FINAL REPORT

Elastic Target Modeling for Physics-Based Automatic Classification

SERDP Project MR-2649

APRIL 2020

Lane Owsley
Aubrey Espana
University of Washington

Distribution Statement A
This document has been cleared for public release

SERDP
DOD ▪ EPA ▪ DOE

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

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.
PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE *(DD-MM-YYYY)*	**2. REPORT TYPE**	**3. DATES COVERED** *(From - To)*
30/04/2020	SERDP Final Report	9/7/2016 - 3/7/2020

4. TITLE AND SUBTITLE

Elastic Target Modeling for Physics-Based Automatic Classification

5a. CONTRACT NUMBER
16-C-0009

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Aubrey Espana
Lane Owsley

5d. PROJECT NUMBER
MR-2649

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
University of Washington
Applied Physics Laboratory
1013 NE 40th St.
Seattle, WA 98105-6698

8. PERFORMING ORGANIZATION REPORT NUMBER
MR-2649

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Strategic Environmental Research and Development Program (SERDP)
4800 Mark Center Drive, Suite 16F16
Alexandria, VA 22350-3605

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

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
MR-2649

12. DISTRIBUTION/AVAILABILITY STATEMENT
DISTRIBUTION STATEMENT A. Approved for public release: distribution unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT This document is a final report for MR-2649, focused on developing methods for relating the elastic physics of the targets of interest to the scattered returns and leveraging the spectral and temporal structure of those returns to improve automatic classification. The primary results presented in the report demonstrate two different approaches to relating observables in scattered returns to the underlying internal physics. First, we model the returns from a target manufactured for this project to break down the observed structures in the return and categorize those structures by the type of associated elastic behavior. We compare the modeled responses to the data collected in the ClutterEx17 experiment, both to validate the modeling and to understand the types of variation to which a feature set must be robust. Second, we analyze two existing classification systems used in studies involving simulated data and data collected in various underwater environments for the purpose of identifying the features currently being used, testing system robustness, and to identify opportunities for improvement. Finally, we consider the implications of the movement to a downlooking sonar in the development of feature sets and classification architectures. The overall assessment of the report is that the modeling approach can be successfully used to identify observables which can be understood in terms of the underlying physics and used as a basis for targeted feature sets. We describe how the results and conclusions obtained in this effort will direct upcoming follow-on work.

15. SUBJECT TERMS
Elastic Target Modeling, Physics-Based Automatic 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**			Lane Owsley
UNCLASS	UNCLASS	UNCLASS	UNCLASS	103	**19b. TELEPHONE NUMBER** *(Include area code)* 206-685-3592

Acknowledgements

Dr. Warren Fox (APL-UW) led the classification analysis effort described in Section 6.1 of this report.

The original field-data classification work was performed by Dr. Kevin Williams and Dr. Daniel Plotnick (both at APL-UW) under funding from the Office of Naval Research, Code 32, Ocean Engineering and Marine Systems Program.

The original TIERSWAT classification work was performed by Dr. Jack McLaughlin (APL-UW) under project MR-2505 (PI: Kargl).

Farfield modeling for targets in sediment was done by Dr. Steve Kargl (APL-UW).

Data collection for this work was done under ClutterEx17, jointly funded by SERDP under project MR-2505 (PI: Kargl) and the Office of Naval Research, Ocean Engineering and Marine Systems Program (Stack).

Materials analysis for the science target was performed by Dr. Jermaine Kennedy (NSWC-PCD).

ABSTRACT

This document is a final report for MR-2649, focused on developing methods for relating the elastic physics of the targets of interest to the scattered returns and leveraging the spectral and temporal structure of those returns to improve automatic classification. The primary results presented in the report demonstrate two different approaches to relating observables in scattered returns to the underlying internal physics. First, we model the returns from a target manufactured for this project to break down the observed structures in the return and categorize those structures by the type of associated elastic behavior. We compare the modeled responses to the data collected in the ClutterEx17 experiment, both to validate the modeling and to understand the types of variation to which a feature set must be robust. Second, we analyze two existing classification systems used in studies involving simulated data and data collected in various underwater environments for the purpose of identifying the features currently being used, testing system robustness, and to identify opportunities for improvement. Finally, we consider the implications of the movement to a down-looking sonar in the development of feature sets and classification architectures. The overall assessment of the report is that the modeling approach can be successfully used to identify observables which can be understood in terms of the underlying physics and used as a basis for targeted feature sets. We describe how the results and conclusions obtained in this effort will direct upcoming follow-on work.

Contents

List of Acronyms **1**

1 Introduction **8**

2 A Discussion On Terms **10**

3 Science Target: Goals, Fabrication and Data Collection **11**

4 Physics Decomposition **14**
4.1 Separation of Elastic and Inelastic Components of Response 14
4.2 Physics Investigation Framework 16
4.3 Survey of Elastic Response . 26
4.4 Internal Physics Constraints 36

5 Evaluation of Clutterex Data **44**

6 Classification System Analysis **57**
6.1 Field-Collected Data Classification System Analysis 57
Overview of the Classification Experiment 57
Classifier Experiment Analysis 58
Relating to Internal Physics 64
Commentary on Application to a Down-Looking Sonar 64
6.2 TIERSWAT Classification System Analysis 68
GMM Component Usage: Orientation-Specific Modeling and Clutter Characterization 70
GMM Component Usage: Visualizing Key Spectral Features 75

7 Down-Looking Sonar: Implications for Feature Sets and Classification Architecture **82**

8 Conclusions **87**
8.1 Follow-On Work . 88

Appendices **96**

List of Acronyms

AC	acoustic color
AUC	area under the curve
APL-UW	Applied Physics Laboratory, University of Washington
AR	autoregressive
ARMA	autoregressive moving-average
BayEx14	Bay Experiment 2014 (St. Andrew Bay)
ClutterEx17	Clutter Experiment 2017 (Gulf of Mexico)
COMSOL	commercial FE simulator/solver
CSAS	circular synthetic aperture sonar
CNN	convolutional neural network
ESTCP	Environmental Security Technology Certification Program
FE	finite element (as in finite element modeling)
GMM	Gaussian mixture model
GUI	Graphical user interface
IPR	in-progress review
NATO	North Atlantic Treaty Organization
NRL	Naval Research Laboratory
NSWC-PCD	Naval Surface Warfare Center - Panama City Division
NURC	NATO Undersea Research Centre, La Spezia, Italy
ONR	Office of Naval Research
PDF	probability density function
PML	perfectly matched layer
ROC	receiver operating characteristic
RVM	relevance vector machine
SAS	synthetic aperture sonar
SERDP	Strategic Environmental Research and Development Program
SEED	SERDP Exploratory Development
SVM	support vector machine
SWAT	Personal Computer Shallow Water Acoustic Tool-Set for reverberation modeling
TIER	target in the environment response
TIERSWAT	simulated dataset created using FE models propagated using TIER, with SWAT reverberation
TNO	Netherlands Organization for Applied Scientific Research
TREx13	Target and Reverberation Experiment 2013 (Gulf of Mexico)
UXO	unexploded ordnance

List of Figures

1 Stages of science target fabrication 12
2 Science target being placed on seafloor for Clutterex data collection 13
3 Separation of elastic and inelastic response using the modeling technique developed under this project. Once the full-physics model is created, the elastic-only scattering is created by reproducing the internal movement of the target without any incident wave. The coherent difference between the two is the inelastic or shape response. 17
4 An example acoustic color image from the physics-investigation GUI used to investigate the return from the air-filled science target. The user has selected three points for which to visualize the internal physics; the results are shown in Figures 6 through 15. . 18
5 Visualizing internal physics of an axisymmetric target by unwrapping the physics on a curved surface. In this example, both figures display the same information. The surface pressure is represented in true spatial coordinates in subfigure (a); in subfigure (b) we use the x-axis to represent degrees from the front (where 'front' is defined based on a plane wave traveling in a direction in the x-z plane as illustrated by the purple arrow in subfigure (a). 19
6 The unwrapped complex displacements (in xyz coordinates) for Point 1 from Figure 4. The black lines indicate wavefront periods and the white lines are half-periods. 20
7 The unwrapped complex displacements (in surface normal-tangential coordinates) for Point 1 from Figure 4 . 20
8 The unwrapped complex axial strains (in xyz coordinates) for Point 1 from Figure 4 21
9 The unwrapped complex axial strains (in surface normal-tangential coordinates) for Point 1 from Figure 4 . 21
10 The unwrapped complex shear strains (in xyz coordinates) for Point 1 from Figure 4 22
11 The unwrapped complex shear strains (in surface normal-tangential coordinates) for Point 1 from Figure 4 . 22
12 The unwrapped complex displacements (in xyz coordinates) for Point 2 from Figure 4 23
13 The unwrapped complex displacements (in surface normal-tangential coordinates) for Point 2 from Figure 4 . 23
14 The unwrapped complex shear strains (in xyz coordinates) for Point 3 from Figure 4 24
15 The unwrapped complex shear strains (in surface normal-tangential coordinates) for Point 3 from Figure 4 . 24
16 The complete set of points whose internal physics will be categorized in the survey of Section 4.3 . 25
17 The subset of observables due to bending, which takes the form of movement in the x-direction dominating movement in the y-direction. Internal physics of these points are shown in Figure 18 . 27
18 The normal x-displacement (real part only) for the points in Figure 17, whose internal physics are dominated by movement in the x-direction over the y-direction 27
19 The subset of observables due to a regular grid of standing waves dominated by movement normal to the target surface. Internal physics of these points are shown in Figure 20 . 28
20 The normal displacement (real part only) for the points in Figure 19, whose internal physics are dominated by normal standing waves in a regular grid 29

21 The subset of observables due to a regular grid of standing waves dominated by movement normal to the target surface but not (exclusively) in a regular grid. Internal physics of these points are shown in Figure 22 30
22 The normal displacement (real part only) for the points shown in shown in Figure 21, which are dominated by movement normal to the target surface but not (exclusively) in a regular grid. . 31
23 The subset of observables due to axial waves. Internal physics of these points are shown in Figure 24 . 32
24 The normal displacement and axial-tangent displacement for a selection the points shown in shown in Figure 23 33
25 The subset of observables due to helical waves. Internal physics of these points are shown in Figure 26 . 34
26 The *xyz* displacement (real, imaginary, and phase) for the points shown in shown in Figure 25. The phase information (and the GUI animation tool) indicate that most of these waves are moving. . 35
27 The subset of observables due to behaviors that are a combination of the other classes or do not fit into any of the classes. Internal physics of these points are shown in Figure 28 . 36
28 The *xyz* displacement (real, imaginary, and phase) for a selection of the points shown in shown in Figure 25. The internal physics for these points are either combinations of the other classes or do not fit into any of the classes. 37
29 A measure on the aspect/frequency plane where points are represented in terms of their ratio of maximum x-displacement to maximum y-displacement. Pure bending modes will be high in this measure. 39
30 The upper figure shows the complete elastic model farfield response for the science target. The lower figure is the response for a model whose physics have been constrained by a y-direction spring to allow only the types of movement identified in Figure 29 . 40
31 A measure on the aspect/frequency plane where points are represented in terms of their ratio of maximum normal-displacement to maximum tangential-displacement 41
32 A measure on the aspect/frequency plane where points are represented in terms of their ratio of maximum tangential-plane shear stress overall to maximum tangential-plane shear stress along a particular line (see text and Figure 15) 42
33 The upper figure shows the complete elastic model farfield response for the science target. The lower figure is the response for a model whose physics have been modified by adjusting the shear stiffness along an axis where certain classes of observables have no tangential shear strain. . 42
34 Measures on the aspect/frequency plane where points are represented in terms of their maximum displacement in different directions 43
35 Measurement of the wall thickness and diameter of the cut target (both halves were measured to ensure consistency). The left figure shows wall thickness at different points along the target axis, and indicates a 1.5 mm drift in the alignment of the inner mold with the outer mold. The right figure shows target diameter, which indicates a 2 mm drift along the axis. Based on these results, the model parameters were updated to the average wall thickness of 1.38 cm. 45
36 Deconstruction of target for materials testing and geometry validation 46
37 Comparison of model and collected data for two filler materials 48

38 Binary acoustic color plots for the model (top row) and data (bottom row) for the air-filled shell (left column) and water-filled shell (right column). 49
39 Binary acoustic color plots for the model (top row) and data (bottom row) for the air-filled shell (left column) and water-filled shell (right column). 50
40 A comparison of the model (top row) and data (middle row) over different ranges of angles, or 'VIEWS', for the air-filled shell. 50
41 A comparison of the model (top row) and data (middle row) over different ranges of angles, or 'VIEWS', for the water-filled shell. 51
42 Acoustic color for the air-filled shell for 3-10 kHz band, showing the model for a flat bottom (top-left), data for a flat bottom (top-right), data for a large rippled bottom (bottom-left) and data for a raked bottom (bottom-right). 53
43 Binary AC plots for the air-filled shell for 3-10 kHz band, showing the model for a flat bottom (top-left), data for a flat bottom (top-right), data for a large rippled bottom (bottom-left) and data for a raked bottom (bottom-right). The threshold value is set to -10 dB. . 53
44 Acoustic color for the water-filled shell for 3-10 kHz band, showing the model for a flat bottom (top-left), data for a flat bottom (top-right), data for a large rippled bottom (bottom-left) and data for a raked bottom (bottom-right). 54
45 Binary AC plots for the water-filled shell for 3-10 kHz band, showing the model for a flat bottom (top-left), data for a flat bottom (top-right), data for a large rippled bottom (bottom-left) and data for a raked bottom (bottom-right). The threshold value is set to -10 dB. . 54
46 Acoustic color for the air-filled shell for 10-30 kHz band, showing the model for a flat bottom (top-left), data for a flat bottom (top-right), data for a large rippled bottom (bottom-left) and data for a raked bottom (bottom-right). 55
47 Binary AC plots for the air-filled shell for 10-30 kHz band, showing the model for a flat bottom (top-left), data for a flat bottom (top-right), data for a large rippled bottom (bottom-left) and data for a raked bottom (bottom-right). The threshold value is set to -15 dB. . 55
48 Acoustic color for the water-filled shell for 10-30 kHz band, showing the model for a flat bottom (top-left), data for a flat bottom (top-right), data for a large rippled bottom (bottom-left) and data for a raked bottom (bottom-right). 56
49 Binary AC plots for the water-filled shell for 10-30 kHz band, showing the model for a flat bottom (top-left), data for a flat bottom (top-right), data for a large rippled bottom (bottom-left) and data for a raked bottom (bottom-right). The threshold value is set to -15 dB. . 56
50 Example of separation of full time series data vs. target aspect into early-time, relating primarily to shape scattering, and late-time, relating primarily to elastic target response, in the lower panels. Corresponding frequency-aspect data are displayed in the upper panels. (Courtesy Dr. Daniel Plotnick and Dr. Kevin Williams) 59
51 Example ROC curves for one of the k-folds partitions. 60
52 Histogram of swath centers where full classifier fails but elastic classifier succeeds. 61
53 Full, shape, and elastic acoustic color plots: swath center at 125 deg., range = 15 m. 62
54 Full, shape, and elastic acoustic color plots: swath center at 125 deg., range = 30 m. 62
55 Full, shape, and elastic acoustic color plots: swath center at 125 deg., range = 40 m. 62
56 Full, shape, and elastic acoustic color plots: swath center at 285 deg., range = 15 m. 63
57 Full, shape, and elastic acoustic color plots: swath center at 285 deg., range = 30 m. 63
58 Full, shape, and elastic acoustic color plots: swath center at 285 deg., range = 40 m. 63

59 Expanded view of target response from Figure 53: Full and elastic from measured data, and TIER model result. The white circle in the TIER results indicates the point in frequency and aspect at which we will investigate the internal physics. 65

60 Internal physics (total displacement in cylindrical coordinates) for aluminum UXO replica at 15 m range. First column shows displacement in r dimension, second column in ϕ dimension, and third column in z. Descending rows show real, imaginary, magnitude, and phase of the components, respectively. 66

61 Internal physics (total displacement in Cartesian coordinates) for aluminum UXO replica at 15 m range. First column shows displacement in x dimension, second column in y dimension, and third column in z. Descending rows show real, imaginary, magnitude, and phase of the components, respectively. 67

62 Stills from the GUI animation of one layer of the target, used to aid in visualising the physics of Figure 61 . 68

63 Illustration of a simple GMM-based classification system for a one-dimensional feature space. The top figure shows two GMMs (target and clutter), each represented as the weighted sum of two Gaussian components. An unknown-class pattern with a particular feature value can be classified by comparing the relative densities of the two models at that particular feature value; the ratio of the target-model density to the clutter-model density is shown in the bottom subfigure. Patterns with high target likelihood will be classified as targets, with the particular alarm threshold dependent on the desired detection/false-alarm requirements of the system. 69

64 Usage of Gaussian Mixture Model components for three different target types. The model has 15 total Gaussians to represent the distribution of ping return spectra; each 'x' in the figure indicates the dominant component index for a particular ping return, with the y axis indicating the incident angle of the ping on the target. Different colors represent different surveys, with hotter colors indicating objects at increasing distance from the sonar. It can be observed that the model is effectively subdivided according to incident angle and target type, with distinct components being used to model particular target types at particular angles. 71

65 Usage of Gaussian Mixture Model components for three different clutter types. The model has 15 total Gaussians to represent the distribution of ping return spectra; each 'x' in the figure indicates the dominant component index for a particular return, with the y axis indicating the incident angle of the ping on the target. Different colors represent different surveys, with hotter colors indicating objects at increasing distance (decreasing grazing angle) from the sonar. As in the target models in Figure 64, it can be observed that the model is effectively subdivided according to incident angle and clutter type, with distinct components being used to model particular clutter types at particular angles. New clutter types are thus likely to be poorly represented by this model. . 72

66 Receiver operating curves (capture rate vs. false alarm rate) for TIERSWAT acoustic-color classification under differing assumptions of clutter-characterization sufficiency. The black curve shows the results for the MR-2505 GMM system, which assumes all clutter object types in the test set are available for training, and achieves nearly perfect performance under that assumption. The red curve assumes no training data are available for the specific clutter types being tested against, and shows a significant performance degradation (see text for discussion of assumptions and resulting algorithm modification). . 74

67 Receiver operating curves (capture (detection) rate vs. false alarm rate) for TIER-SWAT acoustic-color classification under differing assumptions of clutter-characterization sufficiency. The black curve shows the results for the MR-2505 GMM system, which assumes all clutter object types in the test set are available for training, and achieves nearly perfect performance under that assumption. The red curve assumes no training data are available for the specific clutter types being tested against, and shows a significant performance degradation The green curve shows a result from a system modified so that each UXO type is modeled separately, and that models vary as a function of incident aspect angle. The simple modification allows the recovery of a significant amount of the performance loss (see text for details of the algorithms and motivation). . 76

68 Visualizing the GMM representation of 155mm Howitzer acoustic color (AC) information by replacing each ping return's spectrum with the best-fit GMM component center. The AC for six interrogations of Howitzers are shown in the top images, located so that the same target aspect angle in each AC are aligned horizontally (the dotted line represents broadside). The lower images show the corresponding best-fit GMM components for the ping returns in the AC above them. The key spectral structures used in classification can be visualized, and it can be observed that the same structures are identified in each interrogation at a given aspect angle and that relatively large extents of aspect angle are characterized by particular spectral structures. . 78

69 Visualizing the GMM representation of 105mm Howitzer acoustic color (AC) information by replacing each ping return's spectrum with the best-fit GMM component center (see Figure 68 caption for explanation). 79

70 Visualizing the GMM representation of 105mm Howitzer acoustic color (AC) information by replacing each ping return's spectrum with the best-fit GMM component center (see Figure 68 caption for explanation). The clutter AC representation differs from the target representation in both the component spectra which are used and the evolution of component usage as a function of angle; both differences can be exploited in separating clutter from UXO. 80

71 Visualizing the GMM representation of 105mm Howitzer acoustic color (AC) information by replacing each ping return's spectrum with the best-fit GMM component center (see Figure 68 caption for explanation). The clutter AC representation differs from the target representation in both the component spectra which are used and the evolution of component usage as a function of angle; both differences can be exploited in separating clutter from UXO. 81

72 Cartoon describing the geometric relationship between a down-looking sonar traveling in the positive-y direction (along the blue line) above a seafloor defined as the xy plane. The position of the target is defined by the x-position in meters off of the source path, and by the target orientation in degrees. 82

73 Observed target aspect for a down-looking source-receiver making the traverse illustrated in Figure 72 from y=-10 to y=+10 meters. Each sub-figure represents a different target position, and each line represents a different target orientation at that position. Extreme variation in aspect traverse as a function of these two variables can be observed. Simulations of the response acoustic color are shown in Figure 74. 84

74 Simulated acoustic color responses corresponding to a science target observed by a down-looking sonar given the locations/orientations from Figure 74. The x axes are frequency and the y axes are position of the source along the traverse. 86

75 Ray diagrams for a cylinder at broadside or spherical target: (a) proud target with 4 paths included; (b) burial depth eliminates path 4 and only paths 1-3 are present; (c) burial depth eliminates paths 2-4 and only path 1 remains; and (d) full buried target with only a refracted path. (Figure courtesy Dr. Steve Kargl) 88

76 A circlular SAS image of a 155 mm UXO produced using high-frequency data. Image processing techniques could be used to produce an orientation estimate for use in subsequent classification algorithms. (Figure courtesy Dr. Timothy Marston, from MR-2439 final IPR). . 89

77 A mock-up of a proposed GUI that would enable post-mission analysis and follow-up in a physics-based framework. The upper left figure shows a SAS image of a complete scanned scene, with boxes indicating detected objects and their classification results. An object of uncertain identity (the yellow box) has been flagged for investigation, and a data representation (acoustic color) of the scattered return is displayed in the upper right, below which is the return which would be predicted for the hypothesized class. The user can investigate the underlying physics corresponding to the observables in the current orientation (#1 and #2 on the lower right). Further, the system can analyze the complete response of the hypothesized class at all angles (lower left) to identify the physics behaviors which would produce the most definitive return (#3 in lower right) and propose a follow-up scan path designed to maximize classification confidence. . 93

1 Introduction

The motivating goal of this research was to develop modeling methods and physics insights to improve the state of the art in sonar-based automatic detection classification of elastic targets such as UXO. In particular, in this effort we have developed methods for relating the scattered field produced by an interrogated object to the structure and internal physics of that object so that we could better understand when particular return components will be observable and develop feature sets and classification algorithms tailored to these components. In this report, we describe motivations for, data collection efforts supporting, and directions we have taken toward this goal.

In **Section 3**, we describe the scientific target that was created under this effort, including the motivations for such a task and the fabrication process. The target was fielded in the separately funded ClutterEx17 data collection, which we briefly review as well.

In the remaining sections of this report, we describe two different directions we are taking toward the goal of identifying features which will be of value in automatic UXO classification and relating them to the underlying physics. The first approach can be thought of as a "bottom-up" paradigm which begins by modeling the internal physics and developing an understanding of how that affects the target signature. This method is presented in **Section 4**, where we describe a decomposition of the modeled return from our science target. We review the tools we have developed to do this decomposition, discuss categories of observables identified using these tools, and present a survey of the observables within each of these categories.

Section 5 presents the ClutterEx17 data collected using the science target fabricated under this effort and compares it to the models described in Section 4. The original models were not a good fit to the observed data, and we describe the target composition analysis that went into improving the model. The current results are a better fit in terms of our ability to use the models to relate components of the observed return, but the differences that exist highlight the importance of our goal of developing a feature set which is robust to the sources of variability underlying this discrepancy. For this purpose, we also look at variations in the return due to changes in the sediment surface.

Our second approach to identifying useful features and relating them to the underlying physics is described in **Section 6**. We view this as a "top-down" paradigm where we start with an existing complete classification system and analyze how and when it takes advantage of target physics and how it can be improved and made more robust. We study two existing classification systems, one operating on field-collected and one operating on simulated data. Our first study begins with a performance analysis of a baseline classification system which has undergone basic modifications to take advantage of the elastic/inelastic division of the scattered returns. By studying the specific conditions under which this rudimentary modification provides improvement and the additional information which the modified system is able to take advantage of, we are able to identify particular observables in the return which might be better captured with specifically-designed feature sets and an understanding of the underlying physics that produce them. Our analysis of the second system leads to insights on the possibility for performance improvement by doing target orientation estimation. The analysis also makes clear that system concepts based on the assumption of complete clutter characterization have the potential for significant performance degradation in realistic unknown-clutter scenarios, and thus a robust classification system will require an architecture that does not require this assumption.

In **Section 7**, we address the implications of moving to a down-looking sonar system on the choices

of feature sets and classification architecture. While this configuration can provide some significant advantages, it also exacerbates the variability that can be expected in the acoustic response as a function of target orientation. We present some simulations to illustrate this point, and discuss how the need for a physics-based understanding of this variability is even more critical in this scenario.

We conclude in **Section 8** by identifying the key results of the project and the implications for future work. We end with a description of the goals for our follow-on effort to begin in 2020.

2 A Discussion On Terms

In the context of automatic classification, the term *feature* often leads to confusion as it is used (sometimes interchanegably in the same discussion) to refer to multiple concepts. In this report, we adopt the following terminology to distinguish among these concepts:

1. ***Physical characteristic***: A quality of an object to be classified, which may serve to distinguish it from other objects; *e.g.* the length of a UXO or its material composition
2. ***Physics behavior***: A way in which the object interacts with the interrogation method to produce collectable data; *e.g.* an acoustic resonance produced by a particular wave type at a particular frequency
3. ***Observable***: A quality of the object's realization in a particular data representation which may be identified (especially by a human expert) as a potential marker of that object; *e.g.* a strong scattered return at a particular frequency and aspect angle
4. ***Feature***: An extracted value from a particular data representation which will be passed to an automatic classifier; *e.g.* the cross-correlation of a return's acoustic color[1] representation with a class template
5. ***Internal (Classifier) Feature***: A dimension of the object's representation (internal to the classifier itself) in which the class decision boundary is defined

It should be noted that even within the machine learning community, there is not a consensus on terminology, especially as the line between concepts 4 and 5 can be blurry or even absent depending on the specific classification architecture.

[1] We use the term *acoustic color* to refer to the strength of the scattered response as a function of azimuth and frequency.

3 Science Target: Goals, Fabrication and Data Collection

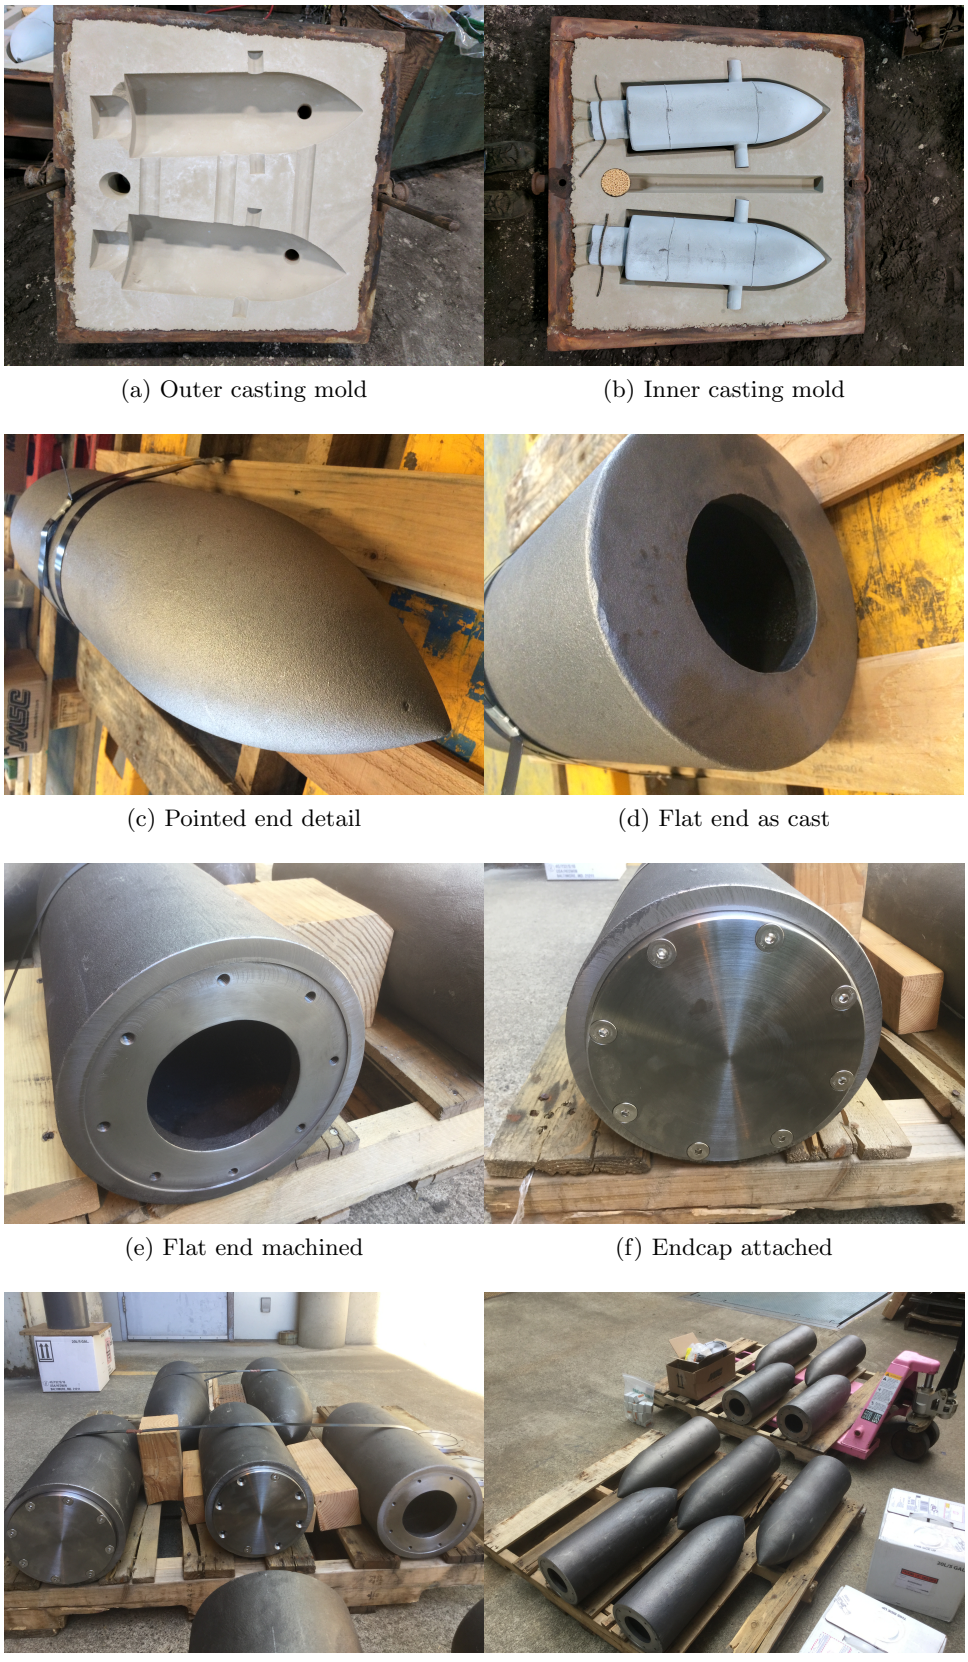
One of the goals of this project was to provide proof-of-concept that observables in the scattered return from an elastic object can be related to the internal physics in such a way that they can be tied to the target geometry, and that we can predict the robustness of features which indicate the presence of such internal physics. For an initial investigation, there is value in simplifying the expected internal physics as much as possible to aid in reproducibility and visual interpretation, while at the same time creating realistic enough responses that results produced in this study would have relevance to UXO or other targets of interest.

To this end, we manufactured a set of science targets for deployment in the ClutterEx17 data collection. A shell was designed that had a shape similar to a 155mm Howitzer shell, but with a smooth exterior, uniform wall thickness, a simple endcap and no internal components (the CAD diagrams for this target are attached to this report in an appendix). The decision was made to have the shells cast from a mold, which greatly reduced the per-unit price relative to a full-machining fabrication and also provided more flexibility in creating a realistic target size and simple geometry. The cast target had an opening at the flat end; the flat end was machined to accommodate screws for an endcap. The machined endcap combined with a gasket to produce a watertight seal; the targets could thus be filled with a variety of substances or left filled with air. Images of the manufacturing and shipping process are shown in Figure 1. The molds have been retained to accommodate any additional need for target manufacture.

Multiple copies of the shell were cast out of steel and deployed during the CLUTTEREX17 sea trial in the Gulf of Mexico off the coast of Panama City, FL in July 2017. The data collection was supported by SERDP project MR-2505, (PI: Steve Kargl). The shells were filled with various fillers and deployed on a sand sediment comprised of different levels of roughness (natural large scale ripples, small ripples, flat). The acoustic scattering was measured for the targets in an isolated state, as well as in the presence of clutter objects. Figure 2 shows some of the targets in the test field. As isolated components are identified within the modeling work, this dataset is used provide the means to validate their presence in a real ocean environment, and examine their robustness to variability in the environment and the presence of clutter. Section 5 of this report describes the validation and the modifications to our initial model that were performed based on the collected data.

The CLUTTEREX17 experiment was conducted at the same site as the TREX13 experiments. The area where the targets are deployed is comprised of a medium to course sand. Historical measurements using a variety of acoustic and direct measurement methods allow us to estimate the sound speed and density to be 1655 m/s and 1912 kg/m^3 [1],[2]. Diver video of the seafloor topography revealed a large scale ripple pattern, whose peaks and troughs run approximately parallel to the source/receiver travel path. The approximate size of the ripples was 1.5" in height, with a period of 24.5".

The targets manufactured for this project are being deployed for data collection using the Multi-Sensor Towbody down-looking sonar system being tested at APL-UW (ESTCP MR18-5004, PI: Williams). The targets and resulting data will part of future classification and corrosion studies (MR20-1443, PI: Owsley).



(a) Outer casting mold

(b) Inner casting mold

(c) Pointed end detail

(d) Flat end as cast

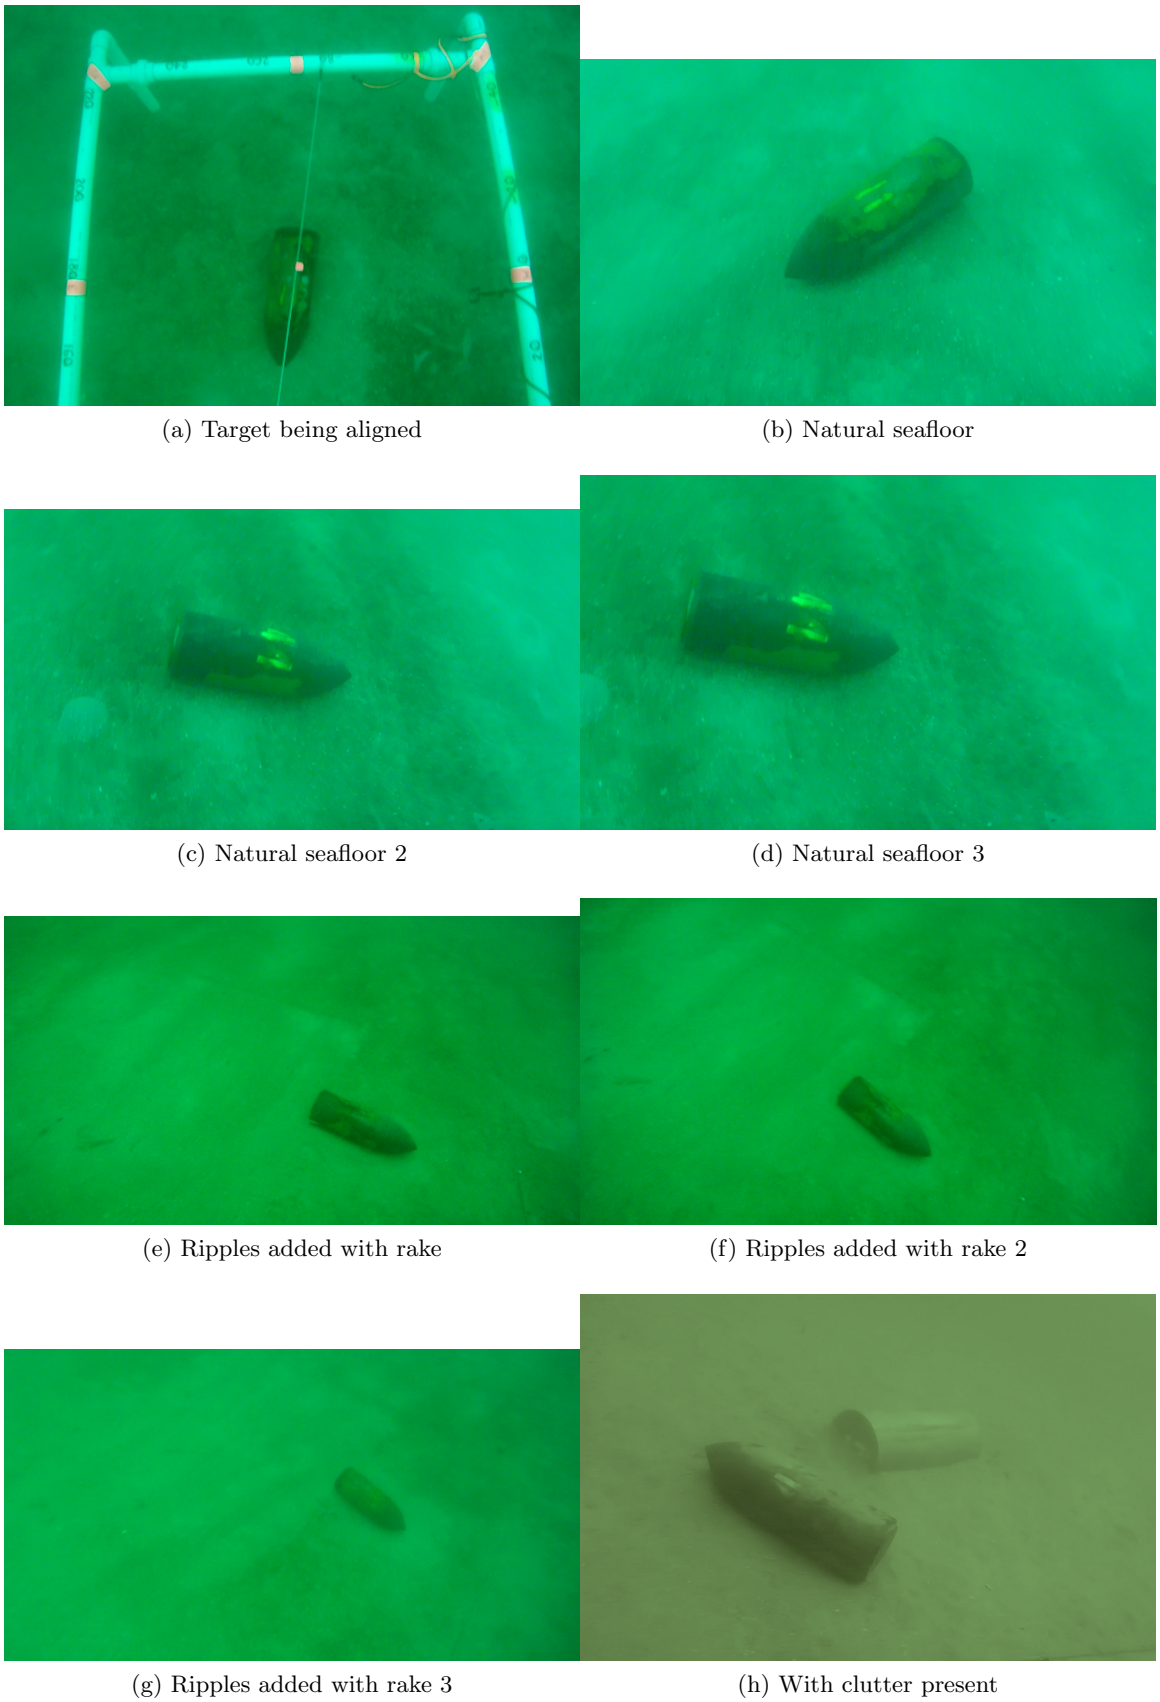
(e) Flat end machined

(f) Endcap attached

(g) On pallets

(h) Ready for shipping

Figure 1: Stages of science target fabrication



(a) Target being aligned

(b) Natural seafloor

(c) Natural seafloor 2

(d) Natural seafloor 3

(e) Ripples added with rake

(f) Ripples added with rake 2

(g) Ripples added with rake 3

(h) With clutter present

Figure 2: Science target being placed on seafloor for Clutterex data collection

4 Physics Decomposition

One of the two primary approaches we have taken in developing feature sets based on the underlying target physics originated from the scattered return predicted by the FE model. By decomposing the complex return into components and categorizing those components based on the modeling of the internal physics, we can separately investigate the effects of variations due to environment and target geometry on each class of physics behavior, and develop features robust to such variation. In this section of the report, we describe work done to decompose and categorize the scattered return. In Section 4.1, we describe a method developed under this work for accurately modeling only the elastic portion of the response. In Section 4.2 we give an overview of the graphical user interface (GUI) we have developed for visualizing the internal physics associated with observables in the return, and provide examples of such visualizations. In Section 4.3, we perform a wide survey of key observables in the 0-10 kHz range of the response and categorize them according to the internal physics. In Section 4.4, we describe methods of modifying the base FE model to isolate classes of physics behavior even when they would normally be co-occurring with other physics.

