

FINAL REPORT

Sensor Phenomenology and Feature Development for Improved Sonar-based Detection & Classification of Underwater UXO

SERDP Project MM-1533

DECEMBER 2009

Eugene Lavelly
BAE Systems

This document has been approved for public release.



Strategic Environmental Research and
Development Program

Report Documentation Page

Form Approved
OMB No. 0704-0188

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

1. REPORT DATE DEC 2009		2. REPORT TYPE N/A		3. DATES COVERED -	
4. TITLE AND SUBTITLE Sensor Phenomenology and Feature Development for Improved Sonar-based Detection & Classification of Underwater UXO				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) BAE Systems				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited					
13. SUPPLEMENTARY NOTES The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT SAR	18. NUMBER OF PAGES 232	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

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.

Contents

1	Objective	7
2	Background	9
3	Materials and Methods	13
3.1	The spectral element method for poroelastic media	13
3.2	T-matrix analysis	13
3.2.1	Acoustic equation and Green function	13
3.2.2	Free-field T -matrix	14
3.2.3	Heterogeneous field T -matrix	16
3.2.4	T -matrix for multiple scatterers in a homogeneous medium	16
3.2.5	T -matrix for multiple scatterers in a heterogeneous medium	17
3.2.6	Basis functions in horizontally stratified media	18
3.2.7	T -matrix for a single homogeneous spherical scatterer	21
3.2.8	T -matrix for an array spheres in a homogeneous background	21
3.2.8.1	Array of small spheres	21
3.2.8.2	More general array of spheres	22
3.2.9	T -matrix for an array of scatterers in a layered background	23
3.2.9.1	Single large scatterer in an array of small scatterers	24
3.2.10	Transmitted and received signals	25
3.2.10.1	Special cases	26
3.2.11	Inverse scattering problems	28
3.2.11.1	Target parameter inversion in a homogeneous background	28
3.2.11.2	Buried target parameter inversion	29
3.2.11.3	Basis function expansions	30
3.2.11.3.1	Green function and plane wave expansions	30
3.2.11.3.2	Re-expansion coefficient identities	32
3.2.11.3.3	Re-expansion of outgoing wave basis functions	32
3.2.12	Green function coefficients for a three layer system	34
3.3	BOSS data processing methods	35
3.3.1	Signal processing	35
3.3.2	Near-field focusing	36
3.3.3	Approximate near-field focusing	36
3.3.4	Image construction	37
3.4	Adjoint tomography	39
3.5	Bayesian approach for BOSS classification analysis (Relevance Vector Machine)	43
4	Results and Discussion	46
4.1	Numerical simulations using the SEM method	46
4.2	Numerical simulation of BOSS data	46
4.2.1	Numerical simulations using the T-matrix method	46
4.2.2	Numerical simulations using the spectral element method	50
4.3	Buried Object Scanning Sonar (BOSS) data	61

4.3.1	BOSS data collections	61
4.3.2	2004 Disk BOSS Trial	61
4.3.3	2006 BMH Sea Trials	62
4.3.4	AUV FEST 2007 BMH Sea Trial	63
4.3.5	AUV FEST 2008 Trial	64
4.4	Summary of target/clutter data descriptors available for classification analysis . . .	65
4.5	Feature analysis of BOSS data from AUV FEST 2008	67
4.5.1	3D beamformed imagery examples from BOSS data collection	68
4.5.2	Histogram analysis of target/clutter feature distributions for AUV FEST 2008 data	69
4.5.3	Feature pair analysis of target/clutter features for AUV FEST 2008 data . .	71
4.5.4	Feature triplet analysis of target/clutter features for AUV FEST 2008 data .	73
4.6	Numerical results for BOSS classification analysis	75
5	Conclusions and Implications for Future Research and Implementation	99
A	Supporting Data	105
A.1	Beamformed imagery (2D)	106
A.2	Beamformed imagery (3D)	113
A.3	Histogram analysis of target/clutter feature distributions for AUV Fest 2008 data .	117
A.4	Feature pair analysis of target/clutter features for AUV Fest 2008 data	136
A.5	Feature triplet analysis of target/clutter features for AUV Fest 2008 data	146
A.6	Statistical realizations for classifier training	165
A.7	BOSS processing results	185
B	List of Scientific/Technical Publications	196
B.1	Papers	196
B.2	Technical abstract for SERDP/ESTCP Annual Symposium (2006)	196
B.3	Technical abstract for SERDP/ESTCP Annual Symposium (2007)	197
B.4	Technical abstract for SERDP/ESTCP Annual Symposium (2008)	198
B.5	Technical abstract for American Geophysical Union Meeting (Winter, 2008)	199

List of Tables

1	Poroelastic parameters considered in adjoint sensitivity formulation.	40
2	Mode wave supported in simulation examples including fundamental waves and their conversions.	40
3	Summary of target and clutter objects of BOSS data taken 2004-2007. These provide sufficient data to test algorithms for target versus clutter discrimination and for target classification.	63
4	Summary of target and clutter objects from AUV Fest 2006 collection for which BOSS data is available for classifier analysis.	65
5	Data for AUV Fest 2007 collection	66
6	Target/clutter data available for the AUV FEST 2008 BOSS collection.	66

7 This table describes in summary form the 50 feature classes derived for each 3D beamformed target or clutter detection. These features define the available components of an input feature vector for use with the relevance vector machine classifier. Here PDF and CDF indicate the probability and cumulative density functions, respectively. 67

8 Summary of feature subsets used for training relevance vector machine classifiers. From a given training set the RVM discovers the optimal subset of feature input vectors for use in a kernel-based classifier design. However, the design of the classifier is dependent on the dimension of the input feature vector. Use of all possible features (in this case, 50), typically leads to poor generalization. Therefore, we experimented with different subsets of features, and obtained the corresponding classifier designs and performance metrics. These subsets, chosen by intuition rather than an optimizer, are indicated in the columns of the table. Each row index corresponds to the feature index and is directly linked to the feature descriptors in Table 7. The columns labeled from 1 to 20 each indicate the feature subset used as as indicated by the bullets, and the bottom row indicates the Figure number for which the performance metrics of the corresponding classifier are plotted. 77

9 Data notation definitions for the BOSS data collections and processing results. . . . 185

List of Figures

1 (a) The scattered wavefield due to an acoustic source is comprised of target and environment returns. The *target scatter* in general consists of a complex superposition of wave modes including direct target scatter, internal target resonances, surface modes, and target-interface resonances all of which are mediated by the target shape, size, acoustic-elastic properties and the incident field (e.g. frequency content). The *environment scatter* include interface scatter, rough interface scatter and volume scatter due to heterogeneity in the fluid and the sediment. (b) Example tessellation of a prolate spheroid target. In high-frequency numerical approaches (e.g. Kirchhoff method) approximate solutions to the Helmholtz boundary integral equation are obtained on each facet. 9

2 SAS data collected at Lake Washington under DARPA-sponsorship for various targets and clutter objects as indicated in the figure. 11

3 This figure shows an example application of our SAS image formation processor to sonar data collected by Coastal Systems Station (Panama City, Florida) using a Northrop-Grumman array. This array contained 11 high frequency (180 kHz) elements and 14 low frequency (20 kHz) elements. For these experiments a series of objects (e.g. large cylinder, ladder, a truncated cone) were placed on a sandy bottom. **Left panel:** This shows the focused image for the high frequency data. A cylinder in the upper right gives a clear shadow. The image was formed using redundant phase centers (RPC) and the range migration algorithm (RMA), but an autofocus step was not required. **Right panel:** Low frequency (20 kHz) SAS image formed using RPC, RMA and autofocus. These results illustrate that objects and background reverberation image characteristics have strong frequency sensitivity. Adapted from [1]. 12

4 For the proposed work we will have access to Buried Object Scanning Sonar data. The sites consist of buried sea mines and clutter objects in various sand and mud environments off of Kaneohe, Hawaii and Panama City, FL. The low frequency data is ideal for buried target interrogation. 12

5 Schematic illustration of background and scattered (outgoing wave) fields. The former represent the field in the absence of the scatterer, while the latter represent the corrections to the field due to the presence of the scatterer. The two fields are expanded in a set of basis functions according to (8) and (9), and related via (11) through the T -matrix. 14

6 Schematic diagram of horizontally stratified geometry. Notation is defined in the text. 19

7 Schematic illustration of the measurement geometry for one large scatterer, in an array of small scatterers, buried in a layered background. 23

8 **Top panel:** Schema of the 2-D model dimensions and source-receiver geometry. The top surface is a free surface whereas the three remaining edges are absorbing boundaries. We also draw possible ray paths for the P-SV wavefield, where PI refers to the fast P wave, PII to the slow P wave and S to the S wave. See Table 2 for nomenclature summary. **Bottom panel:** Velocity seismogram (x-component). 41

9 Plot of the density-normalized P-SV_{PIS+SPI} banana-doughnut kernels for selected poroelastic quantities. The vertical and horizontal axes have the same arbitrary length units of Fig. 8, and refer the same length spans as depicted in that figure. (**top**) - shear modulus kernel, (**middle**) - Bulk modulus, (**bottom**) - Biot coefficient C [29, 4]. See [29] and [4] for details of kernel construction and for explicit kernels for each of the quantities in Table 1. 42

10 This figure shows the source receiver geometry used in the simulation study in which there are 20 receivers on each wing of the indicated AUV with 51 mm spacing, a spherical source displaced 721 mm from the centerline of the receivers in the long dimension of the AUV, and for which the first receiver on each wing is 235 mm from the long-axis centerline of the AUV. This configuration corresponds to the actual experiment geometry used by BOSS in the more recent AUV FEST data collects. 47

11 Resonant spectrum for spherical harmonic degrees $0 \leq l \leq 4$ for a thin elastic spherical shell with radius $R = 50$ cm, and thin shell thickness given by $\Delta R = 10$ cm. The resonance spacing depend on the sphere size, sphere elastic properties, and $\Delta R/R$ 48

12 Time series for single source position and receiver shown Thin-shell target with radius $r = 25$ cm, and shell thickness $\Delta r = 5$ cm. The ringing observed after the first main arrival corresponds to overtones from the resonant mode spectrum, similar to those seen in Fig. 11. A Fourier analysis of this time series would reveal the spectrum, and provide features for use in classification. 49

13 Thicker shell (left) has radius $R = 12$ cm, and shell thickness $\Delta R = 7$ cm, while the thinner shell (right) has values $R = 12$ cm, and shell thickness $\Delta R = 6$ cm. Detailed spectra provide features for discrimination 49

14 Synthesized Time Series for 13 Randomly Distributed Solid and Thin-shell Spheres. All possible multiple interactions are accounted for in this simulation. 50

15 This figure shows the beamformed image for scattering due to a single, thin-shell spherical target, a single source position, and for all 40 receivers shown in Fig. 10. The input time series consists of a center frequency of 8000 Hz, and the source signal has 12,000 Hz bandwidth. The result shows that the scattering intensity is peaked in the true location of the target. However, in constructing this beamformed image, only one source position was used. This fact in combination with the finite bandwidth of the source location is the origin of the smile-like spread of the intensity. The use of additional source positions (as, in fact, is commonly done with BOSS deployments) would yield sharper focus since all of these contributions are coherently added, and volumetric regions containing no scatterers will tend to have destructive interference from the wavefields summed over the sources and the receivers. If the spatial coverage is skewed or non-symmetric, there will tend to be some geometric distortion. However, even in the case of perfect and complete spatial coverage, the scattering location will tend to be a “ball” rather than a point due to the finite bandwidth. Thus, in practice, one might use 100 source positions in the image construction, and each wavefield recording is acquired as the BOSS platform moves at speed 2 m/s and with a source pulse repetition frequency (PRF) of 20 Hz. In this way, beamforming over multiple source positions mitigates against the scattering location ambiguity. 51

16 Simulation of BOSS-like experiment data for a buried metal target (yellow rectangle) in three different sediment layers. The source is indicated by the yellow cross, and the receiver array is indicated by the yellow circles. 53

17 This figure is one of numerous benchmarks we have devised to validate the spectral element method theoretical formulation and numerical implementation. Validation is achieved by comparison to an available analytical solution. This simulation of wave propagation is for a model consisting of two homogeneous poroelastic layers with continuous bulk and shear moduli and discontinuous porosity, as tabulated in the upper right of the figure. The source (cross) is located at the cross and the two receivers are indicated by the circles. The top is a free surface and the remaining three edges are absorbing boundaries. **Upper left:** Snapshot of the vertical-component displacement at $t = 0.9$ s. The direct fast P (a), the reflected fast P (b), the reflected fast P-to-S and the fast P-to-slow P converted (c) waves (which overlap because they have similar wave speeds) can be observed in the upper layer, together with the direct slow P (d), the reflected slow P (e), the reflected slow P-to-S converted (f), and the reflected slow P-to-fast P converted (g) waves. We also observe the reflected fast P wave due to the free surface (h). In the lower layer, the transmitted fast P (i), fast P-to-S and fast P-to-slow P converted (j) waves (which again overlap because they have similar wave speeds) can clearly be identified, together with the transmitted slow P and slow P-to-S converted (k) and slow P-to-fast P converted (l) waves. There are some weak spurious reflections from the absorbing boundary at $x = 0$. **Lower left and lower right:** Vertical-component velocity seismograms at receivers 1 and 2 (Spectral Element Simulation: solid black line, analytical solution: dashed red line). We use domain composition to accommodate the first-order discontinuity in porosity. 54

18 The phase histories (without range compression) for a single channel of the 40-channel BOSS array as the platform tracks over a partially buried, cylindrical target. This data was acquired at AUV FEST 2007 and is one several data samples being used to validate the high-fidelity simulation method developed under this effort. 55

19 Various views of BOSS beamformed imagery for a data cube containing a target scattering response from the AUV FEST 2007 collection. **Upper left:** Top view of a target object shown as the maximum intensity projection (MIP) $P_{xy}(x, y)$ (see eq. 114). **Upper right:** Side view of a target object showing the MIP $P_{yz}(y, z)$ (see eq. 115). **Lower left:** Front view of a target object showing the MIP $P_{xz}(x, z)$ (see eq. 115). 56

20 The model dimensions for simulation are $16 \times 2.5 \times 6$ m. This figure shows a 3-D view of the model with compressional wave speed distribution. 57

21 Side view of the model, where the target in red can be seen. 57

22 Top view of the model showing two different positions of source-receivers vehicle. 58

23 Source waveform represented as a chirp signal sweeping linearly from 2 kHz to 20 kHz. This figures shows the comparison between the actual BOSS waveform for the data (labeled AST), and our simple model for this signal. 58

24 Time series at one receiver over 43 pings. The BOSS platform moves along the target. The background signal has been eliminated. The buried object signature is the only visible signal. 59

25 SAS Beamforming realized by Kent Harbaugh from Applied Signal Technology, Inc. using our 3-D numerical results: this shows a data slice parallel with the wing. 60

26 St. Andrew Bay 10-m water depth sand site data collection from 2006. 62

27 St. Andrew Bay 10-m water depth mud site data collection from 2006. Buried targets are indicated by the red-filled triangles. 62

28 Yankee test site target locations from the 2007 collection. 64

29 Various views of BOSS beamformed imagery for a data cube containing the scattering response for a cylinder shaped object from the AUV FEST 2007 collection. **Upper left:** Top view of a target object shown as the maximum intensity projection (MIP) $P_{xy}(x, y)$ (see eq. 114). **Upper right:** Side view of a target object showing the MIP $P_{yz}(y, z)$ (see eq. 115). **Lower left:** Front view of a target object showing the MIP $P_{xz}(x, z)$ (see eq. 115). 68

30 Distribution histograms for the *geometric center* x coordinate (top panel), y coordinate (top panel) and z coordinate (bottom panel) feature types. In each case the target and clutter features are shown on the left hand and right hand sides, respectively. The features were derived from 3D beamform products using data acquired by the BOSS platform at the AUV Fest 2008. A total of 132 clutter items and 89 targets were processed in generating these distributions. 70

31 Features derived from the targets and clutter in this data set from the BOSS AUV FEST 2008 collection are indicated by the circles and plus signs, respectively. Scatter plots for various feature pairs are shown. **Upper left:** *ellipsoid radius A versus ellipsoid radius B*. **Upper right:** *ellipsoid radius A versus ellipsoid length to width ratio A/B*. **Middle left:** *ellipsoid radius A versus the voxel energy density*. **Middle right:** *ellipsoid radius A versus the peak intensity*. **Lower left:** *ellipsoid radius A versus the first bin of the 3-bin PDF of the relative intensity*. **Lower right:** *ellipsoid radius A versus the third bin of the 3-bin PDF of the relative intensity*. 72

32 3D scatter plots showing training points and optimal relevance vectors for various feature triplets. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors corresponding to known targets are given by the green circles. In addition, feature triplets circled in red are determined to be *relevance vectors* as determined for RVM analysis for the feature subset (4, 5, 20, 26, 29, 30, 32) as described in Sec. 4.6 and the case 1 subset (see Table 8). **Upper panel:** Feature triplet *A, B (best fit ellipsoid radii) and Length to width ratio A/B*. **Middle panel:** Feature triplet *A, B (best fit ellipsoid radii) and Voxel energy density*. **Lower panel:** Feature triplet *A, B (best fit ellipsoid radii) and peak intensity*. 74

33 Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets. 76

34 RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 1 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets. 79

35 RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 2 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets. 80

36 RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 3 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets. 81

37 RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 4 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets. 82

38 RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 5 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets. 83

39 RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 6 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets. 84

40 RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 7 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets. 85

41 RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 8 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets. 86

42 RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 9 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets. 87

43 RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 10 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets. 88

44 RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 11 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets. 89

45 RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 12 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets. 90

46 RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 13 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets. 91

47 RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 14 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets. 92

48 RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 15 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets. 93

49 RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 16 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets. 94

50 RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 17 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets. 95

51 RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 18 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets. 96

52 RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 19 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets. 97

53 RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 20 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets. 98

54 Various views of BOSS beamformed imagery for a data cube containing a target scattering response from the AUV FEST 2007 collection. **Upper left:** Top view of a target object shown as the maximum intensity projection (MIP) $P_{xy}(x, y)$ (see eq. 114). **Upper right:** Side view of a target object showing the MIP $P_{yz}(y, z)$ (see eq. 115). **Lower left:** Front view of a target object showing the MIP $P_{xz}(x, z)$ (see eq. 115). 106

55 Various views of BOSS beamformed imagery for a data cube containing a target scattering response from the AUV FEST 2007 collection. **Upper left:** Top view of a target object shown as the maximum intensity projection (MIP) $P_{xy}(x, y)$ (see eq. 114). **Upper right:** Side view of a target object showing the MIP $P_{yz}(y, z)$ (see eq. 115). **Lower left:** Front view of a target object showing the MIP $P_{xz}(x, z)$ (see eq. 115). 107

56 Various views of BOSS beamformed imagery for a data cube containing a target scattering response from the AUV FEST 2007 collection. **Upper left:** Top view of a target object shown as the maximum intensity projection (MIP) $P_{xy}(x, y)$ (see eq. 114). **Upper right:** Side view of a target object showing the MIP $P_{yz}(y, z)$ (see eq. 115). **Lower left:** Front view of a target object showing the MIP $P_{xz}(x, z)$ (see eq. 115). 108

57 Various views of BOSS beamformed imagery for a data cube containing a target scattering response from the AUV FEST 2007 collection. **Upper left:** Top view of a target object shown as the maximum intensity projection (MIP) $P_{xy}(x, y)$ (see eq. 114). **Upper right:** Side view of a target object showing the MIP $P_{yz}(y, z)$ (see eq. 115). **Lower left:** Front view of a target object showing the MIP $P_{xz}(x, z)$ (see eq. 115). 109

58 Various views of BOSS beamformed imagery for a data cube containing a target scattering response from the AUV FEST 2007 collection. **Upper left:** Top view of a target object shown as the maximum intensity projection (MIP) $P_{xy}(x, y)$ (see eq. 114). **Upper right:** Side view of a target object showing the MIP $P_{yz}(y, z)$ (see eq. 115). **Lower left:** Front view of a target object showing the MIP $P_{xz}(x, z)$ (see eq. 115). 110

59 Various views of BOSS beamformed imagery for a data cube containing a target scattering response from the AUV FEST 2007 collection. **Upper left:** Top view of a target object shown as the maximum intensity projection (MIP) $P_{xy}(x, y)$ (see eq. 114). **Upper right:** Side view of a target object showing the MIP $P_{yz}(y, z)$ (see eq. 115). **Lower left:** Front view of a target object showing the MIP $P_{xz}(x, z)$ (see eq. 115). 111

60 Various views of BOSS beamformed imagery for a data cube containing a target scattering response from the AUV FEST 2007 collection. **Upper left:** Top view of a target object shown as the maximum intensity projection (MIP) $P_{xy}(x, y)$ (see eq. 114). **Upper right:** Side view of a target object showing the MIP $P_{yz}(y, z)$ (see eq. 115). **Lower left:** Front view of a target object showing the MIP $P_{xz}(x, z)$ (see eq. 115). 112

61 3D isocontour of the subvolume for an iso-intensity value given by 1/2 of the maximum intensity value of the data cube, corresponding to the target in Fig. 54. . 113

62 3D isocontour of the subvolume for an iso-intensity value given by 1/2 of the maximum intensity value of the data cube, corresponding to the target in Fig. 55. . 114

63 3D isocontour of the subvolume for an iso-intensity value given by 1/2 of the maximum intensity value of the data cube, corresponding to the target in Fig. 56. . 114

64 3D isocontour of the subvolume for an iso-intensity value given by 1/2 of the maximum intensity value of the data cube, corresponding to the target in Fig. 57. . 115

65 3D isocontour of the subvolume for an iso-intensity value given by 1/2 of the maximum intensity value of the data cube, corresponding to the target in Fig. 58. . 115

66 3D isocontour of the subvolume for an iso-intensity value given by 1/2 of the maximum intensity value of the data cube, corresponding to the target in Fig. 59. . 116

67 3D isocontour of the subvolume for an iso-intensity value given by 1/2 of the maximum intensity value of the data cube, corresponding to the target in Fig. 60. . 116

68 For this feature type we fit a volume associated with a clutter or target item to the *best fitting ellipsoid* characterized by the three principal radii A (top panel), B (middle panel) and C (bottom panel). In each case the target and clutter features are shown on the left hand and right hand sides, respectively. The features were derived from 3D beamform products using data acquired by the BOSS platform at the AUV Fest 2008. A total of 132 clutter items and 89 targets were processed in generating these distributions. 118

- 69 This feature type is given by the unit vector \hat{u}_A describing the direction of the principal radius A of the best fitting ellipsoid volume. The x -component, y -component and z -components are shown, respectively, in the top, middle and bottom panels. In each case the target and clutter features are shown on the left hand and right hand sides, respectively. The features were derived from 3D beamform products using data acquired by the BOSS platform at the AUV Fest 2008. A total of 132 clutter items and 89 targets were processed in generating these distributions. 119
- 70 This feature type is given by the unit vector \hat{u}_B describing the direction of the principal radius B of the best fitting ellipsoid volume. The x -component, y -component and z -components are shown, respectively, in the top, middle and bottom panels. In each case the target and clutter features are shown on the left hand and right hand sides, respectively. The features were derived from 3D beamform products using data acquired by the BOSS platform at the AUV Fest 2008. A total of 132 clutter items and 89 targets were processed in generating these distributions. 120
- 71 This feature type is given by the unit vector \hat{u}_C describing the direction of the principal radius C of the best fitting ellipsoid volume. The x -component, y -component and z -components are shown, respectively, in the top, middle and bottom panels. In each case the target and clutter features are shown on the left hand and right hand sides, respectively. The features were derived from 3D beamform products using data acquired by the BOSS platform at the AUV Fest 2008. A total of 132 clutter items and 89 targets were processed in generating these distributions. 121
- 72 This feature type is given by the cross-sectional areas of the best fitting ellipsoids to the thresholded intensity blobs identified as target/clutter items in the BOSS data. These cross-sectional areas are given as $\pi AB, \pi AC$ and πBC , and are shown, respectively in the top, middle and bottom panels. In each case the target and clutter features are shown on the left hand and right hand sides, respectively. The features were derived from 3D beamform products using data acquired by the BOSS platform at the AUV Fest 2008. A total of 132 clutter items and 89 targets were processed in generating these distributions. 122
- 73 This feature type is given by the length to width ratios $A/C, A/B$ and B/C where A, B and C are the principal radii of the best fitting ellipsoids. These ratios are shown, respectively, in the top, middle and bottom panels. In each case the target and clutter features are shown on the left hand and right hand sides, respectively. The features were derived from 3D beamform products using data acquired by the BOSS platform at the AUV Fest 2008. A total of 132 clutter items and 89 targets were processed in generating these distributions. 123
- 74 This feature type is given by the ellipsoidal volume πABC of the best fitting ellipsoids to each of the thresholded intensity blobs identified as target/clutter items in the BOSS data. The target and clutter distributions for this features are shown on the left hand and right hand sides, respectively. A total of 132 clutter items and 89 targets were processed in generating these distributions. 124

75 This feature type is given by the voxel volume to each of the thresholded intensity blobs identified as target/clutter items in the BOSS data. The target and clutter distributions for this features are shown on the left hand and right hand sides, respectively. A total of 132 clutter items and 89 targets were processed in generating these distributions. 124

76 This feature type is the total energy of the voxel volume corresponding to an identified target or clutter item from the BOSS data. The target and clutter distributions for this features are shown on the left hand and right hand sides, respectively. A total of 132 clutter items and 89 targets were processed in generating these distributions. 125

77 This feature type is the total energy of the voxels contained within the best-fitting ellipsoidal volumes corresponding to an identifier target or clutter items from the BOSS data. The target and clutter distributions for this features are shown on the left hand and right hand sides, respectively. A total of 132 clutter items and 89 targets were processed in generating these distributions. 125

78 This feature type is the voxel energy density for voxels contained within the best-fitting ellipsoidal volumes corresponding to an identified target or clutter items from the BOSS data. The target and clutter distributions for this features are shown on the left hand and right hand sides, respectively. A total of 132 clutter items and 89 targets were processed in generating these distributions. 126

79 This feature type is the voxel energy standard deviation for voxels contained within the best-fitting ellipsoidal volumes corresponding to an identified target or clutter items from the BOSS data. The target and clutter distributions for this features are shown on the left hand and right hand sides, respectively. A total of 132 clutter items and 89 targets were processed in generating these distributions. 127

80 This feature type is the voxel energy standard deviation normalized by mean contained within the best-fitting ellipsoidal volumes corresponding to an identified target or clutter items from the BOSS data. The target and clutter distributions for this features are shown on the left hand and right hand sides, respectively. A total of 132 clutter items and 89 targets were processed in generating these distributions. 127

81 This feature type is the peak intensity voxel contained within the total voxel volume identified with a target or clutter item from the AUV FEST 2008 BOSS data. The target and clutter distributions for this features are shown on the left hand and right hand sides, respectively. A total of 132 clutter items and 89 targets were processed in generating these distributions. 128

82 The feature type shown here is the 3-bin PDF of relative intensity in which the top, middle and bottom panels correspond, respectively, to bins 1, 2 and 3. The intensity voxels were chosen from those contained within the best fitting ellipsoidal volume associated with each target or clutter item. In each case the target and clutter features are shown on the left hand and right hand sides, respectively. The features were derived from 3D beamform products using data acquired by the BOSS platform at the AUV Fest 2008. A total of 132 clutter items and 89 targets were processed in generating these distributions. 129

83 The feature type shown here are bins 1 to 3 of the 9–bin PDF of relative intensity in which the top, middle and bottom panels correspond, respectively, to bins 1, 2 and 3. The intensity voxels were chosen from those contained within the best fitting ellipsoidal volume associated with each target or clutter item. In each case the target and clutter features are shown on the left hand and right hand sides, respectively. The features were derived from 3D beamform products using data acquired by the BOSS platform at the AUV Fest 2008. A total of 132 clutter items and 89 targets were processed in generating these distributions. 130

84 The feature type shown here are bins 4 to 6 of the 9–bin PDF of relative intensity in which the top, middle and bottom panels correspond, respectively, to bins 4, 5 and 6. The intensity voxels were chosen from those contained within the best fitting ellipsoidal volume associated with each target or clutter item. In each case the target and clutter features are shown on the left hand and right hand sides, respectively. The features were derived from 3D beamform products using data acquired by the BOSS platform at the AUV Fest 2008. A total of 132 clutter items and 89 targets were processed in generating these distributions. 131

85 The feature type shown here are bins 7 to 9 of the 9–bin PDF of relative intensity in which the top, middle and bottom panels correspond, respectively, to bins 7, 8 and 9. The intensity voxels were chosen from those contained within the best fitting ellipsoidal volume associated with each target or clutter item. In each case the target and clutter features are shown on the left hand and right hand sides, respectively. The features were derived from 3D beamform products using data acquired by the BOSS platform at the AUV Fest 2008. A total of 132 clutter items and 89 targets were processed in generating these distributions. 132

86 The feature type shown here are bins 1 to 3 of the 9–bin cumulative distribution function (CDF) of relative intensity in which the top, middle and bottom panels correspond, respectively, to bins 1, 2 and 3. The intensity voxels were chosen from those contained within the best fitting ellipsoidal volume associated with each target or clutter item. In each case the target and clutter features are shown on the left hand and right hand sides, respectively. The features were derived from 3D beamform products using data acquired by the BOSS platform at the AUV Fest 2008. A total of 132 clutter items and 89 targets were processed in generating these distributions. 133

87 The feature type shown here are bins 4 to 6 of the 9–bin cumulative distribution function (CDF) of relative intensity in which the top, middle and bottom panels correspond, respectively, to bins 4, 5 and 6. The intensity voxels were chosen from those contained within the best fitting ellipsoidal volume associated with each target or clutter item. In each case the target and clutter features are shown on the left hand and right hand sides, respectively. The features were derived from 3D beamform products using data acquired by the BOSS platform at the AUV Fest 2008. A total of 132 clutter items and 89 targets were processed in generating these distributions. 134

88 The feature type shown here are bins 7 to 9 of the 9–bin cumulative distribution function (CDF) of relative intensity in which the top, middle and bottom panels correspond, respectively, to bins 7, 8 and 9. The intensity voxels were chosen from those contained within the best fitting ellipsoidal volume associated with each target or clutter item. In each case the target and clutter features are shown on the left hand and right hand sides, respectively. The features were derived from 3D beamform products using data acquired by the BOSS platform at the AUV Fest 2008. A total of 132 clutter items and 89 targets were processed in generating these distributions. 135

89 Scatter plot for the features *ellipsoid radius B* versus the *length to width ratio A/B* for the BOSS AUV FEST 2008 data. Features derived from the targets and clutter in this data set are indicated by the circles and plus signs, respectively. 136

90 Scatter plot for the features *ellipsoid radius B* versus the *voxel energy density* for the BOSS AUV FEST 2008 data. Features derived from the targets and clutter in this data set are indicated by the circles and plus signs, respectively. 137

91 Scatter plot for the features *ellipsoid radius B* versus the *peak intensity* for the BOSS AUV FEST 2008 data. Features derived from the targets and clutter in this data set are indicated by the circles and plus signs, respectively. 137

92 Scatter plot for the features *ellipsoid radius B* versus the *first bin of the 3-bin PDF of the relative intensity* for the BOSS AUV FEST 2008 data. Features derived from the targets and clutter in this data set are indicated by the circles and plus signs, respectively. 138

93 Scatter plot for the features *ellipsoid radius B* versus the *third bin of the 3-bin PDF of the relative intensity* for the BOSS AUV FEST 2008 data. Features derived from the targets and clutter in this data set are indicated by the circles and plus signs, respectively. 138

94 Scatter plot for the features *length to width ratio A/B* versus the *voxel energy density* for the BOSS AUV FEST 2008 data. Features derived from the targets and clutter in this data set are indicated by the circles and plus signs, respectively. . . . 139

95 Scatter plot for the features *length to width ratio A/B* versus the *peak intensity* for the BOSS AUV FEST 2008 data. Features derived from the targets and clutter in this data set are indicated by the circles and plus signs, respectively. 140

96 Scatter plot for the features *length to width ratio A/B* versus the *first bin of the 3-bin PDF of the relative intensity* for the BOSS AUV FEST 2008 data. Features derived from the targets and clutter in this data set are indicated by the circles and plus signs, respectively. 140

97 Scatter plot for the features *length to width ratio A/B* versus the *third bin of the 3-bin PDF of the relative intensity* for the BOSS AUV FEST 2008 data. Features derived from the targets and clutter in this data set are indicated by the circles and plus signs, respectively. 141

98 Scatter plot for the features *voxel energy density* versus the *peak intensity* for the BOSS AUV FEST 2008 data. Features derived from the targets and clutter in this data set are indicated by the circles and plus signs, respectively. This feature appears to be a useful discriminant for target versus clutter for at least some targets. It is clear that simultaneously strong peak intensities and voxel energy densities are strongly coordinated to targets. This is because actual targets provide a strong scattering contrast relative to natural and man-made, non-target clutter. 142

99 Scatter plot for the features *voxel energy density* versus the *first bin of the 3-bin PDF of the relative intensity* for the BOSS AUV FEST 2008 data. Features derived from the targets and clutter in this data set are indicated by the circles and plus signs, respectively. 142

100 Scatter plot for the features *voxel energy density* versus the *third bin of the 3-bin PDF of the relative intensity* for the BOSS AUV FEST 2008 data. Features derived from the targets and clutter in this data set are indicated by the circles and plus signs, respectively. 143

101 Scatter plot for the features *peak intensity* versus the *first bin of the 3-bin PDF of the relative intensity* for the BOSS AUV FEST 2008 data. Features derived from the targets and clutter in this data set are indicated by the circles and plus signs, respectively. 144

102 Scatter plot for the features *peak intensity* versus the *third bin of the 3-bin PDF of the relative intensity* for the BOSS AUV FEST 2008 data. Features derived from the targets and clutter in this data set are indicated by the circles and plus signs, respectively. 145

103 Scatter plot for the features *first bin of the 3-bin PDF of the relative intensity* versus the *third bin of the 3-bin PDF of the relative intensity* for the BOSS AUV FEST 2008 data. Features derived from the targets and clutter in this data set are indicated by the circles and plus signs, respectively. 145

104 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *A (ellipsoid radius), B (ellipsoid radius) and first bin in 3-bin probability distribution of relative intensity*, The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles. 146

105 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *A (ellipsoid radii), B (ellipsoid radii) and last bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles. 147

106 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *A (ellipsoid radii), Length to width ratio A/B and Voxel energy density. and first bin in 3-bin probability distribution of relative intensity*, The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles. 147

107 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *A (ellipsoid radii), Length to width ratio A/B, and Peak intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles. 148

108 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *A (ellipsoid radii), Length to width ratio A/B, and first bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles. 148

109 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *A (ellipsoid radii), Length to width ratio A/B, and last bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles. 149

110 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *A (ellipsoid radii), Voxel energy density and Peak intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles. 150

111 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *A (ellipsoid radii), Voxel energy density and first bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles. 151

112 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *A (ellipsoid radii), Voxel energy density and last bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles. 151

113 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *A (ellipsoid radii), Peak intensity and first bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles. 152

114 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *A (ellipsoid radii), Peak intensity and last bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles. 152

115 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *A (ellipsoid radii), first bin in 3-bin probability distribution of relative intensity and last bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles. 153

116 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *B (ellipsoid radii), Length to width ratio A/B and Voxel energy density*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles. . 153

117 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *B (ellipsoid radii), Length to width ratio A/B and Peak intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles. . 154

118 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *B (ellipsoid radii), Length to width ratio A/B and first bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles. 155

119 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *B (ellipsoid radii), Length to width ratio A/B and last bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles. 155

120 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *B (ellipsoid radii), Voxel energy density and Peak intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles. 156

121 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *B (ellipsoid radii), Voxel energy density and first bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles. 156

122 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *B (ellipsoid radii)*, *Voxel energy density* and *last bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles. 157

123 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *B (ellipsoid radii)*, and *Peak intensity* and *first bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles. 157

124 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *B (ellipsoid radii)*, and *Peak intensity* and *last bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles. 158

125 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *B (ellipsoid radii)*, *first bin in 3-bin probability distribution of relative intensity* and *last bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles. 158

126 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *Length to width ratio A/B*, *Voxel energy density* and *Peak intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles. 159

127 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *Length to width ratio A/B*, *Voxel energy density* and *first bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles. 159

128 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *Length to width ratio A/B*, *Voxel energy density* and *last bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles. 160

129 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *Length to width ratio A/B*, and *Peak intensity*. and *first bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles. 160

130 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *Length to width ratio A/B*, and *Peak intensity*. and *last bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles. 161

131 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *Length to width ratio A/B*, *first bin in 3-bin probability distribution of relative intensity* and *last bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles. 161

132 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *Voxel energy density*, *Peak intensity* and *first bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles. 162

133 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *Voxel energy density*, *Peak intensity* and *last bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles. 163

134 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *Voxel energy density*, *first bin in 3-bin probability distribution of relative intensity* and *last bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles. 163

135 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *Peak intensity*, *first bin in 3-bin probability distribution of relative intensity* and *last bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles. 164

- 136 Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets. 166
- 137 Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets. 167
- 138 Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets. 168

- 139 Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets. 169
- 140 Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets. 170
- 141 Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets. 171

- 142 Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets. 172
- 143 Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets. 173
- 144 Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets. 174

- 145 Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets. 175
- 146 Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets. 176
- 147 Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets. 177

- 148 Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets. 178
- 149 Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets. 179
- 150 Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets. 180

- 151 Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets. 181
- 152 Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets. 182
- 153 Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets. 183

154 Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets. 184

155 This figure shows P_{xy} (maximum value over all z planes at each (x,y) point) for an entire run at the mud site. Total water depth is approximately 10 m. 185

156 Similar to Fig. 155 but for a different run. 186

157 This figure shows P_{xy} (maximum value over all z planes at each (x,y) point) for an entire run at the sand site. The platform was at a ‘high’ altitude above the seabed (typically 5 to 6 m). Total water depth is approximately 10 m. Exploded views of this run are shown in Figs. 158 and 159. 186

158 Exploded view of Fig. 157 for the *first half* of the run. 186

159 Exploded view of Fig. 157 for the *second half* of the run. 187

160 This figure shows P_{xy} (maximum value over all z planes at each (x,y) point) for an entire run at the sand site. The platform was at a ‘low’ altitude above the seabed (typically 2 to 3 m). Total water depth is approximately 10 m. Exploded views of this run are shown in Figs. 161 and 162. 187

161 Exploded view of Fig. 160 for the *first half* of the run. 187

162 Exploded view of Fig. 160 for the *second half* of the run. 188

163 Similar to Fig. 160 but for a second run. 188

164 Exploded view of Fig. 163 for the *first half* of the run. 188

165 Exploded view of Fig. 163 for the *second half* of the run. 189

166 This figure shows a detailed view of $B(\mathbf{x}_F)$ (eq. 111) for a particular target at the mud site. This view shows slices on the z -plane at depth indices 16 and 17. 189

167 Similar to Fig. 166 but for depth slices 18 and 19. 190

168 Similar to Fig. 166 but for depth slices 20 and 21. 191

169 Similar to Fig. 166 but for depth slices 22 and 23. 192

170 Similar to Fig. 166 but for depth slices 24 and 25. 193

171 Similar to Fig. 166 but for depth slices 26 and 27. 194

172 Similar to Fig. 166 but for depth slices 28 and 29. 195

List of Acronyms

AST	Applied Signal Technology, Inc.
ATR	Automatic Target Recognition
BOSS	Buried Object Scanning Sonar
CONOPS	Concepts of Operation
CUBIT™	Sandia National Laboratory Geometry and Mesh Generation Toolkit
FM	Frequency Modulated
MIP	Maximum Intensity Projection
MVDR	Minimum Variance Distortionless Response
ONR	Office of Naval Research
RMA	Range Migration Algorithm
RPC	Redundant Phase Centers
RVM	Relevance Vector Machine
SEM	Spectral Element Method
SAB	St. Andrews Bay (Florida)
SAS	Synthetic Aperture Sonar
SAX99	Sediment Acoustics Experiment 1999
SAX04	Sediment Acoustics Experiment 2004
SERDP	Strategic Environmental Research and Development Program
SGS	Sequential Gaussian Simulation
SlowSAS	Slow Speed Synthetic Aperture Sonar
SS	Sandy Site
UUV	Unmanned Underwater Vehicle
UXO	Unexploded Ordnance

Acknowledgements

BAE Systems Advanced Information Technologies and the principal investigator (Eugene Lively) would like to acknowledge Dr. Anne Andrews and Dr. Herb Nelson for their guidance of this effort (Contract No. W912HQ-06-C-0006). Dr. Andrews served as the munitions area PM during the first half of this contract, and Dr. Nelson then succeeded her. In addition, the PI would like to thank the subcontractors to this effort for the substantial contributions that they provided. Dr. Matthew Nelson and Dr. Kent Harbaugh from Applied Signal Technology, Inc. supported BOSS data processing and classification analysis. Professor Jeroen Tromp and Dr. Christina Morency from Princeton University developed the 2D and 3D poroelastic spectral element simulation method under this research effort. Dr. Peter Weichman (BAE Systems AIT) supported the T-matrix analysis reported here. I would like to thank Dr. Steve Schock for making the BOSS data used in this study available to Applied Signal Technology, Inc. I would like to thank Mr. Tony Buschur and Mr. Peter Knowles and the other SERDP support staff for their help in adhering to SERDP program requirements. I would like to Ms. Daphne Curran, Ms. Katherine Anthony for their internal financial oversight of this program and Dr. Gil Ettinger for his management oversight. Finally, I would like to thank SERDP for supporting research in the underwater munitions area, and SERDP's Scientific Advisory Board members for their inputs at the various program reviews.

Abstract

This effort has examined the problem of detection and classification of buried munitions in underwater environments. We have focused on the use of low frequency sonar since high frequency acoustic waves are strongly attenuated by sediments. The focus of this effort has been to process low-frequency data collected from the Buried Object Scanning Sonar (BOSS) into 3D imagery using beamforming, and to develop target/clutter classifiers that use 3D features extracted from this imagery. The principal sonar data sources are BOSS deployments at various shallow water sites. Morphological processing was applied to the derived imagery for feature input into a relevance vector machine classifier. Since ground truth was available, it was possible to compute performance metrics in the form of ROC curves. To enable a systematic understanding of the influence of the environment on target responses, we have developed a poroelastic spectral element method for BOSS data simulations using 2D and 3D models. The classification results establish that buried targets have a high probability of detection with the Buried Object Scanning Sonar. However, features from target imagery responses are easily confused with those of clutter and munitions debris due to their incomplete separation. Small subsets of possible imagery features show the best performance, and various examples are shown. We provide a theoretical development for the estimation of structural acoustic resonance features from BOSS-like data. Future classification performance gains with the sonar modality will likely rely on the combined use of imagery- and resonance-based features.

Executive Summary

The principal objectives of this research effort are to evaluate the effectiveness of synthetic aperture sonar (SAS¹) for detection and classification of underwater munitions, characterize the influence of sediment heterogeneity on buried target detection and discrimination, identify processing methods to improve performance of detection systems, and provide recommendations for future algorithm and system development. In the remainder of this summary we provide a description of the environmental problem addressed, the scientific questions that we have explored, the cumulative results received to date, and potential future applications of the research.

Environmental problems addressed Military training and weapons testing activities in the past have left the legacy of UXO (Unexploded Ordnance) at a number of sites. This problem is of even greater concern for those sites designated for base realignment and closure and at Formerly Used Defense Sites. Particularly difficult is the characterization and remediation of those sites where UXO is found in underwater environments. Many active and former military installations have ordnance ranges and training areas that include adjacent water environments. (e.g., ponds, lakes, rivers, estuaries, and coastal ocean areas). Wartime activities, dumping, and accidents have also generated significant UXO contamination in coastal and inland waters. Dredging projects frequently encounter UXO, and potential hazards to the public from underwater ordnance encounters are beginning to arise. Much of the U.S. underwater contamination has occurred near military practice and test ranges, and potential hazards to the public from underwater ordnance encounters are beginning to arise.

Presently, there exists no effective capability to survey these underwater areas and map the location of UXO for site characterization, and little understanding of the UXO or clutter characteristics from which to establish performance requirements. Factors such as small target size, target burial, natural and man-made clutter and target signature modification due to target-environment acoustic coupling all impact sensor performance. See the recent underwater UXO workshop² for additional information.

This SERDP research program is designed to determine the potential effectiveness of SAS for detection and classification of UXO in complex, underwater environments. A major goal is to acquire a detailed understanding of the influence of sediment heterogeneity on SAS data and the corresponding influence on the performance of detection and classification algorithms. The approach we adopt is three-fold using (i) processing of real SAS data collected for buried and surface target fields in underwater environments, (ii) high fidelity simulations of SAS data for targets embedded in 3D realizations of sediment properties and (iii) post-processing of these data sets for design and performance evaluation of future automatic target recognition (ATR) algorithms, discovery of target features that are robust and repeatable, and insights for successful survey design. These data sets provide will provide rich coverage over a variety of experimental operating

¹The term *synthetic aperture* is taken to include any style of coherent processing of data streams collected from moving sensor platforms. Hence, this encompasses (i) the traditional synthetic aperture sonar (SAS) processing most commonly applied for collection trajectories along a linear path (with suitable motion compensation for small departures from the path) and (ii) generalized beamforming in which data is coherently processed for arbitrary collection geometries e.g., as in backprojection, but for which the source and receiver positions are known.

²SERDP and ESTCP. Final report, SERDP and ESTCP workshop on technology needs for the characterization, management, and remediation of military munitions in underwater environments. Technical report, October, 2007.

conditions and environments, and therefore, can be analyzed to yield a detailed physical and statistical understanding of SAS effectiveness for detection and classification. An additional goal is to quantify the performance of near-optimal classification algorithms using real data acquired under other programs with the Buried Object Scanning Sonar (BOSS), and features derived from 3D SAS beamforming products generated using this data. A final goal is to optimize the design of these algorithms to maximize the ROC performance metric for a selected operating point.

Scientific questions explored Sonar is a natural candidate for UXO detection in shallow water due to its wide-area surveillance capability and target sensitivity. However, sonar signature interpretation is complicated by a number of factors including (i) natural and man-made clutter, (ii) data dependence on viewing geometry and target state, (iii) environmental heterogeneity and wave propagation complexity, (iii) coupling of target response with the environment, and (iv) sensor positioning and motion compensation requirements. This study addresses the above factors with a special concentration on items (i) – (iii), and uses these findings for improved design of classification algorithms.

The evaluation for sonar-based buried target detection and discrimination was achieved by creating and analyzing a comprehensive catalogue of SAS processing results for shallow water environments. The questions explored included the detectability of buried targets with low frequency sonar, the identification of features useful for image-based discrimination, the design of classifiers using these features, and simulation methods useful for understanding target scattering phenomenology in complex environments.

We used BOSS SAS data collected in AUV measurement campaigns (separately funded from this effort), phase histories from high-fidelity simulations (based on the Spectral Element Method and on the T-Matrix method) and environment models derived from ONR-sponsored environment characterization efforts (e.g. the *Seismic Acoustic Experiments* in 1999 – SAX99, which were specifically designed to improve understanding required for detection and classification of objects buried in sediments).

Cumulative results achieved under this effort Our technical approach has yielded advances in sonar modeling capability, SAS data products derived from processing of BOSS data from various collections, a feature database for target and clutter derived from BOSS data and classification tools that exploit this data to achieve discrimination capability. In summary form, these results include

- Development of the theory for poroelastic wave propagation in a form suitable for implementation with ‘weak’ forms of the governing equations of motion including the spectral element method (SEM). This effort led to several publications e.g., [27, 29, 4].
- Numerical implementation of the 2D and 3D poroelastic formulations of the SEM. Comparison of the simulation to various benchmarks demonstrated a correspondence of these simulations with the exact analytical results for the selected end-member cases for which such results were available.
- A new *SPECFEM2D* package was released in 2009 and incorporated the poroelastic and adjoint capabilities developed under this effort. The source code has been made publically

available and it can be downloaded from the following web site:

<http://www.geodynamics.org/cig/software/packages/seismo/specfem2d/>

- A 3D version of the code with the poroelastic upgrade is expected to be released in the future. It will be known as *SPECFEM3D_Sesame*, which includes CUBIT™ compatibility.
- Generalized derivation of the T-matrix scattering formalism for free-field targets and for targets buried in a multi-layered medium. In addition, an inverse theory formalism for estimation of T-matrix coefficients was posed, and selected numerical simulations for the BOSS source/receiver geometry were performed.
- Processing of data from three separate BOSS data collections (AUV FESTs 2006, 2007 and 2008). Processing results included 3D beamformed imagery and features extracted from target and clutter detections in the imagery. Fifty feature types were extracted including geometric, intensity and statistical descriptors.
- Development of a Relevance Vector Machine classifier using extracted features from the BOSS beamformed data products. The classifier used the features to label detections in the generic binary categories of target versus clutter. Access to ground truth knowledge of the object detections enabled the statistical training of the classifier as well as computation of performance metrics. In this case classifiers and ROC curves were developed for various subsets of the available feature classes. Restricted feature set sizes were used to improve generalization performance.

Potential future applications of the research The research results achieved to date have numerous potential applications. The poroelastic extension to the spectral element method has been published in various forums [27, 29, 4]. The publically available code can be used to investigate influence of background propagation models and heterogeneity on wave propagation and scattering from targets within these environments. Investigators can use this data, for example, to design classifiers and identify optimal data collection strategies. The drawback to the method is the computation time, so questions or applications that depend on it should be designed accordingly. The SEM is well-suited to investigate scientific questions that require the high-level of fidelity supported by the code, and for which the issue of computation time is not a paramount concern. For example, it can be used to assess the validity of approximating poroelastic media with acoustic or acoustic/elastic models, and to characterize how scattering physics will differ for targets embedded in such different types of media. Examples of this were shown in [27] for the case of line array recordings for a target embedded in a two layer medium with differing layer types and material parameters. A second appropriate SEM application would be to compute off-line scattering properties for use in an on-line data exploitation system, as discussed in more detail below.

Our combined work in the SEM and the T-matrix methods provides the opportunity to use the T-matrix approach for on-line applications, but with enhanced fidelity by computing key inputs off-line using the SEM. The SEM code can be used to simulate high fidelity scattered fields for a detailed CAD model of a target of interest. The developed T-matrix theory is developed under the assumption that the target is entirely confined to a single layer of a multi-layered medium,

but is otherwise quite general. Using the inverse theory that we developed in Sec. 3.2.11, T-matrix coefficients can be estimated from the scattered data (given sufficient receiver coverage), and these coefficients can then be used on-line for very rapid synthesis of sonar time series for use in either a automatic target recognition algorithm, optimal survey design exercise, assessment of system designs or parameter selection, etc. Further, the synthetic time series may be further processed to yield 3D beamforming products from which features may be extracted. The utility of these features for classification or discrimination can then be assessed in combination with features from real data for improved classifier design. The power of this approach is that the T-matrix coefficients are *intrinsic* to the target. Additional factors that affect observed recordings such as target/sensor geometry, source waveform content, medium layering, etc., are *extrinsic*. Once the T-matrix coefficients are available from the off-line computation, synthetic series can be rapidly simulated on-line for a wide range of extrinsic parameters e.g., number and material properties of layers, source and receiver positions, and so forth. The theory for this is defined in detail in Sec. 3.2, and simple numerical results are shown in Sec. 4.2.1.

We have created a database of features derived from target and clutter objects using the BOSS acquisition system. This data set can be used by others for development of advanced classifier algorithms, and the data will be provided to SERDP with a specified format for this purpose. In addition, we have generated a database of 3D beamformed image products for the BOSS data collections. The features we derived were obtained from processing of these image products (see Table 7 for a summary description of these features). However, additional feature types that we have not considered may be derived from the image products, and these may help improve classifier performance. Additional features may be derived from 3D tomographic estimates of the sediment properties (see Sec. 3.4 and [29, 28]). Features derived from the canonical beamformed products may be used to spatially cue where tomographic inferences should be spatially culled for feature construction. Morphological processing and tomographic estimation both represent promising additional research possibilities for the future. The kernels for tomography have been derived and numerically implemented, but the approach has not yet been applied to real data.

We have developed and demonstrated a classifier algorithm for buried targets using the feature database mentioned above. We developed the algorithm based on the relevance vector machine and characterized its performance for numerous feature subsets (20 in total). These results are shown in Sec. 4.6. There is considerable additional work that can be pursued including the application of a wrapper around the training algorithm itself that adaptively optimizes the feature subset. This is basically a combinatorial optimization problem involving the discrete index selection from a large superset. Various algorithms are applicable to this including greedy methods such as sequential forward method, the sequential backward method, and global methods such as genetic algorithms. The key here is training of the RVM is very fast (less than one minute using a single 2.1 GHz chip), and therefore, the model space can be adequately explored.

To generate the feature classes described above, we applied morphological processing to the 3D beamformed image products. The latter were generated using the image formation algorithms described in Sec. 3.3. However, there were a number of simplifying assumptions used in the described algorithm, and there is the opportunity to enhance the quality and sharpen the focus of the derived data products. This can follow from the inclusion of any number of approaches including exploitation of available environmental information (e.g., by inverting the BOSS data itself for the best fitting 1D sediment model, and to use this in the specification of the beamforming weighting coefficients). An additional strategy to improve image product data is to apply more advanced

beamformers than the simple weighted delay and sum technique that we have used. Candidates for more advanced beamformers include the Capon and Minimum Variance Distortionless Response (MVDR) methods [36].

Finally, the results achieved in this program and the future research directions suggested above should be of interest to Mine Counter-measures (MCM) programs at the ONR. The work is physics-based, and new methodologies have been introduced making application to mine targets relatively straightforward.

1 Objective

Our objective has been to generate a large catalogue of SAS images using real data from multiple data collections acquired elsewhere with the BOSS system, and synthetic data from sonar phase history calculations.

Forward simulation methods that are commonly used include Kirchhoff-based approaches and various implementations of T-matrix theories. These methods also are the basis of the simulation approaches supported in PC SWAT [34]. While very useful for insight to target responses, the methods require either homogeneous or layered environment models and therefore do not readily capture target-environment interactions arising from natural clutter. Thus, an objective of this effort is to extend sonar modeling capability using a state-of-the-art geophysical simulation method known as the spectral element method (SEM) (see, e.g. [20]). The SEM captures arbitrary target/environment complexity and accurately models all wave phenomena (resonant modes, surface waves, diffractions, specular scattering, target-environment coupling) over the low and high frequency regimes of interest. This enables careful characterization of realistic environment imprints on target signatures and signature variability in SAS imagery. Shallow-water environment models can be derived from the outcomes of the extensive SAX99 experiments ([24, 9]).

Our research focus is SAS-based target recognition of buried munition objects in shallow water environments. The latter consists of the fluid column, the water-sediment interface, and the embedding sediments. The stratification of these components and the heterogeneity within them originating from depositional, biological, oceanographic and other processes strongly control wave behavior and the corresponding sonar observations. Target scattering processes are coupled to propagation phenomena in the environment, and will therefore obscure and/or modify target signatures predicted for targets emplaced in idealized media (e.g. a homogeneous two-layer half-space). Further, target responses and wave-field interactions with the environment can change significantly depending on the frequency content of the active source signal. Generalized synthetic aperture sonar processing approaches are applied to SAS phase history observations to yield spatial reflectivity maps with cross-range resolution far greater than can be achieved with real aperture processing alone. In addition, SAS processors attempt to compensate for effects that would otherwise degrade the image resolution (e.g. platform position perturbations from assumed straight-line trajectories). Nonetheless, there will always be a data imprint of shallow water complexity on target signatures due the intrinsic acoustic/elastic coupling of targets and the environment. This will vary widely according to the target type and emplacement, environment conditions, sensor array design and source bandwidth, collection conditions, measurement campaign protocol and type of SAS-processor applied. Natural and man-made clutter objects (both buried and proud) pose an additional challenge to the discrimination problem. The success of sonar-based discrimination algorithms will depend on the distinctness of UXO target and clutter object signatures (e.g. their separation in feature space), their stability/measurability in the complex shallow water environment and the suitability of detection/discrimination algorithm design given the above.

In summary, our objectives are to provide

- a knowledge-base consisting of SAS imaging data products, features and discrimination measures derived from the SAS products, and shallow-water environment models and UXO target models represented with hexahedral elements (required by the SEM below)
- a high-fidelity spectral element numerical method for modeling wave propagation and target

scattering in the complex shallow-water environment (and linked to the knowledge-base); the purpose of this is to permit a thorough understanding of the effect of realistic environment complexity on target signatures, to enable numerous data realizations consistent with the statistics of the environment characterizations, to assess the corresponding observability and stability of target signatures/features, and to provide a valuable resource for future discrimination algorithm development and

- a comprehensive analysis synthesizing findings from the real and synthetic data results including recommendations for future directions.

Our technical approach combines informed use of both real and simulated data inputs to assess sonar-based target detection/discrimination capability. Practical data processing of real data sets invariably requires use of simplifying assumptions, especially for propagation effects in complex environments. The influence of the environment and these assumptions is difficult to untangle for real collections. In terms of target classification, any practical deployment will most likely use predicted target features for a target buried in a simplified background medium. By performing controlled high-fidelity simulation experiments, it will be possible to precisely quantify the influence of simplifications used in the practical data processors on the ultimate classification performance. For example, realistic simulations can be performed for a fully realized complex shallow water environment with a buried target. The processing sequence (such as that described for BOSS below) can then be applied to the synthetic data set using homogeneous background assumptions. In this way, robustness of features used for classification with respect to environmental and target variability can then be precisely characterized. Numerous stochastic realizations of such environments can then be used to provide a meaningful statistical understanding.

In summary, our technical execution plan consisted of the following key steps:

1. Process data from the BOSS collections for phenomenology understanding and empirical characterization of performance. Data was processed from known target fields in which clutter and targets could be separately identified on the basis of prior knowledge.
2. Extract features from BOSS image products and derive classifiers using these features.
3. Perform high-fidelity numerical simulations of the acoustic phase histories for standard SAS data collection experiments including BOSS. The spectral element method was used for this purpose. Perform data product generations (e.g., SAS images) using the phase histories from the resulting simulation.
4. The original objective was to perform SEM simulations for many different stochastic realizations of the propagation models, and for each corresponding SAS image, determine the target signatures and features. The objective of this is to explicitly determine the coupled medium-target influence on the SAS image, to determine the stability of target-related features and the degree to which they can be identified and extracted in the presence of the environment complexity and variability. However, the extended development time of the SEM for poroelastic media, and the computation time required for each simulation prevented this particular objective from being fully realized.

2 Background

Over the past decade there has been a dedicated research effort for the development of detection and classification methods for underwater objects in shallow and littoral waters. The sonar modality has been a method of choice due to its far range, wide-area coverage capability and diagnostic value. Applications have included real-aperture side-looking sonar (SLS) and the higher resolution synthetic aperture sonar. Example theoretical and observational studies of acoustic responses for surface and buried targets include [8, 30, 6, 7, 5, 25, 16, 13, 17, 2, 22, 21, 19, 23, 34, 33, 14]. Figure 1 depicts some of the major wave processes that are involved. Many of these studies have focused on the sea mine discrimination problem; however, considerations for the UXO problem are

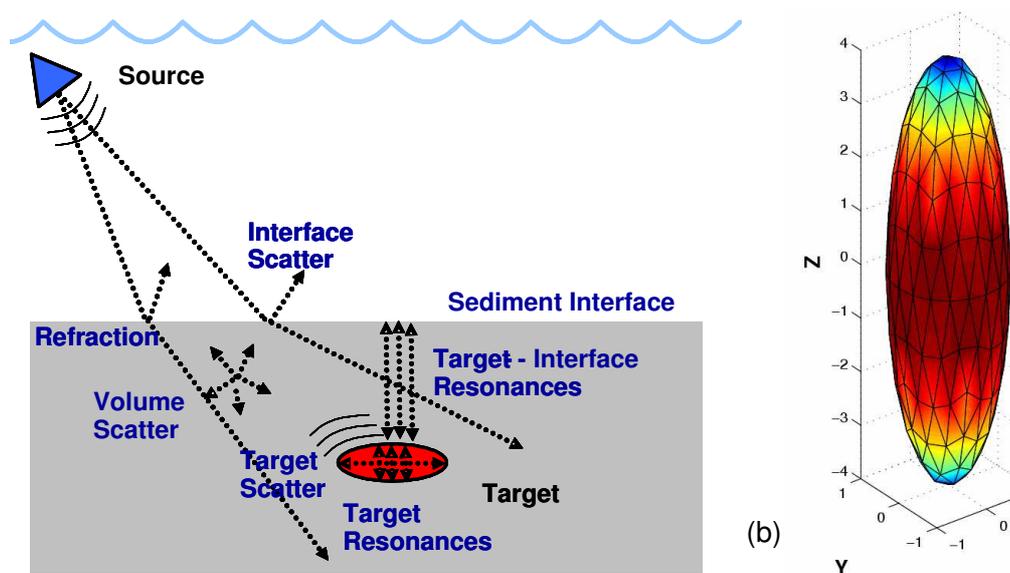


Figure 1: (a) The scattered wavefield due to an acoustic source is comprised of target and environment returns. The *target scatter* in general consists of a complex superposition of wave modes including direct target scatter, internal target resonances, surface modes, and target-interface resonances all of which are mediated by the target shape, size, acoustic-elastic properties and the incident field (e.g. frequency content). The *environment scatter* include interface scatter, rough interface scatter and volume scatter due to heterogeneity in the fluid and the sediment. (b) Example tessellation of a prolate spheroid target. In high-frequency numerical approaches (e.g. Kirchhoff method) approximate solutions to the Helmholtz boundary integral equation are obtained on each facet.

roughly similar. Stripmap and spotlight are the two major types of SAS imaging [11], [12], [15]. In the stripmap mode the beam is held orthogonal to the track of the sonar platform. In the spotlight mode the beam is steered toward a specific area during the entire integration time, resulting in an increase in resolution. In practice, the stripmap mode is most commonly used. More generally, beamforming over synthetic apertures may be applied, and it is this processing style that has been the focus of our processing thus far.

The challenges of SAS imaging for detection/discrimination for UXO targets in shallow water environments derives from a number of sources. These include (i) natural and man-made clutter

objects (both surface and buried), (ii) the complex coupling acoustic/elastic coupling between targets and the propagation environment, (iii) the requirement of a robust SAS processing system for actual data applications and (iv) the need for a discrimination theory well-matched to the complexity of the environment and the targets. Our work focuses on elements (i)-(iii) to provide a comprehensive assessment using state-of-the art technology. A natural outcome will be input for item (iv) i.e. features and processing for a future recognition system. Resolution of the clutter problem will rely on exploiting of the intrinsic sensitivity of the sensor response to target versus clutter objects. Mitigation of the complicating effect of environment imprints on target signatures will rely on first characterizing the nature of these imprints, and either discovering feature classes that are robust as possible to the environment-induced signature variability, or development of methods to deconvolve the effect of the environment on the measurements. Finally, design of an effective SAS-based target recognition system will rely on careful synthesis and extension of findings from elements (i)-(iii) above.

There are at least three data sets that are of interest for continuing study of underwater UXO, and these are the DARPA, CSS and BOSS data sets as we now describe.

DARPA data DARPA funded a three-phased SAS research and development program from 1994-2000. In phase two, SAS-compatible data were recorded in the waters of Washington State and processed by one of our team members (AST). The data were collected in four areas spanning a range of depths, ranges, bottom types, and in both fresh and salt water. Targets ranged from test frames of corner reflectors and metal spheres (resolution targets) to mine-like targets to downed airplanes. In Puget Sound near Carkeek Park, a salt water region, the bottom type was sand and rock and data include upslope views of resolution targets and a simulated mine at long ranges. In Puget Sound near Spring Beach, a salt water region, the bottom type was coarse sand and most data were collected of upslope views of targets. Note that the upslope views often demonstrate multi-path returns from targets resulting in multiple images of targets. Dabob Bay, a salt water region, had a silty, fine sand bottom and various test targets include a WWII mine. In Lake Washington, a fresh water region with a muddy bottom, we have data of a test frame (see Fig. 2) and a sunken PB4Y-2 airplane at various ranges. We can count the wing struts on the airplane and imaged it from a 1 km range with an acoustic path that included a bottom and thermal layer bounce in each direction (e.g., [1]). These early data may be made compatible with AST's PROSAS™, but may not be ideal candidates due to the non-integrated navigation equipment used on the towfish, however, we can process the data using AST's development version SAS processor. Sound velocity profiles were recorded during the trials and could aid our study of the effects of the environment on the recorded sonar data.

CSS data: In 1998 and 1999, AST was funded by ONR to process data collected by the (then) Coastal Systems Station SAS system. AST has in-house several data sets recorded in St. Andrews Bay and in the Gulf of Mexico (both medium sand bottoms) from sea trials in 1998 and 1999; additional data sets should be available from the Naval Surface Warfare Center - Panama City (NSWC PC). These data sets contain targets including fluid filled spheres, ladders, cylinders, and mine-like shapes (see Fig. 3) at two frequencies, 180 kHz and 20 kHz. Some interesting differences between the two images indicate how target response is frequency dependent. For example, the cylinder object (upper right corner) in the high frequency image (left panel) is bright along its whole length,

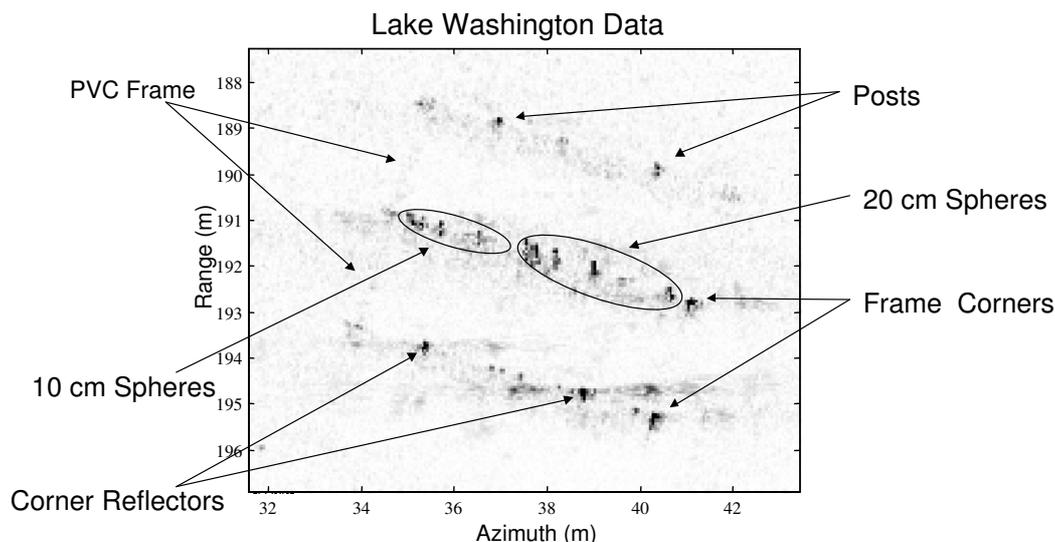


Figure 2: SAS data collected at Lake Washington under DARPA-sponsorship for various targets and clutter objects as indicated in the figure.

and a distinct shadow is visible. However, only its ends are bright in the low frequency image (right panel), and no shadow is present. Similarly, the long shadow behind the truncated cone (just below the cylinder) also disappears for the low frequency image.

Buried Object Scanning Sonar (BOSS) data: Florida Atlantic University has recently developed the BOSS system [32]. This 3-19 kHz, sediment volume imaging sonar operates in conventional, tomographic and synthetic aperture modes of operation. Several data sets for fields consisting of buried sea mines and clutter objects have been collected in sand and mud environments off of Panama City, FL, Boca Raton, FL, and Kaneohe, Hawaii (Fig. 4). These collections have been the focus of the data analysis presented in this report. The data is useful given our objectives since the low frequency BOSS transducer enables sediment penetration. In addition, ground truthing (clutter and target knowledge) is available for a number of collections, and this supports classifier development and performance characterization of the classifier.

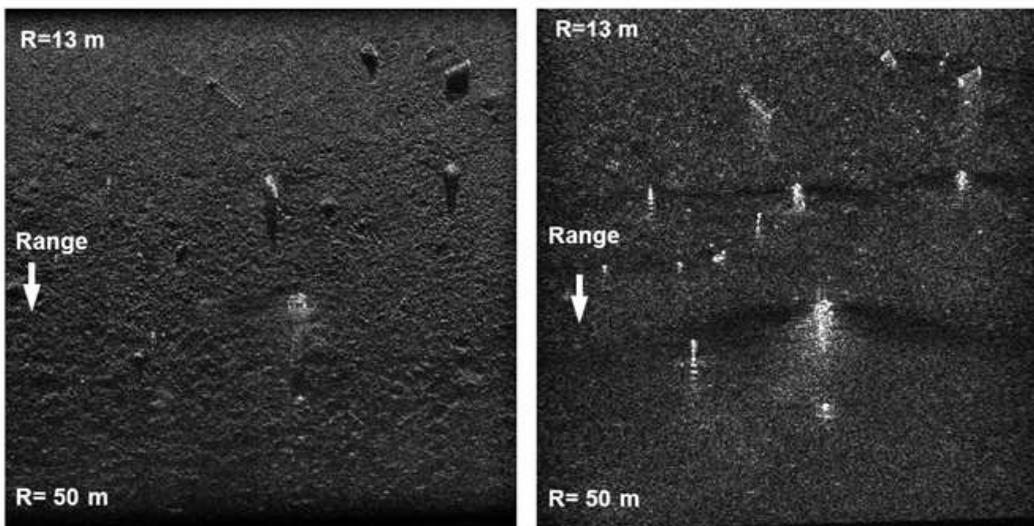


Figure 3: This figure shows an example application of our SAS image formation processor to sonar data collected by Coastal Systems Station (Panama City, Florida) using a Northrop-Grumman array. This array contained 11 high frequency (180 kHz) elements and 14 low frequency (20 kHz) elements. For these experiments a series of objects (e.g. large cylinder, ladder, a truncated cone) were placed on a sandy bottom. **Left panel:** This shows the focused image for the high frequency data. A cylinder in the upper right gives a clear shadow. The image was formed using redundant phase centers (RPC) and the range migration algorithm (RMA), but an autofocus step was not required. **Right panel:** Low frequency (20 kHz) SAS image formed using RPC, RMA and autofocus. These results illustrate that objects and background reverberation image characteristics have strong frequency sensitivity. Adapted from [1].

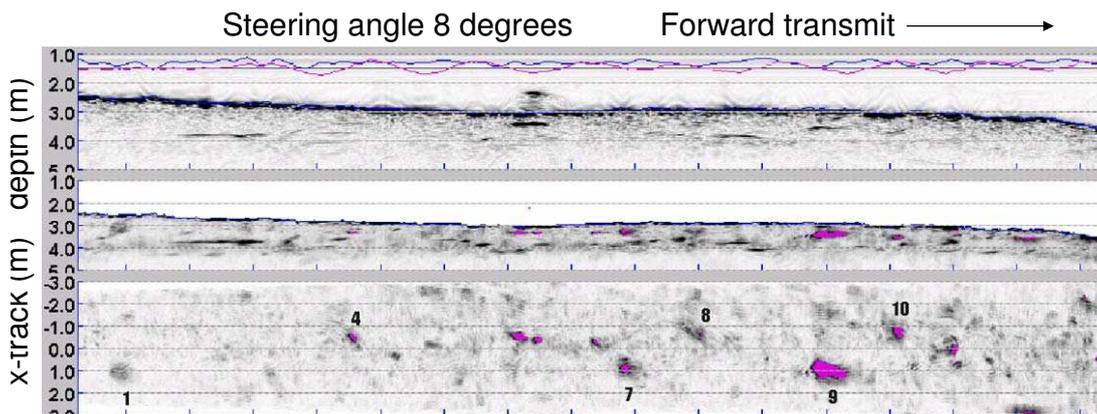


Figure 4: For the proposed work we will have access to Buried Object Scanning Sonar data. The sites consist of buried sea mines and clutter objects in various sand and mud environments off of Kaneohe, Hawaii and Panama City, FL. The low frequency data is ideal for buried target interrogation.

3 Materials and Methods

In this section we address the methodological aspects of the program effort including development of simulation methods (Secs. 3.1 and 3.2 for SEM and T-matrix methods, respectively), BOSS data processing methods (Sec. 3.3), early results on the use of adjoint tomography for developing new feature sets (Sec. 3.4), and the basis of the Bayesian classifier that we apply to BOSS data features (Sec. 3.5).

3.1 The spectral element method for poroelastic media

An extensive research investigation was performed to develop the Spectral Element Method for propagation in poroelastic sediment media. These theoretical finds and numerical methods are reported in various publications [27, 29, 4] and will not be repeated here.

3.2 T-matrix analysis

The following sub-sections develop in detail the forward problem for T-matrix scattering calculations, and the inverse problem for T-matrix coefficients. Numerical results for the forward problem are presented in a later section.

