. Amon, K. Jig, S. Liskey, and T. Yoder, "Structural Acoustic UXO Detection and Identification in Marine Environments – Final Report to SERDP MR-2103", April 10, 2014.
11. M. Tipping, "Sparse Bayesian Learning and the Relevance Vector Machine," *Journal of Machine Learning Research*, vol. 1, pp. 211-244, 2001.
12. Z. J. Waters, H.J. Simpson, A. Sarkissian, S. Dey, B.H. Houston, J.A. Bucaro, and T. Yoder, "Bistatic, Above-critical Angle Scattering Measurements of Fully Buried Unexploded Ordnance (UXO) and Clutter," *J. Acous. Soc. Am.*, vol. 132, pp. 3076–3085 (2012).
13. S. Dey and D.K. Datta, "A Parallel hp-FEM Infrastructure for Three Dimensional Structural Acoustics," *Int. J. Numer. Methods Eng.* 68, 583-603 (2006).
14. Saikat Dey, Angie Sarkissian, Eris S. Mestreau, Brian H. Houston, and Larry Kraus, "Three-Dimensional Structural-acoustics Modeling and its Validation for Free-field and Littoral Environments," *J. Acoust. Soc. Am.* 125, 2702 (2009).
15. Joseph A. Bucaro, Zachary J. Waters, Brian H. Houston, Harry J. Simpson, Angie Sarkissian, Saikat Dey, and Timothy J. Yoder, "Acoustic identification of buried underwater unexploded ordnance using a numerically trained classifier," *J. Acoust Soc Am* 132, 3614-1617 (2012).
16. Paul J. Carroll, "Underwater (UW) Unexploded Ordnance (UXO) Multi-Sensor Data Base (MSDB) Collection," Final report SERDP Project MM-1507, July 2009.

17. Robert A. Leasko, Charles L. Bernstein, Richard Holtzapple, and Jesse I. Angle, “Munitions Detection using Unmanned Underwater Vehicles Equipped with Advanced Sensors,” Interim Report, ESTCP Project MR-201103, June 29, 2012.
18. I. Ulusoy and C. M. Bishop, “Generative versus discriminative methods for object recognition,” in Computer Vision and Pattern Recognition, 297 IEEE Computer Society conference (2005), Vol. 2, pp. 258–265.
19. E. G. Williams, Fourier acoustics: Sound radiation and nearfield acoustic holography. (Academic Press, San Diego, 1999).
20. J.A. Bucaro, A. Sarkissian, B.H. Houston, M. Saniga, H. Simpson, Z.J. Waters, D. Amon, E. Williams, N. Valdivia, T. Yoder, K. Jig, and S. Liskey, “Structural Acoustic UXO Detection and Identification in Marine Environments – Interim Report for SERDP MR-2103 Follow-On”, July 30, 2015.

APPENDIX

The winged BOSS-40 system (photograph shown in Fig. A1) is designed to scan for buried underwater objects using two 1m hydrophone arrays mounted as wings of an autonomous underwater vehicle (AUV). An under-view and head-on view of the system is shown in the




Figure A1. Photograph of BOSS in the Water.




Figure A2. Plan and head-on view of BOSS.

diagram in Fig. A2. The AUV system used here is the Reliant Bluefin 12 vehicle designed and fabricated by Bluefin Robotics. (See Figure A3.) Each wing contains 20 hydrophone channels, yielding a 40-channel strip array used to collect the sonar returns. This system uses an acoustic projector that transmits a FM pulse over the frequency band 3-20 kHz. The winged BOSS is smaller and more mobile than the previous BOSS vehicles, as the wing arrays tend to produce less drag than the large circular array.

In order to improve the resolution of target imagery, the wing BOSS utilizes time-delay focusing extended to hydrophone data collected over several transmissions. With synthetic aperture processing, the along-track resolution of target imagery improves with

distance traveled while forming the synthetic aperture. The use of synthetic aperture processing also allows the along-vehicle dimension of the array to be significantly reduced compared to a real array thereby reducing the hydrophone array drag and surface area and increasing the ease with which BOSS can be deployed on an AUV.




Figure A3. Photograph of BOSS prior to complete assembly showing some of the internal components.

Three-dimensional SAS imagery is generated using a navigation solution based on measurements from a Doppler Velocity Log (DVL) and an Inertial Measurement Unit (IMU) to time-delay and coherently sum matched-filtered phase histories from subsurface focal points over a large number of pings. The focused data consist of a large set of three dimensional SAS data cubes created by a sliding window of ping intervals, where adjacent data cubes have greater than 90% overlap. By using navigation/registration information, these data cubes are fused into a single large three dimensional dataset, in which each voxel's intensity is the maximum intensity of the co-registered voxels across all original data cubes. For improved image contrast, the intensity of the specular seafloor return (a shallow swath of voxels beneath the platform) is spatially nullified with an automated process employing measured backscatter statistics. Fig. A4 shows mosaics of top-view maximum intensity projections for data runs collected over the TREX13 target field in the Gulf of Mexico. These results were obtained by R. Holtzapple and N. Pineda as reported in their "Quick Look Analysis" document following the June 2013 BOSS exercises.

As can be seen in the top portion of Fig. A4, all proud and buried targets in the field of view of these 3 meter altitude west-east flights are detected via the plan view images displayed in this figure. The three 2-D images generated from these are shown for four of the buried NRL targets in the lower part of the figure. In each of these, the plan view image is repeated in the upper left

while the images with depth (x,z and y,z) are shown in the lower left and upper right images, respectively. As can be seen, the N5 and N6 (x,z) images correctly show horizontally buried targets while those for N7 and N8 correctly show targets buried nose up at about 30° and 60° angles.




Figure A4. Preliminary 3-D images obtained from June 23, 2014 BOSS data by Richard Holtzapple and Nick Pineda.