4.1 Separation of Elastic and Inelastic Components of Response

APL-UW has used COMSOL to develop solid/acoustic FE models for underwater elastic targets interrogated with acoustic energy [3]. For a complete representation of the physics a researcher would model the pressure field created by the interrogating source as p_{inc}. The scattered response is p_{scat}, so it follows that the total acoustic pressure field is $p_{tot} = p_{inc} + p_{scat}$.

When we simulate the complete physics of a shell interrogated with a plane wave, the pressure appears as a fluid load on the exterior boundary of the shell:

$$\mathbf{F} = -p_{tot}\mathbf{N}, \tag{1}$$

where $\mathbf{N}$ is the outward-facing normal for each point on the exterior of the shell. The solid mechanics physics of the shell's interior are then coupled back into the water by the constraint that the normal component of the water acceleration be equal to that of the solid acceleration at every point on the solid-fluid interface:

$$\mathbf{a}_{fluid} \cdot \mathbf{N} = \mathbf{a}_{solid} \cdot \mathbf{N}. \tag{2}$$

Given these straightforward equations and appropriate boundary conditions (including perfectly matched layers (PMLs) [4] at the exterior of the modeled water to prevent reflection of the scattered wave), we produce an estimate of the scattered response. We view modeling of individual physical mechanisms contributing to the complete response as a process of turning off particular manifestations of the incident pressure wave or of the physics representation of the geometry itself.

The most conceptually straightforward restriction would seem to be that which isolates the elastic and inelastic portions of the response. In practice, however, there are different ways to approach this goal which produce slightly different results. In our current work we have developed what we believe to be the most physically justifiable approach to this isolation.

One method is to define the inelastic response as the scattered response produced by an inelastic object with the same shape as the target of interest, that is,

$$\mathbf{a}_{fluid} \cdot \mathbf{N} = 0 \tag{3}$$

and to define the elastic response as the residual produced when the inelastic response is subtracted from the complete response. This approach was taken by España *et al.* [5], and has the advantage of being easy to compute. However, it has the potential to produce artifacts in the elastic response because the inelastic response estimate (by definition) assumes all the energy present in the acoustic wave remains in the water, which means it is an overestimate of the portion of the response due to inelastic behavior and subtracting it from the complete response produces an underestimate of the elastic response. Further, the manner in which this response is 'underestimated' is highly dependent on the geometry and acoustic properties and cannot be compensated for in a straightforward fashion.

As a result, we took a different approach to modeling the elastic response in some of our previous work [3] and in the SEED effort which preceded the current investigation. In this case, we modeled the elastic response in a separate step by stimulating the internal movement directly rather than as some effect of incident pressure in the fluid. To reproduce the entirety of the elastic response, we can reduce the forcing function only to the direct load on the target

$$\mathbf{F} = -p_{inc}\mathbf{N}. \tag{4}$$

By investigating the response due to this fluid load alone, we were in effect "shaking" the target directly and isolating the elastic response without ever simulating the incident wave in the acoustic domain. Previously, the approximation of the total pressure field load as that of the incident pressure field was more problematic in some places than others, due to the modeling capabilities at the time. An obvious case is the fluid region at the rear of the target (relative to the source), where the effects of the incident pressure are greatly reduced (and in any case not correctly modeled) by p_{inc}. However, recent advances in the modeling approach by Dr. Marten Nijhof at TNO and Dr. Aubrey España at APL-UW [6] allowed us much more flexibility in customizing the loading region, which means that we could more accurately model effects of target shadowing on the fluid loading and produce a better estimate of the portion of the response due to the elastic behavior of the target. It was noted, though that this approach suffered from three weaknesses:

- it required a separate and complicated calculation of the shadow region for each angle of incidence of a plane wave
- the true pressure field was not simply shadowed/unshadowed but could be a gradual (and complex) transition depending on the target geometry
- as a result, the internal physics of this approach differed from the internal physics of the complete model

In the current work, we have developed a new approach to isolating elastic and inelastic behavior which addresses all of these issues. This method requires a two-stage approach in which we first model the scattered response due to the complete elastic and inelastic physics. We then model the scattered response that would be produced by an elastic object with identical internal physics (i.e. displacement) in the absence of an incident wave. This is the true elastic response, and the inelastic response can be produced by subtracting that result from the complete response. No

approximations to the pressure field are required and the internal physics are (by definition, of course) identical to the complete response. The inelastic response estimate is also more accurate as it does not suffer from the overestimation issue raised previously.

The results of this elastic-isolation process as applied to the science target can be seen in Figure 3. The full-physics response is shown in the top figure, and the portion of the response due to elastic behavior is shown in the middle subfigure. The bottom figure is a coherent subtraction of the two for validation; it can be observed that the broadside and flat-end responses have strong broad-band energy, with interference patterns at other angles consistent with expectations. In the remainder of the physics investigations described in this report, we will primarily use the elastic-only response to guide our focus on observables of potential value for classification.

4.2 Physics Investigation Framework

The remainder of Section 4 describes our investigation into the relationship between observables in the scattered returns and the internal physics which produce them. To aid in this investigation, under this project we have produced a GUI-based software system which allows for visual exploration and quantitative surveying of all the modeled physics (the physics having been pre-calculated using the COMSOL models generated in this project). This system, implemented in Matlab, allows the user to view the acoustic color representation of the response of a free-field model of a target and then click on any observables of interest and dig down into the modeled internal physics which produced that observable. Further, the system allows for processing to be done on the entirety of the modeled physics across all aspect angles and frequencies for the purpose of identifying mathematically regions of the frequency-aspect space in which certain classes of physics behavior may be occurring. All the figures in Section 4 have been generated using this system.[2]

An example of the use of this system begins with the acoustic color plot in Figure 4, representing the farfield response of the science target in freefield. Using the separation method described in Section 4.1, we have isolated the portion of the response due to the target's elastic behavior (though the complete response could also be used). The user has chosen to investigate the physics of three points on the frequency-azimuth plane, which were automatically marked on the acoustic color plot as the user clicked on them. With each click, plots representing all of the internal physics selected by the user are generated. The user has the option of selecting any or all options among:

- displacements, accelerations, stresses, and strains
- real component, imaginary component, magnitude, and phase of the above quantities
- representation in global (xyz) coordinates or target coordinates (normal to the surface, tangential in the axial direction, and tangential in the circumferential direction)

As will be seen in the following discussions, each of these quantities and representations may be illuminating in identifying classes of physics behaviors.

Given the current focus on axisymmetric targets with relatively smooth surfaces, we have chosen a method of data visualization referred to as 'unwrapping' in which we represent the data along a

[2]The results in Section 6.1 of this report are presented in the same format but were generated outside the GUI system because they were modeled using a method that produces a different format and because they represent targets on the seafloor, which requires additional data processing to account for internal physics superposition due to multipath.

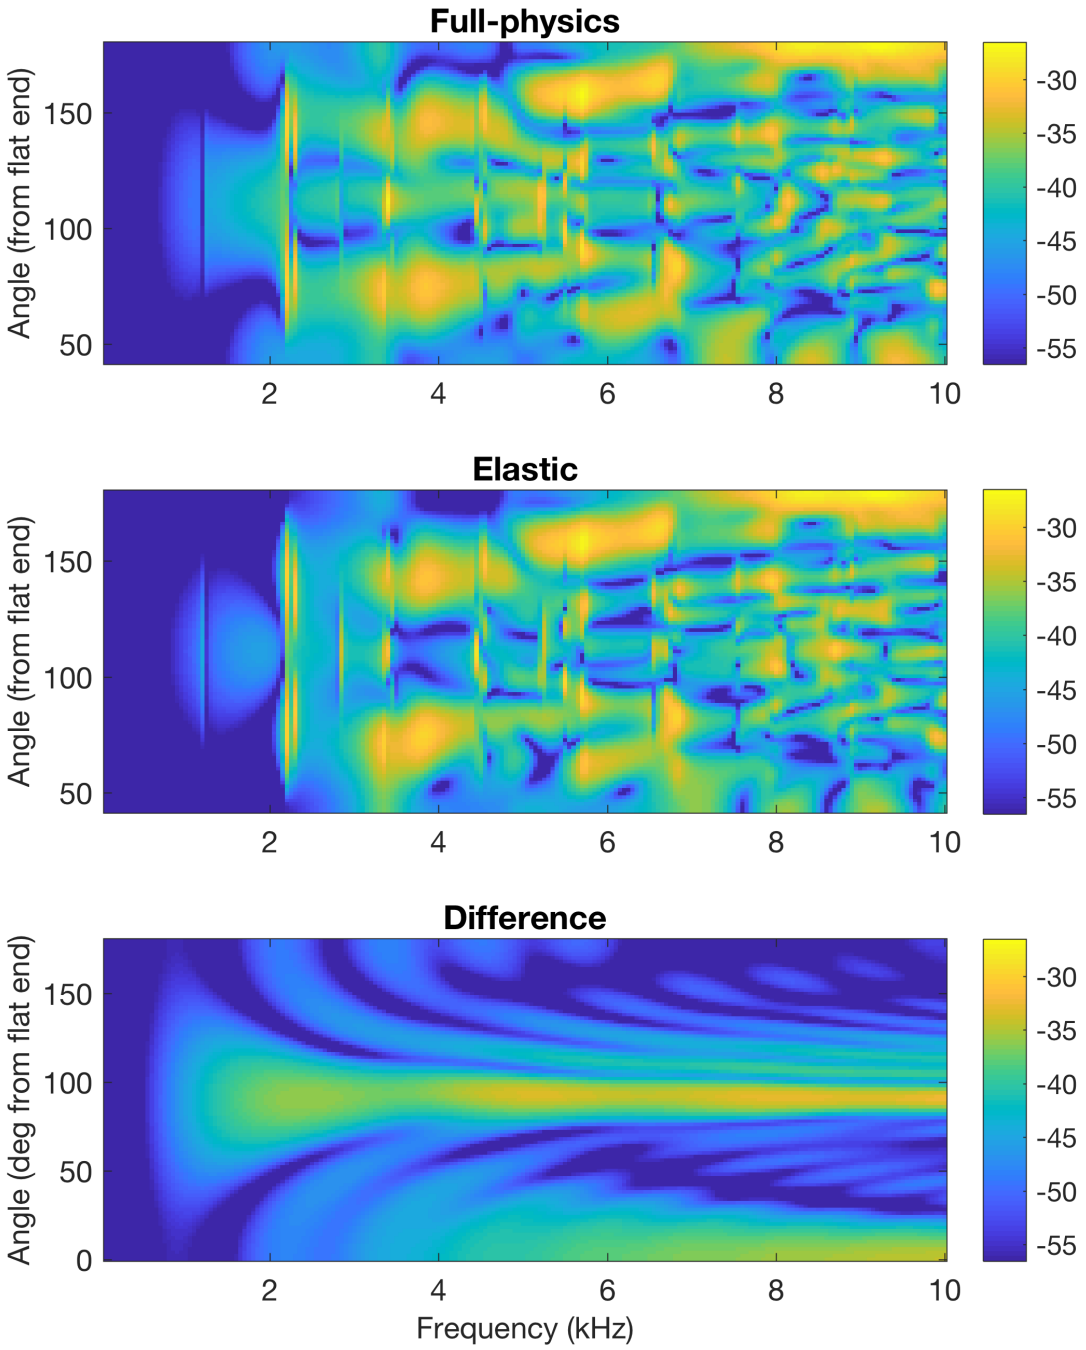


Figure 3: Separation of elastic and inelastic response using the modeling technique developed under this project. Once the full-physics model is created, the elastic-only scattering is created by reproducing the internal movement of the target without any incident wave. The coherent difference between the two is the inelastic or shape response.

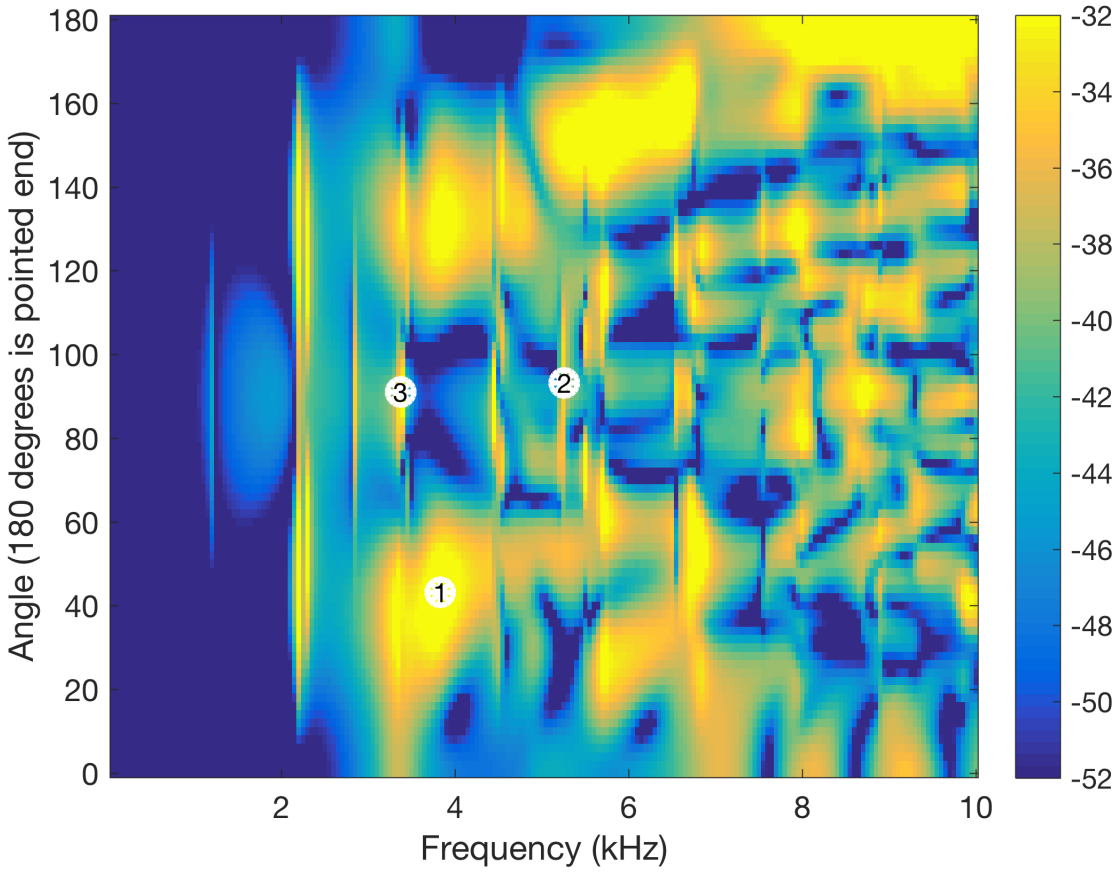
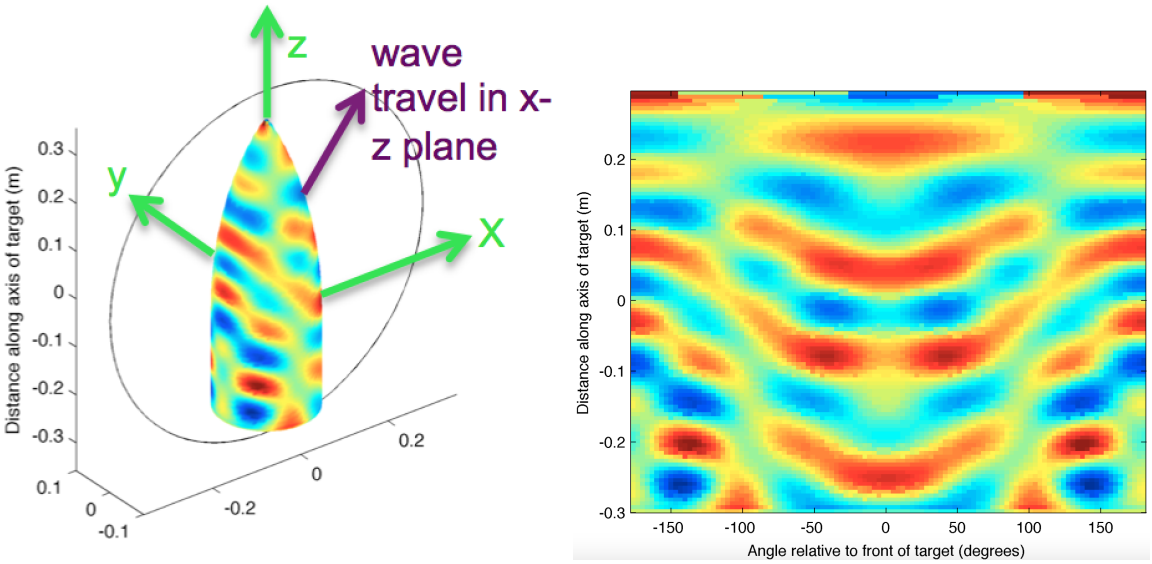


Figure 4: An example acoustic color image from the physics-investigation GUI used to investigate the return from the air-filled science target. The user has selected three points for which to visualize the internal physics; the results are shown in Figures 6 through 15.



(a) Surface pressure: true geometry

(b) Surface pressure: unwrapped visualization

Figure 5: Visualizing internal physics of an axisymmetric target by unwrapping the physics on a curved surface. In this example, both figures display the same information. The surface pressure is represented in true spatial coordinates in subfigure (a); in subfigure (b) we use the x-axis to represent degrees from the front (where 'front' is defined based on a plane wave traveling in a direction in the x-z plane as illustrated by the purple arrow in subfigure (a).

particular axisymmetric 2-D manifold in a plane where the x-axis is circumferential angle and the y-axis is distance along the surface of the object as we move in the axial direction. This can be visualized in Figure 5, where we show the same data (in this case pressure on the object surface) in true spatial coordinates and again in an unwrapped fashion.[3]

The visualizations produced by clicking on the points in Figure 4 is shown in Figures 6 through 15. We have shown the full range of visualizations of displacement and strain for Point 1 and select visualizations for Points 2 and 3. The specific information obtained from these visualizations (and associated automated processing) will be discussed in Sections 4.3 and 4.4. Note the black and white lines superimposed on the figures; these indicate the planewave wavefronts, with one black line per period and white lines at the half-period.

In our investigation, we use this system as part of the process of identifying classes of physics behaviors as follows:

1. Identifying an observable in the acoustic color plane which is strong enough to be potentially useful for classification (or which has been identified based on existing classification results as useful)

2. Looking at the target's internal physics behavior at an incident angle and frequency where the observable is strong

[3]While the unwrapped representation is excellent for viewing many physics behaviors, other representations might be required for understanding other scenarios. For example, investigation of changes in behavior as a function of different depths in the object, or on objects which lack axisymmetry, would require different visualizations.

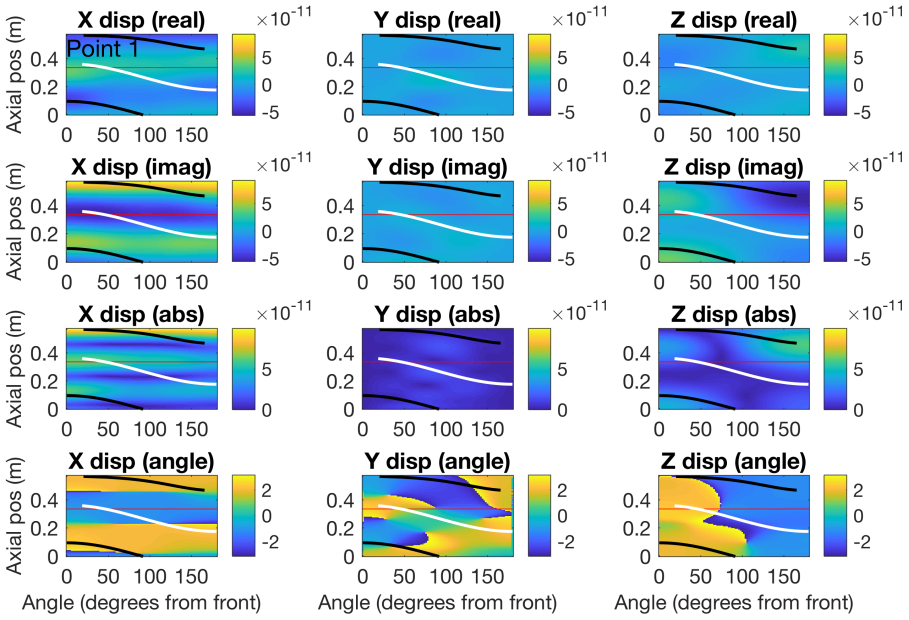


Figure 6: The unwrapped complex displacements (in *xyz* coordinates) for Point 1 from Figure 4. The black lines indicate wavefront periods and the white lines are half-periods.

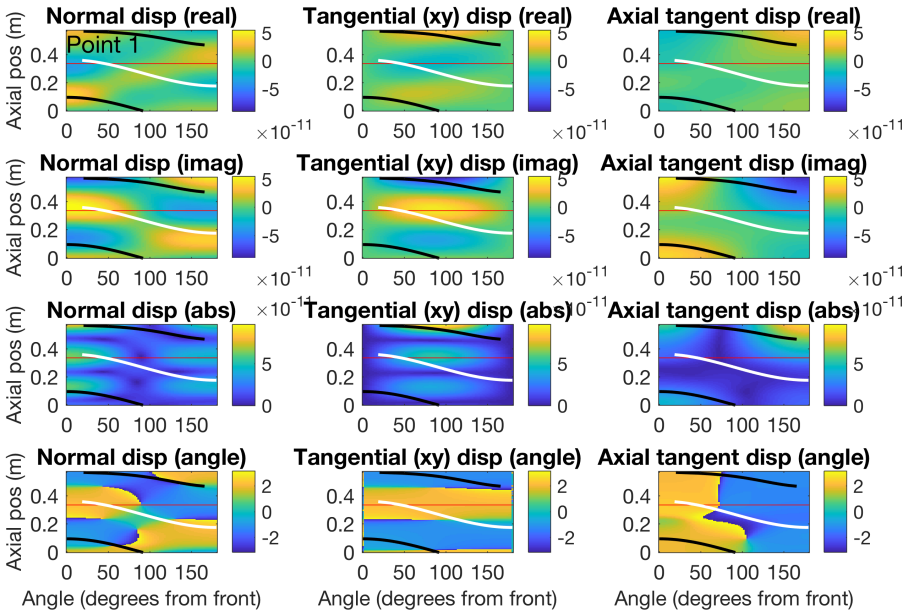


Figure 7: The unwrapped complex displacements (in surface normal-tangential coordinates) for Point 1 from Figure 4

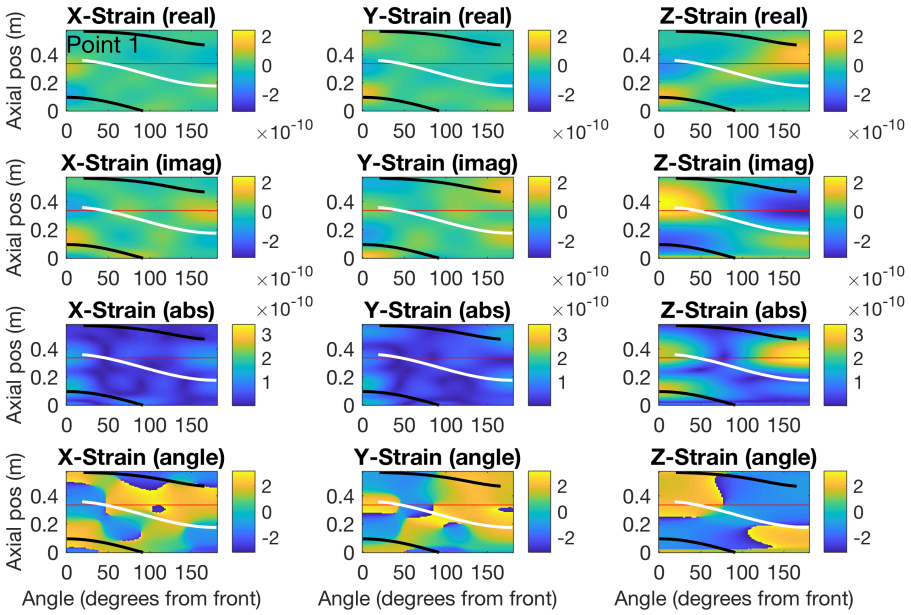


Figure 8: The unwrapped complex axial strains (in *xyz* coordinates) for Point 1 from Figure 4

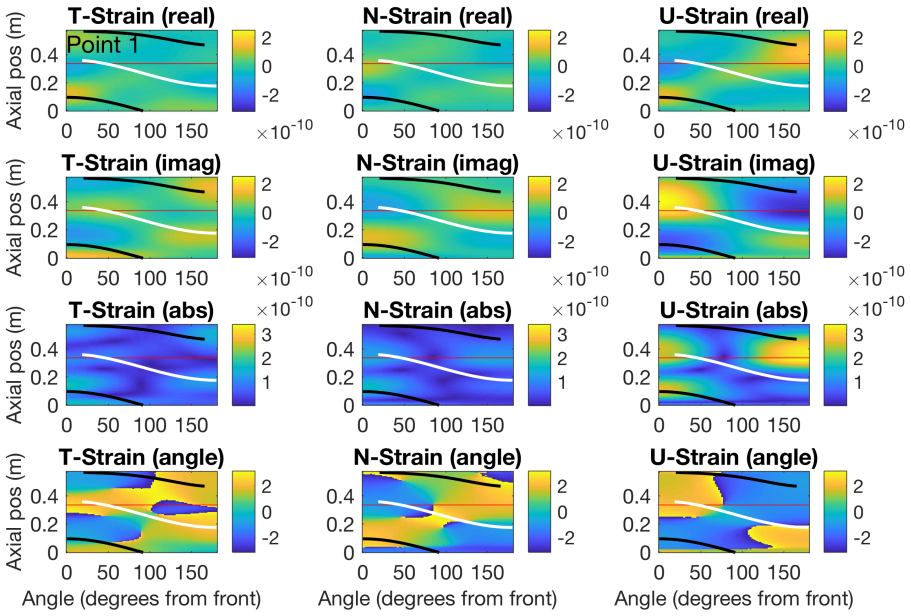


Figure 9: The unwrapped complex axial strains (in surface normal-tangential coordinates) for Point 1 from Figure 4

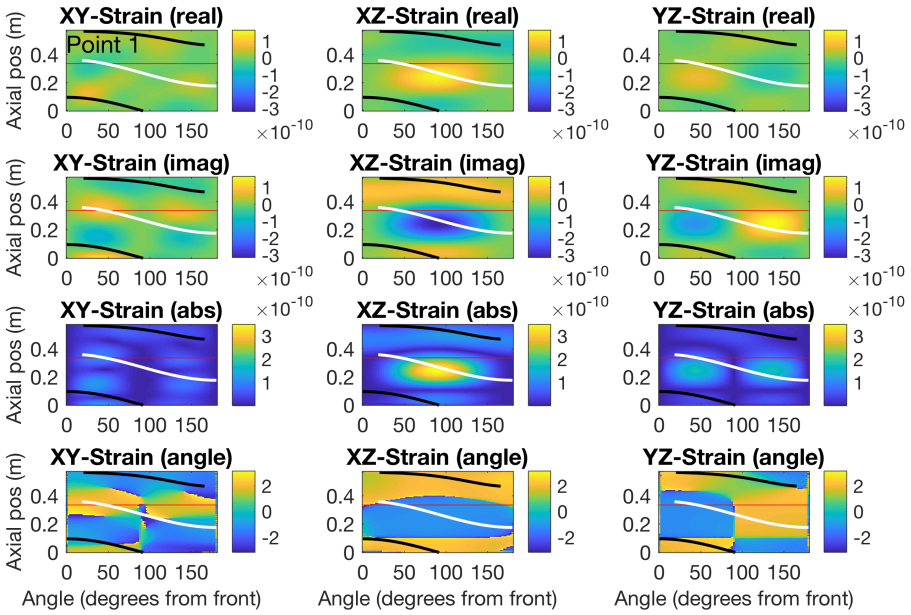


Figure 10: The unwrapped complex shear strains (in *xyz* coordinates) for Point 1 from Figure 4

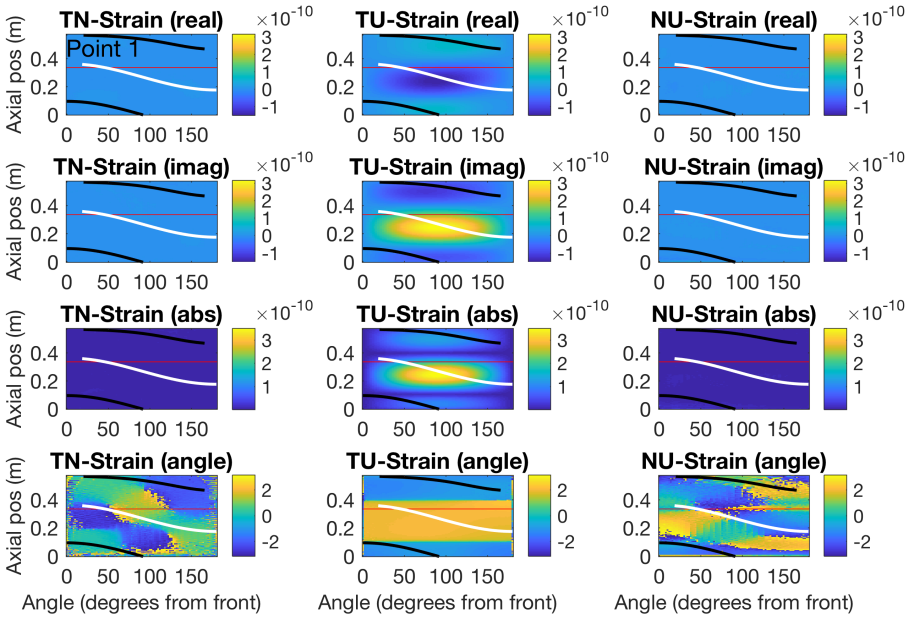


Figure 11: The unwrapped complex shear strains (in surface normal-tangential coordinates) for Point 1 from Figure 4

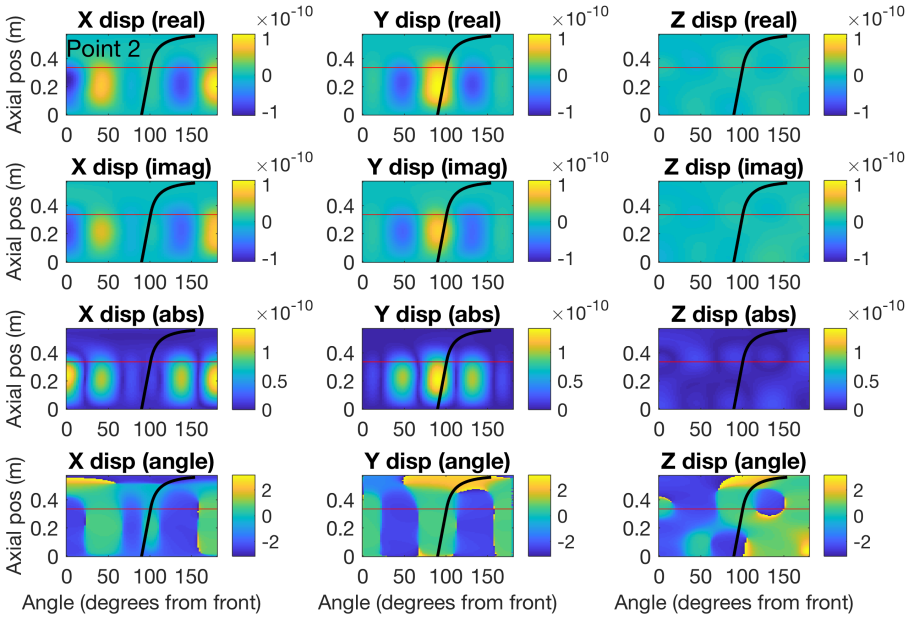


Figure 12: The unwrapped complex displacements (in *xyz* coordinates) for Point 2 from Figure 4

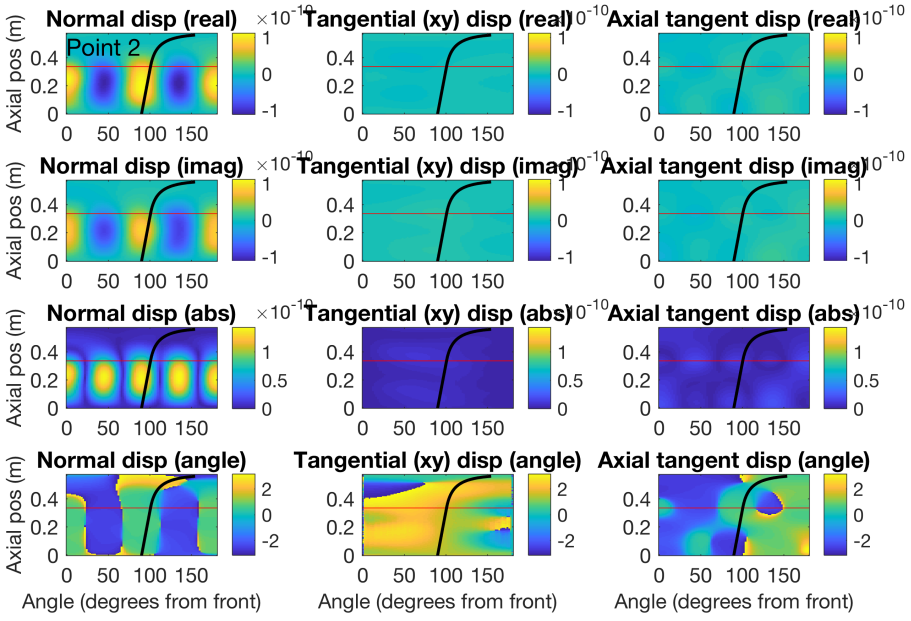


Figure 13: The unwrapped complex displacements (in surface normal-tangential coordinates) for Point 2 from Figure 4

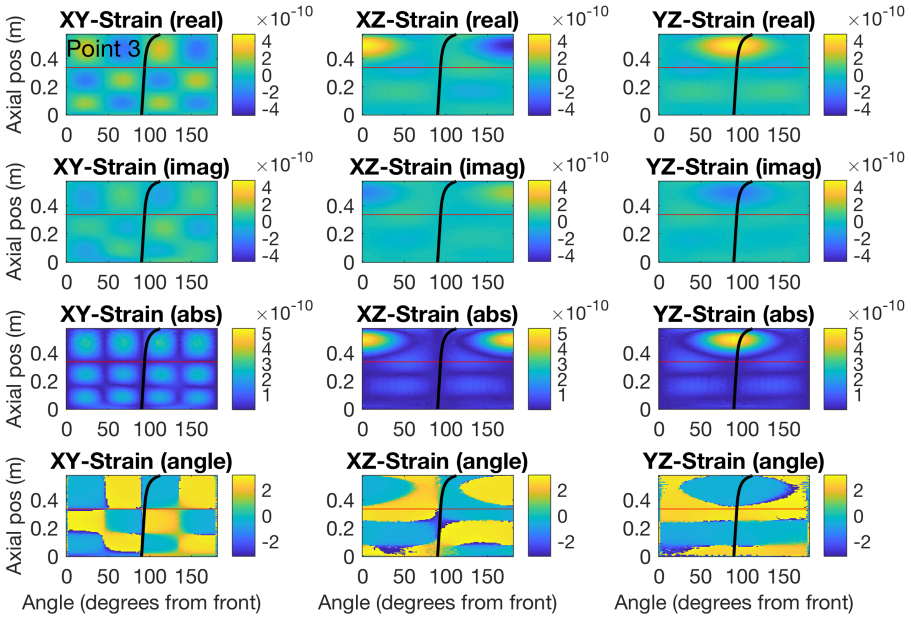


Figure 14: The unwrapped complex shear strains (in *xyz* coordinates) for Point 3 from Figure 4

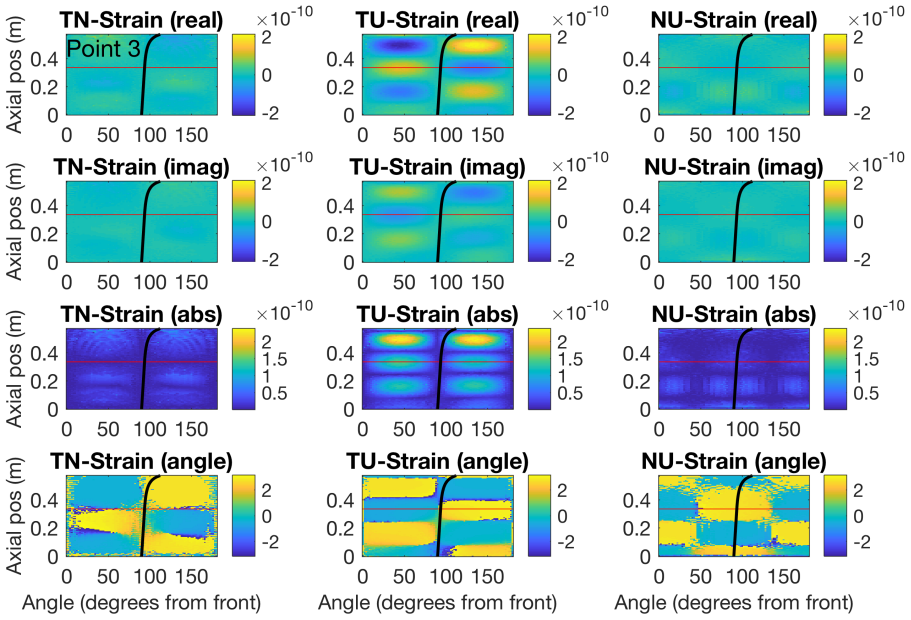


Figure 15: The unwrapped complex shear strains (in surface normal-tangential coordinates) for Point 3 from Figure 4

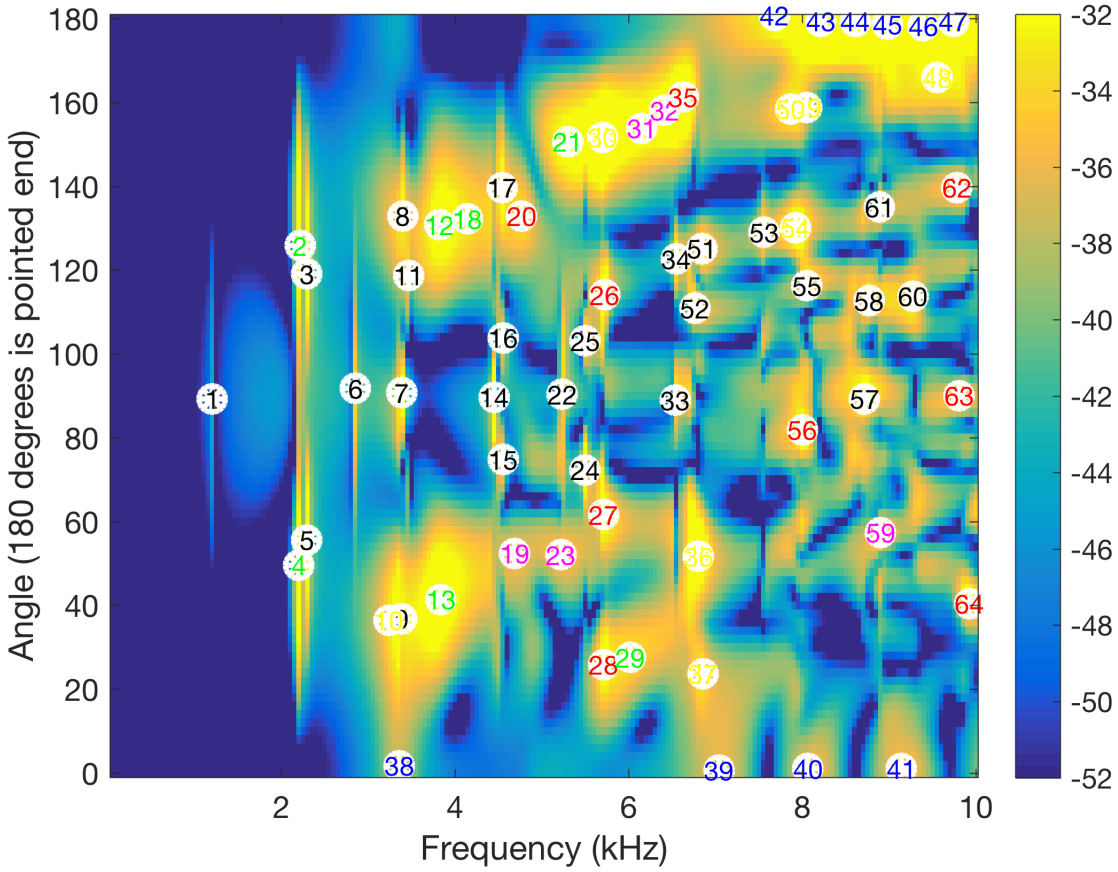


Figure 16: The complete set of points whose internal physics will be categorized in the survey of Section 4.3

3. Understanding how the observed physics can be related to the specific target composition (geometry and materials)
4. Identifying a class of physics behavior underlying the observable, and creating a mathematical measure which can be applied to the internal physics models of all incident angles and frequencies to identify when this class of physics occurs in isolation
5. As necessary, creating a new model which has constraints on the physics which attempt to isolate the entirety of the true return which is due to this class of behaviors, even when they are normally co-located with other physics behaviors

The remainder of Section 4 of this report provides examples of the results obtained under this approach.