3.2.1 Acoustic equation and Green function

The acoustic equation of motion for the pressure field $\phi(\mathbf{x}, t)$ takes the form

$$\frac{1}{\rho c^2} \partial_t^2 \phi - \nabla \cdot \left(\frac{1}{\rho} \nabla \phi \right) = S \quad (1)$$

where $S(\mathbf{x}, t)$, $\rho(\mathbf{x})$, and $c(\mathbf{x})$ are the source, mass density and sound speed fields, respectively. In the frequency domain this takes the form

$$-\nabla \cdot \left(\frac{1}{\rho} \nabla \hat{\phi} \right) - \frac{\omega^2}{\rho c^2} \hat{\phi} = \hat{S}, \quad (2)$$

in which Fourier transform pair $\phi, \hat{\phi}$ is defined by

$$\begin{aligned} \hat{\phi}(\mathbf{x}, \omega) &= \int dt \phi(\mathbf{x}, t) e^{i\omega t} \\ \phi(\mathbf{x}, t) &= \int \frac{d\omega}{2\pi} \hat{\phi}(\mathbf{x}, \omega) e^{-i\omega t}, \end{aligned} \quad (3)$$

and similarly for S . The formal solution to (2) is

$$\hat{\phi}(\mathbf{x}) = \int d^3x' g(\mathbf{x}, \mathbf{x}') \hat{S}(\mathbf{x}'), \quad (4)$$

in which the Green function g satisfies,

$$\nabla \cdot \left[\frac{1}{\rho} \nabla g(\mathbf{x}, \mathbf{x}') \right] + \frac{\omega^2}{\rho c^2} g(\mathbf{x}, \mathbf{x}') = -\delta(\mathbf{x} - \mathbf{x}'). \quad (5)$$

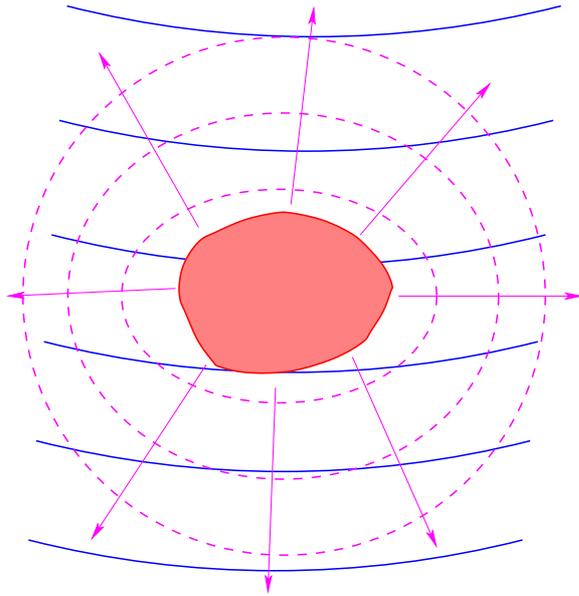


Figure 5: Schematic illustration of background and scattered (outgoing wave) fields. The former represent the field in the absence of the scatterer, while the latter represent the corrections to the field due to the presence of the scatterer. The two fields are expanded in a set of basis functions according to (8) and (9), and related via (11) through the T -matrix.

In an infinite homogeneous medium, one obtains the free-space Green function

$$g_0(\mathbf{x} - \mathbf{x}') = \frac{\rho}{4\pi|\mathbf{x} - \mathbf{x}'|} e^{ik|\mathbf{x} - \mathbf{x}'|}, \quad (6)$$

with wavenumber $k = \omega/c$.

In general (1) will fail if there are non-fluid elements in the medium, such as solid elastic scatterers. However, in what follows we will only require the Green function for the background fluid medium for which (1) is valid everywhere.

3.2.2 Free-field T -matrix

Consider first a target lying in an infinite homogeneous medium with density ρ and sound speed c . We separate the field into background and scattered components (see Fig. 5),

$$\hat{\phi}(\mathbf{x}) = \hat{\phi}_b(\mathbf{x}) + \hat{\phi}_s(\mathbf{x}), \quad (7)$$

in which $\hat{\phi}_b$ is the field generated by the given source in the absence of the scatterer, and $\hat{\phi}_s$ corrects for the scattered waves. Let the scatterer be centered on a point \mathbf{x}_s . It is assumed that $\hat{\phi}_b$ is smooth and regular in a neighborhood of \mathbf{x}_s , including the entire volume occupied by the scatterer, and admits a convergent expansion,

$$\hat{\phi}_b(\mathbf{x}) = \sum_n A_n \xi_n(\mathbf{x} - \mathbf{x}_s) \quad (8)$$

in terms of a conveniently chosen set of basis functions $\xi_n(\mathbf{x})$ which are all regular at the origin. Similarly, the scattered field is assumed to admit a convergent expansion,

$$\hat{\phi}_s(\mathbf{x}) = \sum_n B_n \psi_n(\mathbf{x} - \mathbf{x}_s), \quad (9)$$

outside a sufficiently large neighborhood of the scatterer. The functions ψ_n will generally be singular somewhere inside this neighborhood. Both sets of basis functions are solutions to the homogeneous acoustic equation (2) with $\hat{S} \equiv 0$. For targets that are not too far from spherical, convenient basis function choices are the spherical waves,

$$\begin{aligned} \xi_{lm}(\mathbf{x}) &= j_l(kx)Y_{lm}(\theta, \phi) \\ \psi_{lm}(\mathbf{x}) &= h_l^{(1)}(kx)Y_{lm}(\theta, \phi), \end{aligned} \quad (10)$$

in which (x, θ, ϕ) are spherical coordinates for \mathbf{x} centered on the origin, $j_l(z)$ are the spherical Bessel functions (which are indeed regular, varying as z^l near the origin), $h_l^{(1)}(z)$ the spherical Hankel functions (which indeed diverge as $1/z^{l+1}$ at the origin), and Y_{lm} are the spherical harmonics [18]. For future reference, we note that

$$\begin{aligned} Y_{lm}(\theta, \phi) &= \beta_{lm} P_l^m[\cos(\theta)] e^{im\phi} \\ \beta_{lm} &= \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}} \end{aligned} \quad (11)$$

in which $P_l^m(u)$ are the associated Legendre functions [18]. We also note the useful identities,

$$\begin{aligned} Y_{lm}^*(\theta, \phi) &= Y_{lm}(\theta, -\phi) \\ &= (-1)^m Y_{l,-m}(\theta, \phi) \end{aligned} \quad (12)$$

$$\begin{aligned} Y_{lm}(\pi - \theta, \pi + \phi) &= (-1)^m Y_{lm}(\pi - \theta, \phi) \\ &= (-1)^{l+m} Y_{lm}(\theta, \pi + \phi) \\ &= (-1)^l Y_{lm}(\theta, \phi) \end{aligned} \quad (13)$$

$$P_l^m(-u) = (-1)^{l+m} P_l^m(u). \quad (14)$$

The latter two correspond to spatial inversion, $\mathbf{x} \rightarrow -\mathbf{x}$.

The T -matrix represents the linear relationship

$$B_n = \sum_{n'} T_{nn'}(k) A_{n'} \quad (15)$$

between the A and B coefficients that must emerge from a full solution to the acoustic equation in the presence of the scatterer. For a spherically symmetric scatterer,

$$T_{lm;l'm'}(k) = T_l(k) \delta_{ll'} \delta_{mm'} \quad (16)$$

is diagonal and independent of m in the spherical wave basis, and exact solutions are available for a variety of examples. In matrix form, (15) is written simply as $\mathbf{B} = \mathbf{T}\mathbf{A}$.

Since (7)–(10) refer only to the solutions in the background medium, the relation (15) holds even if (2) fails to apply within the target, for example if it is an elastic body rather than a fluid body. The T -matrix is, of course, much more difficult to compute in such cases.

3.2.3 Heterogeneous field T -matrix

Suppose now that the target resides in a medium that is homogenous in the target neighborhood, but not necessarily globally homogeneous. Examples to be discussed below include horizontally stratified media, which may or may not contain other compact scatterers. Local homogeneity implies that we may continue to expand the background field (which excludes the effects of the target, but includes the effects of all other inhomogeneities in the medium) locally near \mathbf{x}_s in the form (8). However, the scattered field now has an expansion

$$\hat{\phi}_s(\mathbf{x}) = \sum_n B_n \Psi_n(\mathbf{x} - \mathbf{x}_s; \mathbf{x}_s) \quad (17)$$

in which Ψ_n is a solution to the global acoustic equation, and the residual \mathbf{x}_s dependence arises because the system is no longer translation invariant. The expansion is uniquely defined by writing

$$\Psi_n(\mathbf{x}; \mathbf{x}_s) = \psi_n(\mathbf{x}) + \Delta\psi_n(\mathbf{x}; \mathbf{x}_s), \quad (18)$$

where the correction $\Delta\psi_n$ arises from reflections of the scattered wave ψ_n that return to the neighborhood of \mathbf{x}_s . As such, $\Delta\psi_n$ is assumed to be smooth in the neighborhood of \mathbf{x}_s , and have a convergent re-expansion,

$$\Delta\psi_n(\mathbf{x}; \mathbf{x}_s) = \sum_n U_{nn'}(\mathbf{x}_s) \xi_{n'}(\mathbf{x}), \quad (19)$$

in terms of the regular basis functions. Examples of the matrix \mathbf{U} will be presented below.

The free field T -matrix now relates the *full* ξ_n component of the field to the scattered field. In matrix form one obtains,

$$\mathbf{B} = \mathbf{T}(\mathbf{A} + \mathbf{U}\mathbf{B}), \quad (20)$$

with solution,

$$\mathbf{B} = \mathcal{T}\mathbf{A}, \quad \mathcal{T} = (\mathbf{T}^{-1} - \mathbf{U})^{-1}. \quad (21)$$

Through (21) we have succeeded in expressing the full solution to the scattering problem in terms of the free field scattering problem, together with the global properties of the background medium. The existence of such local scattering relations relies on the spatial locality of the gradients in the acoustic equations (1) and (2).

3.2.4 T -matrix for multiple scatterers in a homogeneous medium

Consider next a collection of N targets, centered on points $\mathbf{x}_s^{(j)}$, $j = 1, 2, \dots, N$, in a homogeneous medium, each with its own T -matrix $\mathbf{T}^{(j)}$. The background field is now the field generated by a given source when *all* targets are absent. Since the medium is homogeneous, the basis functions $\xi_n(\mathbf{x} - \mathbf{x}_s^{(j)})$, $\psi_n(\mathbf{x} - \mathbf{x}_s^{(j)})$ in each neighborhood may be chosen simply as translates of each other. We now require the re-expansion coefficients of the scattered waves from each target about every other. Let us define the matrix $\mathbf{U}(\mathbf{x}')$ by

$$\psi_n(\mathbf{x} - \mathbf{x}') = \sum_{n'} U_{nn'}(\mathbf{x}') \xi_{n'}(\mathbf{x}). \quad (22)$$

If $A_n^{(j)}$ are background coefficients appropriate to the neighborhood of $\mathbf{x}_s^{(j)}$, and $B_n^{(j)}$ are the corresponding exact scattering coefficients, then one obtains for each j the self-consistency condition,

$$\mathbf{B}^{(j)} = \mathbf{T}^{(j)} \left[\mathbf{A}^{(j)} + \sum_{j'(\neq j)} \mathbf{U}^{(jj')} \mathbf{B}^{(j')} \right], \quad (23)$$

in which $\mathbf{U}^{(jj')} \equiv \mathbf{U}(\mathbf{x}_s^{(j')} - \mathbf{x}_s^{(j)})$. Let us define the super-matrices

$$\mathbf{T} = \begin{bmatrix} \mathbf{T}^{(1)} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{T}^{(2)} & \dots & \mathbf{0} & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{T}^{(N-1)} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{T}^{(N)} \end{bmatrix}$$

$$\mathbf{U} = \begin{bmatrix} \mathbf{0} & \mathbf{U}^{(12)} & \dots & \mathbf{U}^{(1,N-1)} & \mathbf{U}^{(1,N)} \\ \mathbf{U}^{(21)} & \mathbf{0} & \dots & \mathbf{U}^{(2,N-1)} & \mathbf{U}^{(2,N)} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathbf{U}^{(N-1,1)} & \mathbf{U}^{(N-1,2)} & \dots & \mathbf{0} & \mathbf{U}^{(N-1,N)} \\ \mathbf{U}^{(N1)} & \mathbf{U}^{(N2)} & \dots & \mathbf{U}^{(N,N-1)} & \mathbf{0} \end{bmatrix}. \quad (24)$$

Equation (23) then takes the supermatrix form (20), in which \mathbf{A}, \mathbf{B} are now column vectors constructed from all of the $\mathbf{A}^{(j)}, \mathbf{B}^{(j)}$. With these definitions, the formal solution still takes the form (21).

3.2.5 T -matrix for multiple scatterers in a heterogeneous medium

Let us finally consider the most general problem of a collection of targets in a heterogeneous medium, each of which lies in a homogeneous neighborhood about its center $\mathbf{x}_s^{(j)}$. The object is now to find the formal solution to the scattering problem given only the individual matrices $\mathbf{T}^{(j)}$ and the background properties of the medium if none of the scatterers are present. The major difference from the results of Sec. 3.2.4 is that each such neighborhood in general has a different local density $\rho^{(j)}$ and sound speed $c^{(j)}$, and there is a self-scattering contribution from each target.

From (18), we define the local basis functions

$$\Psi_n^{(j)}(\mathbf{x}; \mathbf{x}_s^{(j)}) = \psi_n^{(j)}(\mathbf{x}) + \Delta\psi_n^{(j)}(\mathbf{x}; \mathbf{x}_s^{(j)}), \quad (25)$$

in which $\psi_n^{(j)}$ depends on j through the local values of $\rho^{(j)}$ and $c^{(j)}$. Similarly, the regular basis functions $\xi_n^{(j)}(\mathbf{x} - \mathbf{x}_s^{(j)})$ will now depend on j . For $j \neq j'$ we define the matrix $\mathbf{U}^{(jj')} = \mathbf{U}(\mathbf{x}_s^{(j)}, \mathbf{x}_s^{(j')})$ via

$$\Psi_{n'}^{(j')}(\mathbf{x} - \mathbf{x}_s^{(j')}; \mathbf{x}_s^{(j')}) = \sum_n U_{nn'}^{(jj')} \xi_n^{(j)}(\mathbf{x} - \mathbf{x}_s^{(j)}), \quad j \neq j', \quad (26)$$

while for $j = j'$ only the reflected wave contributes:

$$\Delta\psi_{n'}^{(j)}(\mathbf{x} - \mathbf{x}_s^{(j)}; \mathbf{x}_s^{(j)}) = \sum_n U_{nn'}^{(jj)} \xi_n^{(j)}(\mathbf{x} - \mathbf{x}_s^{(j)}), \quad j = j'. \quad (27)$$

Thus, $\mathbf{U}^{(jj')}$ describes the re-expansion of the outgoing waves from scatterer j' in the vicinity of scatterer j . If $A_n^{(j)}$ are background coefficients appropriate to the neighborhood of $\mathbf{x}_s^{(j)}$, and $B_n^{(j)}$ are the corresponding exact scattering coefficients, then one obtains the self-consistency condition,

$$\mathbf{B}^{(j)} = \mathbf{T}^{(j)} \left[\mathbf{A}^{(j)} + \sum_{j'} \mathbf{U}^{(jj')} \mathbf{B}^{(j')} \right], \quad (28)$$

in which the diagonal term $j' = j$ is now included in the sum. The formal solution still takes the form (21), but the supermatrix \mathbf{U} is now nonzero along the diagonal:

$$\mathbf{U} = \begin{bmatrix} \mathbf{U}^{(11)} & \mathbf{U}^{(12)} & \dots & \mathbf{U}^{(1,N-1)} & \mathbf{U}^{(1,N)} \\ \mathbf{U}^{(21)} & \mathbf{U}^{(22)} & \dots & \mathbf{U}^{(2,N-1)} & \mathbf{U}^{(2,N)} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathbf{U}^{(N-1,1)} & \mathbf{U}^{(N-1,2)} & \dots & \mathbf{U}^{(N-1,N-1)} & \mathbf{U}^{(N-1,N)} \\ \mathbf{U}^{(N1)} & \mathbf{U}^{(N2)} & \dots & \mathbf{U}^{(N,N-1)} & \mathbf{U}^{(NN)} \end{bmatrix}. \quad (29)$$

3.2.6 Basis functions in horizontally stratified media

Let the background medium parameters depend only on the vertical coordinate z . The Green function is then horizontally translation invariant, and may be conveniently written in terms of its horizontal Fourier transform,

$$\begin{aligned} g(\mathbf{x}, \mathbf{x}') &= \int \frac{d^2q}{(2\pi)^2} \hat{g}(z, z'; q) e^{i\mathbf{q} \cdot (\mathbf{r} - \mathbf{r}')} \\ &= \int_0^\infty \frac{q dq}{2\pi} \hat{g}(z, z'; q) J_0(q|\mathbf{r} - \mathbf{r}'|), \end{aligned} \quad (30)$$

in which $\mathbf{x} = (\mathbf{r}, z)$ and \hat{g} depends only on the magnitude $q = |\mathbf{q}|$. If, in addition, the background consists of L layers, with boundaries at $d_1 > d_2 > \dots > d_{L-1}$, so that the parameters are piecewise constant functions of z (see Fig. 6), then \hat{g} takes the form of a superposition of counter-propagating plane waves within each layer:

$$\begin{aligned} \hat{g}(z, z'; q) &= a_u(q) e^{-l_u(q)|z-z'|} \delta_{uv} \\ &+ \sum_{\sigma, \tau = \pm 1} a_{u\sigma; v\tau}(q) e^{\sigma l_u(q)z + \tau l_v(q)z'}, \end{aligned} \quad (31)$$

in which z lies in layer u , z' lies in layer v ($u, v = 1, 2, \dots, L$), and the vertical wavevector is given by,

$$l_u(q) = \sqrt{q^2 - k_u^2} = -iq_{z,u}(q). \quad (32)$$

Here $k_u = \omega/c_u$, and ρ_u, c_u are the parameter values in layer u . We choose the sign of the square root so that $\text{Re}(l_u) \geq 0$ whereby $\sigma, \tau = +1$ leads to exponential decay for large negative z, z' . If $k_u^2 - q^2$ is real and positive, then we choose $q_{z,u} = \sqrt{k_u^2 - q^2} > 0$ so that $\sigma, \tau = +1$ yield downward propagating waves. The diagonal term

$$a_u(q) = \frac{\rho_u}{2l_u(q)} \quad (33)$$

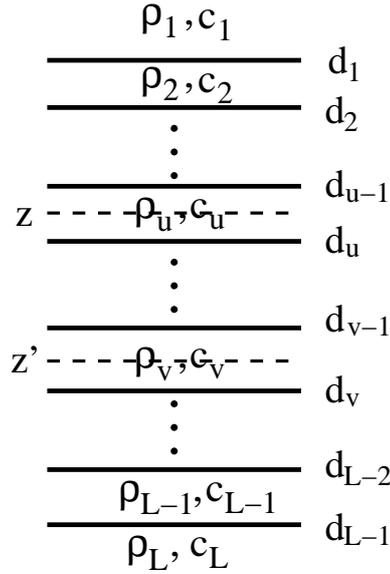


Figure 6: Schematic diagram of horizontally stratified geometry. Notation is defined in the text.

corresponds to the horizontal Fourier transform of the free Green function (6), which may also be written in the form $g_0(\mathbf{x} - \mathbf{x}') = (ik\rho/4\pi)h_0^{(1)}(k|\mathbf{x} - \mathbf{x}'|)$. The coefficients $a_{u\sigma;v\tau}$ all vanish for a homogeneous (single layer) system but rapidly increase in complexity with the number of layers. However, computationally efficient analytic forms for them exist³ [38], and the three layer forms are displayed in App. 3.2.12. Symmetry of \hat{g} under interchange of z, z' leads to the corresponding symmetry $a_{u\sigma;v\tau} = a_{v\tau;u\sigma}$.

Equation (31) represents an extension of the free Green function through all of space, and includes all multiple scattering from and between the layers. In order to similarly extend the higher order Hankel functions, one needs a corresponding horizontal Fourier transform for them. Let us define the horizontal Fourier transform of the outgoing wave basis,

$$\begin{aligned} \psi_{lm}(\mathbf{x}) &= h_l(kx)Y_{lm}(\theta, \phi) \\ &= \int \frac{d^2q}{(2\pi)^2} \hat{\psi}_{lm}(z; \mathbf{q}) e^{i\mathbf{q}\cdot\mathbf{r}}, \end{aligned} \quad (34)$$

so that the free Green function is related to ψ_{00} via

$$\begin{aligned} g_0(\mathbf{x} - \mathbf{x}') &= \frac{ik\rho}{\sqrt{4\pi}} \psi_{00}(\mathbf{x} - \mathbf{x}') \\ \hat{\psi}_{00}(z - z'; \mathbf{q}) &= \frac{\sqrt{4\pi}}{ik\rho} \hat{g}_0(z, z'; q), \end{aligned} \quad (35)$$

where we note that $Y_{00} = 1/\sqrt{4\pi}$. One obtains [31] (see Sec. 3.2.11.3.1 for a derivation):

$$\hat{\psi}_{lm}(z; \mathbf{q}) = i^{-(l+1)} \frac{4\pi}{k} Y_{lm}(\theta_{Q\sigma}, \phi_Q) \frac{1}{2l(q)} e^{-l(q)|z|}, \quad (36)$$

³The required scalar Green function is denoted $\hat{\gamma}^\mu(z, z'; q)$ there, with coefficients $a_{u\sigma;v\tau}^\mu, a_u^\mu$, and the mapping $\mu \rightarrow \rho$ and $\epsilon k^2 \rightarrow k_u^2$ should be made to obtain the scalar Green function $\hat{g}(z, z'; q)$ required in the present work.

in which $\sigma = -\text{sgn}(z)$, and $\theta_{Q\sigma}(\mathbf{q}), \phi_Q(\mathbf{q})$ are the polar coordinate angles of the vector $\mathbf{Q}_\sigma = [\mathbf{q}, -i\sigma l(q)] = [\mathbf{q}, -\sigma q_z(q)]$, corresponding to an upward propagating wave for $z > 0$ and a downward propagating wave for $z < 0$:

$$\begin{aligned}\cos(\theta_{Q\sigma}) &= -\sigma q_z(q)/k = -\sigma \sqrt{1 - (q/k)^2} \\ &= -i\sigma l(q)/k = -i\sigma \sqrt{(q/k)^2 - 1} \\ \sin(\theta_{Q\sigma}) &= q/k.\end{aligned}\quad (37)$$

As seen, for complex k (as in an absorptive medium), or $q > k$ the polar angle $\theta_{Q\sigma}$ becomes complex. However, since the spherical harmonics are polynomials in $\cos(\theta_{Q\sigma})$ and $\sin(\theta_{Q\sigma})$, the analytic continuation is trivial (see the corresponding discussion in App. 3.2.11.3.1). Note also that $\theta_{Q,-\sigma} = \pi - \theta_{Q\sigma}$, and recall the corresponding identity (13).

Although not actually required in what follows, one may also define the horizontal Fourier transform for the regular basis functions,

$$\begin{aligned}\xi_{lm}(\mathbf{x}) &= j_l(kx)Y_{lm}(\theta, \phi) \\ &= \int \frac{d^2q}{(2\pi)^2} \hat{\xi}_{lm}(z; \mathbf{q}) e^{i\mathbf{q}\cdot\mathbf{r}}.\end{aligned}\quad (38)$$

Then one obtains the identity (see App. 3.2.11.3),

$$\begin{aligned}\hat{\xi}_{lm}(z; \mathbf{q}) &= i^{-l} \frac{2\pi}{k} H(k - q) \frac{1}{2q_z(q)} \sum_{\sigma=\pm 1} e^{-i\sigma q_z(q)z} Y_{lm}(\theta_{Q\sigma}, \phi_Q) \\ &= i^{-l} \frac{2\pi}{k} H(k - q) Y_{lm}(\theta_{Q+}, \phi_Q) \frac{1}{2q_z(q)} [e^{-iq_z(q)z} + (-1)^{l+m} e^{iq_z(q)z}],\end{aligned}\quad (39)$$

which corresponds to two counter-propagating waves. Here the Heaviside function $H(s) = 1$ for $s > 0$, and vanishes for $s < 0$ [which then limits the integral (38) to real values of q_z]. As written, this form is valid only for real k , but it may be extended to complex k by replacing in the second line of (38) (a) k by $|k|$ everywhere, including $q_z(q) = \sqrt{|k|^2 - q^2}$, and (b) $e^{i\mathbf{q}\cdot\mathbf{r}} \rightarrow e^{i\zeta\mathbf{q}\cdot\mathbf{r}}$, where $\zeta = k/|k|$ is the phase of k .

The layered medium basis functions $\Psi_{lm}(\mathbf{x} - \mathbf{x}'; \mathbf{x}')$ are now obtained simply by extending (34) using (31): one obtains

$$\Psi_{lm}(\mathbf{x} - \mathbf{x}'; \mathbf{x}') = \int \frac{d^2q}{(2\pi)^2} e^{i\mathbf{q}\cdot(\mathbf{r}-\mathbf{r}')} \hat{\Psi}_{lm}(z, z'; \mathbf{q}) \quad (40)$$

with the Fourier kernel given by,

$$\begin{aligned}\hat{\Psi}_{lm}(z, z'; \mathbf{q}) &= i^{-(l+1)} \frac{4\pi}{k_v \rho_v} \left[\delta_{uv} a_v(q) e^{\sigma l_v(q)(z-z')} Y_{lm}(\theta_{Q\sigma}^v, \phi_Q) \right. \\ &\quad \left. + \sum_{\sigma', \tau'=\pm 1} a_{u\sigma'; v\tau'}(q) e^{\sigma' l_u(q)z + \tau' l_v(q)z'} Y_{lm}(\theta_{Q, -\tau'}^v, \phi_Q) \right]\end{aligned}\quad (41)$$

in which $\sigma = -\text{sgn}(z - z')$, the superscript on $\theta_{Q\sigma}^v$ indicates that k_v should be substituted for k in (37), and, again, z' lies in layer v while z lies in layer u . Thus, the differential operations on \mathbf{x}' [not on \mathbf{x} , as that would not lead to a new solution to the acoustic equation (2) due to the \mathbf{x} dependence of the parameters on the left hand side] that lead to the factor $Y_{lm}(\theta_{Q\sigma}^v, \phi_Q)$ in the diagonal term, must lead to the same factor, with $-\tau'$ replacing σ , in the non-diagonal terms.

3.2.7 T -matrix for a single homogeneous spherical scatterer

For a spherical scatterer the T -matrix is diagonal in the spherical wave expansion basis (10):

$$T_{lm,l'm'}(k) = T_l(k)\delta_{ll'}\delta_{mm'}. \quad (42)$$

For a homogeneous sphere of radius a , with interior parameters ρ_s, c_s (and exterior parameters ρ, c), the full solution takes the form

$$\phi(\mathbf{x}) = Y_{lm}(\theta, \phi) \begin{cases} A_l j_l(k_s x), & x \leq a \\ j_l(kx) + T_l h_l^{(1)}(kx), & x > a, \end{cases} \quad (43)$$

which also serves to precisely define the T -coefficients. Continuity of ϕ and $\rho^{-1}\partial_x\phi$ at $x = a$ produce the conditions

$$\begin{bmatrix} j_l(ka) \\ j_l'(ka) \end{bmatrix} = \begin{bmatrix} j_l(k_s a) & -h_l^{(1)}(ka) \\ (\rho c/\rho_s c_s)j_l'(k_s a) & -h_l^{(1)'}(ka) \end{bmatrix} \begin{bmatrix} A_l \\ T_l \end{bmatrix}, \quad (44)$$

with solution

$$T_l(k) = -\frac{j_l(k_s a)j_l'(ka) - (\rho c/\rho_s c_s)j_l'(k_s a)j_l(ka)}{j_l(k_s a)h_l^{(1)'}(ka) - (\rho c/\rho_s c_s)j_l'(k_s a)h_l^{(1)}(ka)}. \quad (45)$$

In the limit of a small scatterer, $ka \ll 1$, one may use the asymptotic forms [18],

$$j_l(z) \approx \frac{z^l}{(2l+1)!!}, \quad h_l(z) \approx -i\frac{(2l-1)!!}{z^{l+1}}, \quad (46)$$

which, when substituted into (45), yield

$$T_l(k) \approx \frac{l j_l(k_s a) - (\rho/\rho_s)k_s a j_l'(k_s a)}{(l+1)j_l(k_s a) + (\rho/\rho_s)k_s a j_l'(k_s a)} \times \frac{i(k a)^{2l+1}}{(2l+1)!!(2l-1)!!}. \quad (47)$$

For strong contrast, $k_s/k \gg 1$ and $k_s a$ of order unity, the prefactor on the right hand side of (47) may have poles, corresponding to the resonances of a sphere confined in a perfectly rigid medium. For finite ka , these poles move below the real axis, indicating a finite decay rate due to radiation damping.

3.2.8 T -matrix for an array spheres in a homogeneous background

3.2.8.1 Array of small spheres We see from (47) that, for small spheres, the s -wave contribution, $T_0(k) \sim ka$, dominates, and the larger l scattering contributions may be dropped. In this case, the individual \mathbf{T} and \mathbf{U} matrices in (24) and (29) are single elements. For N small spheres in an infinite homogeneous space one obtains

$$\mathbf{T} = \begin{bmatrix} T_0^{(1)}(k) & 0 & \dots & 0 & 0 \\ 0 & T_0^{(2)}(k) & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & T_0^{(N-1)}(k) & 0 \\ 0 & 0 & \dots & 0 & T_0^{(N)}(k) \end{bmatrix} \quad (48)$$

in which the dependence of $T_0^{(j)}(k)$ on the index j is through the sphere parameters ρ_j, c_j, a_j (and we assume, clearly, that $ka_j \ll 1$ for all j). The inverse of this matrix is computed trivially.

Similarly, only the zeroth order outgoing wave $\psi_{00}(\mathbf{x}) = h_0^{(1)}(kx)/\sqrt{4\pi}$ contributes to the \mathbf{U} matrix:

$$\mathbf{U} = \begin{bmatrix} 0 & \psi_{00}(\mathbf{x}_s^{(1)} - \mathbf{x}_s^{(2)}) & \dots & \psi_{00}(\mathbf{x}_s^{(1)} - \mathbf{x}_s^{(N-1)}) & \psi_{00}(\mathbf{x}_s^{(1)} - \mathbf{x}_s^{(N)}) \\ \psi_{00}(\mathbf{x}_s^{(2)} - \mathbf{x}_s^{(1)}) & 0 & \dots & \psi_{00}(\mathbf{x}_s^{(2)} - \mathbf{x}_s^{(N-1)}) & \psi_{00}(\mathbf{x}_s^{(2)} - \mathbf{x}_s^{(N)}) \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \psi_{00}(\mathbf{x}_s^{(N-1)} - \mathbf{x}_s^{(1)}) & \psi_{00}(\mathbf{x}_s^{(N-1)} - \mathbf{x}_s^{(2)}) & \dots & 0 & \psi_{00}(\mathbf{x}_s^{(N-1)} - \mathbf{x}_s^{(N)}) \\ \psi_{00}(\mathbf{x}_s^{(N)} - \mathbf{x}_s^{(1)}) & \psi_{00}(\mathbf{x}_s^{(N)} - \mathbf{x}_s^{(2)}) & \dots & \psi_{00}(\mathbf{x}_s^{(N)} - \mathbf{x}_s^{(N-1)}) & 0 \end{bmatrix}. \quad (49)$$

The A coefficients are, in turn given by the values of the background field,

$$\hat{\phi}_b(\mathbf{x}) = \frac{ik\rho}{4\pi} \int_{V_T} d^3x' h_0^{(1)}(k|\mathbf{x} - \mathbf{x}'|) S(\mathbf{x}') \quad (50)$$

at the sphere centers:

$$\mathbf{A} = \begin{bmatrix} \hat{\phi}_b(\mathbf{x}_s^{(1)}) \\ \hat{\phi}_b(\mathbf{x}_s^{(2)}) \\ \vdots \\ \hat{\phi}_b(\mathbf{x}_s^{(N-1)}) \\ \hat{\phi}_b(\mathbf{x}_s^{(N)}) \end{bmatrix}. \quad (51)$$

For a finite array of spherical wave emitters, $S(\mathbf{x})$ is a superposition of discrete compact sources, and (50) also reduces to a discrete sum. Recalling that the B coefficients refer to the superposition of outgoing waves from each target, the final signal at an observation point \mathbf{x} is,

$$\hat{\phi}(\mathbf{x}) = \frac{1}{\sqrt{4\pi}} \sum_{j=1}^N B^{(j)} h_0^{(1)}(k|\mathbf{x} - \mathbf{x}_s^{(j)}|). \quad (52)$$

For a receiver array, the signal will be a superposition,

$$\hat{\phi}_R = \int_{V_R} d^3x R(\mathbf{x}) \hat{\phi}(\mathbf{x}), \quad (53)$$

in which $R(\mathbf{x})$ is the receiver sensitivity function.

3.2.8.2 More general array of spheres More generally, if ka is not small, let us suppose that we keep angular momentum contributions $0 \leq l \leq l_j$ for sphere j . The matrix \mathbf{T} in (24) now consists of diagonal blocks $\mathbf{T}^{(j)}$ of dimension $1 + 3 + 5 + \dots + (2l_j + 1) = (l_j + 1)^2$. With each $0 \leq l \leq l_j$ is associated a diagonal sub-block of dimension $2l + 1$ with the values $T_l^{(j)}(k)$ along the diagonal. Note that, purely on symmetry grounds, this structure is maintained for any spherically-symmetric scatterer (including, for example, an elastic solid or shell), as long as one substitutes in place of (45) the correct forms for the $T_l(k)$.

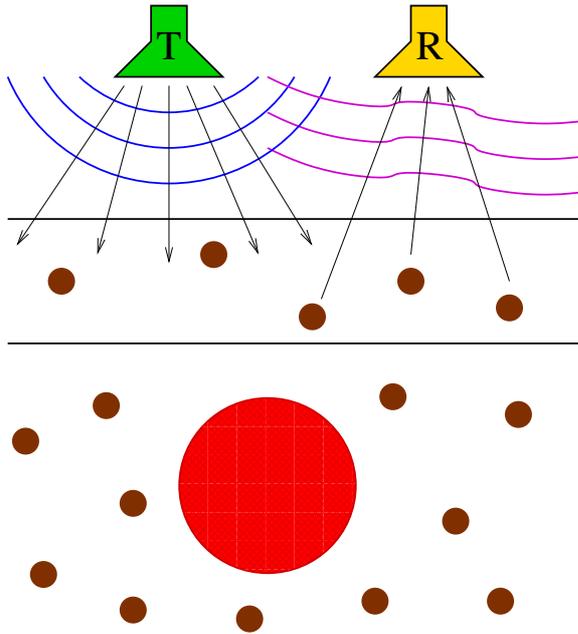


Figure 7: Schematic illustration of the measurement geometry for one large scatterer, in an array of small scatterers, buried in a layered background.

Correspondingly, the matrix $\mathbf{U}^{(jj')}$ is of dimension $(l_j + 1)^2 \times (l_{j'} + 1)^2$. Its components are

$$U_{l_m; l'_m}^{(jj')} = c_{l_m}^{(l'_m)} [k(\mathbf{x}_s^{(j')} - \mathbf{x}_s^{(j)})], \quad (54)$$

in which the c -coefficients are given by (95) (with $\mathbf{x} \rightarrow \mathbf{x} - \mathbf{x}_s^{(j)}$, and $\mathbf{x}' \rightarrow \mathbf{x}_s^{(j')} - \mathbf{x}_s^{(j)}$). A special case is that of a single large sphere immersed in an array of small spheres. Then only $l_1 > 0$, and $\mathbf{U}^{(1j')}$ and $\mathbf{U}^{(j1)}$ are, respectively, column and row vectors of dimension $(l_1 + 1)^2$. All others are the single elements $\psi_{00}(\mathbf{x}_s^{(j)} - \mathbf{x}_s^{(j')})$.

3.2.9 T -matrix for an array of scatterers in a layered background

Consider now the \mathbf{U} matrix entries for a layered background. We assume that each scatterer lies sufficiently far from its layer boundaries that a sufficiently large neighborhood of it exists within that layer on which the basis function expansion coefficients may be defined and related unambiguously through the T -matrix.

Although the re-expansion of the diagonal term in (31) gives rise to the analytic series (54) [with (95)], the non-diagonal terms must be handled numerically via direct integration. Using the general expression (88) for the expansion coefficients of the series (85), the corresponding

expansion (26) of the layered medium outgoing wave basis function (40) has coefficients,

$$\begin{aligned}
 U_{lm;l'm'}^{(jj')} &= i^{l-l'-1} \frac{(4\pi)^2}{k_v \rho_v} \int \frac{d^2 q}{(2\pi)^2} e^{i\mathbf{q} \cdot (\mathbf{r}_s^{(j)} - \mathbf{r}_s^{(j')})} \\
 &\quad \sum_{\sigma, \tau = \pm 1} a_{u\sigma;v\tau}(q) e^{\sigma l_u(q) z_s^{(j)} + \tau l_v(q) z_s^{(j')}} Y_{l'm'}^v(\theta_{Q,-\tau}^v, \phi_Q) Y_{lm}^u(\theta_{Q\sigma}^u, \phi_Q) \\
 &\quad + (1 - \delta_{jj'}) \delta_{uv} c_{lm}^{(l'm')} [k_v (\mathbf{x}_s^{(j')} - \mathbf{x}_s^{(j)})] \\
 &= i^{l-l'-1} i^{m'-m} \beta_{lm} \beta_{l'm'} e^{i(m'-m)\phi^{(jj')}} \frac{8\pi}{k_v \rho_v} \int_0^\infty q dq J_{m'-m}(q |\mathbf{r}_s^{(j)} - \mathbf{r}_s^{(j')}|) \\
 &\quad \times \sum_{\sigma, \tau = \pm 1} a_{u\sigma;v\tau}(q) e^{\sigma l_u(q) z_s^{(j)} + \tau l_v(q) z_s^{(j')}} P_{l'}^{m'}[\tau \sqrt{1 - (q/k_v)^2}] P_l^m[-\sigma \sqrt{1 - (q/k_u)^2}] \\
 &\quad + (1 - \delta_{jj'}) \delta_{uv} c_{lm}^{(l'm')} [k_v (\mathbf{x}_s^{(j')} - \mathbf{x}_s^{(j)})], \tag{55}
 \end{aligned}$$

and these comprise the individual entries in the \mathbf{U} -matrix (29). Here, as described in App. 3.2.11.3.1, for $q^2 > k_u^2$, where $\theta_{Q\sigma}^u$ becomes complex [see (37)], one interprets $Y_{lm}^*(\theta_{Q\sigma}, \phi_Q) = Y_{lm}(\theta_{Q\sigma}, -\phi_Q) = (-1)^m Y_{l,-m}(\theta_{Q\sigma}, \phi_Q)$ [see (12)]. Note that if $j = j'$ the diagonal (δ_{uv}) term is absent, and, since $J_m(0) = \delta_{m0}$, the remaining (non-diagonal) term is nonzero only for $m = m'$. Here, $z_s^{(j')}$ is in layer v , $z_s^{(j)}$ is in layer u , $J_m(z) = (-1)^m J_{-m}(z)$ are the usual cylindrical Bessel functions, $P_l^m(z)$ are the associated Legendre functions, the coefficients β_{lm} were defined in (11), and $\theta^{(jj')}$, $\phi^{(jj')}$ are the polar coordinate angles corresponding to the separation vector $\mathbf{x}_s^{(j)} - \mathbf{x}_s^{(j')}$ [which corresponds to $-\mathbf{x}'$ in (95)]. We use here, according to (37), $\cos(\theta_{Q\sigma}^u) = -\sigma \sqrt{1 - (q/k_u)^2} = i\sigma l_u(q)/k_u$, and $\cos(\theta_{Q,-\tau}^v) = \tau \sqrt{1 - (q/k_v)^2} = -i\tau l_v(q)/k_v$. The connection to $l(q)$ exhibited here provides the analytic continuation for $q > k$. Since $P_l^m(z)$ is of the form $(1 - z^2)^{m/2}$ times a polynomial of degree $l - m$ in z (even for $l - m$ even, and odd for $l - m$ odd) [18], the result here is $(q/k)^m$ times a polynomial of degree $l - m$ in $\pm il(q)/k$ (with the addition of appropriate indices understood here).

3.2.9.1 Single large scatterer in an array of small scatterers As an example, as illustrated in Fig. 7, we construct explicit forms of the \mathbf{U} matrix for a single large scatterer ($j = 1$) in an array of N_s of small ones ($j = 2, 3, \dots, N_s + 1$). The latter are treated in the s -wave approximation. The entries in (29) are as follows. For $j, j' \neq 1$ the sub-blocks are single elements,

$$\begin{aligned}
 U_{00;00}^{(jj')} &= \sqrt{4\pi} \Psi_{00}(\mathbf{x}_s^{(j)} - \mathbf{x}_s^{(j')}; \mathbf{x}_s^{(j')}) = \frac{4\pi}{ik_v \rho_v} g(\mathbf{x}_s^{(j)}, \mathbf{x}_s^{(j')}) \tag{56} \\
 &= \delta_{uv} (1 - \delta_{jj'}) h_0^{(1)}(k_v |\mathbf{x}_s^{(j)} - \mathbf{x}_s^{(j')}|) \\
 &\quad + \frac{2}{ik_v \rho_v} \int_0^\infty q dq J_0(q |\mathbf{r}_s^{(j)} - \mathbf{r}_s^{(j')}|) \sum_{\sigma, \tau = \pm 1} a_{u\sigma;v\tau} e^{\sigma l_u(q) z_s^{(j)} + \tau l_v(q) z_s^{(j')}} ,
 \end{aligned}$$

where we have noted that $P_0^0(z) \equiv 1$. Next, the $j = j' = 1$ block has dimension $(l_1 + 1)^2 \times (l_1 + 1)^2$, where l_1 is the maximum angular momentum index chosen for the large scatterer. The entries are

given, according to (55), by,

$$U_{lm;l'm'}^{(11)} = \delta_{mm'} i^{l-l'-1} \beta_{lm} \beta_{l'm'} \frac{8\pi}{k_v \rho_v} \int_0^\infty q dq \sum_{\sigma, \tau = \pm 1} a_{v\sigma;v\tau}(q) e^{(\sigma+\tau)l_v(q)z_s^{(1)}} \times P_l^m[\tau \sqrt{1 - (q/k_v)^2}] P_l^m[-\sigma \sqrt{1 - (q/k_v)^2}]. \quad (57)$$

These are the same coefficients as those that appear in Ref. [31], where the case of a single elastic sphere was studied.

Finally, when exactly one of j, j' is unity, one obtains:

$$\begin{aligned} U_{00;lm}^{(j1)} &= i^{m-l-1} \beta_{lm} e^{im\phi^{(j1)}} \frac{4\sqrt{\pi}}{k_v \rho_v} \int_0^\infty q dq J_m(q|\mathbf{r}_s^{(j)} - \mathbf{r}_s^{(1)}|) \\ &\quad \times \sum_{\sigma, \tau = \pm 1} a_{u\sigma;v\tau}(q) e^{\sigma l_u(q)z_s^{(j)} + \tau l_v(q)z_s^{(1)}} P_l^m[\tau \sqrt{1 - (q/k_v)^2}] \\ &\quad + \delta_{uv} \sqrt{4\pi} h_l^{(1)}(k_v |\mathbf{x}_s^{(j)} - \mathbf{x}_s^{(1)}|) Y_{lm}(\theta^{(j1)}, \phi^{(j1)}) \\ U_{lm;00}^{(1j)} &= i^{m-l-1} \beta_{lm} e^{-im\phi^{(j1)}} \frac{4\sqrt{\pi}}{k_v \rho_v} \int_0^\infty q dq J_m(q|\mathbf{r}_s^{(j)} - \mathbf{r}_s^{(1)}|) \\ &\quad \times \sum_{\sigma, \tau = \pm 1} a_{u\sigma;v\tau}(q) e^{\sigma l_u(q)z_s^{(j)} + \tau l_v(q)z_s^{(1)}} P_l^m[\tau \sqrt{1 - (q/k_v)^2}] \\ &\quad + \delta_{uv} \sqrt{4\pi} h_l^{(1)}(k_v |\mathbf{x}_s^{(j)} - \mathbf{x}_s^{(1)}|) Y_{lm}^*(\theta^{(j1)}, \phi^{(j1)}) \end{aligned} \quad (58)$$

in which, again, $\theta^{(jj')} = \pi - \theta^{(j'j)}$, $\phi^{(jj')} = \pi + \phi^{(j'j)}$ are the polar coordinate angles associated with the vector $\mathbf{x}_s^{(j)} - \mathbf{x}_s^{(j')}$, and we have noted the identities (13) and (14). In both cases (with the aid of the symmetry $a_{u\sigma;v\tau} = a_{v\tau;u\sigma}$), it has been assumed that $\mathbf{x}_s^{(1)}$ lies in layer v , while $\mathbf{x}_s^{(j)}$ lies in layer u . The diagonal terms are most simply derived directly from (88) using $Y_{00} = 1/\sqrt{4\pi}$ and (36).

Note that since the boundary conditions at infinity are such that there can only be an upward propagating wave in the first (highest) layer, and a downward propagating wave in the N th (lowest) layer, one has $a_{u\sigma;N-} = a_{u\sigma;1+} = 0$ [plus all symmetry-related vanishings, as seen explicitly in (100)]. Thus, for example, in a situation where all of the scatterers lie in the lowest layer, the sum over σ, τ in (55)–(58) reduces to the single term with $a_{N+;N+}$.

3.2.10 Transmitted and received signals

Once the matrices \mathbf{T} and \mathbf{U} are constructed, and the matrix \mathcal{T} is computed via (21), one is finally in a position to solve the original acoustic problem.

First, one needs to compute the \mathbf{A} coefficients for a given source $\hat{S}(\mathbf{x})$. From (4), (30), (31) and (85)–(88), one obtains for a layered medium,

$$\begin{aligned} A_{lm}^{(j)} &= 4\pi i^l \int \frac{d^2q}{(2\pi)^2} \sum_{\sigma, \tau = \pm 1} \hat{S}[\mathbf{q}, i\tau l_v(q)] a_{u\sigma;v\tau}(q) Y_{lm}^*(\theta_{Q\sigma}^u, \phi_Q) e^{i[\mathbf{q}, -i\sigma l_u(q)] \cdot \mathbf{x}_s^{(j)}} \\ &\quad + \delta_{uv} \rho_v \begin{cases} 4\pi i^l \int \frac{d^3Q}{(2\pi)^3} \frac{\hat{S}(\mathbf{Q})}{Q^2 - k_v^2 - i\eta} Y_{lm}^*(\theta_Q, \phi_Q) e^{i\mathbf{Q} \cdot \mathbf{x}_s^{(j)}}, & \text{(Fourier space form)} \\ ik_v \int d^3x' \hat{S}(\mathbf{x}') h_l^{(1)}(k_v |\mathbf{x}' - \mathbf{x}_s^{(j)}|) Y_{lm}^*(\theta^{(lj)}, \phi^{(lj)}), & \text{(Real space form)} \end{cases} \end{aligned} \quad (59)$$

where the real space form of the diagonal term will probably be most appropriate for compact sources. Here $\eta \rightarrow 0^+$ is a positive infinitesimal that selects the outgoing wave Green free function, the source is in layer v , the scatterer in layer u , $\theta^{(lj)}$, $\phi^{(lj)}$ are the polar coordinate angles associated with the vector $\mathbf{x}' - \mathbf{x}_s^{(j)}$, and we have defined the 3D Fourier transform of the source function:

$$\hat{S}(\mathbf{Q}) \equiv \hat{S}(\mathbf{q}, q_z) = \int d^3x' \hat{S}(\mathbf{x}') e^{-i\mathbf{Q}\cdot\mathbf{x}'}. \quad (60)$$

In the second (diagonal) term in (59), \mathbf{Q} is real, and there is no ambiguity in the definitions of the polar coordinate angles θ_Q, ϕ_Q associated with the vector \mathbf{Q} . In the first (off-diagonal) term, one interprets, as usual, $Y_{lm}^*(\theta_{Q\sigma}^u, \phi_Q) = (-1)^m Y_{l,-m}(\theta_{Q\sigma}^u, \phi_Q)$ irrespective of whether $\theta_{Q\sigma}^u$ is real or complex [see (37) and the discussion in App. 3.2.11.3.1].

With (59) as input to (21), the received signal emerges in the form,

$$\hat{\phi}_R = \sum_{j,l,m} B_{lm}^{(j)} \hat{\phi}_{R,lm}^{(j)} \quad (61)$$

with individual basis function and scatterer coefficients,

$$\begin{aligned} \hat{\phi}_{R,lm}^{(j)} &= \int d^3x \hat{R}(\mathbf{x}) \Psi_{lm}^{(j)}(\mathbf{x} - \mathbf{x}_s^{(j)}; \mathbf{x}_s^{(j)}) \\ &= i^{-(l+1)} \frac{4\pi}{k_v \rho_v} \int \frac{d^2q}{(2\pi)^2} \sum_{\sigma,\tau=\pm 1} \hat{R}[-\mathbf{q}, i\sigma l_u(q)] a_{u\sigma;v\tau}(q) Y_{lm}(\theta_{Q,-\tau}^v, \phi_Q) e^{-i[\mathbf{q}, i\tau l_v(q)]\cdot\mathbf{x}_s^{(j)}} \\ &\quad + \delta_{uw} \begin{cases} i^{-(l+1)} \frac{4\pi}{k_v} \int \frac{d^3Q}{(2\pi)^3} \frac{\hat{R}(-\mathbf{Q})}{Q^2 - k_v^2 - i\eta} Y_{lm}(\theta_Q, \phi_Q) e^{i\mathbf{Q}\cdot\mathbf{x}_s^{(j)}}, & \text{(Fourier space form)} \\ \int d^3x' \hat{R}(\mathbf{x}') h_l^{(1)}(k_v |\mathbf{x}' - \mathbf{x}_s^{(j)}|) Y_{lm}(\theta^{(lj)}, \phi^{(lj)}), & \text{(Real space form)} \end{cases} \end{aligned} \quad (62)$$

in which $\hat{R}(\mathbf{x})$ is the receiver sensitivity function, with Fourier transform defined analogously to (60), the scatterer is in layer v , the receiver in layer u , and (40), with (41), has been used to represent the basis functions. For the diagonal (δ_{uw}) term, the second equality in (84) has been used to obtain the Fourier space form. Once again, the real space form is probably most appropriate for compact receivers.

3.2.10.1 Special cases Consider now a few special cases. For an isotropic point source, $\hat{S}(\mathbf{x}) = S_0 \delta(\mathbf{x} - \mathbf{x}_T)$, hence $\hat{S}(\mathbf{Q}) = S_0 e^{-i\mathbf{Q}\cdot\mathbf{x}_T}$, equation (59) reduces to

$$\begin{aligned} A_{lm}^{(j)} &= 2i^{l-m} S_0 \beta_{lm} e^{-im\phi^{(Tj)}} \int_0^\infty q dq J_m(q|\mathbf{r}_T - \mathbf{r}_s^{(j)}|) \\ &\quad \sum_{\sigma,\tau=\pm 1} a_{u\sigma;v\tau}(q) P_l^m[-\sigma\sqrt{1 - (q/k_u)^2}] e^{\sigma l_u(q)z_s^{(j)} + \tau l_v(q)z_T} \\ &\quad + \delta_{uw} i S_0 k_v \rho_v h_l^{(1)}(k_v |\mathbf{x}_T - \mathbf{x}_s^{(j)}|) Y_{lm}^*(\theta^{(Tj)}, \phi^{(Tj)}), \end{aligned} \quad (63)$$

where β_{lm} is defined in (11), and $\theta^{(Tj)}, \phi^{(Tj)}$ are the polar coordinate angles associated with the vector $\mathbf{x}_T - \mathbf{x}_s^{(j)}$.

Similarly, for an isotropic point receiver, $\hat{R}(\mathbf{x}) = R_0\delta(\mathbf{x} - \mathbf{x}_R)$, hence $\hat{R}(\mathbf{Q}) = R_0e^{-i\mathbf{Q}\cdot\mathbf{x}_R}$, equation (62) reduces to

$$\begin{aligned}\hat{\phi}_{R,lm}^{(j)} &= R_0\Psi_{lm}^{(j)}(\mathbf{x}_R; \mathbf{x}_s^{(j)}) \\ &= i^{m-l-1}\frac{2R_0}{k_v\rho_v}\beta_{lm}e^{im\phi^{(Rj)}}\int_0^\infty qdqJ_m(q|\mathbf{r}_R - \mathbf{r}_s^{(j)}|) \\ &\quad \sum_{\sigma,\tau=\pm 1} a_{u\sigma;v\tau}(q)P_l^m[\tau\sqrt{1 - (q/k_v)^2}]e^{\sigma l_u(q)z_R + \tau l_v(q)z_s^{(j)}} \\ &\quad + \delta_{uv}R_0h_l^{(1)}(k_v|\mathbf{x}_R - \mathbf{x}_s^{(j)}|)Y_{lm}(\theta^{(Rj)}, \phi^{(Rj)}),\end{aligned}\quad (64)$$

where $\theta^{(Rj)}, \phi^{(Rj)}$ are the polar coordinate angles associate with the vector $\mathbf{x}_R - \mathbf{x}_s^{(j)}$.

Consider next a phased vertical line transmitter array:

$$\begin{aligned}\hat{S}(\mathbf{x}) &= S_0e^{ik_T(z-z_T)}\delta(\mathbf{r} - \mathbf{r}_T), \\ &\quad -L_T/2 \leq z - z_T \leq L_T/2 \\ \hat{S}(\mathbf{Q}) &= S_0e^{-i\mathbf{Q}\cdot\mathbf{x}_T}\frac{\sin[(k_T - q_z)L_T/2]}{(k_T - q_z)/2},\end{aligned}\quad (65)$$

If one writes $k_T = k_v \cos(\alpha)$, the array generates a cylindrical wave in the direction corresponding to polar angle α . The corresponding coefficients $A_{lm}^{(j)}$ are computed simply by inserting an extra factor,

$$\hat{S}[\mathbf{q}, i\tau l_v(q)]/S_0 = \frac{\sin\{[k_T - i\tau l_v(q)]L_T/2\}}{[k_T - i\tau l_v(q)]/2},\quad (66)$$

into the summand of the first (non-diagonal) term in (63). The second (diagonal) term is probably most efficiently evaluated by directly integrating the corresponding term in (63) to obtain:

$$\begin{aligned}A_{lm}^{(j,\text{diag})} &= \delta_{uv}iS_0k_v\rho_v\int_{-L_T/2}^{L_T/2} d\zeta e^{ik_T\zeta} \\ &\quad \times h_l^{(1)}(k_v|\mathbf{x}_T(\zeta) - \mathbf{x}_s^{(j)}|)Y_{lm}^*[\theta^{(Tj)}(\zeta), \phi^{(Tj)}],\end{aligned}\quad (67)$$

in which $\mathbf{x}_T(\zeta) - \mathbf{x}_s^{(j)} = \mathbf{x}_T - \mathbf{x}_s^{(j)} + \zeta\hat{\mathbf{z}}$, and $\theta^{(Tj)}(\zeta), \phi^{(Tj)}$ are the polar coordinate angles associated with this vector (the latter being independent of ζ).

Similarly, for a phased vertical receiver line array:

$$\begin{aligned}\hat{R}(\mathbf{x}) &= R_0e^{ik_R(z-z_R)}\delta(\mathbf{r} - \mathbf{r}_R), \\ &\quad -L_R/2 \leq z - z_R \leq L_R/2 \\ \hat{R}(\mathbf{Q}) &= R_0e^{-i\mathbf{Q}\cdot\mathbf{x}_R}\frac{\sin[(k_R - q_z)L_R/2]}{(k_R - q_z)/2},\end{aligned}\quad (68)$$

If one writes $k_R = k_u \cos(\beta)$, the array is most sensitive to waves arriving from directions corresponding to polar angle α . The corresponding coefficients $\phi_{R,lm}^{(j)}$ are computed simply by inserting an extra factor,

$$\hat{R}[-\mathbf{q}, i\sigma l_u(q)]/R_0 = \frac{\sin\{[k_R - i\sigma l_u(q)]L_R/2\}}{[k_R - i\sigma l_u(q)]/2},\quad (69)$$

into the summand of the first (non-diagonal) term in (64). The second (diagonal) term is probably, again, most efficiently evaluated by directly integrating the corresponding term in (64) to obtain:

$$\begin{aligned} \hat{\phi}_{R,lm}^{(j,\text{diag})} &= \delta_{uv} R_0 \int_{-L_R/2}^{L_R/2} d\zeta e^{ik_R \zeta} \\ &\times h_l^{(1)}(k_v |\mathbf{x}_R(\zeta) - \mathbf{x}_s^{(j)}|) Y_{lm}[\theta^{(Rj)}(\zeta), \phi^{(Rj)}], \end{aligned} \quad (70)$$

in which $\mathbf{x}_R(\zeta) - \mathbf{x}_s^{(j)} = \mathbf{x}_R - \mathbf{x}_s^{(j)} + \zeta \hat{\mathbf{z}}$, and $\theta^{(Rj)}(\zeta), \phi^{(Rj)}$ are the polar coordinate angles associated with this vector (the latter being independent of ζ).

3.2.11 Inverse scattering problems

The theory presented thus far is intended to support forward modeling of scattered signals from known targets in a known background. Here we briefly discussed an associated inverse problem, namely extraction of T -matrix coefficients from measurements on a known target, e.g., for later use in a forward modeling calculation. Such measurement-based model inputs are an important tool when the target complexity exceeds ones computational ability to directly compute the desired coefficients (or has too many unknown parameters to enable robust modeling). For example, an elastic target may possess a spectrum of internal resonances with very sensitive dependence on its precise internal structure, and on the parameters of the surrounding medium. A series of measurements may provide a better fingerprint of the target response than a low fidelity computation. In what follows, bulk absorption will be neglected and we will neglect reflections from the experimental system boundaries.

3.2.11.1 Target parameter inversion in a homogeneous background For a target in an infinite homogeneous background, one typically considers plane waves, $\hat{\phi}_b(\mathbf{x}) = A_0 e^{i\mathbf{K} \cdot \mathbf{x}}$, of various frequencies incident from various directions. Using the spherical harmonic basis (10) and the plane wave identity (81), one obtains the incoming wave coefficients,

$$A_{lm} = 4\pi A_0 i^l Y_{lm}^*(\theta_K, \phi_K). \quad (71)$$

From (9) and (15), at sufficiently large distance from the target, $Kx \gg 1$, where the asymptotic form $h_l(z) \approx (-i)^{l+1} e^{iKz}/z$ for the Hankel function is valid, the scattered wave takes the form,

$$\begin{aligned} \hat{\phi}_s(\mathbf{x}; \mathbf{K}) &\approx 4\pi A_0 \frac{e^{iKx}}{ix} \sum_{l,m;l',m'} i^{l'-l} T_{lm;l'm'}(K) \\ &\times Y_{l'm'}^*(\theta_K, \phi_K) Y_{lm}(\theta, \phi). \end{aligned} \quad (72)$$

If, for given wavenumber magnitude $K = \omega/c$, one were able to perform a dense set of measurements at fixed distance x_R , for all incoming plane wave and outgoing spherical wave directions, one may extract the desired coefficients by integration:

$$\begin{aligned} T_{lm;l'm'}(K) &= \frac{i^{l-l'+1}}{4\pi A_0} x_R e^{-iKx_R} \int d\Omega \int d\Omega_K \\ &\times Y_{l'm'}(\theta_K, \phi_K) Y_{lm}^*(\theta, \phi) \hat{\phi}_s(\mathbf{x}; \mathbf{K}). \end{aligned} \quad (73)$$

In the absence of full coverage, measurements on a finite grid of points \mathbf{x}_j , \mathbf{K}_j may be used to invert (72) under the assumption that a suitably truncated finite subset of the T -coefficients dominate the sum.

In the special case of spherically symmetric target, (16) and the addition theorem for spherical harmonics [18] allows one to reduce (72) and (73) to the form,

$$\begin{aligned}\hat{\phi}_s(\mathbf{x}; \mathbf{K}) &= A_0 \frac{e^{iKx}}{ix} \sum_l (2l+1) T_l(K) P_l(u_{Kx}) \\ T_l(K) &= \frac{ix_R e^{-iKx_R}}{(2l+1)A_0} \int_{-1}^1 du \hat{\phi}_s[\mathbf{x}(u); \mathbf{K}] P_l(u)\end{aligned}\quad (74)$$

in which P_l are the usual Legendre polynomials, $u_{Kx} = \hat{\mathbf{K}} \cdot \hat{\mathbf{x}} = \cos(\theta_{Kx})$ is the cosine of the angle between the two unit vectors, and the integral in the second line is over the great circle $\hat{\mathbf{Q}} \cdot \hat{\mathbf{x}}(u) = u$. If one assumes that $T_l(K)$ is effectively zero for $l > l_{\max}$, then one may evaluate the remaining T -coefficients exactly via Gauss-Legendre integration:

$$T_l(K) = \frac{ix_R}{(2l+1)A_0} e^{-iKx_R} \sum_{j=1}^{l_{\max}} w_j \hat{\phi}_s[\mathbf{x}(u_j); \mathbf{K}] P_l(u_j) \quad (75)$$

in which u_j, w_j are the requisite points and weights for l_{\max} -point integration [37]. It should not be too surprising that only l_{\max} measurements are required, but by carefully selecting the observation points $\mathbf{x}(u_j)$, (75) provides an explicit inversion formula.

3.2.11.2 Buried target parameter inversion Since target scattering parameters may depend strongly on the background parameters (through the appropriate surface boundary conditions), measurements in different background media may be required. However, laboratory measurements with target, transmitter, and receivers all surrounded by sediment, for example, may be very difficult to perform. Instead, one might hope to infer the T -matrix parameters from a two-layer geometry, in which the target is buried in a homogeneous sediment, while the transmitter and receiver lie in the homogeneous fluid above it. The boundary between the layers will be taken as $z = 0$, and we will again neglect bulk absorption and scattering from the boundaries of the experimental volume.

Consider therefore a downward-propagating plane wave, $A_0 e^{i\mathbf{K} \cdot \mathbf{x}}$, in the upper layer with wavevector $\mathbf{K} = (\mathbf{k}, -k_z)$, $k_z > 0$ (and magnitude $|\mathbf{K}| = k_1 = \omega/c_1$). Since there are only two layers, this generates a single downward propagating wave in the lower layer,

$$\hat{\phi}_b(\mathbf{x}) = A_0 a_{2+;1-}(k) e^{i\mathbf{K}_2 \cdot \mathbf{x}}, \quad (76)$$

with two-layer amplitude $a_{2+;1-}(k) = a_{1-;2+}(k) = 1/(\chi_1 + \chi_2)$ given by (101), wavevector $\mathbf{K}_2 = (\mathbf{k}, -k_{z,2})$, and $k_{z,2} = \sqrt{(\omega/c_2)^2 - |\mathbf{k}|^2} > 0$. The A -coefficients therefore take the form (71) with \mathbf{k} replaced by \mathbf{k}_2 . The incident background amplitudes are therefore,

$$A_{lm} = A_0 e^{i\mathbf{K}_2 \cdot \mathbf{x}_s} a_{2+;1-}(q) i^l Y_{lm}(\theta_K^2, \phi_K), \quad (77)$$

in which θ_K^2, ϕ_K are the polar angles associated with the wavevector \mathbf{K}_2 . From (41), the scattered field takes the form,

$$\hat{\phi}_s(\mathbf{x}) = \sum_{l,m} B_{lm} \Psi_{lm}(\mathbf{x} - \mathbf{x}_s; \mathbf{x}_s), \quad (78)$$

in which,

$$\begin{aligned}\Psi_{lm}(\mathbf{x} - \mathbf{x}_s; \mathbf{x}_s) &= \int \frac{d^2q}{(2\pi)^2} \hat{\Psi}_{lm}(z, z_s; \mathbf{q}) e^{i\mathbf{q} \cdot (\mathbf{r} - \mathbf{r}_s)} \\ \hat{\Psi}_{lm}(z, z_s; \mathbf{q}) &= i^{-(l+1)} \frac{4\pi}{k_2 \rho_2} Y_{lm}(\theta_{Q-}, \phi_Q) \\ &\times e^{-iq_{z,2}(q)l_2(q)z_s} \begin{cases} a_{1-;2+}(q) e^{iq_{z,1}(q)z}, & z > 0 \\ a_{2+;2+}(q) e^{-iq_{z,2}(q)z}, & z < 0 \end{cases},\end{aligned}\quad (79)$$

where $a_{2+;2+}(q) = (\chi_2 - \chi_1)/2\chi_2(\chi_1 + \chi_2)$, and $\theta_{Q\pm}, \phi_Q$ are the polar coordinates associated the wavevectors $\mathbf{Q}_{2\pm} = [\mathbf{q}, \mp q_{z,2}(q)]$ (with magnitude $|\mathbf{Q}_{2\pm}| = k_2 = \omega/c_2$) in the lower layer. The $z > 0$ form is relevant for receivers in the upper layer. Finally, the B -coefficients are obtained from (21) with \mathbf{U} matrix provided by the re-expansion (55) of the lower layer ($z < 0$) form of (79):

$$\begin{aligned}U_{lm;l'm'} &= i^{l-l'-1} \frac{(4\pi)^2}{k_2 \rho_2} \int \frac{d^2q}{(2\pi)^2} a_{2+;2+}(q) e^{-2iq_{z,2}(q)z_s} \\ &\times Y_{l'm'}(\theta_{Q-}, \phi_Q) Y_{lm}^*(\theta_{Q+}, \phi_Q) \\ &= \delta_{mm'} i^{l-l'-1} \beta_{lm} \beta_{l'm'} \frac{8\pi}{k_2 \rho_2} \\ &\times \int_0^\infty q dq a_{2+;2+}(q) e^{-2iq_{z,2}(q)z_s} \\ &\times P_l^m[\sqrt{1 - (q/k_2)^2}] P_{l'}^m[-\sqrt{1 - (q/k_2)^2}],\end{aligned}\quad (80)$$

which is just (57) specialized to the two layer system.

Given the position of the target and the background parameters, one may in principle accurately compute the A coefficients, the scattering functions Ψ_{lm} , and thereby the U -matrix coefficients. Given a sufficiently dense set of measurements, a subset of the B -coefficients may then be found by inverting (78), perhaps through a SVD approach [37]. Equation (20) then specifies a set of linear equations for a suitably truncated set of T -matrix elements, perhaps also soluble through a SVD approach. It is likely that with transmitter and receiver positions only in the upper layer, there will be ambiguities in the inversion, i.e., the data is sensitive to only certain linear combinations of B and T coefficients, even for only low order values of l . It may be necessary to have at least a few receivers placed in the lower layer.

3.2.11.3 Basis function expansions

3.2.11.3.1 Green function and plane wave expansions We make extensive use of the following identity [18],

$$e^{i\mathbf{Q} \cdot \mathbf{x}} = 4\pi \sum_{l,m} i^l j_l(Qx) Y_{lm}(\theta, \phi) Y_{lm}^*(\theta_Q, \phi_Q), \quad (81)$$

which relates a plane wave to the spherical wave basis functions, where (Q, θ_Q, ϕ_Q) represents the spherical coordinate representation of the wavevector \mathbf{Q} . For real \mathbf{Q} (81) is unambiguous, but we also require this identity for various complex wavevectors.

First, if one is able to write $\mathbf{Q} = Q\hat{\mathbf{Q}}$, where Q may be complex, but $\hat{\mathbf{Q}}$ is a real unit vector (thus, $\mathbf{Q} \cdot \mathbf{Q} = Q^2$), then (81) remains valid with θ_Q, ϕ_Q being the polar coordinate angles associated with $\hat{\mathbf{Q}}$. It is in this form that it leads to the identity (39): simply multiply both sides by $Y_{lm}(\theta_Q, \phi_Q)$, and integrate over all solid angles of \mathbf{Q} . The two counter-propagating waves come from the two hemispheres $\mathbf{Q}_{\pm} = (\mathbf{q}, \pm\sqrt{|k|^2 - q^2})$ after changing variables from (θ_Q, ϕ_Q) to $\mathbf{q} = |k| \sin(\theta_Q)[\cos(\phi_Q), \sin(\phi_Q)]$.

Second, and more importantly, we require (81) for complex wavevectors of the form $\mathbf{Q}_{\pm} = (\mathbf{q}, \pm\sqrt{k^2 - q^2})$, where \mathbf{q} is a real 2D vector (so that the azimuthal angle ϕ_Q remains well defined), and $q_z(q) = \sqrt{k^2 - q^2} = il(q)$ may be complex, either because k is complex (an absorptive medium leads to $\text{Im}(k) > 0$, so that $h_l^{(1)}(kx) \approx i^{-(l+1)}e^{ikx}/kx$ decays exponentially for large x), or because $q^2 > k^2$. For $q^2 > k^2$ we interpret $\theta_{Q\sigma}$ according to (37) and $Y_{lm}^*(\theta_Q, \phi_Q) \equiv Y_{lm}(\theta_{Q\sigma}, -\phi_Q) = (-1)^m Y_{l,-m}(\theta_{Q\sigma}, \phi_Q)$ [see (12)], so that the complex conjugation does not actually act on θ_Q .

Using this second form, the identity (36) may be derived from the Green function expansion [18],

$$\begin{aligned} h_0^{(1)}(k|\mathbf{x} - \mathbf{x}'|) &= \frac{e^{ik|\mathbf{x} - \mathbf{x}'|}}{ik|\mathbf{x} - \mathbf{x}'|} \\ &= 4\pi \sum_{l,m} j_l(kx_{<}) h_l^{(1)}(kx_{>}) Y_{lm}(\theta, \phi) Y_{lm}^*(\theta', \phi'), \end{aligned} \quad (82)$$

which provides the re-expansion (22) for the lowest order basis function ψ_{00} . Here $x_{>} = \max\{x, x'\}$ and $x_{<} = \min\{x, x'\}$. Since \mathbf{x}, \mathbf{x}' are always real (although k may be complex), there is no ambiguity here in the definitions of the spherical harmonics. The more general result for the higher order ψ_{lm} will be addressed in Sec. 3.2.11.3.3 below. The connection to (36) begins with the Fourier relation,

$$\begin{aligned} h_0^{(1)}(k|\mathbf{x} - \mathbf{x}'|) &= \frac{4\pi}{ik} \int \frac{d^3Q}{(2\pi)^3} \frac{e^{i\mathbf{Q} \cdot (\mathbf{x} - \mathbf{x}')}}{Q^2 - k^2 - i\eta} \\ &= \frac{4\pi}{ik} \int \frac{d^2q}{(2\pi)^2} \frac{e^{i\mathbf{Q}_\sigma \cdot (\mathbf{x} - \mathbf{x}')}}{2l(q)}, \end{aligned} \quad (83)$$

where $\eta \rightarrow 0^+$ is a positive infinitesimal (required for real k), $\mathbf{Q}_\sigma(\mathbf{q}) = [\mathbf{q}, -i\sigma l(q)]$, and $\sigma = -\text{sgn}(z)$. In the second equality, the q_z integral has picked up the pole at $\sqrt{k^2 - q^2 + i\eta}$ in the upper half plane for $z > 0$, and at $-\sqrt{k^2 - q^2 + i\eta}$ in the lower half plane for $z < 0$. By expanding the \mathbf{x}' dependence using (81) [noting that the polar coordinate angles associated with $-\mathbf{x}'$ are $\pi - \theta'$, $\pi + \phi'$, and using (13)], and comparing the resulting series to (82) (under the assumption $x' < x$), one identifies,

$$\begin{aligned} h_l^{(1)}(kx) Y_{lm}(\theta, \phi) &= \frac{4\pi}{k} i^{-(l+1)} \int \frac{d^2q}{(2\pi)^2} \frac{e^{i\mathbf{Q}_\sigma \cdot \mathbf{x}}}{2l(q)} \\ &\quad \times Y_{lm}(\theta_{Q\sigma}, \phi_Q) \\ &= \frac{4\pi}{k} i^{-(l+1)} \int \frac{d^3Q}{(2\pi)^3} \frac{e^{i\mathbf{Q} \cdot \mathbf{x}}}{Q^2 - k^2 - i\eta} \\ &\quad \times Y_{lm}(\theta_Q, \phi_Q). \end{aligned} \quad (84)$$

where $\theta_{Q\sigma}, \phi_Q$ are the polar coordinate angles associated with \mathbf{Q}_σ , defined in and above (37). The first equality in (84) coincides precisely with (36).

3.2.11.3.2 Re-expansion coefficient identities The local expansion of a regular solution $\phi(\mathbf{x})$ to the homogeneous acoustic equation is accomplished using the orthogonality properties of the spherical harmonics. Thus, if one seeks an expansion

$$\phi(\mathbf{x}) = \sum_{l,m} c_{lm} j_l(kx) Y_{lm}(\theta, \phi), \quad (85)$$

then one identifies

$$c_{lm} j_l(kx) = \int d\Omega \phi(\mathbf{x}) Y_{lm}^*(\theta, \phi). \quad (86)$$

The choice of reference radius x here is arbitrary, so long as the sphere of radius x lies entirely within a locally homogeneous region.

Alternatively, if one represents $\phi(\mathbf{x})$ in terms of its horizontal Fourier expansion,

$$\phi(\mathbf{x}) = \int \frac{d^2q}{(2\pi)^2} \sum_{\sigma=\pm 1} \hat{\phi}_\sigma(\mathbf{q}) e^{i\mathbf{Q}_\sigma \cdot \mathbf{x}}, \quad (87)$$

then by using (81) to represent the plane waves, one obtains

$$\begin{aligned} c_{lm} &= 4\pi i^l \int \frac{d^2q}{(2\pi)^2} \sum_{\sigma=\pm 1} \hat{\phi}_\sigma(\mathbf{q}) Y_{lm}^*(\theta_{Q\sigma}, \phi_Q) \\ &= 4\pi i^l \int \frac{d^2q}{(2\pi)^2} [\hat{\phi}_+(\mathbf{q}) + (-1)^{l+m} \hat{\phi}_-(\mathbf{q})] \\ &\quad \times Y_{lm}^*(\theta_{Q+}, \phi_Q), \end{aligned} \quad (88)$$

in which, for complex $\theta_{Q\sigma}$, we interpret $Y_{lm}^* = (-1)^m Y_{l,-m}$ as described above.

3.2.11.3.3 Re-expansion of outgoing wave basis functions In general the integral (88) must be performed numerically. Consider as a special case, however, the re-expansion of the outgoing wave basis functions $\psi_{lm}(\mathbf{x})$:

$$\begin{aligned} \psi_{lm}(\mathbf{x} - \mathbf{x}') &= i^{-(l+1)} \frac{4\pi}{k} \int \frac{d^2q}{(2\pi)^2} \frac{1}{2l(q)} Y_{lm}(\theta_{Q\sigma}, \phi_Q) \\ &\quad \times e^{i\mathbf{Q}_\sigma \cdot (\mathbf{x} - \mathbf{x}')} \\ &= \sum_{l',m'} c_{l'm'}^{(lm)}(k\mathbf{x}') j_{l'}(kx) Y_{l'm'}(\theta, \phi), \end{aligned} \quad (89)$$

in which $\sigma = -\text{sgn}(z - z')$. By expanding the \mathbf{x} -dependence using (81), one obtains

$$\begin{aligned} c_{l'm'}^{(lm)}(k\mathbf{x}') &= (-1)^{m'} i^{l'-l-1} \frac{(4\pi)^2}{k} \int \frac{d^2q}{(2\pi)^2} \frac{1}{2l(q)} \\ &\quad \times e^{-i\mathbf{Q}_\sigma \cdot \mathbf{x}'} Y_{lm}(\theta_{Q\sigma}, \phi_Q) Y_{l',-m'}(\theta_{Q\sigma}, \phi_Q). \end{aligned} \quad (90)$$

where we have explicitly substituted $Y_{lm}^* = (-1)^m Y_{l,-m}$. We now make use of the identity (see Eq. (16.89) of Ref. [26]),

$$\begin{aligned} Y_{l_1 m_1}(\theta, \phi) Y_{l_2 m_2}(\theta, \phi) \\ = \sum_{l=|l_1-l_2|}^{l_1+l_2} D_{l_1 m_1; l_2 m_2}^{l, m_1+m_2} Y_{l, m_1+m_2}(\theta, \phi), \end{aligned} \quad (91)$$

in which the D coefficients are related to the Clebsch-Gordon coefficients

$$C_{l_1 m_1, l_2 m_2}^{lm} \equiv \langle l_1 l_2 m_1 m_2 | l_1 l_2 l m \rangle$$

(which can be nonzero only if $m = m_1 + m_2$) via

$$D_{l_1 m_1, l_2 m_2}^{lm} = \sqrt{\frac{(2l_1 + 1)(2l_2 + 1)}{4\pi(2l + 1)}} C_{l_1 0; l_2 0}^{l 0} C_{l_1 m_1, l_2 m_2}^{lm} \quad (92)$$

which therefore also can be nonzero only if $m = m_1 + m_2$. The Clebsch-Gordon coefficients have the symmetry properties [26],

$$\begin{aligned} C_{l_1 m_1; l_2 m_2}^{lm} &= (-1)^{l-l_1-l_2} C_{l_2 m_2; l_1 m_1}^{lm} \\ &= (-1)^{l-l_1-l_2} C_{l_1, -m_1; l_2, -m_2}^{lm}, \end{aligned} \quad (93)$$

which implies the corresponding symmetries,

$$D_{l_1 m_1, l_2 m_2}^{lm} = D_{l_2 m_2, l_1 m_1}^{lm} = (-1)^{l-l_1-l_2} D_{l_1, -m_1; l_2, -m_2}^{lm}. \quad (94)$$

Inserting (91) into (90), one obtains

$$\begin{aligned} c_{l' m'}^{(lm)}(k\mathbf{x}') &= 4\pi(-1)^{m'} \sum_{l''=|l-l'|}^{l+l'} i^{l'-l-l''} D_{lm; l', -m'}^{l'', m-m'} \\ &\quad \times h_{l''}^{(1)}(kx') Y_{l'', m-m'}(\theta', \phi'), \end{aligned} \quad (95)$$

which expresses the coefficients as a finite series in the original outgoing wave basis functions. The series (82) is recovered when $l = m = 0$ since in this case only the term $l'' = l'$ survives, and since $Y_{00} = 1/\sqrt{4\pi}$, one obtains $D_{00; l', -m'}^{l', -m'} = 1/\sqrt{4\pi}$ as well.

The expansion (89) is expected to converge for $x' < x$. For $x' > x$ one should interchange $\mathbf{x} \leftrightarrow \mathbf{x}'$. Using the identity (13), one obtains

$$\begin{aligned} \psi_{lm}(\mathbf{x} - \mathbf{x}') &= (-1)^l \psi_{lm}(\mathbf{x}' - \mathbf{x}) \\ &= (-1)^l \sum_{l', m'} c_{l' m'}^{(lm)}(k\mathbf{x}) j_{l'}(kx') Y_{l' m'}(\theta', \phi'), \end{aligned} \quad (96)$$

in which $c_{l' m'}^{(lm)}(k\mathbf{x})$ is given by (95) with the appropriate substitutions.

3.2.12 Green function coefficients for a three layer system

We display here the results for the Green function a coefficients for the case of a three layer system, i.e., a single finite width layer sandwiched between two semi-infinite half-spaces. Let the boundaries be at $d_1 = d$ and $d_2 = -d$. Let c_u , ρ_u , $k_u = \omega/c_u$, and $l_u(q) = \sqrt{q^2 - k_u^2}$, denote the quantities in the upper half space ($u = 1$), the finite layer ($u = 2$), and the lower half space ($u = 3$), respectively. It should be recalled that, according to (32), for $q^2 - k_u^2$ real and negative we evaluate the square root on the lower side of the branch cut, defining $l_u(q) = -i\sqrt{k_u^2 - q^2} \equiv -iq_z(q)$. It is convenient to define the combination,

$$\chi_u = l_u/\rho_u, \quad (97)$$

along with,

$$\begin{aligned} N_\chi &= (\chi_1 - \chi_2)(\chi_2 + \chi_3)e^{4dl_2} + (\chi_1 + \chi_2)(\chi_2 - \chi_3) \\ \bar{N}_\chi &= (\chi_3 - \chi_2)(\chi_2 + \chi_1)e^{4dl_2} + (\chi_3 + \chi_2)(\chi_2 - \chi_1) \\ D_\chi &= (\chi_1 + \chi_2)(\chi_2 + \chi_3)e^{4dl_2} + (\chi_2 - \chi_1)(\chi_3 - \chi_2). \end{aligned} \quad (98)$$

In terms of (98) the, diagonal coefficients (33) are given by,

$$a_u(q) = \frac{1}{2\chi_u}, \quad (99)$$

and the off-diagonal coefficients $a_{u\sigma;v\tau}(q)$ are given by, [38]:

$$\mathbf{a}(q) = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & e^{2l_1d} \frac{1}{2\chi_1} \frac{N_\chi}{D_\chi} & e^{(l_1+3l_2)d} \frac{\chi_2+\chi_3}{D_\chi} & e^{(l_1+l_2)d} \frac{\chi_2-\chi_3}{D_\chi} & e^{(l_1+2l_2+l_3)d} \frac{2\chi_2}{D_\chi} & 0 \\ 0 & e^{(l_1+3l_2)d} \frac{\chi_2+\chi_3}{D_\chi} & \frac{1}{2\chi_2} e^{2l_2d} \frac{(\chi_2-\chi_1)(\chi_2+\chi_3)}{D_\chi} & \frac{1}{2\chi_2} \frac{(\chi_2-\chi_1)(\chi_2-\chi_3)}{D_\chi} & e^{(l_2+l_3)d} \frac{\chi_2-\chi_1}{D_\chi} & 0 \\ 0 & e^{(l_1+l_2)d} \frac{\chi_2-\chi_3}{D_\chi} & \frac{1}{2\chi_2} \frac{(\chi_2-\chi_1)(\chi_2-\chi_3)}{D_\chi} & \frac{1}{2\chi_2} e^{2l_2d} \frac{(\chi_2+\chi_1)(\chi_2-\chi_3)}{D_\chi} & e^{(3l_2+l_3)d} \frac{\chi_2+\chi_1}{D_\chi} & 0 \\ 0 & e^{(l_1+2l_2+l_3)d} \frac{2\chi_2}{D_\chi} & e^{(l_2+l_3)d} \frac{\chi_2-\chi_1}{D_\chi} & e^{(3l_2+l_3)d} \frac{\chi_2+\chi_1}{D_\chi} & \frac{1}{2\chi_3} e^{2l_3d} \frac{\bar{N}_\chi}{D_\chi} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}, \quad (100)$$

in which the rows are labeled by u, σ and the columns by v, τ in the sequence $(1+, 1- | 2+, 2- | 3+, 3-)$. The symmetry relation $a_{u\sigma;v\tau} = a_{v\tau;u\sigma}$ is evident.

By taking the limit $d \rightarrow 0$, the two layer results follow in the form

$$\mathbf{a}(q) = \left(\begin{array}{cc|cc} 0 & 0 & 0 & 0 \\ 0 & \frac{\chi_1-\chi_2}{2\chi_1(\chi_1+\chi_2)} & \frac{1}{\chi_1+\chi_2} & 0 \\ \hline 0 & \frac{1}{\chi_1+\chi_2} & \frac{\chi_2-\chi_1}{2\chi_2(\chi_1+\chi_2)} & 0 \\ 0 & 0 & 0 & 0 \end{array} \right), \quad (101)$$

in which all of the middle layer parameters drop out as expected, and the layer 3 parameters have now been relabeled with the index 2. The rows and columns are labeled in the sequence $(1+, 1- | 2+, 2-)$.

3.3 BOSS data processing methods

The basic processing element for BOSS image formation is 3D beamforming using a series of overlapped synthetic apertures. The beamformer is a time-domain processor that corrects the locations of each of the BOSS receiver array elements at every ping as it accumulates a fixed-length synthetic along-track aperture. The processor then applies near-field beamforming to each 3-d image pixel location, then shifts the synthetic aperture forward and repeats the process. A level, uniform, 3-d grid encompassing the entire data run is laid out, and the raw image data from the beamformer output is then mapped to the final output grid locations. This removes all redundant or overlapping beamformed data, and results in an image that accurately represents the vehicle track and maneuvers. We describe these steps in detail in Secs. 3.3.1 - 3.3.4. In Sec. A.7 we show various views of the resulting data products (derived from the raw data described in Sec. 4.3.1).

3.3.1 Signal processing

The BOSS transmits a linear FM pulse $x(t)$ with a sweep rate b [Hz/s] and instantaneous frequency $f_c = f_c + bt$, and with rectangular amplitude modulation i.e.,

$$x(t) = \begin{cases} A \cos(2\pi f_c t + \pi b t^2) & -T/2 \leq t \leq T/2 \\ 0 & |t| > T/2 \end{cases}$$

where the rectangular enveloped is slightly tapered to prevent generation of harmonics. The sonar processor performs real time correlation to compress the sonar return $Y_n(f)$

$$Y_n(f) = X(f)H_{Tx}(f)H_n(f)H_{Tr}(f) + N(f) \quad (102)$$

recorded at the n th hydrophone to a zero phase wavelet with notation convention

$$X(f) = \int_{-\infty}^{\infty} x(t) \exp^{i2\pi ft} dt \quad (103)$$

and where

- $H_{Tx}(f)$ = transmitting voltage response of the projector,
- $H_{Tr}(f)$ = receiving response (transfer function) of the array,
- $H_n(f)$ = frequency response of the seabed including targets,
- $N(f)$ = noise spectrum (ambient, electronic and self-noise).

The output of the correlation filter is the *analytic signal*, given by

$$s_n(t) = \int_0^{\infty} y_n(\tau) f(t + \tau) d\tau \quad (104)$$

where $f(t)$ is the correlation replica and $y_n(t)$ is the output of hydrophone n . The replica is chosen to compensate for the frequency response of the projectors and receivers. In other words, the correlation filter is designed so an echo from an ideal target appears to have a white spectrum. Thus, we obtain the correlation replica as

$$f(t) = \int_{-\infty}^{\infty} \frac{X(f)}{H_{Tx}(f)H_{Tr}(f)} \exp^{-i2\pi ft} df. \quad (105)$$

This result is used to compute the analytic signal in equation (104) and provides the signal-processed input required for the beamforming algorithm that we now describe.

3.3.2 Near-field focusing

Following the matched-field processing described in Sec. 3.3.1 (resulting in the analytical signals $s_n(t)$), we then focus the data by coherent summation of this data over selected physical and synthetic apertures. The objective of this operation is to migrate the scattered echo responses to their physical source location (corresponding to scattering heterogeneities or discontinuities in the propagation medium). Clearly then, for any given focal location \mathbf{x}_f in the medium, data from multiple aspects are processed to obtain the scattering amplitude. After each source transmission (sonar ping), a 3D matrix of image pixels is calculated from the last M pings of the synthetic aperture, starting at ping m_1 and ending at ping m_2 . This yields the image product $A(\mathbf{x}_F, m_2)$ which represents the n -phone scattering response from the 3D focal point \mathbf{x}_F in the seabed:

$$A_M(\mathbf{x}_F, m_2) = \left| \frac{1}{NM} \sum_{n=1}^N \sum_{m=m_2-M+1}^{m_2} s_{n,m}(t_{n,m,F}) c_{n,m}(\mathbf{x}_F, \mathbf{x}_m^T, \mathbf{x}_{m,n}^R) \right|. \quad (106)$$

Here, $s_{n,m}(t_{n,m,F})$ is the value of the analytic signal sampled at the arrival time of sound from point \mathbf{x}_F for the two way path from transducer location \mathbf{x}_m^T (at the m th ping) to the focal point at \mathbf{x}_F to hydrophone n . In addition, $c_{n,m}(x, y, z)$ is a weight coefficient that corrects for two way spherical spreading, boundary losses, refraction, attenuation, etc., computed with respect to a reference model.

The degree of complexity in computing $c_{n,m}$ will depend on the degree of fidelity desired and the information that is available (or can be estimated). The simplest approach is to assume the ray paths connecting the transmitter to the target/scene and the target/scene back to the receiver are straight-lines. This, of course, is consistent only for an environment propagation model with a constant sound speed. This assumption will not generally be satisfied for longer ranges or shallow grazing angles. Vertical stratification in the sound-speed profile will lead to ray refraction. In addition, ray interactions with the surface and bottom result in multipath arrivals which can lead to ghosting effects. Phase-coherent processing schemes must, in general, account for platform uncertainties or perturbation of environmental propagation parameters that alters the acoustic phase of the sonar echo via an appropriate processing model. Failure to do so leads to image degradation (e.g., reduced SNR in image formation, target mis-location, reduced cross-range resolution, larger side-lobes levels - which is an especially large problem for high-clutter conditions) and blurring. All of these factors will compromise image interpretation. The experiment protocol for BOSS partially mitigates some of these issues, but they will become increasingly significant with increasing range.

3.3.3 Approximate near-field focusing

The focusing construction in eq. (106) assumes prior knowledge of the seafloor properties as embodied in the weighting coefficient $c_{m,n}$. Of course, in practice, the reference model used to compute $c_{m,n}$ will be a simplistic approximation to the real medium since typically only limited information will be available. Perhaps the simplest approximation to $c_{m,n}$ is to account for spherical spreading and to ignore all other effects. Assuming then a homogeneous medium, spherical spreading occurs along the straight-line propagation paths connecting the active source to the desired scattering point, and the scattering point back to the receiver array element. In this case, $c_{n,m}$

is written as

$$c_{n,m}(\mathbf{x}_F, \mathbf{x}_m^T, \mathbf{x}_{m,n}^R) = R(\mathbf{x}_m^T, \mathbf{x}_F) \times R(\mathbf{x}_F, \mathbf{x}_{m,n}^R) \quad (107)$$

in which

$$R(\mathbf{x}_m^T, \mathbf{x}_F) \equiv \text{path length from the transmitter to } \mathbf{x}_F, \quad (108)$$

$$R(\mathbf{x}_F, \mathbf{x}_{m,n}^R) \equiv \text{path length from } \mathbf{x}_F \text{ to the } n\text{th receiver} \quad (109)$$

for ping m . From equation (106), the scattering amplitude becomes

$$A_M(\mathbf{x}_F, m_2) \cong \left| \frac{1}{NM} \sum_{n=1}^N \sum_{m=m_2-M+1}^{m_2} s_{n,m}(t_{n,m,F}) R(\mathbf{x}_m^T, \mathbf{x}_F) R(\mathbf{x}_F, \mathbf{x}_{m,n}^R) \right|. \quad (110)$$

The scattering amplitudes computed in this way can then be further processed to yield various image constructs as we now discuss.

3.3.4 Image construction

Image formation from the scattering amplitude matrices $A_M(\mathbf{x}_F, m_2)$ may be approached in a number of ways. In our first construct we compute a multi-aspect image $B(\mathbf{x}_F)$ by saving for each desired focal point \mathbf{x}_F the maximum scattering amplitude value $A(\mathbf{x}_F, m_i)$ drawn from a sequence of Q overlapping scattering amplitude matrices such that $m_i \in \{m_2 - Q + 1, \dots, m_2\}$. In other words, the processing from the focusing operation will have generated scattering amplitude intensities for a sequence of 3D volume samples, and each of these volume samples contain partial overlaps of co-located sampling points (depending on the processing parameters chosen and the kinematic trajectory of the sounder). With this choice, the image expression may be written

$$B(\mathbf{x}_F) \equiv \max\{A_M(\mathbf{x}_F, m_2 - Q + 1), \dots, A_M(\mathbf{x}_F, m_2)\} \quad (111)$$

where the maximum is chosen from one of Q scattering amplitude matrices, and \mathbf{x}_F is geolocation of the imaging point. Implicitly, the single final image over the collection of imaging points \mathbf{x}_F will be composed from inputs of possibly many matrices, and the contributing matrix may vary from point to point. However, it is reasonable to expect that there will be small-wavelength spatial correlations (over \mathbf{x}_F with respect to the identity of the source matrix. This operation yields a 3D matrix which may be visualized in various ways (e.g., horizontal or vertical slices, isocontours, etc.). In addition, this image product may be post-processed to derive various features (e.g., binary partitions of target/no-target volumes, histograms of scattering intensities within a target scattering volume, etc.). Alternatively to equation (111), a *tomographic image* may be obtained as

$$T(\mathbf{x}_F) = \sqrt{\frac{1}{M} \sum_{m=1}^M A^2(\mathbf{x}_F, m)} \quad (112)$$

where the sum is taken for a fixed collocated point \mathbf{x}_F . Finally, in one more construct, orthogonal views of the seabed may be computed using *Maximum Intensity Projection* (MIP) mapping may

be defined as follows:

$$P_{xy}(x, y) = \max_{z_i \in \mathbf{x}_F} \{B(x, y, z_i)\} \quad \text{for fixed } (x, y), \quad (113)$$

$$P_{xz}(x, z) = \max_{y_i \in \mathbf{x}_F} \{B(x, y_i, z)\} \quad \text{for fixed } (x, z), \quad (114)$$

$$P_{yz}(y, z) = \max_{x_i \in \mathbf{x}_F} \{B(x_i, y, z)\} \quad \text{for fixed } (y, z). \quad (115)$$

3.4 Adjoint tomography

We have investigated adjoint tomography as a means to estimate poroelastic medium parameters from BOSS data. The objective is to develop 3D image products for these images, and from these derive new feature classes similar to those used in the traditional beamformed images but with the new data. In principle, this could represent a significant augmentation of the standard intensity derived feature classes.

To enable this program, it is necessary to derive the so-called *poroelastic finite-frequency sensitivity kernels* based upon adjoint methods. This work is reported in [29] and [28], where it is shown that the adjoint equations are similar to the forward Biot equations. The only difference involves the choice of the adjoint source. The key point about these kernels is that they quantify the sensitivity of seismic observables to structural parameters, and therefore form the basis of tomographic inversions. Thus, we have defined a series of new sensitivity kernels attached to diverse key poroelastic parameters, which can be used to infer potentially valuable information for discrimination.

Defining the data misfit between an observed and predicted sonar time series as $\delta\chi$, the sensitivity kernels $K_{p_i}(\mathbf{x})$ [29, 28] for parameter p_i are defined in the following sense:

$$\delta\chi = \int \sum_{i=1}^N \delta \ln p_i(\mathbf{x}) K_{p_i}(\mathbf{x}) d^3\mathbf{x} \quad (116)$$

where the poroelastic parameters p_i are enumerated in Table 1.

For example, the kernel for fluid density is given by

$$K_{\rho}(\mathbf{x}) = \int_0^T \rho(\mathbf{x}) \mathbf{u}_s^*(\mathbf{x}, T - t) \cdot \partial_t^2 \mathbf{u}_s(\mathbf{x}, t). \quad (117)$$

Kernels for the remaining terms in [29] and [4], and where we also provide a comprehensive presentation of the adjoint theory, and the basis of the adjoint tomographic process. In the remainder of this subsection we simply consider a 2D propagation scenario in a poroelastic medium, and we provide visualization of the various adjoint kernels (such as eq. 117) that enter into into the data misfit definition in 117.

We consider a simple 2-D poroelastic SEM simulation to generate banana-doughnut kernels (a common name for the adjoint kernels)⁴ properties associated with cross-correlation traveltimes measurements. The 2-D model has an aspect ratio of 2:5 in depth versus width, and is homogeneous. The source was simulated using a Ricker wavelet source time function. The source-receiver geometry and a schematic illustration of the P-SV body-wave arrivals are displayed in Fig. 8 together with the velocity recorded at the receiver. Since we are modeling poroelastic media there are multiple types of wave modes and wave conversions (at interfaces) that are supported. A subset of these and the notation for them are enumerated in Table 2

We construct an adjoint source by windowing the PI-S + S-PI arrival recorded at the receiver. The interaction of the forward field and the adjoint field gives rise to a series of kernels associated

⁴Banana-doughnut kernels are so-named in the geophysical literature because of their resemblance to bananas. They indicate the “zone of sensitivity” for a ray i.e., they define the spatial region that define the behavior and properties of a ray’s characteristics. In a medium with increasing velocity with depth, ray bending occurs, and the ray launched downward eventually refracts, bends upwards, and returns to a receiver. This ray and the spatial kernel defining it are what appear to have a banana-like shape for many cases.

Parameter	Description
p_1	solid density
p_2	fluid density
p_3	porosity
p_4	tortuosity
p_5	solid bulk modulus
p_6	fluid bulk modulus
p_7	frame bulk modulus
p_8	fluid viscosity
p_9	frame shear modulus
p_{10}	fast P-wave speed
p_{11}	slow P-wave speed
p_{12}	S-wave speed

Table 1: Poroelastic parameters considered in adjoint sensitivity formulation.

S	transverse or shear wave mode
P	longitudinal or P-wave mode
PI	fast P-wave mode
PII	slow P-wave mode
PI-PII	Phase due to conversion of fast P-wave mode to slow P-wave mode
S-PII	Phase due to conversion of shear mode to slow P-wave mode
PI-S	Phase due to conversion of fast P-wave mode to shear mode
PII-PII	Phase due to reflection of slow P-wave mode from an interface
S-S	Phase due to reflection of a shear mode from an interface
PI-PI	Phase due to reflection of fast P-wave mode from an interface
PII-PI	Phase due to conversion of slow P-wave mode to fast P-wave mode
PII-S	Phase due to conversion of slow P-wave mode to a shear mode
S-PI	Phase due to conversion of a shear mode to a fast P-wave mode

Table 2: Mode wave supported in simulation examples including fundamental waves and their conversions.

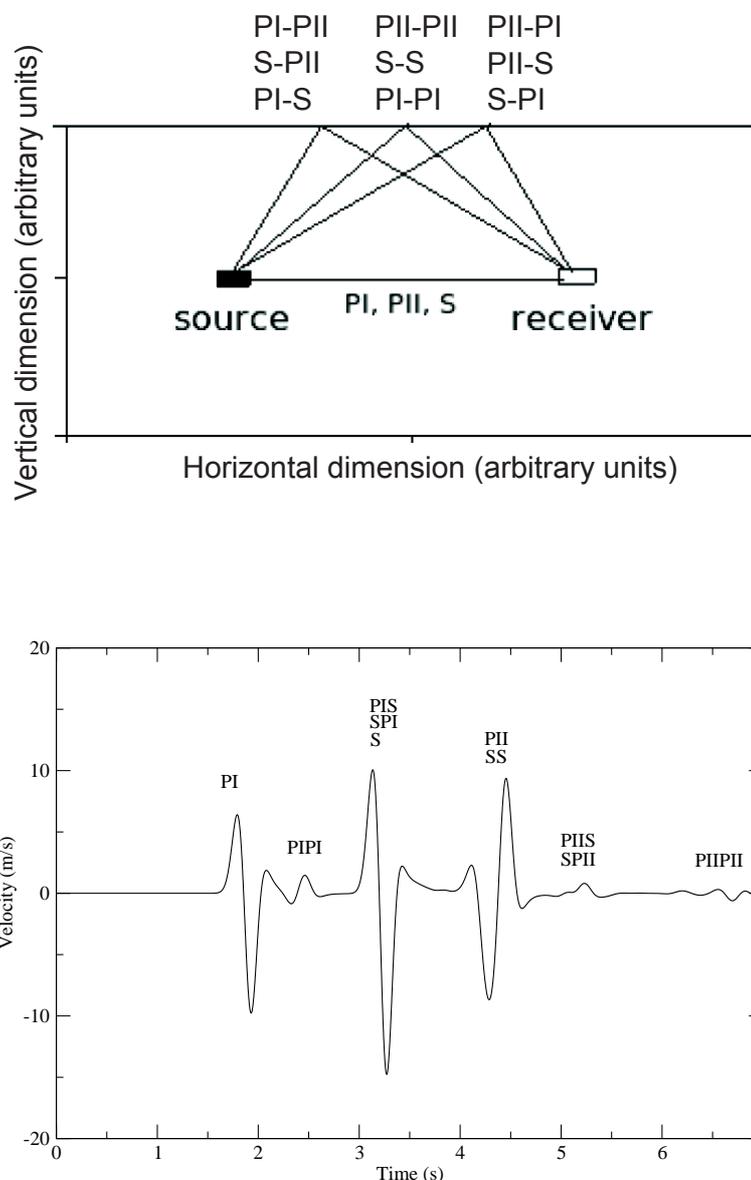


Figure 8: **Top panel:** Schema of the 2-D model dimensions and source-receiver geometry. The top surface is a free surface whereas the three remaining edges are absorbing boundaries. We also draw possible ray paths for the P-SV wavefield, where PI refers to the fast P wave, PII to the slow P wave and S to the S wave. See Table 2 for nomenclature summary. **Bottom panel:** Velocity seismogram (x-component).

with the PI-S + S-PI arrival. An example subset of kernels are displayed in Fig. 9. Similarly, we construct an adjoint source by windowing the SS arrival recorded at the receiver. The interaction of the forward field and the adjoint field gives rise to a series of kernels associated with the SS arrival. These and other kernels are derived and displayed in [29] and [4].

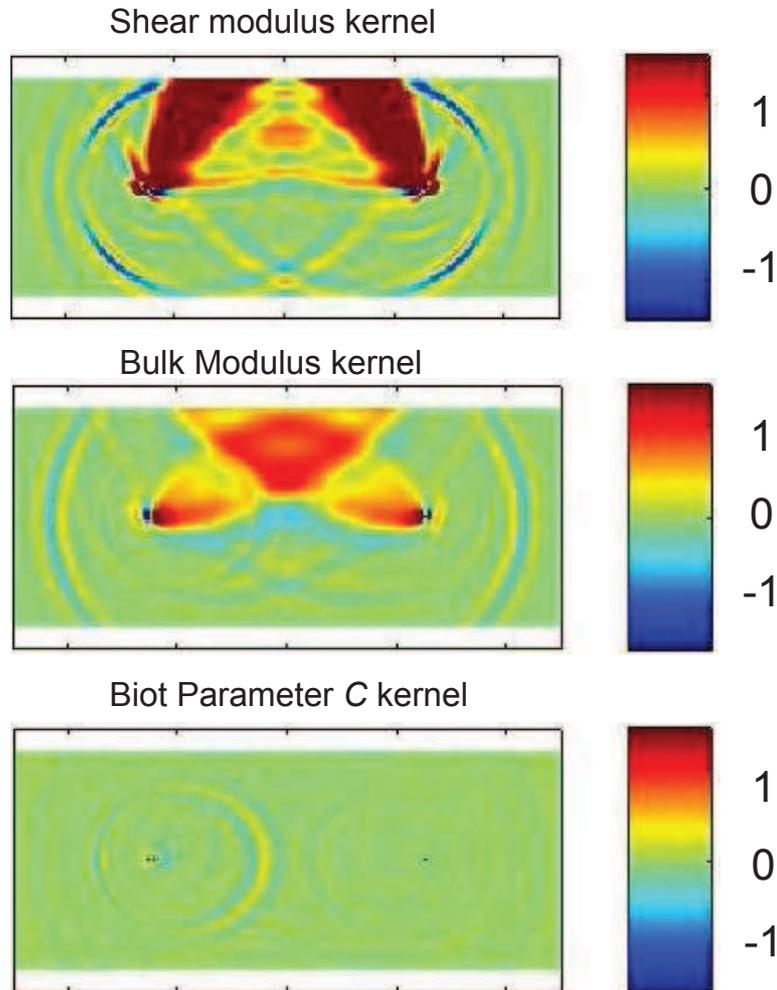


Figure 9: Plot of the density-normalized $P-SV_{PIS+SPI}$ banana-doughnut kernels for selected poroelastic quantities. The vertical and horizontal axes have the same arbitrary length units of Fig. 8, and refer the same length spans as depicted in that figure. **(top)** - shear modulus kernel, **(middle)** - Bulk modulus, **(bottom)** - Biot coefficient C [29, 4]. See [29] and [4] for details of kernel construction and for explicit kernels for each of the quantities in Table 1.

3.5 Bayesian approach for BOSS classification analysis (Relevance Vector Machine)

A key objective of this effort has been to use BOSS data collections and high fidelity simulations of BOSS data to develop an effective classification and detection algorithm for buried underwater munitions. The classifier can have several designs depending on the mission objective, and on the available data for training the classifier. The classifier should admit feature vector inputs that can readily be extracted from raw waveforms or beamformed data, and the outputs (i.e., class labels, probability of correct classification) should correspond, to the extent possible, to the needs of the end-user. An example output is the binary target/clutter decision. Of course, the training of the algorithm will vary depending on the prior selection of what constitutes a clutter item and a target item. If greater specificity is required e.g., bomb, artillery shell or the corresponding sub-varieties of these target items, then the output dimensionality of the classifier algorithm increases accordingly, and there is a greater burden on the part of the algorithm designer to collect the correct descriptors. These comments assume that *supervised learning* is applied.

The decisive factors for algorithm performance are the choice and number of input features, the choice and number of output classes, the quality, diversity and size of the available training exemplars, and the algorithm design itself. An additional key factor is the ‘distance’, in multi-dimensional space, between features corresponding to the output classes. This reflects the intrinsic sensitivity of the collected data to the varying target and clutter items. The combination of these factors define the key performance metrics of the classification system (characterized, for example, by a receiver operator characteristic (ROC)). A key requirement is good *generalization* of the classifier algorithm i.e., the algorithm performance as applied to data inputs outside of the training set. These and other issues are discussed in a vast literature, and an excellent exposition can be found in the text by Bishop [3].

We use in the following the Relevance Vector Machine (RVM) [35, 3] to achieve near-optimal classification performance with respect to the available inputs. We report in Sec. 4.6 on the numerical results obtained from analysis of the AUV FEST 2008 data, and in the remainder of this section we describe the salient aspects of the RVM that drove our selection of this technique for the classification analysis.

In supervised learning we are given a set of examples of input vectors (feature vectors) $\mathcal{A} = \{x_n\}_{n=1}^N$ derived from the data, along with corresponding values $\mathcal{B} = \{t_n\}_{n=1}^N$ that could be numerical values of an unknown function (for regression), or class labels (e.g., target or clutter, target type, etc.) for classification. From this ‘training set we wish to learn a model of the dependency of the values or labels on the feature inputs. The objective is to develop a model with which to make accurate predictions of the value t for a previously unseen value x . In real-world data, the presence of noise (in regression) and class overlap (in classification) implies that the principal modelling challenge is to avoid overfitting of the training set.

Stated probabilistically, we seek to determine a parameterized model for the conditional probability

$$P(B|A) = f(A; \mathbf{w}) \quad (118)$$

where f is a parameterized functional form, A is an instance of the input feature vector, B is an instance of the output label or regression prediction, and \mathbf{w} are the parameters to be determined. More generally, the function form f is not necessarily fixed. For example, one possible form of f is

a weighted, linear summation over basis functions, each of which may have nonlinear dependency on the input feature vector. In the case of RVM, there may be a basis function associated with each input training vector, but the initial assumed functional form may be highly sub-optimal. An informed training process would decide which basis functions to retain, and which to delete. This is an estimation problem. We may take the variable \mathbf{w} to denote a vector of all the ‘adjustable’ parameters in the model, and its dimension will depend on the optimal form of f . In the most general sense, the optimization problem is over the space of permissible functions f , and the space of parameters \mathbf{w} associated with each f . For the case of the linear sum representation, non-zero values of w_i indicate which basis functions are retained.

Following a conventional route, given a set \mathcal{D} of N examples of our variables, $\mathcal{D} = \{A_n; B_n\}_{n=1}^N$, an approach is to maximize an objective function that characterizes the accuracy of the model for \mathcal{D} with respect to (f, \mathbf{w}) . The output for a given A , is a prediction of B by evaluating $f(A; \mathbf{w})$ with parameters \mathbf{w} . The form f and the values \mathbf{w} are discovered during the ‘training’ phase, and the predictions using previously unseen samples are made during the ‘test’ phase.

Of course, if our model f is made too complex i.e., there are too many basis terms in f and the corresponding dimension of \mathbf{w} is too large, then the model is at risk of being over-specialized with respect to the observed data D used for training. This yields the classic problem of ‘fitting the noise’ rather than discovering the true underlying distribution $P(B|A)$. We have adopted the RVM for our approach since it presents a principled Bayesian inference paradigm in which the posterior distribution for \mathbf{w} is economically determined. A given functional form for f is initially specified, and which is likely too complex for good generalization properties. However, by the Bayesian approach, the distribution for \mathbf{w} is discovered, and small values for individual components of \mathbf{w} reduce the complexity of f by suppressing the influence of the associated terms i.e., effectively zeroing them out. This process begins by specifying a ‘prior’ distribution $p(\mathbf{w})$ on \mathbf{w} before we observe the data. The Bayesian approach exemplifies ‘Occam’s Razor’ which is commonly used in geophysical inverse problems. It automatically ‘integrates out’ all irrelevant variables since there is an automatic preference for simple models that sufficiently explain the data. This feature avoids the need to develop cross-validation procedures.

Since the theory of RVM is described in great detail in [35] and [3], we do not reproduce that exposition here. However, we do provide the minimal background needed to understand the basic numerical outputs of the BOSS classification results and diagnostics presented in the following sections.

The model we seek defines a conditional distribution for an output variable (which could be a continuous variable in the case of regression, or a discrete label in the case of classification). We denote the output variable as t and the input vector as \mathbf{x} , and which takes the form

$$p(t|\mathbf{x}, \mathbf{w}, \beta) = \mathcal{N}(t|y(\mathbf{x}), \beta^{-1}) \tag{119}$$

where β is the inverse noise variance, and the mean is given as the linear model

$$y(\mathbf{x}; \mathbf{w}) = \sum_{i=1}^M w_i \psi_i(\mathbf{x}) = \mathbf{w}^T \Phi(\mathbf{x}) \tag{120}$$

with fixed, nonlinear basis functions ψ_i , which will typically include a constant term to correct for

bias. A specialization of the form in eq. 120 is given by

$$y^j(\mathbf{x}, \mathbf{w}) = \sum_{n=1}^N w_n K^{(j)}(\mathbf{x}, \mathbf{x}_n) + w_0 \quad (121)$$

where K is a kernel function (e.g., Gaussian, Cauchy, polynomial, etc.) effectively defining one basis function for each example in the training set, and w_0 is the bias term corresponding to a constant basis function. In the case of our application to BOSS, an instance of an input training vector \mathbf{x}_n for the algorithm would correspond to valuations of a subset of the 50 possible features listed in Table 7.

Suppose we are given a set of N observations of the input vector \mathbf{x} , which we denote collectively by the data matrix \mathbf{X} whose n^{th} row is \mathbf{x}_n^T with $1 \leq n \leq N$. The corresponding target values are given by $\mathbf{t} = (t_1, \dots, t_N)^T$. The likelihood function is given by

$$p(\mathbf{t}|\mathbf{X}, \mathbf{w}, \beta) = \prod_{n=1}^N p(t_n|\mathbf{x}_n, \mathbf{w}, \beta^{-1}) \quad (122)$$

The key feature of the RVM is that it utilizes a prior distribution over the parameter vector \mathbf{w} , and a separate hyperparameter α_i for each of the weight parameter w_i is assumed instead of a single shared hyperparameter. We typically consider a zero-mean Gaussian prior, and thus, the weight prior takes the form

$$p(\mathbf{w}|\boldsymbol{\alpha}) = \prod_{i=1}^M \mathcal{N}(w_i|0, \alpha_i^{-1}) \quad (123)$$

where α_i represents the precision of the corresponding parameter w_i , and $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_M)^T$. In the RVM, when the evidence is maximized with respect to these hyperparameters, a significant number of them go to infinity, and the corresponding weight parameters have posterior distributions with means and variances that are concentrated at zero. The basis functions associated with these parameters therefore play no role in the predictions made by the model and are effectively pruned out, resulting in a sparse model. In the case of models of the form in eq. 120, the inputs \mathbf{x}_n corresponding to the remaining non-zero weights are called *relevance vectors*.

4 Results and Discussion

In this section we discuss the specific experiments that were performed along with results of the data analysis. The principal thrusts areas that we discuss include: (i) numerical simulations using the SEM (Sec. 4.1), (ii) numerical simulations of BOSS data using both the SEM and T-matrix methods (Sec. 4.2), feature analysis (Secs. 4.4 and 4.5) of the BOSS data summarized in Sec. 4.3, and the detailed classification results (Sec. 4.6). The classification analysis focused on the AUV FEST 2008 data since this data set was the largest available, and since it also contained the most complete list of separately labeled target and clutter items.

4.1 Numerical simulations using the SEM method

Extensive numerical simulations using the SEM developed under this program have been published [27, 29, 4] and will not be repeated here. However, it is worth noting that the results reveal the rich phenomenology of the sonar response induced by poroelastic properties of the sediment medium. Also, the numerical results were validated by comparison to analytical results for benchmark problems.

4.2 Numerical simulation of BOSS data

We have pursued numerical simulations of BOSS waveforms and BOSS beamformed image products using two distinct methods, namely, the Spectral Element Method (SEM) and the T-matrix method (sec. 4.2.1 below). The theory for the latter method was presented in Sec. 3.2.1 and the poroelastic SEM theory was published as [27].

4.2.1 Numerical simulations using the T-matrix method

Simulations were performed for a source-receiver geometry consistent with that of the BOSS platform actually used in several of the AUV FEST collections. The geometry used is shown in Fig. 10.

As described in Sec. 3.2.1, the T-matrix method may be used to synthesize the scattered field due to a single scatterer, interactions between multiple scatterers and interactions between the single or multiple scatterers and the boundary interfaces of a layered medium. The physics of the scattering process are encoded by the T-matrix coefficients $T_{nn'}$. The latter are fundamental inputs to the approach, and must be derived by other means e.g., analytical solution, numerical solution or by inversion from real data (as described, for example, in Sec. 3.2.11). We have implemented the T-matrix coefficients for spherical elastic shells, solid elastic spheres and acoustic (fluid) spheres using the analytical results described in [10]. The T-matrix coefficients define scattering amplitudes and their frequency-dependence implicitly define the resonance spectrum of a particular target. The spectrum can provide a ‘fingerprint’ for target types. Further, for aspherical targets, modes will be excited differentially depending on incident azimuth and elevation, and this characteristic variation can also serve as an identifying characteristic. Figure 11 is an example resonance spectrum for a thin-shell spherical target. In this the sphere radius is 50 cm, and the shell thickness is 10 cm. The modal vibration shapes are described by spherical harmonic basis functions $Y_l^m(\theta, \phi)$. These vary with degree l and are independent of m due to the spherical symmetry.

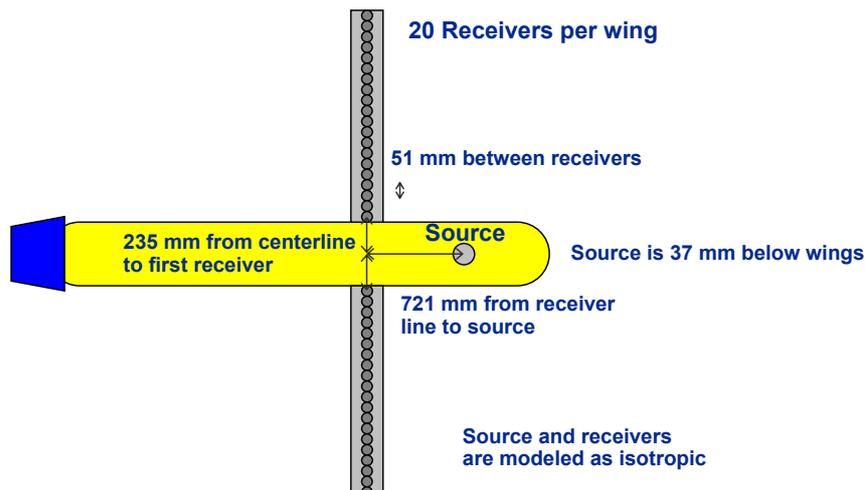


Figure 10: This figure shows the source receiver geometry used in the simulation study in which there are 20 receivers on each wing of the indicated AUV with 51 mm spacing, a spherical source displaced 721 mm from the centerline of the receivers in the long dimension of the AUV, and for which the first receiver on each wing is 235 mm from the long-axis centerline of the AUV. This configuration corresponds to the actual experiment geometry used by BOSS in the more recent AUV FEST data collects.

Shown are the coefficients as a function of frequency for harmonic degrees $0 \leq l \leq 4$ where $l = 0$ is the fundamental mode. Since this is a purely spherical ‘breathing’ mode it is also known as the *s mode*.

The time series for a thin-shell target sphere displaced 5 m along the center line from the source position of the BOSS AUV geometry is shown in Fig. 12, and illustrates the ringing phenomenon due to the resonant nature of the spherical shell target.

In contrast to the calculation for the target in Fig. 11, Fig. 13 shows an example of a thick-shell resonant spectrum. In this case the target is also smaller in dimension so that the lowest frequency modes have higher frequencies than in the thin-shell example (since the two-way travel time or round-trip circumferential travel time is faster).

To illustrate the character of time series for the case of multiple scatterers in the medium, we show in Fig. 14 the time series resulting due to 13 randomly distributed solid and thin-shell spheres. In addition, all multiple interactions between the spheres are included. In this case direct and interfering arrivals are observed, and the various time delays observed in the wavetrain correspond to the relative source/receiver and target geometries. A homogenous medium was assumed and the sound speed was taken as 1500 m/s.

In Fig. 15 we apply the beamform processing described in Sec. 3.3 to a set of 40 synthetic time series corresponding to the 40 receivers shown in Fig. 10. Each of these series were similar to the synthetic phase history shown in Fig. 12. In this calculation only a single source position was used⁵. In the image we observed that the intensity is peaked at the true target location. However, we also observe a hyperbolic ‘smile’, and this shows the ambiguity in the target location as a consequence of the restricted experiment aperture. By using additional source positions, we would

⁵In the real data examples in this report sources were used typically over a 10 m length with a 10 cm ping interval

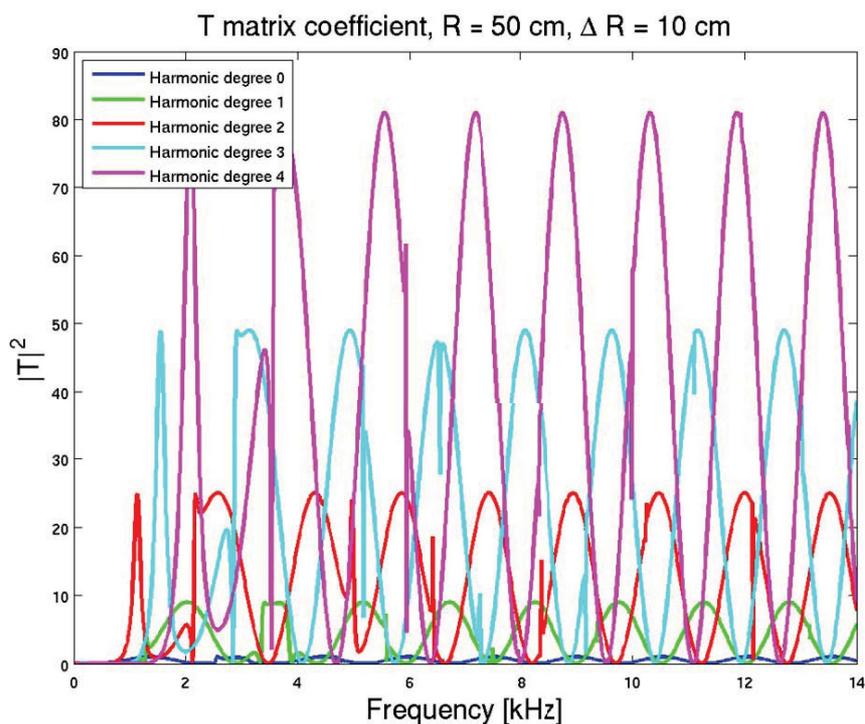


Figure 11: Resonant spectrum for spherical harmonic degrees $0 \leq l \leq 4$ for a thin elastic spherical shell with radius $R = 50$ cm, and thin shell thickness given by $\Delta R = 10$ cm. The resonance spacing depend on the sphere size, sphere elastic properties, and $\Delta R/R$.

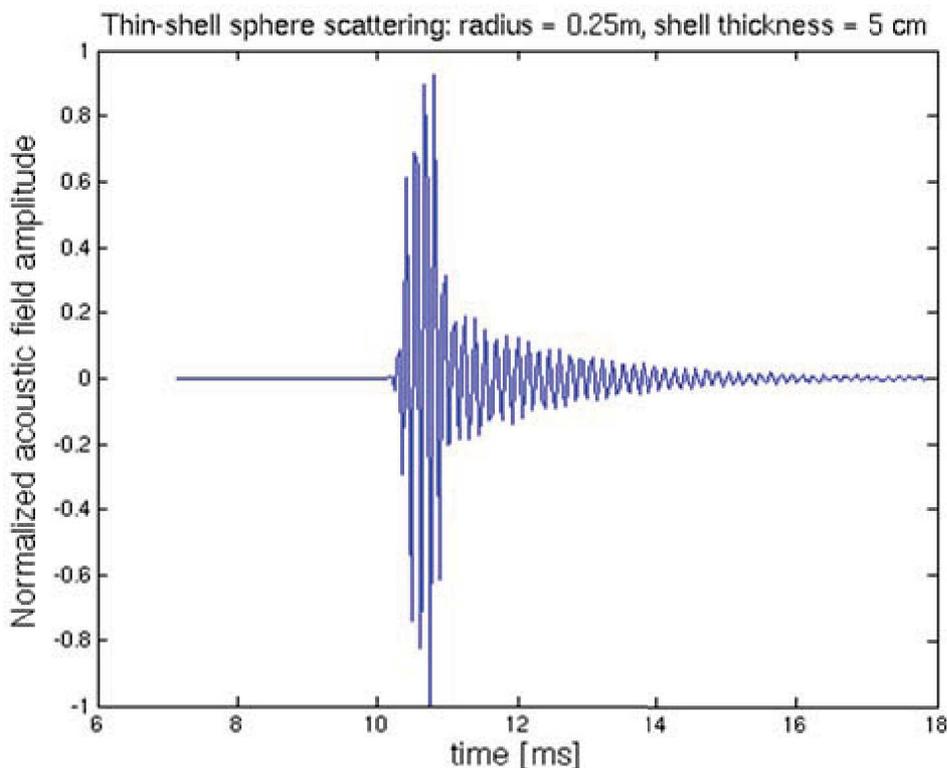


Figure 12: Time series for single source position and receiver shown Thin-shell target with radius $r = 25$ cm, and shell thickness $\Delta r = 5$ cm. The ringing observed after the first main arrival corresponds to overtones from the resonant mode spectrum, similar to those seen in Fig. 11. A Fourier analysis of this time series would reveal the spectrum, and provide features for use in classification.

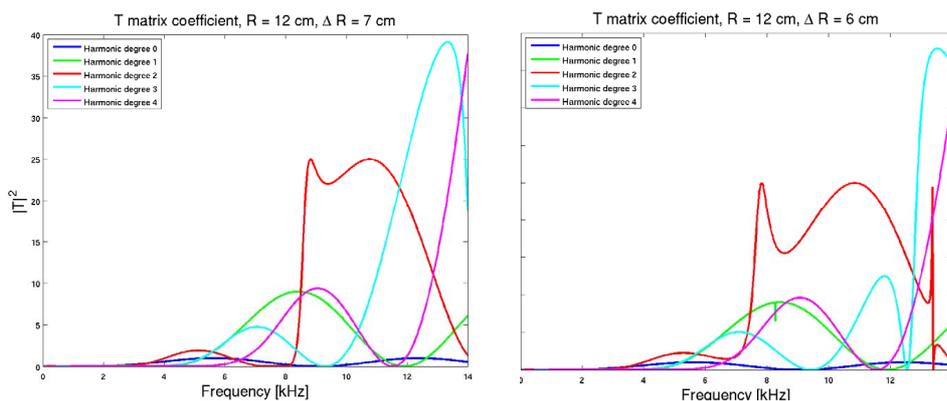


Figure 13: Thicker shell (left) has radius $R = 12$ cm, and shell thickness $\Delta R = 7$ cm, while the thinner shell (right) has values $R = 12$ cm, and shell thickness $\Delta R = 6$ cm. Detailed spectra provide features for discrimination

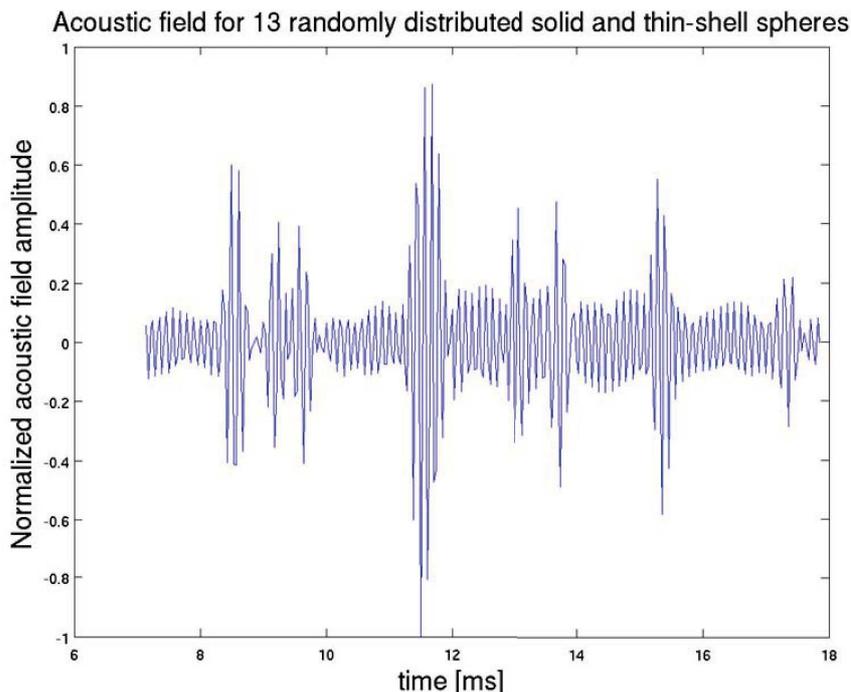


Figure 14: Synthesized Time Series for 13 Randomly Distributed Solid and Thin-shell Spheres. All possible multiple interactions are accounted for in this simulation.

discover that returns from the true target location constructively interfere while returns from ‘false’ positions do not. Hence, the collective effect is to suppress the ‘smile’ and to enhance the true point scattering location.

4.2.2 Numerical simulations using the spectral element method

We show in Fig. 16 an application of the SEM modeling capability to simulate the scattering response for a buried target for a data experiment that mimics the collection style of the BOSS system. We designed three models to evaluate the signature of a purely elastic buried object (yellow rectangle in the figure) in three types of environments:

- (1) Model 1: acoustic layer on top of a poroelastic medium with a porosity gradient and no viscous damping.
- (2) Model 2: acoustic layer on top of a poroelastic medium with a porosity gradient and viscous damping (with fluid viscosity $\eta_f = 10^{-4}$ Pas, and isotropic permeability $k = 10^{-11}$ m²).
- (3) Model 3: acoustic layer on top of an elastic medium.

Figure 16 shows differential acoustic/seismic waveforms due to the insonification for each of the bottom environments. The differential calculation is taken between waveforms synthesized with and without the buried target, and they therefore provide a direct signature of the target. The results show significant differences in the response depending on the environment. In further downstream

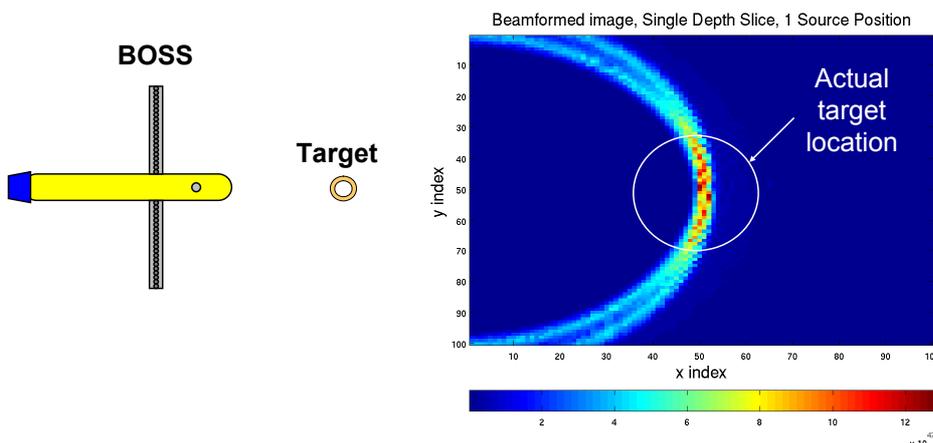
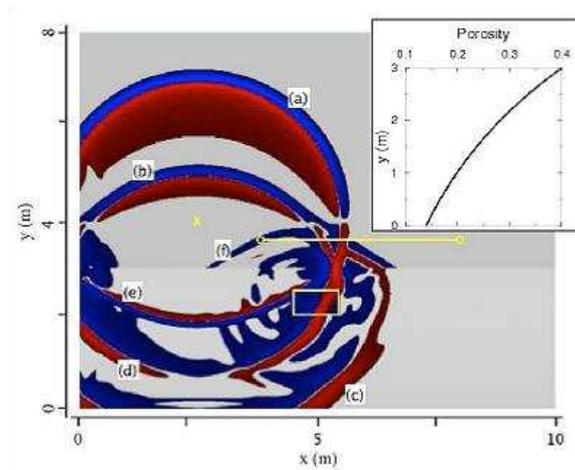


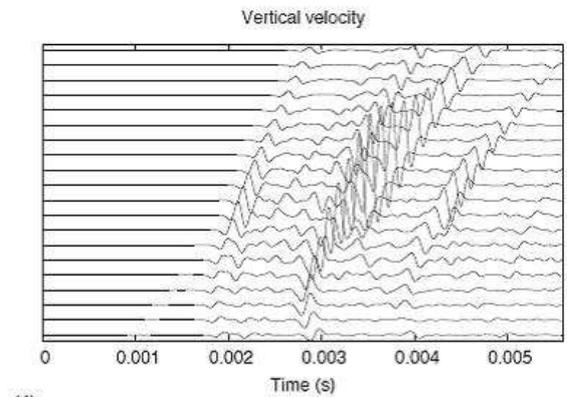
Figure 15: This figure shows the beamformed image for scattering due to a single, thin-shell spherical target, a single source position, and for all 40 receivers shown in Fig. 10. The input time series consists of a center frequency of 8000 Hz, and the source signal has 12,000 Hz bandwidth. The result shows that the scattering intensity is peaked in the true location of the target. However, in constructing this beamformed image, only one source position was used. This fact in combination with the finite bandwidth of the source location is the origin of the smile-like spread of the intensity. The use of additional source positions (as, in fact, is commonly done with BOSS deployments) would yield sharper focus since all of these contributions are coherently added, and volumetric regions containing no scatterers will tend to have destructive interference from the wavefields summed over the sources and the receivers. If the spatial coverage is skewed or non-symmetric, there will tend to be some geometric distortion. However, even in the case of perfect and complete spatial coverage, the scattering location will tend to be a “ball” rather than a point due to the finite bandwidth. Thus, in practice, one might use 100 source positions in the image construction, and each wavefield recording is acquired as the BOSS platform moves at speed 2 m/s and with a source pulse repetition frequency (PRF) of 20 Hz. In this way, beamforming over multiple source positions mitigates against the scattering location ambiguity.

processing, waveform feature descriptors of the response may be derived as a function of frequency (e.g., acoustic color), target type and aspect angle, and others. Of course, these waveforms may be processed into 2D and 3D image products from which additional features may also be extracted. The differential acoustic/seismic traces for Models 1 (top panel) and 2 (middle panel) illustrate the impact of viscous damping on the slow compressional waves, which are clearly suppressed in Model 2. The signature of the object in differential seismograms for elastic Model 3 (bottom panel) is noticeably different from that in the poroelastic models.

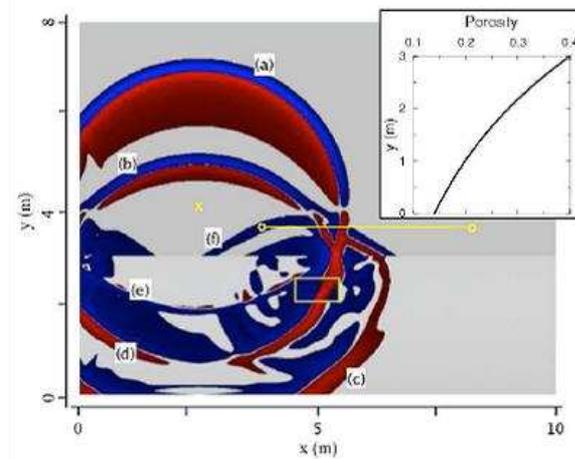
The acoustic/elastic target properties were taken as follows: density – 4000 kg/m³, bulk modulus – 63.6 GPa, shear modulus – 8 GPa, P-wave speed – 4309 m/s, and S wave speed – 1414 m/s. We use a Ricker source time function with a dominant frequency of 5 kHz. The dimensions of the simulation are 10 m × 8 m, the source (cross) is located in the acoustic domain at $\mathbf{x}_s = (2.5, 4.0)$, and the 20 receivers (circles) are placed near the bottom of this domain in a line array with receiver elements evenly located between $x_{r1} = (4.0, 3.5)$ and $x_{r20} = (8.0, 3.5)$. All the edges in the SEM simulation are absorbing boundaries. We can observe the direct P (a) and the reflected P (b) waves in the acoustic domain, the transmitted fast P (c), the P-to-S converted (d), and the fast P-to-slow P converted (e) waves in the poroelastic domain, plus waves reflected by the elastic object (f).



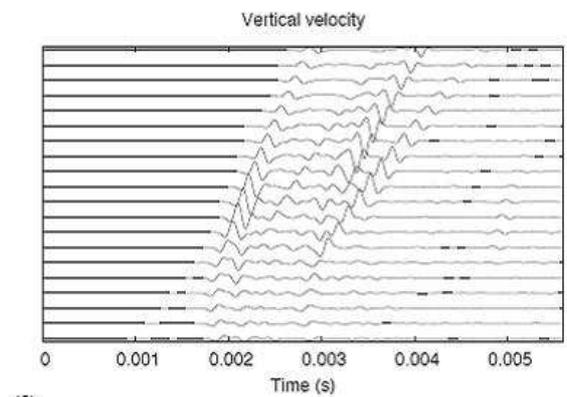
(1)



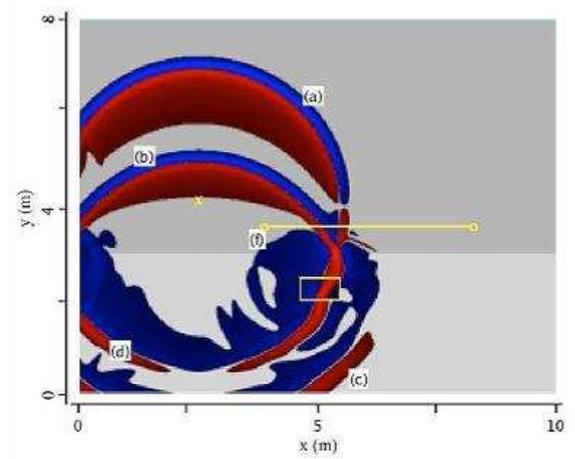
(4)



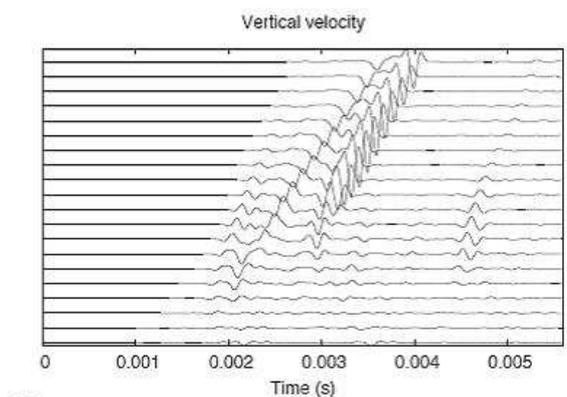
(2)



(5)



(3)



(6)

Figure 16: Simulation of BOSS-like experiment data for a buried metal target (yellow rectangle) in three different sediment layers. The source is indicated by the yellow cross, and the receiver array is indicated by the yellow circles.

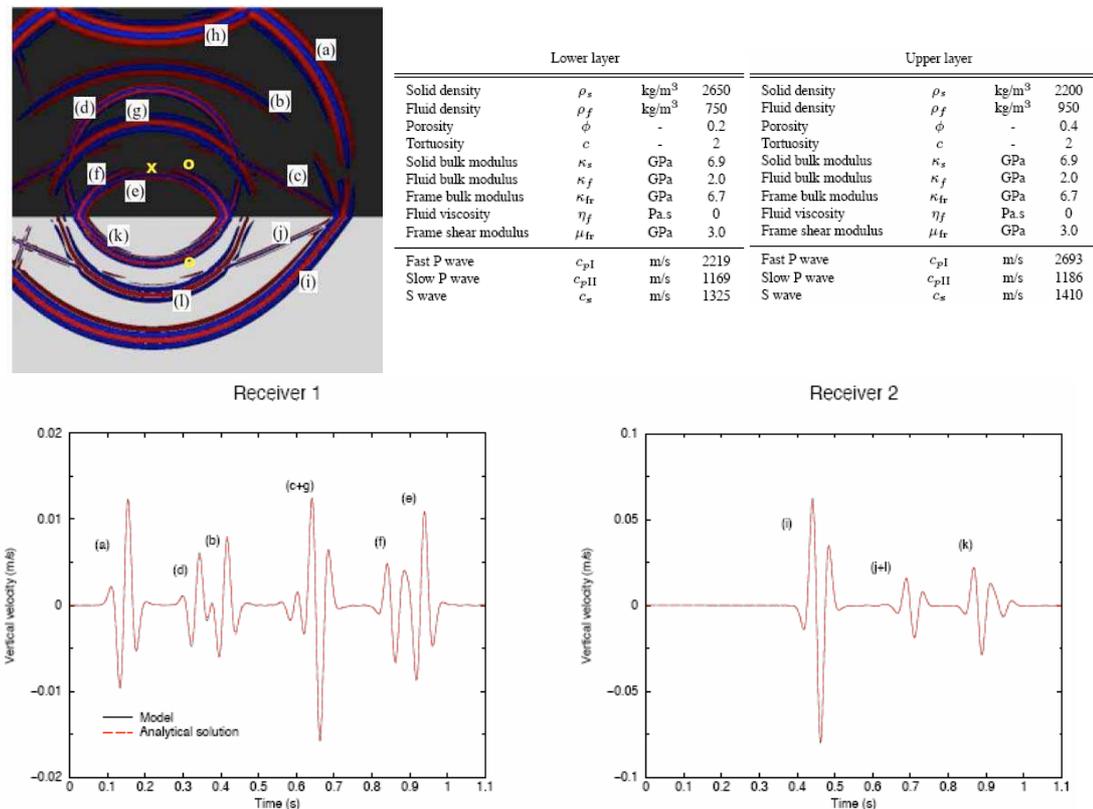


Figure 17: This figure is one of numerous benchmarks we have devised to validate the spectral element method theoretical formulation and numerical implementation. Validation is achieved by comparison to an available analytical solution. This simulation of wave propagation is for a model consisting of two homogeneous poroelastic layers with continuous bulk and shear moduli and discontinuous porosity, as tabulated in the upper right of the figure. The source (cross) is located at the and the two receivers are indicated by the circles. The top is a free surface and the remaining three edges are absorbing boundaries. **Upper left:** Snapshot of the vertical-component displacement at $t = 0.9$ s. The direct fast P (a), the reflected fast P (b), the reflected fast P-to-S and the fast P-to-slow P converted (c) waves (which overlap because they have similar wave speeds) can be observed in the upper layer, together with the direct slow P (d), the reflected slow P (e), the reflected slow P-to-S converted (f), and the reflected slow P-to-fast P converted (g) waves. We also observe the reflected fast P wave due to the free surface (h). In the lower layer, the transmitted fast P (i), fast P-to-S and fast P-to-slow P converted (j) waves (which again overlap because they have similar wave speeds) can clearly be identified, together with the transmitted slow P and slow P-to-S converted (k) and slow P-to-fast P converted (l) waves. There are some weak spurious reflections from the absorbing boundary at $x = 0$. **Lower left and lower right:** Vertical-component velocity seismograms at receivers 1 and 2 (Spectral Element Simulation: solid black line, analytical solution: dashed red line). We use domain composition to accommodate the first-order discontinuity in porosity.

An example of BOSS phase histories from AUV FEST 2007 that we are currently modeling is shown in Fig. 18. Modeling of such data requires knowledge of the source strength, receiver response to a unit source as a function of frequency (i.e., transfer function), pulse repetition frequency, precise source-receiver geometry, UUV kinematics, UUV pose with respect to the environment, and knowledge of the environment itself including sediment background model as well as the target/clutter field. These data have been gathered to support the simulations. The beamformed

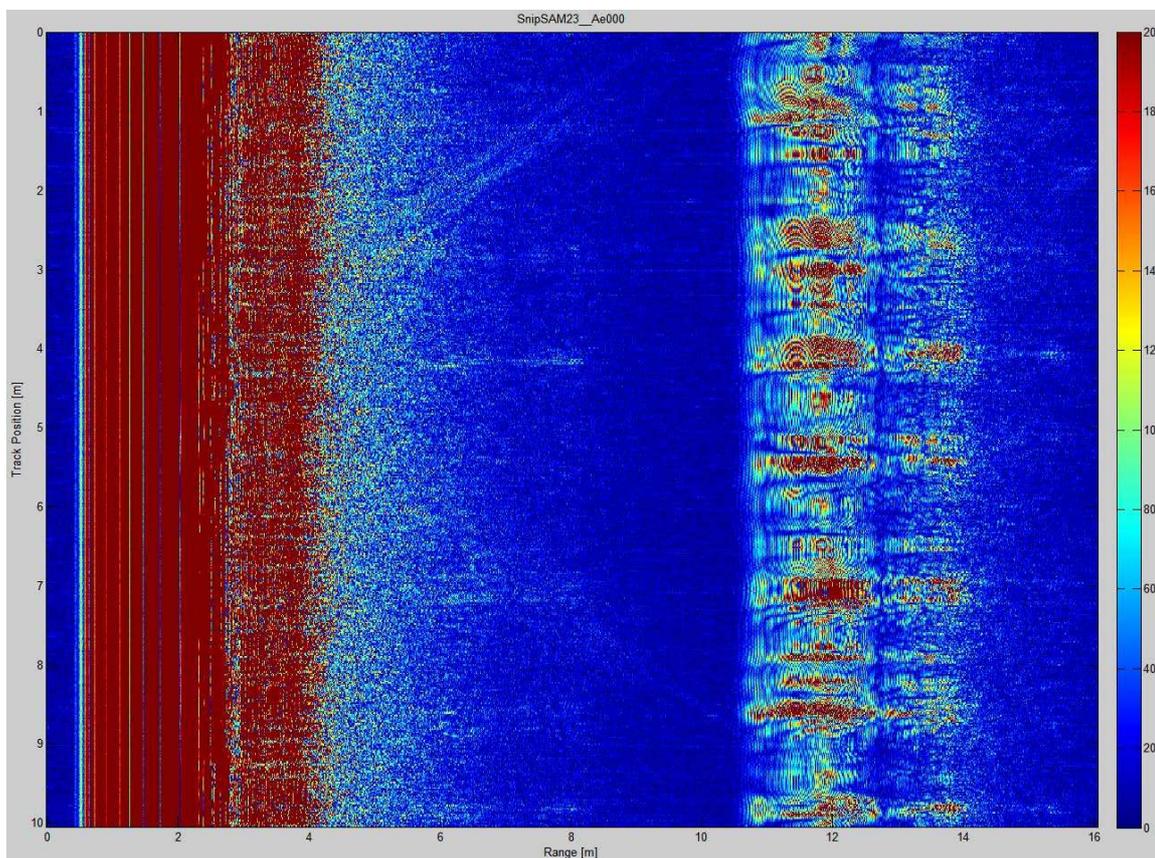


Figure 18: The phase histories (without range compression) for a single channel of the 40-channel BOSS array as the platform tracks over a partially buried, cylindrical target. This data was acquired at AUV FEST 2007 and is one several data samples being used to validate the high-fidelity simulation method developed under this effort.

image for the target in this data snippet is shown in Fig. 19. The modeling objective is to match the waveforms and/or the beamformed image products.