4.3 Survey of Elastic Response

In this section we describe a manual survey of the observables in the portion of the return due to the elastic behavior of the target. The complete set of points discussed in this report is shown in Figure 16. As this investigation is performed, we identify potential classes of behavior which can potentially be related to the object's geometry and composition and targeted for representation in a feature set. Several such classes have been identified using this approach; additionally other observables have been identified which may be the result of other physics behaviors or some combination of the existing identified classes. In summarizing each of these classes, we show the subset of the physics quantities (such as displacement and strain) which illustrate the characteristics of that class. It should be noted that the boundaries between these categories are not absoulute and in some cases points could fit the definitions of more than one category. We would like to emphasize that, although we are attempting to characterize as much of the energy in the scattered returns as possible, it is not a necessary condition for performance improvement that the entirety of the response be understood in terms of the underlying physics. To the extent that a few key observables can be represented with a feature set that allows for variation with environmental conditions, detection and classification gains can be achieved.

In Figure 17 are the subset of the surveyed points that represent bending modes. To understand how the internal physics can be identified as a bending mode, consider the displacement physics shown in Figure 6. The incoming plane wave is travelling in a direction on the x-z plane, as shown in Figure 5. The qualities that define bending are: a) very little movement in the y-direction (orthogonal to wave travel) and b) similar movement patterns for points at the same location along the target's main axis. The former can be seen in the low y-displacement magnitude in Figure 6, and the latter can by observed by noting the stripes in the x-displacement section of that figure. In Figure 18, we show the internal physics for the complete set of identified bending observables; for brevity we show only the x-displacement. The position and strength of bending modes are dependent on (and thus can provide information about) target length and internal composition.

The second category of internal physics with which we have identified observables is that of a grid of standing waves where displacement is primarily in the direction normal to the surface. Such physics occur when the incoming wave acoustic wave causes displacement patterns that align with the fundamental geometry of the target. The point whose displacements are shown in Figure 13 is such a case; note the regular grid of the normal displacements and the much smaller magnitude of the displacements in other directions. In Figure 19 we show the set of points on observables that match this definition; the internal physics (normal displacement only) of these points are shown in Figure 20.

The next category of internal physics is similar to the previous one in that it is dominated by displacements in the direction normal to the surface, but the displacement patterns are more complex. The points in this category are shown in Figure 21, and the normal components of the displacements are shown in Figure 22.

Incoming waves closer to end-on create the next category of observables, which are waves travelling primarily in the axial direction. This set of points is shown in Figure 23, and the corresponding internal physics are shown in Figure 24. We note that there is significant movement not just in the normal direction but in the tangential direction aligned with the major target axis; most of the points in this category are travelling rather than standing waves.

The final survey category is that of helical waves; the set of points categorized in this way are

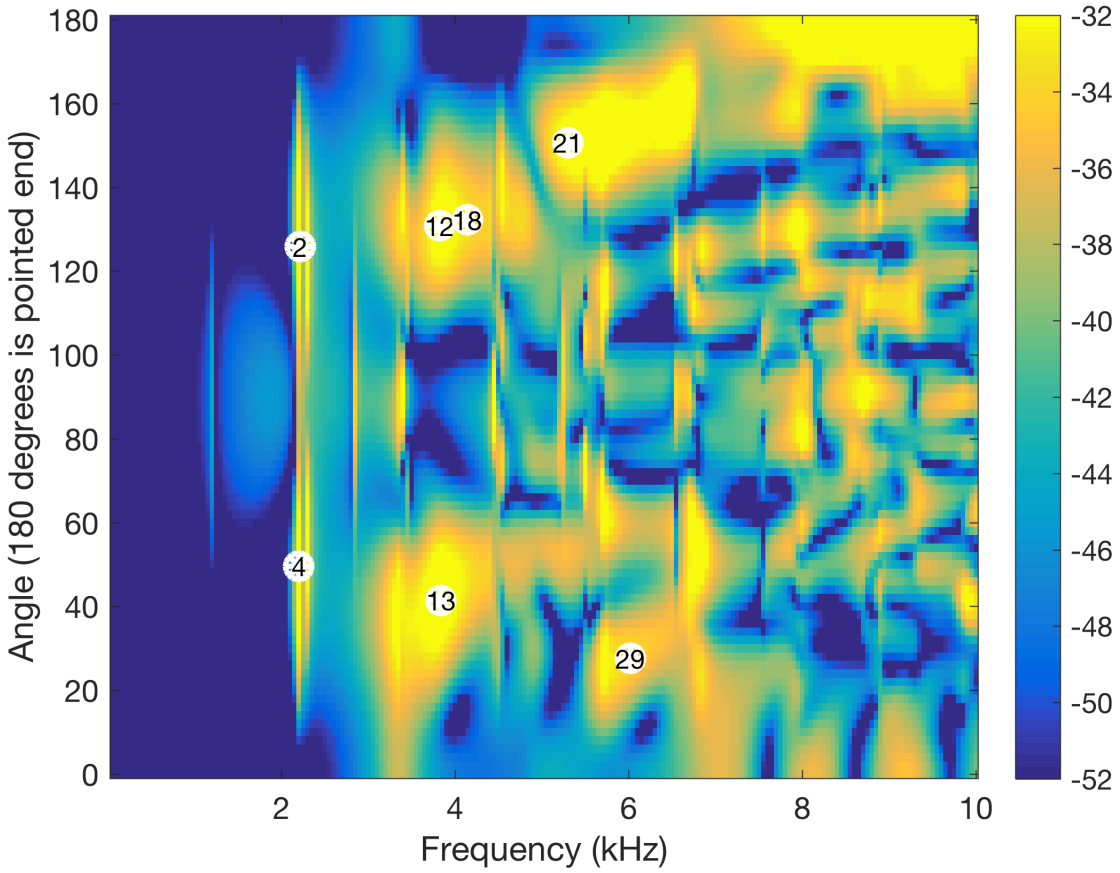


Figure 17: The subset of observables due to bending, which takes the form of movement in the x-direction dominating movement in the y-direction. Internal physics of these points are shown in Figure 18

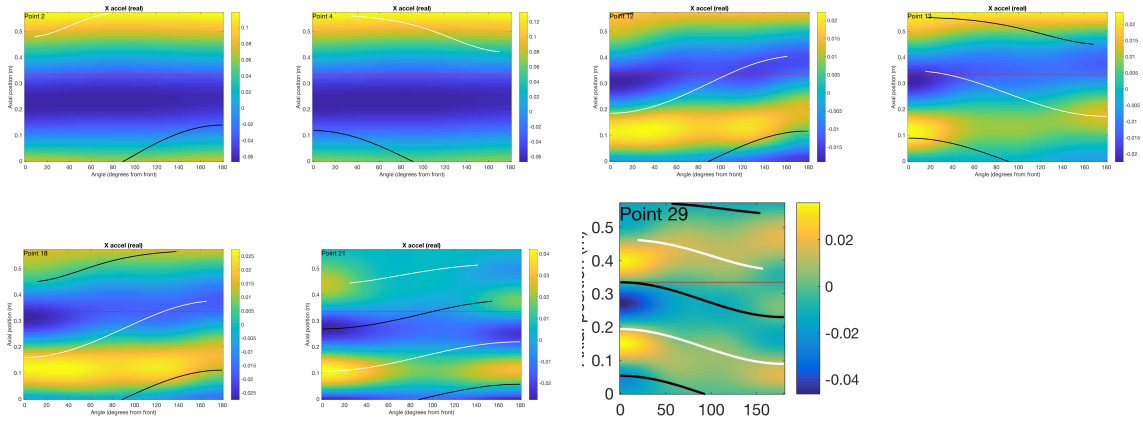


Figure 18: The normal x-displacement (real part only) for the points in Figure 17, whose internal physics are dominated by movement in the x-direction over the y-direction

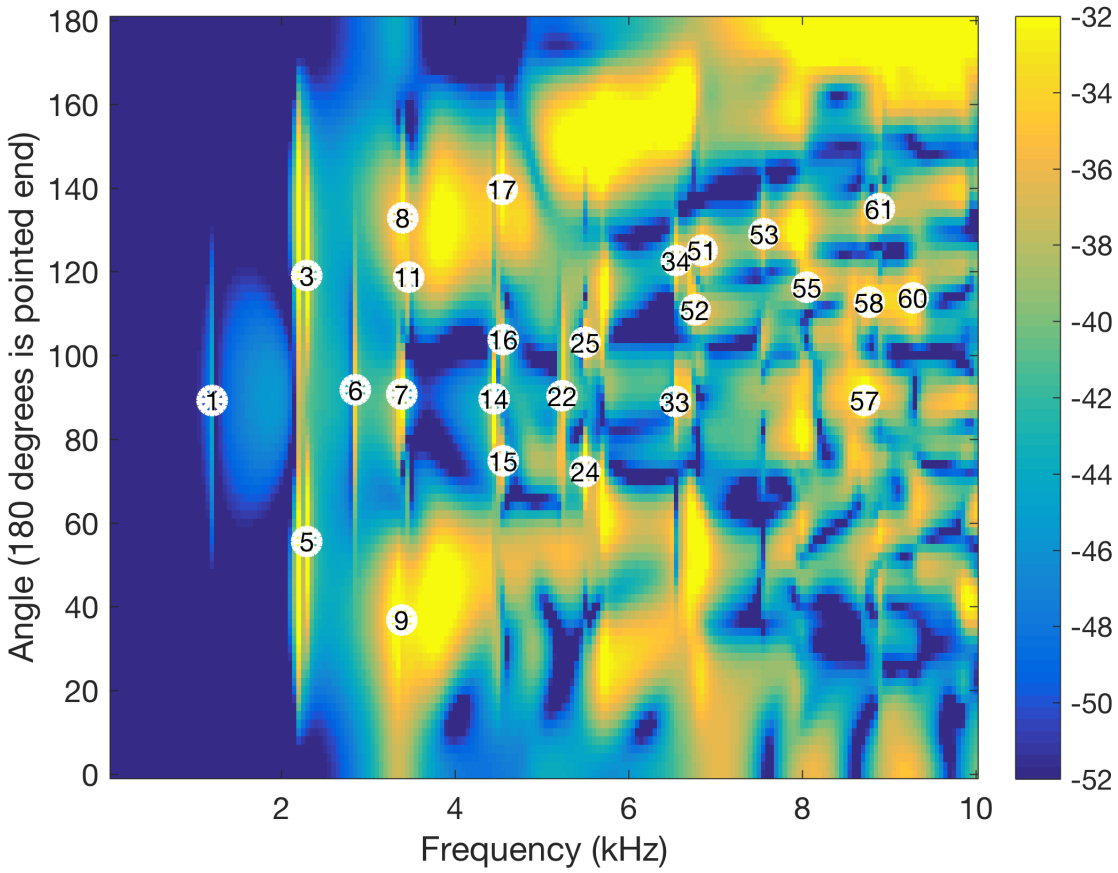


Figure 19: The subset of observables due to a regular grid of standing waves dominated by movement normal to the target surface. Internal physics of these points are shown in Figure 20

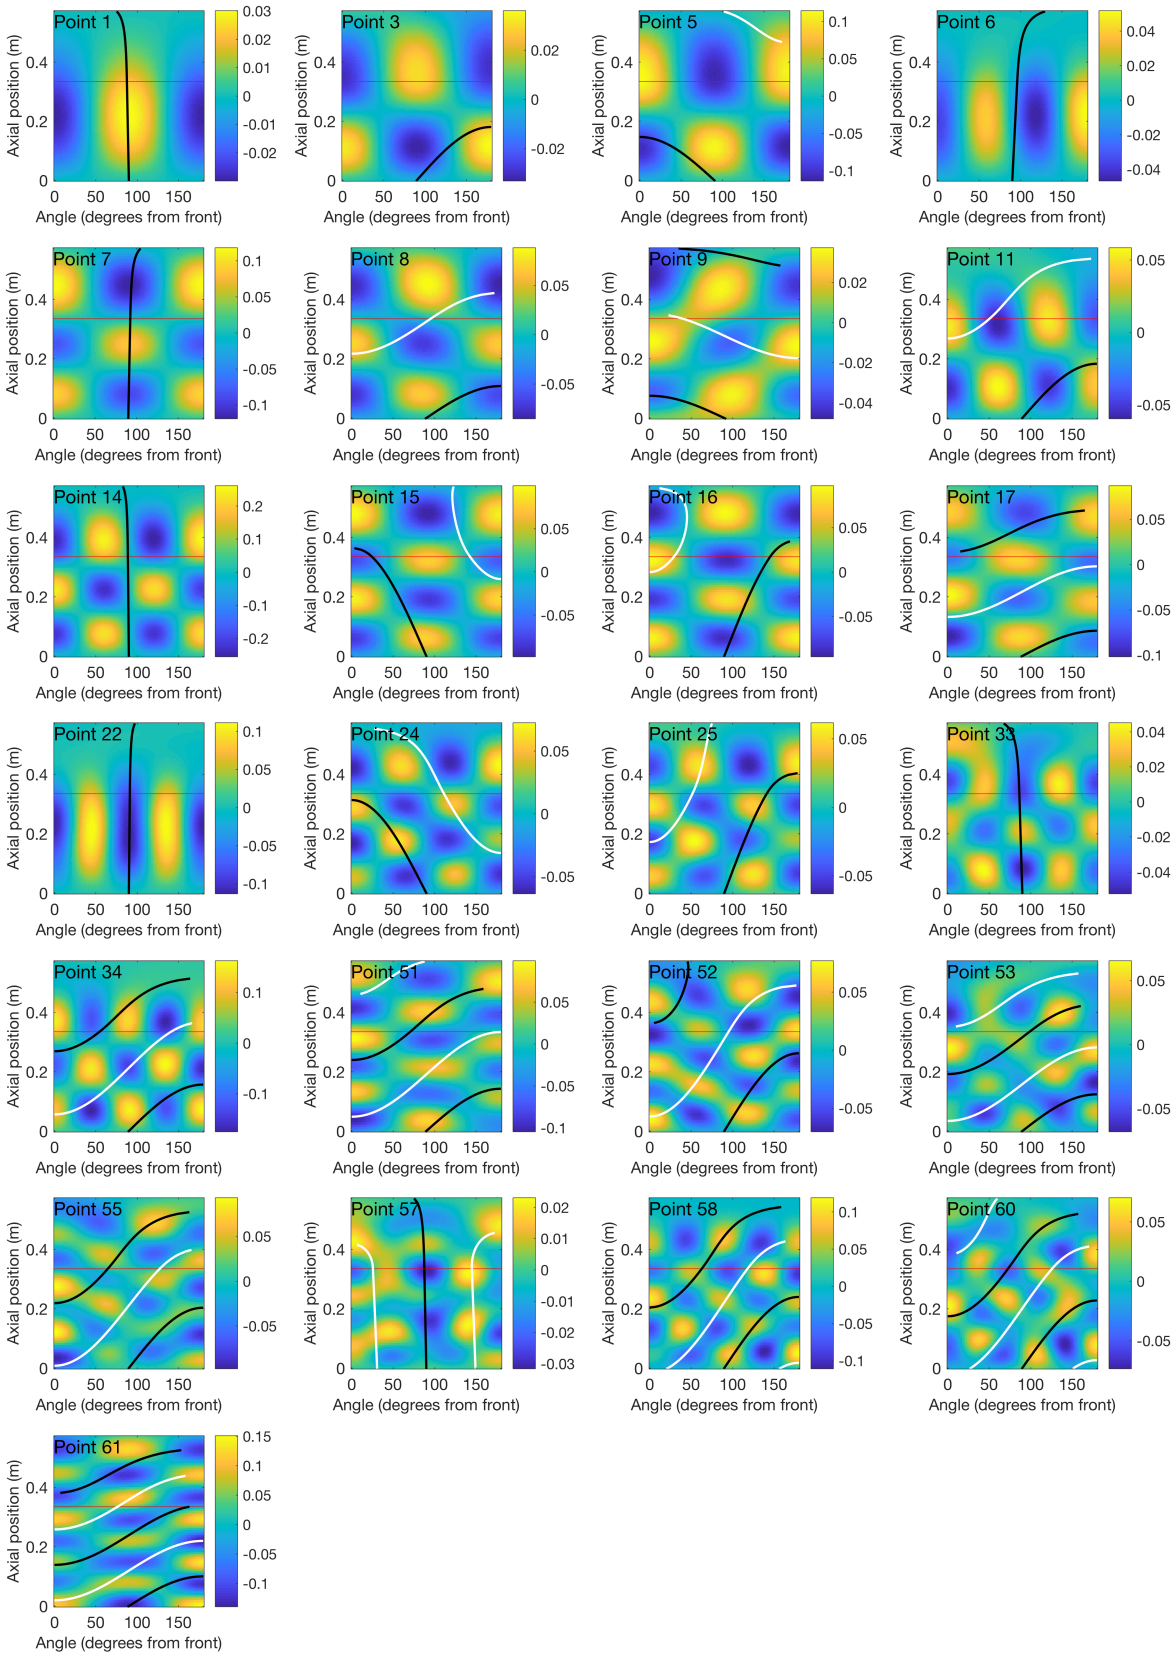


Figure 20: The normal displacement (real part only) for the points in Figure 19, whose internal physics are dominated by normal standing waves in a regular grid

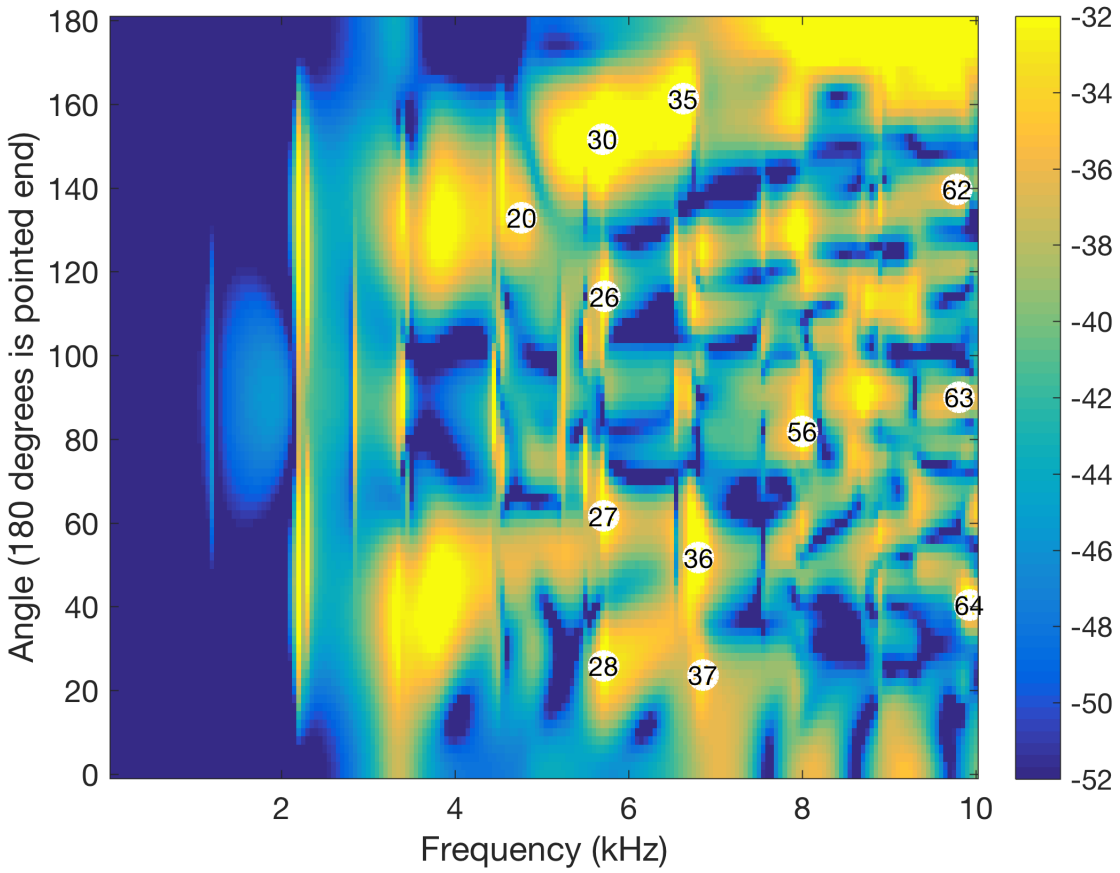


Figure 21: The subset of observables due to a regular grid of standing waves dominated by movement normal to the target surface but not (exclusively) in a regular grid. Internal physics of these points are shown in Figure 22

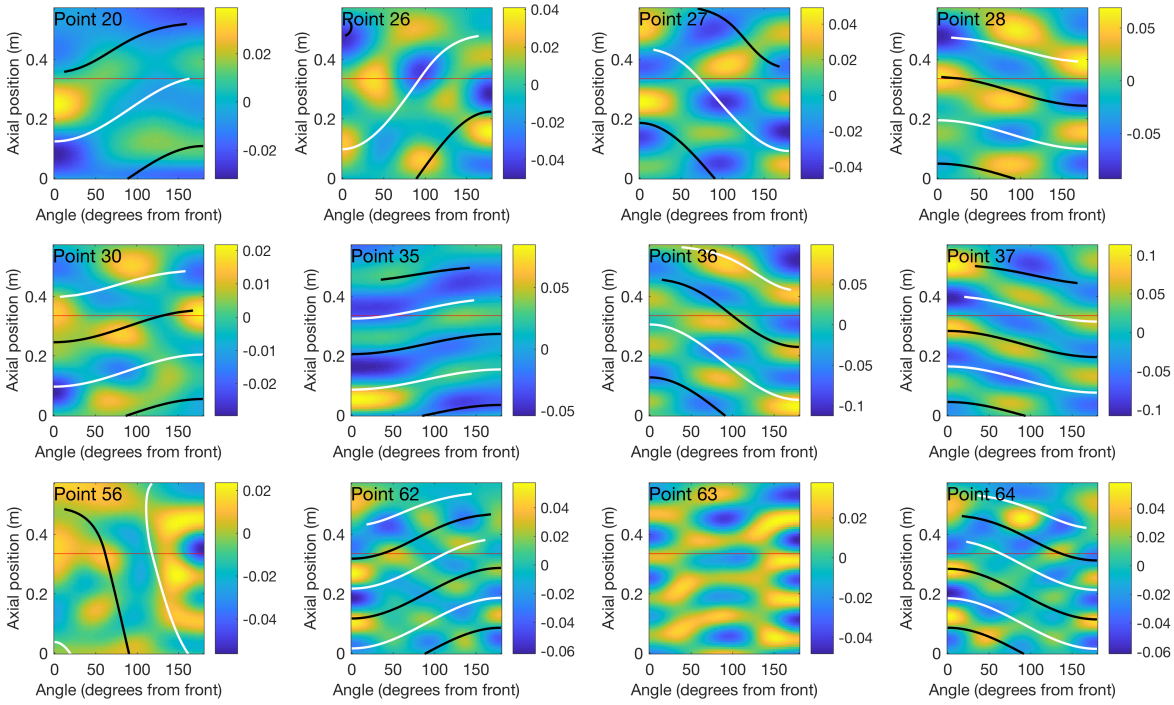


Figure 22: The normal displacement (real part only) for the points shown in shown in Figure 21, which are dominated by movement normal to the target surface but not (exclusively) in a regular grid.

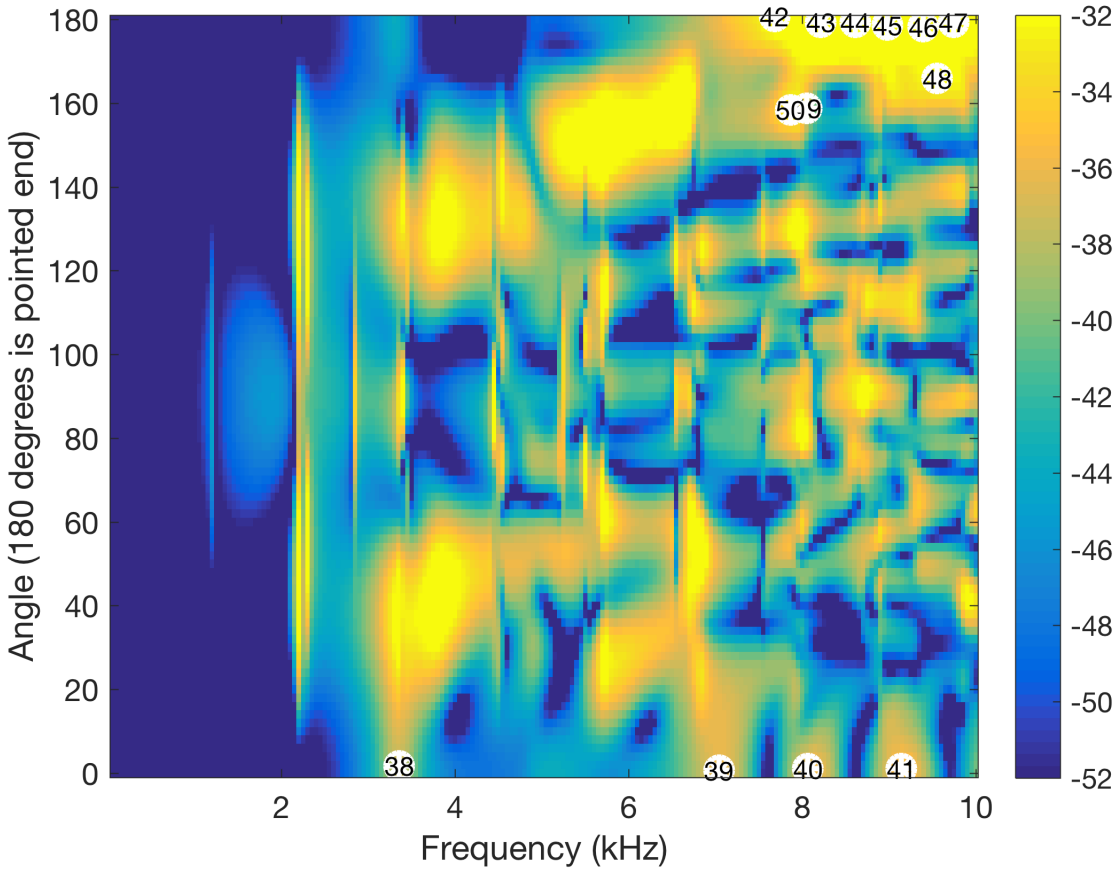


Figure 23: The subset of observables due to axial waves. Internal physics of these points are shown in Figure 24

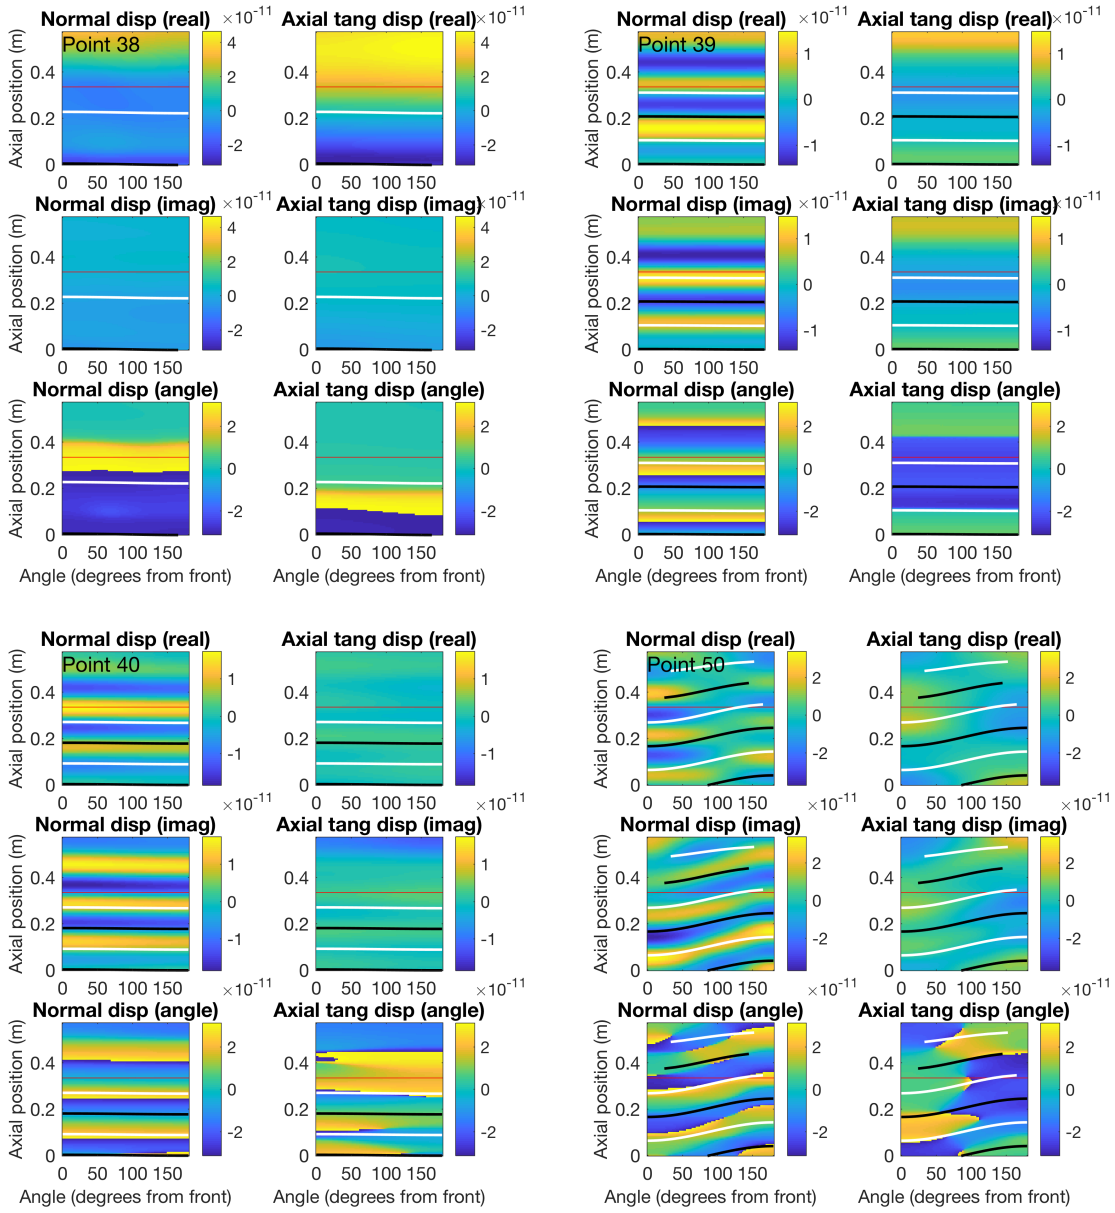


Figure 24: The normal displacement and axial-tangent displacement for a selection the points shown in shown in Figure 23

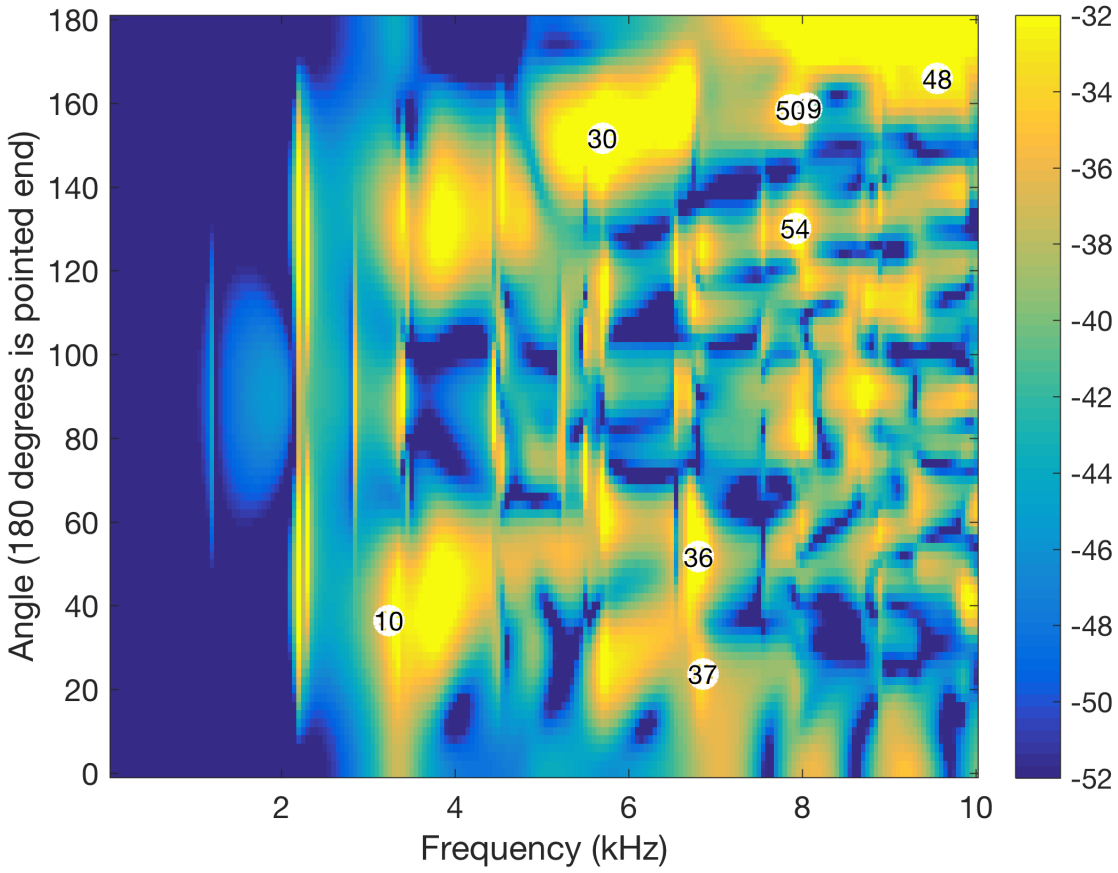


Figure 25: The subset of observables due to helical waves. Internal physics of these points are shown in Figure 26

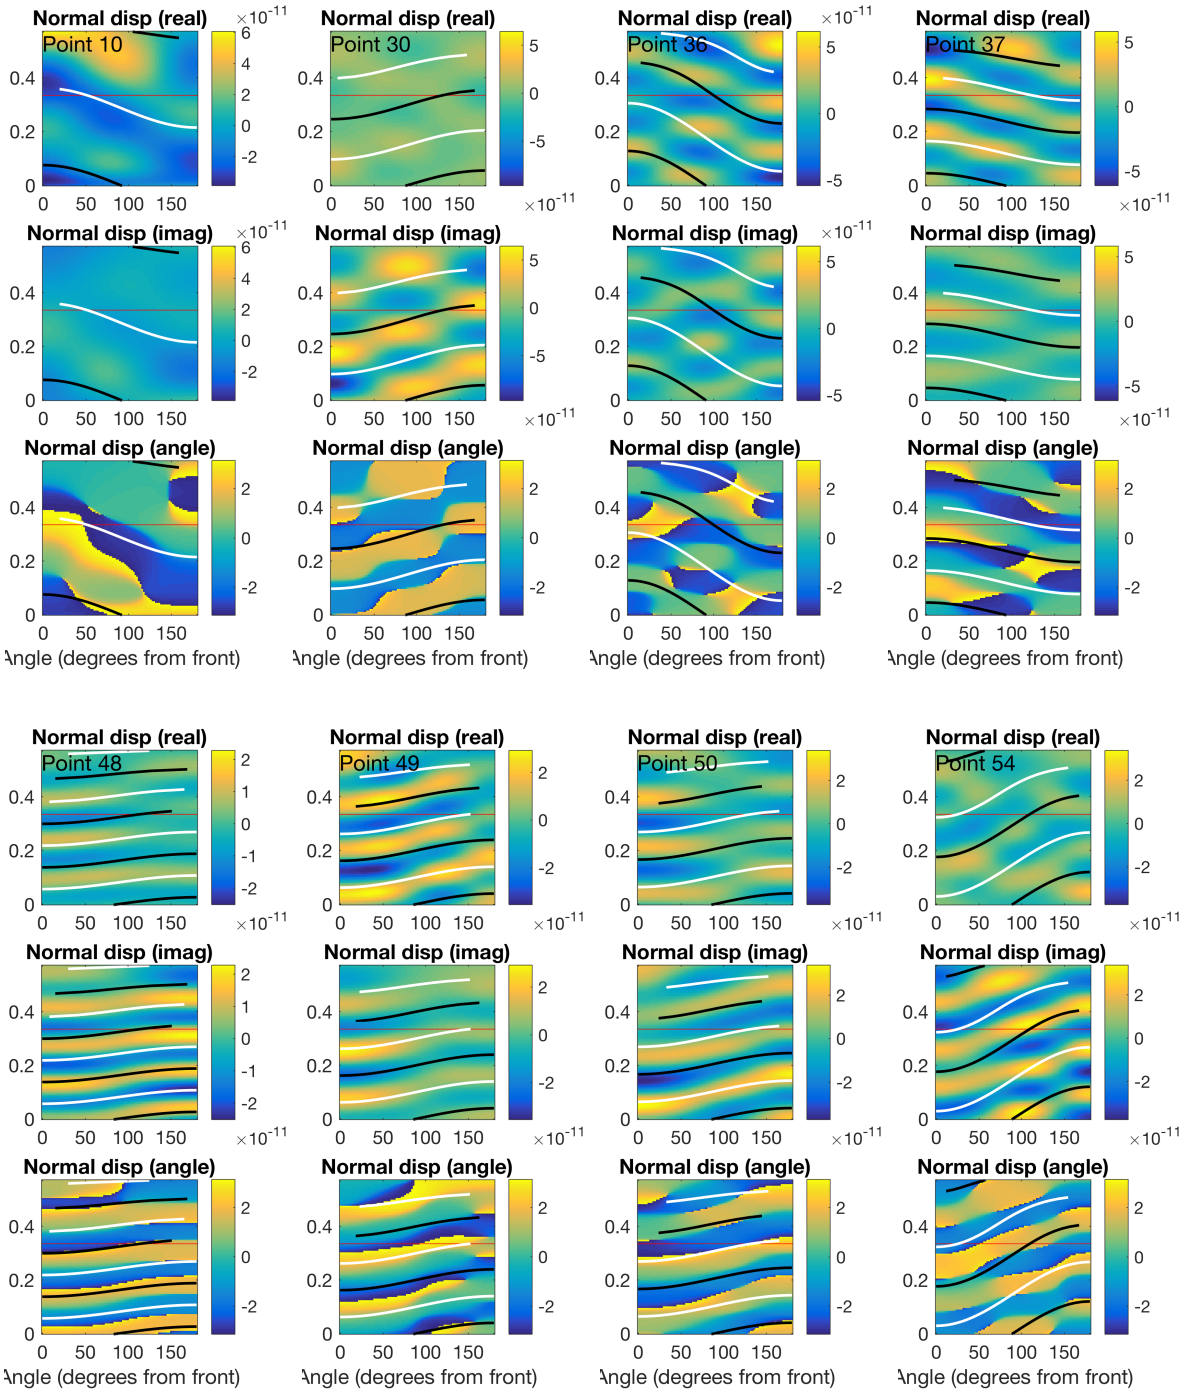


Figure 26: The *xyz* displacement (real, imaginary, and phase) for the points shown in shown in Figure 25. The phase information (and the GUI animation tool) indicate that most of these waves are moving.