Our initial BOSS data modeling efforts are for cylindrical targets because these targets represent a common ordnance class. Because of the low frequencies required to insonify deep into the sediment, a SAS approach designed to detect buried targets will necessarily be low resolution, typically 10-20 cm. Hence, the measurable differences between the various bomb-shaped and torpedo-shaped targets when fully buried are minimal. A classifier robust enough to operate against buried targets in multiple environmental and clutter conditions will find it difficult to distinguish the specific cylindrical target type in all but the most pristine of conditions. Thus, examining the features needed to differentiate a cylindrical target from common environmental clutter is sufficient

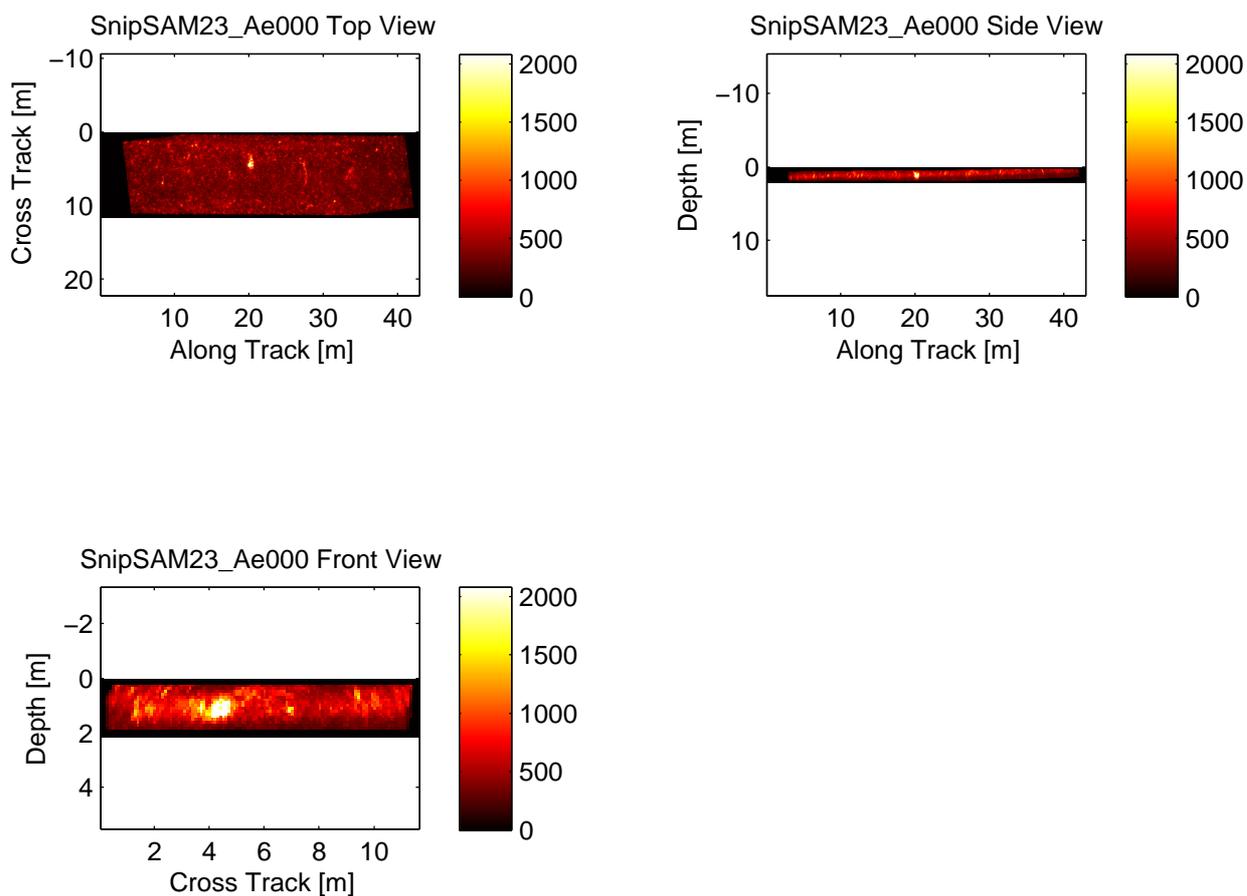


Figure 19: Various views of BOSS beamformed imagery for a data cube containing a target scattering response from the AUV FEST 2007 collection. **Upper left:** Top view of a target object shown as the maximum intensity projection (MIP) $P_{xy}(x, y)$ (see eq. 114). **Upper right:** Side view of a target object showing the MIP $P_{yz}(y, z)$ (see eq. 115). **Lower left:** Front view of a target object showing the MIP $P_{xz}(x, z)$ (see eq. 115).

to classify any ordnance of similar length and composition.

In the following we present results of a 3D simulation performed using the simulation package SPECFEM3D developed under this program for an elastic sediment layer. The schematic design of the model geometry is displayed in Figs. 20 – 22. The model dimensions are $16 \times 2.5 \times 6$ m. We use a source/receivers geometry reflecting the UUV and BOSS sensor geometry, and that was shown in (Fig. 10). We model a source waveform close to that secured by our AST partner, as a chirp signal sweeping linearly from 2 kHz to 20 kHz (Fig. 23).

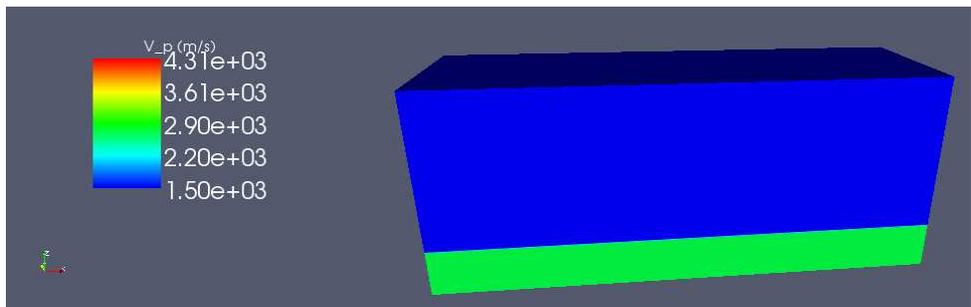


Figure 20: The model dimensions for simulation are $16 \times 2.5 \times 6$ m. This figure shows a 3-D view of the model with compressional wave speed distribution.

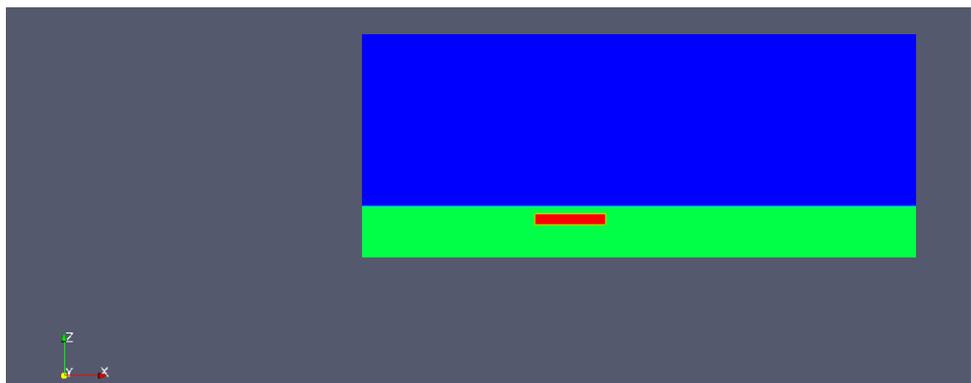


Figure 21: Side view of the model, where the target in red can be seen.

In Fig. 24, the time series at one receiver over 43 pings are displayed. The background signal has been eliminated, and thus, we only see the signature of the buried object.

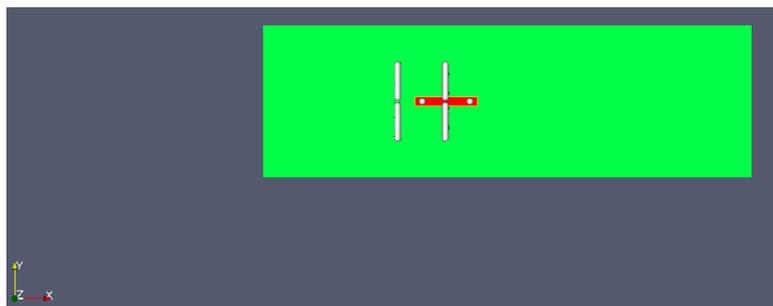


Figure 22: Top view of the model showing two different positions of source-receivers vehicle.

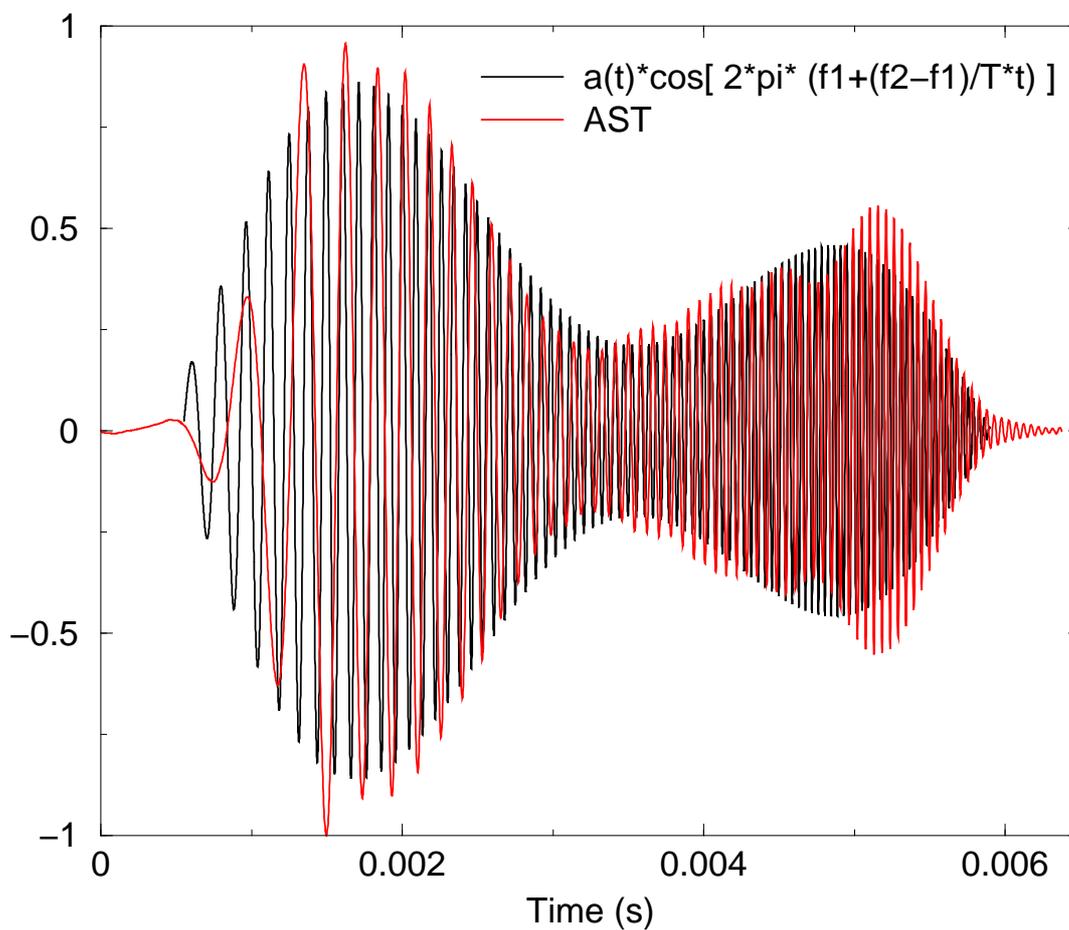


Figure 23: Source waveform represented as a chirp signal sweeping linearly from 2 kHz to 20 kHz. This figure shows the comparison between the actual BOSS waveform for the data (labeled AST), and our simple model for this signal.

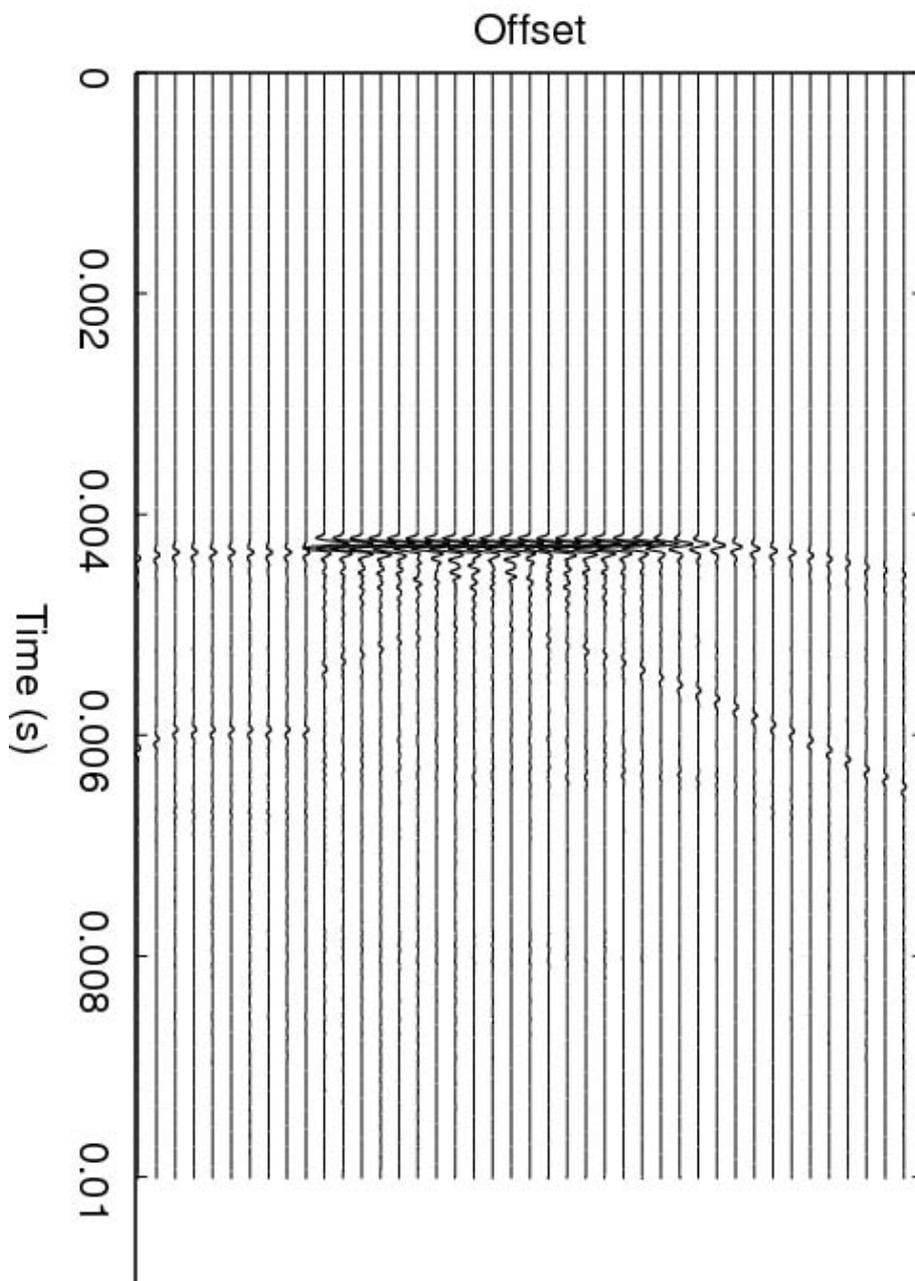


Figure 24: Time series at one receiver over 43 pings. The BOSS platform moves along the target. The background signal has been eliminated. The buried object signature is the only visible signal.

Figure 25 shows an example of the BOSS beamform processor applied to SEM-generated synthetic traces over the 40 receivers in the model and for a sequence of 43 pings as the BOSS platform moves over a linear trajectory at 2 m/s.

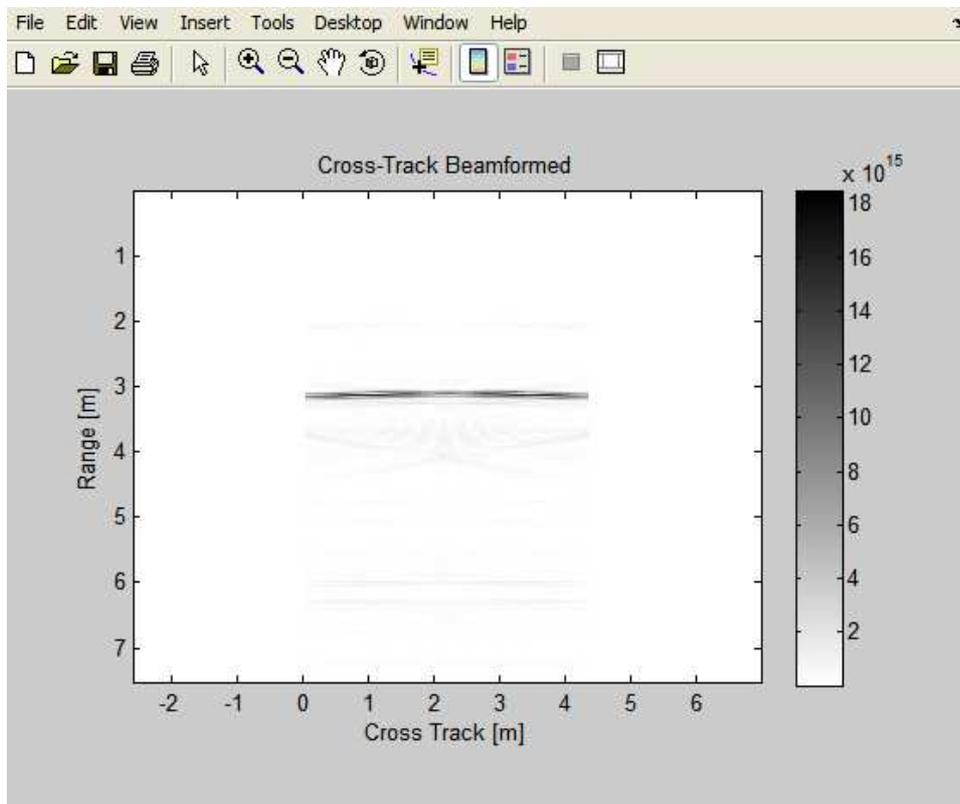


Figure 25: SAS Beamforming realized by Kent Harbaugh from Applied Signal Technology, Inc. using our 3-D numerical results: this shows a data slice parallel with the wing.

4.3 Buried Object Scanning Sonar (BOSS) data

We present a general overview of the BOSS data in Sec. 4.3.1 and a summary of the specific collections in Secs. 4.3.2 - 4.3.5.

4.3.1 BOSS data collections

We have acquired from Florida Atlantic University sonar collected in the BOSS⁶ experiments for our initial data analysis. The BOSS collection set was chosen as it fairly extensive and the sites contain a number of emplaced UXO-like targets. In addition, the site-properties and sensor deployments are relevant to the shallow water conditions of interest including sand and mud media and shallow water depths.

All of the data presented in what follows was collected in St. Andrew's Bay (SAB) near Panama City, Florida by Florida Atlantic University (FAU). The data collections were over two different target fields – one at a 'sand site' just inside SAB, south of Panama City, and the other at a 'mud site' in the northern part of SAB, directly West of the airport. At these sites the unmanned underwater vehicle (UUV) was running at constant-depth, and the seabed was slightly sloping. These target fields were long and narrow, and contained a variety of target shapes – some inert mine shapes and some other clutter. For each site the collections were run repeatedly parallel to the length of the field, with sequential small offsets from the centerline on each run. Thus, this collection plan yielded many looks at the same targets from slightly different aspects.

The receiver array element were on the underside of 'wings' on either side of the UUV and the source was directly under the UUV body. The projected acoustic source signal is an FM chirp sweeping 2 kHz to 20 kHz. Since none of the targets were 'deeply' buried, the lower end of this spectrum was not necessary for sufficient sediment penetration. In fact, in our processing⁷ it was found that the image appearance was improved when the lower frequencies were suppressed. Thus, all the images shown below were generated with using a bandwidth of 10-18 kHz.

The four key BOSS data collections available for analysis are (a) the 2004 Disk BOSS Trial, (b) the 2006 BMH Sea Trials, (c) the AUV FEST 2007 BMH Sea Trial and (d) the AUV FEST 2008 Trial. These are described, respectively, in Secs. 4.3.2 – 4.3.5.

We have processed (beamformed imagery available for the 2006, 2007 and 2008 AUV FEST data collections. In addition, we have derived feature sets for target/clutter items from each set of imagery. In terms of the classification analysis presented in the later sections of this report, we consider only the 2008 data. Results for 2006 and 2007 data sets will be delivered at a later time.

4.3.2 2004 Disk BOSS Trial

The 2004 collection was performed for an array of receivers arranged in a disk geometry. Although this data is available, it was acquired with an earlier version of the BOSS system, and we therefore find it less suitable than the three collections described below, and therefore, we do not pursue it further.

⁶See <http://www.oe.fau.edu/CHIRP/boss.html>

⁷The processing was performed by Applied Signal Technology using an executable provided to us by FAU.

4.3.3 2006 BMH Sea Trials

The 2006 collections were performed for a UUV platform with a wing geometry, and in which 20 receivers were arranged on each of the wings. The site conditions are sands and muds, and the burial depths range from proud targets to buried a few meters deep.

The 2006 data collection was performed by Florida Atlantic University (FAU) in St. Andrew’s Bay (SAB) near Panama City. The collections were over two different target fields – one at a ‘sand site’ just inside SAB, south of Panama City, and the other at a ‘mud site’ in the northern part of SAB, directly West of the airport.

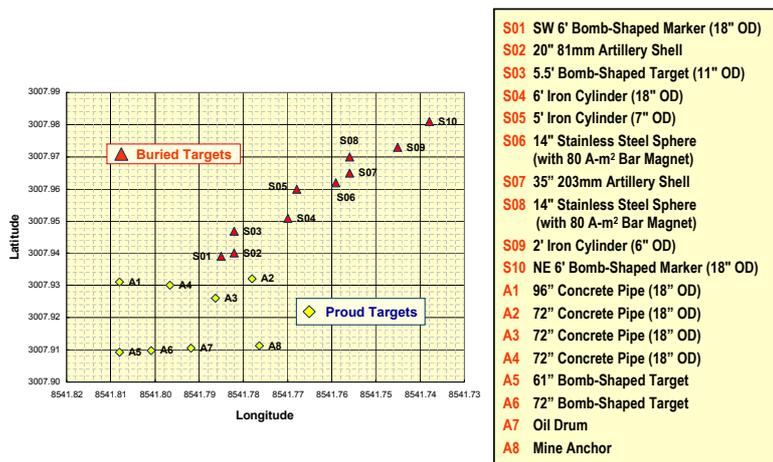


Figure 26: St. Andrew Bay 10-m water depth sand site data collection from 2006.

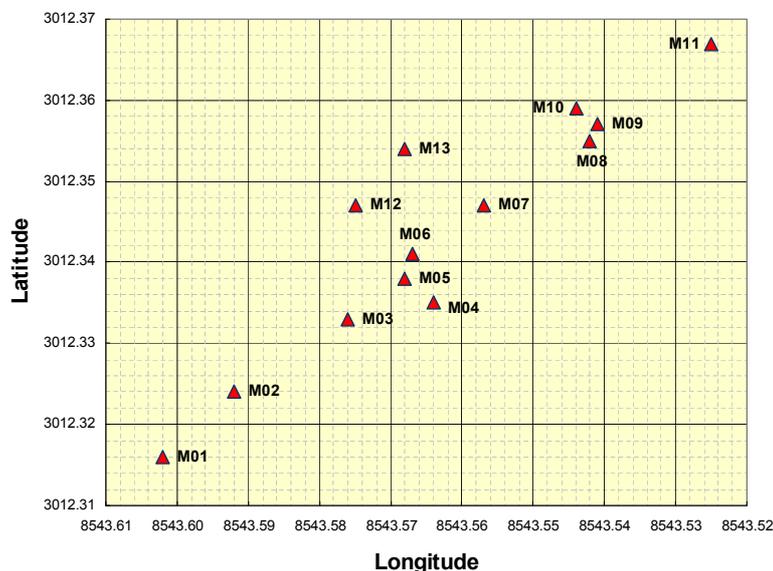


Figure 27: St. Andrew Bay 10-m water depth mud site data collection from 2006. Buried targets are indicated by the red-filled triangles.

At these sites the UUV hosting the BOSS system was running at constant-depth, and the seabed

was slightly sloping. These target fields were long and narrow, and contained a variety of target shapes – some inert mine shapes and some other clutter. For each site the collections were run repeatedly parallel to the length of the field, with sequential small offsets from the centerline on each run. Thus, this collection plan yielded many looks at the same targets from slightly different aspects. Figure 26 shows the geographic distribution at the *sand site* for many of the labeled targets (both proud and buried) and known clutter objects in Table 3. Figure 27 shows the clutter object distribution at the mud site. The identity of most of these is indicated in the table.

Internal Label	Index	Target/Clutter	Instances	Description
S1, S10	0	clutter	4	End marker
S3	1	target	4	5.5ft-long bomb
S4	2	target	7	6ft-long cylinder
	3	target	7	NSWC target
S5	4	target	2	5ft cylinder
S6, S8	101	clutter	25	14in-diameter sphere
S7	102	clutter	4	Bullet-shape object
S9	103	clutter	2	2ft-long cylinder
	104	clutter	1	100mm artillery shell
S2	105	clutter	1	81mm artillery shell
	106	clutter	4	Sand-filled scuba tank
	107	clutter	5	Cement-filled metal tube
	108	clutter	0	Sand-filled pipe
M1	111	clutter		Unaltered pipe
M2	112	clutter		Cylinder A 2in lip
M3	113	clutter		Cylinder B 2in indent
M4	114	clutter		Cylinder C flat finish
M5	115	clutter		Cylinder D 9in hem
M6	116	clutter		Cylinder E lip 30 deg
M7, Y19, Y26	117	clutter	2	Oil drum
M8, Y21, Y28	118	clutter		Clump
Y20, Y27	119	clutter		Tire
Y22	120	clutter		Magnet
	121	clutter		Large Lim cylinder
	122	clutter		Medium Lim cylinder
	123	clutter		Small Lim cylinder
	200	clutter	694	Unidentified clutter object

Table 3: Summary of target and clutter objects of BOSS data taken 2004-2007. These provide sufficient data to test algorithms for target versus clutter discrimination and for target classification.

4.3.4 AUV FEST 2007 BMH Sea Trial

The BOSS 2007 AUV-Fest data collection is very similar to the 2006 collection described above, but with improved control and knowledge of targets, and so we are using targets from this collection validation of our simulation method. In addition, data snippets from this collection are being used to validate the 3D SEM simulation method that we have developed. The available target and clutter distribution for the 2007 collection is shown in Fig. 28.



Figure 28: Yankee test site target locations from the 2007 collection.

The receiver array element were on the underside of ‘wings’ on either side of the UUV and the source was directly under the UUV body. The projected acoustic source signal is an FM chirp sweeping 2 kHz to 20 kHz. Since none of the targets were ‘deeply’ buried, the lower end of this spectrum was not necessary for sufficient sediment penetration. In fact, in our processing it was found that the image appearance was improved when the lower frequencies were suppressed. Thus, for this data set images are generated with using a bandwidth of 10 kHz to 18 kHz.

4.3.5 AUV FEST 2008 Trial

The AUV FEST 2008 site was located off the coast of Newport, RI near the Naval Undersea Warfare Center (NUWC), and the collection was conducted in May, 2008. The AUV (Autonomous Undersea Vehicles) FEST was co-sponsored by NOAA and the Office of Naval Research, and featured a variety of autonomous vehicle technologies including BOSS. Newport was selected because Narragansett Bay offers a realistic environment to evaluate the performance of potential littoral mine countermeasure systems. Also, at the same time it was an ideal location to explore the seabed for historical shipwrecks. It is worthwhile to consider the phenomenology of these features in detail for more informed development of classifiers that depend upon them.

4.4 Summary of target/clutter data descriptors available for classification analysis

We have developed feature data sets from the 2006, 2007 and 2008 AUV FEST data collections described above. In each of these, varying levels of detail are available for the target/clutter descriptions, and there are varying instances of targets and clutter for which data were recorded at each of the sites. With respect to the 2006 data collection, we have a set of 50 features for both target and clutter items. In this data set there are 17 instances of targets and 10 instances of clutter. Descriptions of the target/clutter items in this data set are summarized in Table 4. Although we do have some specificity for some of the targets in this collection, other targets are only named as such, and are not specifically described since they are classified. These classified targets comprise the first 10 entries in the cited table. The geographic distribution of targets in this data set was

Row Index	Target/clutter label	Object	Description	Data structure name
1	y18	target	no label (classified)	Targ2006
2	y18	target	no label (classified)	Targ2006
3	y23	target	no label (classified)	Targ2006
4	y24	target	no label (classified)	Targ2006
5	y23	target	no label (classified)	Targ2006
6	y25	target	no label (classified)	Targ2006
7	y23	target	no label (classified)	Targ2006
8	y25	target	no label (classified)	Targ2006
9	y15	target	no label (classified)	Targ2006
10	y23	target	no label (classified)	Targ2006
11	s4	target	6ft-long cylinder	Targ2006
12	s5	target	5ft-long cylinder	Targ2006
13	s5	target	5ft-long cylinder	Targ2006
14	s4	target	6ft-long cylinder	Targ2006
15	s4	target	6ft-long cylinder	Targ2006
16	s4	target	6ft-long cylinder	Targ2006
17	s4	target	6ft-long cylinder	Targ2006
1	y19	clutter	oil drum	Clut2006
2	y19	clutter	oil drum	Clut2006
3	s3	clutter	unknown	Clut2006
4	s1	clutter	end marker	Clut2006
5	s1	clutter	end marker	Clut2006
6	s1	clutter	end marker	Clut2006
7	s7	clutter	bullet-shape object	Clut2006
8	s6	clutter	14in-diameter sphere	Clut2006
9	s10	clutter	end maker	Clut2006
10	s9	clutter	2ft-long cylinder	Clut2006

Table 4: Summary of target and clutter objects from AUV Fest 2006 collection for which BOSS data is available for classifier analysis.

shown in Fig. 27.

The BOSS data from the 2007 AUV FEST collection consists of 29 targets and 49 clutter

items, as summarized in Table 5 below. There are 50 features that have been extracted for each item. Unfortunately, only the target/clutter labels are available for each contact. The targets are geographically distributed as shown in Fig. 28.

Row Index in .mat file	Target/clutter	Description	Data structure
1	target	no label (not available)	Targ2007
2	target	no label (not available)	Targ2007
...	target	no label (not available)	Targ2007
29	target	no label (not available)	Targ2006
1	clutter	no label (not available)	Clut2007
2	clutter	no label (not available)	Clut2007
49	clutter	no label (not available)	Clut2007

Table 5: Data for AUV Fest 2007 collection

The 2008 field objects are the same types of objects as those listed Table 4 in the 2006 collection. However, individual target/clutter labels for each feature vector is not available at this time. This is because the identity of the target/clutter items is classified. The data itself is not classified, but it would become classified if labels were attached to any of the detections. Instead, only the binary target/clutter descriptor is available. In total, 89 target items and 132 clutter items are included in this data set, as summarized in Table 6 below.

Row Index in .mat file	Target/clutter	Description	Data structure
1	target	no label (not available)	FeatTarg2008
2	target	no label (not available)	FeatTarg2008
...
89	target	no label (not available)	FeatTarg2008
1	clutter	no label (not available)	FeatClut2008
2	clutter	no label (not available)	FeatClut2008
...
132	clutter	no label (not available)	FeatClut2008

Table 6: Target/clutter data available for the AUV FEST 2008 BOSS collection.

Feature index	Feature class	Feature name
1-3	Geometric	Geometric center
4-6	Geometric	Best-fit ellipsoidal radii: A, B, C
7-9	Geometric	Direction of radius vector associated with A
10-12	Geometric	Direction of radius vector associated with B
13-15	Geometric	Direction of radius vector associated with C
16-18	Geometric	Cross-sectional areas: $\pi AB, \pi AC, \pi BC$
19-21	Geometric	Length-to-width ratios: $A/C, A/B, B/C$
22	Geometric	Ellipsoidal volume: πABC
23	Geometric	Voxel volume: $N \Delta X \Delta Y \Delta Z$
24	Intensity	Total energy
25	Intensity	Ellipsoidal energy density
26	Intensity	Voxel energy density
27	Intensity	Voxel energy standard deviation
28	Intensity	Voxel energy s.d. normalized by mean
29	Intensity	Peak intensity
30-32	Probabilistic	3-bin PDF of relative intensity
33-41	Probabilistic	9-bin PDF of relative intensity
42-50	Probabilistic	9-bin CDF of relative intensity

Table 7: This table describes in summary form the 50 feature classes derived for each 3D beamformed target or clutter detection. These features define the available components of an input feature vector for use with the relevance vector machine classifier. Here PDF and CDF indicate the probability and cumulative density functions, respectively.

4.5 Feature analysis of BOSS data from AUV FEST 2008

In this section we provide detailed views into the nature of various morphological features derived from the 3D beamformed intensity images of BOSS data collected for a distribution of targets and clutter at the AUV FEST 2008 data collection site. The features derived from this data are summarized in Table 7, and as discussed later, various subsets of these features are used in the classifier development.

Features may be derived from data inputs in numerous ways including descriptors derived from raw observed waveforms to morphological or statistical properties derived from 2D and 3D image products. Image products are typically formed from beamforming, but other images types are also possible. *Tomographic estimates of material attributes* obtained, for example, from adjoint tomography (see Sec. 3.4) can provide a potentially extremely valuable discriminant. Tomographic maps can be independently derived for such quantities as shear wave speed, fast and slow P-wave speeds, density, porosity, permeability, and other poroelastic properties. These estimates are for the sediment medium and any embedded targets, and provide complementary information to the scattering intensities derived from beamforming. Estimates of the spatial-distributed values for each of the parameter types may be used to segregate natural clutter from targets, and more precisely define the target shape and composition. An additional feature class is *acoustic color* or spectral properties measured as a function of target/sensor aspect and spectral band.

Desirable attributes for feature inputs include robustness to natural variability, and high dis-

crimination power for targets versus man-made and natural clutter. An additional issue for feature identification is not just the class of feature, but tuning parameters used to extract the feature. As an example, various geometric-based and intensity-based features depend on a threshold operation applied to volumetric image products. Knowledge of environment-induced contribution to these features can be used for intelligent threshold selection. The value of the simulations for feature development is that the propagation effects and data collection parameters that enhance or degrade the utility of features for classification can be systematically evaluated.

4.5.1 3D beamformed imagery examples from BOSS data collection

The features cited in Table 7 are derived from 3D beamformed image products. Figure 29 shows an example of this beamformed imagery using projected image intensities for various views (in 2D). Additional example images along with 3D isocontours of scattering intensity are shown in App. A.2.

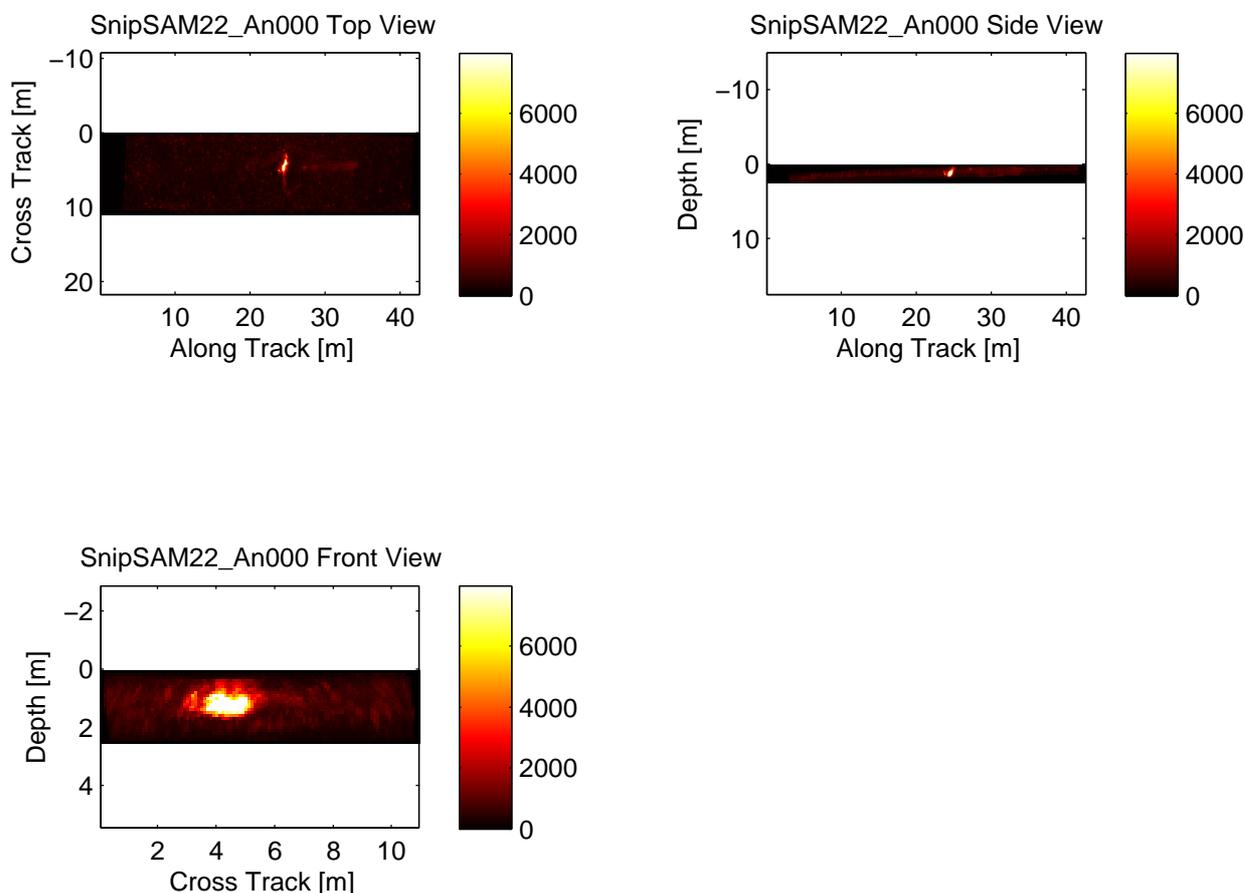


Figure 29: Various views of BOSS beamformed imagery for a data cube containing the scattering response for a cylinder shaped object from the AUV FEST 2007 collection. **Upper left:** Top view of a target object shown as the maximum intensity projection (MIP) $P_{xy}(x, y)$ (see eq. 114). **Upper right:** Side view of a target object showing the MIP $P_{yz}(y, z)$ (see eq. 115). **Lower left:** Front view of a target object showing the MIP $P_{xz}(x, z)$ (see eq. 115).

4.5.2 Histogram analysis of target/clutter feature distributions for AUV FEST 2008 data

We consider here an empirical example of the feature distributions for the target and clutter items available for analysis in the AUV FEST 2008 BOSS data set. As described in Sec. 4.4 and Table 6, there are 89 target exemplars and 132 clutter exemplars in this data set. The 50 features from Table 7 were derived for each of these target/clutter exemplars, and we obtained feature histograms for each of these. As noted earlier, due to reasons of security classification, only the binary target/descriptor labels are available for this data set. Therefore, for each feature, we generated two different feature histograms corresponding to target population and to the clutter population. The histogram results for the first three features in Table 7 are shown in Fig. 30. Histograms for the remaining 47 features in the target and clutter populations are shown in App. A.3.

Of course, for a given feature class, if the feature distribution of the target and clutter populations strongly differ, the feature has the potential to be a strong discriminator input for a classifier algorithm. More commonly, feature distributions for target/clutter items can have significant overlap. However, the hope is that by combining inputs from multiple feature classes, improved discrimination can be achieved.

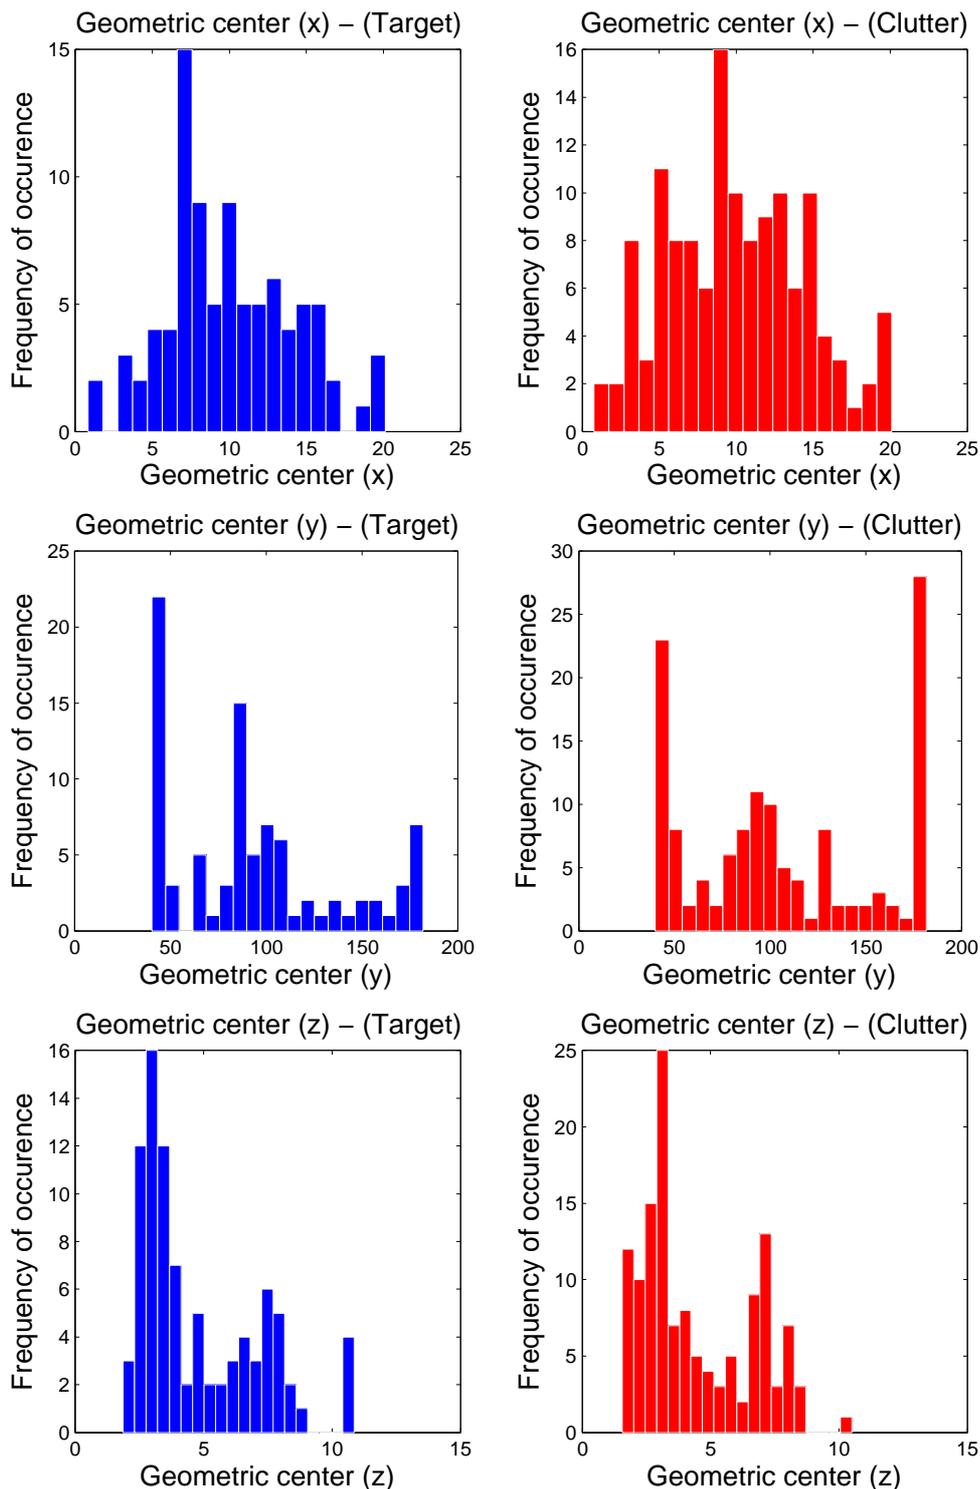


Figure 30: Distribution histograms for the *geometric center* x coordinate (top panel), y coordinate (top panel) and z coordinate (bottom panel) feature types. In each case the target and clutter features are shown on the left hand and right hand sides, respectively. The features were derived from 3D beamform products using data acquired by the BOSS platform at the AUV Fest 2008. A total of 132 clutter items and 89 targets were processed in generating these distributions.

4.5.3 Feature pair analysis of target/clutter features for AUV FEST 2008 data

An informative method to evaluate the discrimination power of features is by examination of the 2D spatial spread of features for various feature pairs. Of course, for feature vectors of large dimensionality, it is impossible to examine all possible feature pairs. However, in classifier construction, we usually use a small subset of the available features. For the case of the feature subset corresponding to features 4, 5, 20, 26, 29, 30 and 32 (see Table 7), we present feature pair plots for all possible feature pairs (from the AUV FEST 2008 data set). Several example feature pairs for all available target/clutter instances are shown in Fig. 31. Appendix A.4 shows the feature pair plots for all of the other cases.

The results in Fig. 31 show good separation for several of the feature classes amongst target/clutter pairs, but within a fixed feature pairs, there remains significant overlap for many of the exemplars. Again, by using a sufficient number of features, the hope is that the feature distance between targets and clutter in the multi-dimensional feature space is sufficient to provide good discrimination capability.

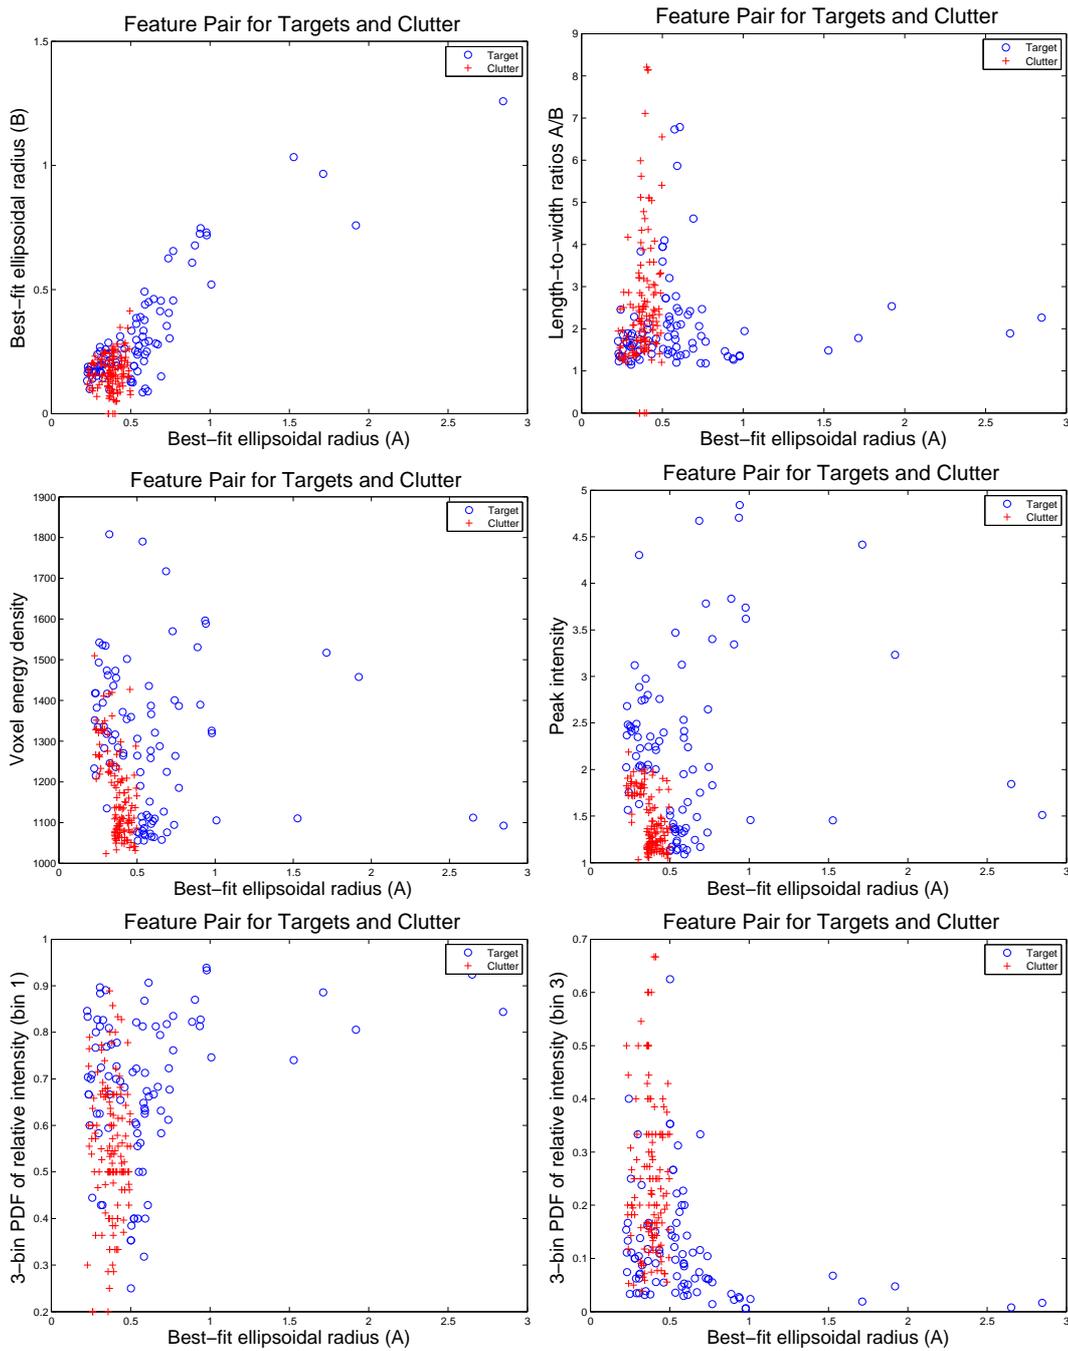


Figure 31: Features derived from the targets and clutter in this data set from the BOSS AUV FEST 2008 collection are indicated by the circles and plus signs, respectively. Scatter plots for various feature pairs are shown. **Upper left:** *ellipsoid radius A* versus *ellipsoid radius B*. **Upper right:** *ellipsoid radius A* versus *ellipsoid length to width ratio A/B*. **Middle left:** *ellipsoid radius A* versus the *voxel energy density*. **Middle right:** *ellipsoid radius A* versus the *peak intensity*. **Lower left:** *ellipsoid radius A* versus the *first bin of the 3-bin PDF of the relative intensity*. **Lower right:** *ellipsoid radius A* versus the *third bin of the 3-bin PDF of the relative intensity*.

4.5.4 Feature triplet analysis of target/clutter features for AUV FEST 2008 data

Complementary to the feature pair analysis presented in Sec. 4.5.3, we consider here feature analysis for *triplets* of features. Extending the feature dimensions to higher dimension can clearly be valuable as the extra dimension(s) can provide improved separability of target/clutter responses. Of course, there is a limit to this as generalization eventually degrades. It is useful to visualize feature separation in three dimensions and so consider all possible feature triplets for the subset of features 4, 5, 20, 26, 29, 30 and 32 (see Table 7). These correspond to the A and B best-fitting ellipsoid radii, the length to width ratio A/B , the voxel energy density within the best-fitting ellipsoid, the peak intensity, and the first and last bins of a 3-bin probability density distribution of relative intensity values for 3D voxels within the best-fitting ellipsoid.

An informative method to evaluate the discrimination power of features is by examination of the 2D spatial spread of features for various feature pairs. Of course, for feature vectors of large dimensionality, it is impossible to examine all possible feature pairs. However, in classifier construction, we usually use a small subset of the available features. For the case of the feature subset corresponding to features 4, 5, 20, 26, 29, 30 and 32 (see Table 7), we present feature pair plots for all possible feature pairs (from the AUV FEST 2008 data set). Several example feature pairs for all available target/clutter instances are shown in Fig. 32. Appendix A.4 shows the feature triplet plots for all of the remaining cases with respect to this feature subset.

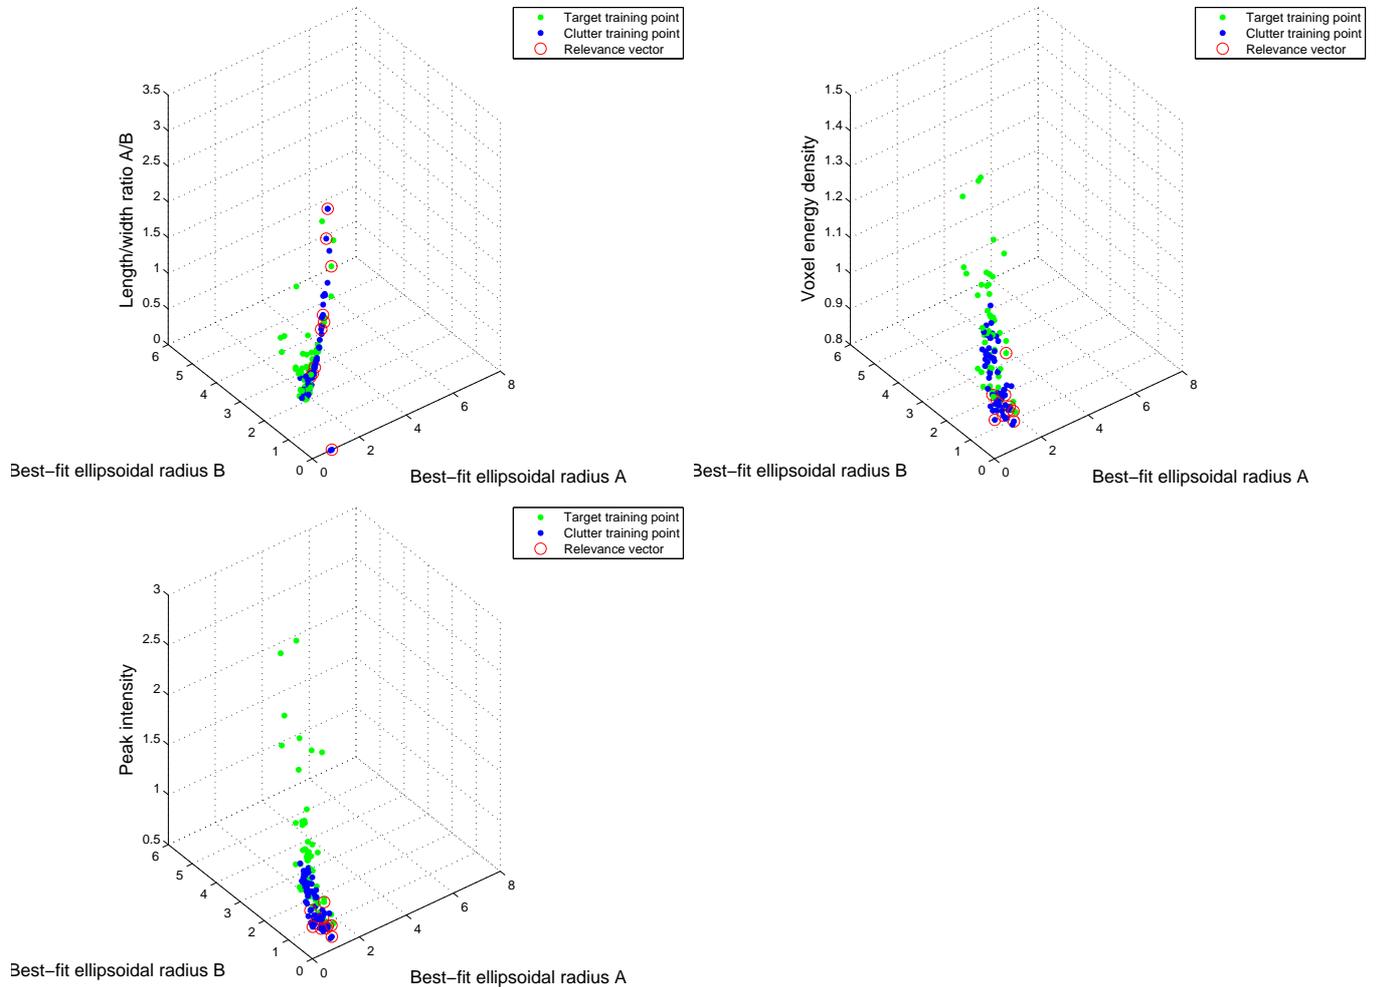


Figure 32: 3D scatter plots showing training points and optimal relevance vectors for various feature triplets. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors corresponding to known targets are given by the green circles. In addition, feature triplets circled in red are determined to be *relevance vectors* as determined for RVM analysis for the feature subset (4, 5, 20, 26, 29, 30, 32) as described in Sec. 4.6 and the *case 1* subset (see Table 8). **Upper panel:** Feature triplet *A*, *B* (*best fit ellipsoid radii*) and *Length to width ratio A/B*. **Middle panel:** Feature triplet *A*, *B* (*best fit ellipsoid radii*) and *Voxel energy density*. **Lower panel:** Feature triplet *A*, *B* (*best fit ellipsoid radii*) and *peak intensity*.

4.6 Numerical results for BOSS classification analysis

In this section we process the features derived from the AUV Fest 2008 data to obtain target/clutter classifications. A detailed examination of these features was provided in Sec. 4.5 and the appendices. The Bayesian analysis that we apply to the feature data for development of the classifier (based on the Relevance Vector Machine [35, 3]) was described in Sec. 3.5.

The sequence of classifier experiments that we performed is summarized in Table 8. Each experiment corresponds to the selection of a distinct subset of features for construction of the corresponding RVM classifier. In total, 20 distinct subsets of features were chosen, and the classifier results for each of these are shown in Figs. 34 – 53 (as indicated in the table). We also note that we used the available AUV FEST 2008 data to form both the training and test data sets. Typically, half of the available input training vectors are chosen to build the classifier, and of these, only a portion are used to form the relevance vectors that define the final classification model. These subsets of input vectors are randomly chosen, and so the classifier performance will vary for each set of inputs. Clearly, to obtain a statistical robust means to choose between different sets of features for the final classifier, there should be numerous realizations of the training inputs. An example of the corresponding results obtained is shown in Fig. 33 for the case of feature subset (1) (see Table 8). Corresponding statistical outcomes for the remaining 19 feature subsets considered are shown in App. A.6.

The sub-plots in each of Figs. 34 – 53 have the same interpretation. We note that each result shown is for a single realization of a random subsampling of the available training vectors. The *upper left* of each figure is a stem plot that shows the feature indices used for the classifier development. The RVM algorithm iterates over a maximum of 100 iterations. The *upper right* of each figure shows the progress of the algorithm for each iteration. The action of the RVM is to obtain a sparse set of relevance vectors from the clutter and target exemplars. For each case it can be seen that the starting set contains ~ 115 exemplars. As the iteration proceeds, this set is winnowed down until the final, optimal set of exemplars is discovered i.e., the relevance vectors. The size of the set will vary depending on the feature subset used. The *middle left* of each figure shows the values of the hyperparameters α_i associated with each weight w_i of the selected relevance vectors. The number of α_i shown, is the same as final number of basis points in the upper right figure. Recall that during the RVM training process, a significant number of the α_i are driven to a very large number (toward infinity) so that the corresponding mean and variance of the associated weight is centered on zero. Only those training vectors associated with α_i that are sufficiently distinct from infinity are retained, and those become the relevance vectors (see Sec. 3.5) and [35, 3] for further details). The *middle right* of each figure shows the receiver operator characteristic (ROC curve) obtained from the derived RVM classifier as applied to the training set. The *lower left* of each figure shows the receiver operator characteristic (ROC curve) obtained from the derived RVM classifier as applied to the test set (which is the complement of the training set with respect to the totality of all available data). Finally, the *lower right* of each figure shows the superposed ROC curves for both the training and the test sets.

Examination of the results in Figs. 34 – 53 reveals that, as expected, some feature subsets perform better than others. However, it should be noted that results in the mean (with respect to an ensemble of training vector sets) as shown in Fig. 33 and in App. A.6 should be used for identification of the superior feature input set. Each feature subset tested (as listed in Table 8) was chosen by intuition. For a more robust result, an optimizer should be used to discover

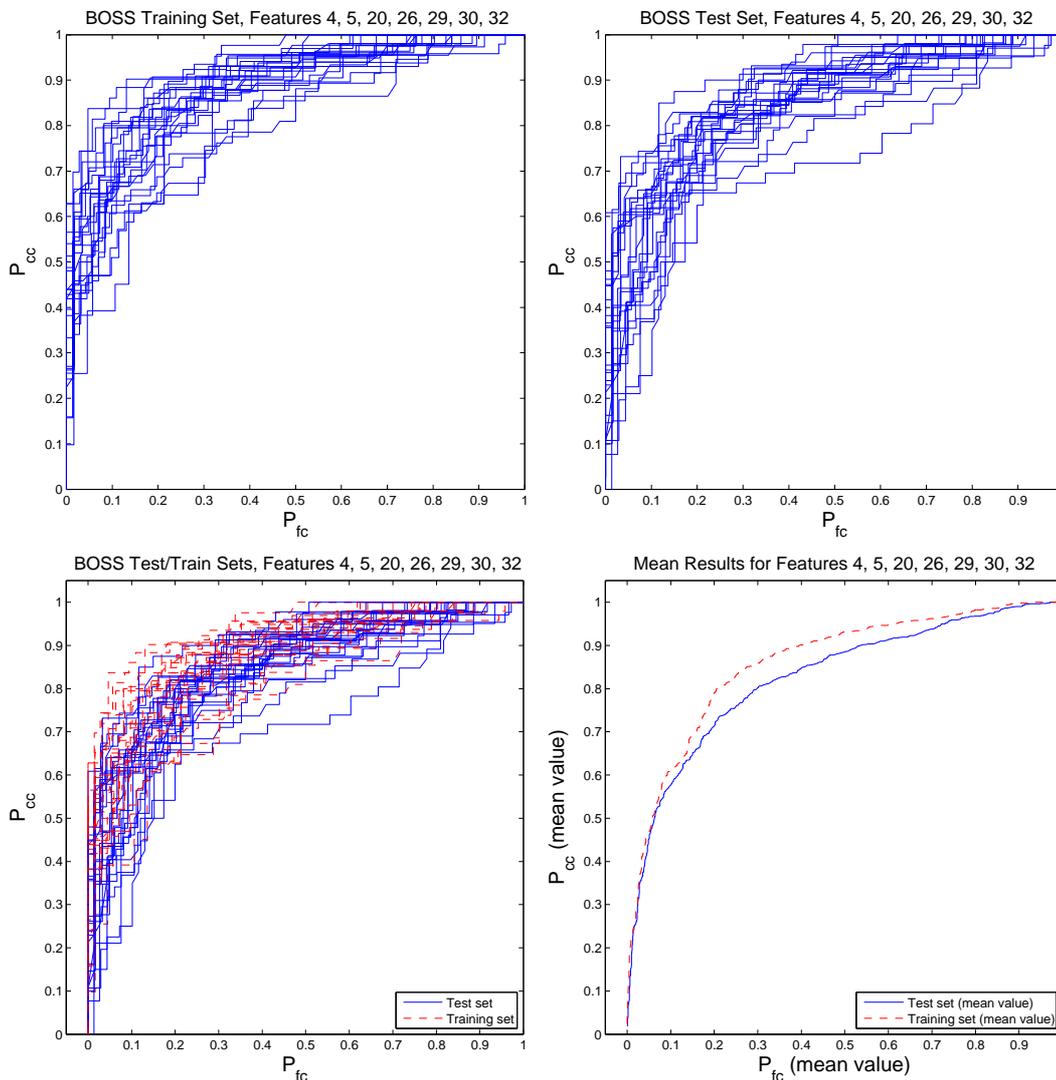


Figure 33: Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets.

Case	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Feature																				
1																				
2																				
3																				
4	•	•	•	•	•	•	•	•	•	•		•		•	•		•	•	•	•
5	•	•	•	•	•	•	•	•	•	•		•		•	•		•	•	•	•
6				•				•	•	•		•							•	•
7																				
8																				
9																				
10																				
11																				
12																				
13																				
14																				
15																				
16									•	•										
17									•	•										
18																				
19																				
20	•	•	•	•	•	•	•		•	•				•	•		•	•	•	•
21																				
22																				
23		•						•				•	•	•	•	•	•	•	•	•
24								•				•	•	•	•	•	•	•	•	•
25												•	•	•	•	•	•	•	•	•
26	•	•	•	•	•	•	•		•	•		•	•	•	•	•	•	•	•	•
27									•	•		•	•	•	•	•	•	•	•	•
28										•		•	•	•	•	•	•	•	•	•
29	•	•	•	•	•	•	•	•	•	•		•	•	•	•	•	•	•	•	•
30	•	•	•	•	•	•	•	•	•	•		•	•	•	•	•	•	•	•	•
31	•	•	•	•	•	•	•	•	•	•		•	•	•	•	•	•	•	•	•
32	•	•	•	•	•	•	•	•	•	•		•	•	•	•	•	•	•	•	•
33																				•
34																				
35																				
36																				
37																				
38																				
39																				
40																				
41					•															
42						•														
43						•														
44						•														
45						•														
46						•														
47						•														
48						•														
49						•														
50						•														
Figure	34	35	36	37	38	39	40	41	42	43	44	45	46	47	48	49	50	51	52	53

Table 8: Summary of feature subsets used for training relevance vector machine classifiers. From a given training set the RVM discovers the optimal subset of feature input vectors for use in a kernel-based classifier design. However, the design of the classifier is dependent on the dimension of the input feature vector. Use of all possible features (in this case, 50), typically leads to poor generalization. Therefore, we experimented with different subsets of features, and obtained the corresponding classifier designs and performance metrics. These subsets, chosen by intuition rather than an optimizer, are indicated in the columns of the table. Each row index corresponds to the feature index and is directly linked to the feature descriptors in Table 7. The columns labeled from 1 to 20 each indicate the feature subset used as indicated by the bullets, and the bottom row indicates the Figure number for which the performance metrics of the corresponding classifier are plotted.

the optimal subset of features. This could be readily implemented as a wrapper around the core algorithms already developed and implemented. Essentially, the feature subset selection problem is a combinatoric optimization problem. However, effective techniques for treating such problems have already been developed. Of course, we are not restricted to the feature set that was used, and we speculate that high-performing features can be discovered, and that is the subject of future research.

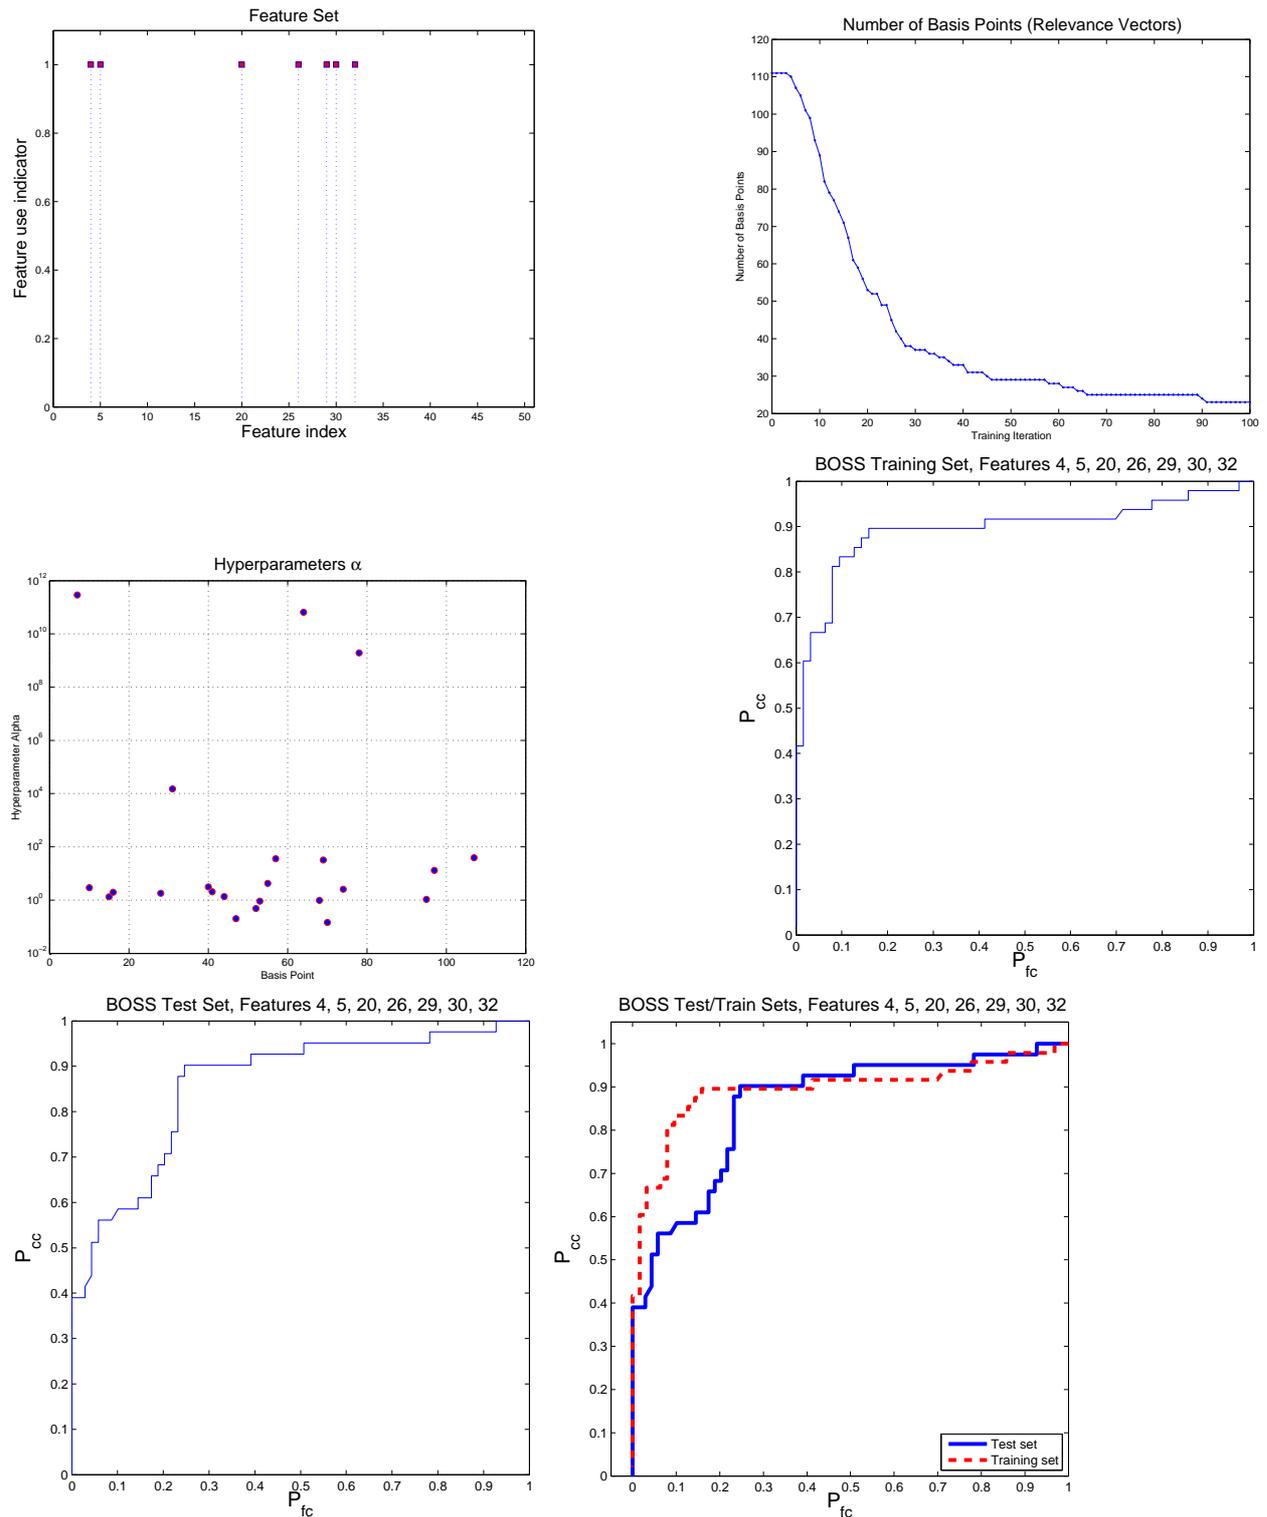


Figure 34: RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 1 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets.