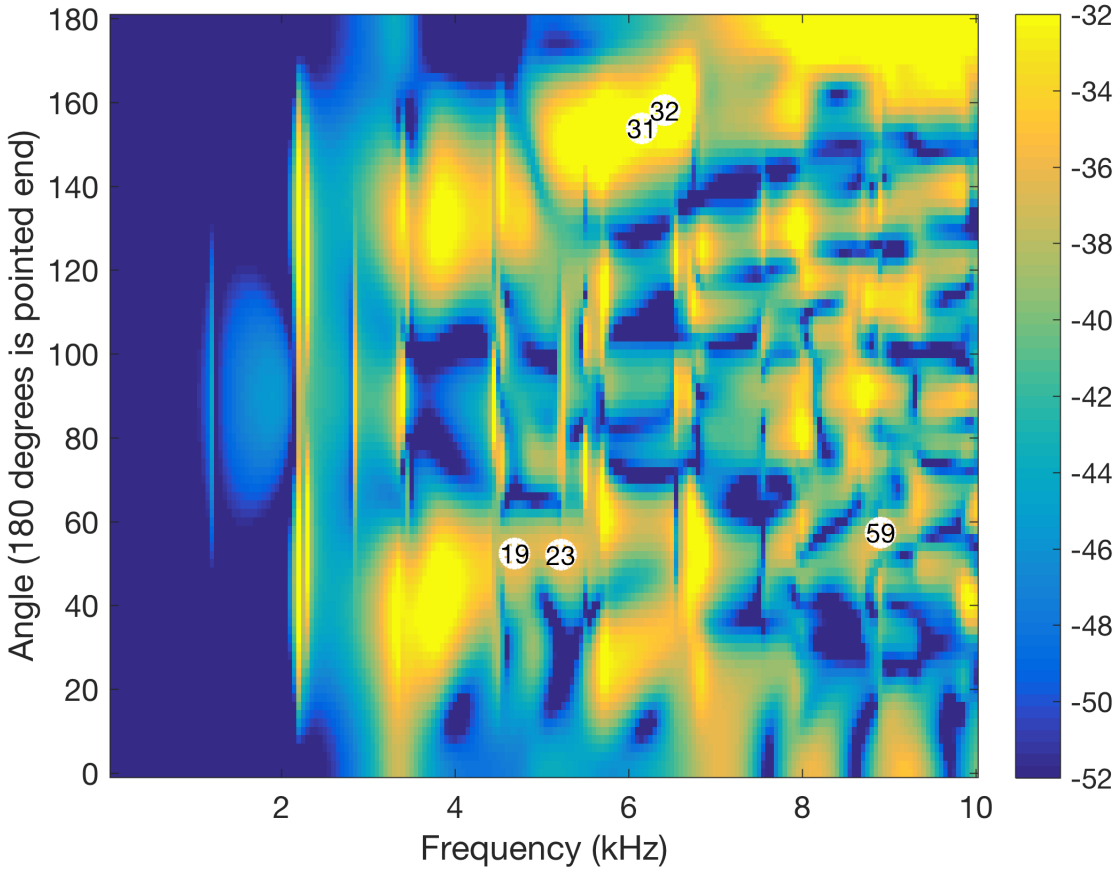


Figure 27: The subset of observables due to behaviors that are a combination of the other classes or do not fit into any of the classes. Internal physics of these points are shown in Figure 28

shown in Figure 25. The internal physics are shown in Figure 26. These are again standing waves and dominated by movement in the normal direction.

The remaining points in the survey are shown in Figure 27, and their internal physics do not fit easily into any of the above categories. They may represent cases where multiple behaviors are occurring simultaneously, transitional forms, or other behaviors not yet identified.

4.4 Internal Physics Constraints

For elastic objects simple enough to be modeled analytically (such as spheres and cylinders), scattered returns can be decomposed and understood in terms of the underlying wave behavior which produces the structure of the returns [5]. To some extent, the knowledge gained in this fashion can be transferred to more complex objects through *a posteriori* interpretation of the returns. However, the geometric deviations between the targets of interest and the simple shapes to which they bear some resemblance has meant that the above-mentioned interpretations have not proven accurate enough to result in any actual utility in improving feature representation or classifier design. As

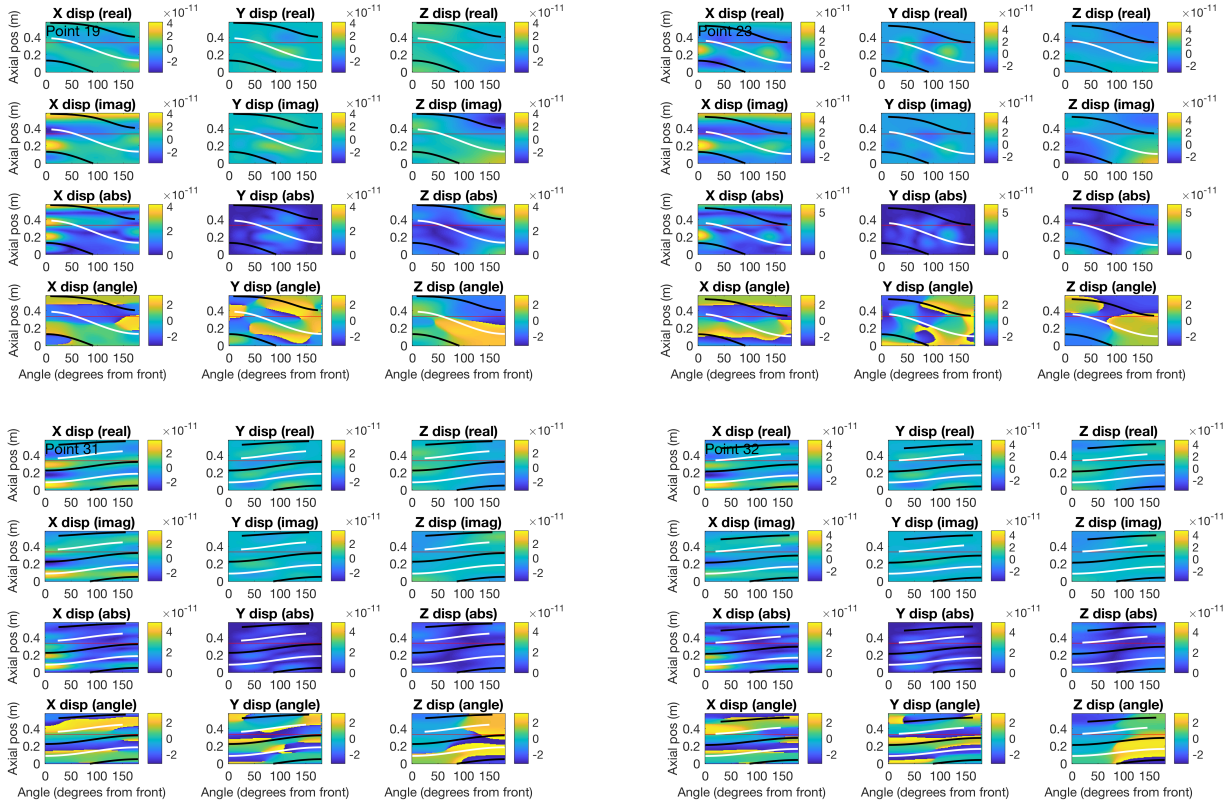


Figure 28: The *xyz* displacement (real, imaginary, and phase) for a selection of the points shown in shown in Figure 25. The internal physics for these points are either combinations of the other classes or do not fit into any of the classes.

shown in Section 4.3, though, this basic physics knowledge of the types of elastic waves which can occur often allows us to categorize the sources of observables by interpreting the modeled internal physics. In other cases, however, more complex internal physics are occurring at the aspect and frequency of the observable, possibly due to multiple physics behaviors being superimposed. To the extent that these classes of observables can be related to particular wave types and the specific geometric measures, their exact manifestations in frequency/azimuth space can be predicted rather than simply observed. As a result, variations due to shell and filler material as well as multipath can be searched over in a classification algorithm. It is not necessary to isolate the classes in the experimental data in order to measure the presence of such observables given a hypothesized target state. In Section 4.4, we investigate modifications to the physics to constrain the types of behaviors which can occur, with a goal of identifying the contributions of a class of behaviors to the response as a whole, even in places where they may be co-located with other behaviors. We have taken the following approach:

1. Identify a potential class of physics behavior, using a combination of basic physics knowledge and the results of the previously-described observable survey
2. Develop a mathematical description of this class which can be used to identify the extent to which interrogation at a particular aspect angle and frequency will produce this class of behavior
3. Apply this measure across all aspects and frequencies to produce a map of the regions where this behavior appears in isolation
4. When possible, develop constraints on the FE model that allow us to identify other regions of the aspect/frequency plane where this class of behavior occurs, even when it would normally occur in conjunction with other behaviors

In the early stages of this effort, we developed a toolset of methods within the COMSOL framework for modifying the physics to isolate classes of behavior. This toolset includes:

1. isolating portions of the return based on acoustic/elastic coupling regions (developed during our SEED effort, MR-2324)
2. isolating the elastic (and inelastic) portion of the returns, as described above in Section 4.1
3. constraining the direction of displacement within the target body
4. modifying the target material stiffness properties

In Section 4.4, we describe isolations of observables in the response based on both displacement constraints (#3) and material stiffness (#4).

Direction of Displacement: Bending Modes

The surveys described in Section 4.3 identified a number of classes of observables that could be described in terms of their directions of displacement. Bending modes, for example, are characterized by minimal movement in the direction orthogonal to the plane defined by the incident wave vector and the major axis of the target. The observables identified in Figure 17 were bending modes according to this definition. Recalling from the geometry in Figure 5 that the y-direction is the direction of minimal movement, we can develop a measure of the aspect/frequency points where bending is occurring in isolation. Figure 29 shows this measure, the ratio of the maximum

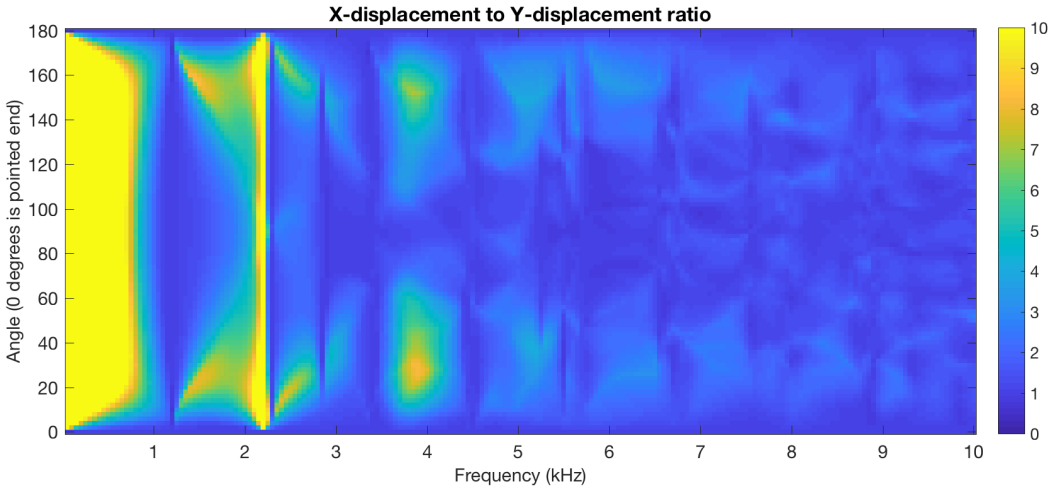


Figure 29: A measure on the aspect/frequency plane where points are represented in terms of their ratio of maximum x-displacement to maximum y-displacement. Pure bending modes will be high in this measure.

x-displacement to the maximum y-displacement, for every point in the frequency/aspect plane. Comparing this to Figure 17, we note that the points identified as bending in the survey also score high on this measure.

To identify other portions of the return which may be resulting from bending but which may not captured by this measure because other behaviors are occurring simultaneously which *do* result in y-displacement, we can constrain the physics simulated by the FE model to reduce such co-occuring behaviors. It is possible to simply eliminate movement in the y direction completely, but it should be noted that even pure bending produces some y movement as the object compresses at the bends; further, the specifics of acoustic/elastic coupling mean that other behaviors besides bending might be necessary in order to initiate the bending. Instead, we utilize a technique known as a *spring foundation* which allows us to apply a force to points in the elastic object which act against displacement, just as happens with a real spring. Unlike a real spring, however, this modeling technique can be applied to complete surfaces or volumes and with different forces in every direction. To isolate bending modes, we have applied a spring foundation in the y direction to all the points on the surface of the target within a central band of the cylindrical portion of the body. The elastic scattered return from this modified model are shown in the bottom of Figure 30, with the un-modified return shown on the top for comparison. We note that the survey observables identified as bending are present in this return, but other observables that were not previously identified as bending are now represented as well.

Direction of Displacement: Normal waves

Several of the classes of behavior identified in the surveys of Section 4.3 were characterized by movement primarily in the direction normal to the target's surface. Figure 19 identified the observables where the normal displacements took place in a regular grid, and Figures 21 and 25 also were dominated by normal movement. Again, we can develop a map of the degree to which each point in the frequency/azimuth plane has this characteristic by looking at the ratio of the maximum normal

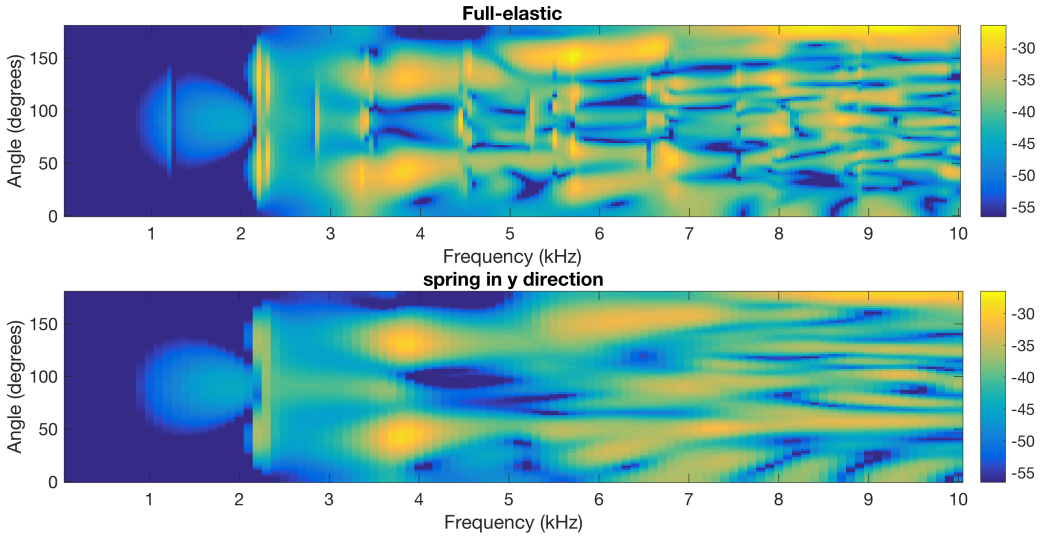


Figure 30: The upper figure shows the complete elastic model farfield response for the science target. The lower figure is the response for a model whose physics have been constrained by a y-direction spring to allow only the types of movement identified in Figure 29

displacement to the maximum tangential displacement. This map is shown in Figure 31, and a comparison with the respective survey point sets indicates that we are capturing this behavior in isolation. Modifying the physics modeling to isolate normal waves may be an area of future work. The spring foundation approach as taken for the bending modes described above has been less successful because the tangential displacement of the standing normal waves is more crucial to the behavior we are trying to isolate, and thus the simple spring foundation distorts the behavior under study even as it eliminates other behaviors.

Isolation by Modification of Material Stiffness

The discussions above have focused on the use of displacement directions as a method of characterizing and categorizing observables, but the internal physics can be studied in other ways as well. In Figures 14 and 15, for example, we have shown the shear strains for a particular point in the observable study. Shear strain is a measure of microscopic deformation that occurs when bending occurs across a particular region. TU shear strains are shear strains in the plane tangent to the object surface, and it can be observed that TU-strains get quite large at certain points in the geometry but stay very low along particular lines. This means that, while all the points along this line may be displacing in the TU plane, they are displacing very close to in tandem with their neighbors in the line so there is little microscopic deformation in that plane.

As above, we can create a measure of the degree to which each point in the frequency/azimuth plane exhibits similar behavior. In this case, we take a ratio of the maximum overall TU-strain to the maximum TU-strain for the subset of points along the center line. The map for this measure is shown in Figure 32, and gives us a new set of points where such strain behavior is isolated in the true return. To identify the observables where this behavior is co-located with other behaviors, we can modify the stiffness matrix of the modeled object. The stiffness matrix defines the relationship

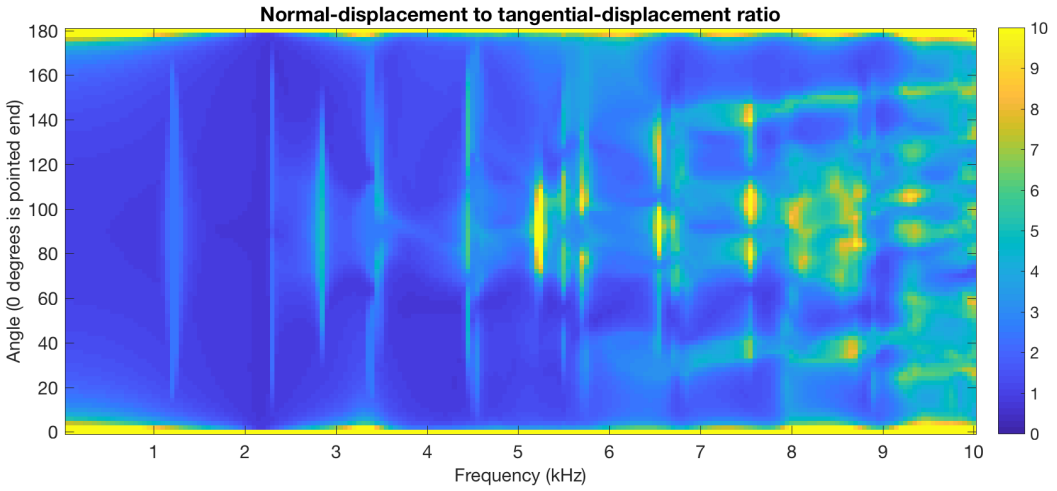


Figure 31: A measure on the aspect/frequency plane where points are represented in terms of their ratio of maximum normal-displacement to maximum tangential-displacement

between stresses applied to microunits of a material and their resulting deformation (strain), and is controllable from within the FE framework. By appropriate modifications of the stiffness matrix along the regions where we have observed low TU shear strain in some observables, we can create a model were even high stresses will produce little or no shear strain so the points will move in tandem. The observables isolated by this physics modification are shown in Figure 33, again with the unmodified-physics elastic response shown above for comparison.

Other Measures of Elastic Behavior

There are, of course, many other measures possible which might be used to isolate classes of observables and guide the developments of feature sets. Figure 34 shows some of the simplest, which are maps of the maximum displacements in the x, y, and surface-normal directions. One advantage of looking at these types of maps over looking at acoustic color plots of farifield returns is that the question of what behaviors are potentially excitable in an elastic object can be separated from what is actually received given a particular receiver location. Downward-looking sonars with receiver arrays might be better served by feature sets which search for signs of a particular internal behavior rather than trying to create and match templates which represent the exact combination of internal behavior observables which will be seen at the assumed target/sonar/sediment geometry under interrogation.

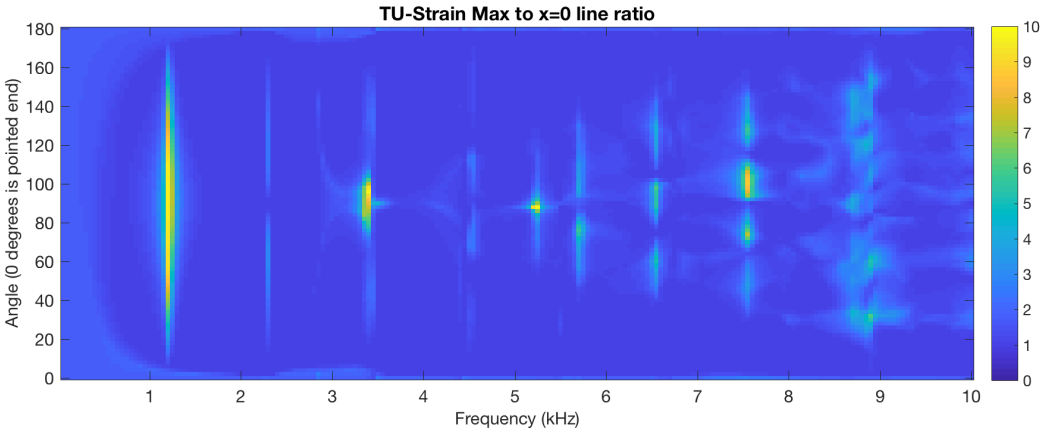


Figure 32: A measure on the aspect/frequency plane where points are represented in terms of their ratio of maximum tangential-plane shear stress overall to maximum tangential-plane shear stress along a particular line (see text and Figure 15)

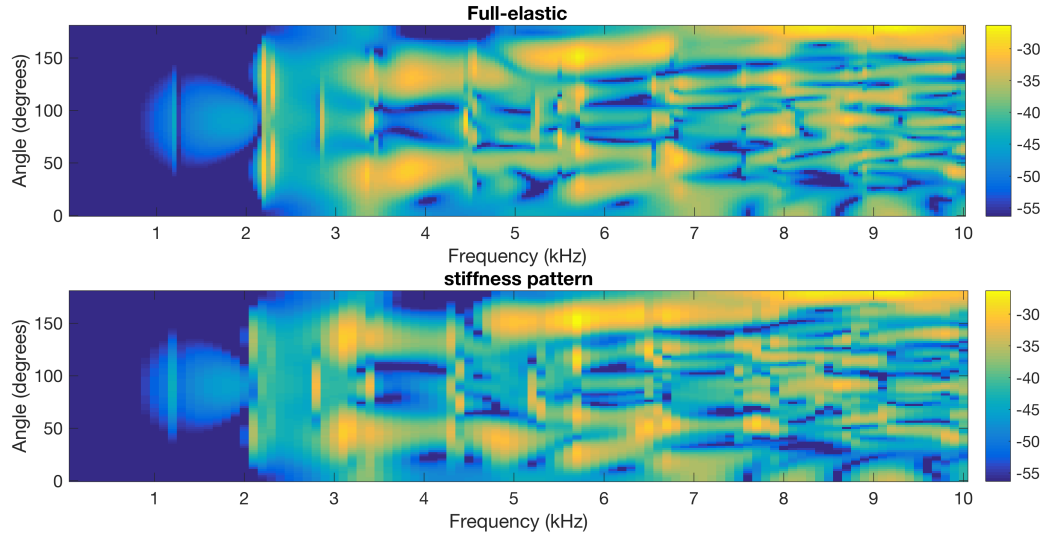
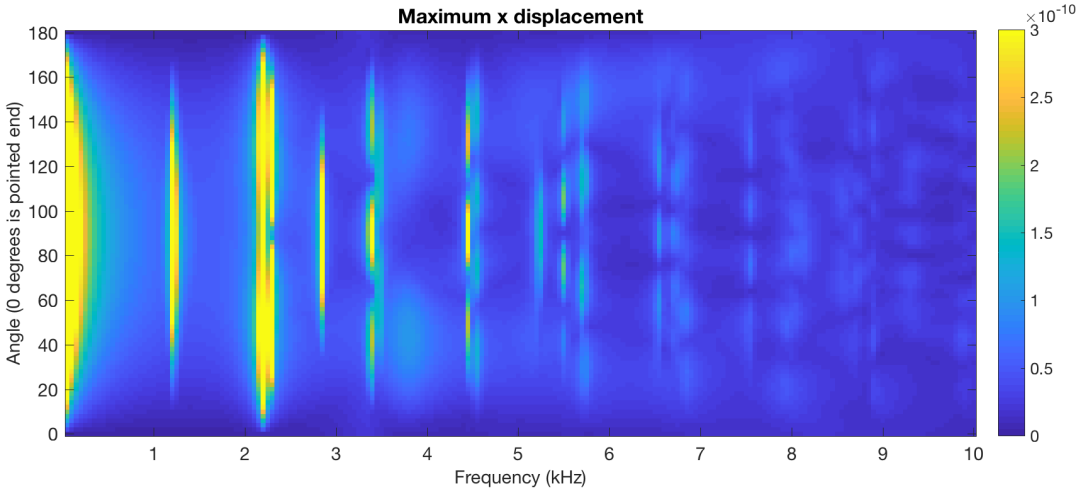
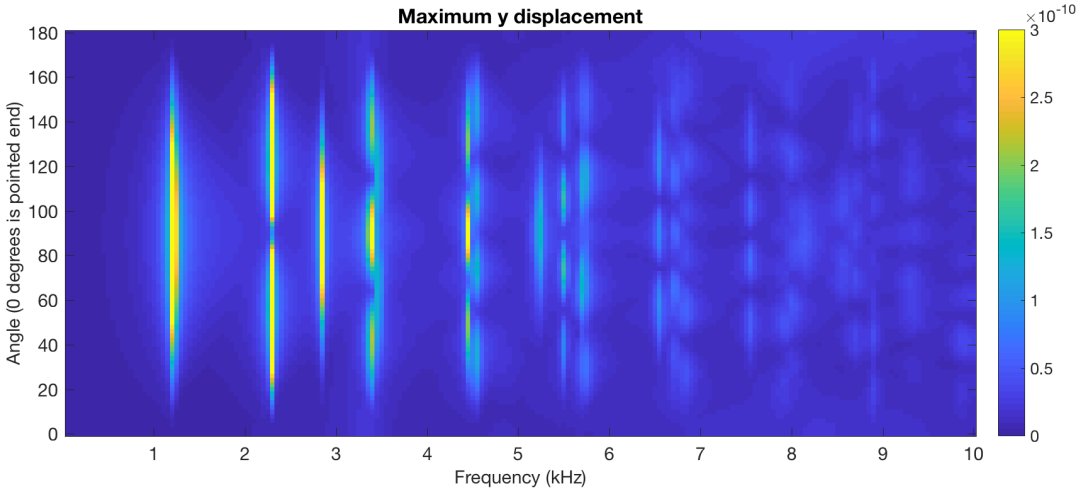


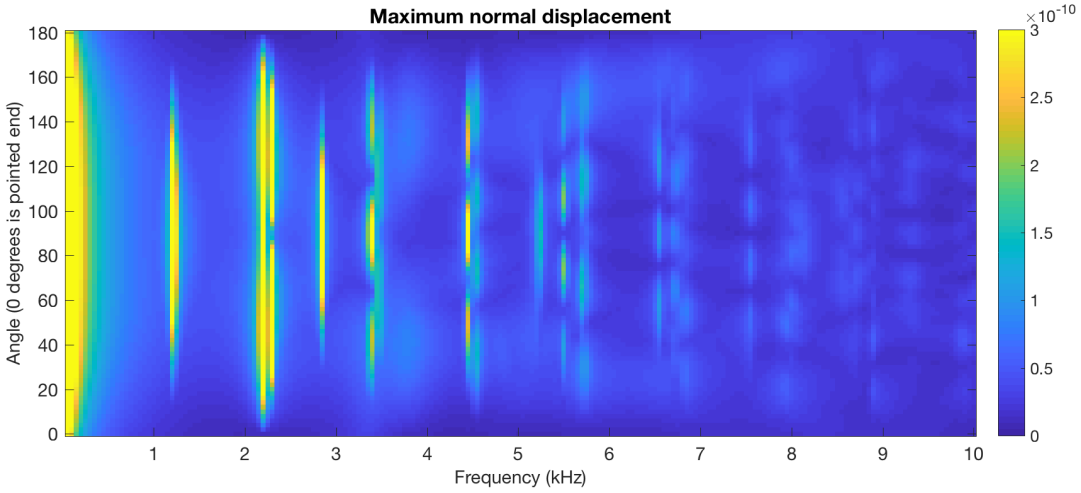
Figure 33: The upper figure shows the complete elastic model farfield response for the science target. The lower figure is the response for a model whose physics have been modified by adjusting the shear stiffness along an axis where certain classes of observables have no tangential shear strain.



(a) Maximum x-displacement



(b) Maximum y-displacement



(c) Maximum displacement in direction normal to surface

Figure 34: Measures on the aspect/frequency plane where points are represented in terms of their maximum displacement in different directions

5 Evaluation of Clutterex Data

The science target whose manufacture was described in Section 3 was interrogated in multiple bottom configurations and with two different filler materials during the ClutterEx17 data collection. The three main goals for this task were:

1. Allow for comparison to enable validation of the FE model
2. Evaluate response variability under nominally equal sediment/burial conditions
3. Add to the database of available responses for future classification system testing

The work done toward model validation involved updates to the model geometry and material properties to lessen model/data deviations seen in the original model. The manufacturing process described in Section 3 contained several potential sources of deviation, including:

1. The ideal geometry as defined in the model may not have been achievable in the casting process, particularly at corners and points
2. The material properties (density, sound speed) of mild steel have variability depending on the specific formula used at the fabrication facility
3. The tolerances in the casting process, particularly as relates to the alignment of the inner mold with the outer mold, may have caused deviations to the axisymmetry of the fielded target
4. The assembly process involved gaskets and screws, which were not modeled in the axisymmetric approach.
5. The surface smoothness of the manufactured target is a function of the casting process and may differ from the model, especially at the regions of mold connection or material flow.

We focused on investigating #1 and #2 in this list; #3 and #4 are outside the scope of what could be modeled in the time-frame of this project. The effects of #5 (surface smoothness) were deemed to be within the range of what could be expected from corrosion or biological fouling as would be expected in an ocean environment; results from Bucaro (MR-2500) [7] have attempted to characterize the effects of such variation.

We began the investigation by cutting one of the manufactured targets in half, as shown in Figure 36a. This enabled careful measurements of shell thickness at points all along the body, as well as measurement of curvature at the corners and other shape information (as in Figure 36b). Based on this, we concluded that the modeled thickness was too large (possibly due to shrinkage in the cooling process) and that there was some deviation due to misalignment of the inner mold with the outer mold but that this was small enough to be optimistic that its effects would be minor. We also made minor adjustments to some of the other model geometry. The measurement results are shown in Figure 35.

In the next phase of investigation, we cut samples out of the target half and machined them into rectangular plates for materials testing (see Figures 36c and 36d). We measured the density (and adjusted the model accordingly) and sent the sample plates to Dr. Jermaine Kennedy (NSWC-PCD) for materials testing. His results indicated deviations from the nominal sound speeds for mild steel, so we further adjusted the model in response.

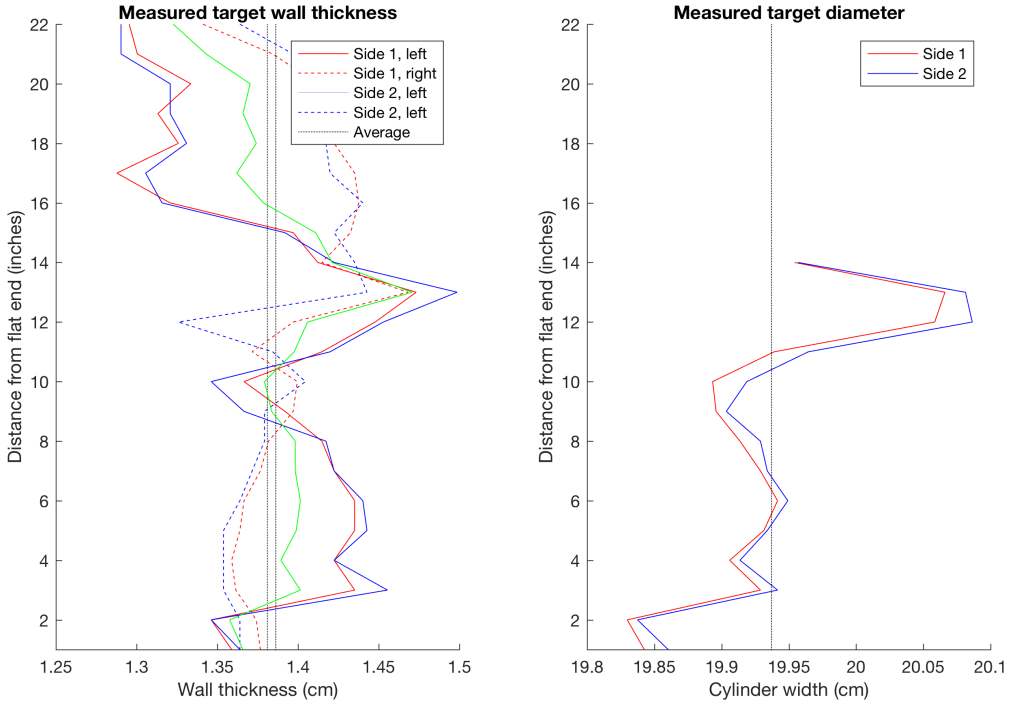
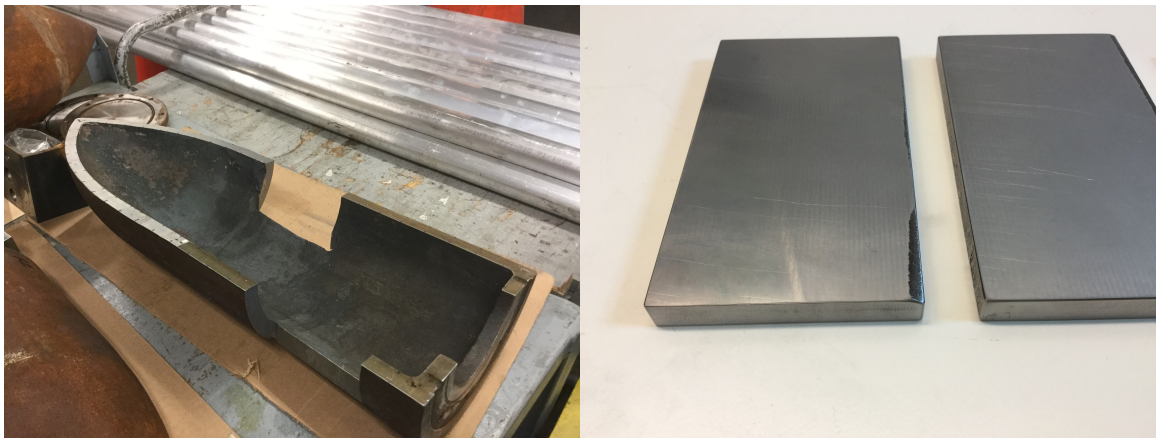


Figure 35: Measurement of the wall thickness and diameter of the cut target (both halves were measured to ensure consistency). The left figure shows wall thickness at different points along the target axis, and indicates a 1.5 mm drift in the alignment of the inner mold with the outer mold. The right figure shows target diameter, which indicates a 2 mm drift along the axis. Based on these results, the model parameters were updated to the average wall thickness of 1.38 cm.



(a) Target cut in half

(b) Detail of geometry for comparison with model



(c) Samples cut out for materials testing

(d) Partially machined samples for speed and density testing

Figure 36: Deconstruction of target for materials testing and geometry validation

The resulting models were much more agreement with the collected data up to 15 kHz, the range being studied in this investigation. Figure 38 shows a model/comparison between collected ClutterEx17 data for proud targets on a smooth seafloor with two different fillers (air and water); we only show frequencies up to 10kHz as this is the range for the previously-described observables study[4]. We note that when examining the model, much of the structure present in the model observables can be seen in the collected data, and that the differences between the two filler types is also reflected in both the model and data. We also note that significant differences in the relative strengths of the observables as well as their exact positions in frequency/azimuth space can be observed. This is an indication that pixel-by-pixel image-matching approaches such as cross-correlation would be challenging and further strengthens the motivation for physics-based feature extraction techniques which can search across this range of variability.

A more quantitative means of comparison is via a binary acoustic color (AC) plot and subsequent integral, and preliminary use of these data products in a UXO classification problem has shown promise under a separately funded ONR project[5]. A threshold value is set, in this case either -10 dB (3-10kHz) or -15 dB (10-30 kHz), and any target strength in the acoustic color that is above this threshold is set to one, anything below becomes zero. Figure 39 illustrates the binary AC plots for the same data and model results of Figure 38. For the bottom row of plots in Figure 39, at each frequency all the 'ones' are summed up over the entire angle range. Although any sort of absolute target strength has been stripped out of these data products, and the subsequent comparisons are admittedly still inherently visual, they offer a means of comparison that is a step closer to being quantitative. As previously noted, upon visual inspection of Figure 38, the models for both air and water capture some of the observables in the data, particularly the structural modes near -50 and 50 deg in the 4-6 kHz region. The models also reflect similar differences between them as that observed in the data. The air-filled shell has more broad, smooth peaks (in frequency), while the water-filled shell has narrower, sharper structure, as reflected in the 1-D integral plots.

Further clues are offered by generating these plots over various angle ranges, which would be more in line with what would be acquired in a mow-the-lawn type of area survey. Figures 40 (air-filled) and 41 (water-filled) show the binary AC plots over 3 angle ranges; left column is near the pointed nose end, middle left column is near broadside and the middle right column is near the flat tail end. The top row shows the model results, the middle row shows the data results. The bottom row is the sum over the angles of each region as a function of frequency, essentially the same plots as Figure 39, but just summed over a smaller number of angles. The right-most column is a sum over the frequency for each of the angle regions, or 'Views,' and gives a quick 1-D representation of how the energy is distributed across the angle space. There are a number of observations that can be made, but one that particularly stands out is that the maximum distribution of energy seems to be focused around broadside, as evidenced by the long-dashed lines (VIEW#2) in the right column of both Figures 40 and 41. This discovery will be explored in a follow-on effort as part of a task to develop an angle estimator.

Our efforts throughout the project were focused on the development of physical interpretation techniques, and thus we chose to focus on a smaller frequency band (2-10 kHz) to decrease the

[4]The *in situ* simulation results were created using a Hybrid FE/Propagation technique, in which the free field scattered pressures and derivatives are calculated on a discrete set of points very close to the target surface, and subsequently used in the discrete form of the Helmholtz Integral [8],[9]. This Helmholtz Integral, when used with the appropriate Green's function for a two-layered medium (water column-over-sand bottom), contains all of the physics for how the sound scatters from an object in contact (or embedded in) the layered medium.

[5]The binary acoustic color plots and 1-D integrals were developed by K. Williams under ONR Contract N00014-12-G-0078/N00014-16-F-3002

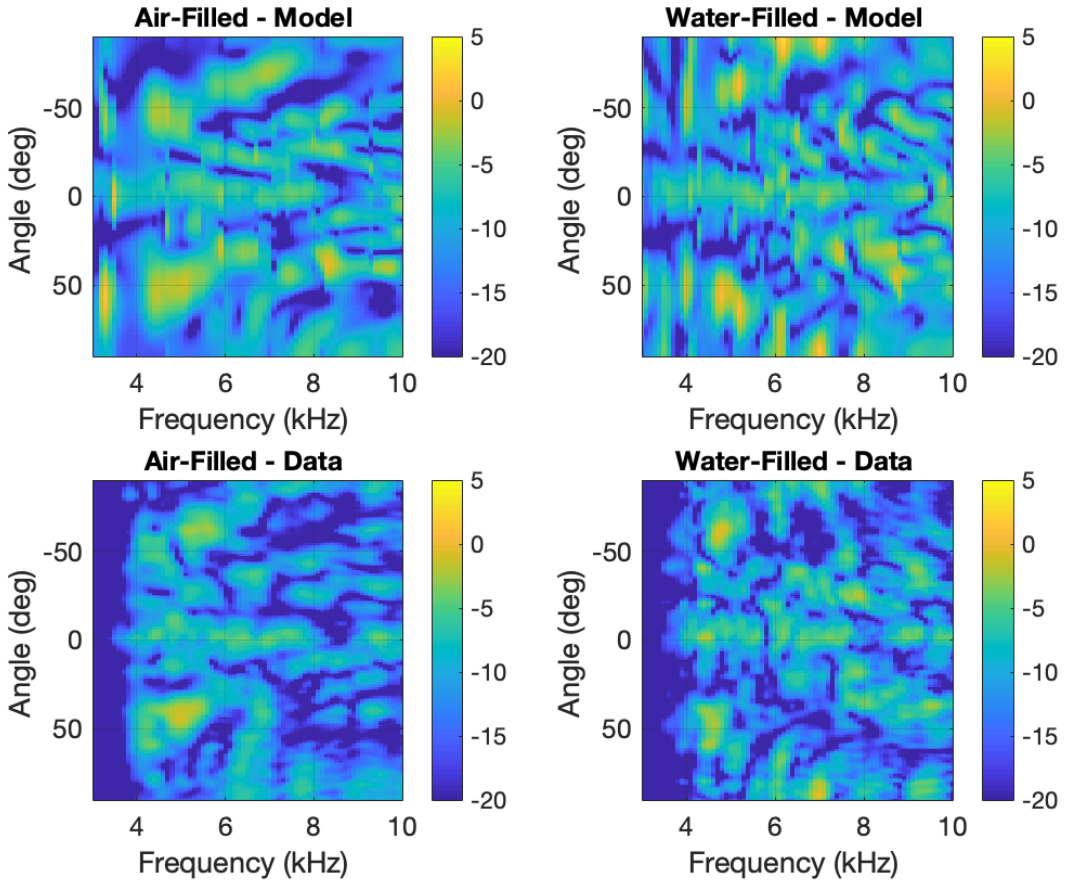


Figure 37: Binary acoustic color plots for the model (top row) and data (bottom row) for the air-filled shell (left column) and water-filled shell (right column).