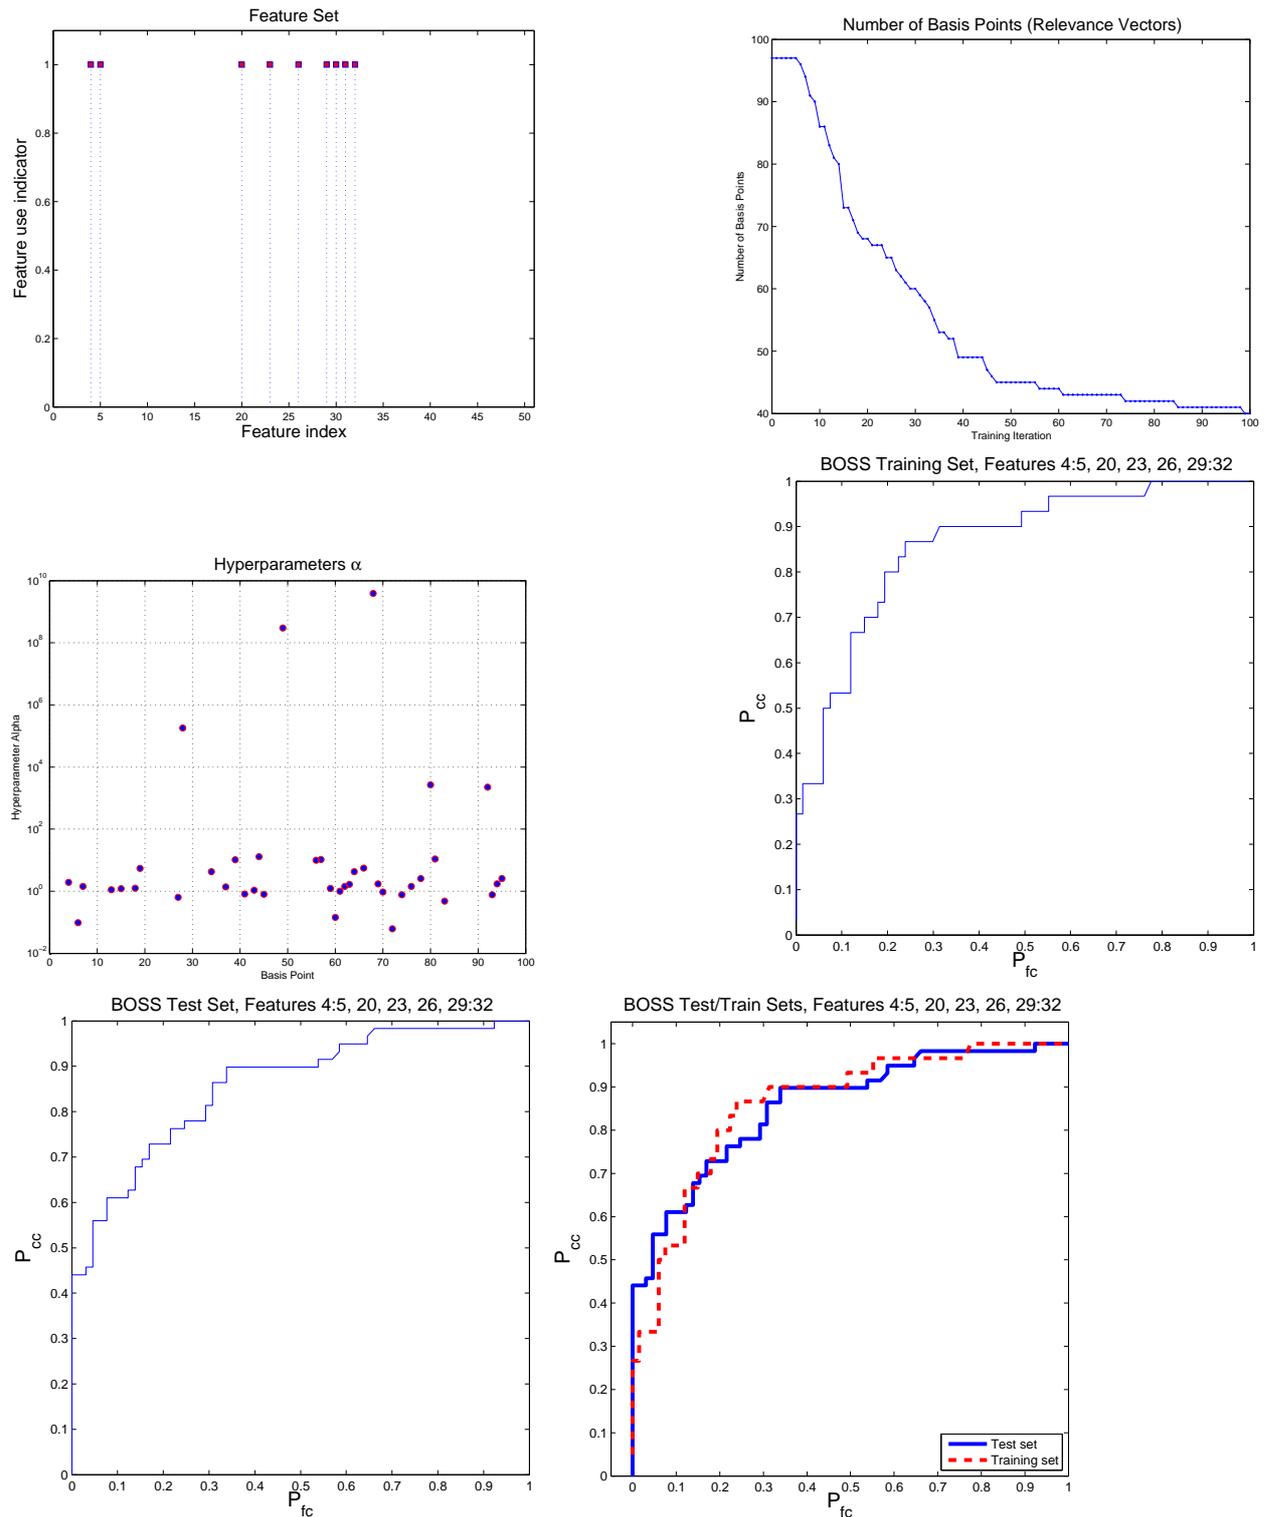


Figure 35: RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 2 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets.

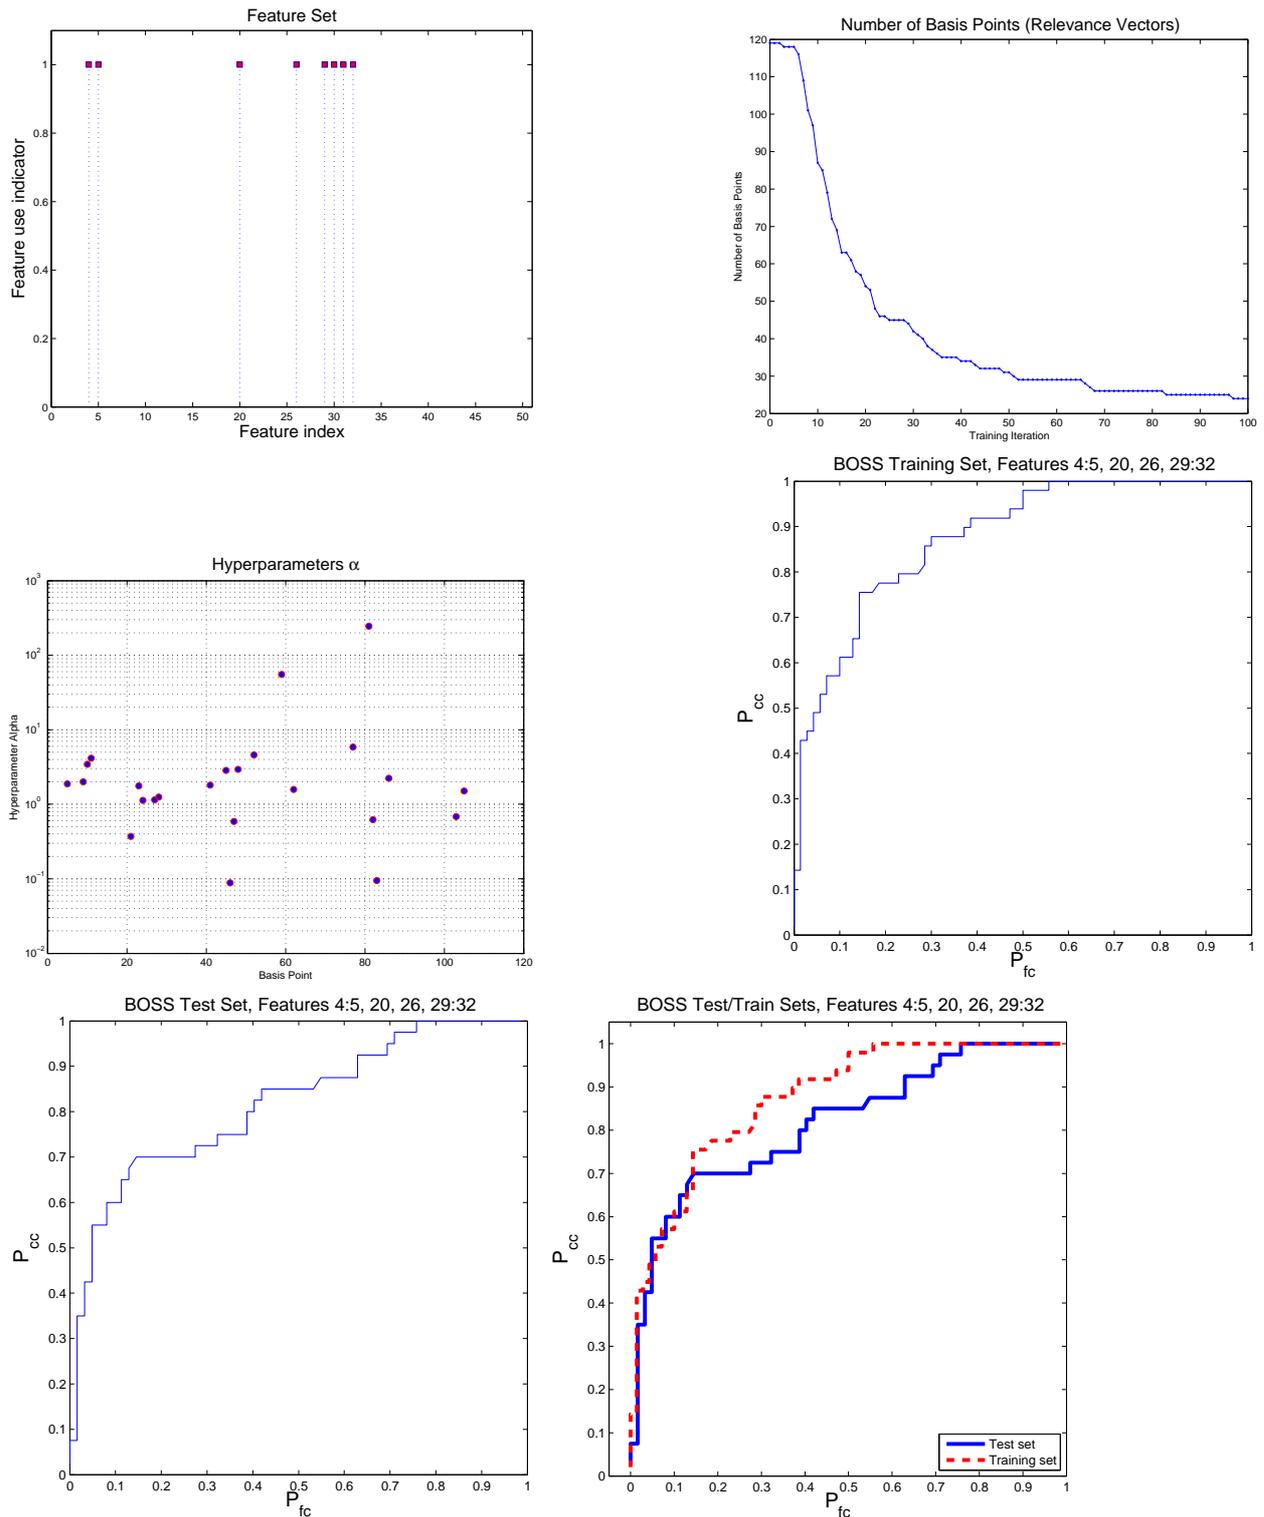


Figure 36: RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 3 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets.

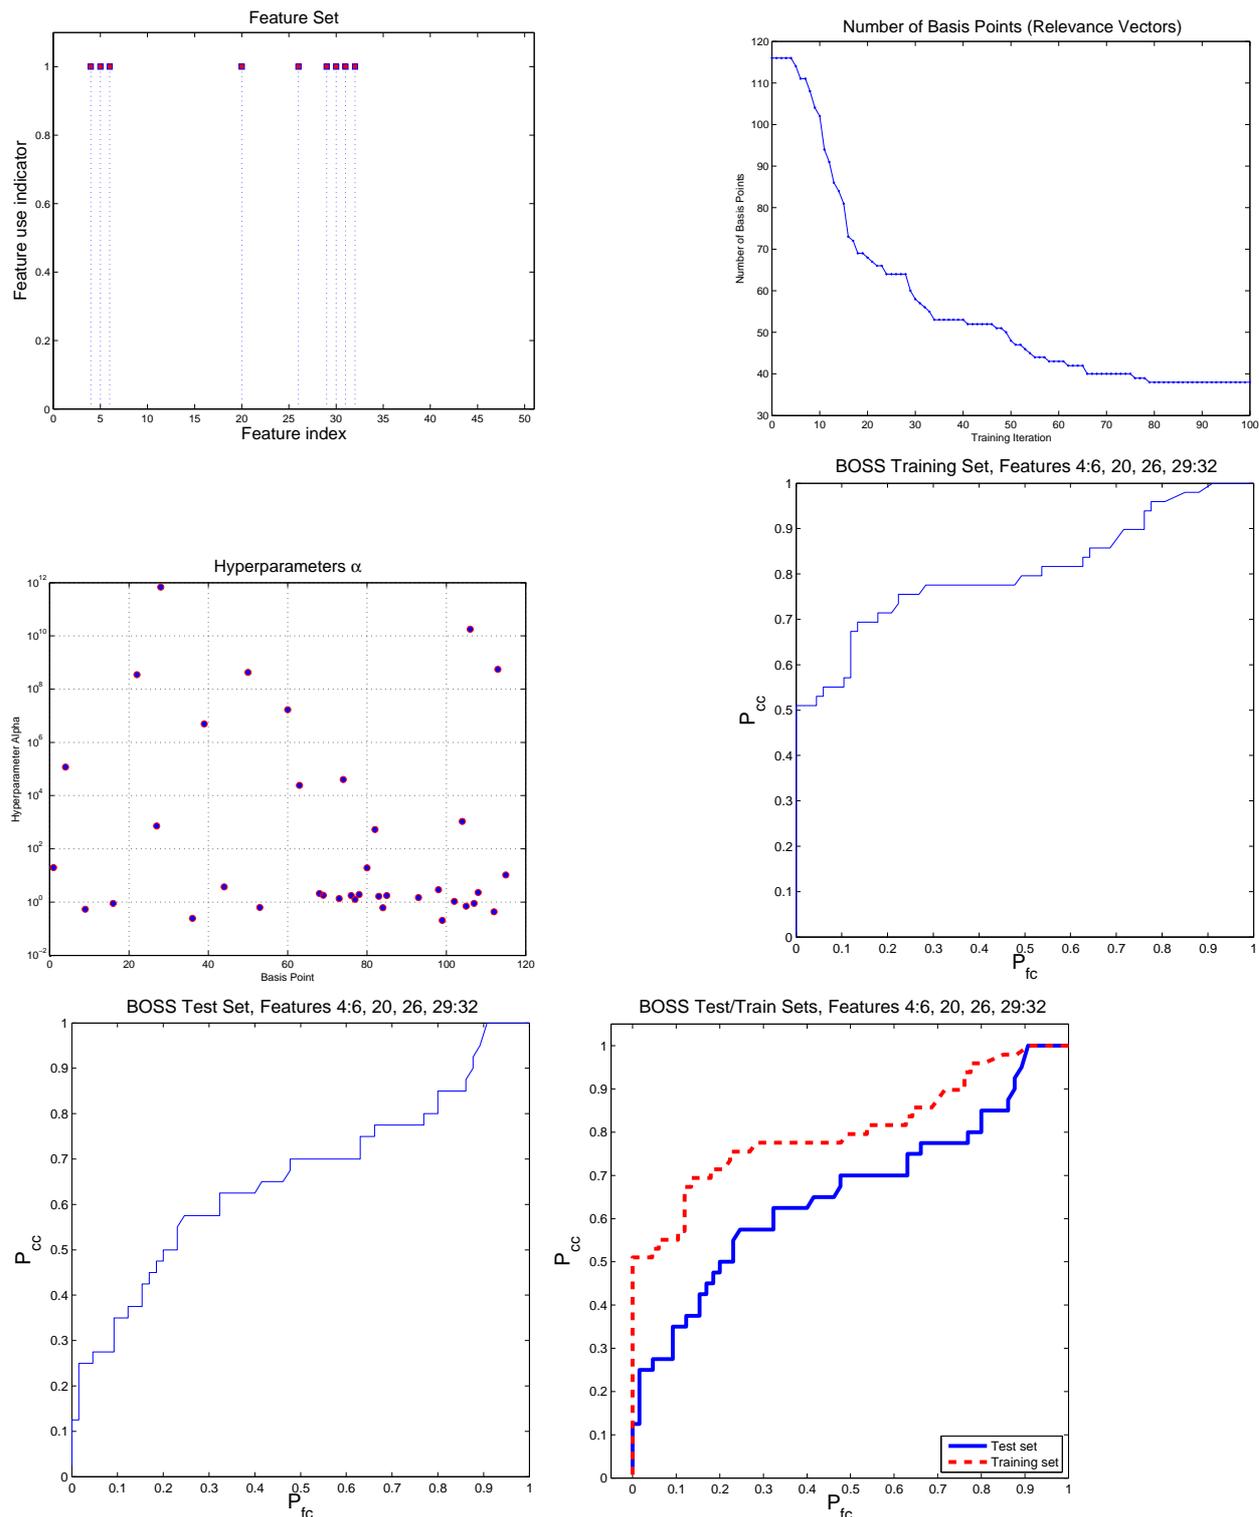


Figure 37: RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 4 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets.

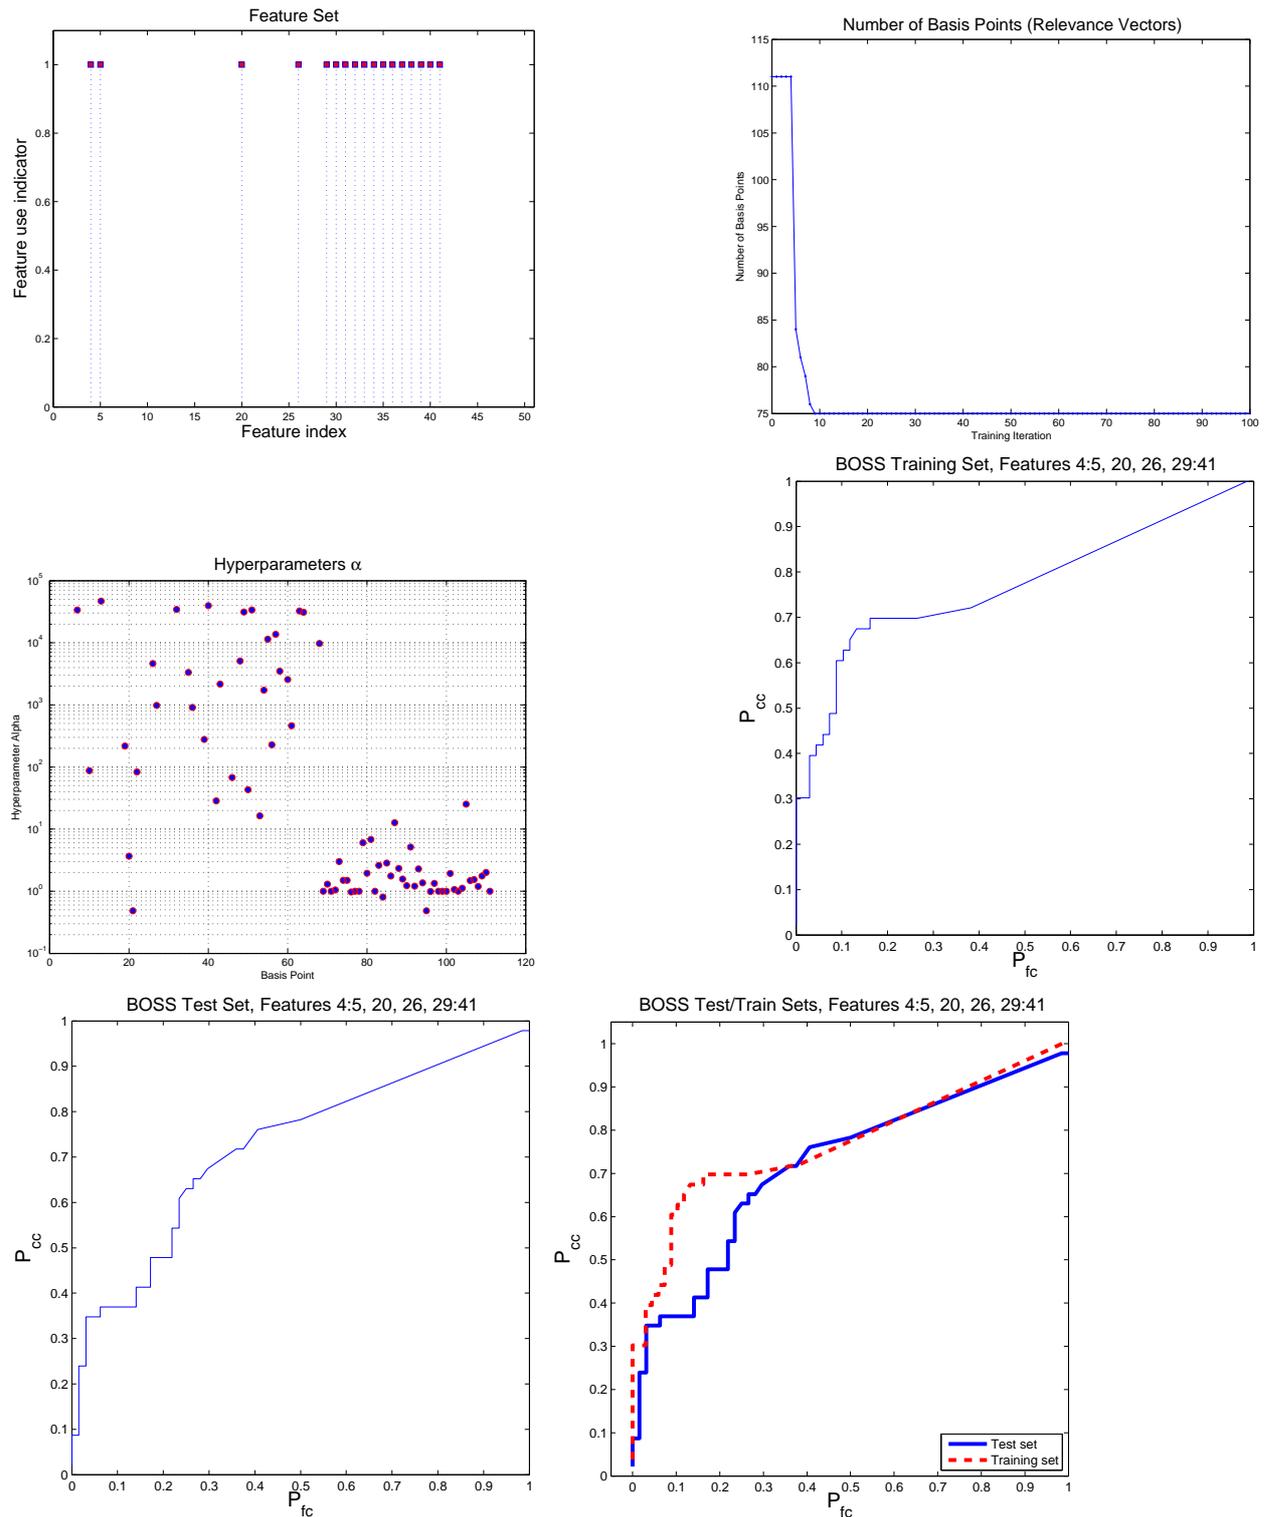


Figure 38: RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 5 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets.

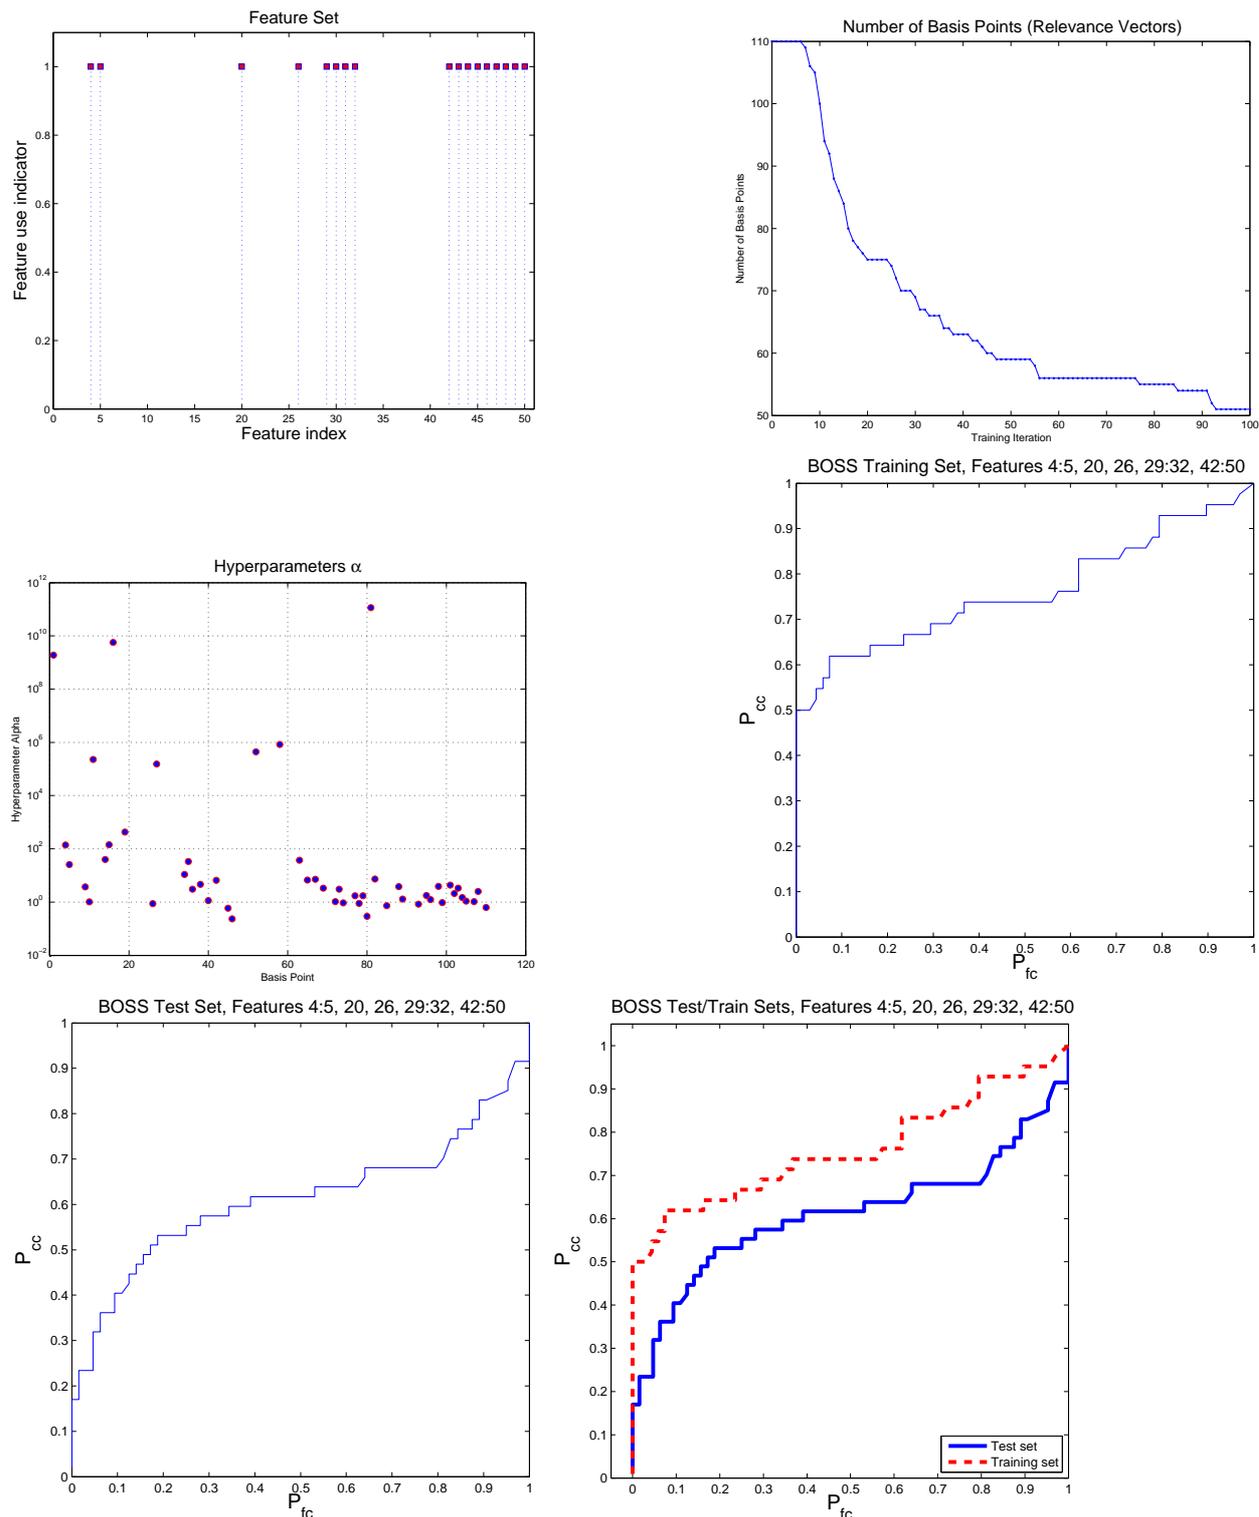


Figure 39: RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 6 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets.

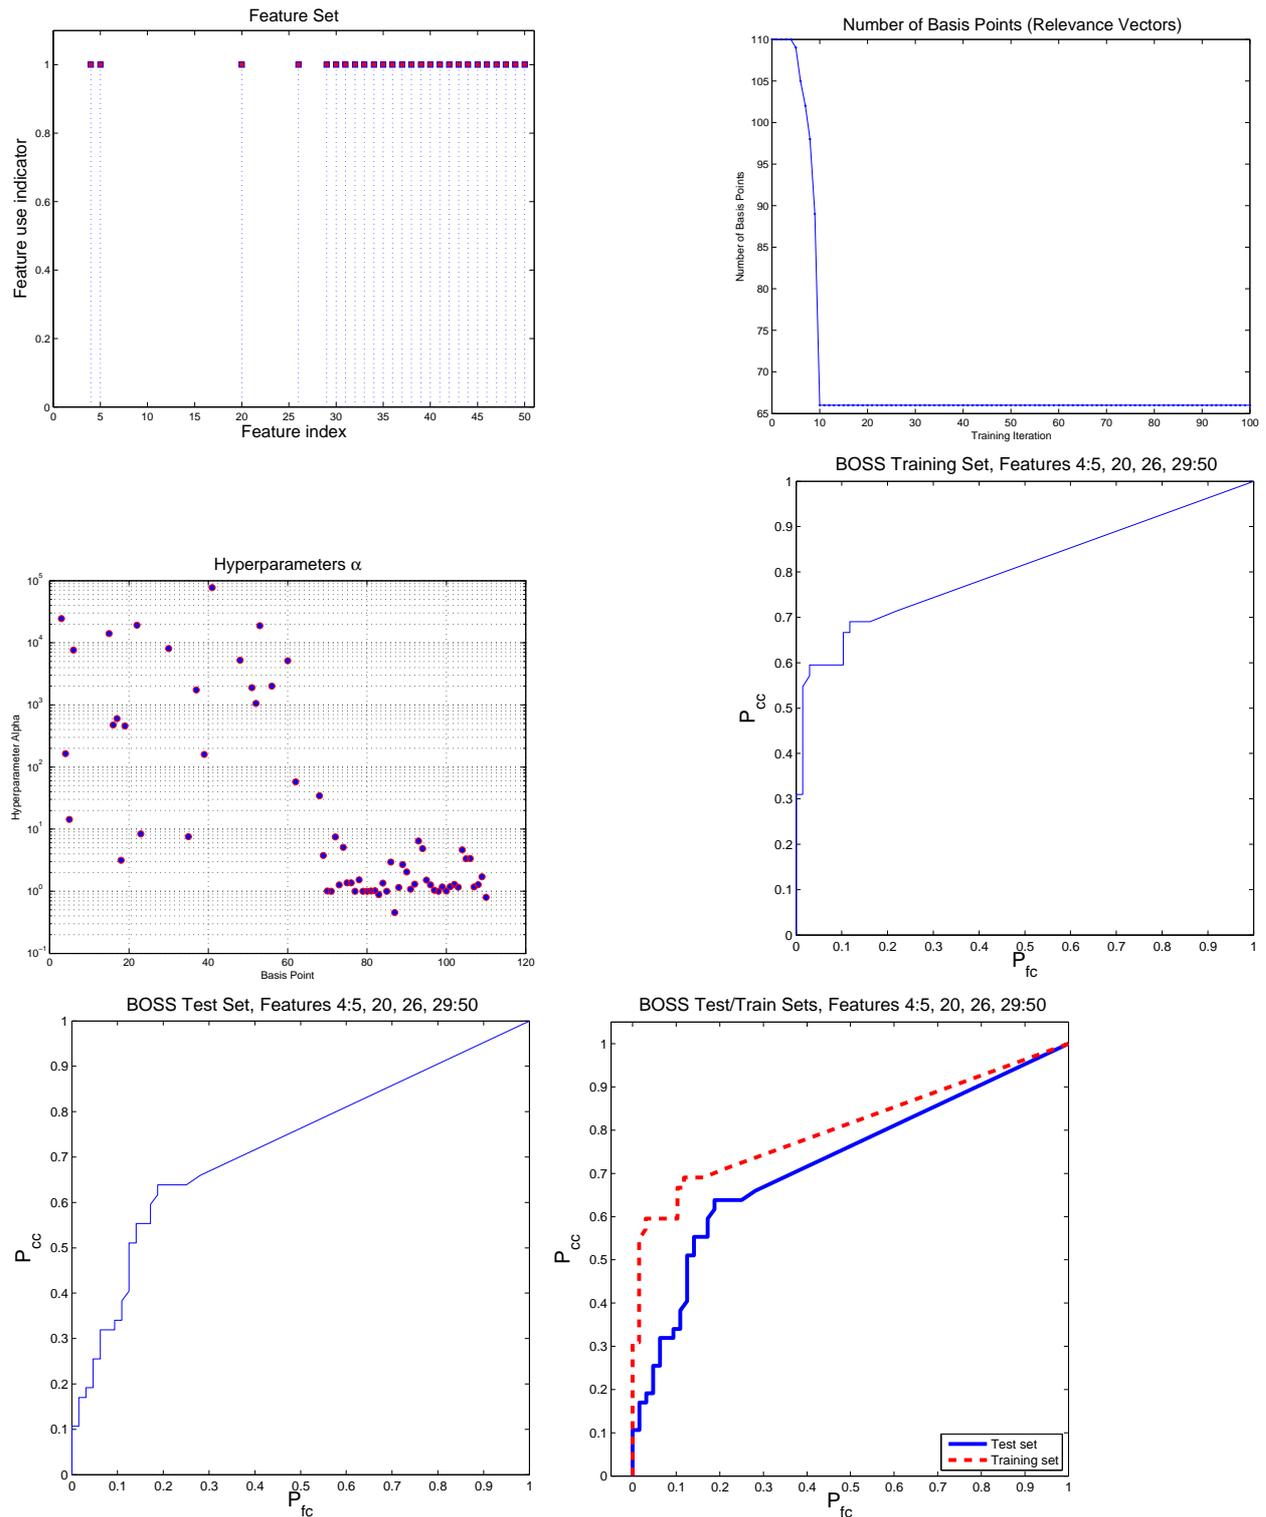


Figure 40: RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 7 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets.

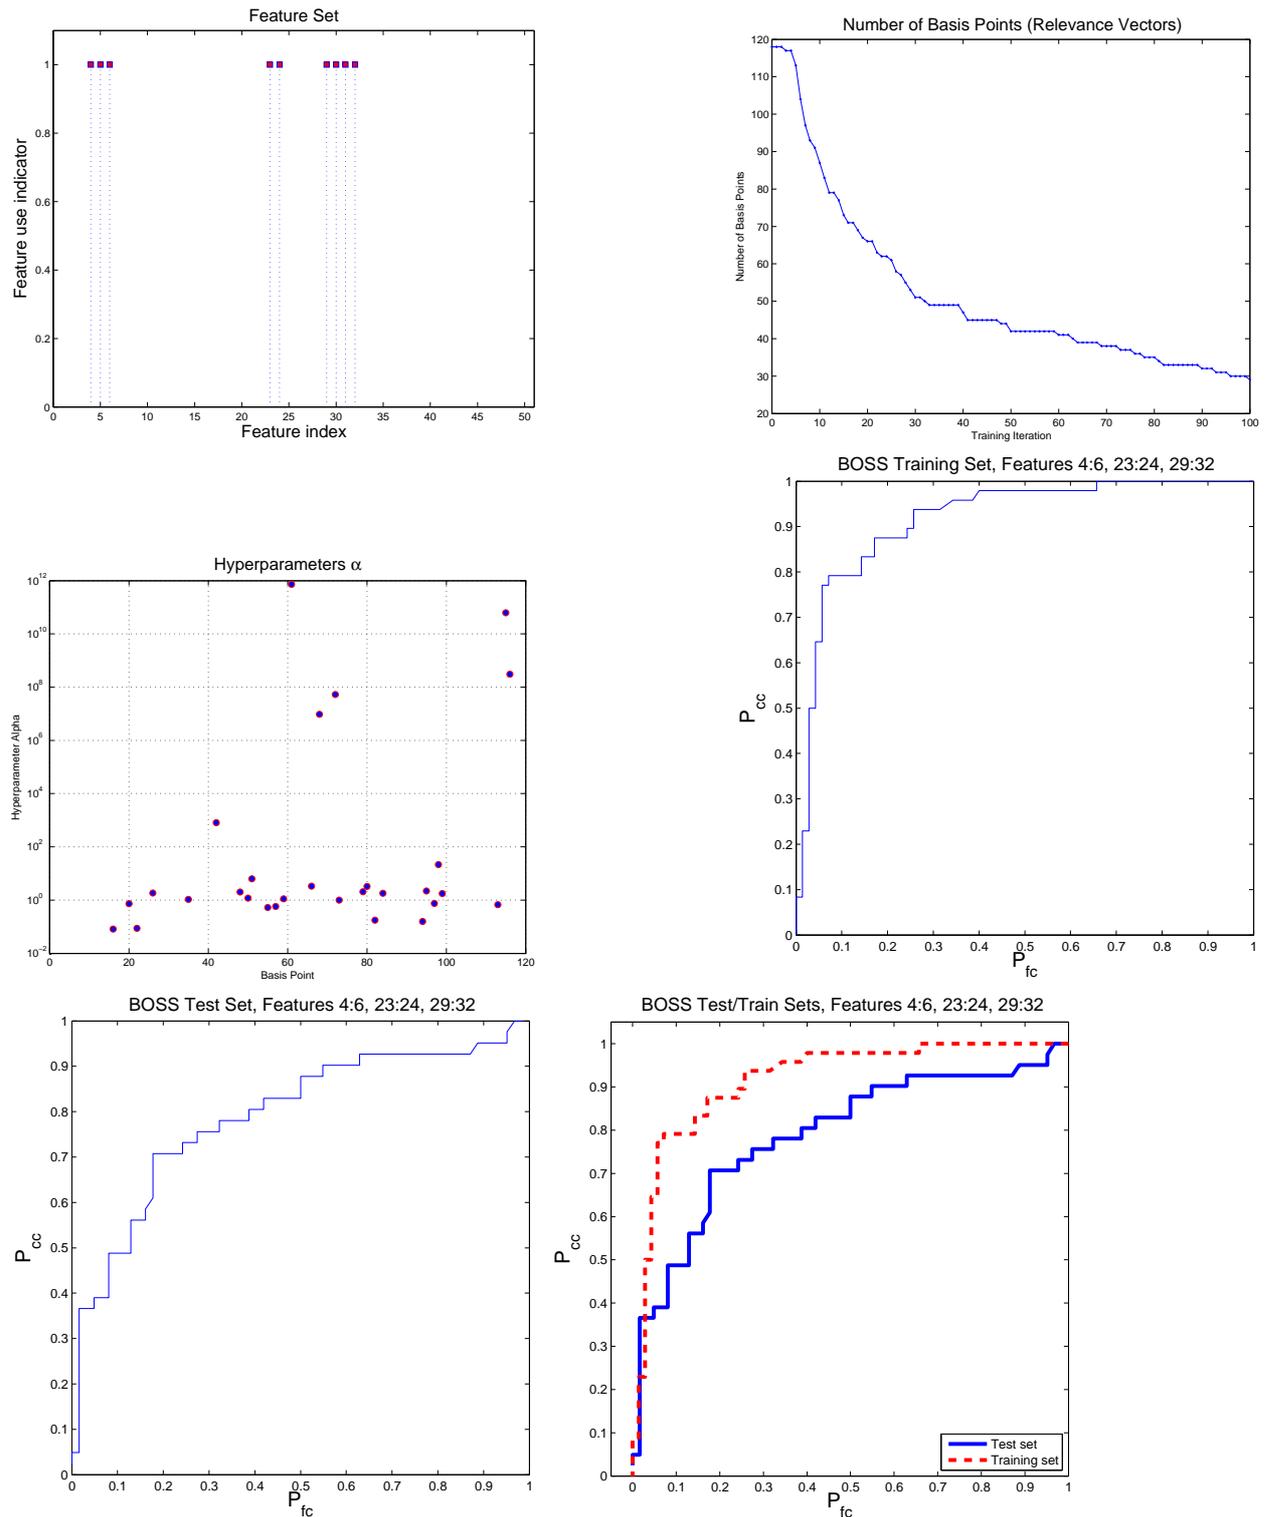


Figure 41: RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 8 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets.

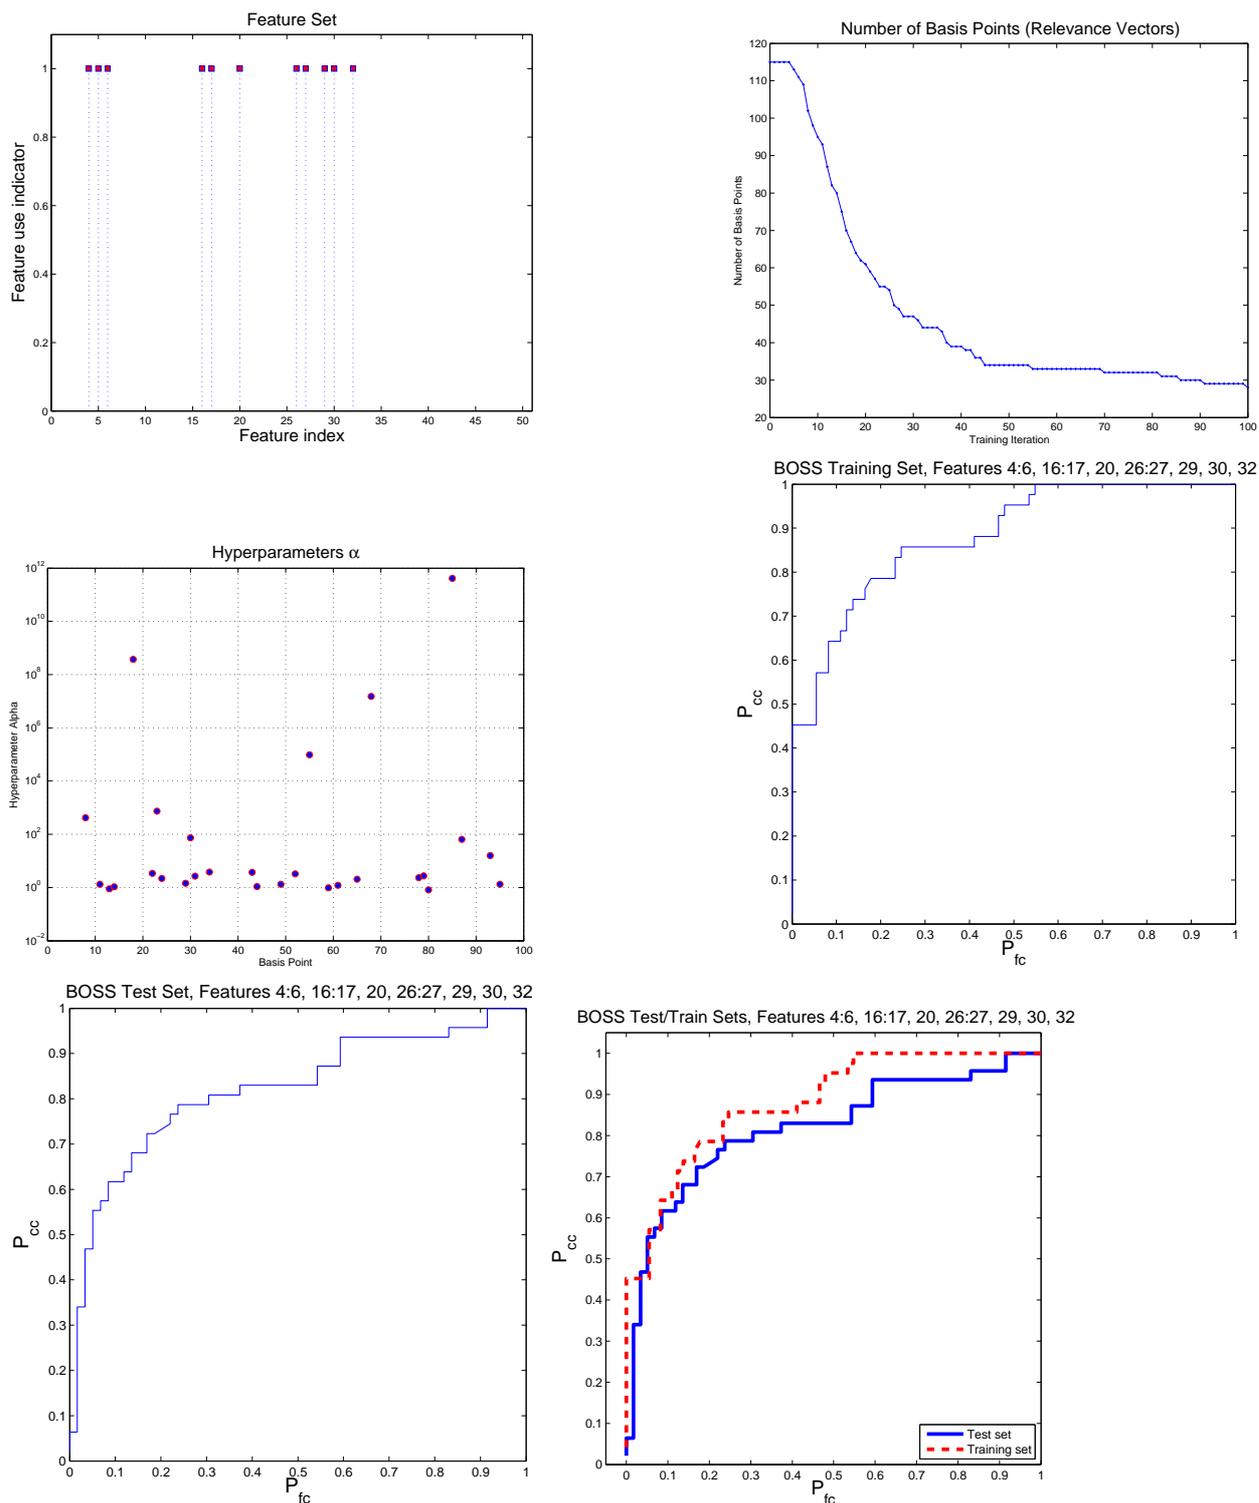


Figure 42: RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 9 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets.

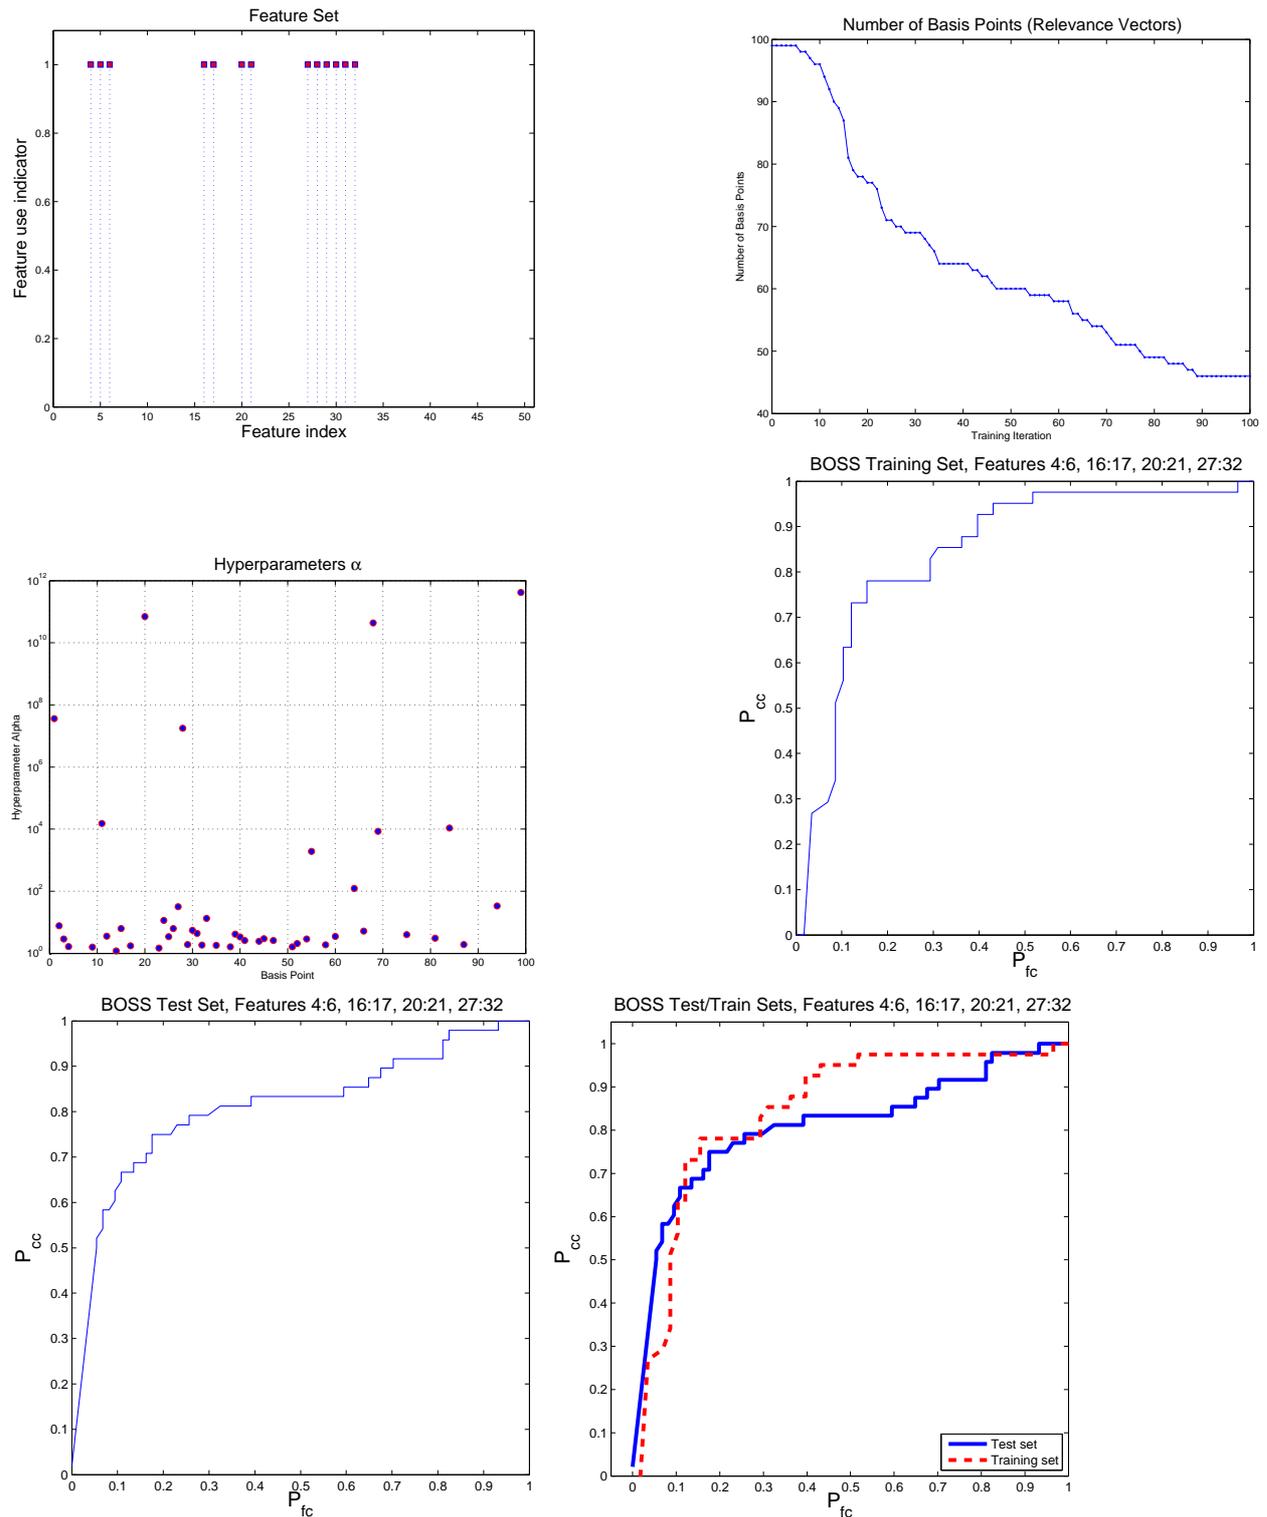


Figure 43: RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 10 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets.

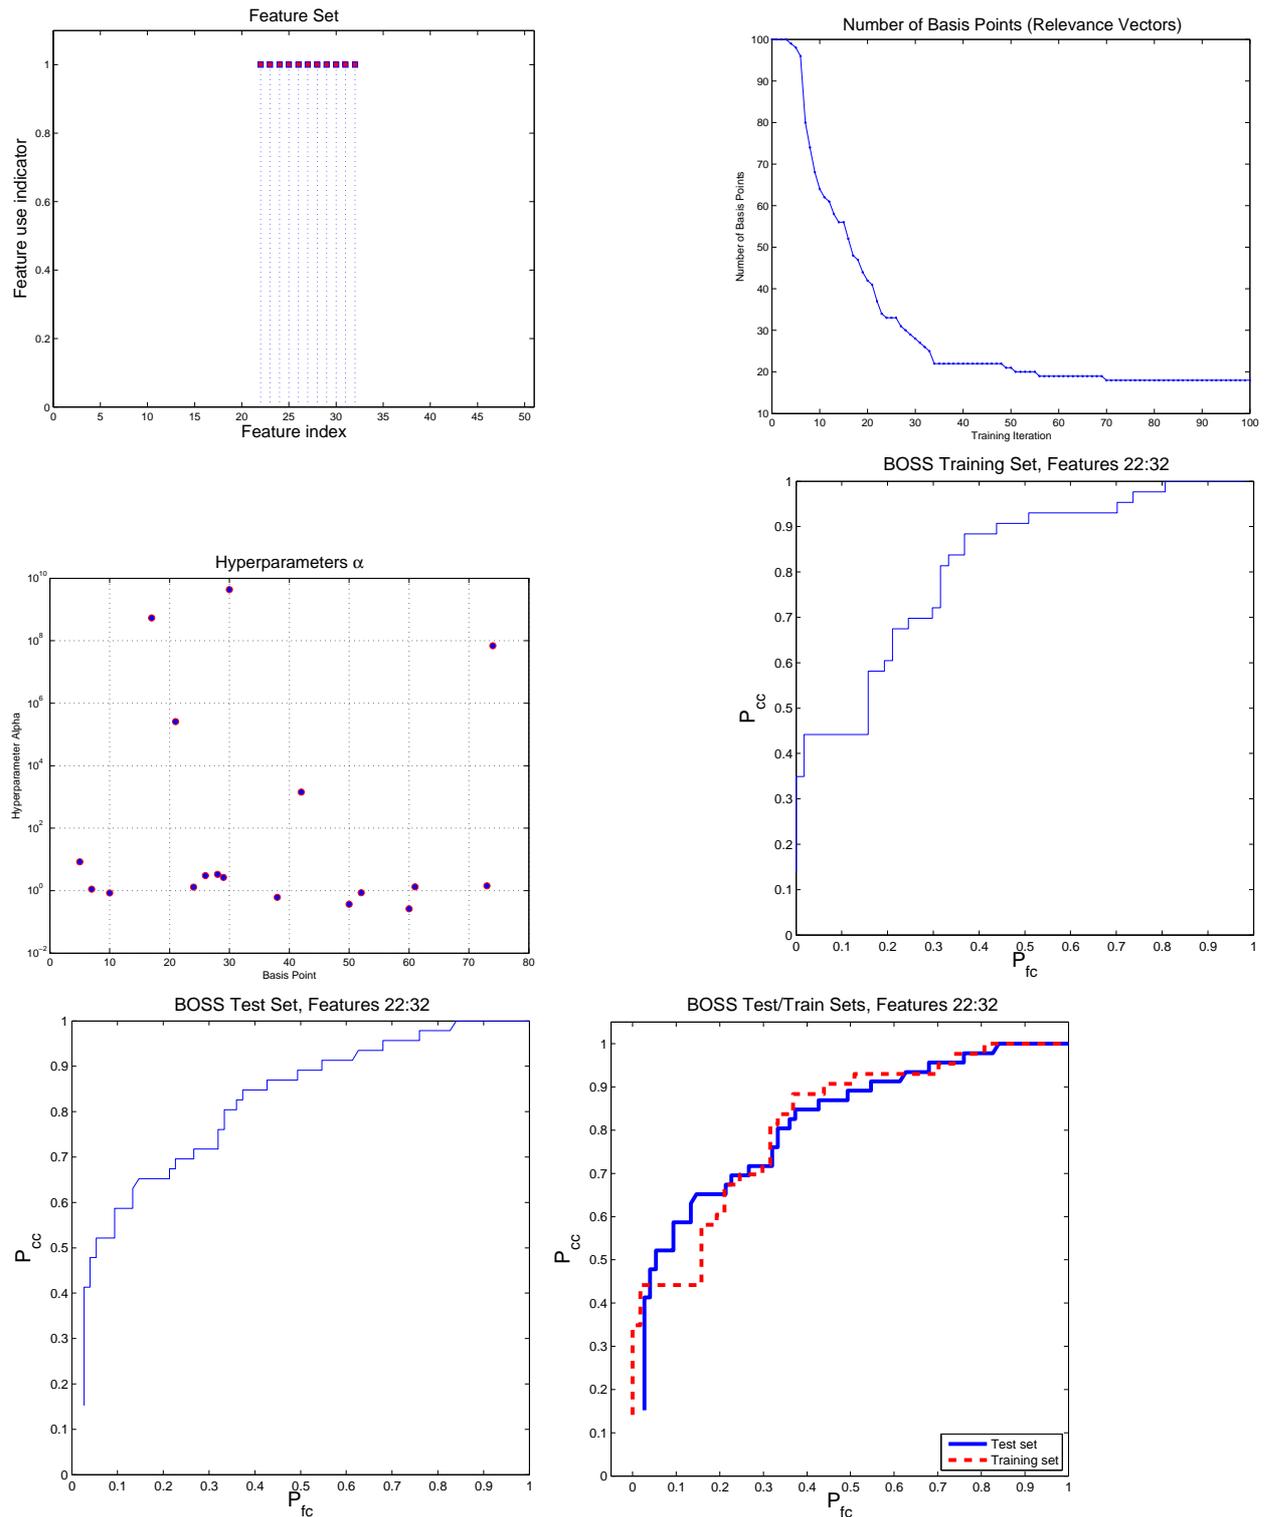


Figure 44: RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 11 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets.

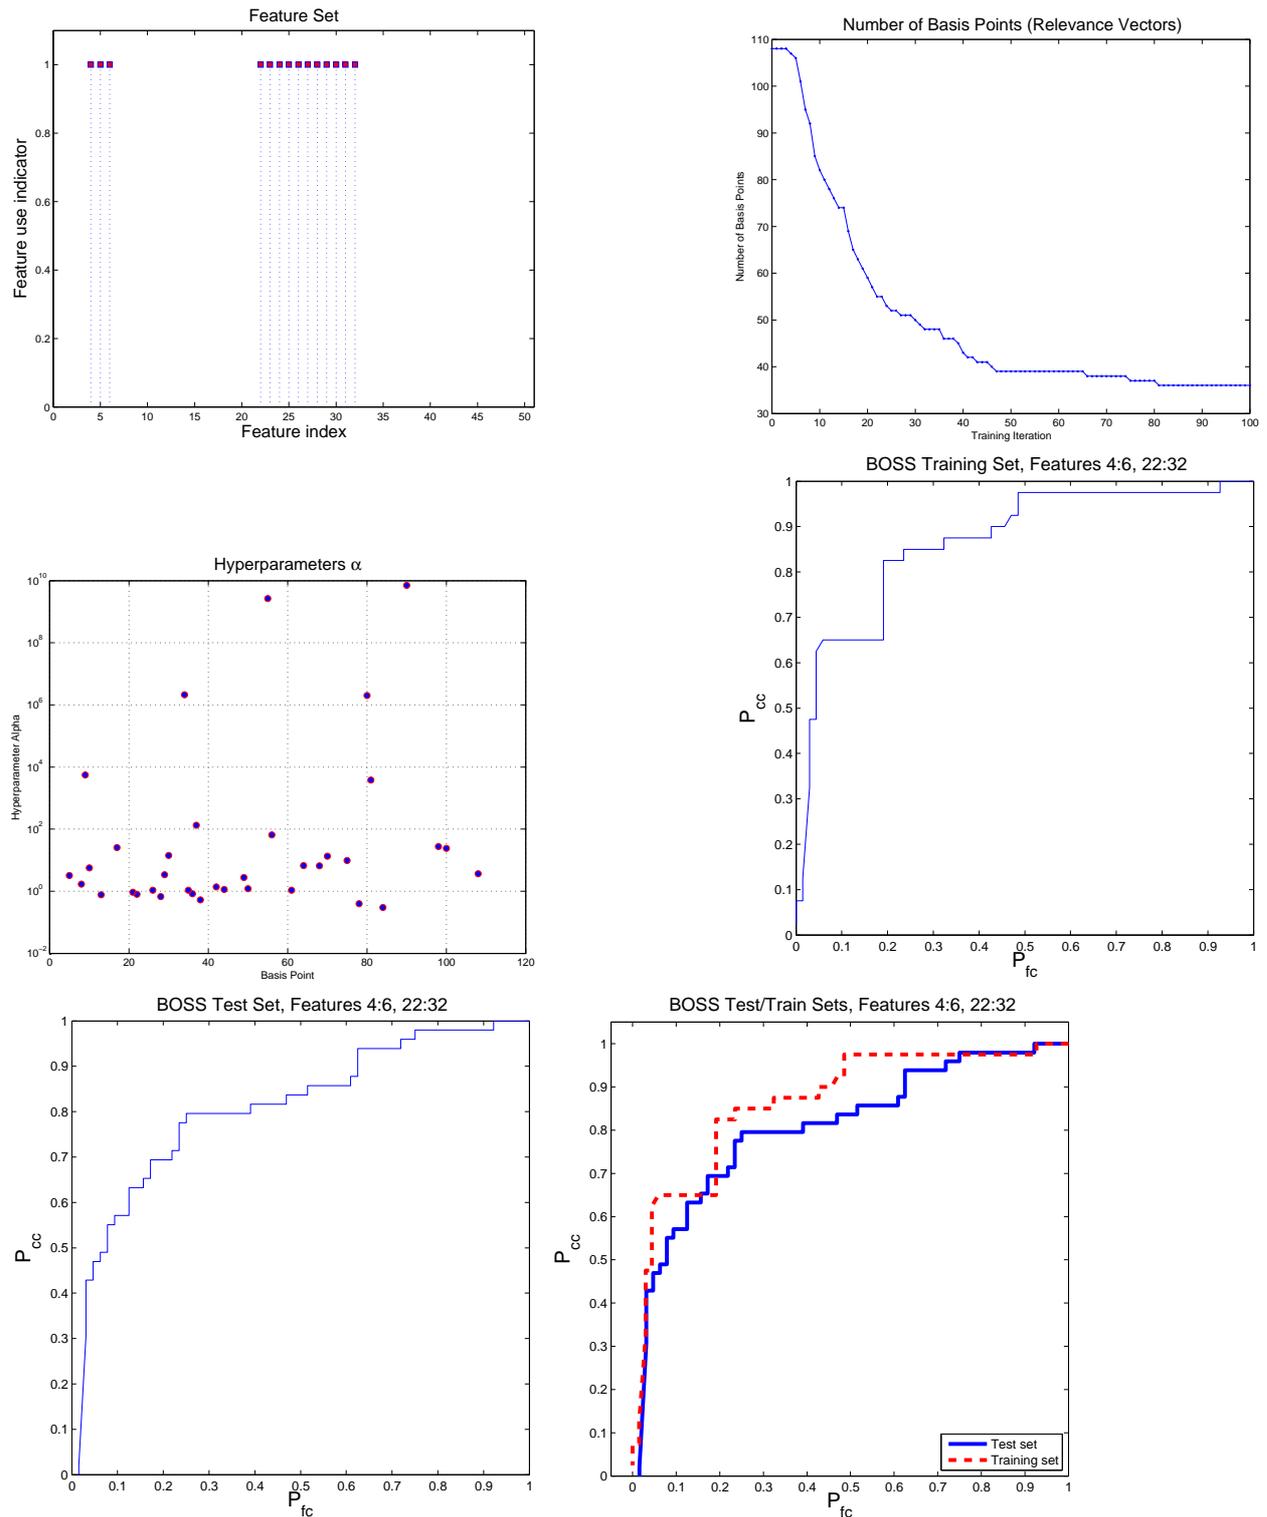


Figure 45: RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 12 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets.

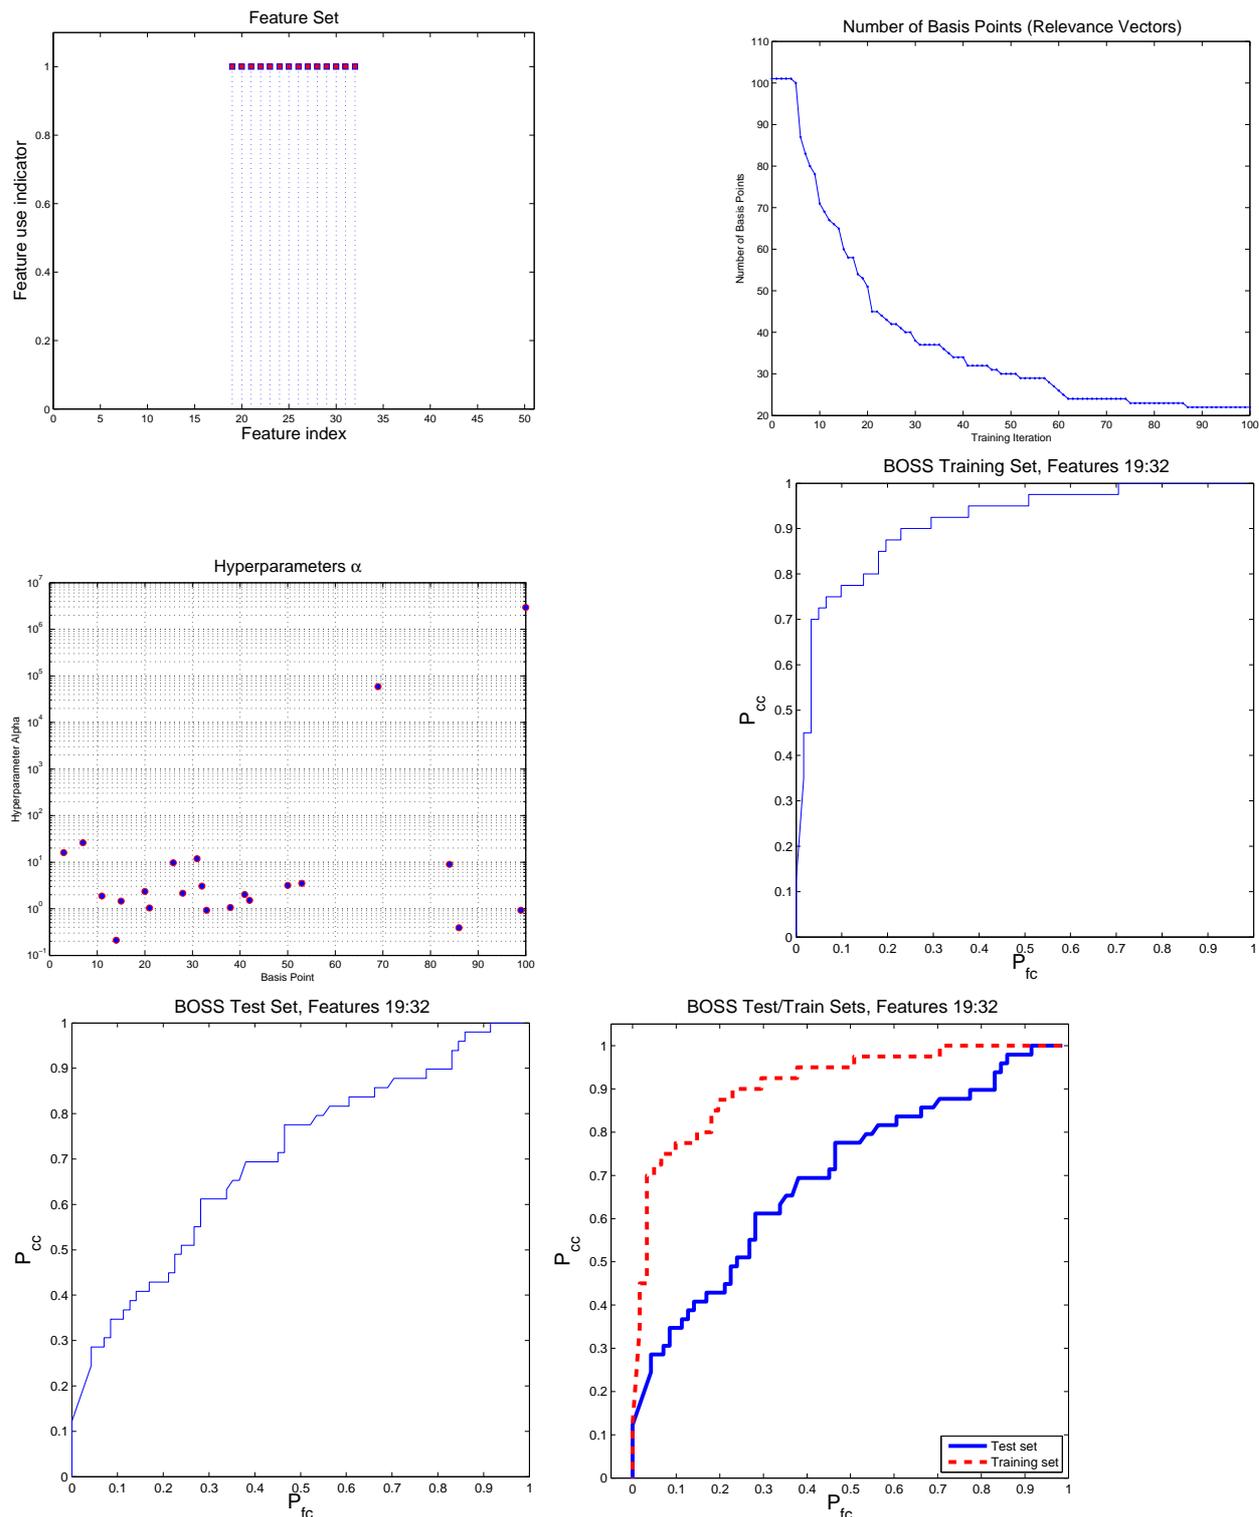


Figure 46: RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 13 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets.