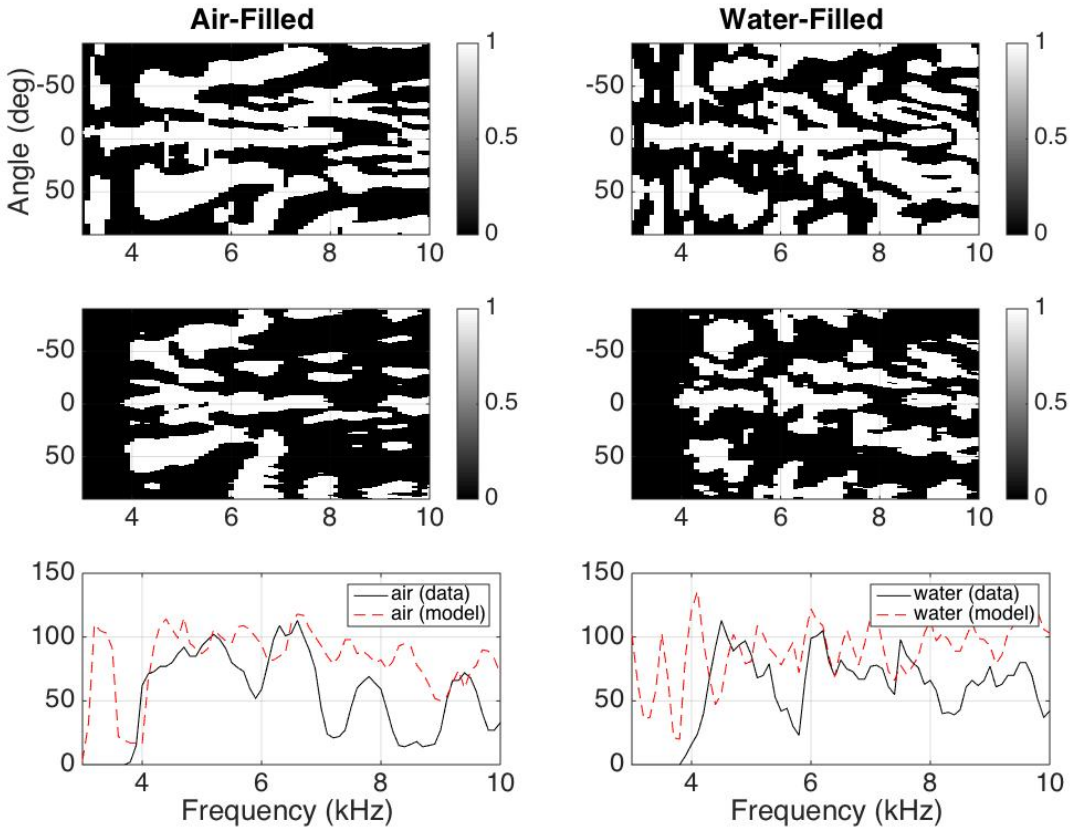


Figure 38: Binary acoustic color plots for the model (top row) and data (bottom row) for the air-filled shell (left column) and water-filled shell (right column).

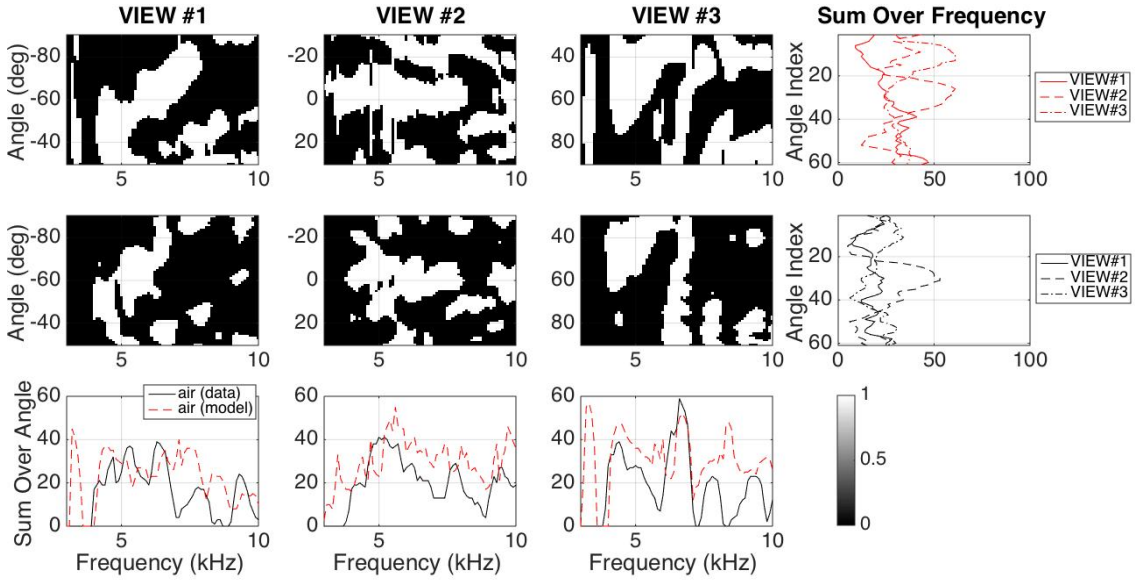


Figure 39: A comparison of the model (top row) and data (middle row) over different ranges of angles, or 'VIEWS', for the air-filled shell.

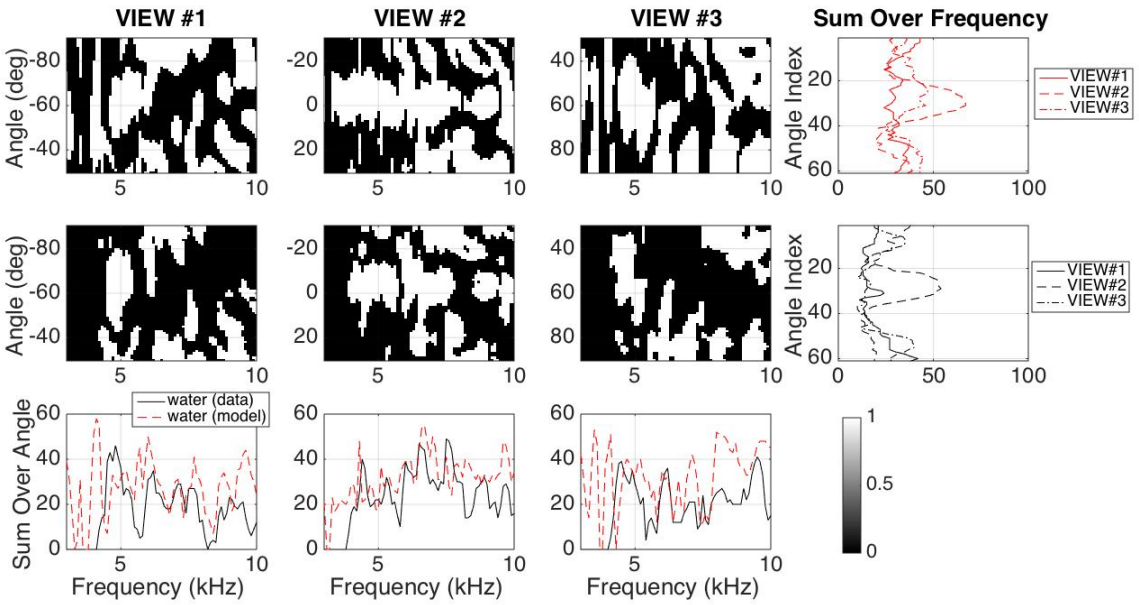


Figure 40: A comparison of the model (top row) and data (middle row) over different ranges of angles, or 'VIEWS', for the water-filled shell.

scope of the problem. That being said, we recognize that fielded UXO remediation systems may operate over a wider band of frequencies, and in fact, the ClutterEx17 data went up to 30kHz. As such, although an in-depth investigation of the data/model agreement in the 10-30 kHz band was out of the scope of this effort, we noted that the data and model agreement was less acceptable in this range. We therefore opted to briefly investigate two potential sources of this disagreement, the first being the material attenuation properties. As the frequency of the incoming sound becomes higher, the attenuation in the steel material becomes non-negligible. The attenuation is highly dependent on the texture of the material, particularly the form, size and percentage of the graphite particles present in the material [10]. The difficulty, however, is that these structures are altered widely by the conditions of the casting and heat-treatment. Measuring the attenuation of our samples experimentally was out of the scope of the test performed by Dr. Kennedy at NSWC-PCD. However, guided by the measured sound speed and density, and the results and discussion of Kamigaki [10], we were able to estimate a value for the attenuation. The inclusion of an attenuation parameter brought the model into closer agreement with the measured result from ClutterEx17, particularly in the 10-30 kHz region. However, the primary source of variation is still believed to be due to the environment, specifically the texture and composition of the ocean bottom, which is discussed in the paragraphs that follow.

As a second stage of analysis, we looked at the potential variability of the returns within a class of seafloor conditions that would be assumed to be identical under modeling approaches such as TIER [11]. As described in Section 3 and shown in Figure 2, multiple interrogations were done on the same target with artificial deviations added to the seafloor to simulate the seafloor variability created by water movement. Figures 42 through 45 show the AC and binary AC for the air-filled and water-filled shells over the 3-10 kHz range, while Figures 46 through 49 show the 10-30 kHz range. Although a more quantitative analysis via the 1-D integral plots was not conducted, there are still a number of observations that can be made via the binary AC's. First, the environmental variability over the 3-10 kHz region in Figure 43 (air-filled) and Figure 45 (water-filled) is minimal.

In contrast, the variation over the higher frequency 10-30 kHz region in Figure 47 (air-filled) and Figure 49 (water-filled) is much larger, particularly near the nose-end of both shells (-90 to -45 deg). Combining the physical interpretation techniques and tools described in Section 4, and the various data representations and comparative measures described in this section (binary AC and 1-D integrals) will be explored in more detail in our upcoming follow-on effort.

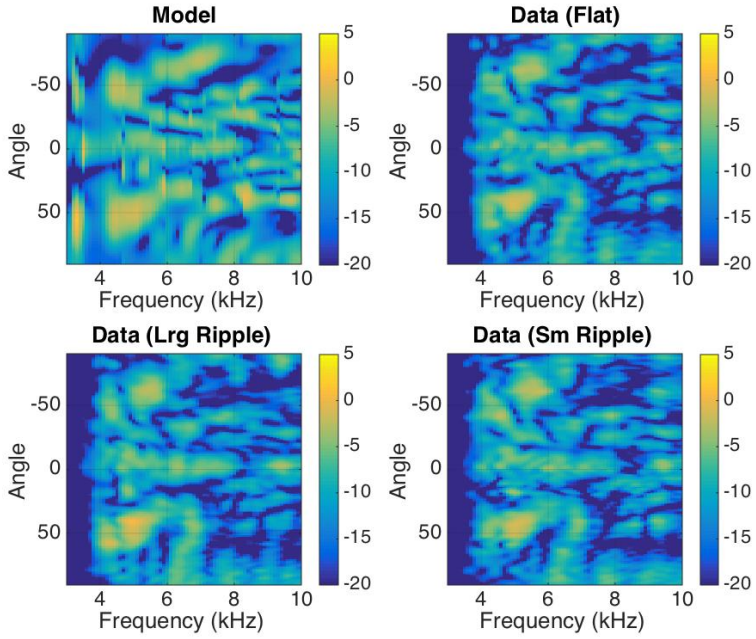


Figure 41: Acoustic color for the air-filled shell for 3-10 kHz band, showing the model for a flat bottom (top-left), data for a flat bottom (top-right), data for a large rippled bottom (bottom-left) and data for a raked bottom (bottom-right).

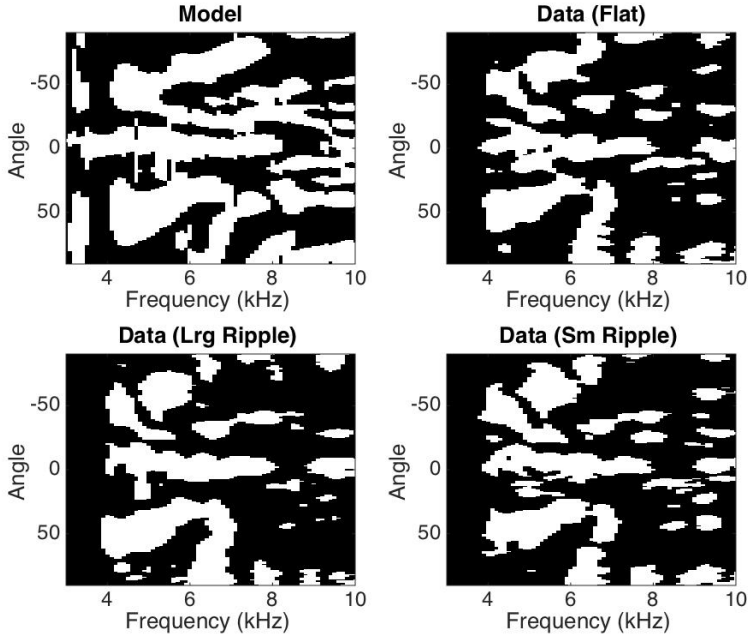


Figure 42: Binary AC plots for the air-filled shell for 3-10 kHz band, showing the model for a flat bottom (top-left), data for a flat bottom (top-right), data for a large rippled bottom (bottom-left) and data for a raked bottom (bottom-right). The threshold value is set to -10 dB.

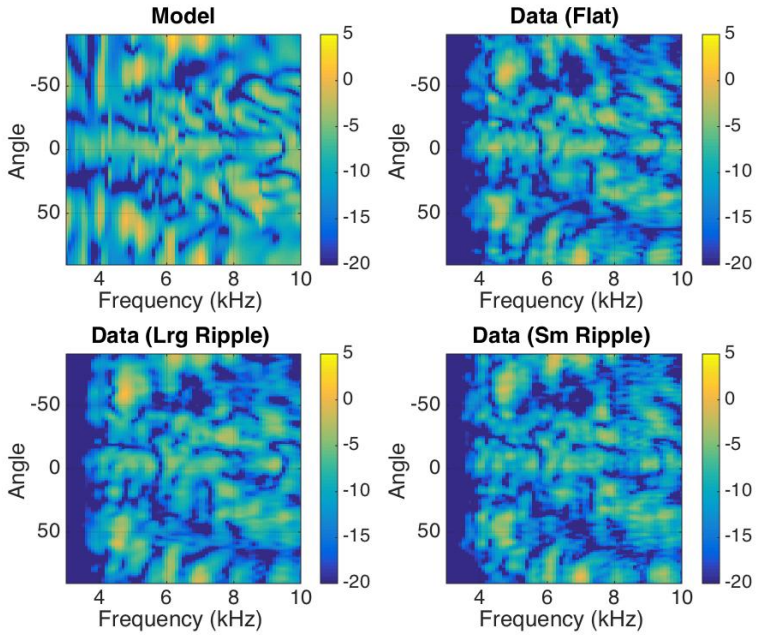


Figure 43: Acoustic color for the water-filled shell for 3-10 kHz band, showing the model for a flat bottom (top-left), data for a flat bottom (top-right), data for a large rippled bottom (bottom-left) and data for a raked bottom (bottom-right).

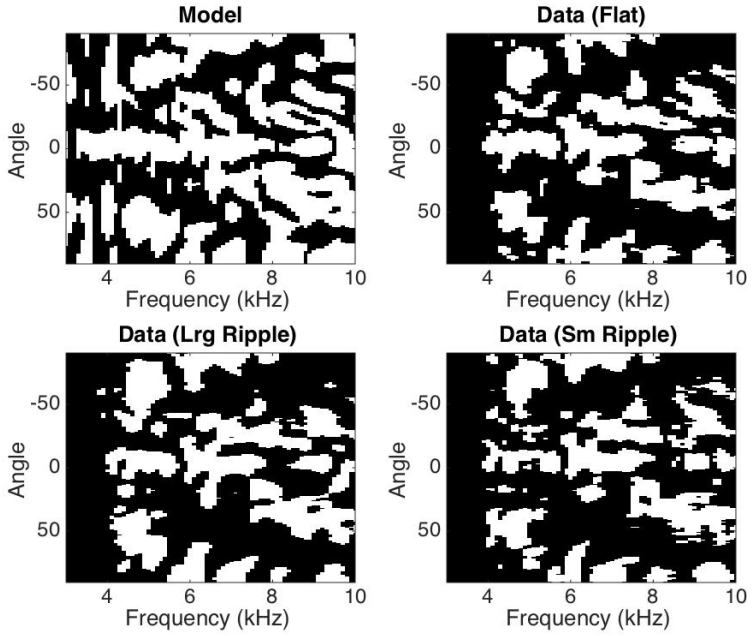


Figure 44: Binary AC plots for the water-filled shell for 3-10 kHz band, showing the model for a flat bottom (top-left), data for a flat bottom (top-right), data for a large rippled bottom (bottom-left) and data for a raked bottom (bottom-right). The threshold value is set to -10 dB.

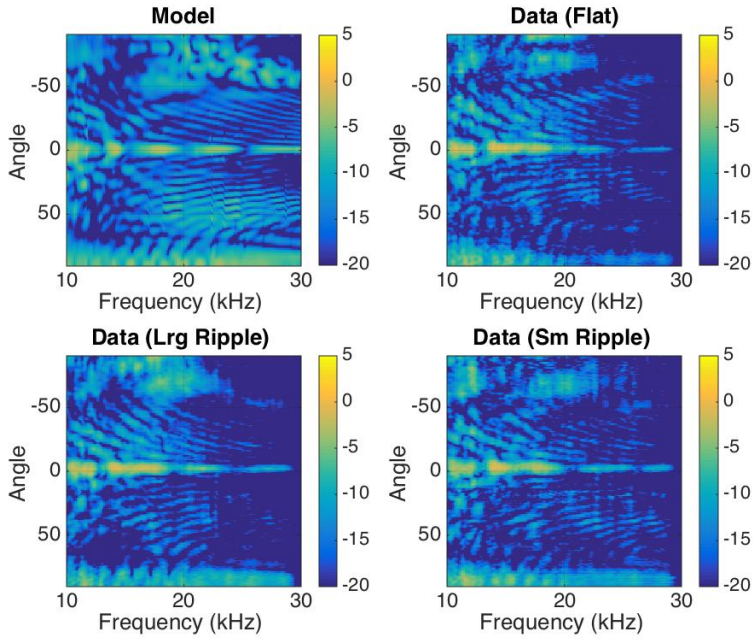


Figure 45: Acoustic color for the air-filled shell for 10-30 kHz band, showing the model for a flat bottom (top-left), data for a flat bottom (top-right), data for a large rippled bottom (bottom-left) and data for a raked bottom (bottom-right).

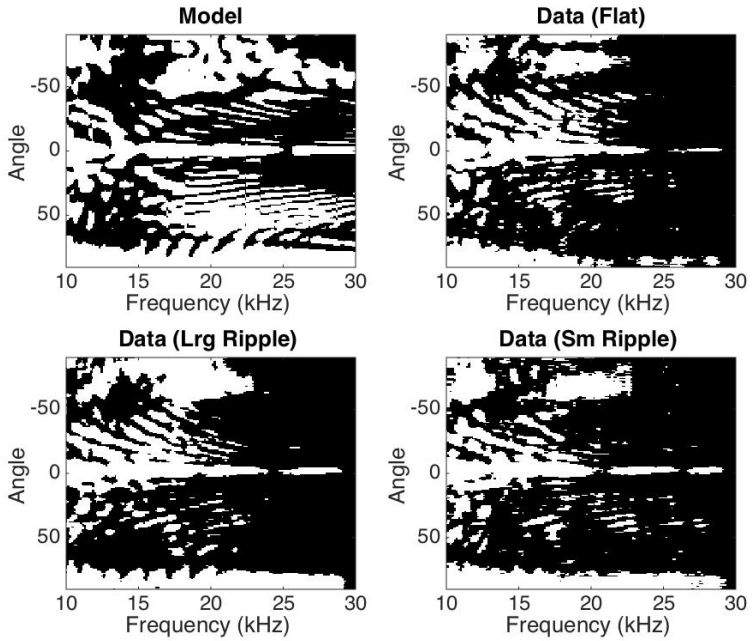


Figure 46: Binary AC plots for the air-filled shell for 10-30 kHz band, showing the model for a flat bottom (top-left), data for a flat bottom (top-right), data for a large rippled bottom (bottom-left) and data for a raked bottom (bottom-right). The threshold value is set to -15 dB.

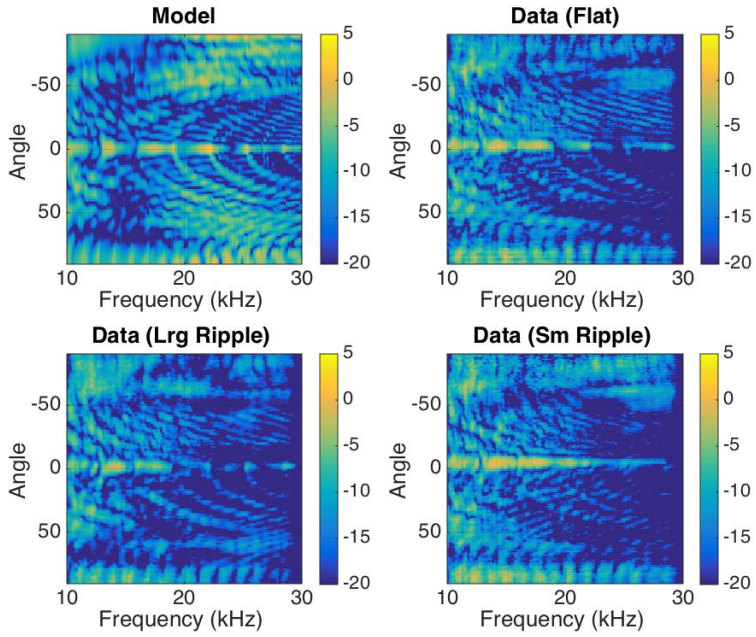


Figure 47: Acoustic color for the water-filled shell for 10-30 kHz band, showing the model for a flat bottom (top-left), data for a flat bottom (top-right), data for a large rippled bottom (bottom-left) and data for a raked bottom (bottom-right).

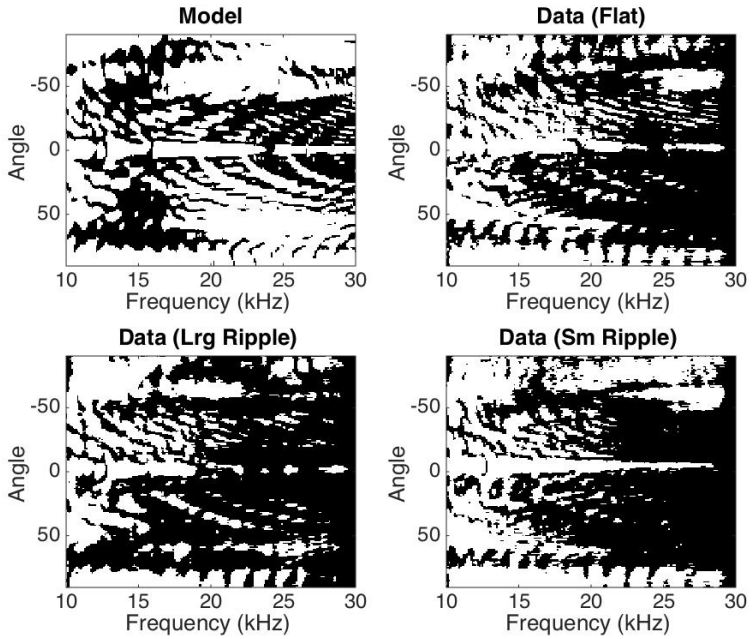


Figure 48: Binary AC plots for the water-filled shell for 10-30 kHz band, showing the model for a flat bottom (top-left), data for a flat bottom (top-right), data for a large rippled bottom (bottom-left) and data for a raked bottom (bottom-right). The threshold value is set to -15 dB.

6 Classification System Analysis

The second major approach we have taken in identifying potentially valuable physically-motivated features was to investigate an existing UXO classification system and a modified version of that system that incorporated physics knowledge in a rudimentary fashion. The goal of this approach is not to improve the existing system but to focus on the specific cases in which the modification improved performance, as a guide to creating feature sets that focused more directly on the sources of that improvement.

6.1 Field-Collected Data Classification System Analysis

A study at APL-UW done by Plotnick and Marston developed a classification system with a goal of utilizing knowledge of scattering physics into a data processing methodology [12]. We have attempted to use some of their results to find hints as to what specific acoustic features an automatic classifier might use to differentiate targets of interest from clutter objects, and to draw insights as to the physical mechanisms of these features, with the eventual goal of using these types of features in a more sophisticated, and hopefully more effective, classification procedure. As development of machine learning algorithms continues to advance, this type of analysis whereby the algorithms themselves can aid humans in an understanding and development loop is becoming more common [13].

Overview of the Classification Experiment

We focus here on one of the results from Plotnick [12], where physics-motivated procedures were used in development of a target classification approach. The data set used was comprised of acoustic color data (backscattered intensity vs. aspect and frequency) covering 360 degrees of aspect over a variety of targets. The data set was assembled from several different at-sea experimental campaigns (NSWC-PCD pond experiments, BayEx, and TREx), and was augmented with simulated data provided by TIER. As well as full coverage in aspect, the data set also included targets at various ranges from the source/receiver combination (in order to have grazing angle diversity), and had buried targets for the steepest grazing angle cases. There were 90 total target instances (combination of target type and range/burial state) in the data set, each with 360 degree aspect coverage.

The classifier structure used in this work is the relevance vector machine (RVM) [14]. The RVM structure itself can potentially lead to physical insights for the classification problem, since the RVM formulation provides a sparse set of basis functions (the relevance vectors) that can appear to be "prototypical" examples of the target or non-target class. This is opposed to the support vector machine (SVM) formulation, where the basis functions are typically more suited to defining the boundary between classes.

The main "feature" fed to the classifier for the set of experiments we are interested in is the result of a cross correlation of acoustic color templates across aspect angle. The experiment was run assuming a range of possible swath widths (from 40 through 140 degrees). For a real sonar system that was running linear tracks through a target field, this could be seen as having a transmitter with horizontal beamwidth the value of the swath width, with squinted receive beams being formed at the various aspects on a target as the sonar transits across a target view.

To train a classifier for a particular swath width, the 360 degree acoustic color data products were broken into segments the length of the swath value under test, with segments overlapping by 20 degrees. For example, for the case of 50 degree swaths, segments were generated covering 0-50 degrees, 20-70 degrees, 40-90 degrees, etc., with 18 segments generated for each target instance. Each of these segments was then considered a data "examplar" to be used in a classifier training/testing procedure. The use of 90 total target instances and 18 swaths per target resulted in data sets of 1620 examplars. A repeated k-folds cross-validation scheme was utilized in the training and testing. The value of k here was five, so that the 1620 exemplars were randonly divided into five sets, with training utilizing four of the sets and testing on the fifth set. This procedure was then repeated such that each set of the five was held out for testing, resulting in five separately trained and tested classifiers. In addition, the k-folds data separation was repeated ten times, resulting in a total of 50 separately trained classifiers.

These classification experiments were run as a function of the various swath widths and for various levels of additive noise. Having access to 50 different realizations for each case also allowed an estimation of mean performance with confidence intervals for each case. The study of Plotnick and Marston judged performance by using area under the ROC curve (AUC) to measure performance. In general, the performance AUC tended to increase as swath width increased, which is likely explained by the increase in information content with increased swath width. Better performance was achieved at lower noise levels, which is to be expected.

Of particular interest to our study was a further breakdown of the data and the resulting separate training of RVM-based classifiers. Using the hypothesis that early-time signal content is dominated by specular (or "shape") reflected energy, and that late-time signal content is dominated by elastic scattering contributions, the data sets were separated at the time series level into nominal shape and elastic components. These data were then converted to frequency-angle space (acoustic color). Figure 50 shows an example of this type of separation, with time series data vs. aspect in the lower panels, and the corresponding frequency domain data vs. aspect angle in the upper panels. Using the scheme described above, three separate sets of classifiers were trained: *full* (before any separation in the time domain), *shape*, and *elastic*. The repeated k-folds scheme was performed on each of the three data sets, and a fourth ensemble *vote* classifier that combined the scores of the other three was also generated.

The results of the analysis described above show that the vote classifier tended to consistently outperform the other three (see Figure 51 and accompanying description; full details are in Plotnick [12]). This implies that there is informational value that a machine learning algorithm can use to separate target classes in the nominal shape-elastic separation methodology, a physically motivated procedure. In the next section we attempt to draw insights from cases where performance improvement is achieved, with the eventual goal of more precisely identifying phenomena in the backscattered signals that can be tied to physical scattering mechanisms specific to the target types we are searching for, such that this information can be leveraged by an automatic classification scheme.

Classifier Experiment Analysis

In the classification experiment described above, there were data from 27 physical object types (at various ranges/grazing angles and burial conditions). For the classifier training and testing, the target class was composed of three of these objects: aluminum and steel UXO replicas, and an

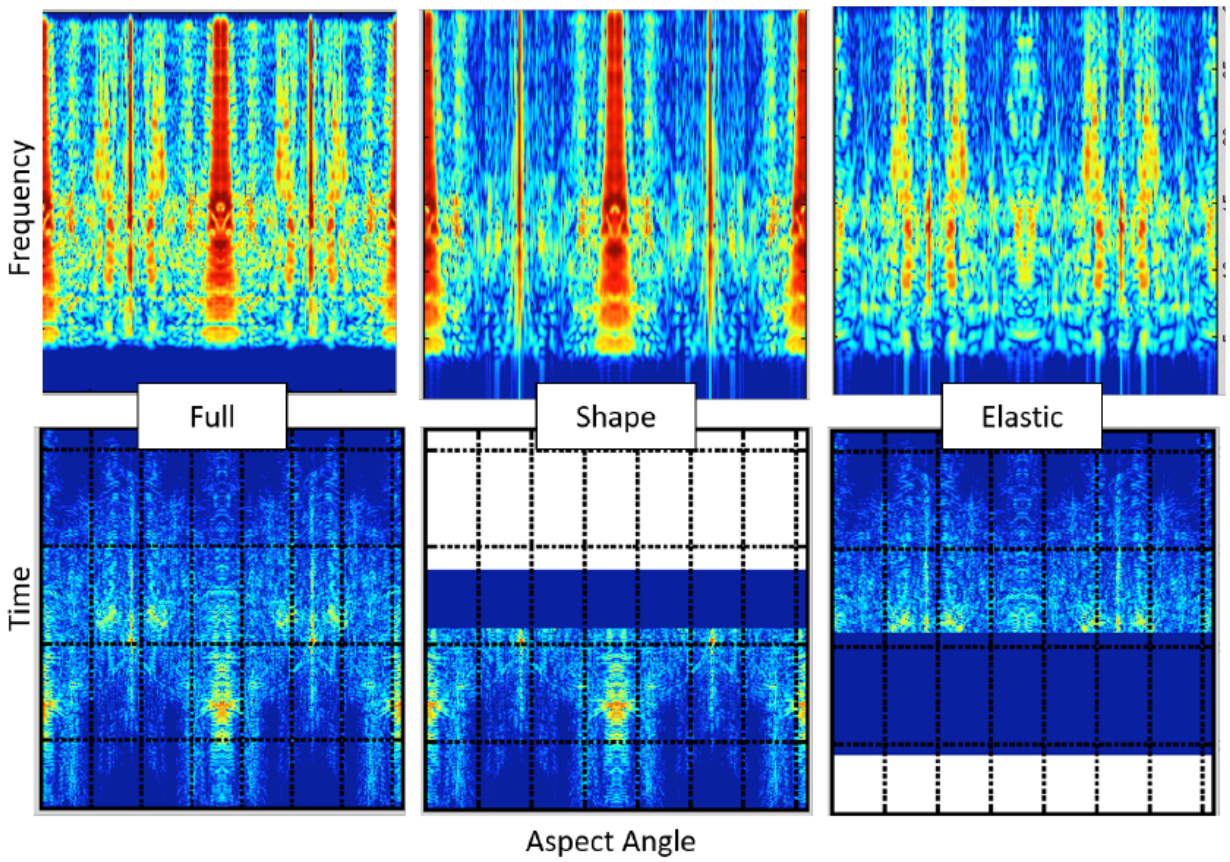


Figure 49: Example of separation of full time series data vs. target aspect into early-time, relating primarily to shape scattering, and late-time, relating primarily to elastic target response, in the lower panels. Corresponding frequency-aspect data are displayed in the upper panels. (Courtesy Dr. Daniel Plotnick and Dr. Kevin Williams)

"original material" UXO. The remainder of the objects were considered the clutter class. We have undertaken some analysis of the classifier results for one of these target types, the aluminum UXO replica.

We were given access to the full training and testing results of these experiments by Plotnick and Williams [12], and have used them under the current SERDP contract to perform the following analyses. This includes the scores of the classifier outputs for each testing result, allowing us to recreate receiver operating characteristic (ROC) curves for each classifier that was trained in the k-folds scheme. Figure 51 shows an example set of these ROC curves for one of the k-folds partitions. In this case there was no additive noise, and the swath width was 50 degrees.

One hypothesis motivating our examination of the internal physics and resulting scattered field of UXO targets is that there is information in the elastic response that could be useful for classification purposes if taken advantage of in appropriate processing strategies. The bulk separation in time into shape-dominated and elastic-dominated regimes is an initial methodology for achieving this. The data described above allowed us to start looking at this possibility in both qualitative and quantitative ways.

We chose a specific operating point, in this case a false positive rate equal to 0.1, and then searched

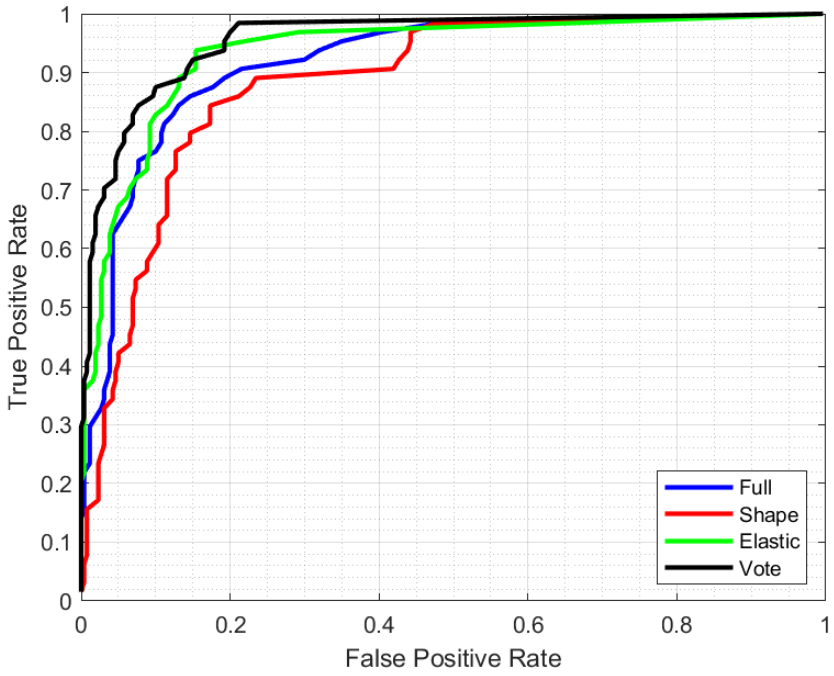


Figure 50: Example ROC curves for one of the *k*-folds partitions.

through all 50 partitioned classifier results for specific cases where the full classifier failed to correctly classify our target of interest (the aluminum UXO replica), but the elastic classifier achieved a correct classification. One way to categorize these results is to sort them by the central angle of the 50 degree swaths, in order to see if certain ranges of aspect angles provide better information for classification in the target elastic responses. A histogram of the results of this sorting can be seen in Figure 52.

In Figure 52, 90 and 270 degree angles represent broadside aspects to the aluminum UXO replica (i.e., directly viewing the longer dimension of the axially symmetric shape), with 180 degrees indicating a nose-on aspect, and 0/360 degrees indicating an aspect looking directly at the larger diameter endcap. The asymmetry of the histogram about 180 degrees is an indication of the variability in field data, as there is no reason to expect different results when looking at the long dimension of an axially symmetric object from opposite sides.

One noticeable trend in the histogram is that there are significantly fewer results where elastic-only processing improves the classification performance near the nose-on aspect. The overall response at this aspect tends to have lower backscattered intensity in general, so this is likely a region of overall poor classification performance.

We now look at specific examples where the elastic-only classifier provides improved classification results compared to its corresponding classifier for the full response. Figures 53 through 55 show comparisons for a swath center of 125 degrees, and for cases of targets at varying range from the sonar system (15, 30, and 40 m, respectively). All color scales are for the same dynamic range spanning 30 dB. In each figure, the full response is in the left panel, shape response in the center panel, and elastic response in the right panel. Here, range is essentially a surrogate for grazing angle with the sea bottom. The nature of the responses can be highly dependent on the grazing

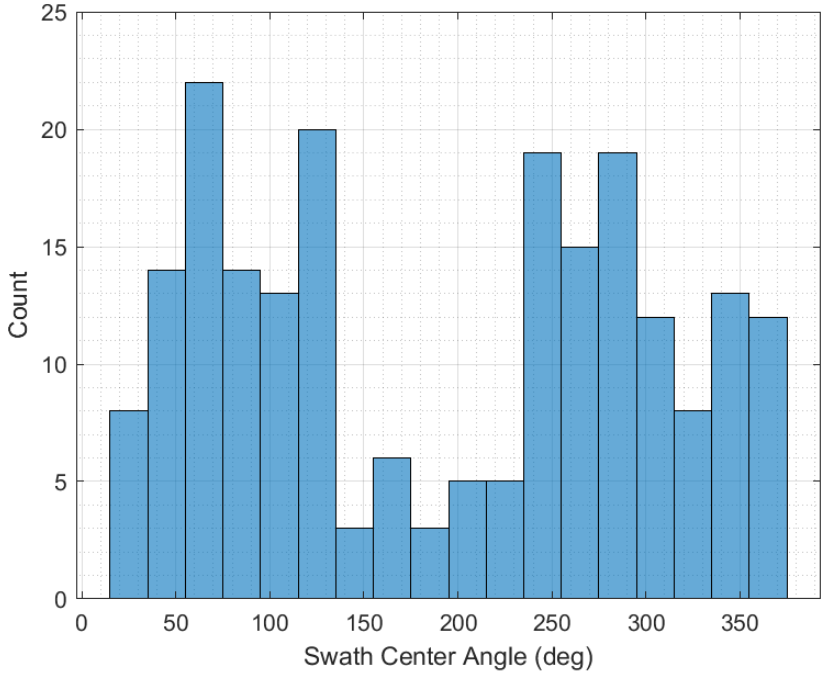


Figure 51: Histogram of swath centers where full classifier fails but elastic classifier succeeds.

angle, since the total incident and reflected fields are coherent combinations of the various paths that travel to and from the target (i.e., direct sonar to target, sonar to seabed target, etc.).

Several interesting insights can be made from this collection of figures. One should keep in mind that the separation into shape and elastic components was performed in the time domain, so that combination of the separated acoustic color plots, in order to arrive back at the full result, would actually need to be done with the underlying complex values and performed coherently in order to be strictly correct. However, looking for areas where energy exists (or doesn't exist) in this intensity domain is a reasonable approximation to the separation process. Also note this swath interval does not contain the direct broadside aspect around 90 degrees, where we typically observe the highest levels of backscattered energy.

Since we know for these cases that the elastic classifier outperformed the full classifier, it is instructive to examine the comparisons between those two acoustic color (AC) products. Of note is that in all three grazing angle cases, there is significant signal energy near 130 degree aspect and near the center of the frequency band in the elastic AC. This energy is also present in the full response, but the fact that shape energy has been removed in the elastic AC indicates that the shape energy may be common across both the target and clutter classes, and the elastic energy that remains is an indicator of our targets of interest.

Figures 56 through 58 show some different examples, this time for a central swath angle of 285 degrees. In this case, the swath does contain the broadside response (270 degrees). Here we see some similarities to the previous example, with regions of consistent energy across grazing angles in the elastic-only response that become more prominent when the shape response is removed, both at broadside aspect (270 degrees) and off-broadside.

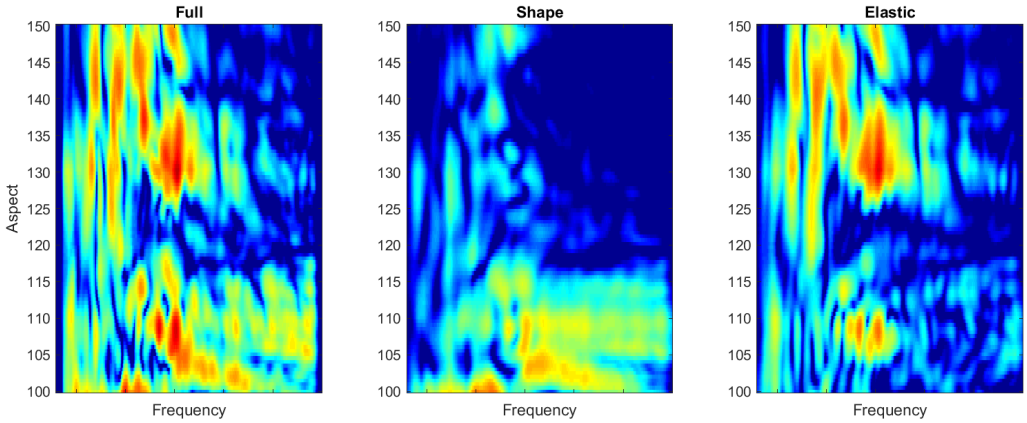


Figure 52: Full, shape, and elastic acoustic color plots: swath center at 125 deg., range = 15 m.

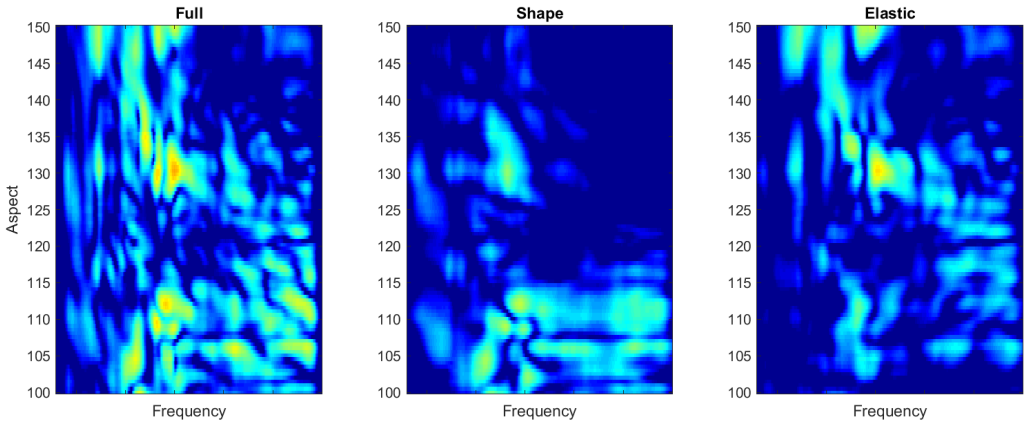


Figure 53: Full, shape, and elastic acoustic color plots: swath center at 125 deg., range = 30 m.

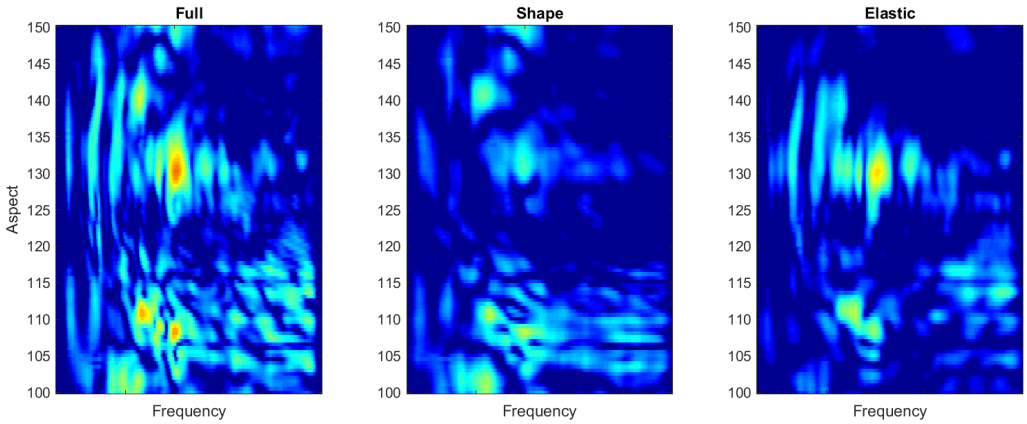


Figure 54: Full, shape, and elastic acoustic color plots: swath center at 125 deg., range = 40 m.

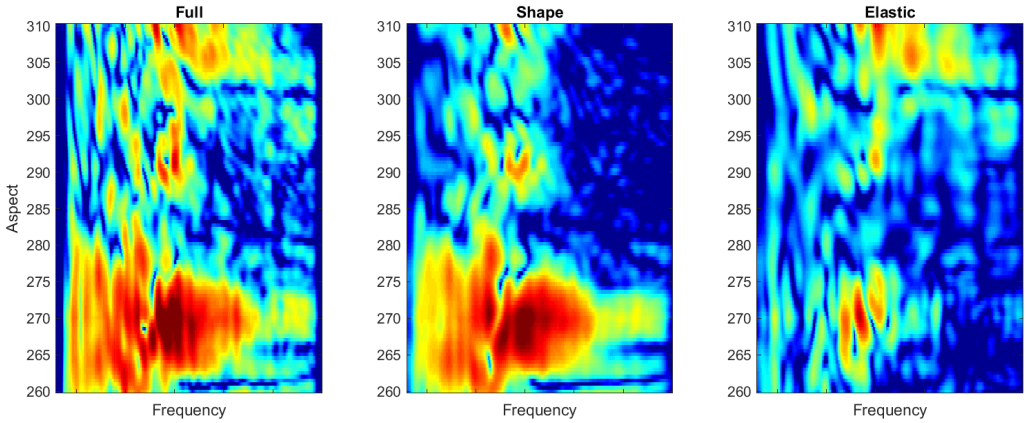


Figure 55: Full, shape, and elastic acoustic color plots: swath center at 285 deg., range = 15 m.

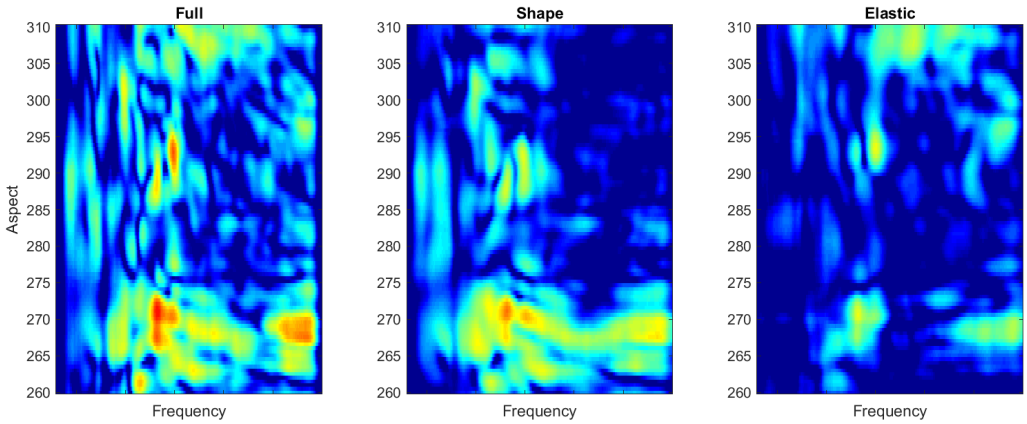


Figure 56: Full, shape, and elastic acoustic color plots: swath center at 285 deg., range = 30 m.

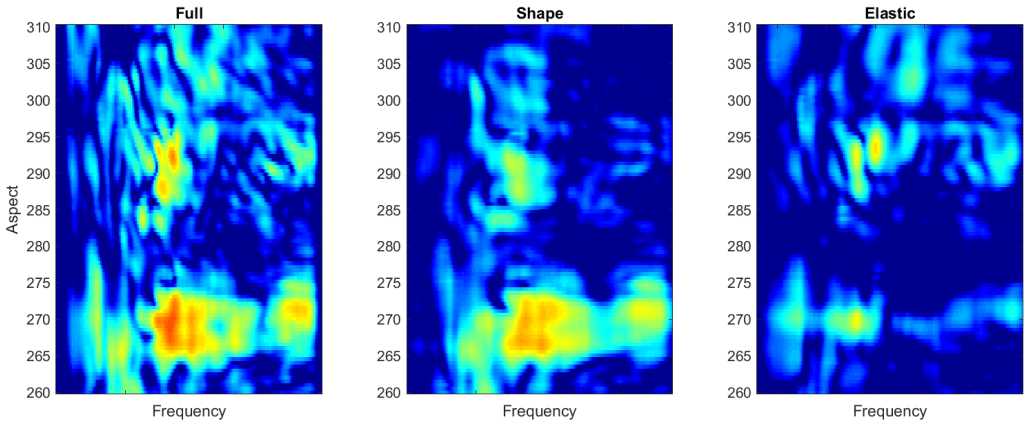


Figure 57: Full, shape, and elastic acoustic color plots: swath center at 285 deg., range = 40 m.

Relating to Internal Physics

This section describes an example of how target internal physics can be explored and tied to measured acoustic phenomena which are indicative of driving improved classification performance. Figure 59 shows the same target as in Figure 53, but over an extended aspect range. The left panel shows the full response, and the center panel shows the elastic response, both from measured data. The rightmost panel shows the output of the TIER model for this target and measurement geometry (15 m range). The most direct comparison can be made between the two full responses (lefmost and rightmost panels). While the modeled and collected data exhibit meaningful differences on a pixel-by-pixel basis, the qualitative match between observables in the modeled data and those in the field-collected data are sufficient to assume that similar internal physics are responsible for the energy seen in the acoustic color plane. This is quite evident in the aspect angles covering roughly 120-150 degrees. This also is a region that contains significant elastic-only energy. This is particularly noticeable in the rightmost panel of Figure 53, where the classifier covering 50 degree swaths was improved by considering elastic-only energy. Using this as a cue, we can investigate the internal physics that may be responsible for this phenomenon. We will look at the point marked with a white circle in the rightmost panel, near 135 degrees in aspect and near the center of the frequency band.

Figure 60 shows an example of the target internal physics that can be examined. Here we have computed particle displacements in cylindrical coordinates (r, ϕ, and z going across the columns), with the descending rows showing real part, imaginary part, magnitude, and phase of the components, respectively. Likewise, Figure 61 shows the displacements for the same target and geometry in Cartesian coordinates (x, y, and z going across the columns, with similar breakdown going down the rows compared to Figure 60). The grid pattern of the standing waves observable in this feature is a validation of the value of the observable categories discussed in Section 4.3. One consideration in interpreting the internal physics of Figures 60 and 61 is that the presence of the seafloor means that the response is a function of two incident waves (unlike that in Section 4.3) and thus the internal physics is more complex. To aid in interpretation of movement patterns, we have developed an animation component of the GUI (stills from which are shown in Figure 62. The specific grid period can be tied in with the target geometry; further, the nature of this wave movement is such that it is affected by the filler material. Investigation of this specific interaction will continue in follow-on work.

Commentary on Application to a Down-Looking Sonar

The data collections that were used to feed the analysis described above were primarily assuming a side-looking sonar system, and accompanying low grazing angle geometries, covering potentially very wide swaths of aspect angle for any given target. As the SERDP/ESTCP program has evolved, the hardware system that will likely be put into practice appears to favor a more down-looking geometry, with an across-track swath width that will keep interrogation of targets and seabed above the critical angle of the sediment. This geometry has implications for acoustic color-based target classification schemes, and the methodologies used for side-looking sonar [12] may not be directly applicable to this new geometry.

However, the system under development also has plans for forming imagery-based data products, such that external shape and orientation of targets can be estimated. The analysis shown above indicates that there may be particular aspect angles for targets of interest where elastic-dominant

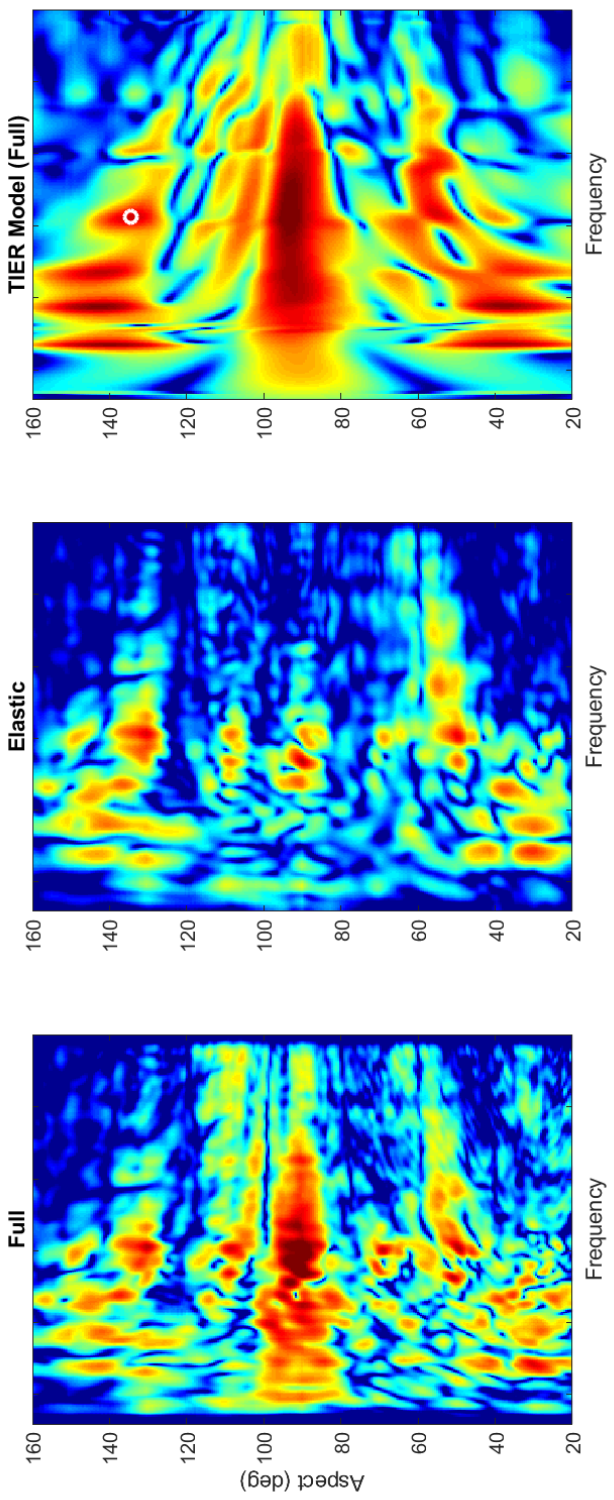


Figure 58: Expanded view of target response from Figure 53: Full and elastic from measured data, and TIER model result. The white circle in the TIER results indicates the point in frequency and aspect at which we will investigate the internal physics.

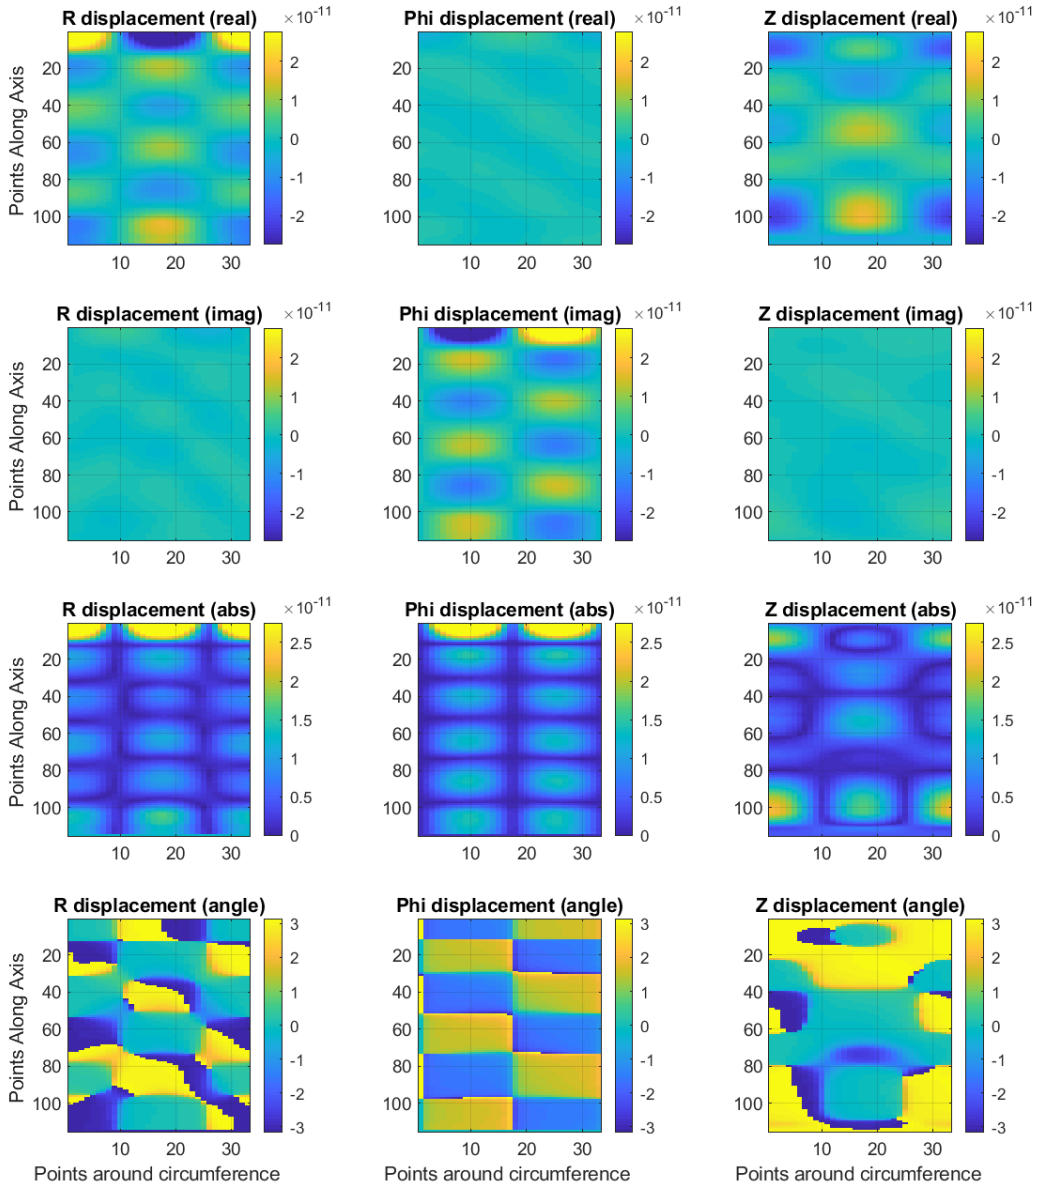


Figure 59: Internal physics (total displacement in cylindrical coordinates) for aluminum UXO replica at 15 m range. First column shows displacement in r dimension, second column in ϕ dimension, and third column in z. Descending rows show real, imaginary, magnitude, and phase of the components, respectively.

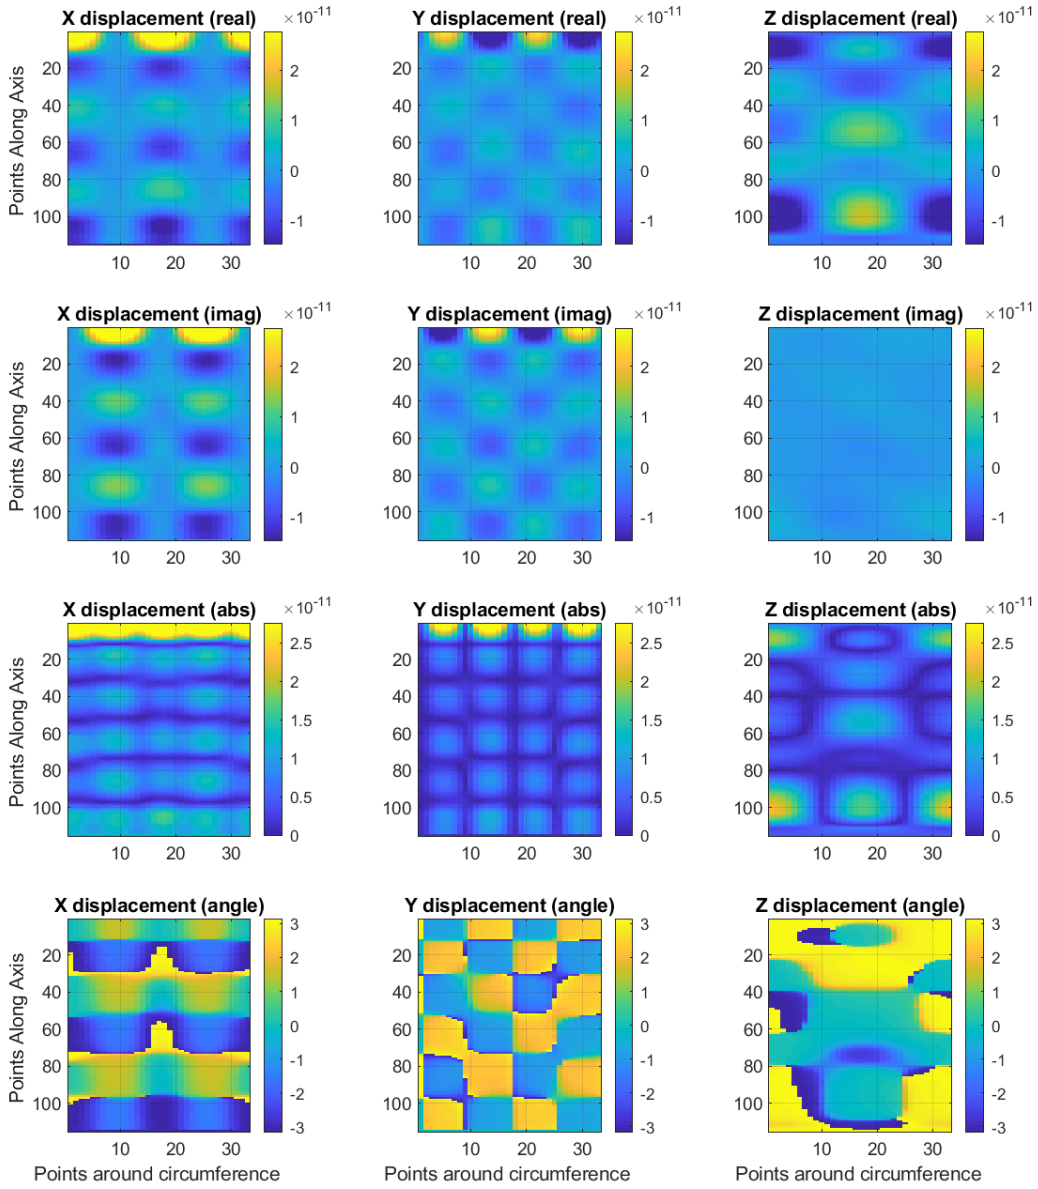


Figure 60: Internal physics (total displacement in Cartesian coordinates) for aluminum UXO replica at 15 m range. First column shows displacement in x dimension, second column in y dimension, and third column in z. Descending rows show real, imaginary, magnitude, and phase of the components, respectively.

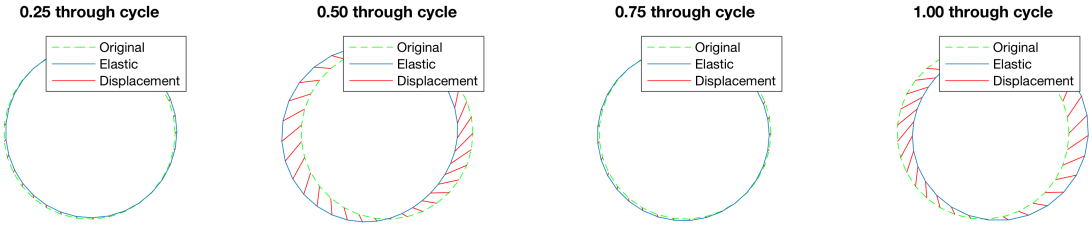


Figure 61: Stills from the GUI animation of one layer of the target, used to aid in visualising the physics of Figure 61

energy can more effectively classify them against non-targets. In the previously developed classifiers, the actual aspect angle at which a target was observed was treated as an unknown. If the imagery-based information can provide an estimate of target size and orientation, this could provide a very powerful tool in combination with the elastic effects seen above to improve classification performance, reduce false alarms, and thereby reduce time and cost for remediation efforts. Further implications of a down-looking sonar will be discussed in Section 7 of this report.

6.2 TIERSWAT Classification System Analysis

We performed ananalysis on a second existing system, one developed for the TIERSWAT[6] Under MR-2505 (PI: Kargl) Dr. Jack McLaughlin performed a classification study on the dataset, which simulated a linear synthetic aperture sonar (SAS) operation in which objects were interrogated across 60 degrees of azimuth in 1-degree pings, with 1-10 kHz frequency band for the transmitted signal [16]. The set of ping returns for a particular interrogated object were classified independently and the overall object label was calculated by combining the ping results via a weighted voting scheme. Both the target (UXO) and clutter object types are listed in Table 1.

The system developed by McLaughlin under MR-2505 was based on Gaussian mixture models (GMMs) [17], a classical method of creating a probability density function (PDF) model representing a distribution across some feature space. As the name implies, the overall distribution model is a sum of individual Gaussian (normal) distribution functions, each with their own mean, covariance, and weight. Figure 63 provides a simple illustration of this technique; the top subfigure shows two GMMs representing distributions along a one-dimensional feature space, with each GMM composed of two separate Gaussian components. In a two-class (e.g. target *vs.* clutter) scenario where both classes have representative training patterns available, GMMs can be used to create a classification system for an unknown pattern represented in the feature space by looking at the relative likelihoods that the two models produced the feature vector representing the unknown pattern. The lower subfigure of Figure 63 shows the relative likelihood that a pattern with a given feature value came from the target model, and the operator can place a threshold for alarming on this likelihood to achieve a desired detection/false-alarm ratio.

The system developed by McLaughlin under MR-2505 was conceptually similar to the simple example just described, except that the feature space was 81-dimensional (since the feature vectors were the 81 values of each ping return's acoustic frequency spectrum). Two models were created,

[6]TIERSWAT: a dataset produced using the TIER simulation target responses combined with environment-specific reverberation simulations produced using PC SWAT (the Personal Computer Shallow Water Acoustic Tool-Set)[15].

Target types	100-mm aluminum replica of UXO 100-mm steel replica of UXO 105-mm bullet-shape UXO, air-filled 105-mm bullet-shape UXO, water-filled 155-mm howitzer with no endcap, water-filled 155-mm howitzer with endcap, air-filled 155-mm howitzer with endcap, water-filled
Clutter types	55 gallon drum, one end opened scuba tank point scatter blobs (clutter) 2:1 aluminum pipe (3/8 inch wall thickness) 2:1 solid aluminum cylinder (2 ft long, 1 ft diameter) 3:1 solid aluminum cylinder (3 ft long, 1 ft diameter)

Table 1: The object types in the TIERSWAT simulated dataset, which consists of 7 target (UXO) types and six clutter types, with a variety of sediment types, orientations, and distances.

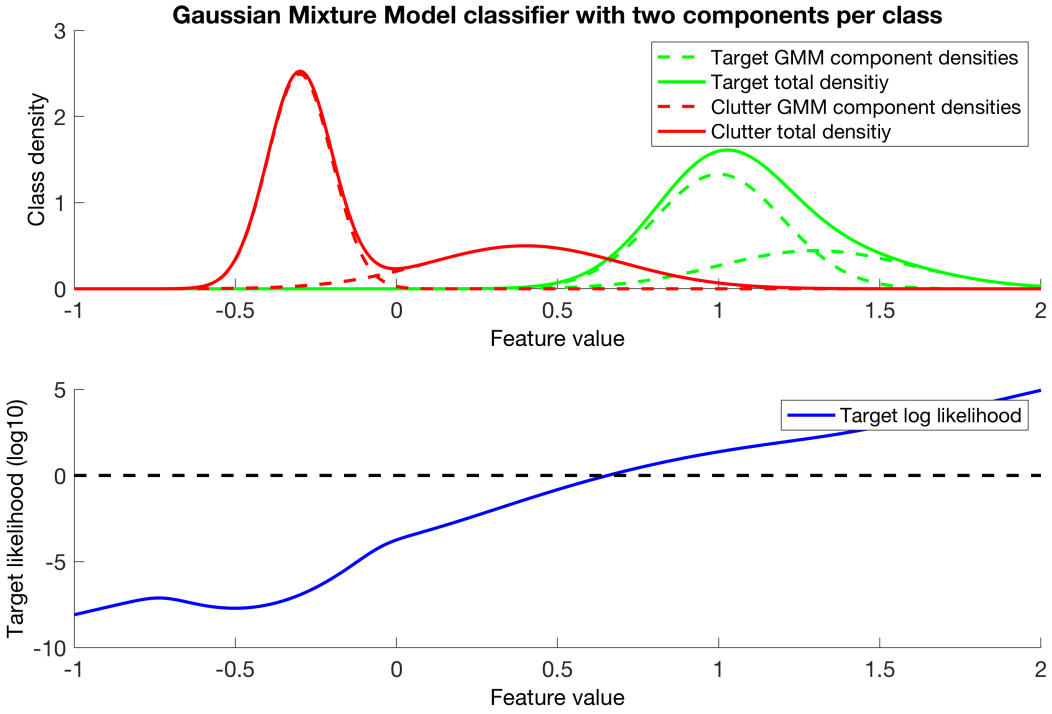


Figure 62: Illustration of a simple GMM-based classification system for a one-dimensional feature space. The top figure shows two GMMs (target and clutter), each represented as the weighted sum of two Gaussian components. An unknown-class pattern with a particular feature value can be classified by comparing the relative densities of the two models at that particular feature value; the ratio of the target-model density to the clutter-model density is shown in the bottom subfigure. Patterns with high target likelihood will be classified as targets, with the particular alarm threshold dependent on the desired detection/false-alarm requirements of the system.

one for the set of target objects and one for a set of clutter objects, with the models representing a probability density estimate of the ping return spectral vectors for each class as calculated using a Gaussian mixture model (GMM) with 15 Gaussian components used to represent the complete distribution of each class.

The scope of work under McLaughlin's MR-2505 study was simply to perform a study to investigate alternative architectures to achieve good automatic classification performance (including multi-class separation) could be achieved on this dataset; as such the classifier was treated as a "black box" and no attempt was made to establish what physics or spectral features were being leveraged to achieve that performance or how robust the performance could be expected to be under differing conditions. The goal of our study was to analyze the performance to understand the features being used and to infer whether the initial classification study performance could be expected to generalize to more realistic conditions.

GMM Component Usage: Orientation-Specific Modeling and Clutter Characterization

One approach to understanding the way the GMM-based system was doing classification is to represent the return from each ping in the test patterns using the index of the GMM component most likely to have produced that ping return's feature vector[7]. Figure 64 shows some results of this analysis for three target types: 155 mm Howitzer shells, 105 mm bullet-shape UXO, and 100 mm steel UXO. The figure shows which Gaussian component most closely matched each return in the test set swaths. It can be observed that the target model, which was explicitly trained only to model the distribution of the target ping return spectra as a whole, is in fact subdividing the problem by using different components to model the targets depending on the incident angle at which the the target is being interrogated. This result suggests that an approach that explicitly estimates orientation and chooses a classification method informed by that estimate could outperform a generic classification approach. As discussed in Section 7, this insight will be particularly critical for a down-looking sonar scenario in which the relationship between source position and incident angle at the target will vary much more greatly as a function of target orientation.

We further note that the GMM is subdividing the classification space by target type; i.e. for a given incident angle a different subset of the components is used to represent the 105 mm UXO than (for example) the 100 mm UXO. This suggests the possibility that models which focus on individual target types would be less complex and more robust than one model which attempted to combine the entirety of the "UXO" class.

A similar best-component representation of clutter ping returns can help to understand the limitations of the approach taken in MR-2505[8]. The classification algorithm being investigated in that analysis assumed a two-class problem in which models are created for both the target distribution and clutter distribution, and the classification decision was based on which model was a better fit. Architectures based on algorithms like SVMs or RVMs, as well as many deep-learning approaches, make similar assumptions. However, a risk of this approach is illustrated in Figure 65, which shows the fit of the clutter model components to several types of clutter objects. Similar to the case with the target models, we again see that the clutter model is subdividing its discrimination power by

[7]Recall that a GMM is a weighted combination of Gaussian distributions, so the likelihood that each distribution would have produced a particular feature vector is the density of that distribution at that particular point times the component's weight.

[8]In this case we are representing the clutter ping returns using components from the clutter GMM model, so the 15 components are not the same as those in the above UXO-class discussion.

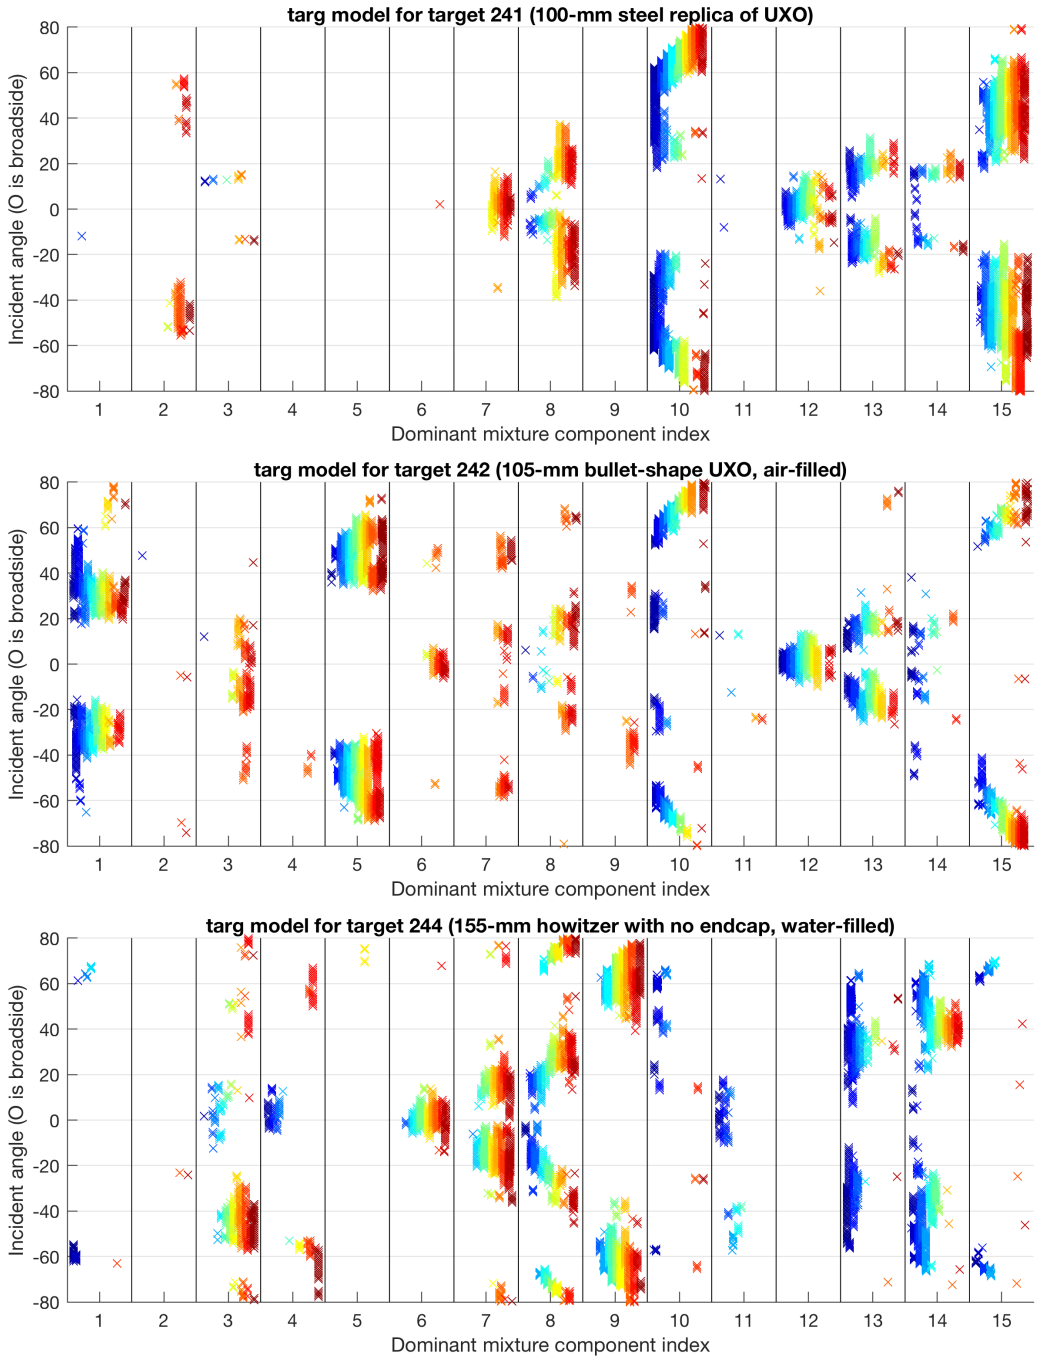


Figure 63: Usage of Gaussian Mixture Model components for three different target types. The model has 15 total Gaussians to represent the distribution of ping return spectra; each 'x' in the figure indicates the dominant component index for a particular ping return, with the y axis indicating the incident angle of the ping on the target. Different colors represent different surveys, with hotter colors indicating objects at increasing distance from the sonar. It can be observed that the model is effectively subdivided according to incident angle and target type, with distinct components being used to model particular target types at particular angles.

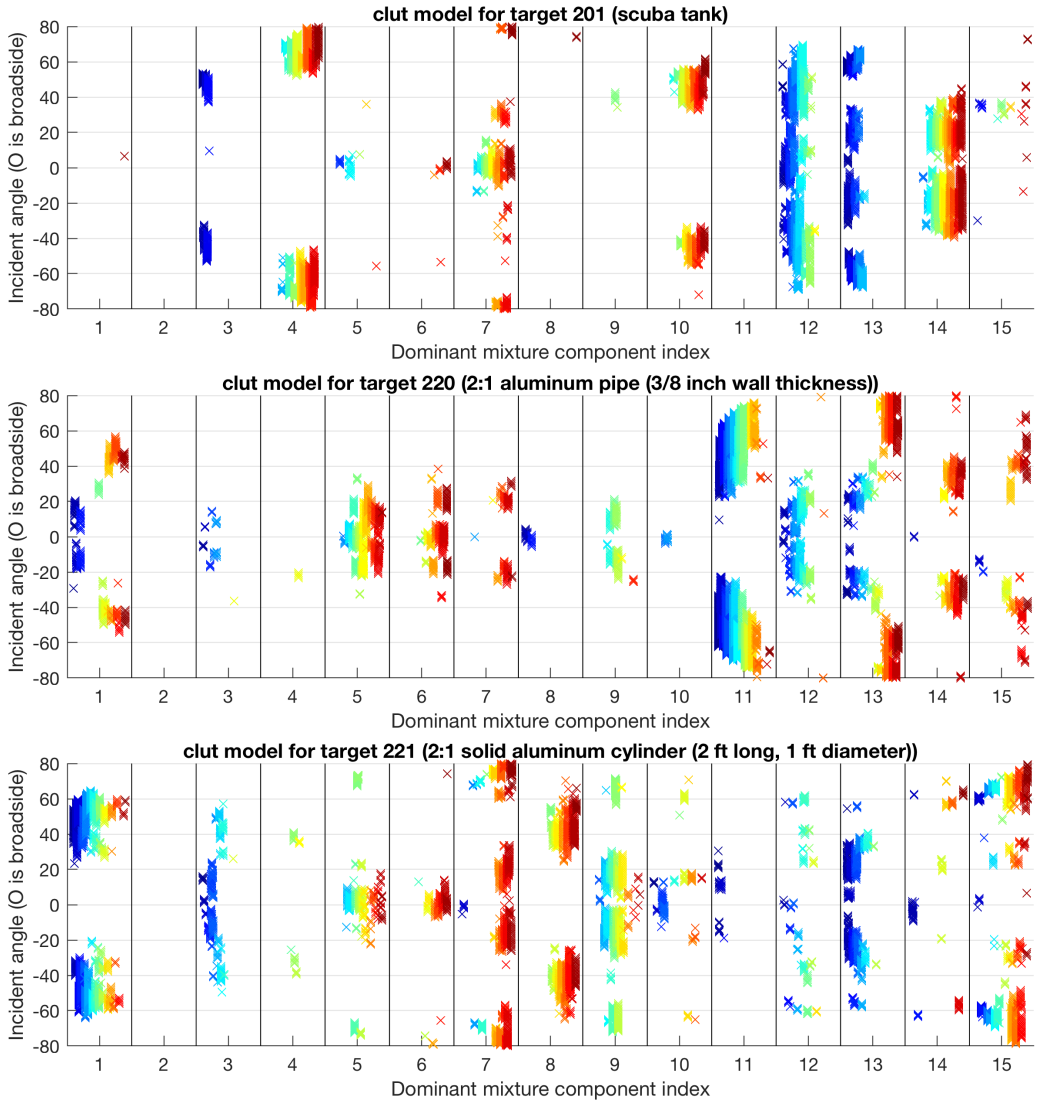


Figure 64: Usage of Gaussian Mixture Model components for three different clutter types. The model has 15 total Gaussians to represent the distribution of ping return spectra; each 'x' in the figure indicates the dominant component index for a particular return, with the y axis indicating the incident angle of the ping on the target. Different colors represent different surveys, with hotter colors indicating objects at increasing distance (decreasing grazing angle) from the sonar. As in the target models in Figure 64, it can be observed that the model is effectively subdivided according to incident angle and clutter type, with distinct components being used to model particular clutter types at particular angles. New clutter types are thus likely to be poorly represented by this model.

using particular components to specialize in identifying each clutter type. A reasonable inference from this result is that the feature set being used does not utilize some universal 'clutter features' but that each of the modeled clutter types is occupying a discrete part of the feature space. As a result, other clutter types not previously modeled could be expected to be a poor fit to the existing clutter model and thus not well-handled by the two-class paradigm as their relative fit between the two models is not a good measure of UXO likelihood.