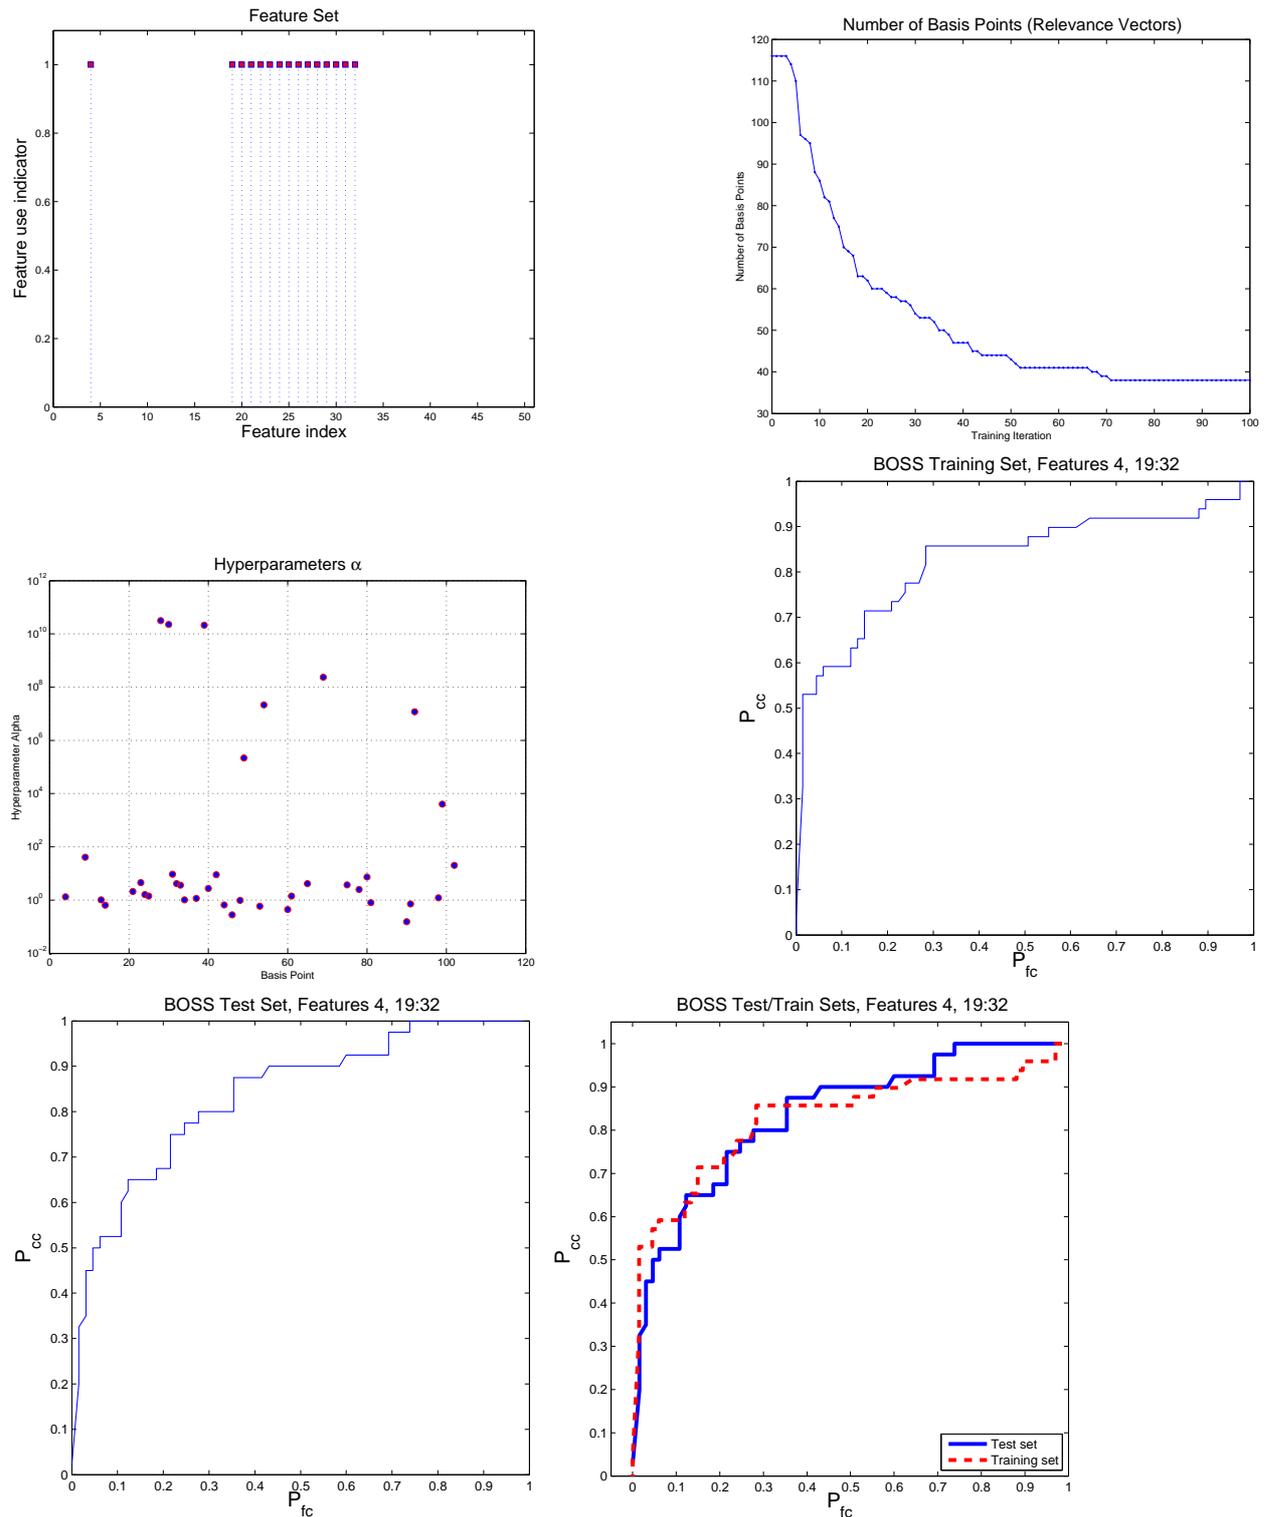


Figure 47: RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 14 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets.

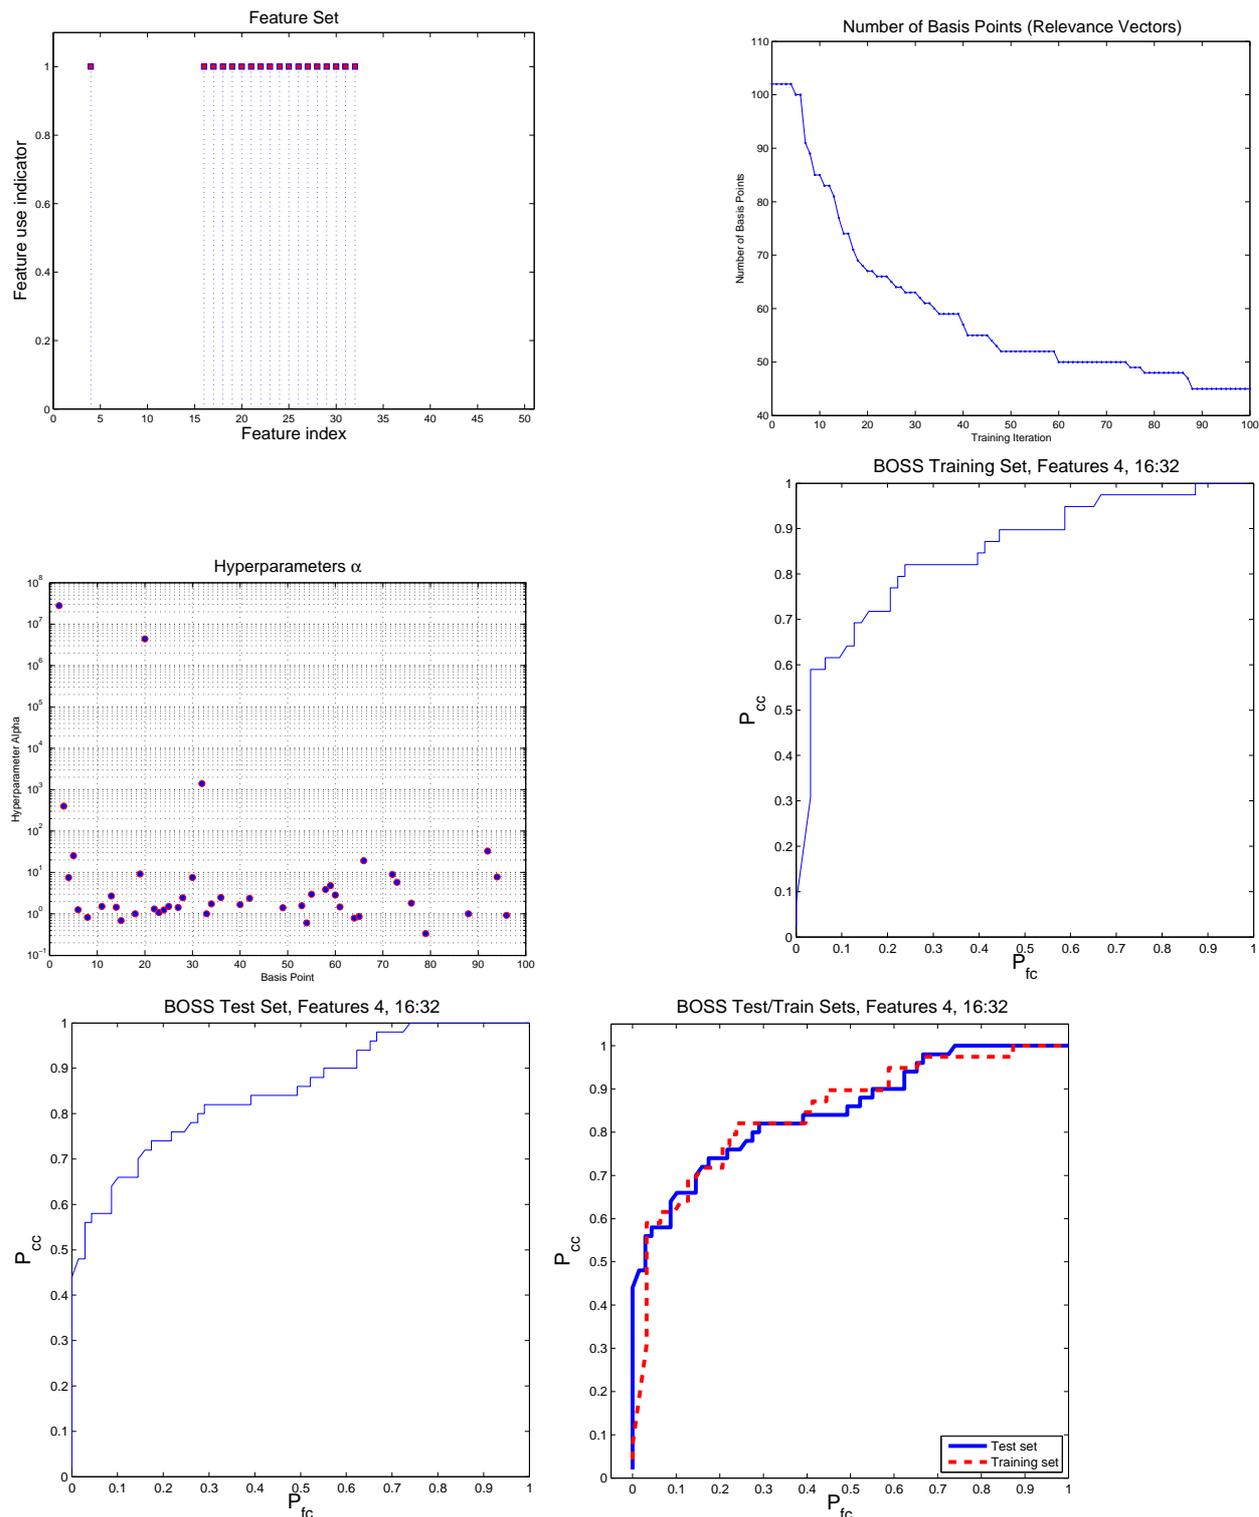


Figure 48: RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 15 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets.

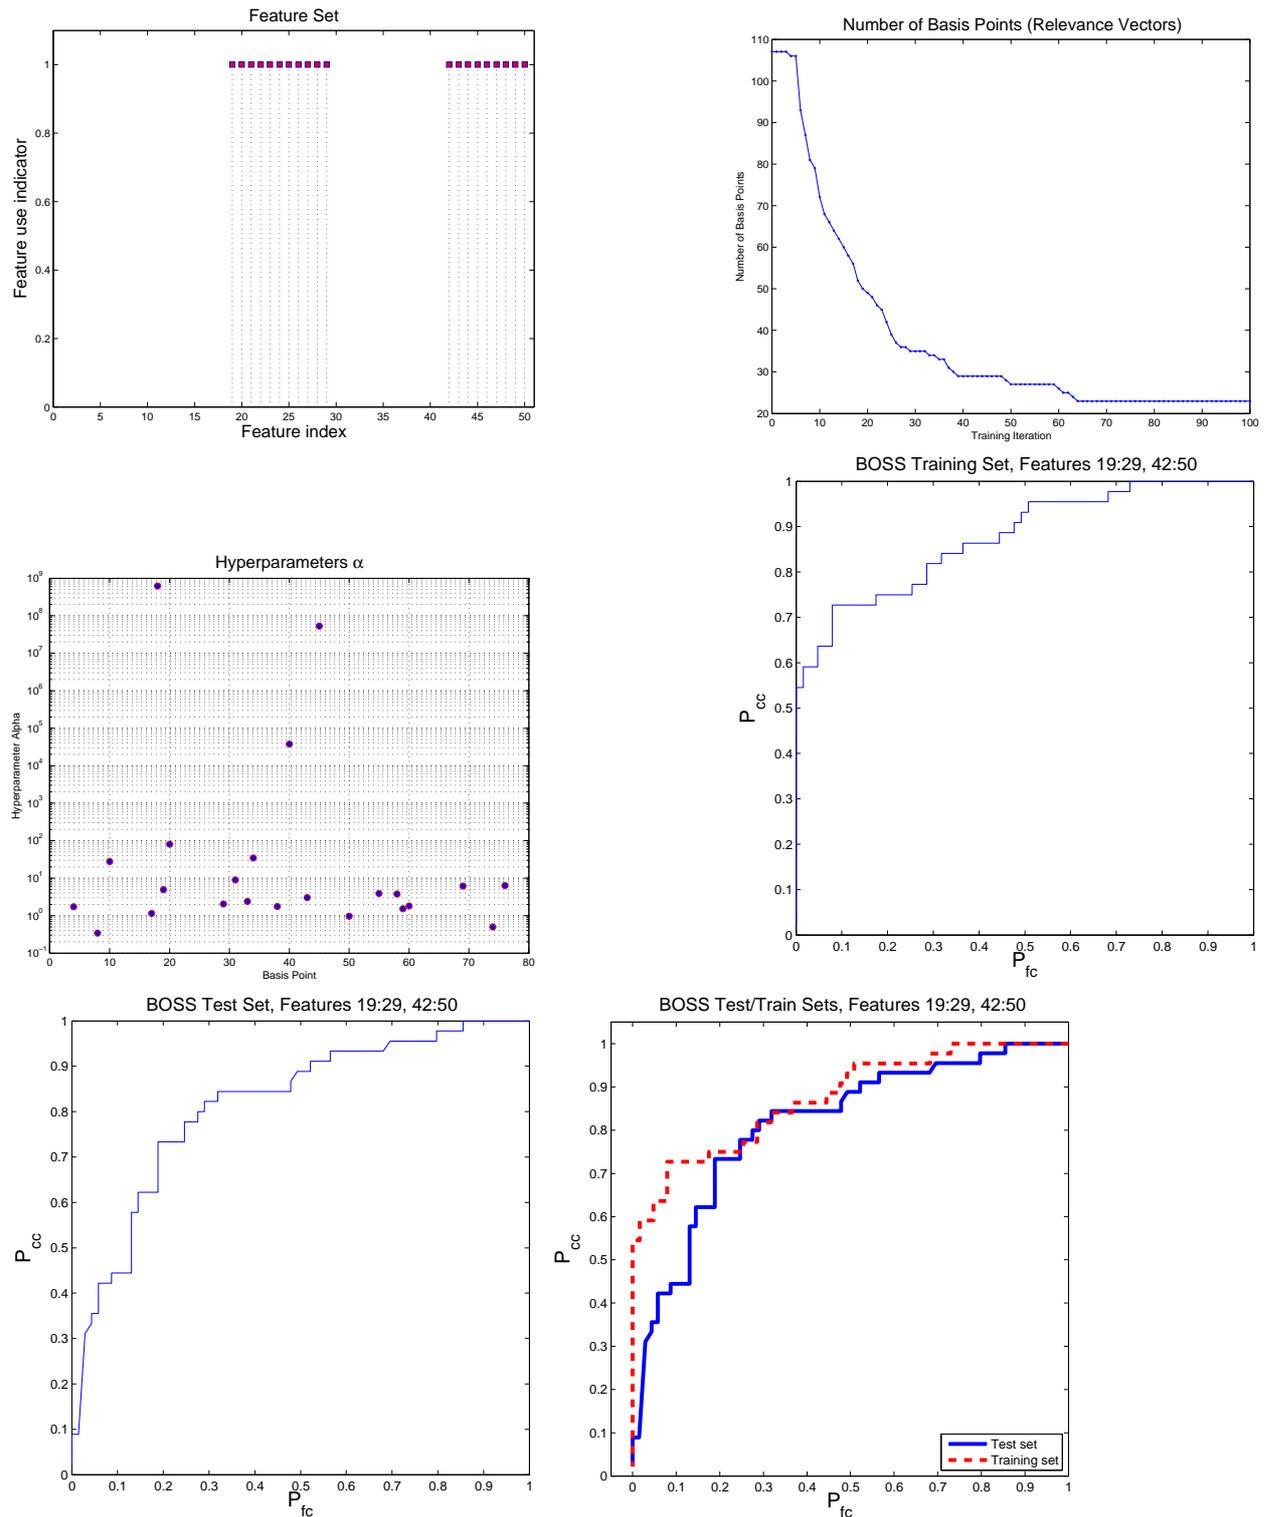


Figure 49: RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 16 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets.

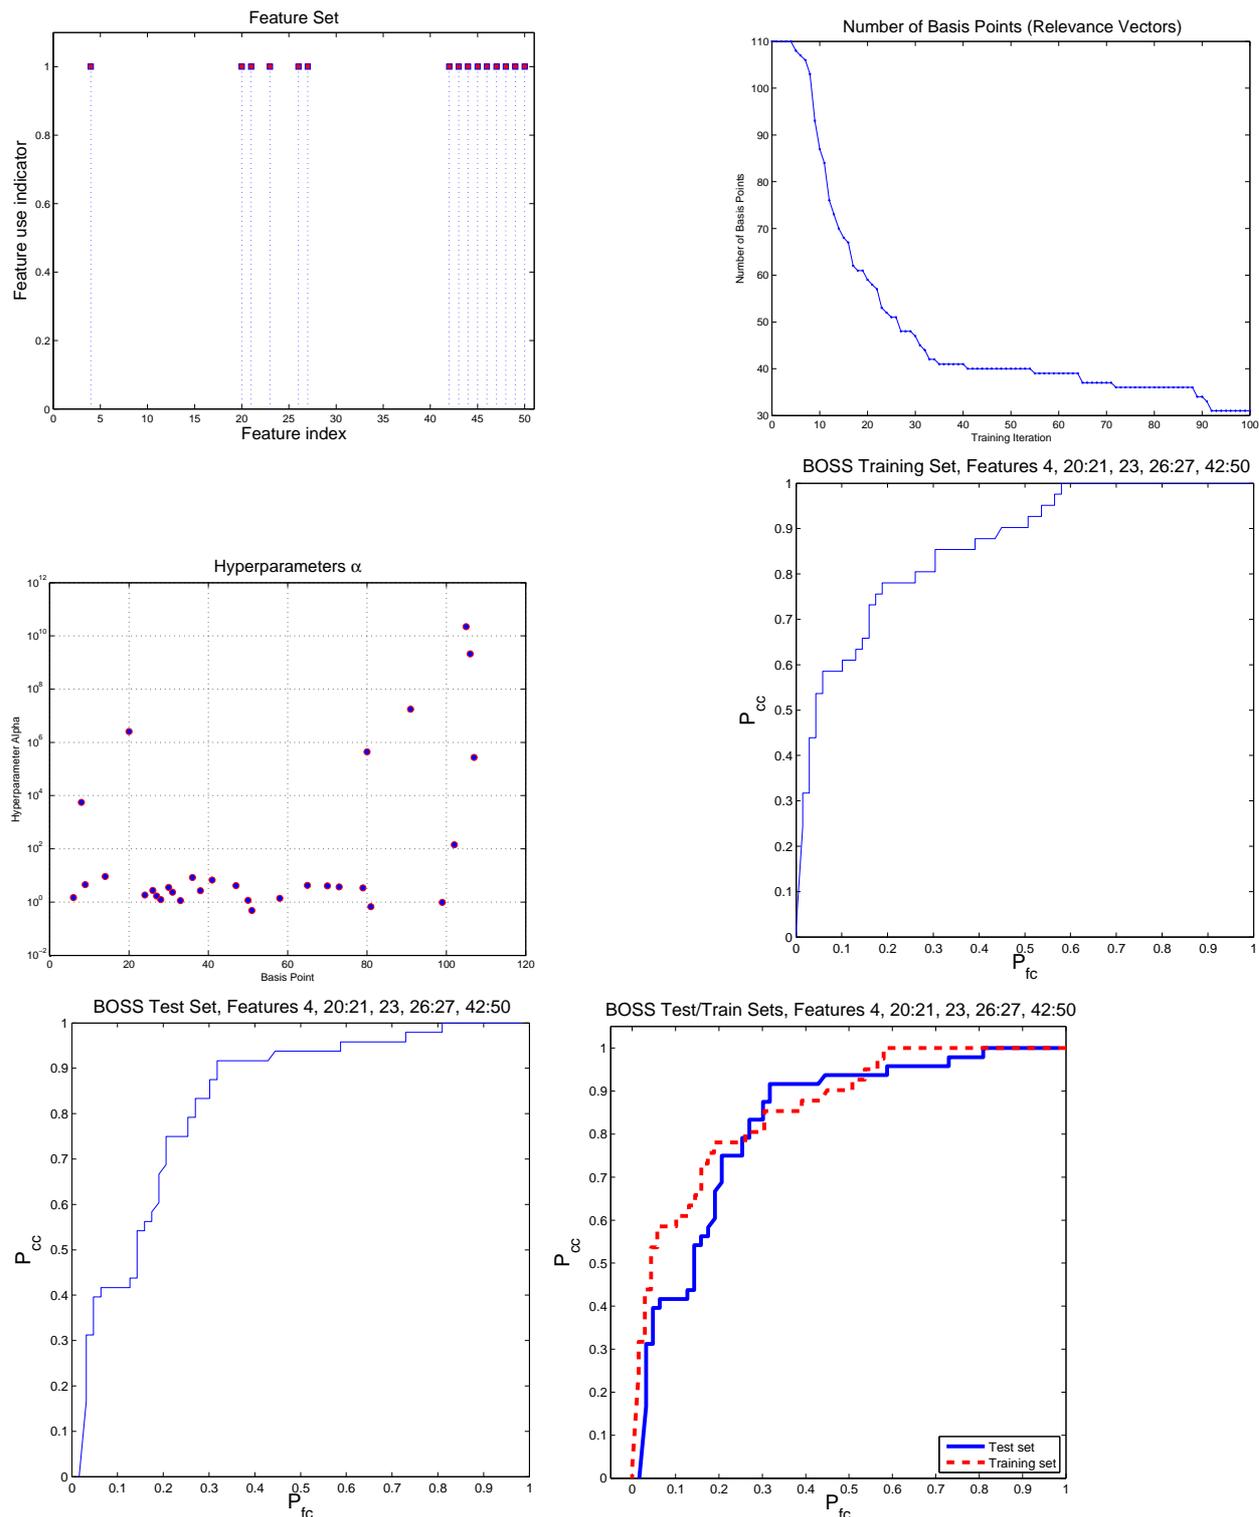


Figure 50: RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 17 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets.

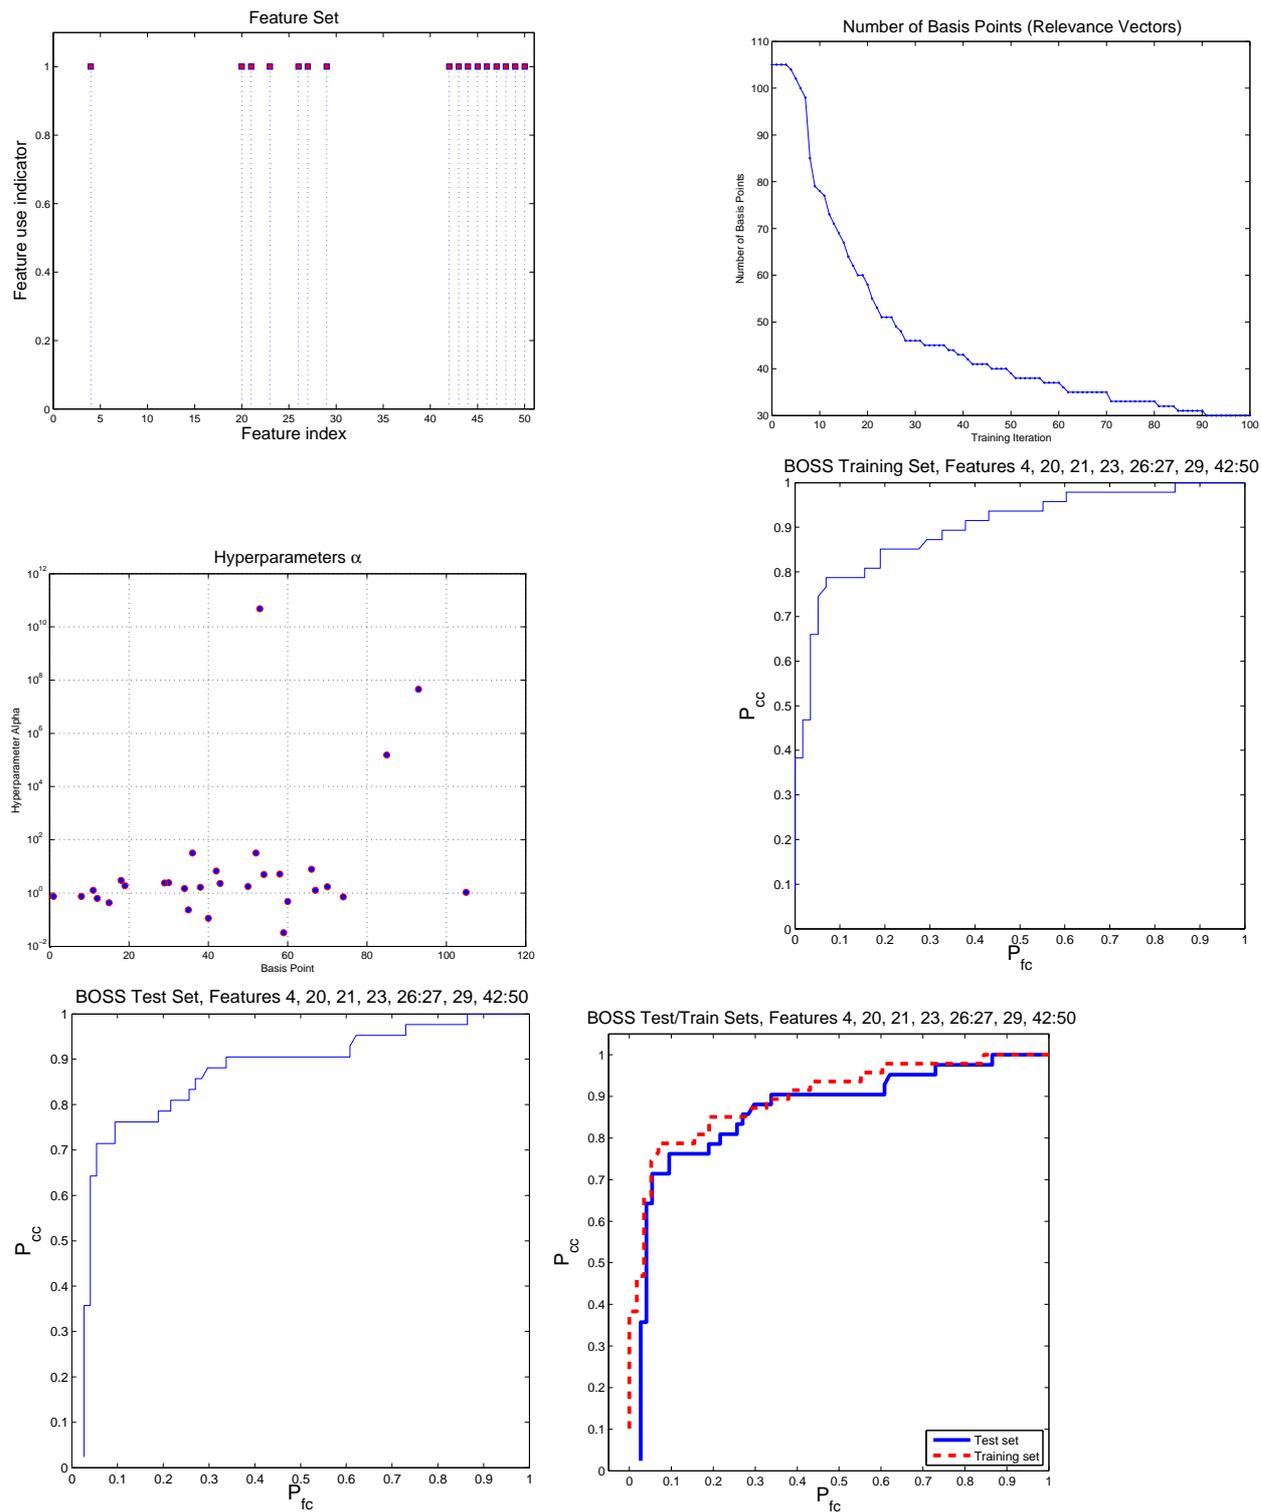


Figure 51: RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 18 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets.

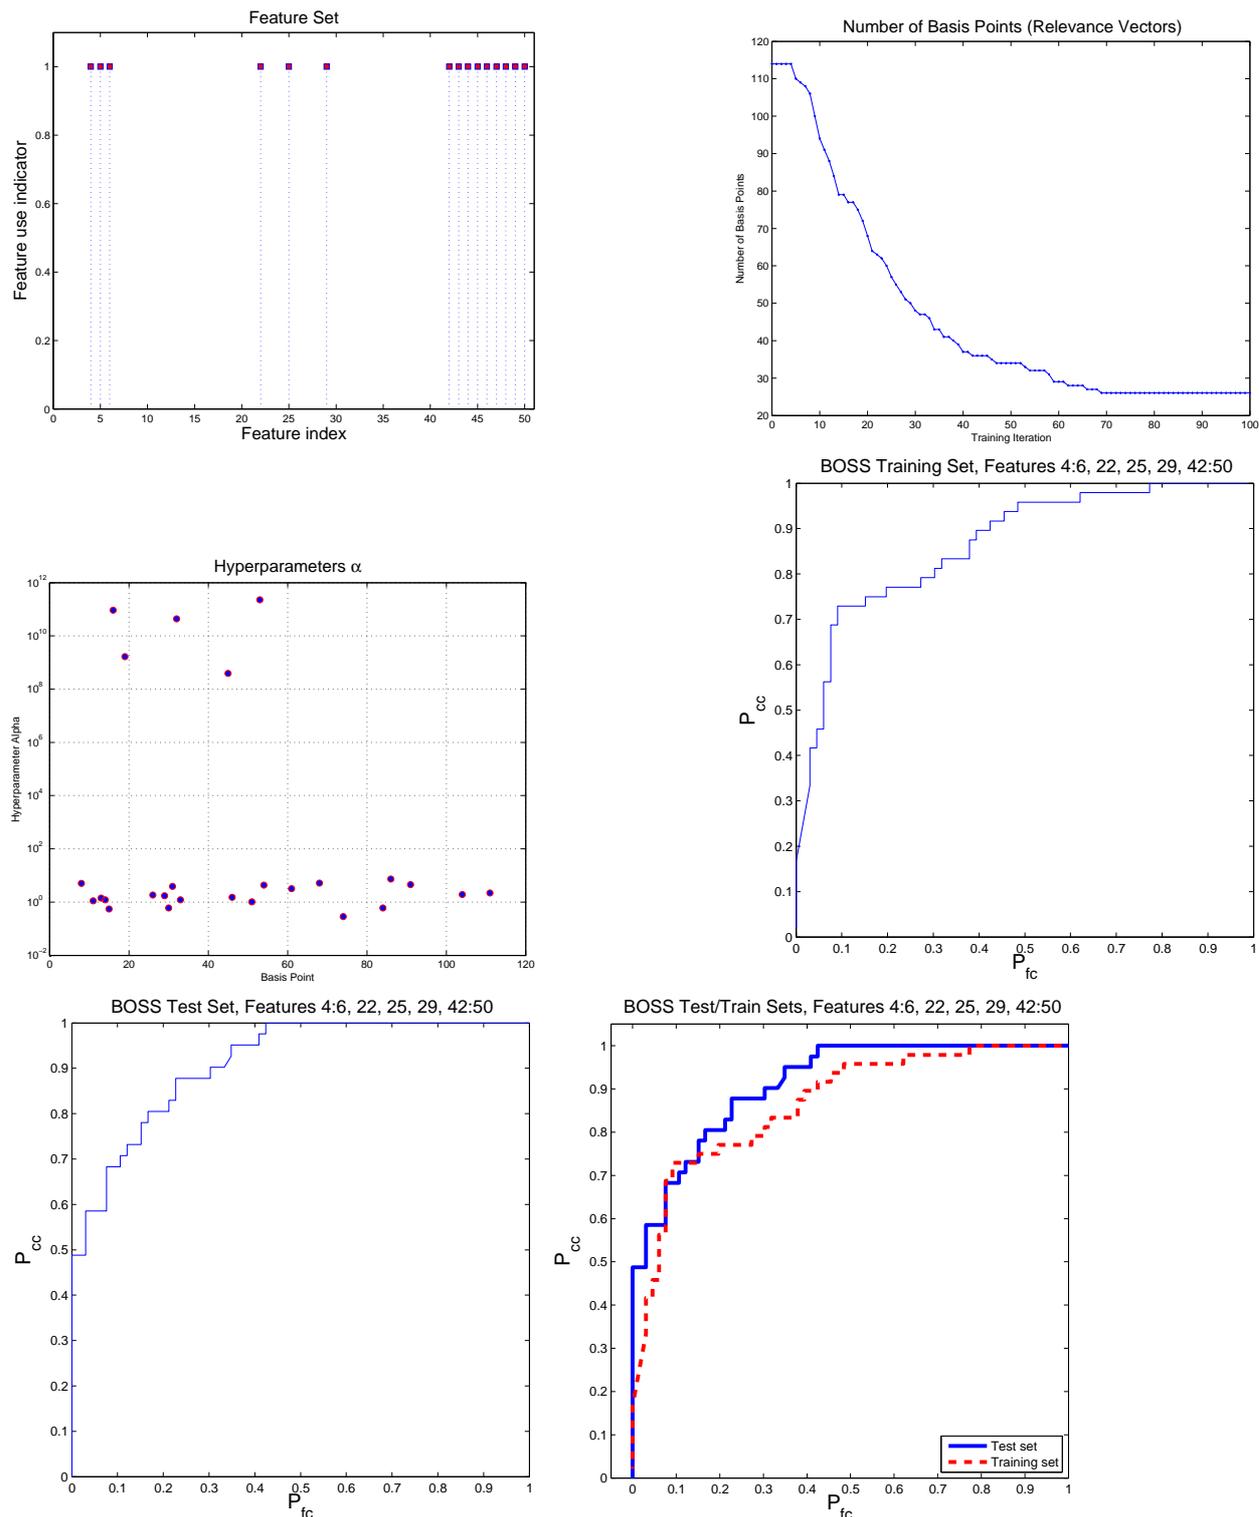


Figure 52: RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 19 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets.

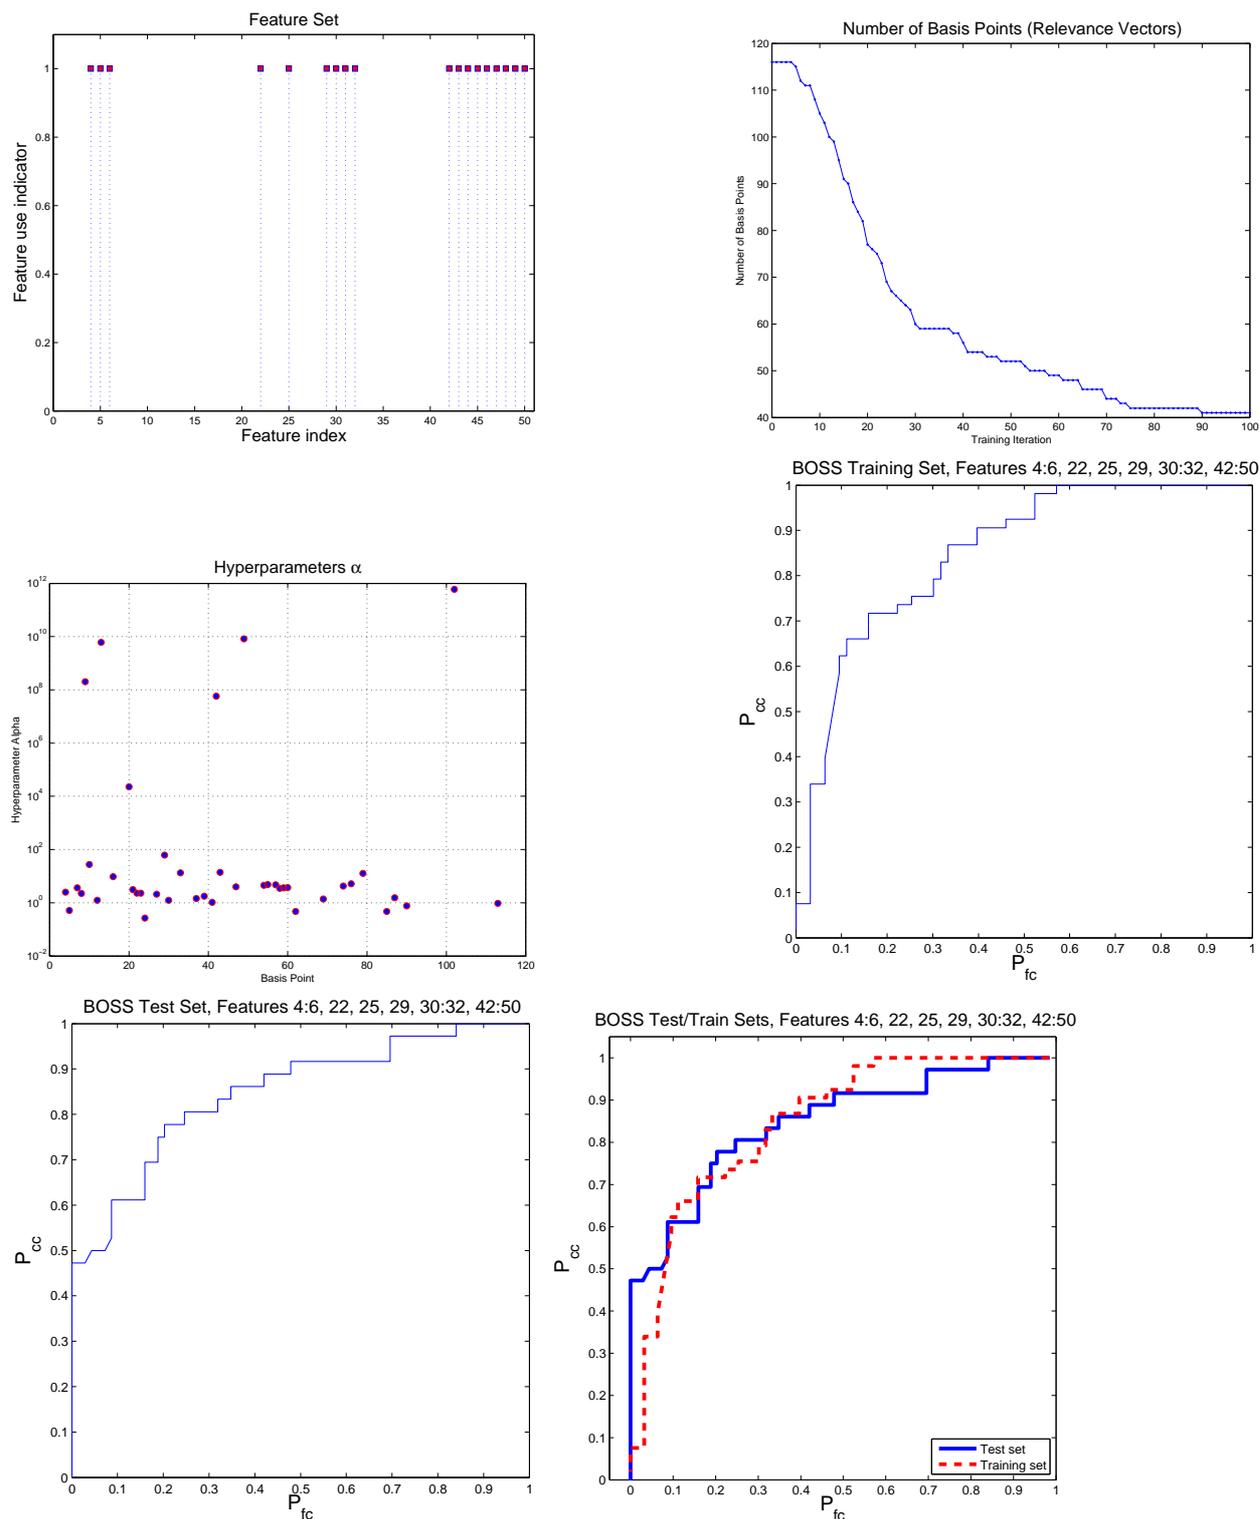


Figure 53: RVM classification results for AUV Fest 2008 BOSS data. **Upper left:** The RVM was trained using the indicated subset of the 50 available features (case 20 in Table 8, see also Table 7). **Upper right:** Number of *relevance vectors* chosen from available training vectors per RVM iteration. **Middle left:** Significant hyperparameters α_i from the final RVM iteration indicating indices of *relevance vectors* chosen from training vector set. **Middle right:** ROC for the training set (half of target/clutter exemplars). **Lower left:** ROC for the test set (training set complement). **Lower right:** ROCs for training *and* test sets.

5 Conclusions and Implications for Future Research and Implementation

Our findings, show that buried targets can be detected by low-frequency sonar, and that there is some discrimination capability provided by use of features derived from 3D beamformed imagery of such data. However, for economical remediation of underwater munitions, our results show that their needs to be continued detailed survey sonar-based improvements for buried targets. This is primarily driven by the facts that (1) targets are hard to detect since sediments attenuate high-frequency waves, and (2) they are hard to discriminate since high-frequency features e.g., shadows are lost. This implies that target imagery responses are easily confused with those of clutter and munitions debris due to incomplete separation of target and clutter features. It will be necessary to continue to improve phenomenology understanding of munitions and site conditions. Object responses are strongly affected by factors extrinsic to the object e.g., sediment properties, object orientation and proximity to sediment interfaces. The compensation of target signatures for the influence of munitions state and site conditions is one approach to enhance classification value of existing and new features.

For future research efforts, we suggest that improved classification can be derived by augmenting the image-based features used in this study with additional feature classes. In particular, following the findings from the structural acoustics community, and the T-matrix results presented here, we recommend the combined use of image-based and resonance-based features. This would require a method to extract resonance features from the data, and the development of a joint classifier that uses both the image- and physics-based features. In addition, methods to calibrate out the effect of the environment need to be developed. A true model-based effort that exploits the resonance features can benefit from the precise characterization of such features for real targets and targets expected in field conditions. A combined approach utilizing carefully controlled laboratory data for such information with application to BOSS or other low-frequency sonar data should provide a step forward in the discrimination problem.

From our findings, and from a survey of other work performed in the underwater munitions area, we have identified the following technology gaps:

- There is a need for adaptive environment calibration for improved image formation (correction for time delay and amplitude effects) and for model-based resonant characteristic estimation (which as shown in the T-matrix development presented here, depends on the 1D background model). Such an approach should be developed to exploit the unique BOSS source and receiver aperture.
- There is a need for a method to invert for intrinsic resonant fingerprint constructively for non-spherical targets using all available data.
- There is a need for a classifier combining 3D image and physics features (use model-based and in-field data training methods).
- There is a lack of high-fidelity models of real-target scattering (especially for high aspect ratio targets), and usable for adaptive, in-field estimation. Usability depends on both fidelity and computational feasibility.

- Effective classifier design will rely on uncertainty analysis for resonance features, and detectability of such features in at-sea conditions (especially aspherical targets).

The features we used for the classification results presented in Sec. 3.5 were described in Sec. 4.5 and summarized in Table 7. The discovery of the optimal subsets of feature types for use in classification analysis will require additional experimentation, and perhaps the application of an optimization wrapper around the core classification algorithm. Results for various manually selected feature subsets were shown in Sec. 3.5. Additional feature classes can also be derived to improve classification metrics. We speculate that features of interest include T-matrix coefficients estimated from the data (see e.g., Sec. 3.2.11) and volumetric estimates of material parameters from adjoint tomography (see e.g., Sec. 3.4).

In summary, the key conclusions and results from the present effort include the following:

1. Buried targets are detectable in the 2D and 3D image products from the BOSS processing results. This means that despite attenuation in the sandy sediments, there is a strong enough scattered return for the frequency range of interest. This finding suggests that it is worthwhile to further characterize and optimize the use of SAS for target detection and classification.
2. We applied the processing sequence described in Sec. 4.3 to generate 2D and 3D image products for BOSS data collected in AUV FESTs 2006, 2007 and 2008. Examples of the resulting image products are shown in App. A.7.
3. Classification analysis was applied to AUV FEST 2008 data since this collection had the most information concerning target/clutter labels for known objects in the test field. The classification results demonstrated the utility of image-based features for target/clutter discrimination, but there is still room for improvement.
4. The Spectral Element Method was developed for target scattering simulation in poroelastic media. Much effort was committed to developing this technique. The publications [27, 29, 4] resulting from the effort provide a methodological advance to the community, and the corresponding code has been made publically available. It is expected that interesting applications of the code will be found in the future.
5. A feature database from 3D imagery was developed, and other researchers can use this database for testing new classification algorithms. An interesting feature class is tomographic estimates of the material parameters in the sediments and objects embedded in the sediments. Adjoint tomography was proposed for this purpose, and was described in Sec. 3.4.
6. In terms of forward simulation, it is necessary to capture all aspects of the system in addition to the specific target scattering phenomenology. The former consists of source properties (location, Chirp waveform, strength, etc.), receiver specification (e.g., geometry of receiver array, individual receiver response, data acquisition protocol (e.g., kinematics of the sonar platform, pulse repetition frequency) and environment specification (e.g., water column depth and sediment properties). These were implemented in both a T-matrix framework and in the SEM implementation.

With respect to the Go/No-Go decision point established by the Scientific Advisory Board for this effort (the ability to extract features for proud and buried targets), this was amply demonstrated in the classification results for the AUV FEST 2008 data, and in the image processing results shown in App. A.7. Numerous targets are visible in the swaths (Figs. 155 - 165) while Figs. 166 - 172 show 2D xy-slices for a single target. From these figures it is very clear that various features can be extracted such as target area regions, intensity values, etc.

Key objectives in this study were to (1) evaluate low frequency sonar for buried target detection; acquire and process a target-rich data set in a relevant littoral environment to guide our phenomenology understanding; and to recommend high-performance classifier designs, (2) develop and an SEM simulation method for high fidelity modeling, (3) apply the SEM code to numerous stochastic realizations of the environment to thoroughly characterize the influence of the environment on target detectability and classification and (4) develop a database of sonar products for use in classifier design. We believe we have succeeded in executing objectives (1), (2) and (4). However, developing the SEM code and theory required more resources and time than we originally expected. Therefore, there were not enough program resources to carry out the Monte Carlo study, or to perform a detailed comparison of predictions to real data. However, the essentially exact correspondence of the code predictions to available analytical results for end-member models encourages us as to the essential correctness of the theory and its numerical implementation. There remains a great deal of interesting and worthwhile work to do regarding the exploitation of the SEM development for furthering the basic understanding of sonar phenomenology in the littorals. In implementing objectives (1) and (4), we were fortunate to gain access to BOSS data sets through Professor Steve Schock at Florida Atlantic University and our AST subcontractor.

References

- [1] R. Chatham D. Marx M. Nelson A. Putney, E. Chang and L.K. Warman. Synthetic aperture sonar - the modern method of underwater remote sensing. In *2001 IEEE Aerospace Conference Proceedings*, pages 1749–1756, Big Sky, Montana, 2001.
- [2] W. L. J. Fox R. Lim A. Tesei, A. Maguer and H. Schmidt. Measurements and modeling of acoustic scattering from partially and completely buried spherical shells. *J. Acoust. Soc. Am.*, 113:1817–1830, 2002.
- [3] C. M. Bishop. In *Pattern Recognition and Machine Learning*. Springer, Singapore, 2007.
- [4] Y. Luo C. Morency and J. Tromp. Finite-frequency kernels for wave propagation in porous media based upon adjoint methods. 2009.
- [5] D.W. Lemonds M. Ericksen D. D. Sternlicht, R.D. Dikeman and S.G. Schock. Rapid surveying of buried targets by hybrid seafloor image and biomimetic signal classification. In *SPIE: Detection and Remediation Technologies for Mines and Minelike Targets, Vol. 4394*, pages 1183–1189, 2001.
- [6] D.W. Lemonds R.D. Dikeman M. T. Korporaal D. D. Sternlicht, A.W. Thompson. Image and signal classification for a buried object scanning sonar. In *Proceedings MTS/IEEE Oceans 2001*, pages 485–490, 2002.
- [7] R.D. Dikeman M. Ericksen D. D. Sternlicht, D.W. Lemonds and S.G. Schock. Detection and classification of buried objects with an adaptive acoustic mine-hunting system. In *Proceedings MTS/IEEE Oceans 2001*, pages 212–219, 2002.
- [8] D.W. Lemonds M. T. Korporaal D.D. Sternlicht, R.D. Dikeman and A. M. Teranishi. Target confirmation architecture for a buried object scanning sonar. In *Proceedings MTS/IEEE Oceans 2003*, pages 512–520, 2003.
- [9] N. P. Chotiros J. T. Christoff K.W. Commander C. F. Greenlaw D. V. Holliday D. R. Jackson J. L. Lopez D. E. McGehee M. D. Richardson J. Piper E. I. Thorsos, K. L. Williams and D. Tang. An overview of sax99: Acoustic measurements. *IEEE J. Oceanic Eng.*, 26:4–25, 2001.
- [10] R. Goodman and R. Stern. Reflection and transmission of sound by elastic spherical shells. *J. Acoust. Soc. Am.*, 34, 1962.
- [11] P.T. Gough and D.W. Hawkins. Imaging algorithms for a stripmap synthetic aperture sonar: minimizing the effects of aperture errors and aperture undersampling. *ieeooceans*, 22:27–39, 1997.
- [12] P.T. Gough and D.W. Hawkins. Unified framework for modern synthetic aperture imaging algorithms. *The International Journal of Imaging Systems and Technology*, 8:343–358, 1997.

- [13] E. Bovio W.L.J. Fox K. LePage N.G. Pace R. Hollett P. Guerrini P.A. Sletner E. Michelozzi B. Moran H. Schmidt, A. Maguer and R. Grieve. Goats98 auv network sonar concepts for shallow water mine countermeasures. Technical Report Report SR-302, SACLANT Undersea Research Centre, 1998.
- [14] R. H. Hackman. Acoustic scattering from elastic solids. In A. D. Pierce and R. N. Thurston, editors, *Underwater Scattering and Radiation, Physical Acoustics, Volume 22*, pages 1–194. Academic Press, San Diego, CA, 1994.
- [15] D.W. Hawkings. *Synthetic aperture imaging algorithms: with application to wide bandwidth sonar*. PhD thesis, Univ. of Canterbury, Christchurch, N.Z., September 1996.
- [16] J.R. Edwards I. Veljkovic and H. Schmidt. Sub-critical isonification of buried elastic targets. In *OCEANS, 2001. MTS/IEEE Conference and Exhibition*, pages 59–62, Honolulu, HI, 2001.
- [17] R. Lim D. Tang K. L. Williams J. L. Lopes, C. L. Nesbitt and E. I. Thorsos. Measurements of subcritical grazing angle detection of targets buried under a rippled sand interface. *J. Acoust. Soc. Am.*, 113:2300, 2003.
- [18] J.D. Jackson. In *Classical Electrodynamics, 2nd Edition*. John Wiley & Sons, New York, NY, 1975.
- [19] D. Tang R. Lim K. Williamns, E. Thorsos and J. Lopes. Model/data comparisons of subcritical grazing angle detection of targets buried under a rippled sand interface. *J. Acoust. Soc. Am.*, 113:2300, 2003.
- [20] D. Komatitsch, J. Ritsema, and J. Tromp. The spectral-element method, Beowulf computing, and global seismology. *Science*, 298:1737–1742, 2002.
- [21] R. Lim. Acoustic scattering by a partially buried three-dimensional elastic obstacle. *J. Acoust. Soc. Am.*, 104:769–782, 1998.
- [22] R. Lim. Effect of bottom ripple on detection of a buried object using shallow-grazing-angle underwater sonar. *J. Acoust. Soc. Am.*, 109:2300, 2001.
- [23] R. Lim and J. J. Simpson. A comparison of benchmark acoustic scattering solutions with tank experiments on buried and partially buried spheres. *J. Acoust. Soc. Am.*, 105:1166, 1999.
- [24] L. D. Bibee P. A. Jumars W. B. Sawyer D. B. Albert R. H. Bennett T. K. Berger M. J. Buckingham N. P. Chotiros P. H. Dahl N. T. Dewitt P. Fleischer R. Flood C. F. Greenlaw D. V. Holliday M. H. Hulbert M. P. Hutnak P. D. Jackson J. S. Jaffe H. Paul Johnson D. L. Lavoie A. P. Lyons C. S. Martens D. E. McGehee K. D. Moore T. H. Orsi J. N. Piper R. I. Ray A. H. Reed R. F. Liko Self J. L. Schmidt S. G. Schock F. Simonet R. D. Stoll D. Tang D. E. Thistle E. I. Thorsos D. J. Walter M. D. Richardson, K. B. Briggs and R. A. Wheatcroft. Overview of sax99: Environmental considerations. *IEEE J. Oceanic Eng.*, 26:26–53, 2001.
- [25] P. Pérez M. Mignotte, C. Collet and P. Bouthemy. Sonar image segmentation using an unsupervised hierarchical markov random field model. *IEEE Trans. Image Processing*, 9, 2000.

- [26] E. Merzbacher. In *Quantum Mechanics, 2nd Edition*. John Wiley & Sons, New York, NY, 1970.
- [27] C. Morency and J. Tromp. Spectral element simulations of wave propagation in porous media. *Geophysical Journal International*, 175, 2008.
- [28] C. Morency and J. Tromp. Spectral-element simulations of wave propagation in porous media: Finite-frequency sensitivity kernels based upon adjoint methods. *Eos Trans. AGU*, 89(53), *Fall Meet. Suppl., Abstract NS13A-1088.*, 2008.
- [29] Luo Y. Morency, C. and J. Tromp. Spectral-element simulations of wave propagation in porous media: Finite-frequency sensitivity kernels based upon adjoint methods. In *POROMECHANICS IV (Proceedings of the Fourth Biot Conference on Poromechanics)*. DEStech Publications, Inc., Pennsylvania, USA, 2009.
- [30] D.D. Sternlicht M.R. Azimi-Sadjadi and D. Lemonds. Multi-aspect acoustic classification of buried objects. In *Proceedings MTS/IEEE Oceans 2003*, pages 478–484, 2003.
- [31] R. H. Hackman R. Lim, J. L. Lopes and D. G. Todoroff. Scattering by objects buried in underwater sediments: theory and experiment. *J. Acoust. Soc. Am.*, 93:1762–1783, 1992.
- [32] J. Wulf J. Sara M. Ericksen S. Schock, A. Tellier. Buried object scanning sonar. *IEEE Oceans*, 26:677–689, 2001.
- [33] G. S. Sammelmann. Acoustic scattering from large aspect ratio elastic targets. Technical Report CSS/TR-02/18, Naval Surface Warfare Center, CSS, October 2002.
- [34] G.S. Sammelmann. Personal computer shallow water acoustic tool-set (pc swat) 7.0: Users' manual. Technical Report CSS Report 02-07, Coastal Systems Station, 2002.
- [35] M. E. Tipping. Sparse bayesian learning and the relevance vector machine. *Journal of Machine Learning Research*, 1:211–244, 2001.
- [36] H. L. Van Trees. In *Optimum Array Processing (Detection, Estimation, and Modulation Theory, Part IV)*. Wiley-Interscience, New York, 2002.
- [37] S. A. Teukolsky W. H. Press, B. P. Flannery and W. T. Vetterling. In *Numerical Recipes, The Art of Scientific Computing*. Cambridge University Press, Cambridge, England, 1986.
- [38] P. B. Weichman. Green function approach to electromagnetic scattering in horizontally stratified media. *BAE Systems AIT Technical Report*, 2007.

A Supporting Data

This section provides data generated under this effort including 2D views of beamformed imagery (Sec. A.1), 3D views of the same imagery in Sec. A.2, histogram analysis of feature classes for targets and clutter in the AUV FEST 2008 data set (Sec. A.3), illustrations of pair-wise target/clutter separation for various feature pairs (Sec. A.4), illustrative target/clutter separations for feature triplets (Sec. A.5), statistical examination of ROC performance curves as a function of random partitions of the available data into training and test sets (Sec. A.6) and examples of processed BOSS imagery in the form of maximum intensity projections for 2D display and as selected depth slices (Sec. A.7).

A.1 Beamformed imagery (2D)

The figures in this section show front, top and side views of the scattering intensities for selected targets from the AUV 2007 collection. Target names were not provided to the performers due to evaluation protocols in the conduct of the AUV Fest 2007 collections. The only data available for public release is the binary designation of target versus clutter. The objects shown here are targets.

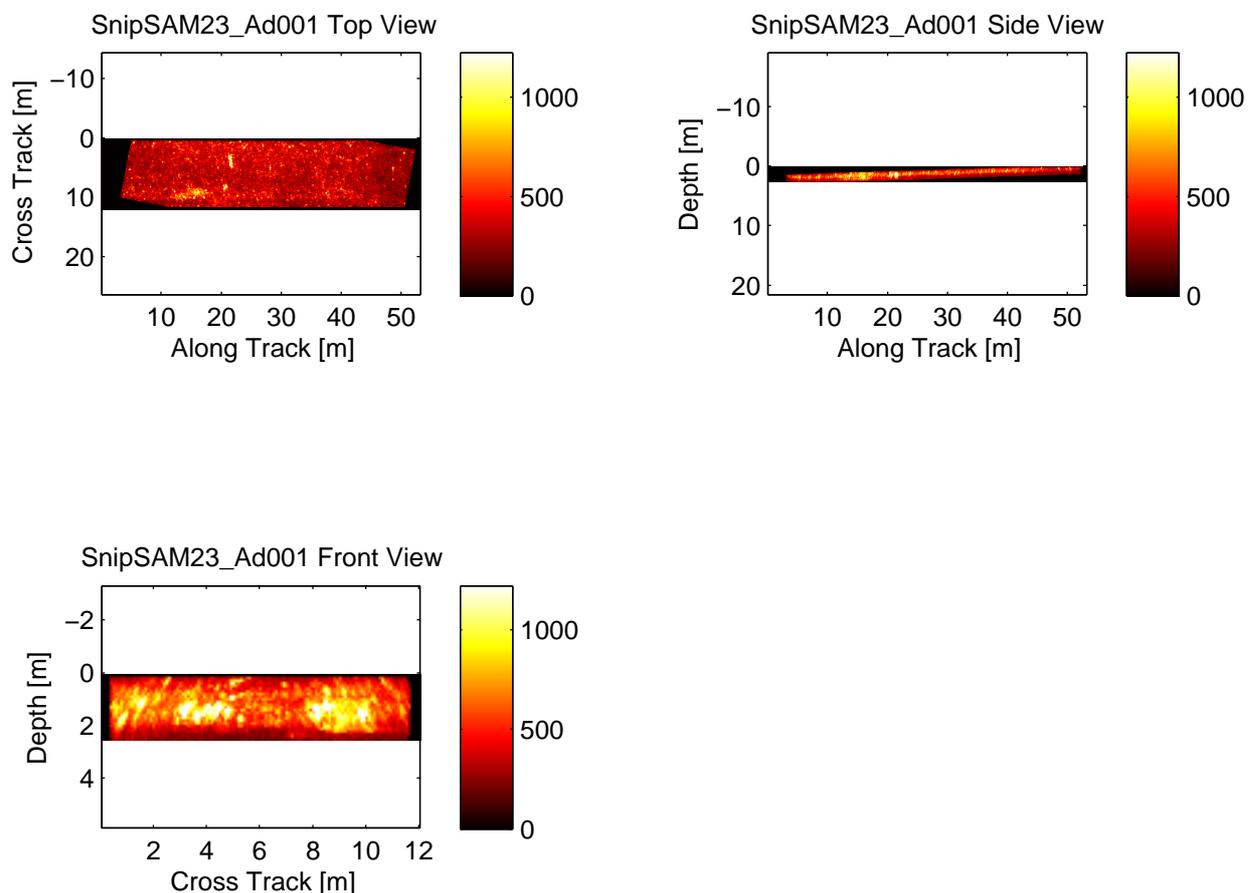


Figure 54: Various views of BOSS beamformed imagery for a data cube containing a target scattering response from the AUV FEST 2007 collection. **Upper left:** Top view of a target object shown as the maximum intensity projection (MIP) $P_{xy}(x, y)$ (see eq. 114). **Upper right:** Side view of a target object showing the MIP $P_{yz}(y, z)$ (see eq. 115). **Lower left:** Front view of a target object showing the MIP $P_{xz}(x, z)$ (see eq. 115).

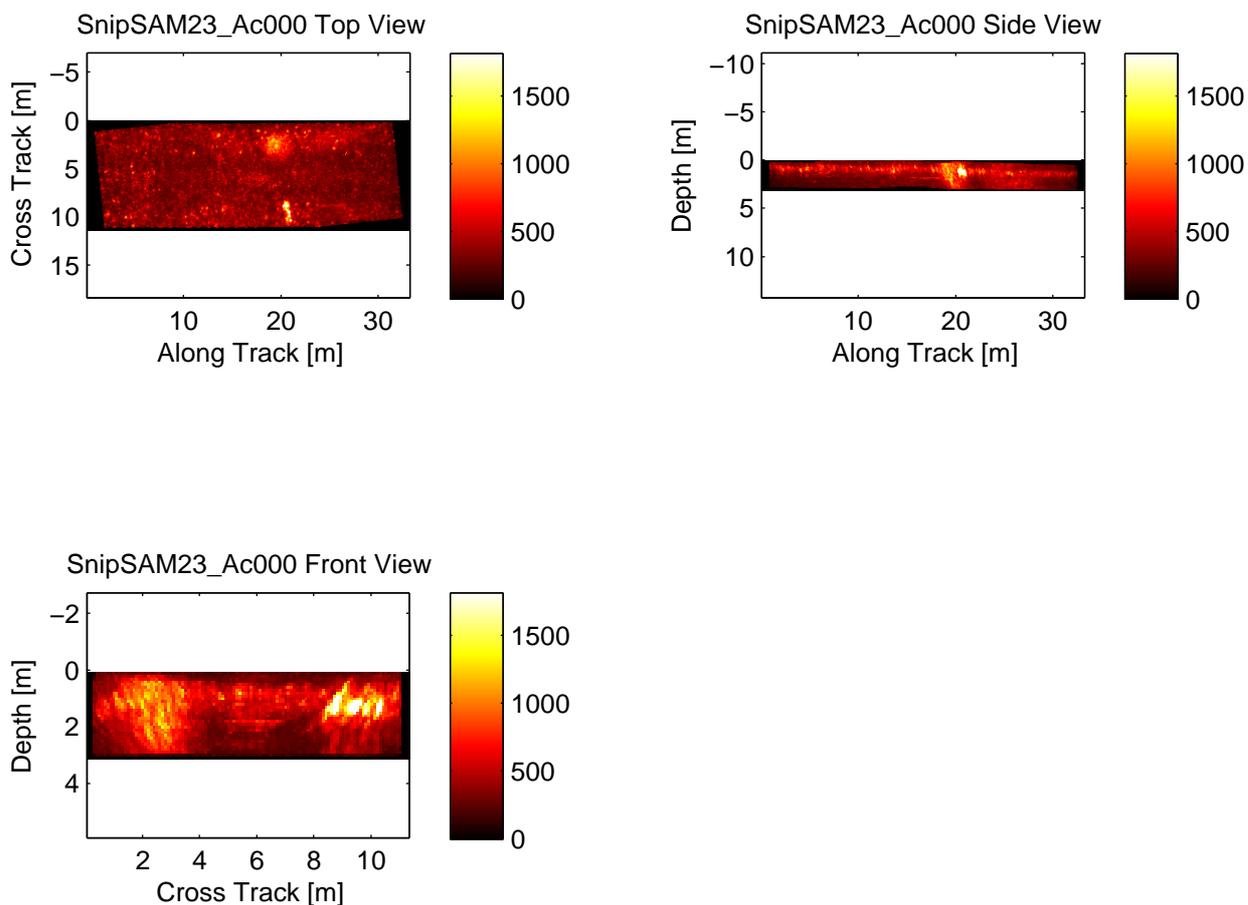


Figure 55: Various views of BOSS beamformed imagery for a data cube containing a target scattering response from the AUV FEST 2007 collection. **Upper left:** Top view of a target object shown as the maximum intensity projection (MIP) $P_{xy}(x, y)$ (see eq. 114). **Upper right:** Side view of a target object showing the MIP $P_{yz}(y, z)$ (see eq. 115). **Lower left:** Front view of a target object showing the MIP $P_{xz}(x, z)$ (see eq. 115).

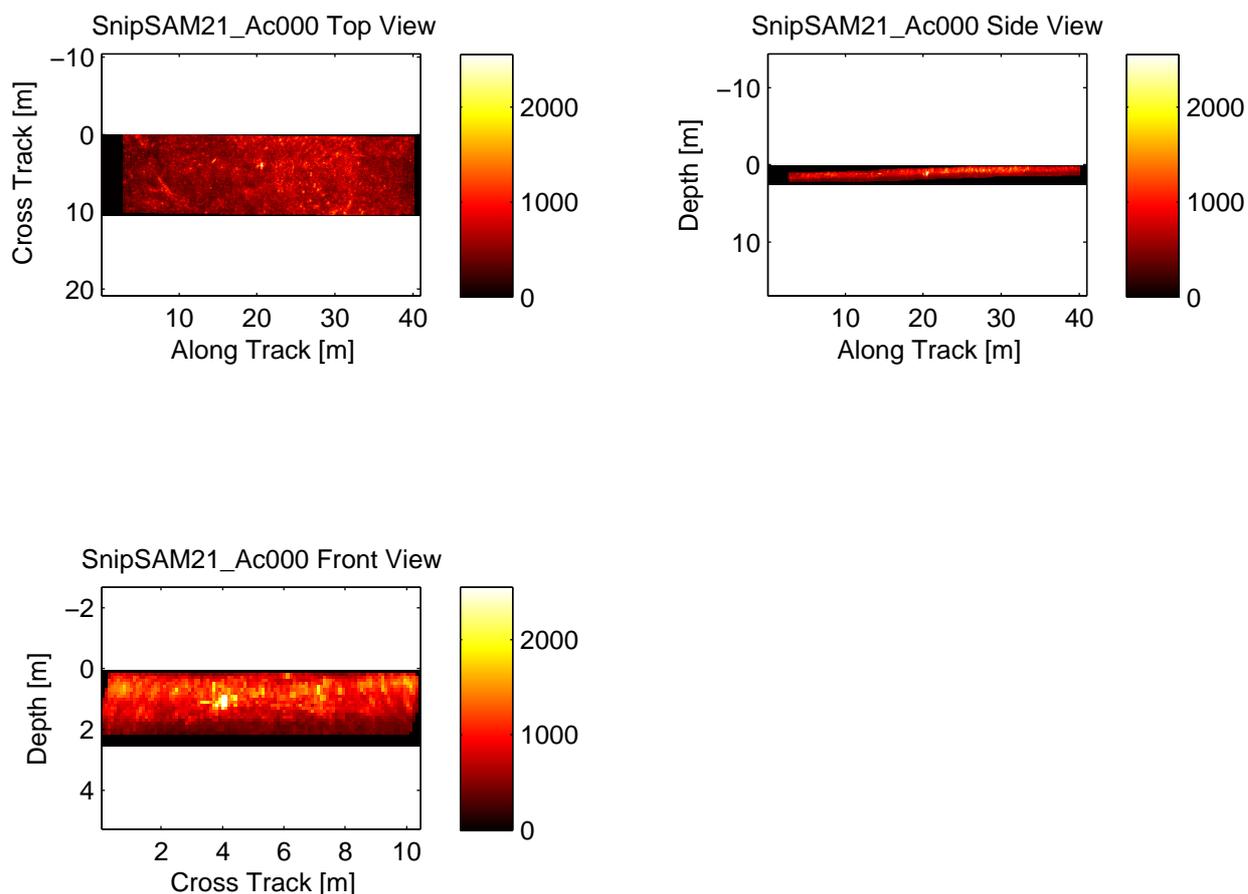


Figure 56: Various views of BOSS beamformed imagery for a data cube containing a target scattering response from the AUV FEST 2007 collection. **Upper left:** Top view of a target object shown as the maximum intensity projection (MIP) $P_{xy}(x, y)$ (see eq. 114). **Upper right:** Side view of a target object showing the MIP $P_{yz}(y, z)$ (see eq. 115). **Lower left:** Front view of a target object showing the MIP $P_{xz}(x, z)$ (see eq. 115).

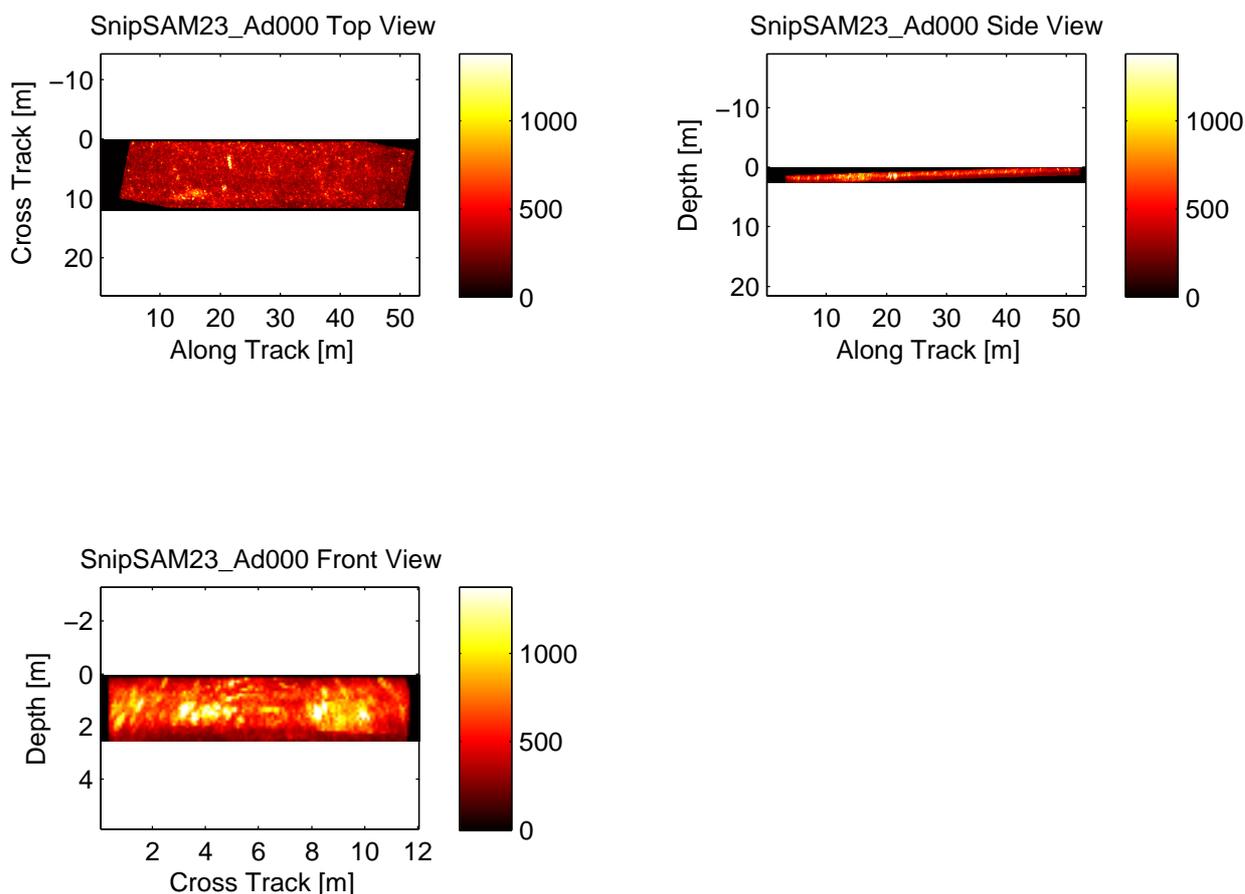


Figure 57: Various views of BOSS beamformed imagery for a data cube containing a target scattering response from the AUV FEST 2007 collection. **Upper left:** Top view of a target object shown as the maximum intensity projection (MIP) $P_{xy}(x, y)$ (see eq. 114). **Upper right:** Side view of a target object showing the MIP $P_{yz}(y, z)$ (see eq. 115). **Lower left:** Front view of a target object showing the MIP $P_{xz}(x, z)$ (see eq. 115).