We designed a simple modification to investigate the implications of the incomplete clutter-class characterization by adapting the MR-2505 classifier to consider only the fit of the data to the UXO model rather than the relative fit of the UXO and clutter models. Specifically, the MR-2505 system was a threshold on

$$\frac{p(\mathbf{v} = \mathbf{v_{test}}|k = \text{UXO})}{p(\mathbf{v} = \mathbf{v_{test}}|k = \text{UXO}) + p(\mathbf{v} = \mathbf{v_{test}}|k = \text{clutter})} \tag{5}$$

as illustrated in the blue curve of Figure 63, whereas the modified version of the classifier that did not assume training samples that were available for the clutter type of interest was a threshold on

$$p(\mathbf{v} = \mathbf{v_{test}}|k = \text{UXO}) \tag{6}$$

where, in both equations, $p()$ is a probability density function, $\mathbf{v_{test}}$ is the feature vector to be classified, and k is the specific class model (so $p(\mathbf{v} = \mathbf{v_{test}}|k = \text{UXO})$ would be the green curve in Figure 63).

The effects of not having a clutter model are stark; the original performance of the MR-2505 system as represented by Equation 5 is shown in the black curve of Figure 66, which shows nearly perfect performance (a capture rate of more than 0.9 with no false alarms, with the false alarm rate rising only to 0.03 to achieve complete UXO detection.) However, the performance of the more realistic version of the system which does not use the clutter model is shown in red, and has a detection rate of only 0.3 before false alarms start appearing, and still hasn't achieved total detection even when the false alarm rate rises to 0.5.

It was not a goal of the classification system analysis to develop an improved or more robust classifier to deal with clutter types which were not represented in the training data. However, one of the observations discussed earlier in Section 6.2 of this report was that the MR-2505 classification system was performing implicit orientation estimation and building simple sub-classifiers for different orientations of different targets. We theorized that simple models would be more robust, which could include robustness against false alarms due to clutter objects on which the system had not been trained. To investigate this hypothesis, we developed a simple adaptation of the MR-2505 GMM classifier which had a separate 15-component model for each UXO type and also varied the probability density estimates as a function of ping aspect angle. In brief, the standard GMM training algorithm [17] is an iterative expectation-maximization (EM) approach in which a set of training vectors are used to estimate component means, covariances, and weights. Given an initial set of component parameters, each iteration consists of two steps:

1. Estimate the membership of each training vector in each component (the likelihood that that component produced the training vector, given that it was produced by one of the components)

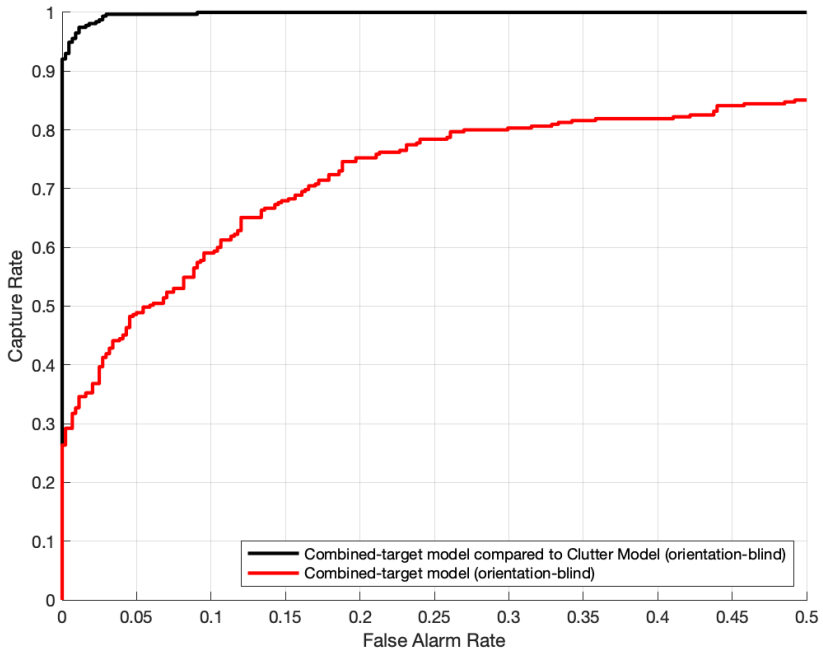


Figure 65: Receiver operating curves (capture rate vs. false alarm rate) for TIERSWAT acoustic-color classification under differing assumptions of clutter-characterization sufficiency. The black curve shows the results for the MR-2505 GMM system, which assumes all clutter object types in the test set are available for training, and achieves nearly perfect performance under that assumption. The red curve assumes no training data are available for the specific clutter types being tested against, and shows a significant performance degradation (see text for discussion of assumptions and resulting algorithm modification).

2. Re-estimate the component means, variances, and weights, with the training vectors contribution to the estimate for a particular component being proportional to their membership in that component

We modified the training algorithm to create a model which varied as a function of hypothesized ping aspect angle. To create the model for a given aspect angle, we began with the model trained on all the ping feature vectors for all aspect angles of a particular UXO type and adapted it through an additional component-weight estimation step in which the membership level of each training vector was modified by the closeness between the model aspect angle and the training vector aspect angle, so training vectors whose aspect angles were significantly different from the model angle were not used to calculate the component weights. To use the modified system in classification, we assumed an accurate orientation estimate[9] and then used the known target range and receiver elevation to choose the appropriate aspect angle for the model. The modified system's classification results are shown in the green curve of Figure 67. It can be seen that a significant fraction of the loss incurred by removing the clutter training data has been recovered; the system achieves a 90% detection rate at a 7.5% false alarm rate. ***We emphasize again that the goal of this work was not to propose an alternate classification system which was robust to unknown clutter types, but merely to investigate whether a simple modification of the existing algorithm to include orientation estimation would be of value. The positive result supports the approach of the follow-on work which we discuss in Section 8 of this report.***

GMM Component Usage: Visualizing Key Spectral Features

Section 6.2 discusses the results and implications of representing the fit of a particular unknown target interrogation by the GMM component indices which best fit the interrogation pings' feature vectors. While this representation provided insight into how the models were subdividing the problem and the importance of aspect angle and target-type specialization, the representation was somewhat abstract and provided little insight into the features being used or the underlying target physics which were being exploited. An alternate method of visualizing the fit of models to target pings is to represent each ping feature vector using the mean vector of the GMM component which best fits that ping. In the MR-2505 TIERSWAT classification study we analyzed, each interrogation had 59 pings with one-degree azimuthal spacing. Figures 68 and 69 show examples of this visualization technique for two different types of UXO, with six interrogations for each type.

The value of this method of representing model fit is that we can clearly observe the spectral features that the model uses to characterize each ping of the interrogation, and this allows us to make a number of inferences. In particular, we note that there is very high correspondence across the six interrogations of which component spectra are used to represent the scattered returns. Further, the extracted spectral structure is consistent across target range and therefore grazing angle, in spite of the resulting multipath interactions changes. Additionally, it can be seen that significant extents of the range of aspect angles for a given UXO type are being characterized by particular spectral shapes. This observation will be used to guide follow-on work by understanding where differences between simulated and field-collected data are critical and when spectral variability due to target condition or burial state will affect classification performance.

[9] In a practical system, this approach would require either a search across all possible target orientations or a method for estimating target orientation. Our follow-on work to begin in FY20 will explore the feasibility of these approaches.

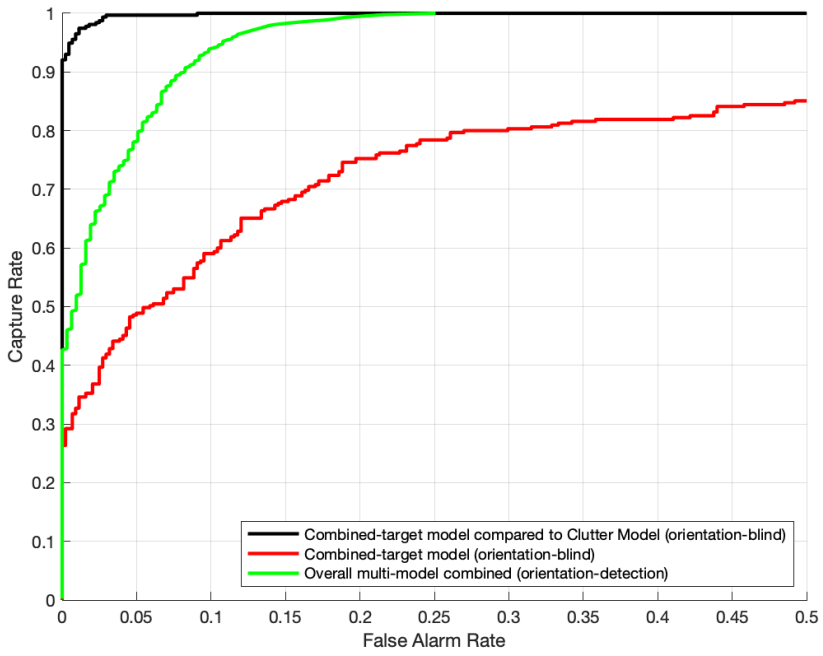


Figure 66: Receiver operating curves (capture (detection) rate vs. false alarm rate) for TIERSWAT acoustic-color classification under differing assumptions of clutter-characterization sufficiency. The black curve shows the results for the MR-2505 GMM system, which assumes all clutter object types in the test set are available for training, and achieves nearly perfect performance under that assumption. The red curve assumes no training data are available for the specific clutter types being tested against, and shows a significant performance degradation The green curve shows a result from a system modified so that each UXO type is modeled separately, and that models vary as a function of incident aspect angle. The simple modification allows the recovery of a significant amount of the performance loss (see text for details of the algorithms and motivation).

The implication of consistency (for a particular UXO type) of component usage for a particular aspect angle across different interrogations and across nearby aspect angles is that the evolution of spectra for a particular target type can be described in a very low-dimensional space. As discussed in Section 7, next-generation sonar systems for UXO interrogation are based on down-looking geometries rather than the side-looking data studied under MR-2505, and we noted that down-looking interrogations will undergo much higher variability in return spectral evolution across pings as a function of target orientation than in the side-looking case. This low-dimensional representation of the key spectral structures and their evolution may be critical in making possible future approaches which rely on searching across the space of variability due to target orientation, burial state, and condition.

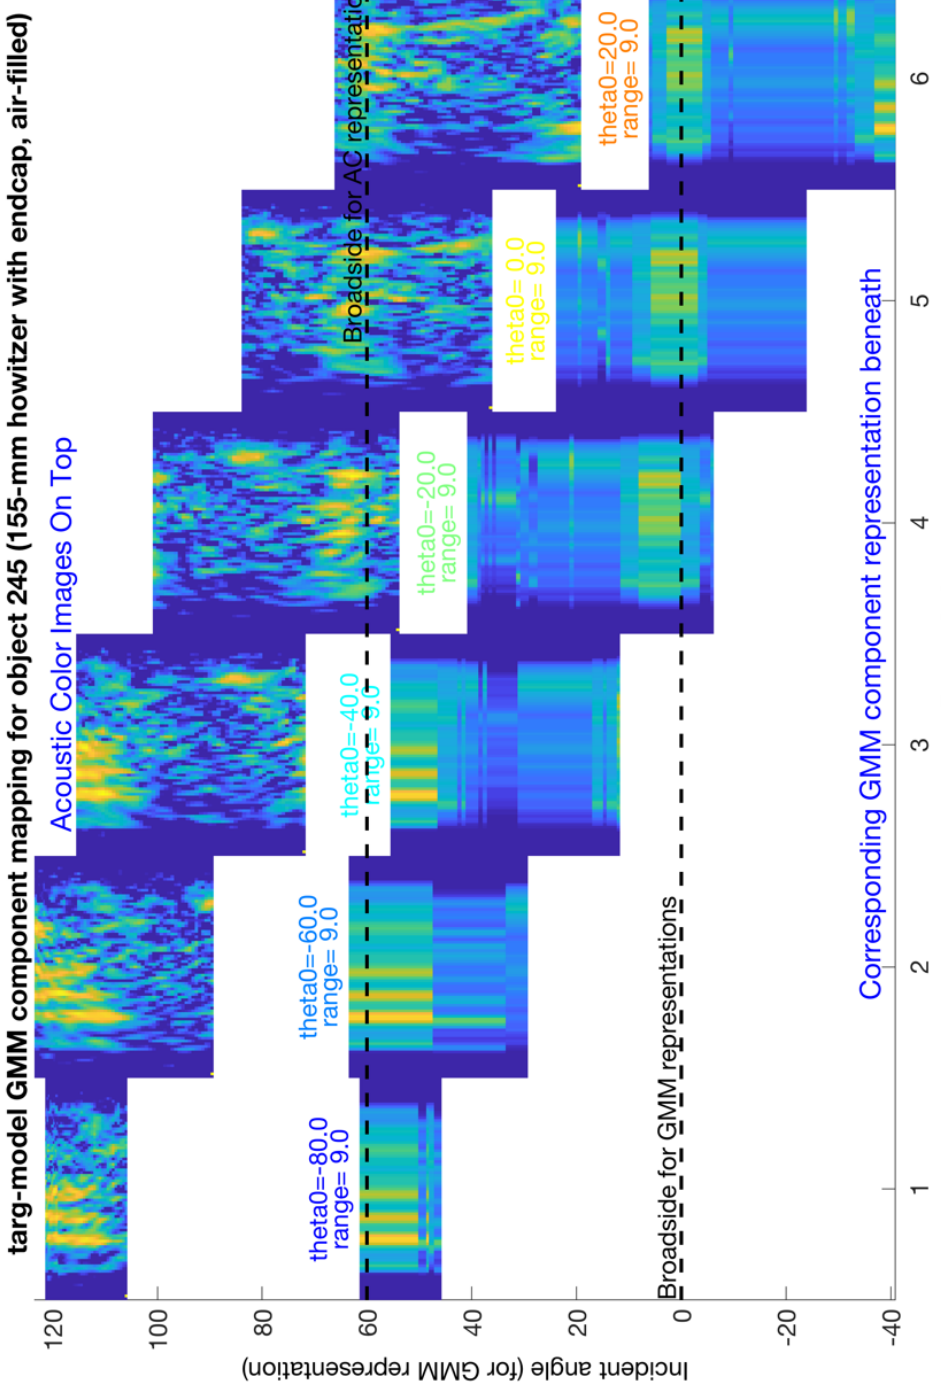


Figure 67: Visualizing the GMM representation of 155mm Howitzer acoustic color (AC) information by replacing each ping return's spectrum with the best-fit GMM component center. The AC for six interrogations of Howitzers are shown in the top images, located so that the same target aspect angle in each AC are aligned horizontally (the dotted line represents broadside). The lower images show the corresponding best-fit GMM components for the ping returns in the AC above them. The key spectral structures used in classification can be visualized, and it can be observed that the same structures are identified in each interrogation at a given aspect angle and that relatively large extents of aspect angle are characterized by particular spectral structures.

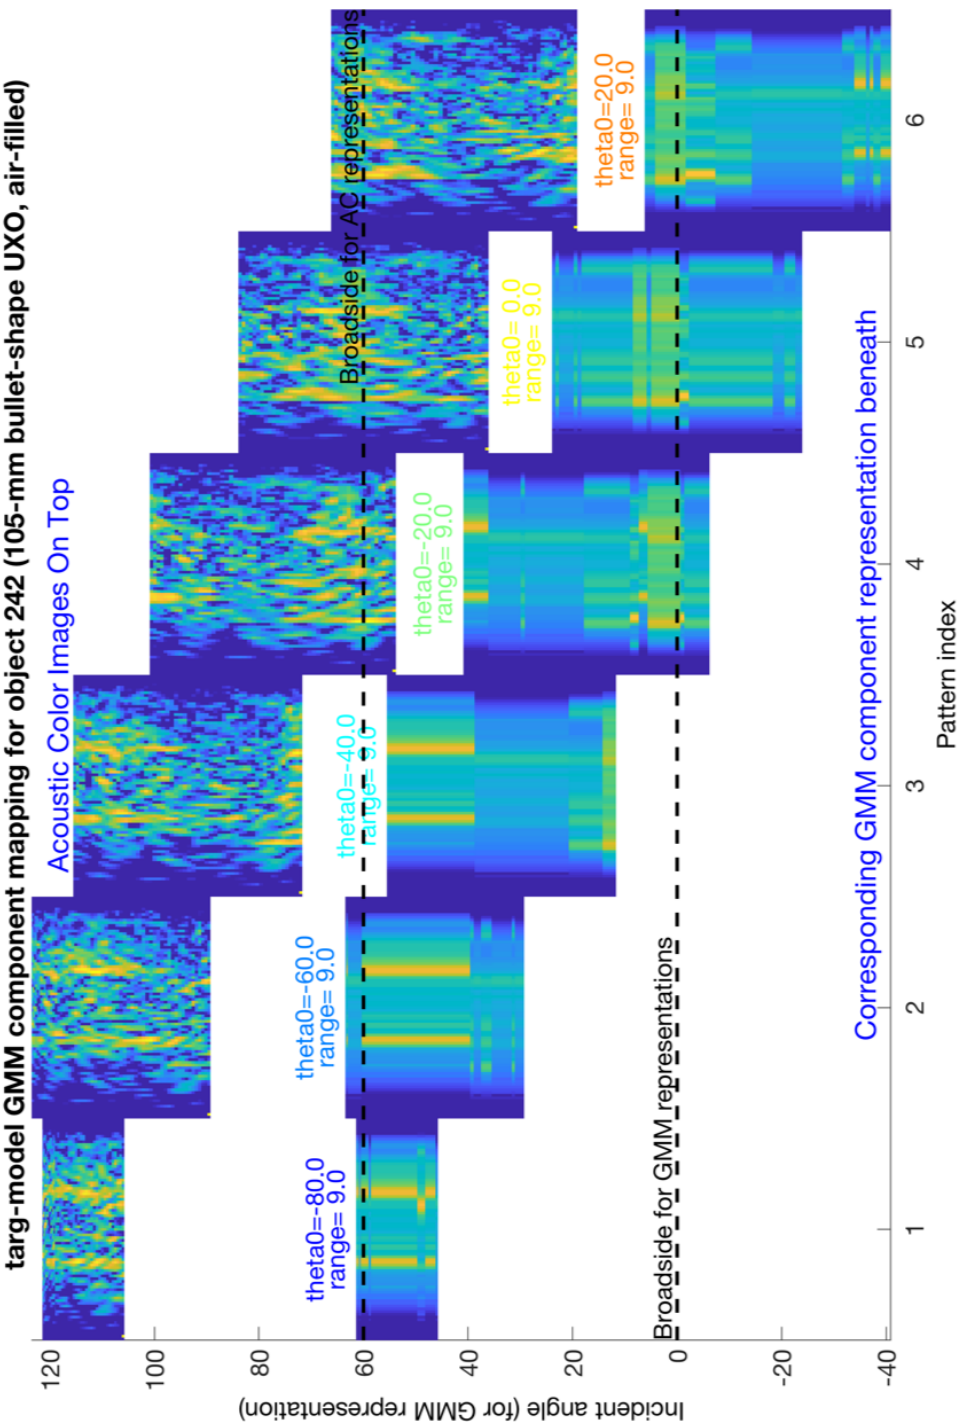


Figure 68: Visualizing the GMM representation of 105mm Howitzer acoustic color (AC) information by replacing each ping return's spectrum with the best-fit GMM component center (see Figure 68 caption for explanation).

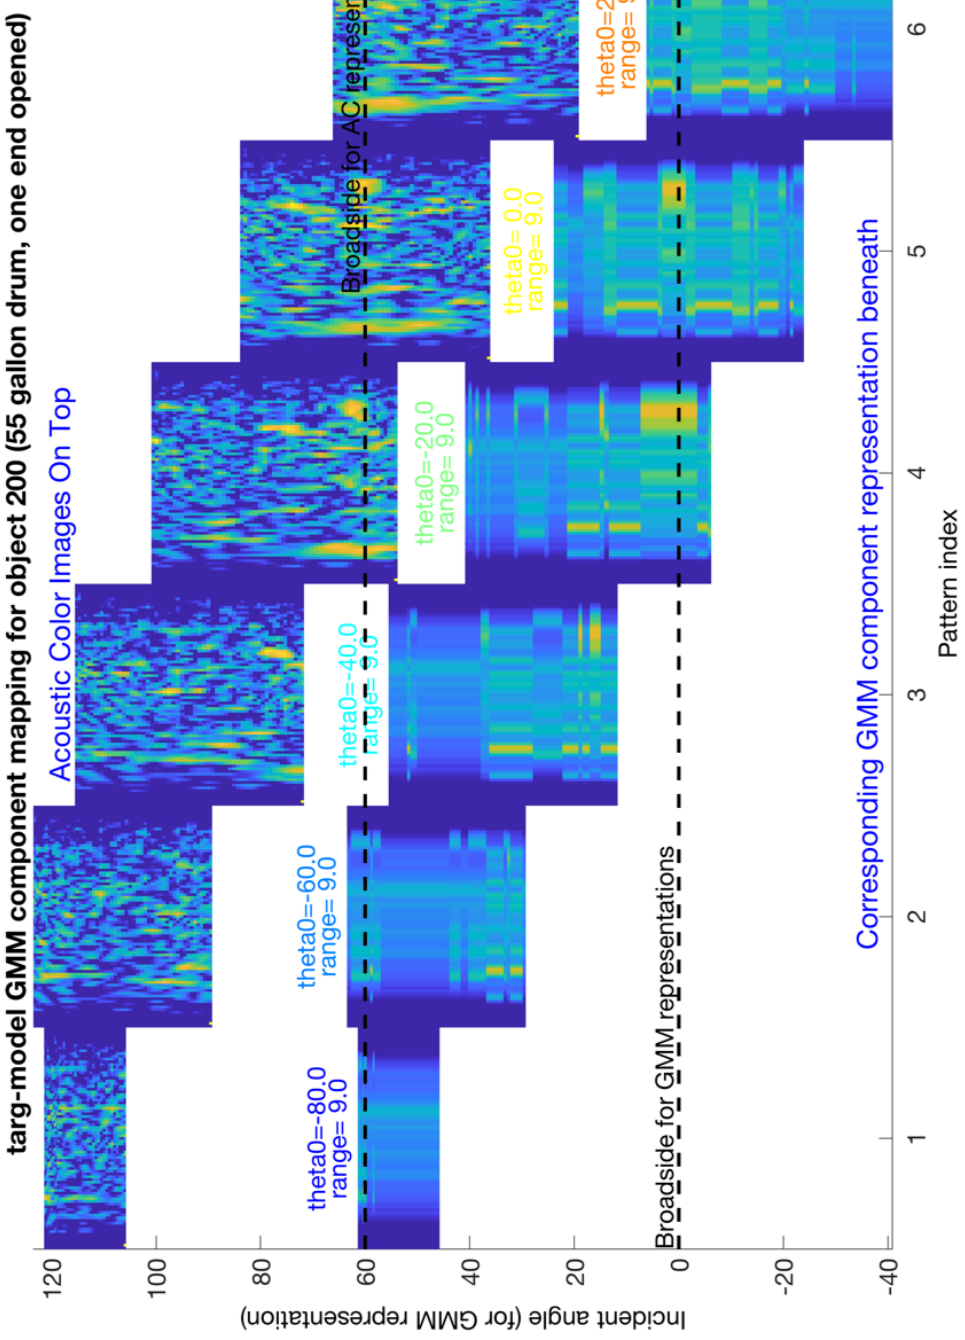


Figure 69: Visualizing the GMM representation of 105mm Howitzer acoustic color (AC) information by replacing each ping return's spectrum with the best-fit GMM component center (see Figure 68 caption for explanation). The clutter AC representation differs from the target representation in both the component spectra which are used and the evolution of component usage as a function of angle; both differences can be exploited in separating clutter from UXO.

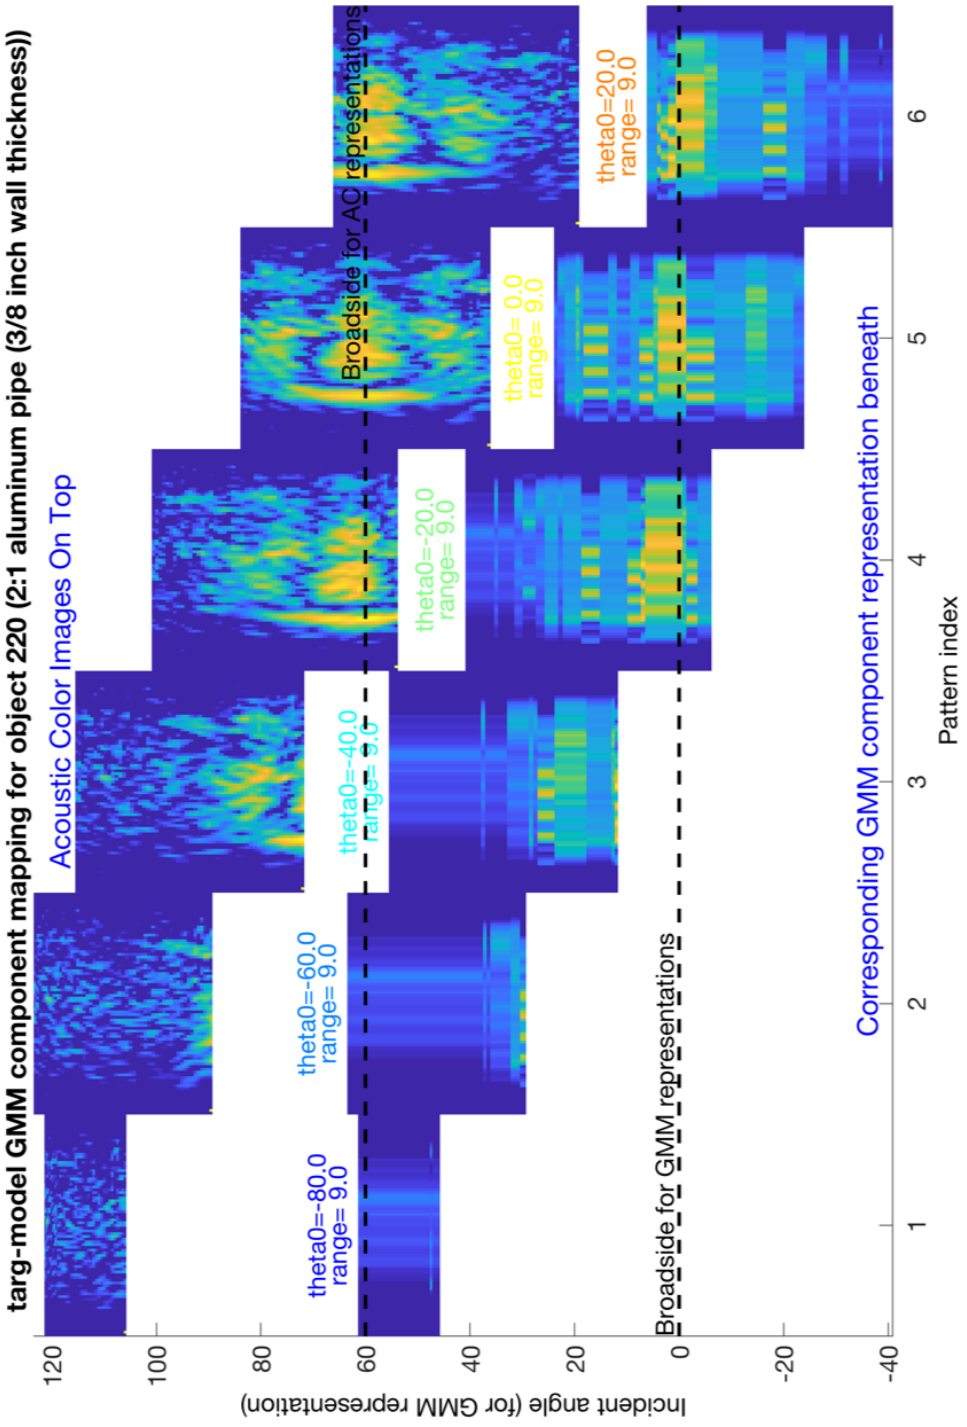


Figure 70: Visualizing the GMM representation of 105mm Howitzer acoustic color (AC) information by replacing each ping return's spectrum with the best-fit GMM component center (see Figure 68 caption for explanation). The clutter AC representation differs from the target representation in both the component spectra which are used and the evolution of component usage as a function of angle; both differences can be exploited in separating clutter from UXO.

7 Down-Looking Sonar: Implications for Feature Sets and Classification Architecture

When the performers first approached the problem of physics-based UXO classification for the SERDP SEED project MR-2324, the paradigm (and the available data) was focused on side-looking sonar systems. This approach has a potential advantage in wide-area coverage, but had a number of disadvantages as well. Some of these were related to the development of potentially sensitive technologies. However, from a modeling perspective one of the biggest challenges in a side-looking sonar approach was that the effects of multipath could make the same target produce a very different return as a function of source/receiver height (which could be controlled), target range (which might not be controllable under many conops) and sediment composition (which is of course not controllable).

Down-looking sonar, on the other hand, does not have the per-ping coverage possibilities, particularly in shallow water. However, for applications like UXO mitigation where time frames are flexible, conops can be adjusted to mitigate this issue. As a result, significant advantages of down-looking approaches can be potentially exploited. Primary among these may be that the multipath issue is severely lessened (though the direct-path bottom return is of course increased). However, a down-looking configuration introduces another key challenge which makes even more crucial a physics-based approach to classifier design such as the one being pursued in this work.

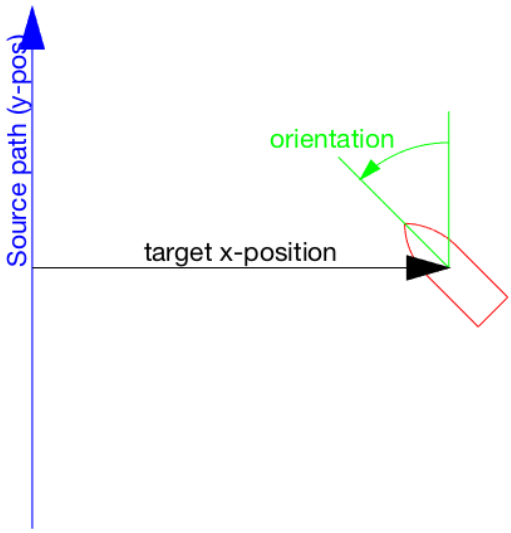


Figure 71: Cartoon describing the geometric relationship between a down-looking sonar traveling in the positive-y direction (along the blue line) above a seafloor defined as the xy plane. The position of the target is defined by the x-position in meters off of the source path, and by the target orientation in degrees.

Consider the source/target scenario illustrated in Figure 72. The source/receiver is moving in a straight line, and the position of the target is defined by the distance of the target from the closest point on this line and the orientation of the target (for this discussion, the target is assumed to

be lying horizontally). In trying to predict the scattered response of the target during acoustic interrogation, a key question is the aspect angle at which the sonar wavevector strikes the target (where, for an example of an axisymmetric target with a flat end and a pointed end, we define -90 degrees as striking the flat end along the major axis and 90 degrees as striking the pointed end along the axis. Figure 73 shows the observed target aspects as the source with an elevation of 5 meters moves along a hypothetical 30 meter path from y=-15m to y=15m and the target is at y=0m. It can be seen that the traverse in aspect angles varies greatly as a function of target x position and orientation (as described in Figure 72. At one extreme is the case where the target is directly in the source path (x=0m) at 90-degree orientation, where only one target aspect is observed during the traverse.

The significance of this variability can be considered in terms of systems developed for side-looking sonars which are based around template-correlation of acoustic color images. In existing systems which define the match between a template acoustic color image and an unknown acoustic color image by the maximum cross-correlation observed across different shifts in angle, the assumption is that the evolution of target aspect as a function of azimuth in the acoustic color is the same in both template and unknown and the only variable which must be searched over is the starting angle. This is equivalent to saying that the two patterns have similar slopes in the representation of Figure 73.

In a side-looking sonar, it is of course known that the aspect evolution as a function of azimuth is grazing-angle-dependent, both because of the effects on multipath interaction and on the changing relationship between azimuth and target aspect. However, the range of grazing angles[10] at which the target may be observed is relatively small (for example, a receiver at 4 meters off the bottom whose near-field was 10 meters would see grazing angle variations of less than 20 degrees) and current results suggest that the expected variations can be covered by using templates collected at a small set of ranges.

With down-looking sonar, the mismatch problem is much more significant. This is because the evolution of target aspect as a function of azimuth is dependent on target orientation, which may be anywhere from 0 to 360 degrees and result in extreme variation in the acoustic color. Figure 74 illustrates this effect by simulating the direct returns that would result from the science target over the range of orientations and positions shown in Figure 73. The variations among these patterns are much more extreme than is observed in side-looking sonar, and would thus require a much larger set of templates to be collected or simulated in order to achieve the same degree of correlation potential. Even then, however, a direct transition of the existing approach would require a significantly more complex subsequent classification algorithm. It should be noted that side-looking sonars can in theory face similar effects if targets are cocked at an angle significantly off of the horizontal.

A physics-based approach to the problem, however, should be able to compensate for these distortions. This approach begins with the understanding that the primary source of the variation among returns which would be collected for the above-described range of target positions is in the angular evolution, and for a given aspect observed in each of these cases the physics are likely to be very similar. In theory, the simplest variation on an existing correlation-based system would be to distort the templates to simulate the range of possible target orientations. However, the correlation component of the classifier is already the most computationally expensive part and adding this new dimension of search would be a huge additional burden. Even in scenarios where additional computation can be accomodated, however, the multipath effects would be different at different

[10] *grazing angle*: the angle between the horizontal and the direction of the target

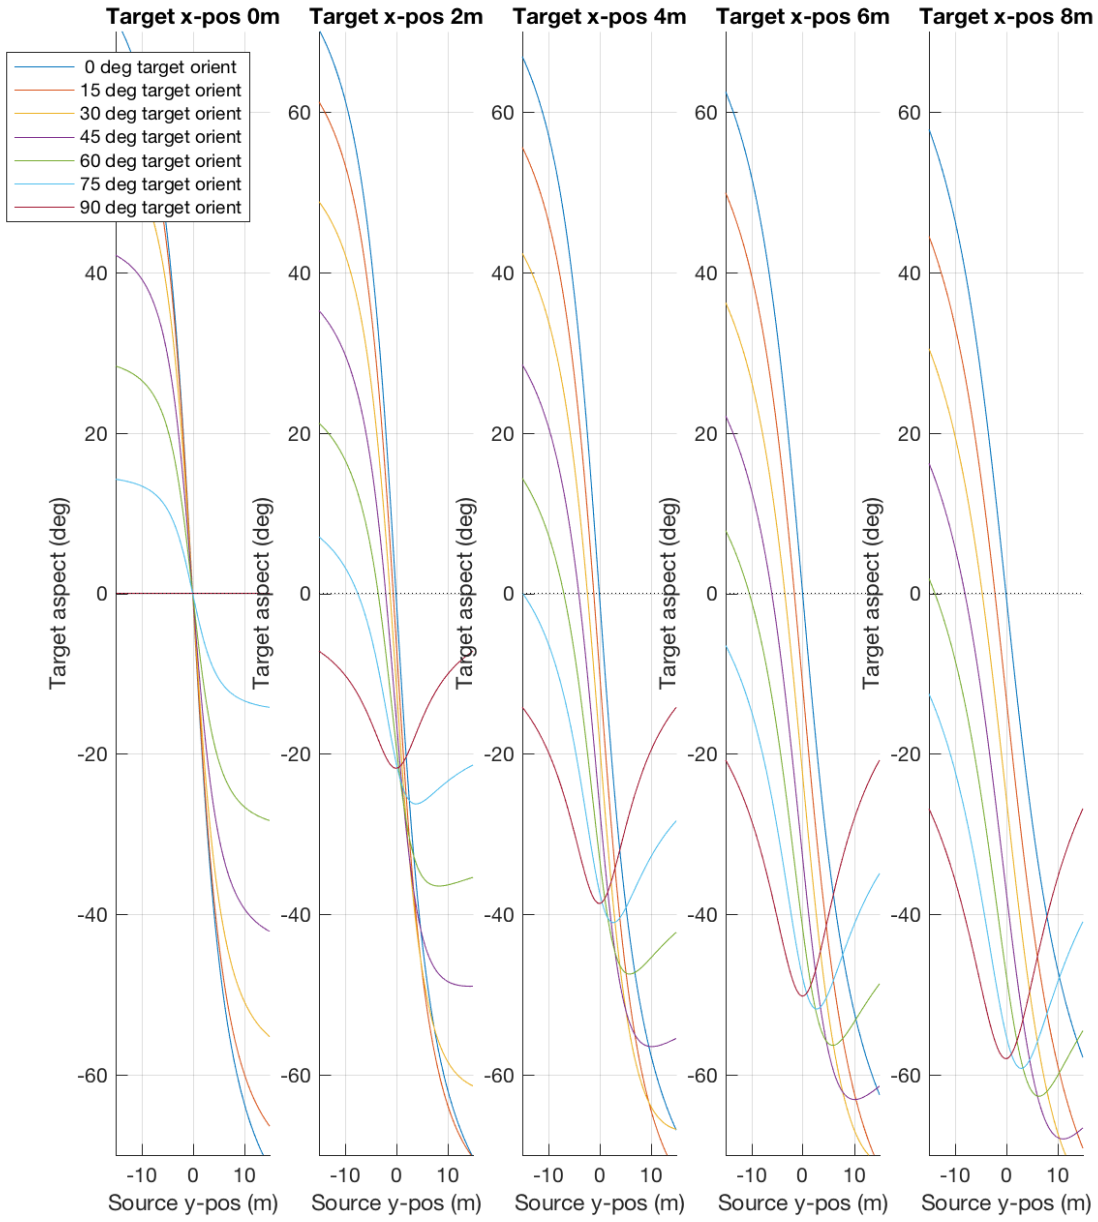


Figure 72: Observed target aspect for a down-looking source-receiver making the traverse illustrated in Figure 72 from y=-10 to y=+10 meters. Each sub-figure represents a different target position, and each line represents a different target orientation at that position. Extreme variation in aspect traverse as a function of these two variables can be observed. Simulations of the response acoustic color are shown in Figure 74.

target x-positions and that would not be accommodated under this framework.

It has been shown, however, that the multipath effect can be modeled accurately with a superposition of freefield responses [11]. As a result, a classification system based not on comparison static templates, but on features whose realization at an arbitrary target aspect could be predicted or measured on the fly, could exactly handle any new geometry through superposition. This approach would be even more powerful if orientation information could be known *a priori*. The system under development has plans for imagery-based data products that would allow for accurate estimation of target size and orientation, and the feature realization should take advantage of that information.