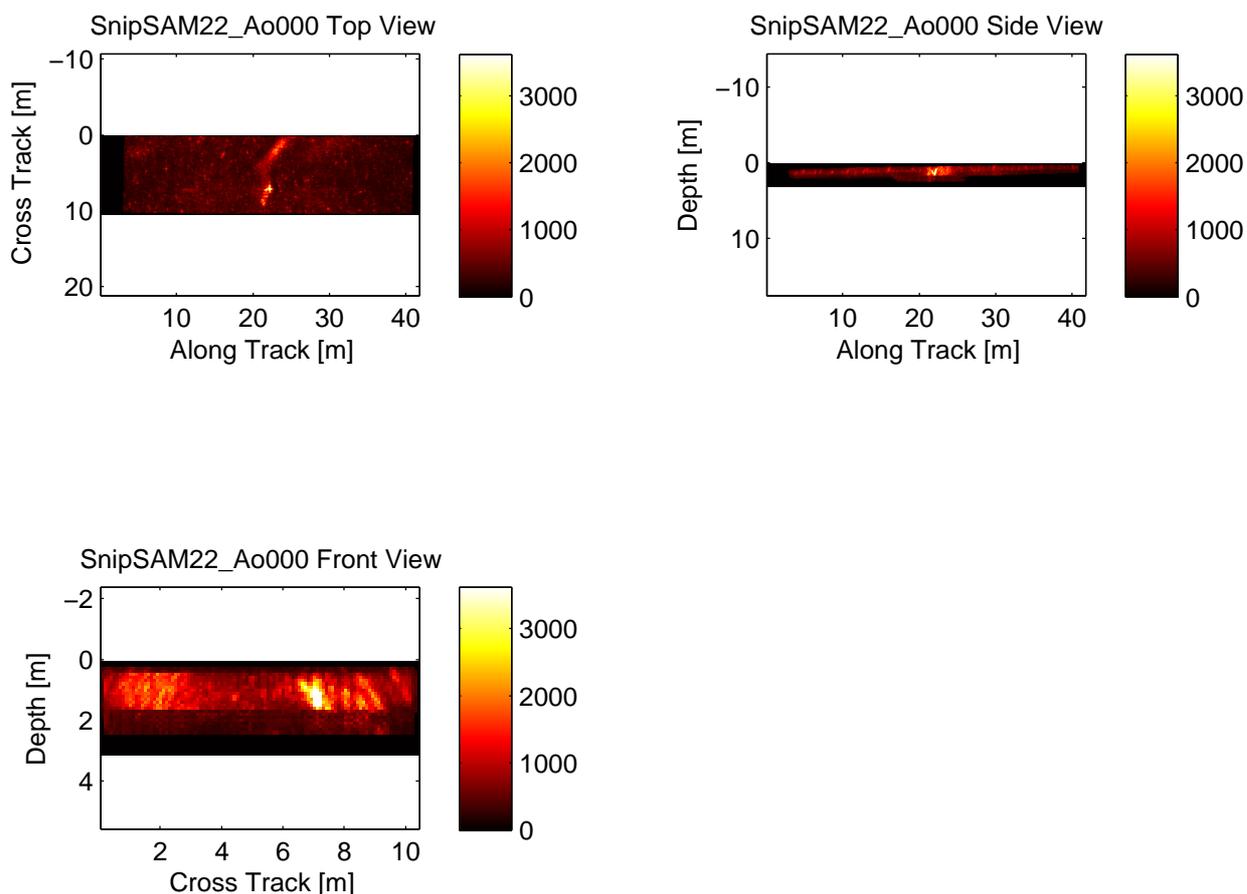


Figure 58: Various views of BOSS beamformed imagery for a data cube containing a target scattering response from the AUV FEST 2007 collection. **Upper left:** Top view of a target object shown as the maximum intensity projection (MIP) $P_{xy}(x, y)$ (see eq. 114). **Upper right:** Side view of a target object showing the MIP $P_{yz}(y, z)$ (see eq. 115). **Lower left:** Front view of a target object showing the MIP $P_{xz}(x, z)$ (see eq. 115).

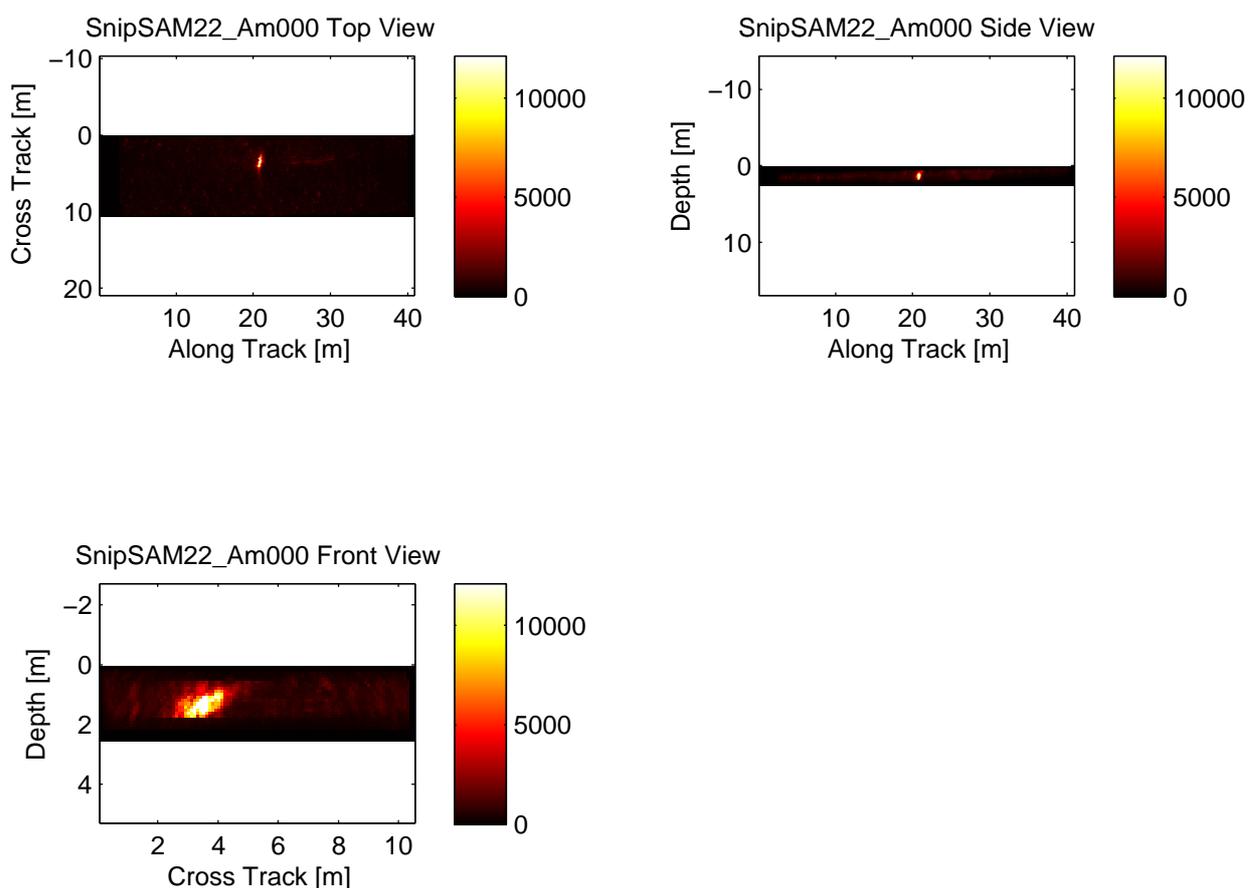


Figure 59: Various views of BOSS beamformed imagery for a data cube containing a target scattering response from the AUV FEST 2007 collection. **Upper left:** Top view of a target object shown as the maximum intensity projection (MIP) $P_{xy}(x, y)$ (see eq. 114). **Upper right:** Side view of a target object showing the MIP $P_{yz}(y, z)$ (see eq. 115). **Lower left:** Front view of a target object showing the MIP $P_{xz}(x, z)$ (see eq. 115).

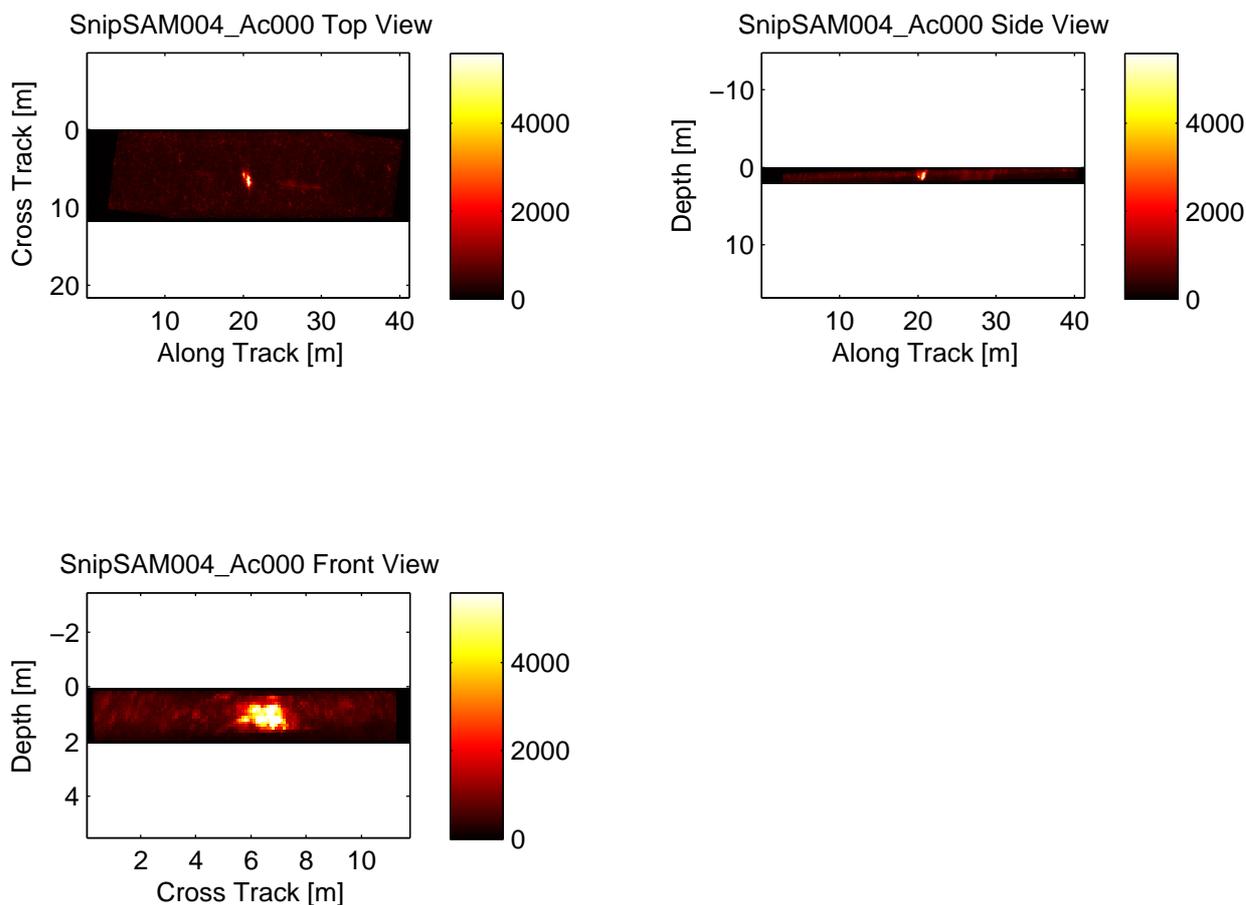


Figure 60: Various views of BOSS beamformed imagery for a data cube containing a target scattering response from the AUV FEST 2007 collection. **Upper left:** Top view of a target object shown as the maximum intensity projection (MIP) $P_{xy}(x, y)$ (see eq. 114). **Upper right:** Side view of a target object showing the MIP $P_{yz}(y, z)$ (see eq. 115). **Lower left:** Front view of a target object showing the MIP $P_{xz}(x, z)$ (see eq. 115).

A.2 Beamformed imagery (3D)

The figures in this section 3D isocontours of scattering intensities for selected targets from the AUV 2007 collection. Target names were not provided to the performers due to evaluation protocols in the conduct of the AUV Fest 2007 collections. The only data available for public release is the binary designation of target versus clutter. The objects shown here are targets.

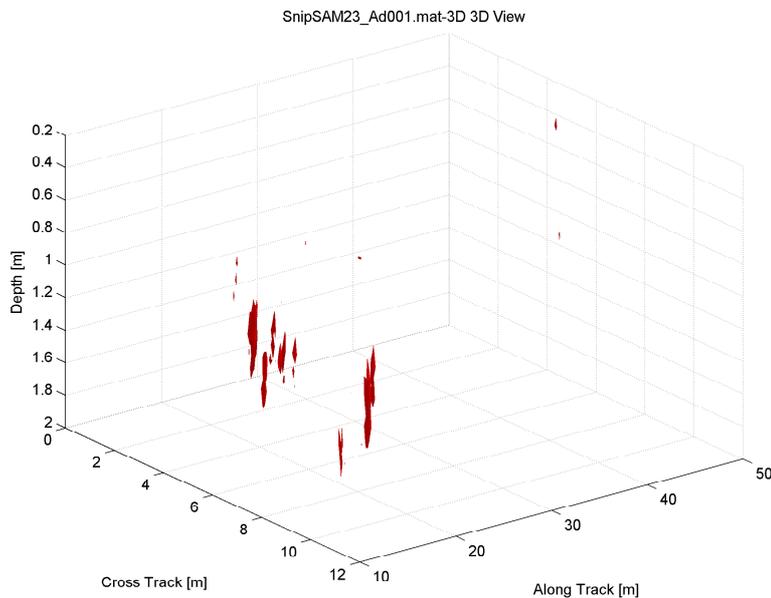


Figure 61: 3D isocontour of the subvolume for an iso-intensity value given by 1/2 of the maximum intensity value of the data cube, corresponding to the target in Fig. 54.

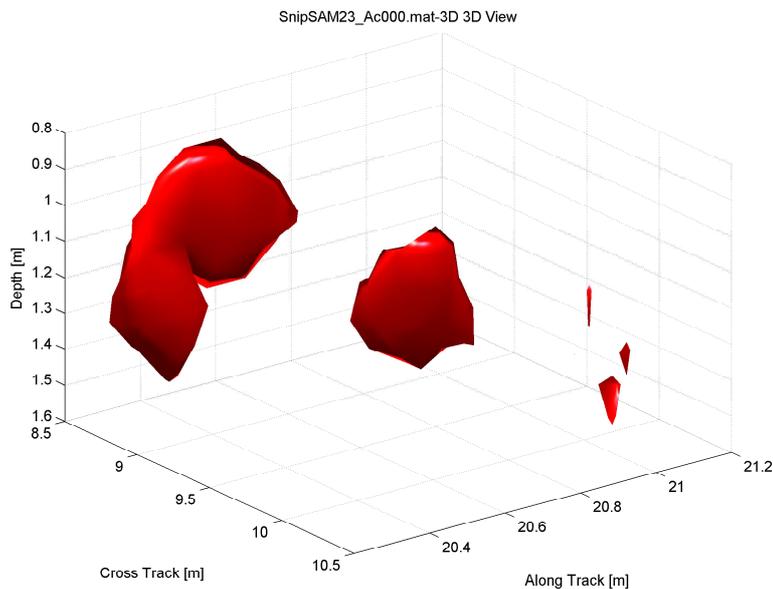


Figure 62: 3D isocontour of the subvolume for an iso-intensity value given by 1/2 of the maximum intensity value of the data cube, corresponding to the target in Fig. 55.

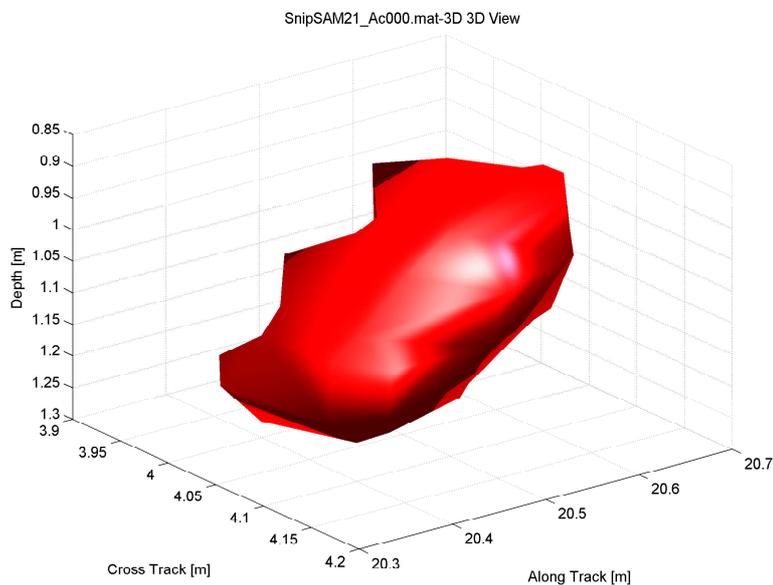


Figure 63: 3D isocontour of the subvolume for an iso-intensity value given by 1/2 of the maximum intensity value of the data cube, corresponding to the target in Fig. 56.

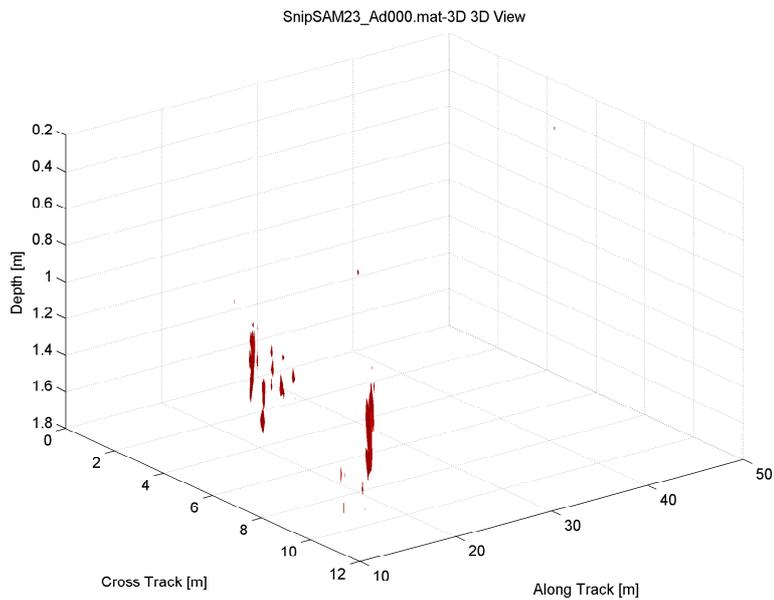


Figure 64: 3D isocontour of the subvolume for an iso-intensity value given by 1/2 of the maximum intensity value of the data cube, corresponding to the target in Fig. 57.

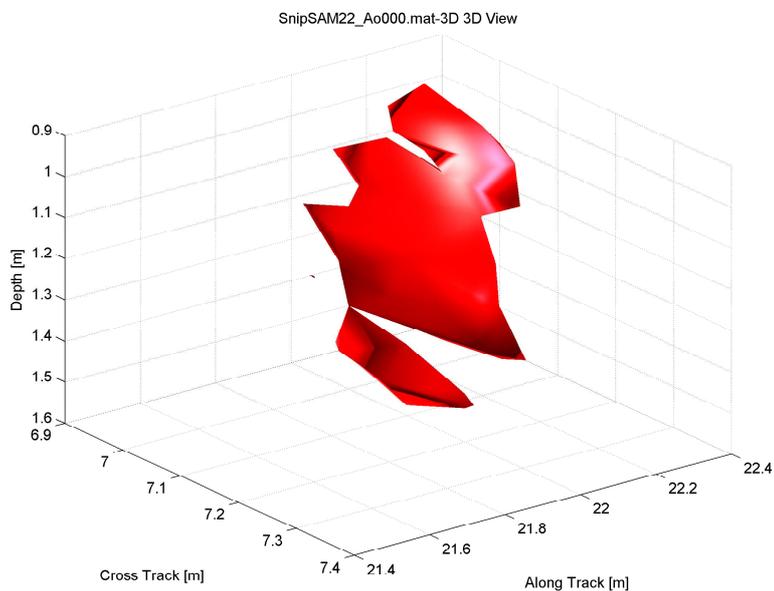


Figure 65: 3D isocontour of the subvolume for an iso-intensity value given by 1/2 of the maximum intensity value of the data cube, corresponding to the target in Fig. 58.

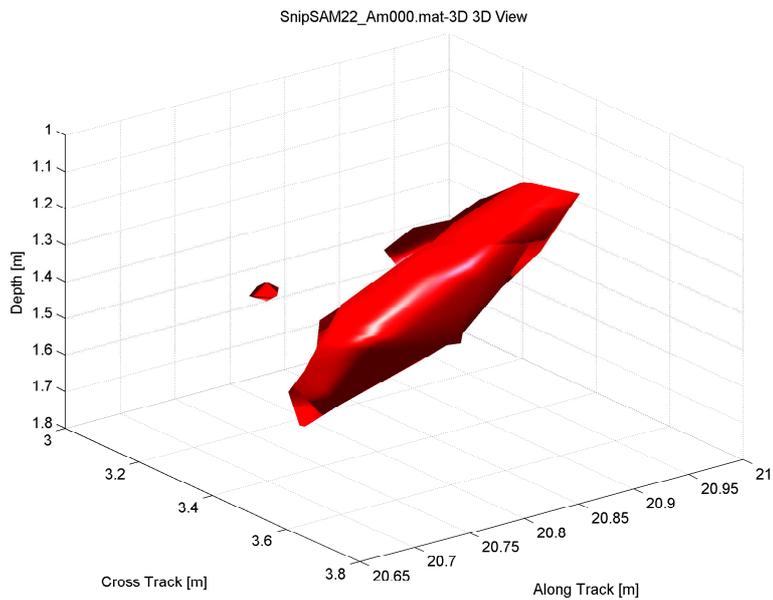


Figure 66: 3D isocontour of the subvolume for an iso-intensity value given by 1/2 of the maximum intensity value of the data cube, corresponding to the target in Fig. 59.

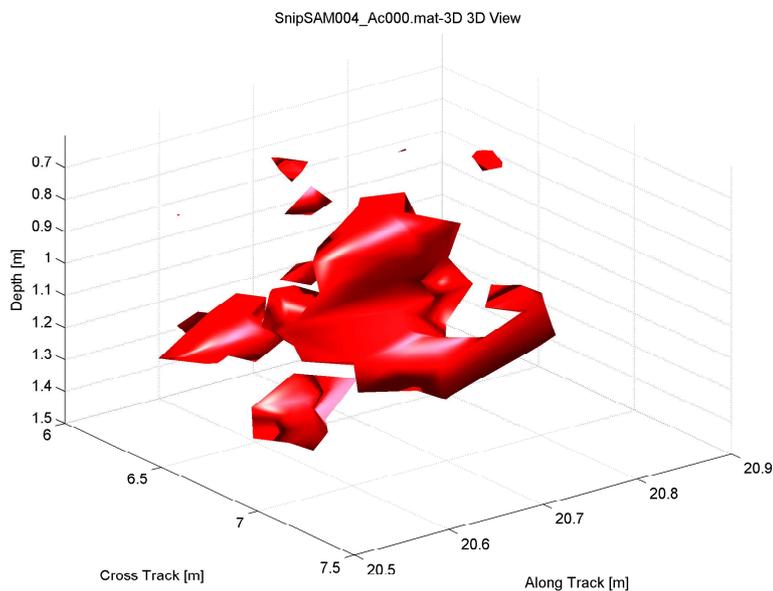


Figure 67: 3D isocontour of the subvolume for an iso-intensity value given by 1/2 of the maximum intensity value of the data cube, corresponding to the target in Fig. 60.

A.3 Histogram analysis of target/clutter feature distributions for AUV Fest 2008 data

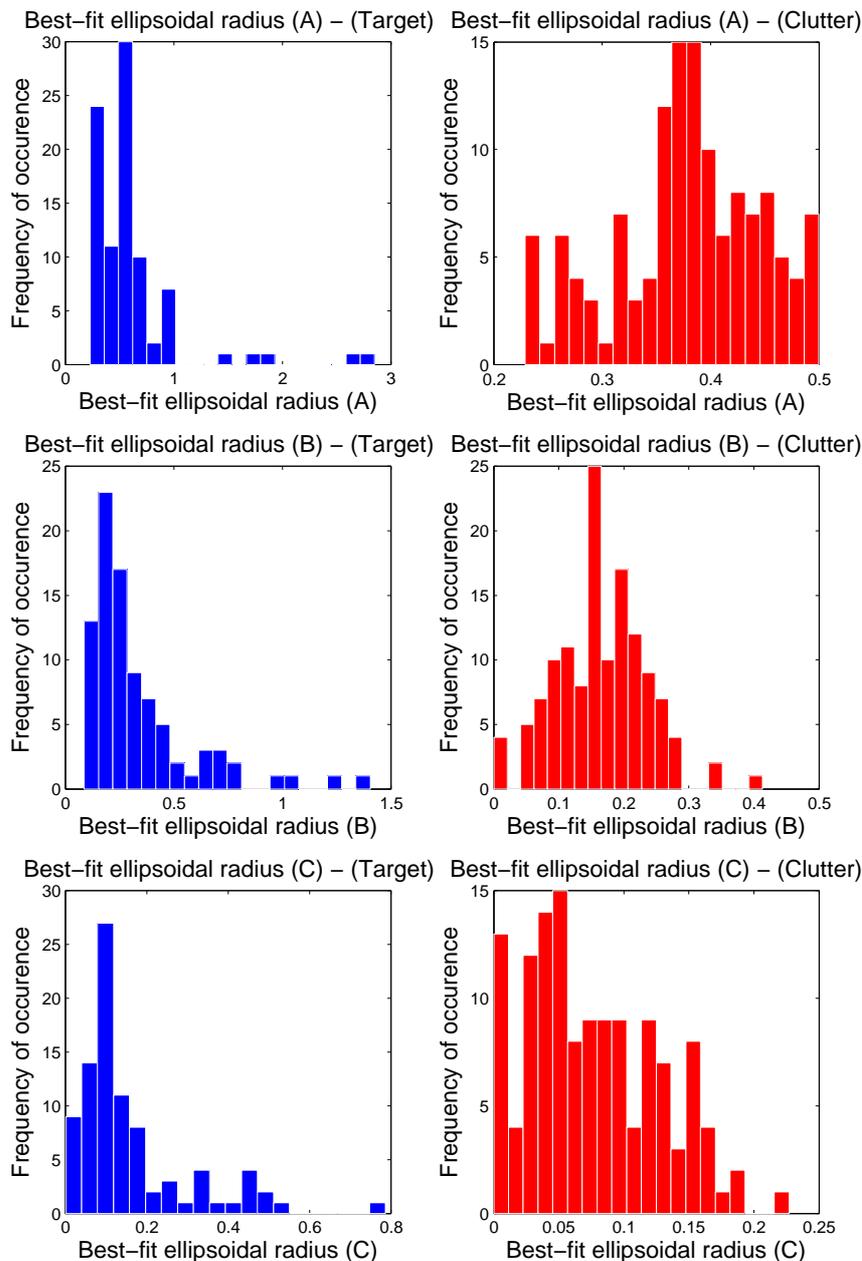


Figure 68: For this feature type we fit a volume associated with a clutter or target item to the *best fitting ellipsoid* characterized by the three principal radii *A* (top panel), *B* (middle panel) and *C* (bottom panel). In each case the target and clutter features are shown on the left hand and right hand sides, respectively. The features were derived from 3D beamform products using data acquired by the BOSS platform at the AUV Fest 2008. A total of 132 clutter items and 89 targets were processed in generating these distributions.

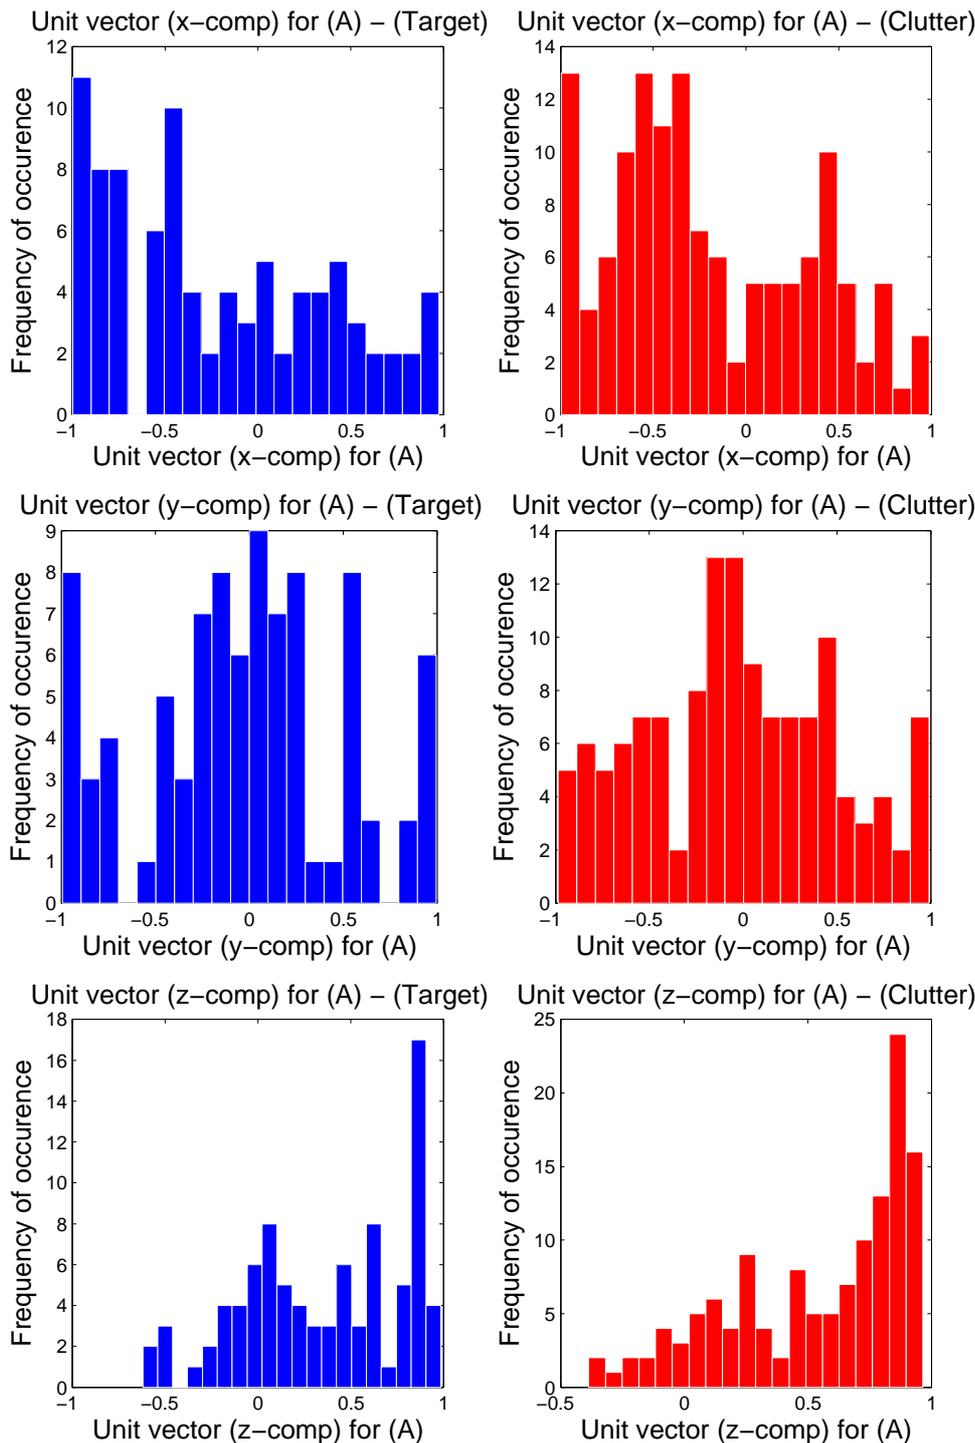


Figure 69: This feature type is given by the unit vector \hat{u}_A describing the direction of the principal radius A of the best fitting ellipsoid volume. The x -component, y -component and z -components are shown, respectively, in the top, middle and bottom panels. In each case the target and clutter features are shown on the left hand and right hand sides, respectively. The features were derived from 3D beamform products using data acquired by the BOSS platform at the AUV Fest 2008. A total of 132 clutter items and 89 targets were processed in generating these distributions.

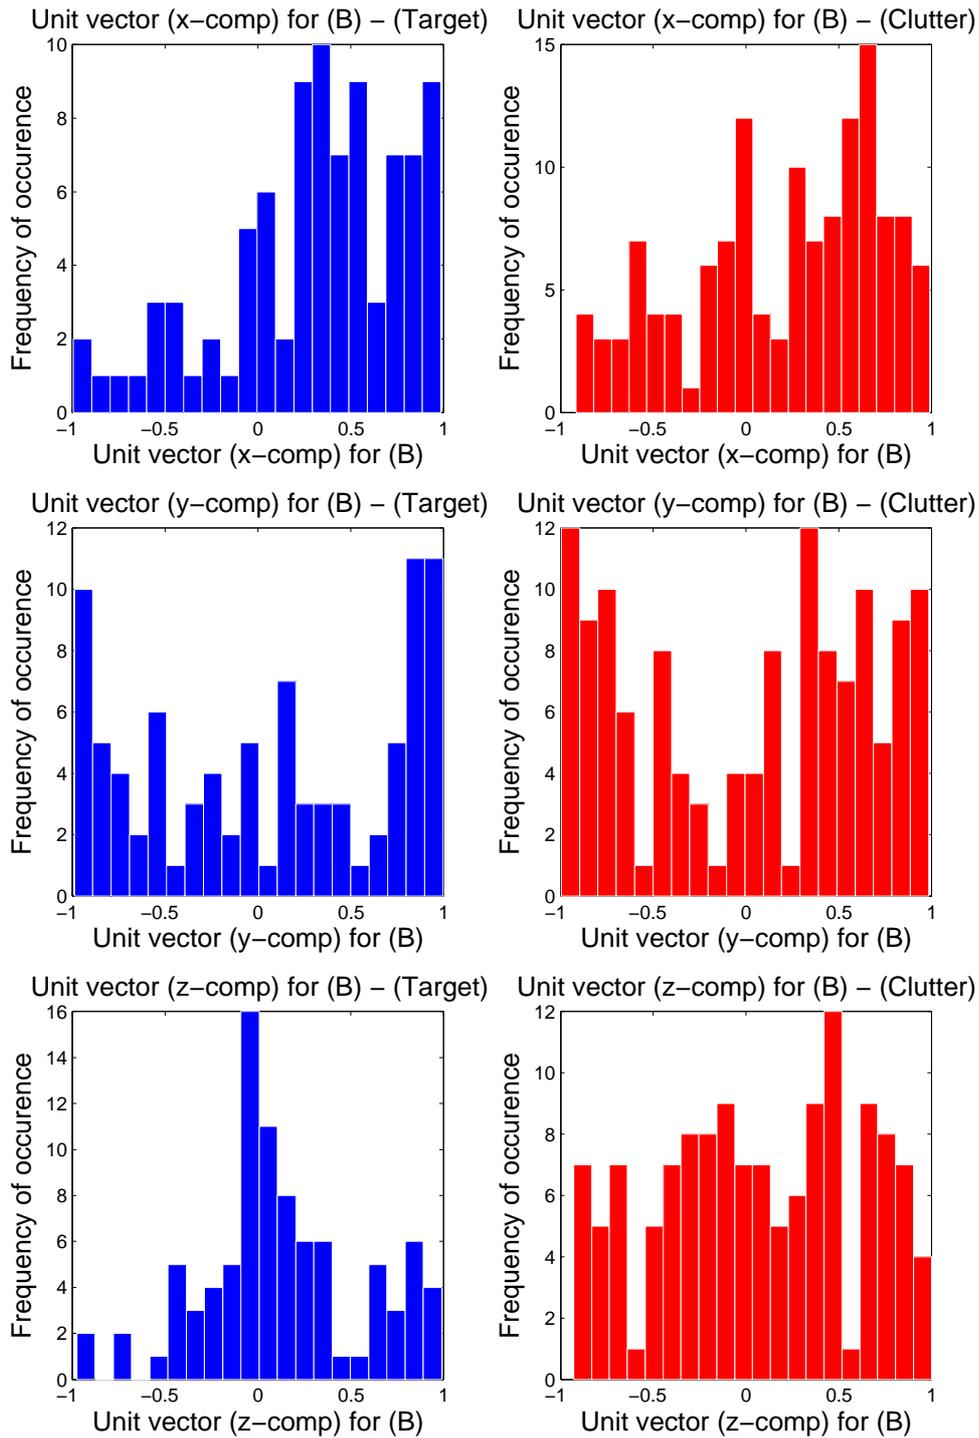


Figure 70: This feature type is given by the unit vector \hat{u}_B describing the direction of the principal radius B of the best fitting ellipsoid volume. The x -component, y -component and z -components are shown, respectively, in the top, middle and bottom panels. In each case the target and clutter features are shown on the left hand and right hand sides, respectively. The features were derived from 3D beamform products using data acquired by the BOSS platform at the AUV Fest 2008. A total of 132 clutter items and 89 targets were processed in generating these distributions.

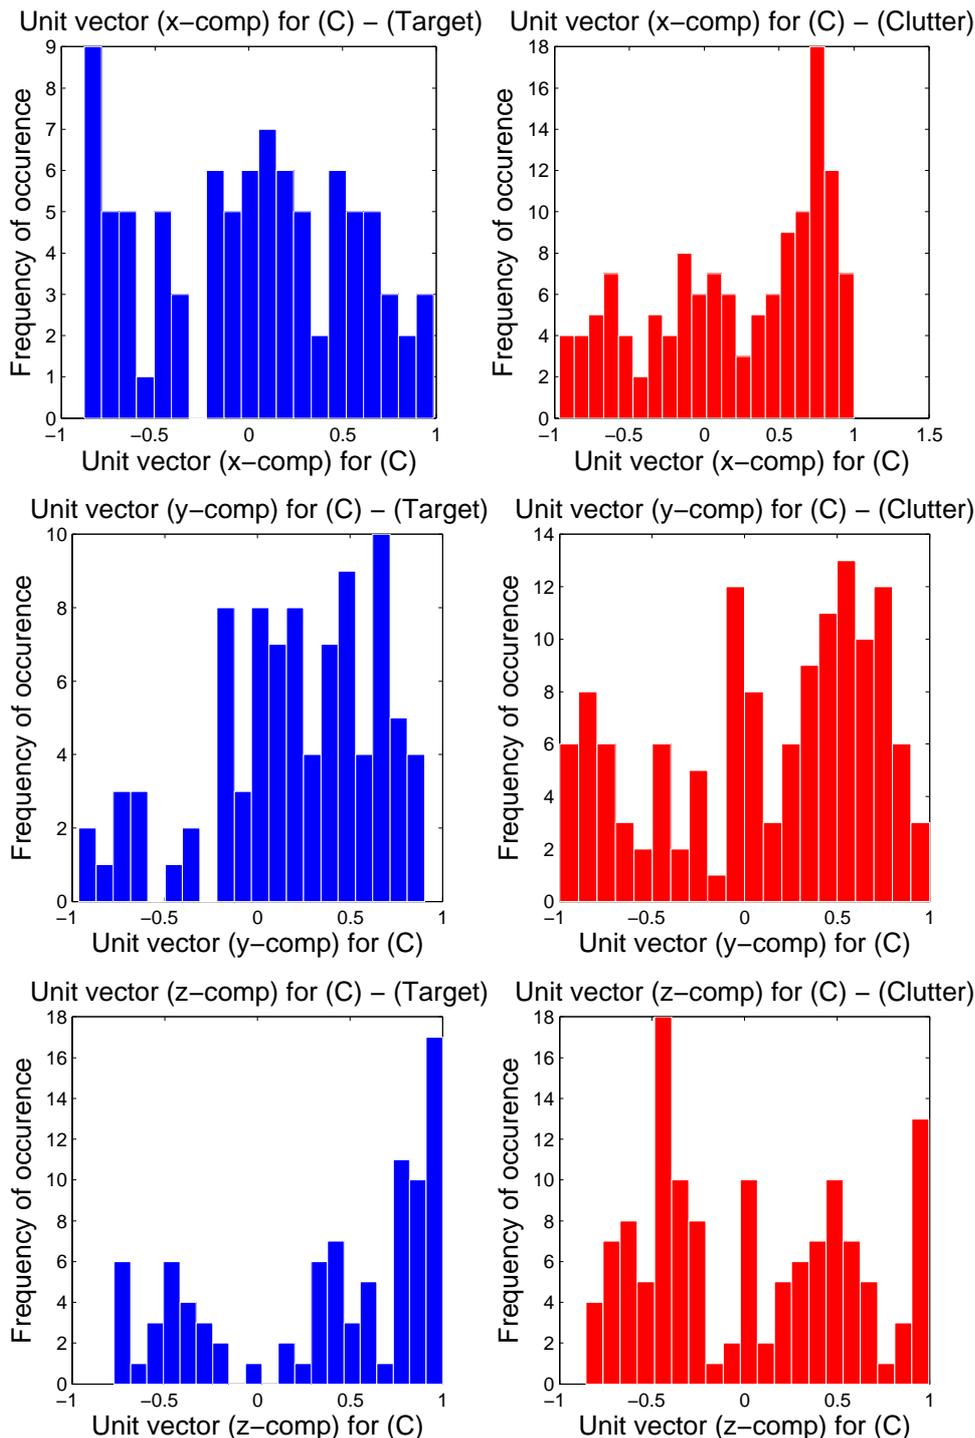


Figure 71: This feature type is given by the unit vector \hat{u}_C describing the direction of the principal radius C of the best fitting ellipsoid volume. The x -component, y -component and z -components are shown, respectively, in the top, middle and bottom panels. In each case the target and clutter features are shown on the left hand and right hand sides, respectively. The features were derived from 3D beamform products using data acquired by the BOSS platform at the AUV Fest 2008. A total of 132 clutter items and 89 targets were processed in generating these distributions.

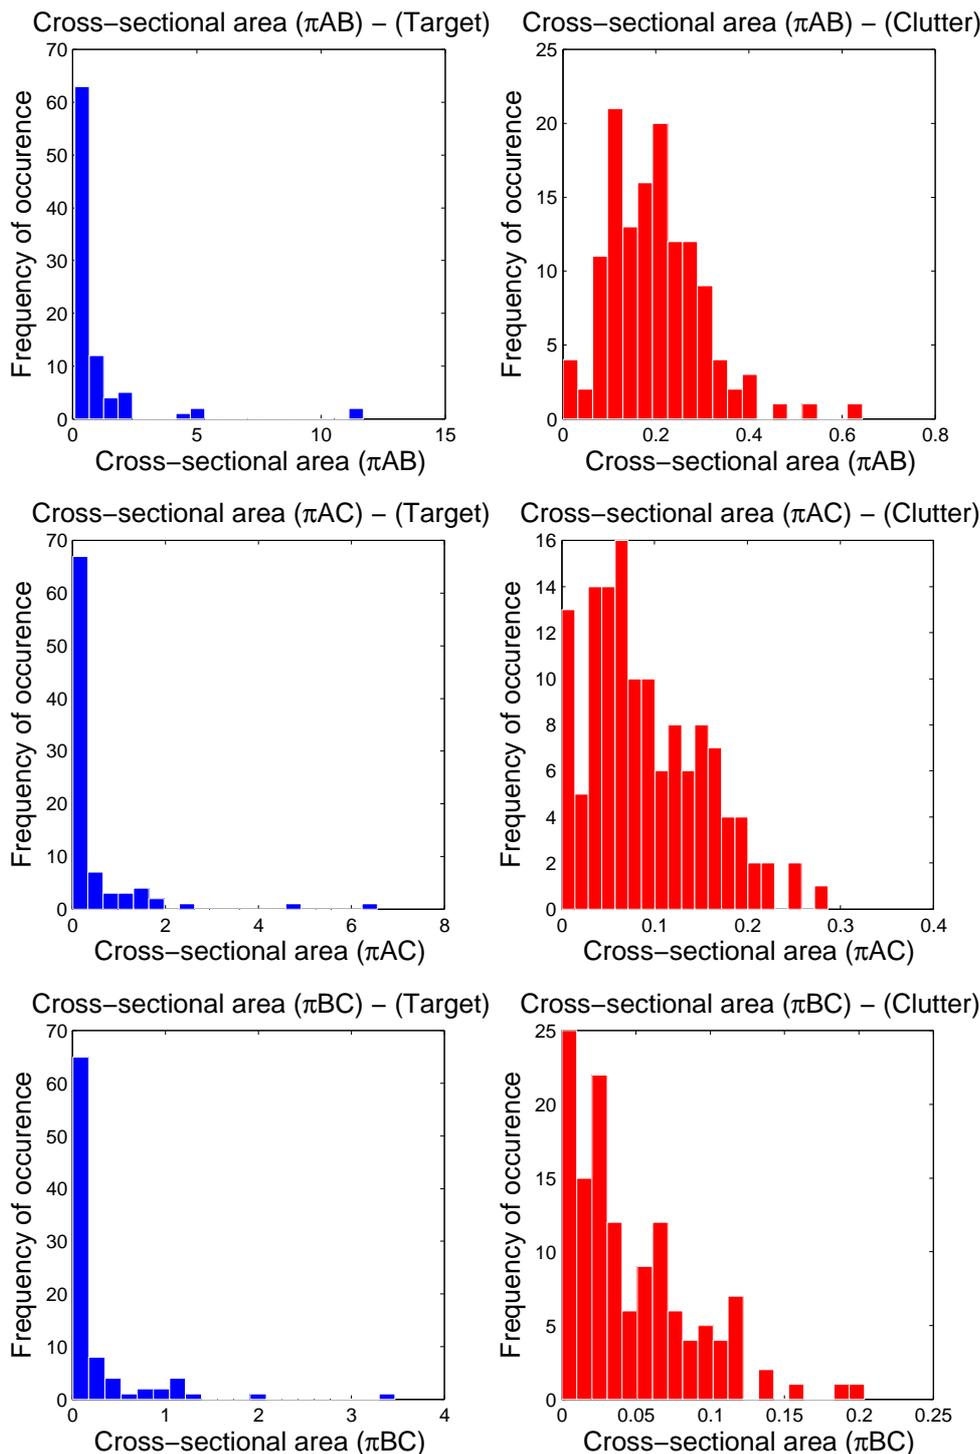


Figure 72: This feature type is given by the cross-sectional areas of the best fitting ellipsoids to the thresholded intensity blobs identified as target/clutter items in the BOSS data. These cross-sectional areas are given as π_{AB}, π_{AC} and π_{BC} , and are shown, respectively in the top, middle and bottom panels. In each case the target and clutter features are shown on the left hand and right hand sides, respectively. The features were derived from 3D beamform products using data acquired by the BOSS platform at the AUV Fest 2008. A total of 132 clutter items and 89 targets were processed in generating these distributions.

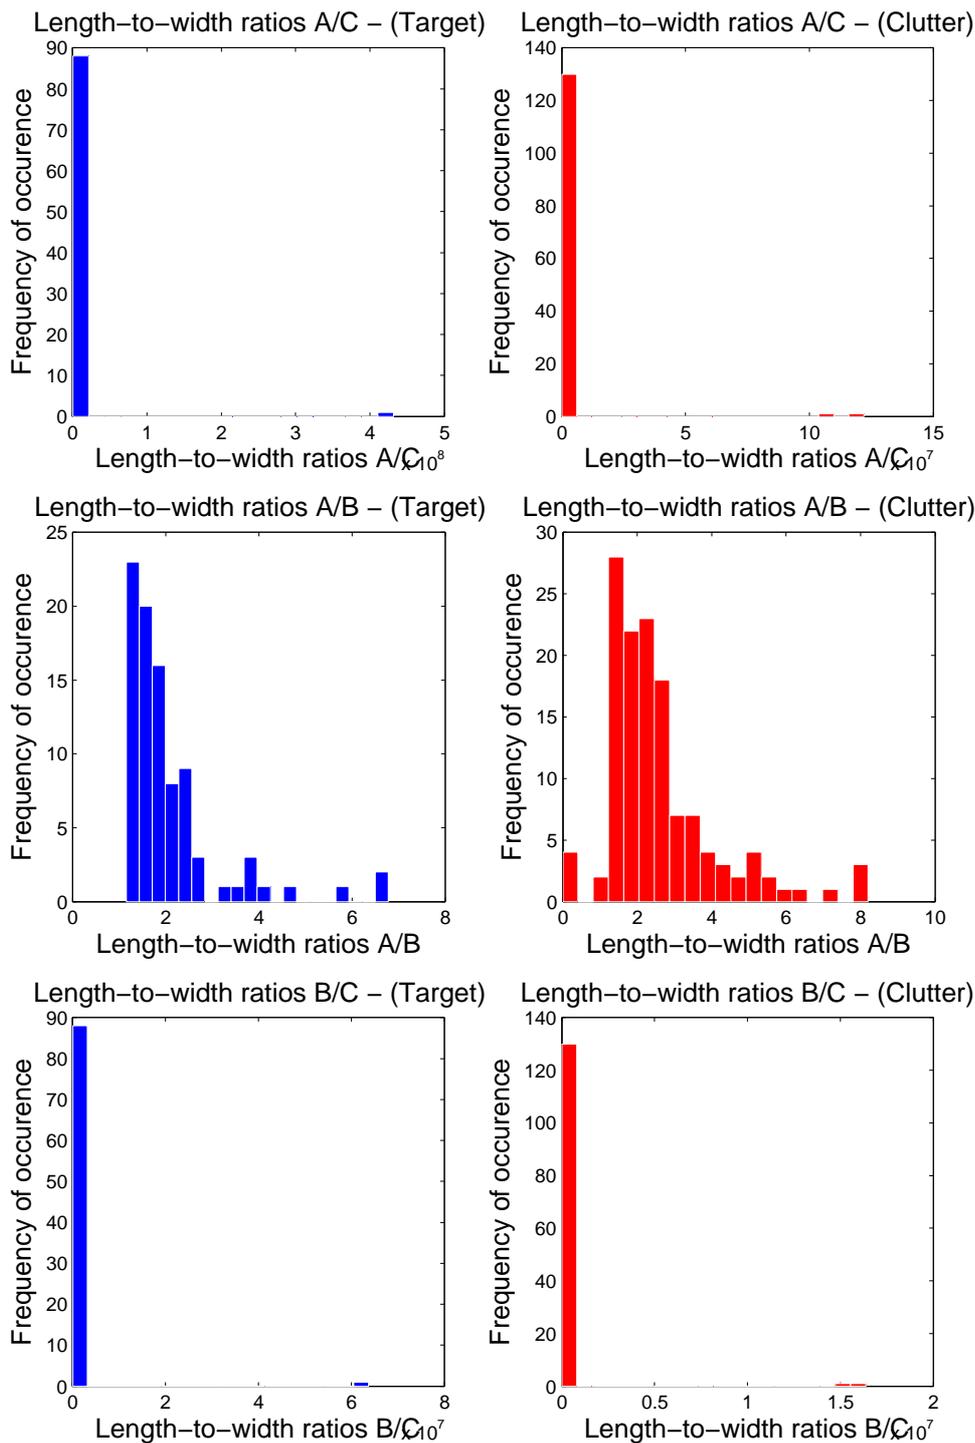


Figure 73: This feature type is given by the length to width ratios A/C , A/B and B/C where A , B and C are the principal radii of the best fitting ellipsoids. These ratios are shown, respectively, in the top, middle and bottom panels. In each case the target and clutter features are shown on the left hand and right hand sides, respectively. The features were derived from 3D beamform products using data acquired by the BOSS platform at the AUV Fest 2008. A total of 132 clutter items and 89 targets were processed in generating these distributions.

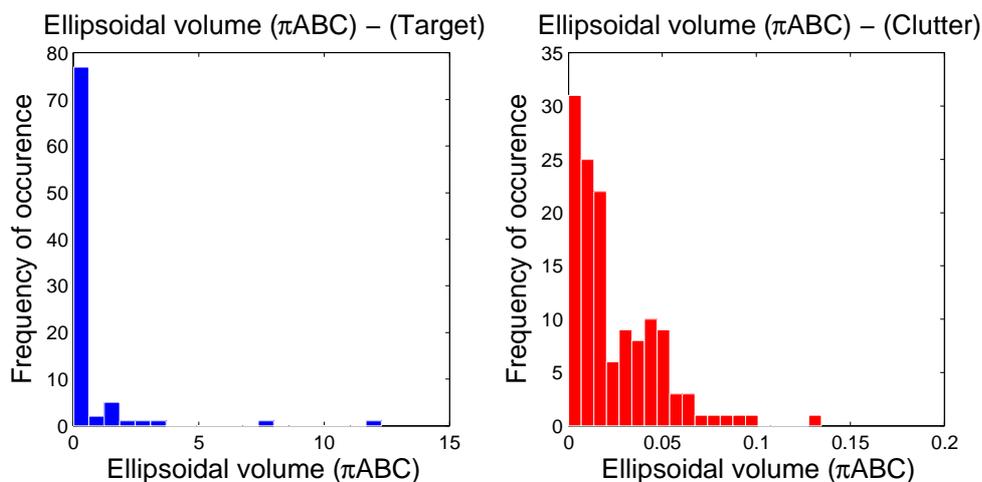


Figure 74: This feature type is given by the ellipsoidal volume πABC of the best fitting ellipsoids to each of the thresholded intensity blobs identified as target/clutter items in the BOSS data. The target and clutter distributions for this features are shown on the left hand and right hand sides, respectively. A total of 132 clutter items and 89 targets were processed in generating these distributions.

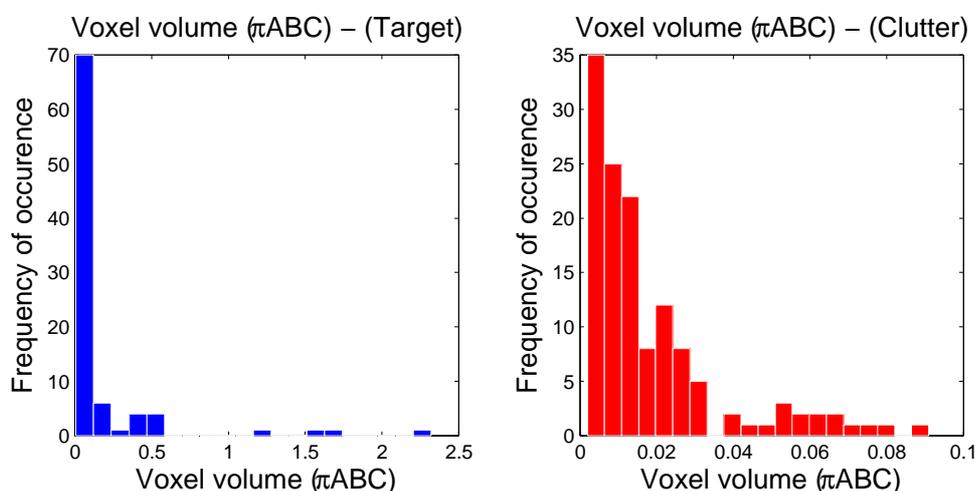


Figure 75: This feature type is given by the voxel volume to each of the thresholded intensity blobs identified as target/clutter items in the BOSS data. The target and clutter distributions for this features are shown on the left hand and right hand sides, respectively. A total of 132 clutter items and 89 targets were processed in generating these distributions.

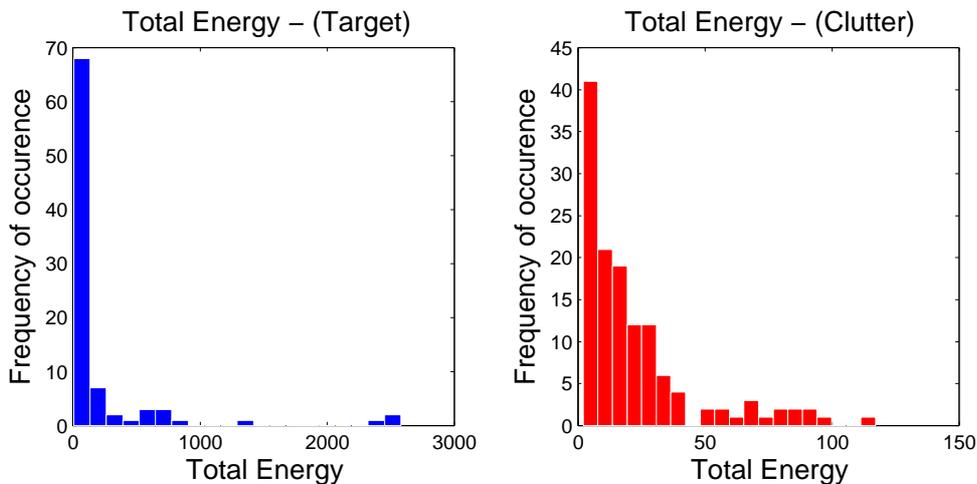


Figure 76: This feature type is the total energy of the voxel volume corresponding to an identified target or clutter item from the BOSS data. The target and clutter distributions for this features are shown on the left hand and right hand sides, respectively. A total of 132 clutter items and 89 targets were processed in generating these distributions.

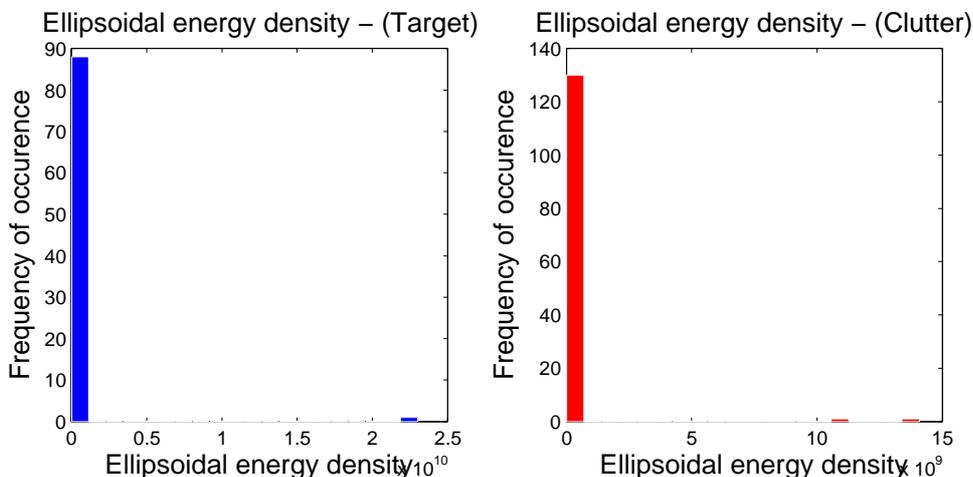


Figure 77: This feature type is the total energy of the voxels contained within the best-fitting ellipsoidal volumes corresponding to an identifier target or clutter items from the BOSS data. The target and clutter distributions for this features are shown on the left hand and right hand sides, respectively. A total of 132 clutter items and 89 targets were processed in generating these distributions.

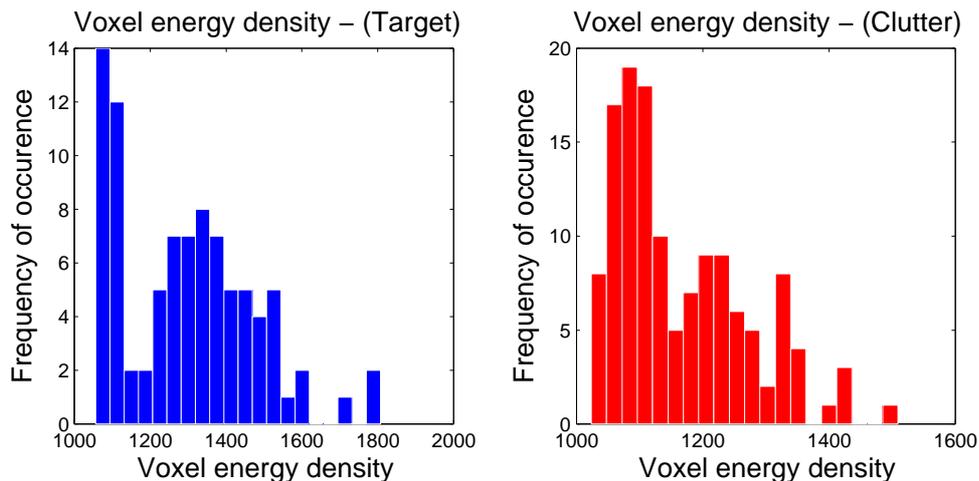


Figure 78: This feature type is the voxel energy density for voxels contained within the best-fitting ellipsoidal volumes corresponding to an identified target or clutter items from the BOSS data. The target and clutter distributions for this features are shown on the left hand and right hand sides, respectively. A total of 132 clutter items and 89 targets were processed in generating these distributions.

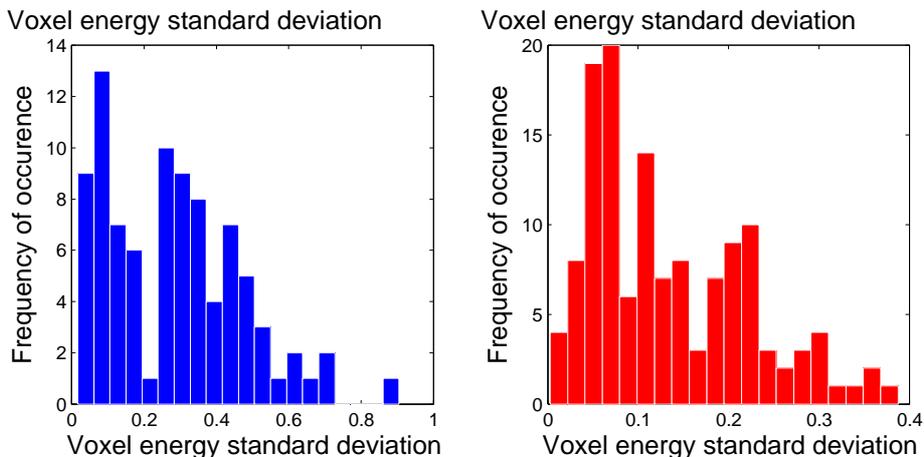


Figure 79: This feature type is the voxel energy standard deviation for voxels contained within the best-fitting ellipsoidal volumes corresponding to an identified target or clutter items from the BOSS data. The target and clutter distributions for this features are shown on the left hand and right hand sides, respectively. A total of 132 clutter items and 89 targets were processed in generating these distributions.

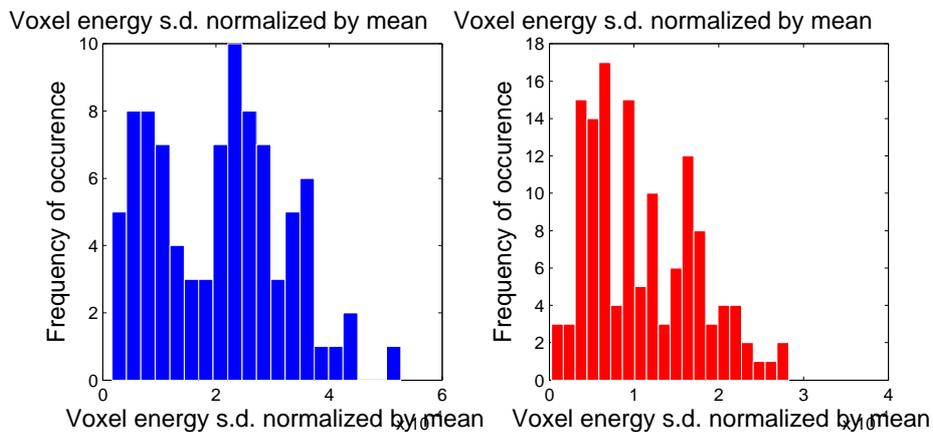


Figure 80: This feature type is the voxel energy standard deviation normalized by mean contained within the best-fitting ellipsoidal volumes corresponding to an identified target or clutter items from the BOSS data. The target and clutter distributions for this features are shown on the left hand and right hand sides, respectively. A total of 132 clutter items and 89 targets were processed in generating these distributions.

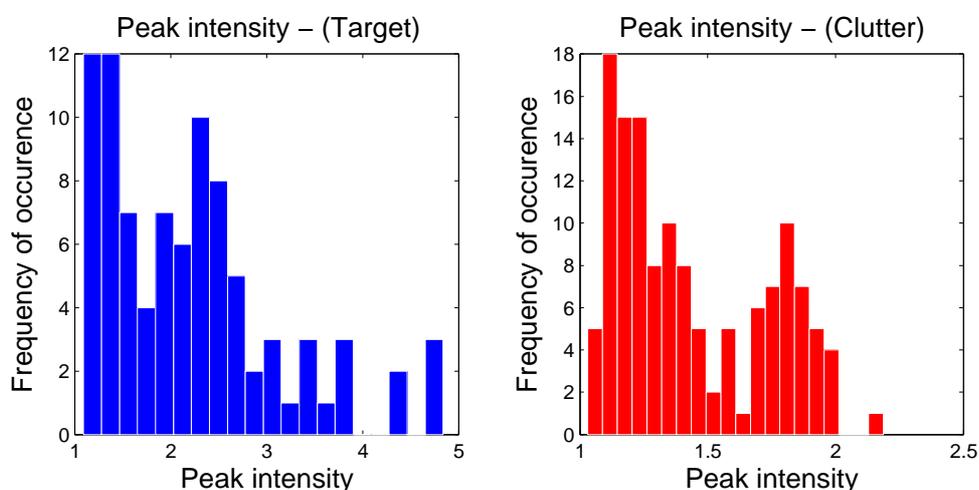


Figure 81: This feature type is the peak intensity voxel contained within the total voxel volume identified with a target or clutter item from the AUV FEST 2008 BOSS data. The target and clutter distributions for this features are shown on the left hand and right hand sides, respectively. A total of 132 clutter items and 89 targets were processed in generating these distributions.

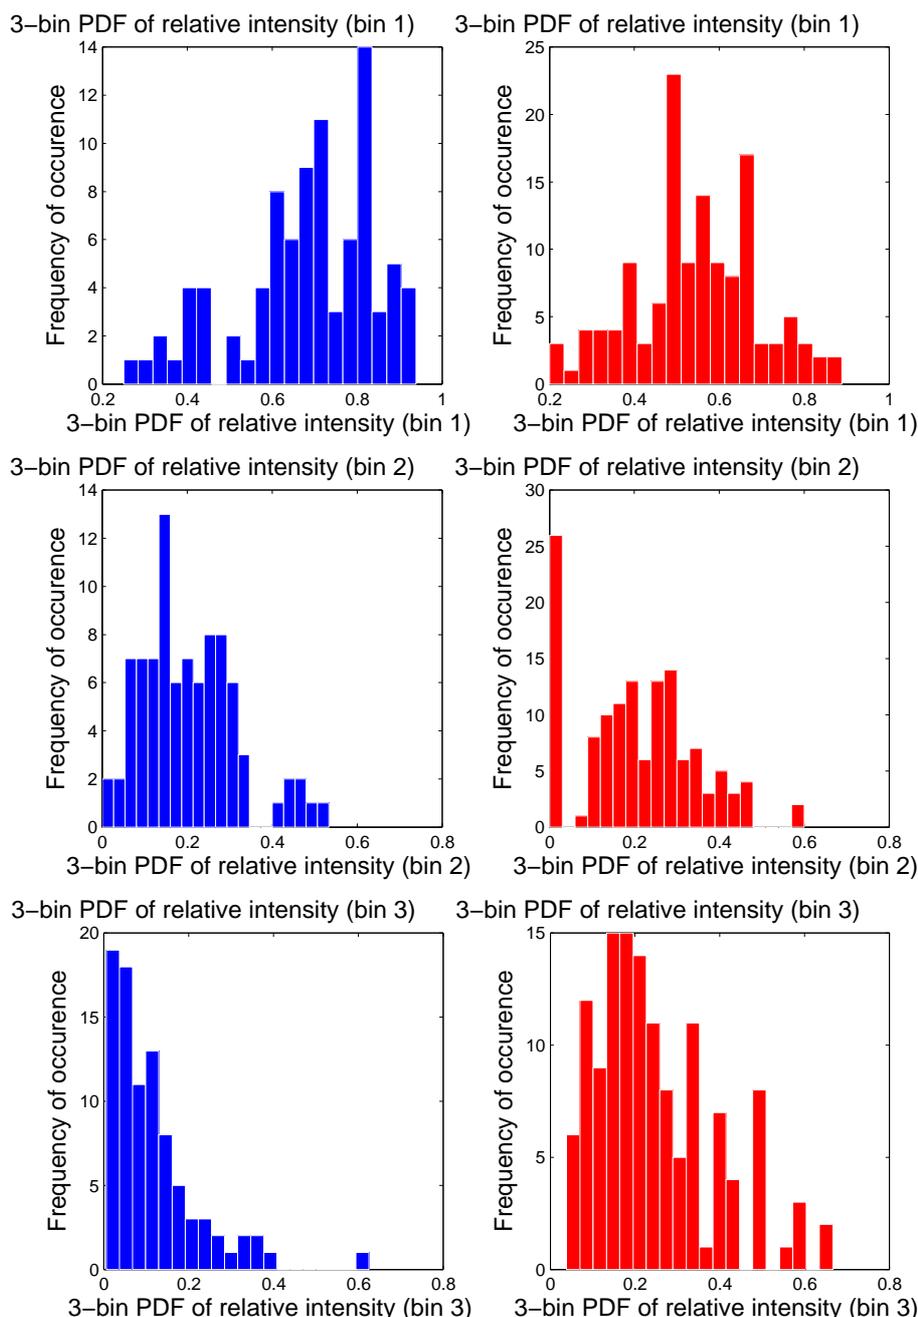


Figure 82: The feature type shown here is the 3-bin PDF of relative intensity in which the top, middle and bottom panels correspond, respectively, to bins 1, 2 and 3. The intensity voxels were chosen from those contained within the best fitting ellipsoidal volume associated with each target or clutter item. In each case the target and clutter features are shown on the left hand and right hand sides, respectively. The features were derived from 3D beamform products using data acquired by the BOSS platform at the AUV Fest 2008. A total of 132 clutter items and 89 targets were processed in generating these distributions.

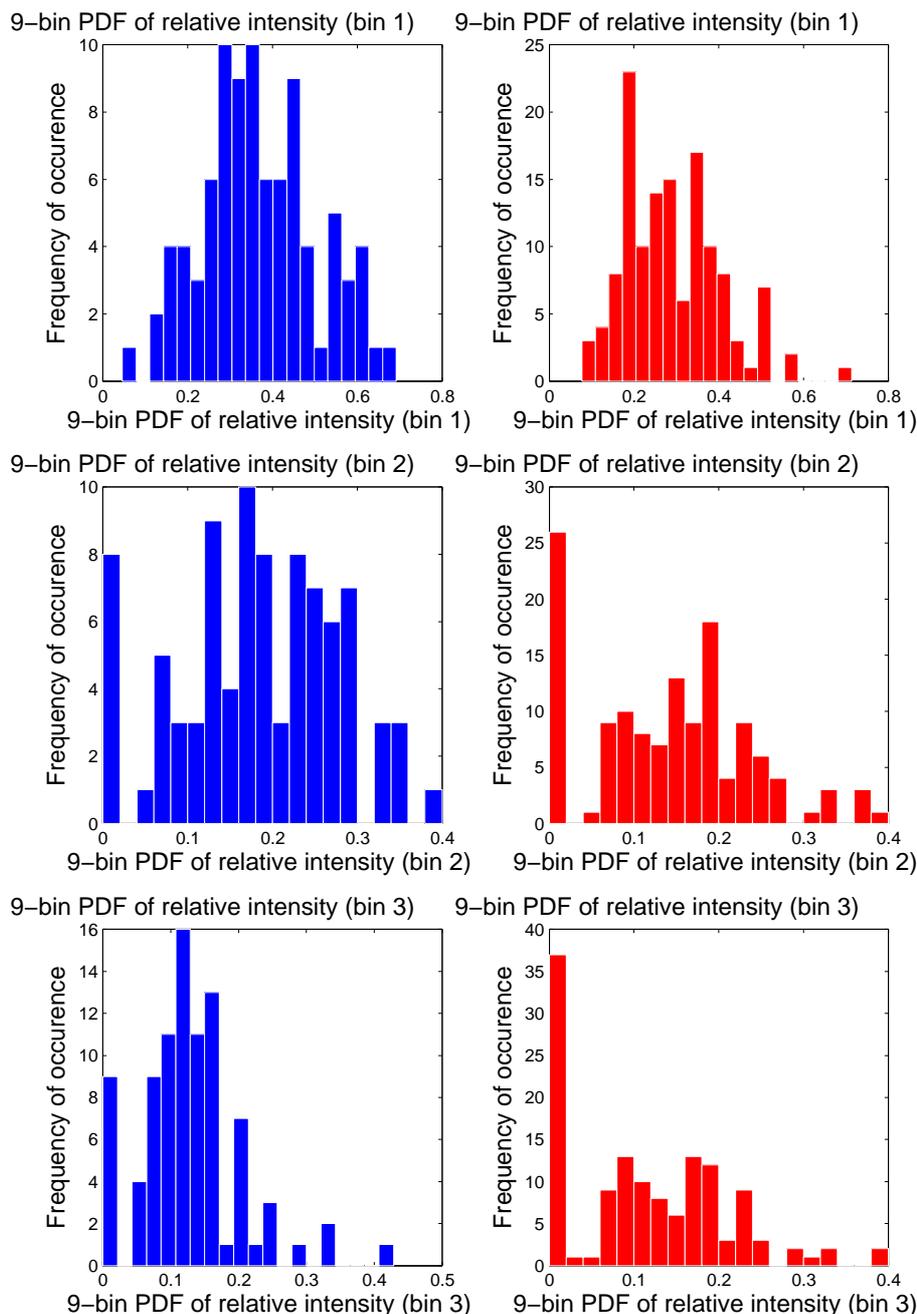


Figure 83: The feature type shown here are bins 1 to 3 of the 9-bin PDF of relative intensity in which the top, middle and bottom panels correspond, respectively, to bins 1, 2 and 3. The intensity voxels were chosen from those contained within the best fitting ellipsoidal volume associated with each target or clutter item. In each case the target and clutter features are shown on the left hand and right hand sides, respectively. The features were derived from 3D beamform products using data acquired by the BOSS platform at the AUV Fest 2008. A total of 132 clutter items and 89 targets were processed in generating these distributions.