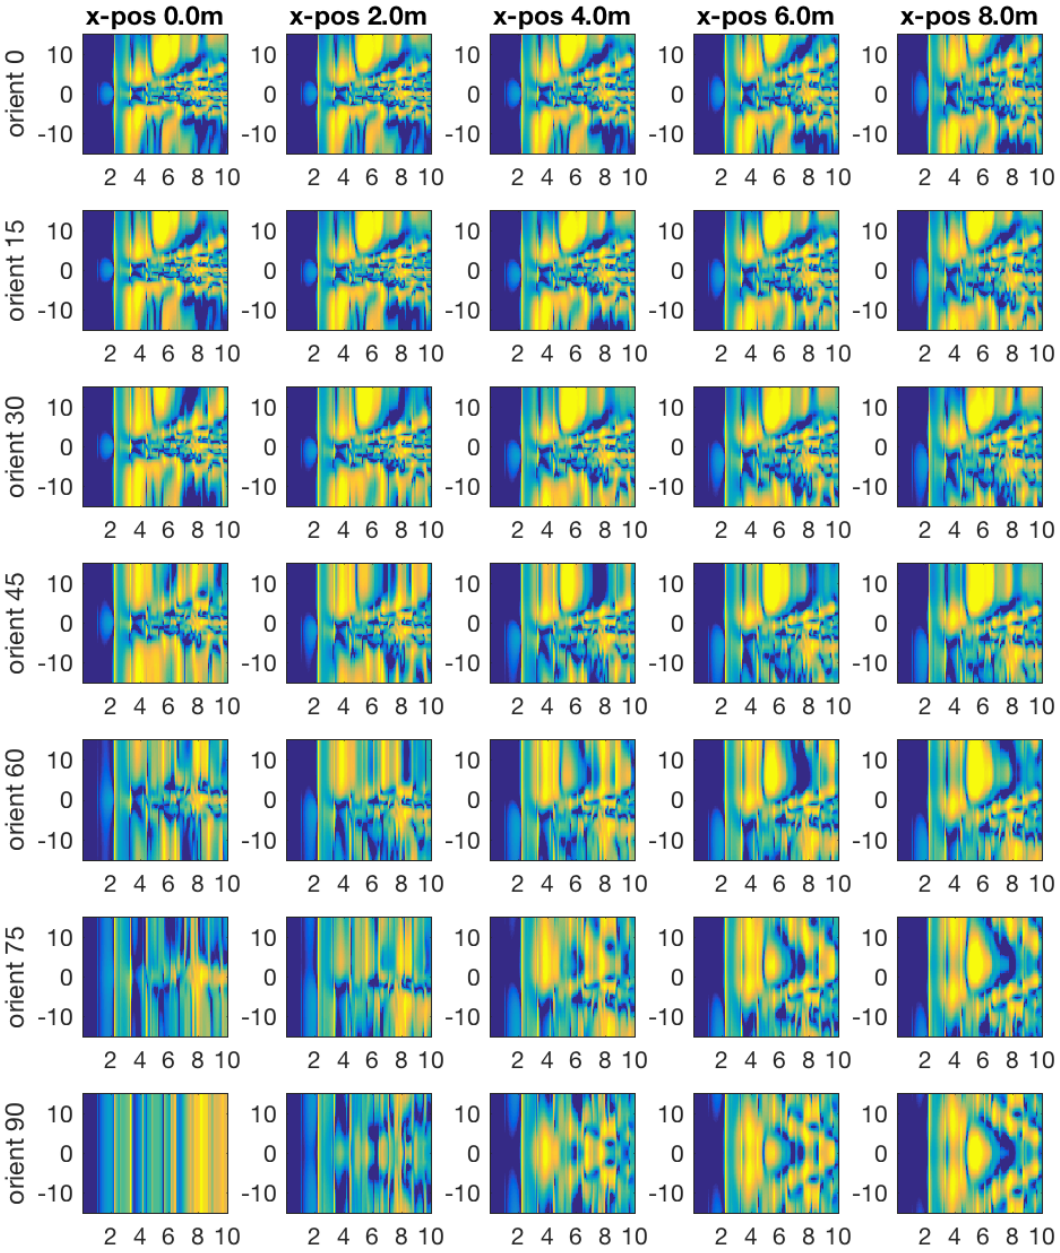


Figure 73: Simulated acoustic color responses corresponding to a science target observed by a down-looking sonar given the locations/orientations from Figure 74. The x axes are frequency and the y axes are position of the source along the traverse.

8 Conclusions

This report summarizes the work done to relate observables in the scattered returns from acoustic interrogation to the underlying geometry and composition of the elastic object being interrogated. We approached this problem from two different directions, both of which yielded important insight into the development of feature sets which are robust to the variations in returns which will be observed in fielded systems.

The first approach we took to identifying relationships between observables and the underlying physics was a "bottom-up" investigation, driven by the FE model itself. We performed a survey of prominent free-field observables in the sub-10kHz range and categorized them according to their underlying physics. We observed that the types of physics behaviors could be related to important properties of the geometry and composition of the target. We developed measures for identifying regions with similar internal physics, and created modified FE models that allowed us to isolate classes of physics behaviors so that we could better understand the range of conditions under which returns due to a particular behavior could be exploited.

The data-collection effort of ClutterEx17 allowed us to assess and validate the FE models under study, both for the science target manufactured in this effort and for existing targets and models used for companion efforts at APL-UW. The first conclusion of our study in this area was that the model accuracy is sufficient for us to be confident that the models give us reliable information about the internal physics behind the observables in the collected data. The second conclusion, however, was that differences between the predicted and collected results are significant enough that current feature-representation approaches will not be able to take advantage of the information available from the model, and physics-based features that are robust to the sources of these variations are necessary.

The second approach we took to the identification of potential features was a "top-down" investigation driven by an analysis of two existing APL-UW classification systems. The first analysis was a classification system for field-collected and a modified version of this system which allowed for a rudimentary exploitation of physics knowledge through separate processing of elastic and inelastic components. Analysis of the cases where performance improved allowed us to identify candidate observables for study. The FE target model allowed us to understand the physics behaviors which are the source for these observables, and again relate them to target structure and composition. The second analysis was of a simulated-data classification system, and our investigation allowed us to draw insights into how to improve robustness by modeling UXO types separately and incorporating target orientation estimates. We also demonstrated that the reality of incomplete clutter characterization must be considered in developing a classification architecture.

Since the origins of this work in the SEED effort, the focus has evolved from a concept based around a monostatic side-look sonar to that of a multistatic down-look sonar. The implications of this change are significant; there are many advantages for this approach but they will require concepts of feature extraction and classification architecture that, while driven by the same underlying physics, are very different from the approach that has been taken in the side-look scenario. The remainder of the work in this effort will be adapted to classification strategies for this new system concept.

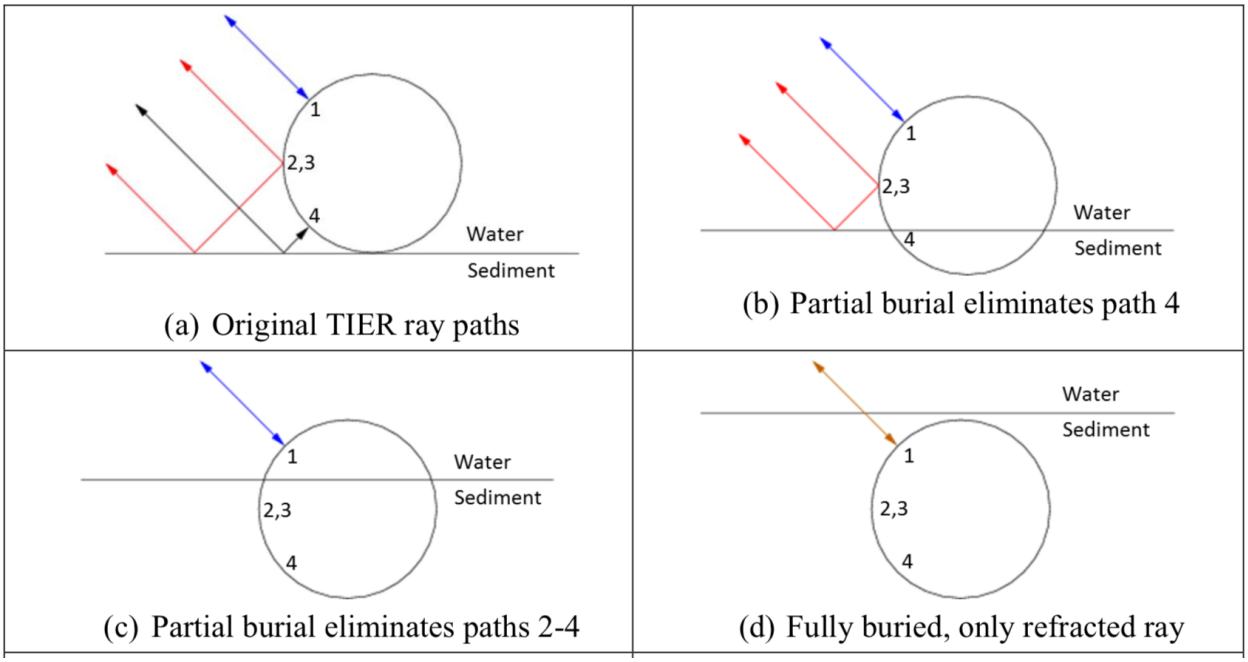


Figure 74: Ray diagrams for a cylinder at broadside or spherical target: (a) proud target with 4 paths included; (b) burial depth eliminates path 4 and only paths 1-3 are present; (c) burial depth eliminates paths 2-4 and only path 1 remains; and (d) full buried target with only a refracted path. (Figure courtesy Dr. Steve Kargl)

8.1 Follow-On Work

The conclusions outlined above have guided the directions for our follow-on work, which has been funded under MR20-1443 (PI: Owsley) and will begin later this year. An APL-UW classification framework is under development ((ESTCP MR18-5004, PI: Williams) and has been tested on ClutterEx17 data sets; it includes a detection algorithm that then extracts the image, the calibrated acoustic color, and time domain data products for each detection. A custom classification scheme that includes a data pre-processor, feature extractor, and trainable classifier has also been generated. The current code then outputs an image of the local scene, showing all contacts and their assigned class labels. The modular nature of the system allows physics-based understanding to be incorporated into various stages of the processing chain as they are developed. Our follow-on contribute to this system in a number of areas, in particular:

Task 1: Classification based on orientation estimation: This task is motivated by the analysis of existing APL-UW classification experiments which suggests that well-performing black-box classifiers are doing implicit orientation estimation and applying different measures to returns based on incident angle at the target rather than simply identifying some universal characteristics of returns that can be observed across orientation, as described in Section 6.2 of this report. Our analyisis suggested that an approach that explicitly estimates orientation and chooses a classification method informed by that estimate could outperform a generic classification approach. This will be particularly critical for a down-looking sonar scenario in which the relationship between source position and incident angle at the target will vary much more greatly as a function of target orientation.

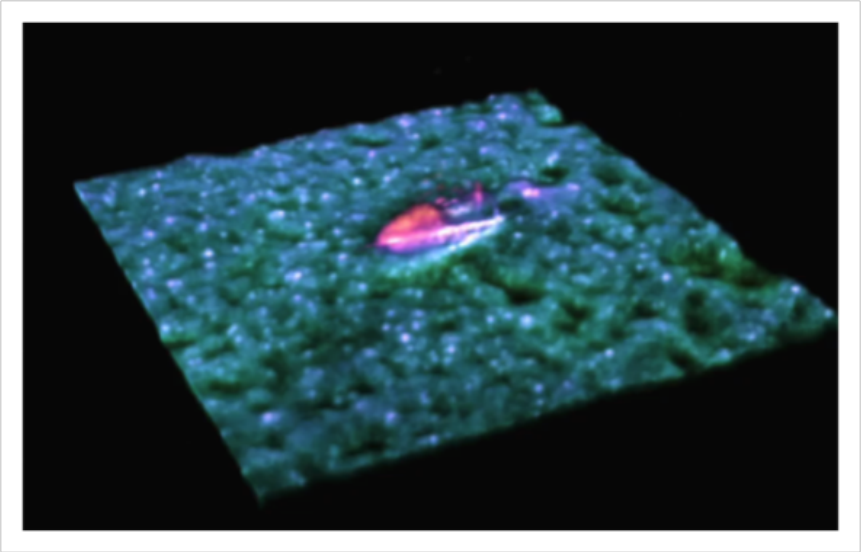


Figure 75: A circlular SAS image of a 155 mm UXO produced using high-frequency data. Image processing techniques could be used to produce an orientation estimate for use in subsequent classification algorithms. (Figure courtesy Dr. Timothy Marston, from MR-2439 final IPR).

For proud or slightly buried targets, image-processing techniques operating on output from the SAS being developed for incorporation into the APL-UW system could potentially provide good orientation estimates. Figure 76 shows an example of a high-frequency circlular SAS image of a proud tilted 155 mm UXO which could be available as an intermediate data product in the APL-UW system for image-processing-based orientation techniques.

However, we first propose to investigate an alternative approach to target orientation estimation that has the possibility of application to more deeply buried targets as well. This approach is motivated by the Target-in-the-Environment Response (TIER) work of Dr. Steve Kargl (partially supported by MR-2231))[11], who has demonstrated that accurate predictions of a target's scattered response can be achieved by considering the complete response to be a combination of the individual responses of waves traveling along a set of ray paths determined by the target's position relative to the sediment layer, as illustrated in Figure 75a. This can be formulated as a linear superposition of versions of the freefield response, with each version scaled and phase-shifted to incorporate the time-delays of the paths and reflection coefficients of the interfaces:

$$P(w) = \left[\frac{B_1}{d_0^2} f_1 e^{i\omega t_0} + \frac{2L(\theta_g)B_2}{d_0 d_1} f_2 e^{i\omega t_1} + \frac{L^2(\theta_g)B_3}{d_1^2} f_3 e^{i\omega t_2} \right] r_0 P_{src}(\omega) \tag{7}$$

where ω is an angular frequency, B_i are beam patterns, d_i are separation distances, f_i are complex scattering amplitudes, t_i are propagation time delays, $L(\theta_g)$ is a reflection coefficient for the water-sediment boundary, θ_g is the local grazing angle of incidence, d_0 is a normalization distance, and $P_{src}(\omega)$ is the spectrum of the transmitted signal. The three terms in the equation represent the direct path (1), the bottom/direct (2,3) and bottom/bottom (4) paths respectively, with the numbers corresponding to the rays in Figure 75. Kargl has also developed a modified version of

this equation that handles partly buried and fully buried targets by inclusion of water/sediment transmission coefficients and removal of ray paths as the geometry dictates, as in Figures 75b, 75c, and 75d.

By modeling the in-environment response as a linear combination of free-field responses, this formulation suggests that a matched filter of the freefield response at the correct orientation would have a strong response to the individual components of Equation 7. Preliminary work by Dr. Timothy Marston [18] provides evidence to support this hypothesis; Dr. Marston observed that field-collected scan data for an air-filled howitzer exhibited significantly higher correlation to a properly-oriented FE model for that object than to models containing other fillers or to potential clutter models of similar size. We propose to leverage this idea by developing a filter bank containing modeled freefield responses $\hat{P}_{k,i}(\omega)$ for each target of interest k at a variety of possible orientations indexed by i. Several research topics will be addressed. The ideal filter would be a function of orientation and material parameters; we will study the fineness of sampling which will be required to provide a useful orientation estimate. Further, we will consider the value of combining estimates from multiple pings and the need to consider the interaction of the terms in Equation 7. We will also report on the additional potential value which could be added through post-mission processing with a larger number of filters than could be processed on-board.

The system will produce a confidence measure for the orientation estimate, as well as a potential recommendation on travel direction for maximally informative follow-up scans. The deliverable will be a classifier which is specialized for particular target orientations when they can be accurately estimated, but will revert to the baseline system when orientation is uncertain.

Task 2: Use of physics-based features for identification of specific target classes: The goal of this task is to develop modules for specific target types that incorporate specific knowledge of the physics underlying the acoustic signatures of those target types. There are several advantages to this approach. First, it enables us to add information about each target type as it is studied, and to tailor the classification algorithm to the relevant features as appropriate. It allows us to be more robust to environmental and target-composition variations as they affect the features of interest. Finally, identification based on the unique features of a particular target type will allow us to be more robust to the presence of clutter types that have not been explicitly modeled. The classification architecture being developed at APL-UW is well-suited to sequential incorporation of such target identifiers as they are developed.

In Section 6.2 of this report, we noted that the results shown in Figure 64 indicate that the black-box target model we analyzed was doing implicit incident angle estimation and subdividing its discrimination power by treating different angles differently. It can be further observed that similar subdivision is occurring for different target types, with particular GMM components specializing in identifying different types of UXO replicas. The classification algorithm being investigated in that analysis assumed a two-class problem in which models are created for both the target distribution and clutter distribution, and make a classification decision based on which model was a better fit. Architectures based on algorithms like SVMs or RVMs make similar assumptions. However, a risk of this approach is illustrated in Figure 65, which shows the fit of the clutter model components to several types of clutter objects. Similar to the case with the target models, we again see that the clutter model is subdividing its discrimination power by using particular components to specialize

in identifying each clutter type. A reasonable inference from this result is that the feature set being used does not utilize some universal 'clutter features' but that each of the modeled clutter types is occupying a discrete part of the feature space. As a result, other clutter types not previously modeled could be expected to be a poor fit to the existing clutter model and thus not well-handled by the two-class paradigm as their relative fit between the two models is not a good measure of UXO likelihood. As we develop identifiers for particular UXO types, we will therefore not just be looking for features and algorithms which distinguish those types from known clutter types but whose presence can be related directly back to the geometry and physics unique to that UXO.

One of the key considerations in development of a feature set and classification architecture will be the effect of down-looking sonar on the acoustic returns. As discussed in Section 7 of this report, a down-looking geometry, while advantageous in many ways, will result in significantly more variability in scattered returns as a function of target orientation than in side-looking sonar systems. A physics-based approach to the problem, however, should be able to compensate for these distortions. This approach begins with the understanding that, as illustrated in Figure 74, the primary source of the variation among returns is in the angular evolution. For a given incident angle observed in each of these cases the physics is likely to be very similar. In theory, the simplest variation on an existing correlation-based system would be to distort the templates to simulate the range of possible target orientations. However, the correlation component of the classifier is already the most computationally expensive part and adding this new dimension of search would be a huge additional burden. Even in scenarios where additional computation can be accommodated, however, the multipath effects would be different at different target x-positions and that would not be accommodated under this framework. The TIER work of Kargl has shown, however, that the multipath effect can be modeled accurately with a superposition of freefield responses [11]. As a result, we propose to develop a classification approach based not on the comparison of static templates, but on features whose realization at an arbitrary target aspect could be predicted or measured on the fly, and could exactly handle any new geometry through superposition.

Task 3: Physics-based assist for post-mission analysis: For ConOps where potential targets have been detected and classified via an automatic classification algorithm, the APL-UW GUI provides a framework for a user to drill down into the collected data and explore a variety of data products derived from the underlying acoustic data collected on those targets. These data products would be displayed with accompanying physics-based representations of those same types of data products, so that users can directly compare the measured data to representations of what would be expected from targets of interest based on a physical understanding of acoustic scattering from the hypothesized target type. This would allow the user to make a more informed classification decision in low-confidence situations, and to identify best options for follow-up.

Figure 77 shows a mock-up of what this GUI might look like. A SAS image showing the surveyed region is shown in the upper left, with boxes indicating detected objects and their classification result. The system has flagged a particular object for which the identification result is of low confidence and shown a zoomed-in image of that object. The user is presented with a data-product representation of the scattered return from the object and, for comparison, a model of what the expected return would be given a hypothesized UXO type. The user can investigate the internal physics that are responsible for the observables in the modeled return. Further, the system will be able to analyze the physics of the complete model for the hypothesized UXO type to determine the

best path for a further scan to make a higher-confidence decision.

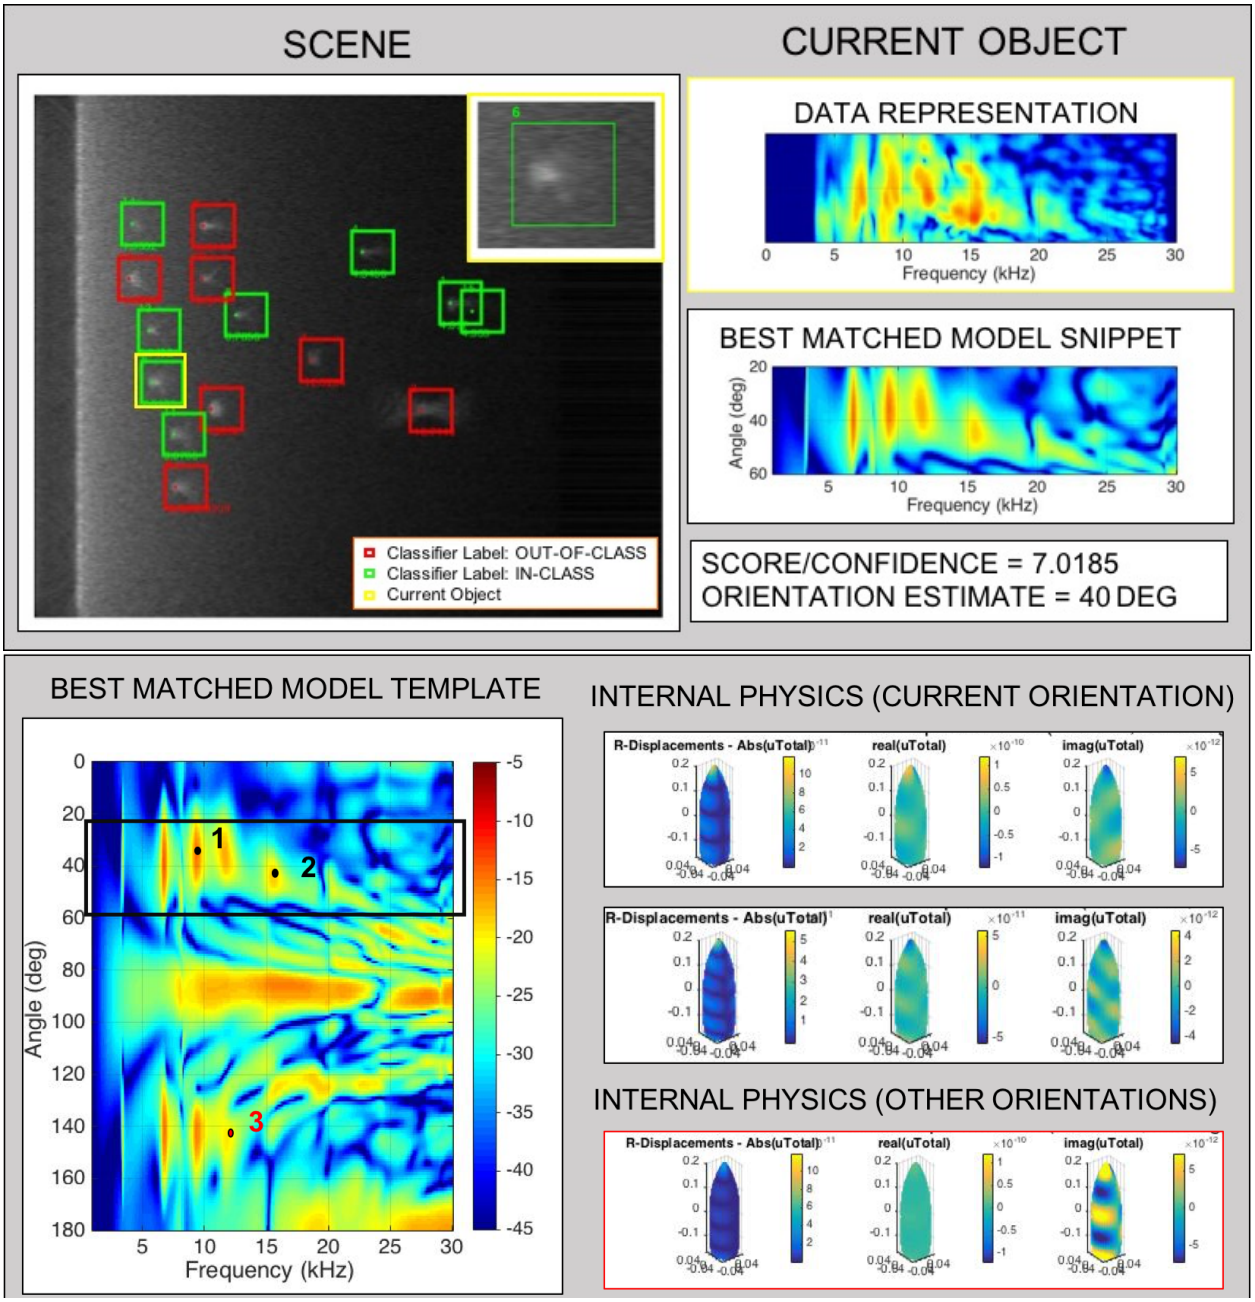


Figure 76: A mock-up of a proposed GUI that would enable post-mission analysis and follow-up in a physics-based framework. The upper left figure shows a SAS image of a complete scanned scene, with boxes indicating detected objects and their classification results. An object of uncertain identity (the yellow box) has been flagged for investigation, and a data representation (acoustic color) of the scattered return is displayed in the upper right, below which is the return which would be predicted for the hypothesized class. The user can investigate the underlying physics corresponding to the observables in the current orientation (#1 and #2 on the lower right). Further, the system can analyze the complete response of the hypothesized class at all angles (lower left) to identify the physics behaviors which would produce the most definitive return (#3 in lower right) and propose a follow-up scan path designed to maximize classification confidence.

References

[1] J. Yang and D. Tang, "Direct measurements of sediment sound speed and attenuation in the frequency band of 2-8 khz at the target and reverberation experiment site," *IEEE Journal of Oceanic Engineering*, vol. 42, pp. 1102–1109, 2017.

[2] B Hefner, "Final report: Inversion of high frequency acoustic data for sediment properties needed for the detection and classification of uxos, serdp project mr-2229," Tech. Rep., Applied Physics Laboratory, University of Washington, 2015.

[3] Lane Owsley and Jack McLaughlin, "Modeling acoustic target response by component," in *Proceedings of MTS/IEEE OCEANS*, 2010.

[4] M. Zampolli, A. Tesei, F.B. Jensen, N. Malm, and J.B. Blottman III, "A computationally efficient finite element model with perfectly matched layers applied to scattering from axially symmetric objects," in *Journal of the Acoustical Society of America*, September 2007, vol. 122, no. 3, pp. 1472–1485.

[5] A.L. España, K.L. Williams, D.S. Plotnick, and P.L. Marston, "Acoustic scattering from a water-filled cylindrical shell: Measurements, modeling, and interpretation," *The Journal of the Acoustical Society of America*, vol. 136, pp. 109–121, 2014.

[6] M.J.J. Nijhof, A. España, and K. Williams, "Efficient calculation of broadband acoustic scattering from a partially, obliquely buried cylinder," *Proc. 1st International Conf. Exhibition Underwater Acoustics*, pp. 697–702, 2013.

[7] JA Bucaro, A Sarkissian, BH Houston, and M Saniga, "Effects of target corrosion/bio fouling on emi and structural acoustic signatures in underwater environments acoustic component final report to serdp mr 2500," Tech. Rep., Naval Research Laboratory Washington United States, 2017.

[8] A.L. España, K.L. Williams, D.S. Plotnick, and P.L. Marston, "Acoustic scattering from a water-filled cylindrical shell: Mode identification and interpretation via finite element and analytical models," in *Proceedings of the Acoustical Society of America*, 2013, vol. 19, pp. 1764–1769.

[9] M. Zampolli, A. L. España, K. L. Williams, S. G. Kargl, E. I. Thorsos, J. L. Kennedy J. L. Lopes, and P. L. Marston, "Low- to mid-frequency scattering from elastic objects on a sand sea floor: Simulation of frequency and aspect dependent structural echoes," *Journal of Computational Acoustics*, vol. 20, no. 2, 2012.

[10] Kazuo Kamigaki, "Ultrasonic attenuation in steel and cast iron," *Science Reports of the Research Institutes, Tohoku University, Series A, Physics, Chemistry and Metallurgy*, vol. 9, pp. 48–77, 1957.

[11] S. Kargl, A. L. España, K. Williams, J. L. Kennedy, and J. Lopes, "Scattering from objects at a water-sediment interface: Experiment, high-speed and high-fidelity models, and physical insight," *IEEE Journal of Oceanic Engineering*, vol. 40, pp. 632–642, 2015.

[12] Daniel S. Plotnick and Timothy M. Marston, "Final Report: ATR training set generation from multi-pass CSAS data and ATR feature sensitivity studies," Award Number N00014-16-F-3014.

[13] David Williams, "Demystifying deep convolutional neural networks for sonar image classification," in *Proceedings of the Underwater Acoustics Conference*, September 2017, Skiathos, Greece.

[14] Michael E. Tipping, "Sparse Bayesian learning and the relevance vector machine," *Journal of Machine Learning Research*, vol. 1, pp. 211–244, 2001.

[15] G. Sammelmann, *Personal Computer Shallow Water Acoustic Tool-Set (PC SWAT) 7.0: Low Frequency Propagation and Scattering*, Storming Media, 2002.

[16] Aubrey España, "Performance Estimation of Underwater Mine Countermeasure Operations," *ONR Contract N00014-12-G-0078/N00014-16-F-3003*.

[17] Richard O Duda, Peter E Hart, and David G Stork, *Pattern classification*, John Wiley & Sons, 2012.

[18] Jermaine Kennedy, "Multipass and Non-Concentric Target CSAS," *SERDP MR-2439 Final Report*, 2016.

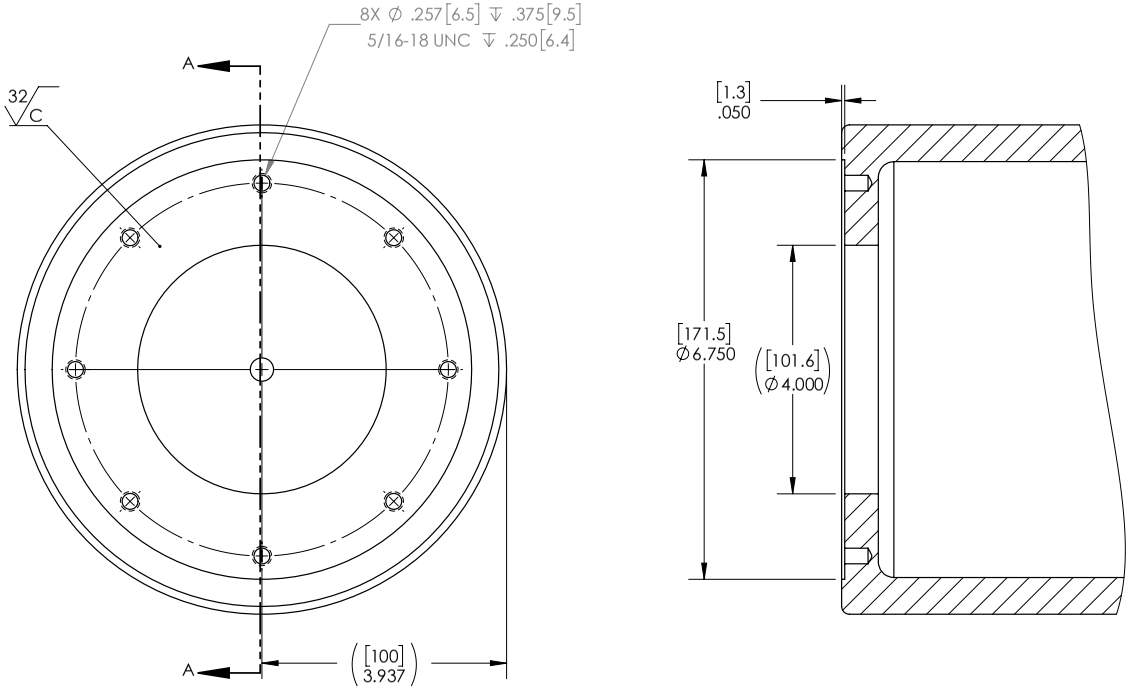
Appendix: Target Fabrication Blueprints

The following three pages contain the blueprints for the science target. The first page is a drawing of the flat endcap; the succeeding two pages are the side view and end view of the body, respectively. The post-deployment target measurement described in Section 5 resulted in slight adjustments in the model geometry (the target walls were reduced from the specification value of 15 mm to 13.8 mm, and the curvature of the pointed end was modified).

DWG NO: 77265 SHEET 1

CAD MAINTAINED: CHANGES SHALL BE INCORPORATED BY THE DESIGN ACTIVITY

REVISIONS				
ZONE	REV.	DESCRIPTION	DATE	APPROVED



SECTION A-A
SCALE 1 : 2

Originator	DYER 206-818-3675
Project No	483500
Task No	LEJ
Qty	10

UNLESS OTHERWISE NOTED

- APPLICABLE STANDARDS: ASME Y14.100
- SURFACE FINISH: 125RMS
- MACHINED FILLETS: R.015 MAX
- REMOVE ALL BURRS, BREAK SHARP EDGES .015 MAX
- DIMENSIONS & SURFACE FINISH ARE AFTER SPECIFIED COATINGS

THIRD ANGLE PROJECTION

DIMENSIONS ARE IN INCHES
WEIGHT IS IN POUNDS

TOLERANCES:
FRACTIONS: ± 1/16
.X: ± .030
.XX: ± .020
.XXX: ± .005
ANGLES: ± 1.0

MATERIAL:
MILD STEEL

WEIGHT: 10.2807

APPROVALS		DATE
CHIEF ENGR:		
SYS ENGR:		
QA/SAFETY:		
SR ENGR:	-	-
ENGINEER:	DYER	1.2616
CHECKED:	-	-
DRAWN:	DYER	5/1/2017

APPLIED PHYSICS LABORATORY
UNIVERSITY OF WASHINGTON
1013 NE 40TH ST. SEATTLE, WA. 98105

TITLE: TARGET CASTING MACHINING

PROJECT FOLDER:

SIZE: B | CAGE CODE: | DWG NO: 77265 | REV: 01

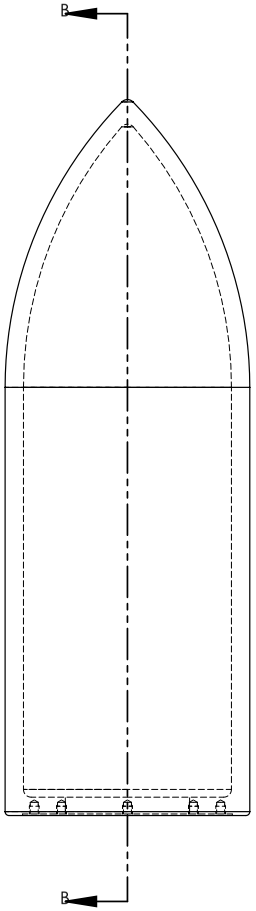
DO NOT SCALE DRAWING

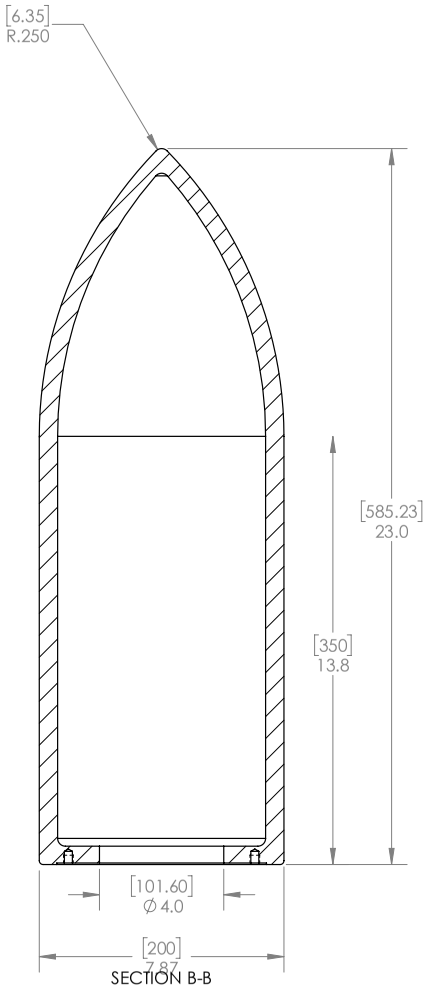
CAD FILE: 77265 TARGET CASTING MACHINING

SHEET 1 OF 2

NEXT ASSY	USED ON
APPLICATION DATA	

CAD MAINTAINED: CHANGES SHALL BE INCORPORATED BY THE DESIGN ACTIVITY





SECTION B-B

TARGET CASTING MACHINING

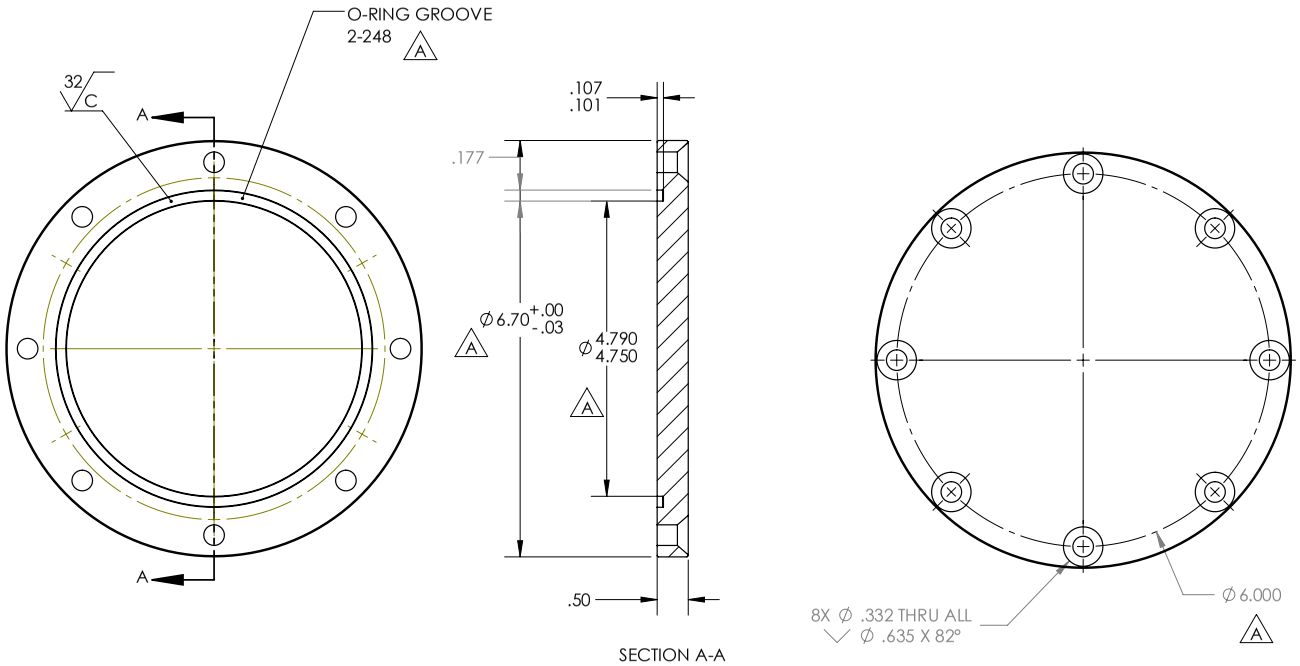
DYER 1.2616

APPLIED PHYSICS LABORATORY UNIVERSITY OF WASHINGTON	SIZE: B	CAGE CODE:	DWG NO: 77265	REV: 01
10.2807	77265 TARGET CASTING MACHINING			SHEET 2 OF 2

CAD MAINTAINED: CHANGES SHALL BE INCORPORATED BY THE DESIGN ACTIVITY

DWG NO: 77266 SHEET 1

REVISIONS				
ZONE	REV.	DESCRIPTION	DATE	APPROVED
	A	INCREASED DIAMETER TO MATCH TARGET CASTING	5/25/2017	
	B	DRAWING UPDATED TO REFLECT ACTUAL	5/25/2017	



Originator	DYER 206-818-3675
Project No	483500
Task No	LEJ
Qty	10

UNLESS OTHERWISE NOTED

- APPLICABLE STANDARDS: ASME Y14.100
- SURFACE FINISH: 125RMS
- MACHINED FILLETS: R.015 MAX
- REMOVE ALL BURRS, BREAK SHARP EDGES .015 MAX
- DIMENSIONS & SURFACE FINISH ARE AFTER SPECIFIED COATINGS

THIRD ANGLE PROJECTION

DIMENSIONS ARE IN INCHES
WEIGHT IS IN POUNDS

TOLERANCES:
FRACTIONS: ± 1/16
.X: ± .030
.XX: ± .020
.XXX: ± .005
ANGLES: ± 1.0

MATERIAL:
1018 COLD ROLLED

WEIGHT: 0.61

APPROVALS		DATE
CHIEF ENGR:		
SYS ENGR:		
QA/SAFETY:		
SR ENGR:	-	-
ENGINEER:	DYER	1.2616
CHECKED:	-	-
DRAWN:	DYER	5/1/2017

APPLIED PHYSICS LABORATORY
UNIVERSITY OF WASHINGTON
1013 NE 40TH ST. SEATTLE, WA. 98105

TITLE: TARGET ENDCAP

PROJECT FOLDER:

SIZE: B | CAGE CODE: | DWG NO: 77266 | REV: B

DO NOT SCALE DRAWING

CAD FILE: 77266 TARGET ENDCAP

SHEET 1 OF 1

NEXT ASSY	USED ON
APPLICATION DATA	