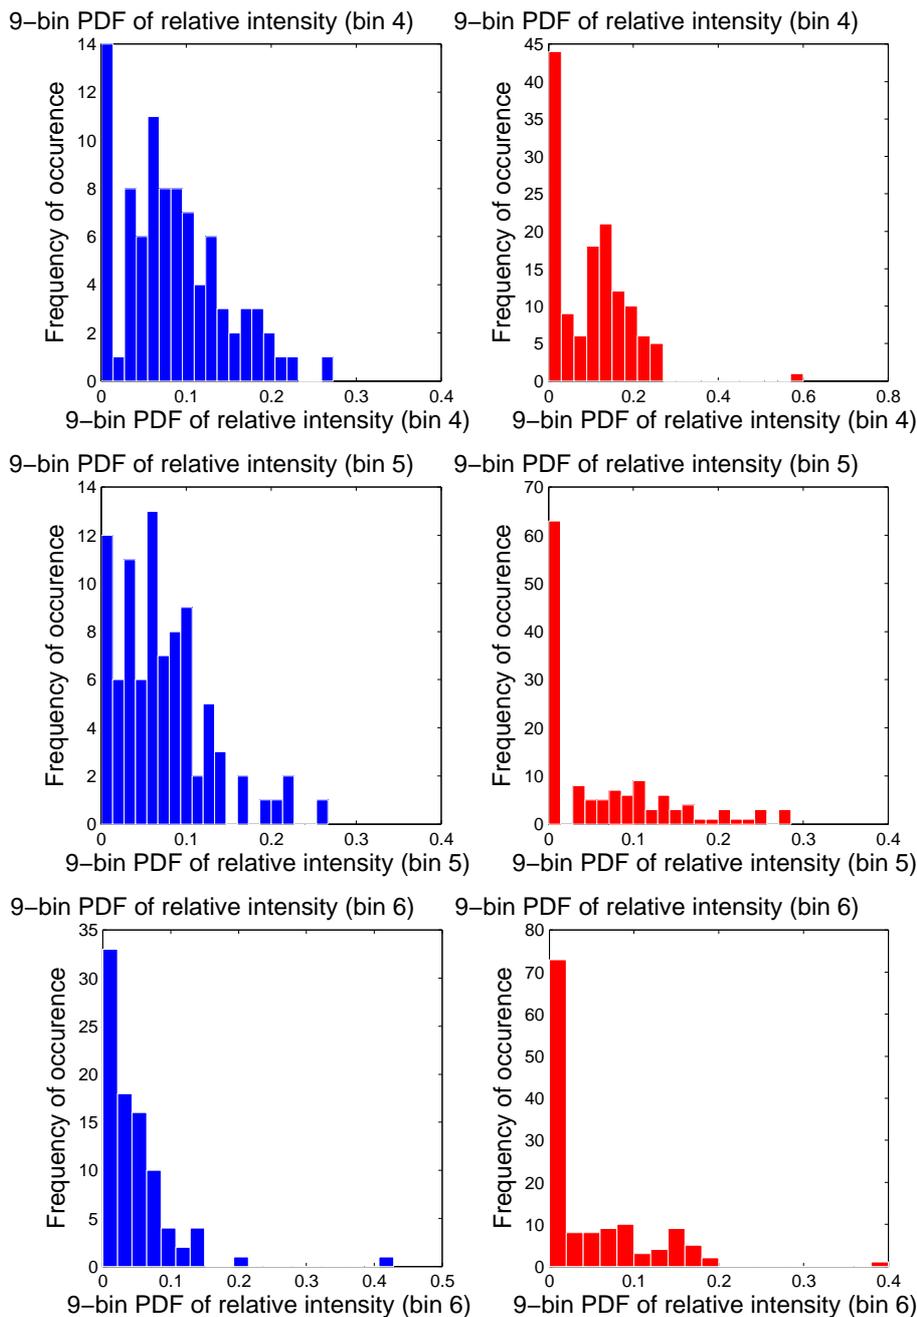


Figure 84: The feature type shown here are bins 4 to 6 of the 9-bin PDF of relative intensity in which the top, middle and bottom panels correspond, respectively, to bins 4, 5 and 6. The intensity voxels were chosen from those contained within the best fitting ellipsoidal volume associated with each target or clutter item. In each case the target and clutter features are shown on the left hand and right hand sides, respectively. The features were derived from 3D beamform products using data acquired by the BOSS platform at the AUV Fest 2008. A total of 132 clutter items and 89 targets were processed in generating these distributions.

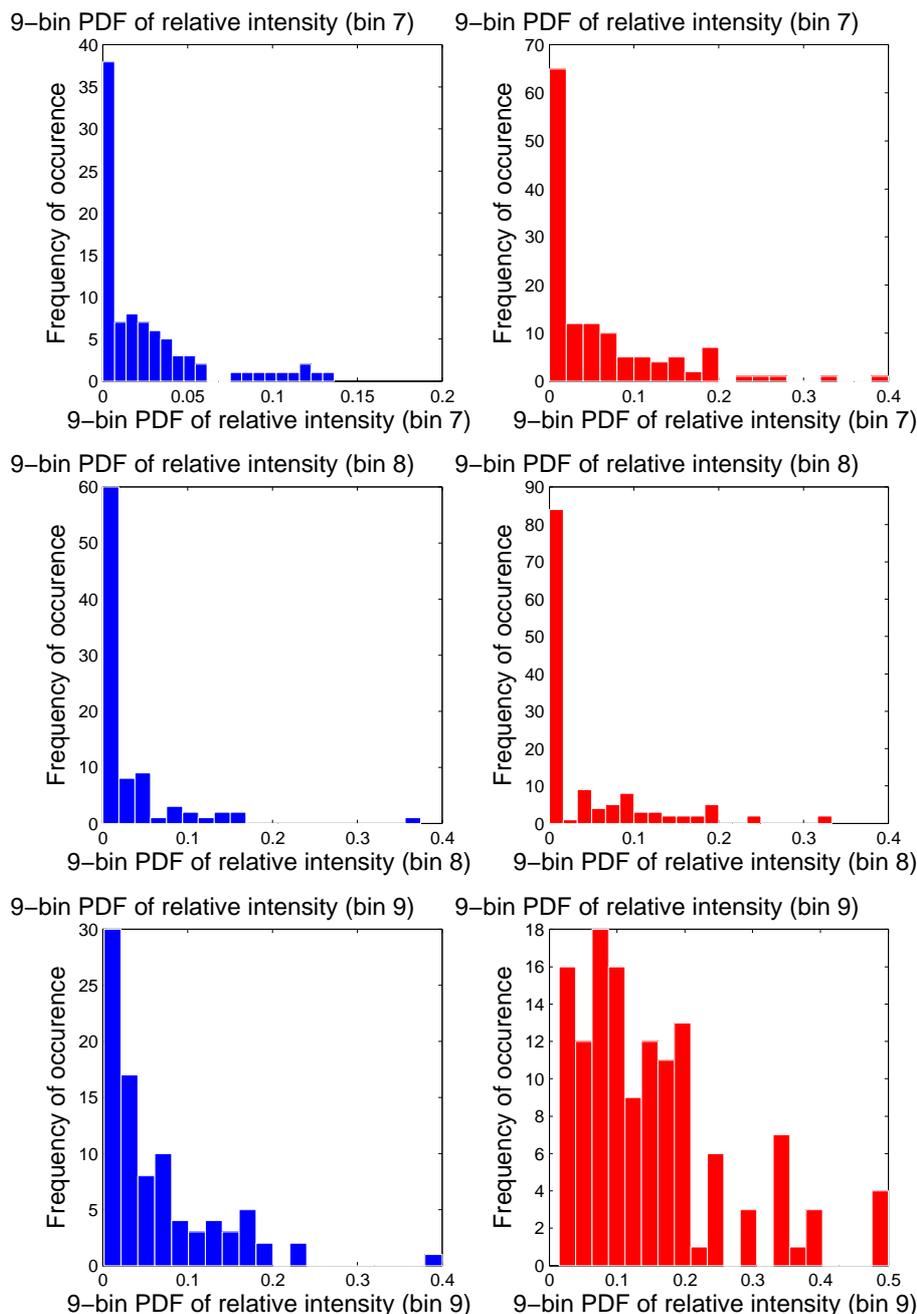


Figure 85: The feature type shown here are bins 7 to 9 of the 9-bin PDF of relative intensity in which the top, middle and bottom panels correspond, respectively, to bins 7, 8 and 9. The intensity voxels were chosen from those contained within the best fitting ellipsoidal volume associated with each target or clutter item. In each case the target and clutter features are shown on the left hand and right hand sides, respectively. The features were derived from 3D beamform products using data acquired by the BOSS platform at the AUV Fest 2008. A total of 132 clutter items and 89 targets were processed in generating these distributions.

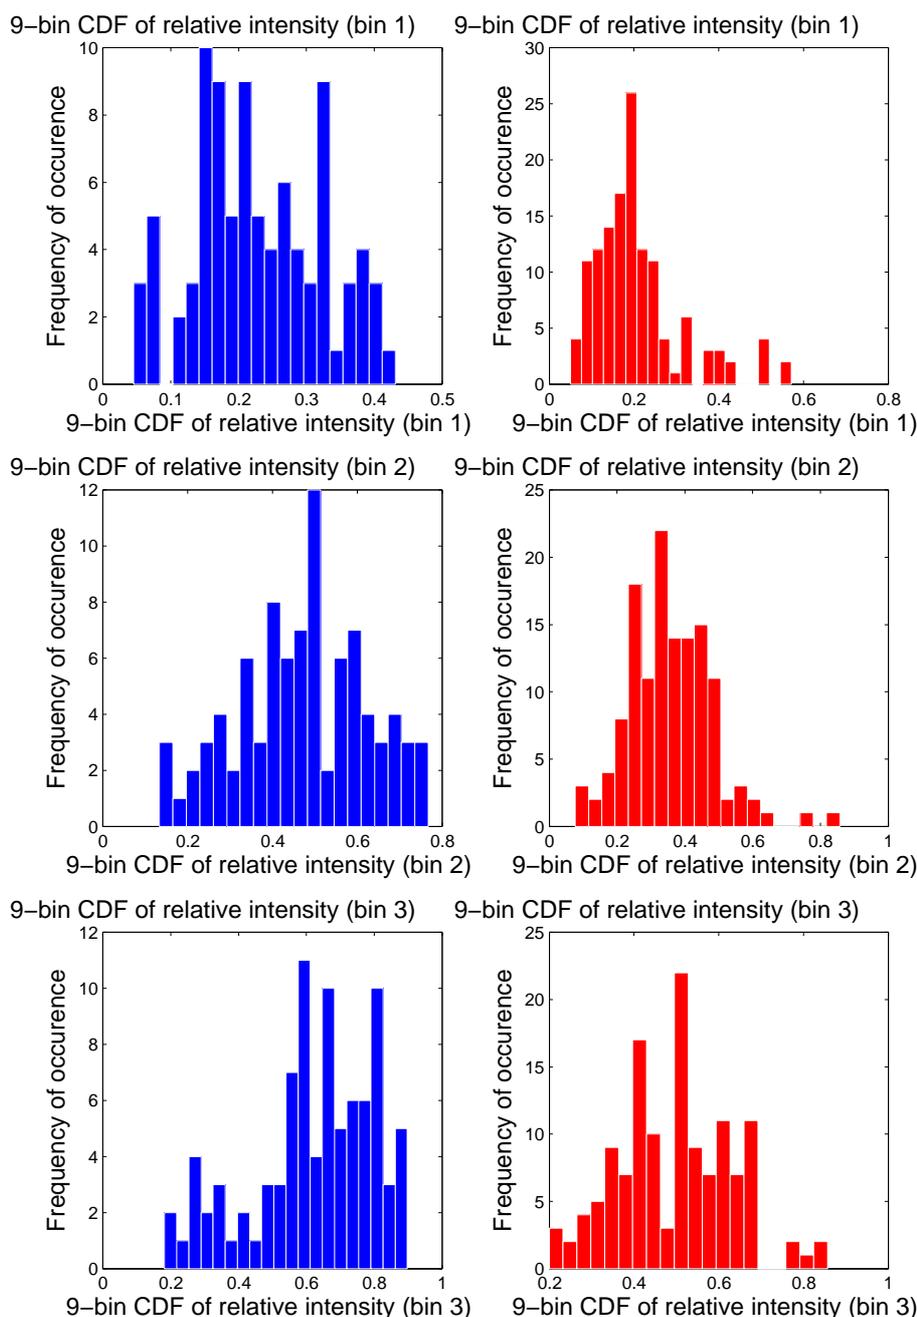


Figure 86: The feature type shown here are bins 1 to 3 of the 9-bin cumulative distribution function (CDF) of relative intensity in which the top, middle and bottom panels correspond, respectively, to bins 1, 2 and 3. The intensity voxels were chosen from those contained within the best fitting ellipsoidal volume associated with each target or clutter item. In each case the target and clutter features are shown on the left hand and right hand sides, respectively. The features were derived from 3D beamform products using data acquired by the BOSS platform at the AUV Fest 2008. A total of 132 clutter items and 89 targets were processed in generating these distributions.

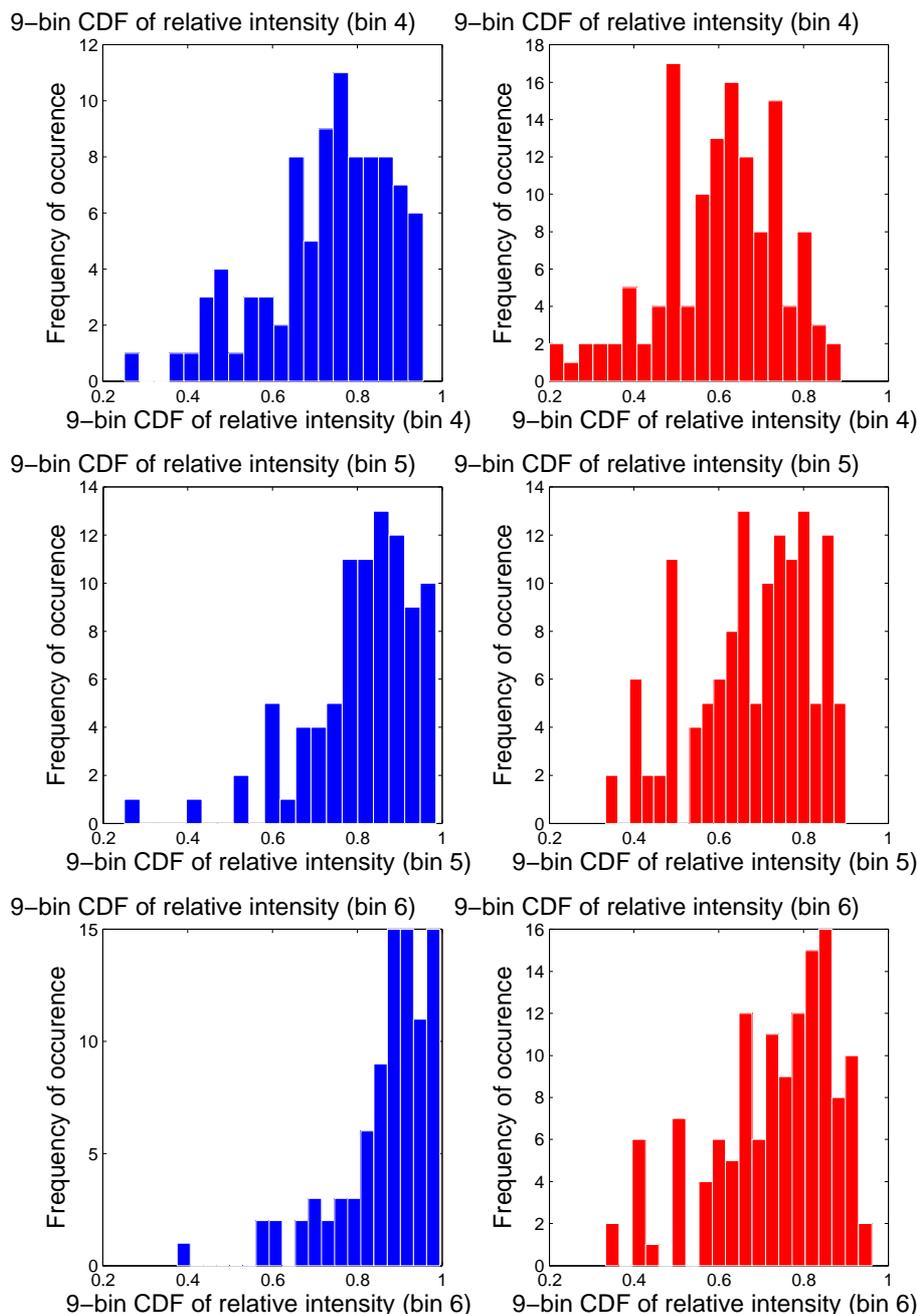


Figure 87: The feature type shown here are bins 4 to 6 of the 9-bin cumulative distribution function (CDF) of relative intensity in which the top, middle and bottom panels correspond, respectively, to bins 4, 5 and 6. The intensity voxels were chosen from those contained within the best fitting ellipsoidal volume associated with each target or clutter item. In each case the target and clutter features are shown on the left hand and right hand sides, respectively. The features were derived from 3D beamform products using data acquired by the BOSS platform at the AUV Fest 2008. A total of 132 clutter items and 89 targets were processed in generating these distributions.

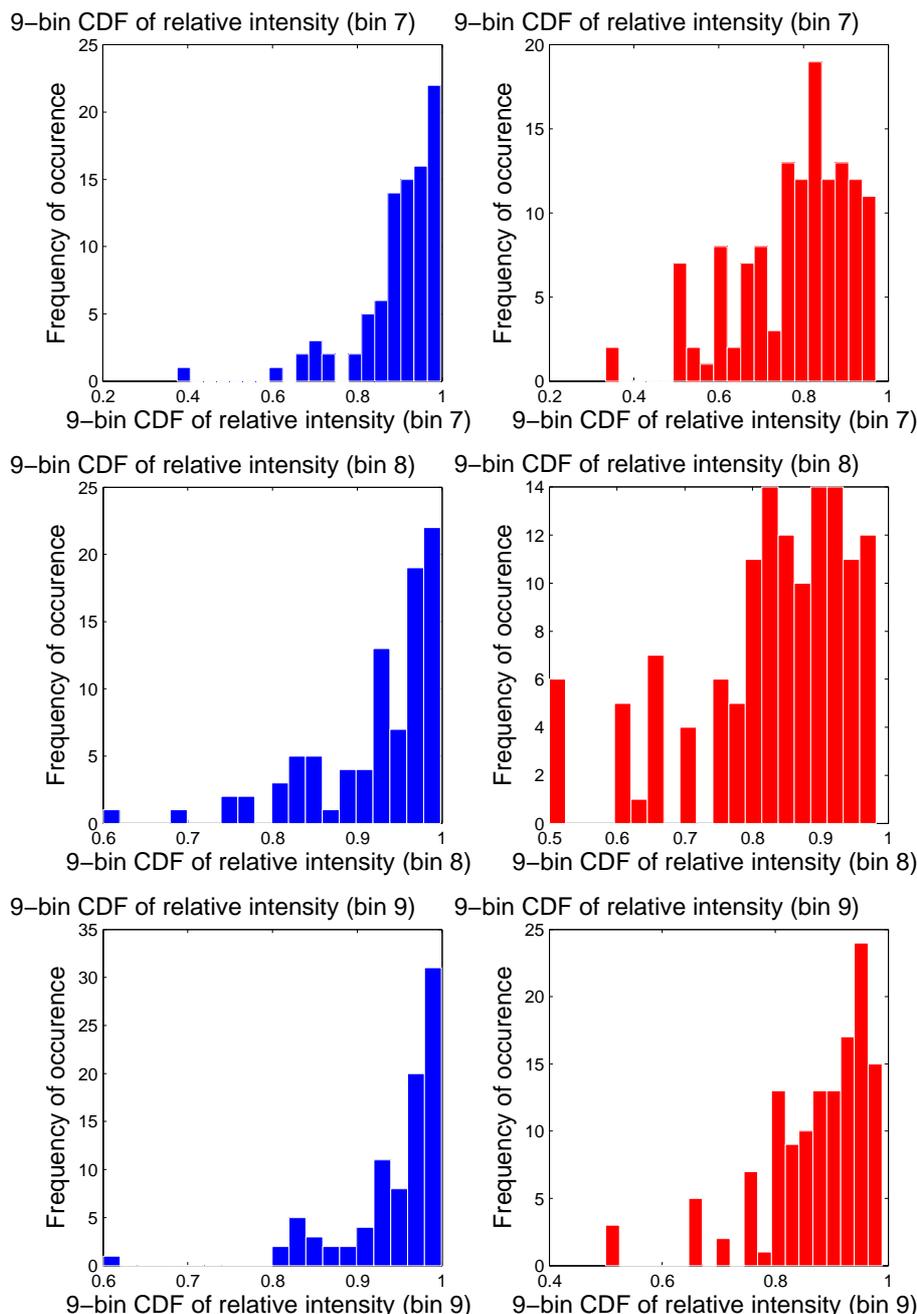


Figure 88: The feature type shown here are bins 7 to 9 of the 9–bin cumulative distribution function (CDF) of relative intensity in which the top, middle and bottom panels correspond, respectively, to bins 7, 8 and 9. The intensity voxels were chosen from those contained within the best fitting ellipsoidal volume associated with each target or clutter item. In each case the target and clutter features are shown on the left hand and right hand sides, respectively. The features were derived from 3D beamform products using data acquired by the BOSS platform at the AUV Fest 2008. A total of 132 clutter items and 89 targets were processed in generating these distributions.

A.4 Feature pair analysis of target/clutter features for AUV Fest 2008 data

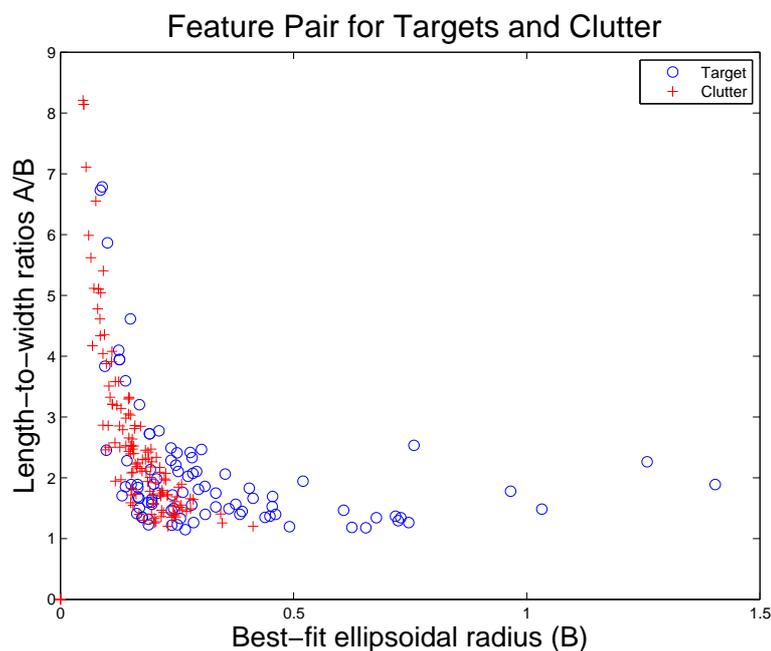


Figure 89: Scatter plot for the features *ellipsoid radius B* versus the *length to width ratio A/B* for the BOSS AUV FEST 2008 data. Features derived from the targets and clutter in this data set are indicated by the circles and plus signs, respectively.

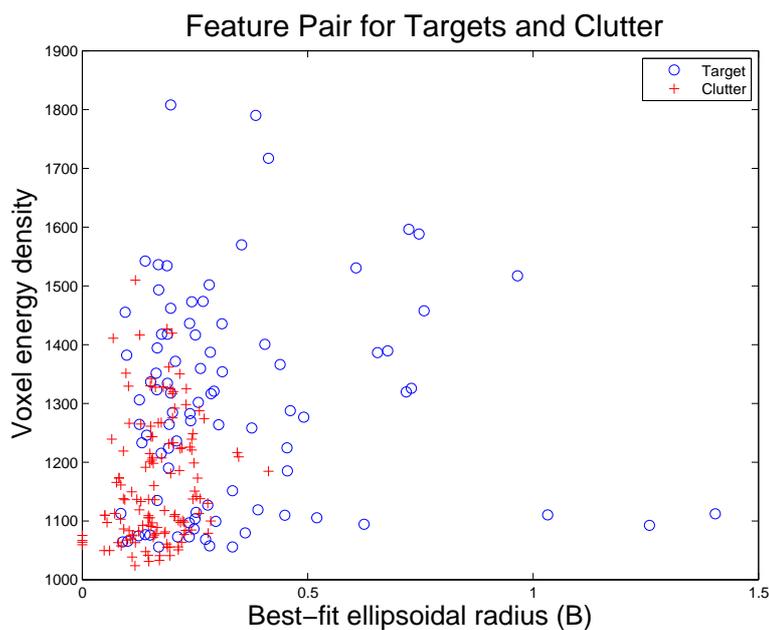


Figure 90: Scatter plot for the features *ellipsoid radius B* versus the *voxel energy density* for the BOSS AUV FEST 2008 data. Features derived from the targets and clutter in this data set are indicated by the circles and plus signs, respectively.

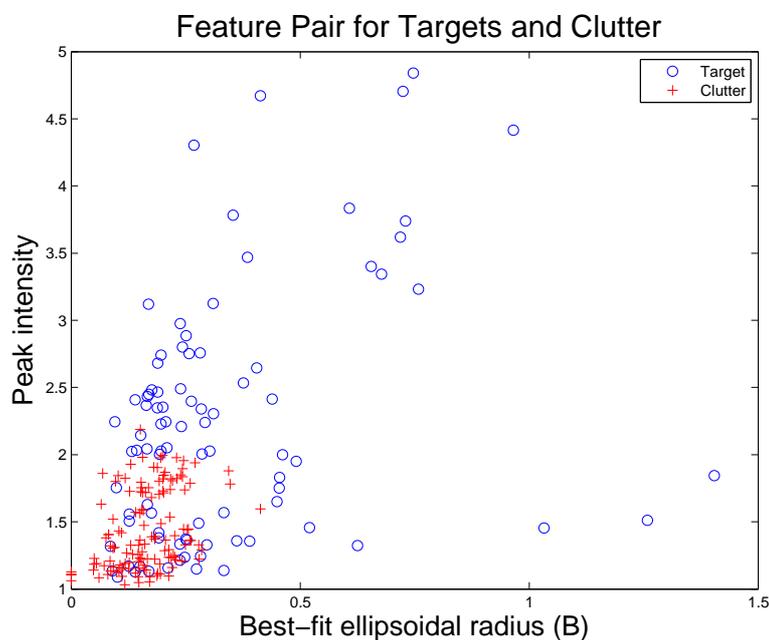


Figure 91: Scatter plot for the features *ellipsoid radius B* versus the *peak intensity* for the BOSS AUV FEST 2008 data. Features derived from the targets and clutter in this data set are indicated by the circles and plus signs, respectively.

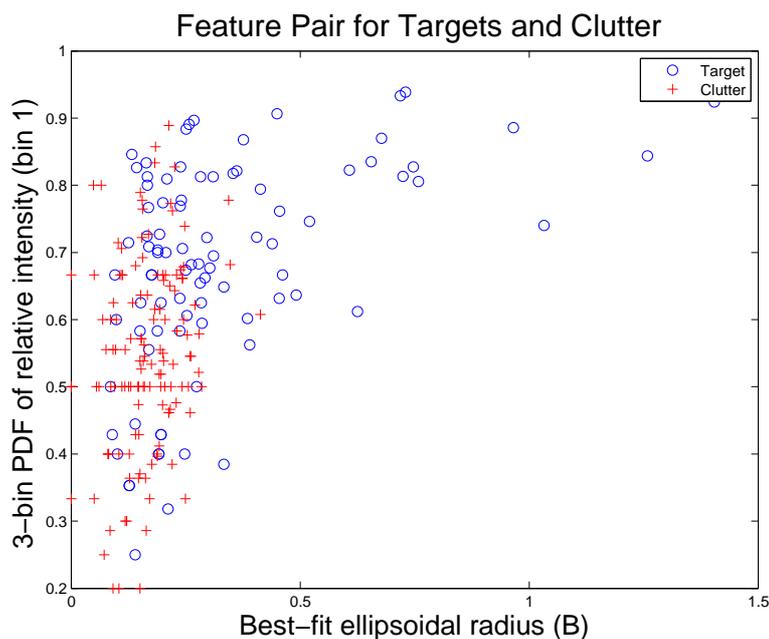


Figure 92: Scatter plot for the features *ellipsoid radius B* versus the *first bin of the 3-bin PDF of the relative intensity* for the BOSS AUV FEST 2008 data. Features derived from the targets and clutter in this data set are indicated by the circles and plus signs, respectively.

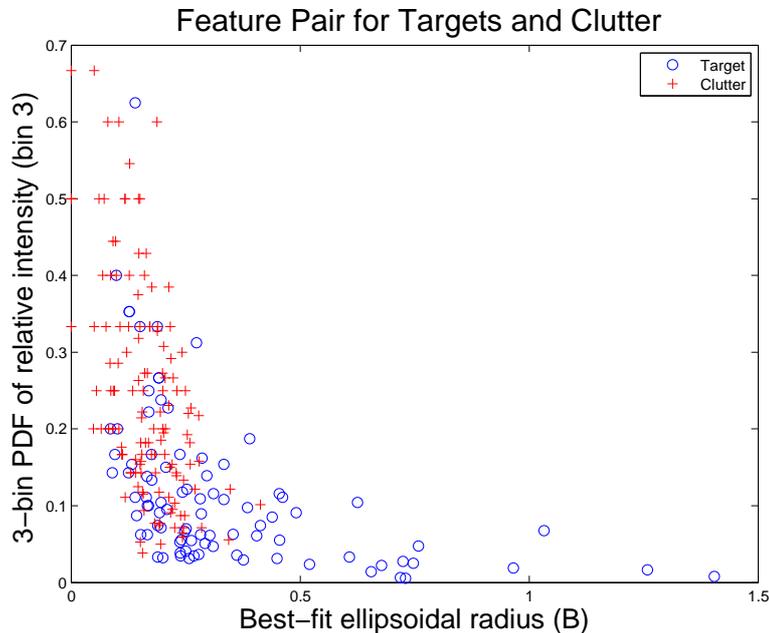


Figure 93: Scatter plot for the features *ellipsoid radius B* versus the *third bin of the 3-bin PDF of the relative intensity* for the BOSS AUV FEST 2008 data. Features derived from the targets and clutter in this data set are indicated by the circles and plus signs, respectively.

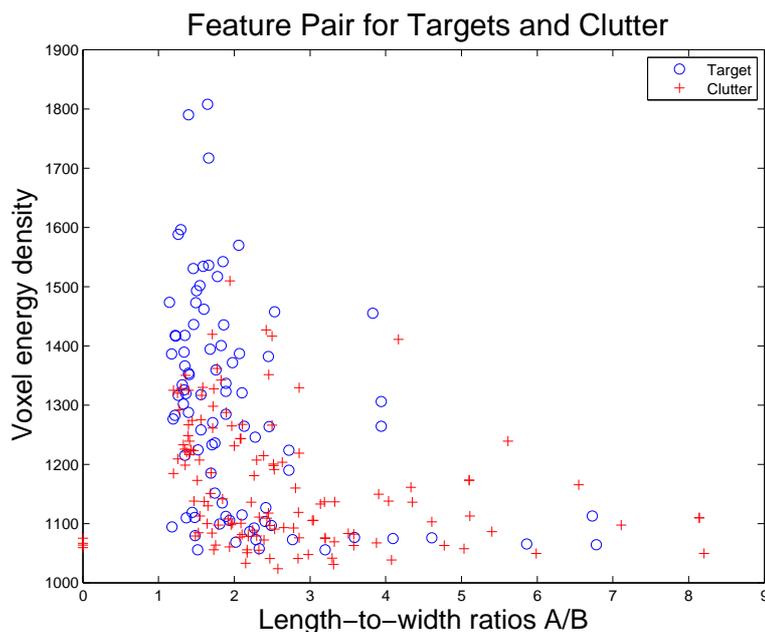


Figure 94: Scatter plot for the features *length to width ratio* A/B versus the *voxel energy density* for the BOSS AUV FEST 2008 data. Features derived from the targets and clutter in this data set are indicated by the circles and plus signs, respectively.

Consistent with Fig. 98, voxel energy densities for targets have a right-skewed component relative to that of clutter.

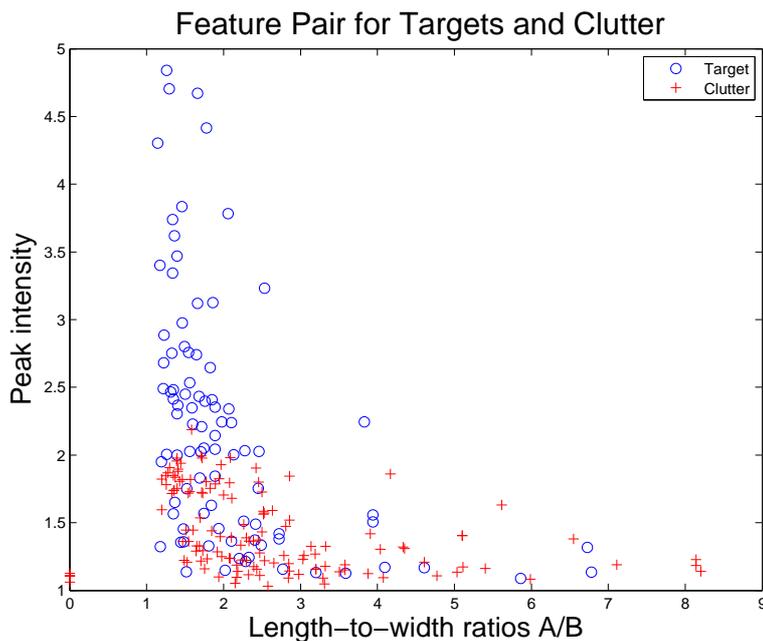


Figure 95: Scatter plot for the features *length to width ratio* A/B versus the *peak intensity* for the BOSS AUV FEST 2008 data. Features derived from the targets and clutter in this data set are indicated by the circles and plus signs, respectively.

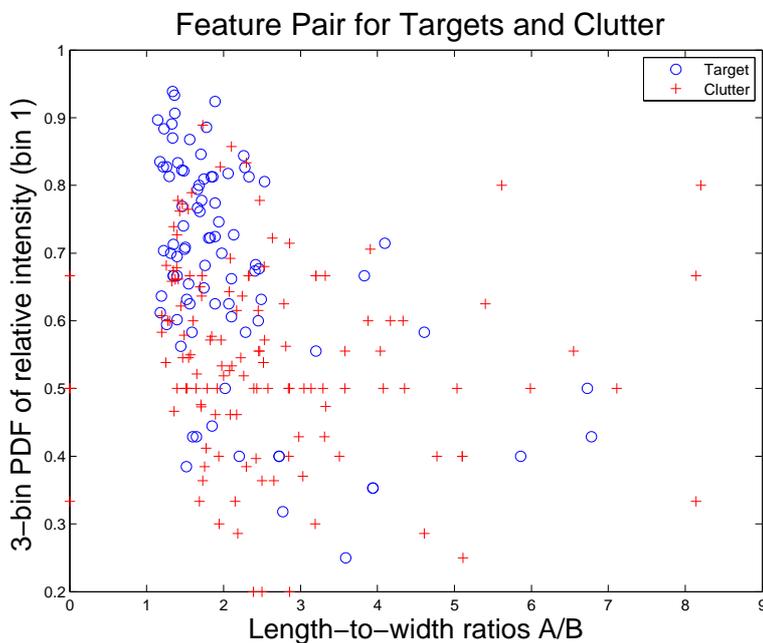


Figure 96: Scatter plot for the features *length to width ratio* A/B versus the *first bin of the 3-bin PDF of the relative intensity* for the BOSS AUV FEST 2008 data. Features derived from the targets and clutter in this data set are indicated by the circles and plus signs, respectively.

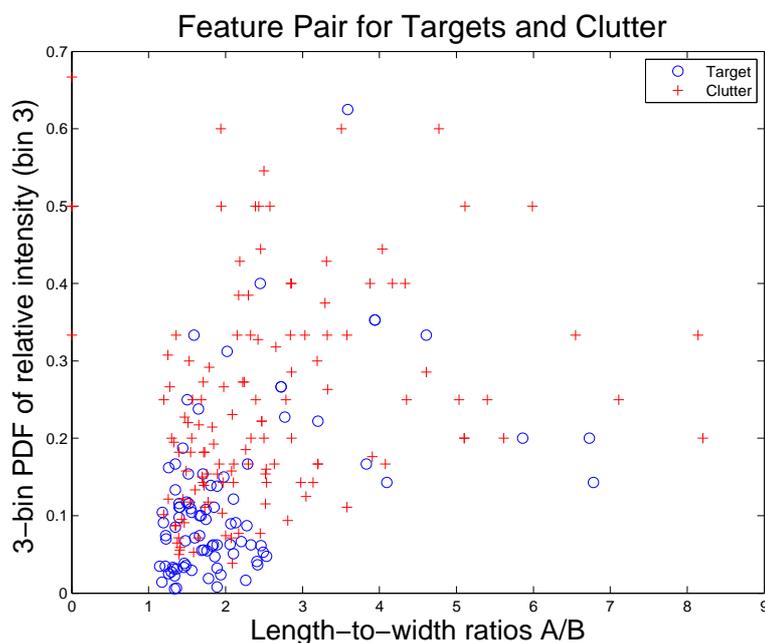


Figure 97: Scatter plot for the features *length to width ratio A/B* versus the *third bin of the 3-bin PDF of the relative intensity* for the BOSS AUV FEST 2008 data. Features derived from the targets and clutter in this data set are indicated by the circles and plus signs, respectively.

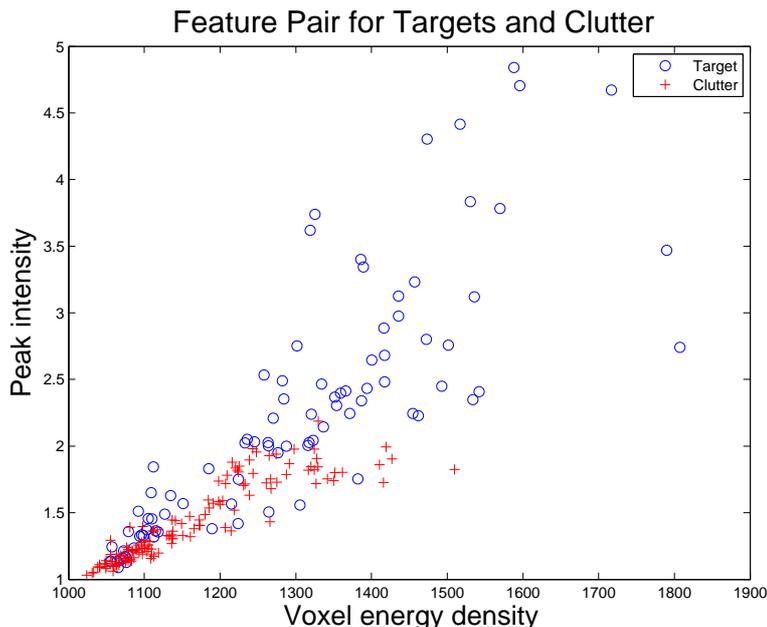


Figure 98: Scatter plot for the features *voxel energy density* versus the *peak intensity* for the BOSS AUV FEST 2008 data. Features derived from the targets and clutter in this data set are indicated by the circles and plus signs, respectively. This feature appears to be a useful discriminant for target versus clutter for at least some targets. It is clear that simultaneously strong peak intensities and voxel energy densities are strongly coordinated to targets. This is because actual targets provide a strong scattering contrast relative to natural and man-made, non-target clutter.

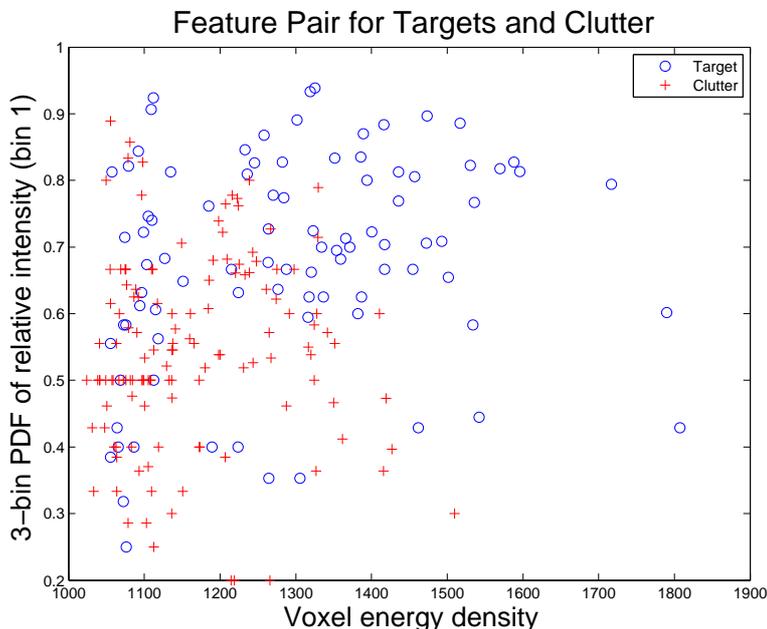


Figure 99: Scatter plot for the features *voxel energy density* versus the *first bin of the 3-bin PDF of the relative intensity* for the BOSS AUV FEST 2008 data. Features derived from the targets and clutter in this data set are indicated by the circles and plus signs, respectively.

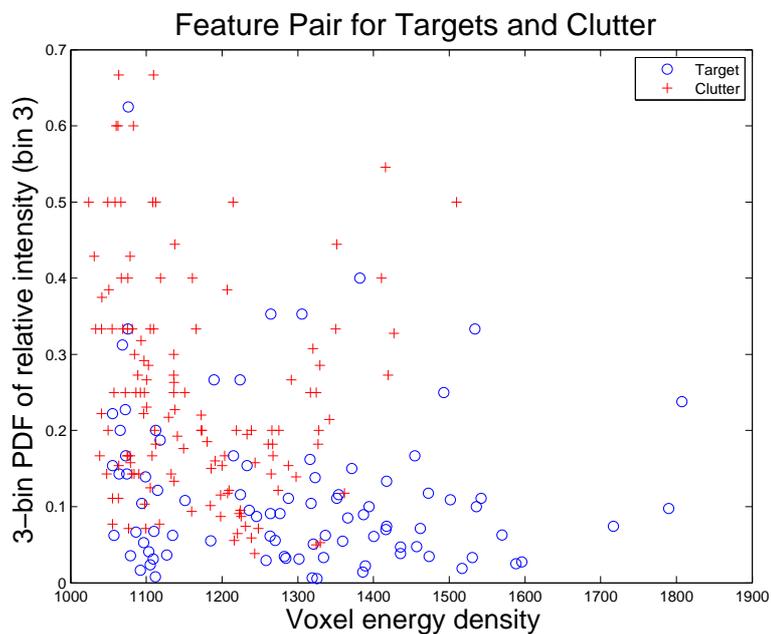


Figure 100: Scatter plot for the features *voxel energy density* versus the *third bin of the 3-bin PDF of the relative intensity* for the BOSS AUV FEST 2008 data. Features derived from the targets and clutter in this data set are indicated by the circles and plus signs, respectively.

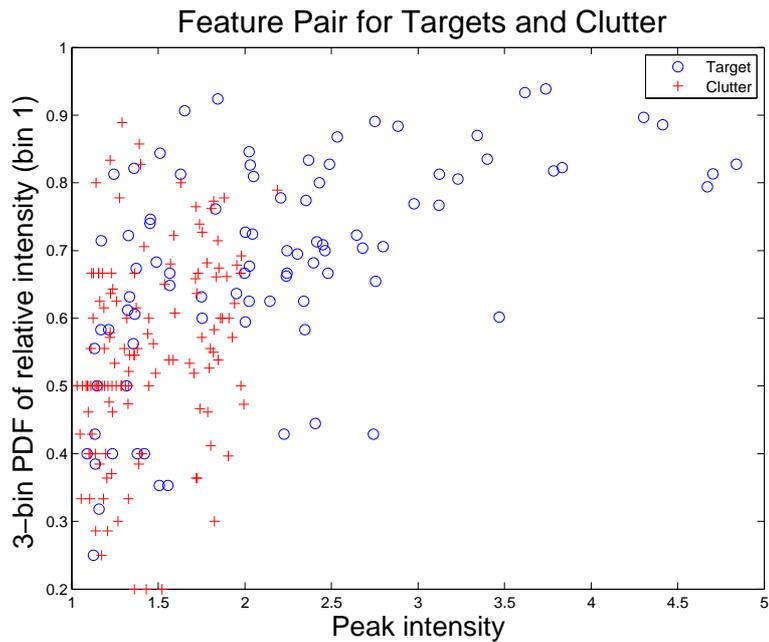


Figure 101: Scatter plot for the features *peak intensity* versus the *first bin of the 3-bin PDF of the relative intensity* for the BOSS AUV FEST 2008 data. Features derived from the targets and clutter in this data set are indicated by the circles and plus signs, respectively.

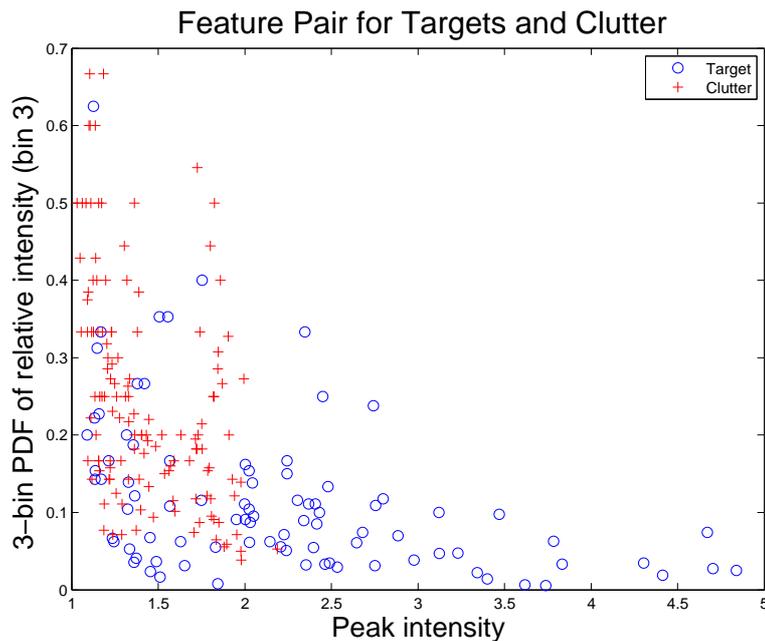


Figure 102: Scatter plot for the features *peak intensity* versus the *third bin of the 3-bin PDF of the relative intensity* for the BOSS AUV FEST 2008 data. Features derived from the targets and clutter in this data set are indicated by the circles and plus signs, respectively.

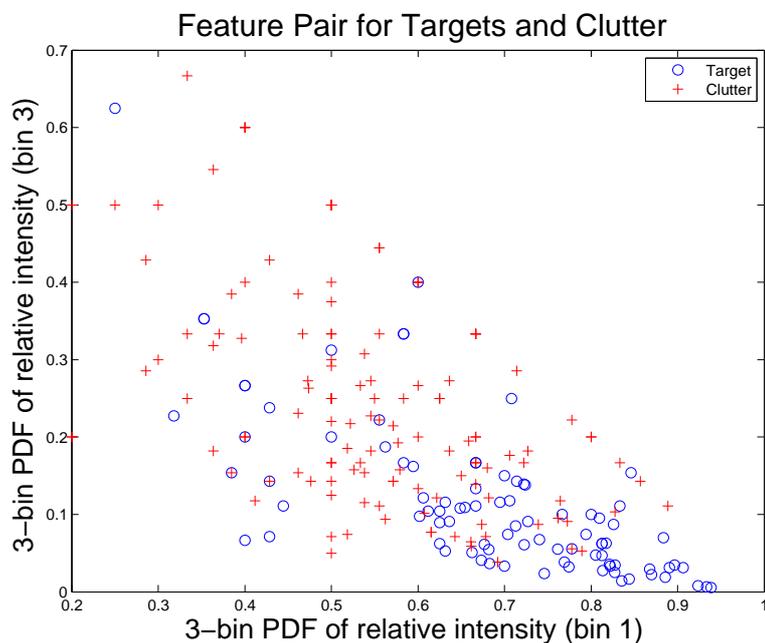


Figure 103: Scatter plot for the features *first bin of the 3-bin PDF of the relative intensity* versus the *third bin of the 3-bin PDF of the relative intensity* for the BOSS AUV FEST 2008 data. Features derived from the targets and clutter in this data set are indicated by the circles and plus signs, respectively.

A.5 Feature triplet analysis of target/clutter features for AUV Fest 2008 data

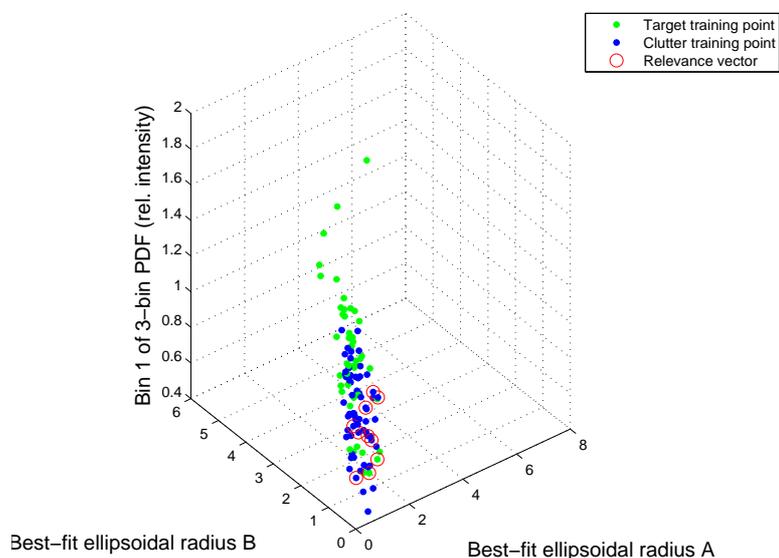


Figure 104: 3D scatter plot showing training points and optimal relevance vectors for the feature triplet A (*ellipsoid radius*), B (*ellipsoid radius*) and *first bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles.

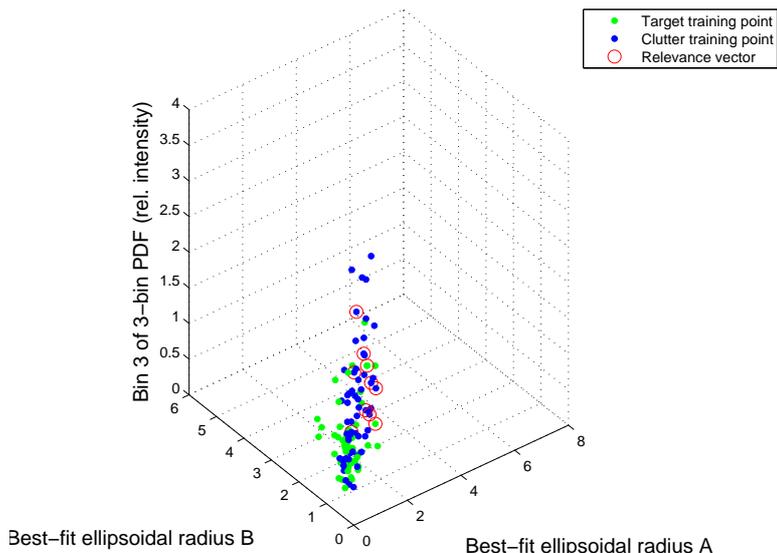


Figure 105: 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *A (ellipsoid radii), B (ellipsoid radii) and last bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles.

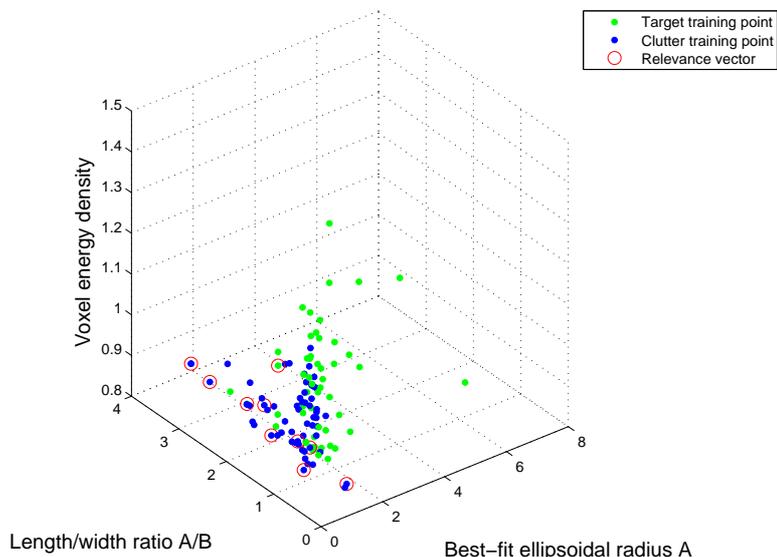


Figure 106: 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *A (ellipsoid radii), Length to width ratio A/B and Voxel energy density*. and *first bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles.

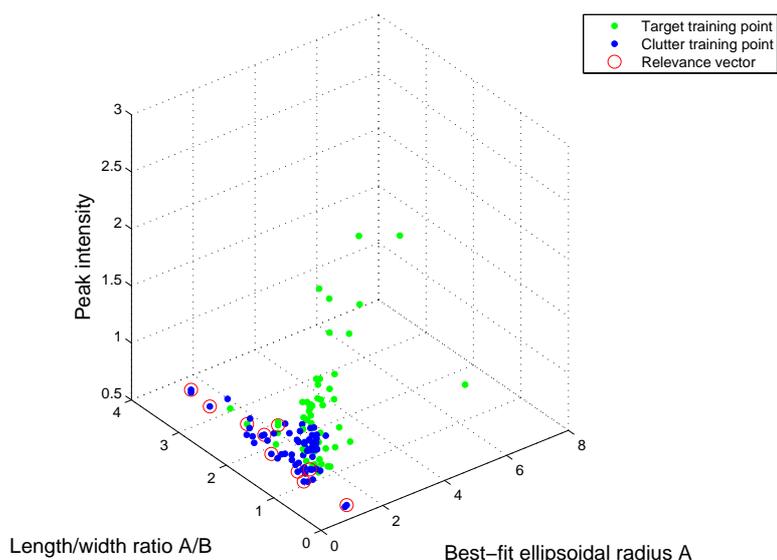


Figure 107: 3D scatter plot showing training points and optimal relevance vectors for the feature triplet A (*ellipsoid radii*), *Length to width ratio A/B*, and *Peak intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles.

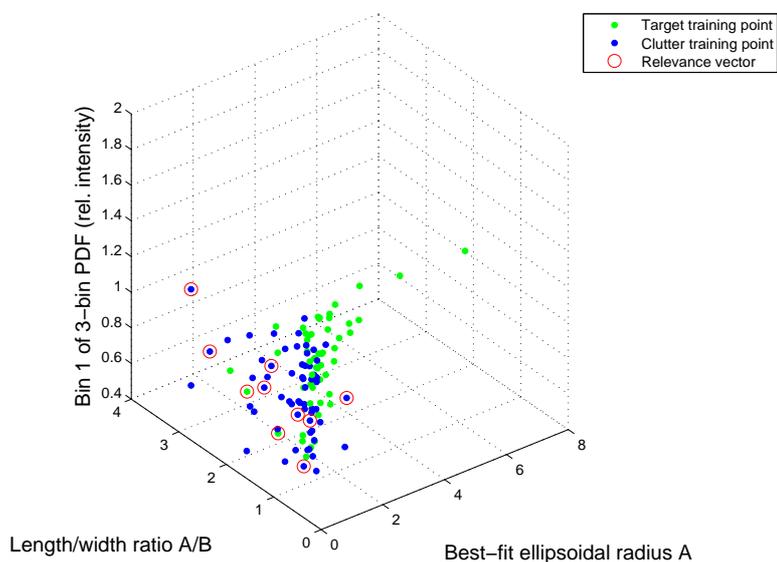


Figure 108: 3D scatter plot showing training points and optimal relevance vectors for the feature triplet A (*ellipsoid radii*), *Length to width ratio A/B*, and *first bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles.

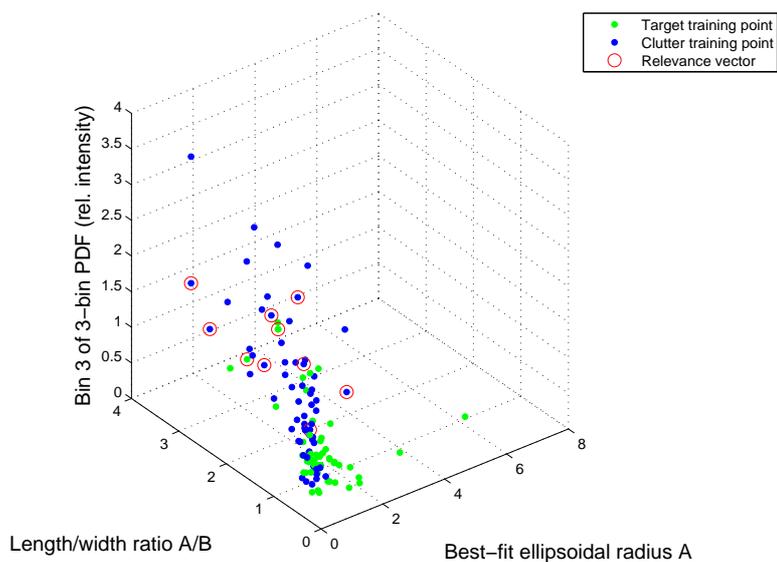


Figure 109: 3D scatter plot showing training points and optimal relevance vectors for the feature triplet A (*ellipsoid radii*), *Length to width ratio A/B*, and *last bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles.

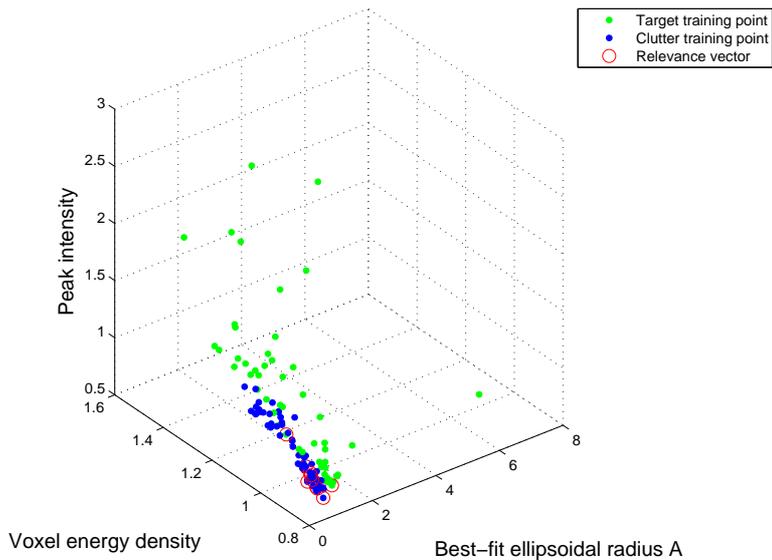


Figure 110: 3D scatter plot showing training points and optimal relevance vectors for the feature triplet A (*ellipsoid radii*), *Voxel energy density* and *Peak intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles.

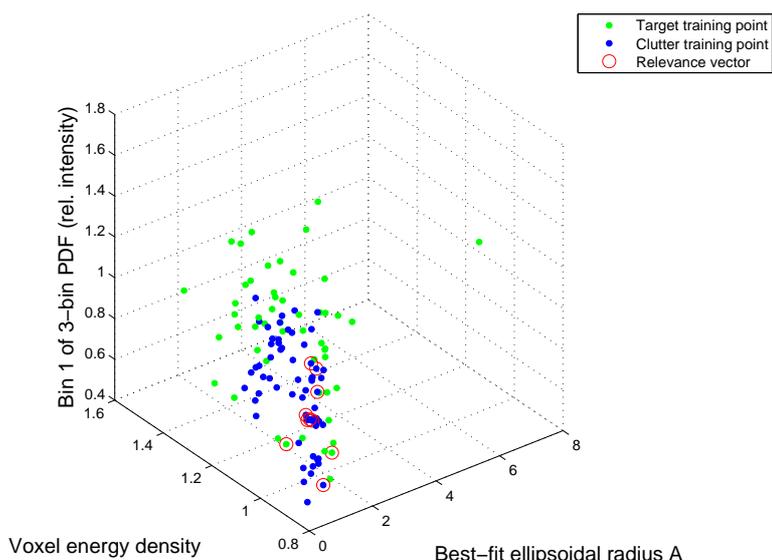


Figure 111: 3D scatter plot showing training points and optimal relevance vectors for the feature triplet A (*ellipsoid radii*), *Voxel energy density* and *first bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles.

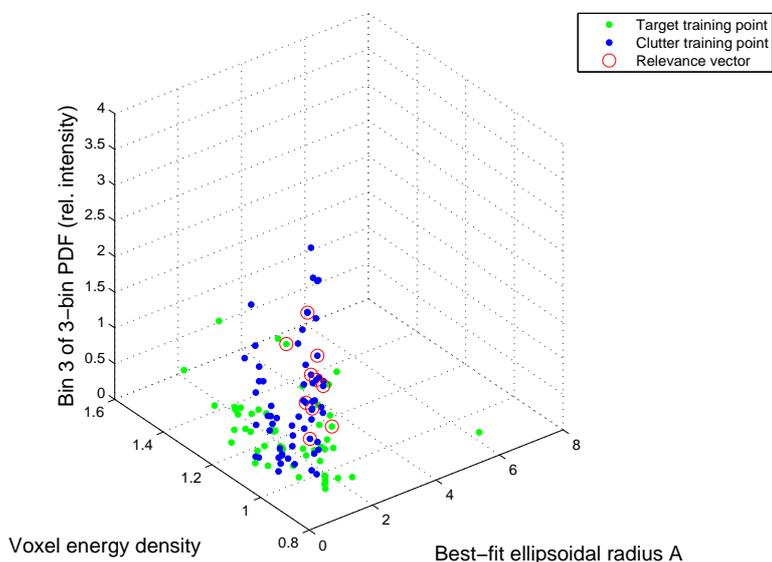


Figure 112: 3D scatter plot showing training points and optimal relevance vectors for the feature triplet A (*ellipsoid radii*), *Voxel energy density* and *last bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles.

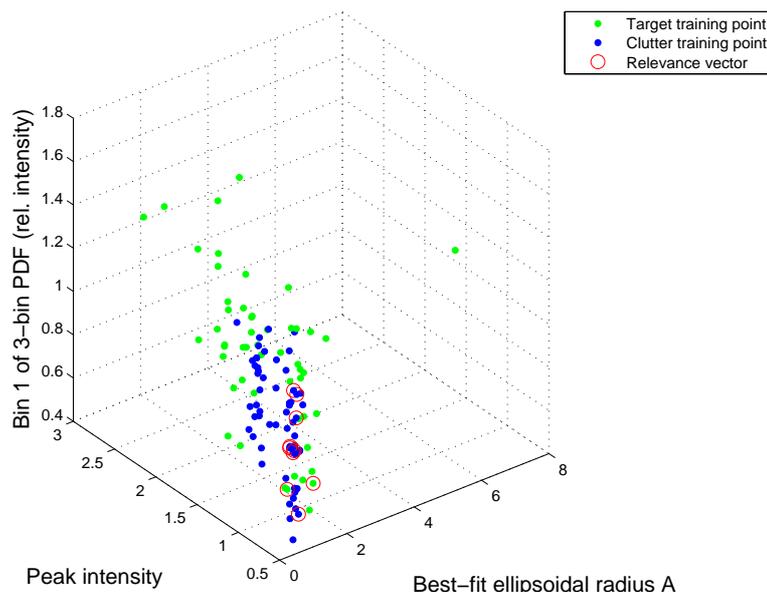


Figure 113: 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *A* (ellipsoid radii), *Peak intensity* and *first bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles.

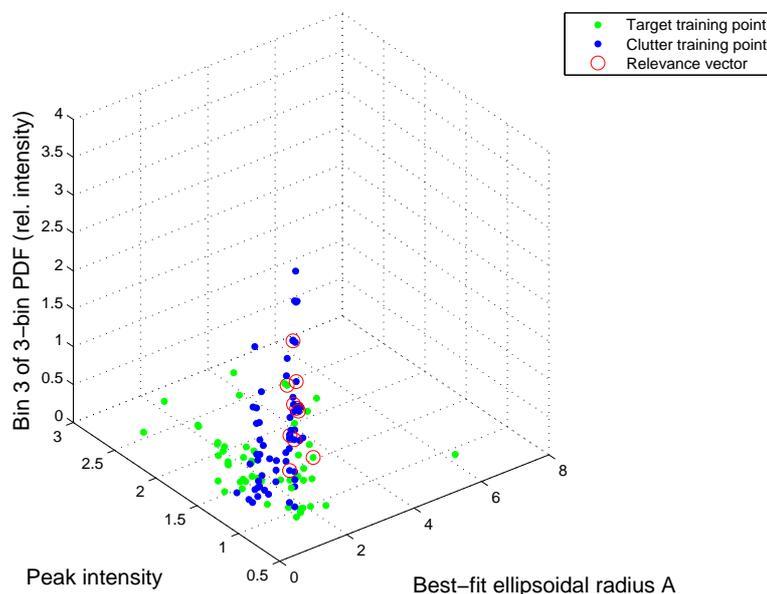


Figure 114: 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *A* (ellipsoid radii), *Peak intensity* and *last bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles.

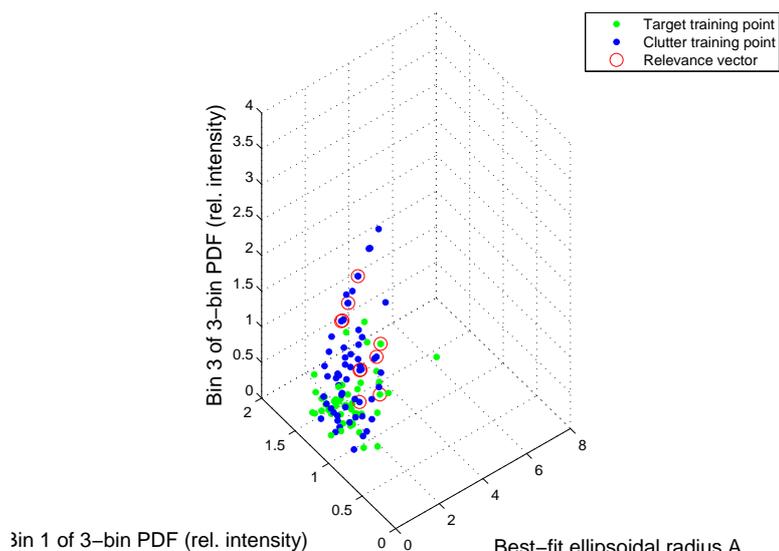


Figure 115: 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *A* (ellipsoid radii), first bin in 3-bin probability distribution of relative intensity and last bin in 3-bin probability distribution of relative intensity. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles.

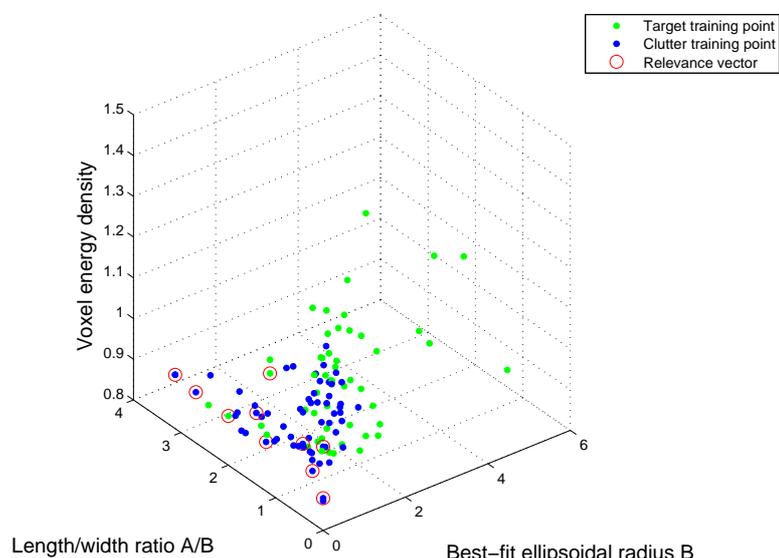


Figure 116: 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *B* (ellipsoid radii), Length to width ratio *A/B* and Voxel energy density. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles.

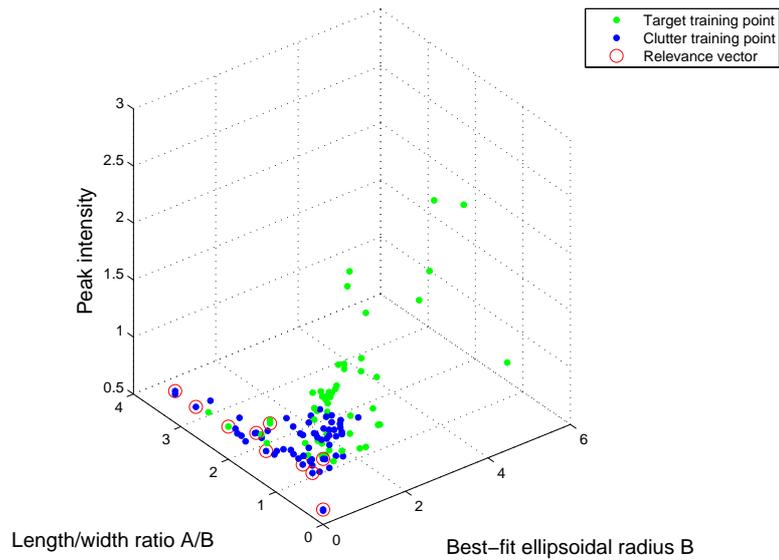


Figure 117: 3D scatter plot showing training points and optimal relevance vectors for the feature triplet B (ellipsoid radii), Length to width ratio A/B and Peak intensity. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles.

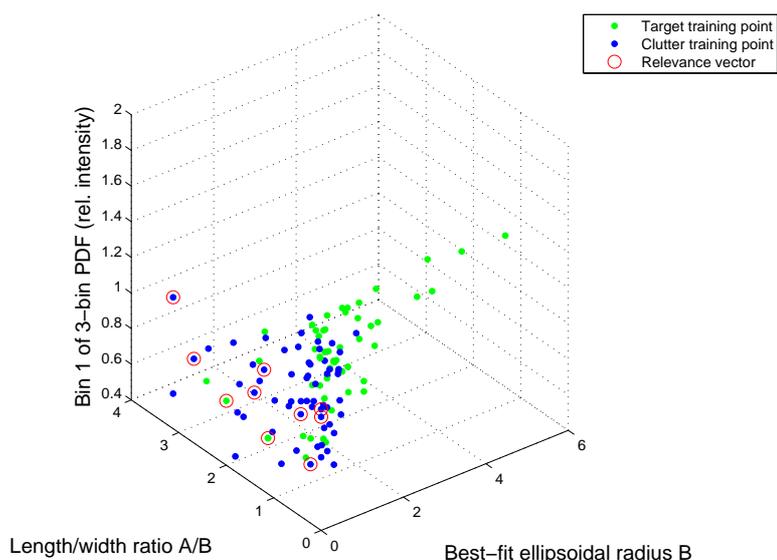


Figure 118: 3D scatter plot showing training points and optimal relevance vectors for the feature triplet B (ellipsoid radii), Length to width ratio A/B and first bin in 3-bin probability distribution of relative intensity. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles.

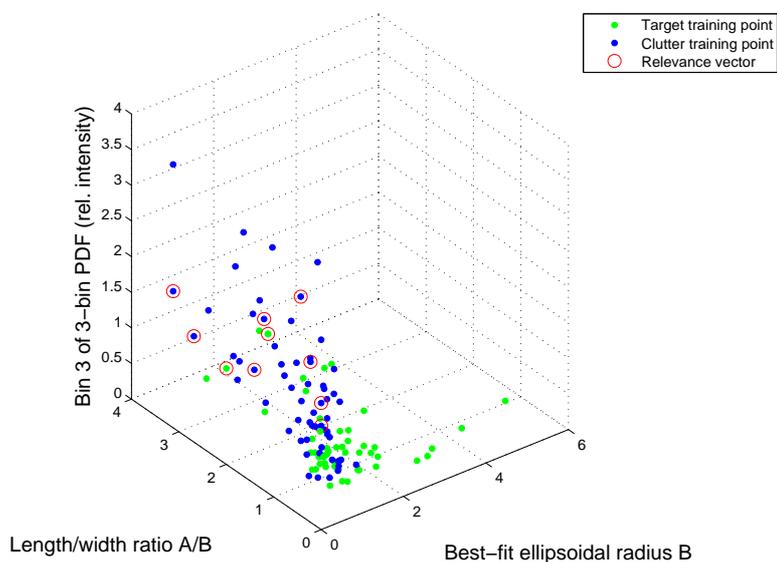


Figure 119: 3D scatter plot showing training points and optimal relevance vectors for the feature triplet B (ellipsoid radii), Length to width ratio A/B and last bin in 3-bin probability distribution of relative intensity. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles.

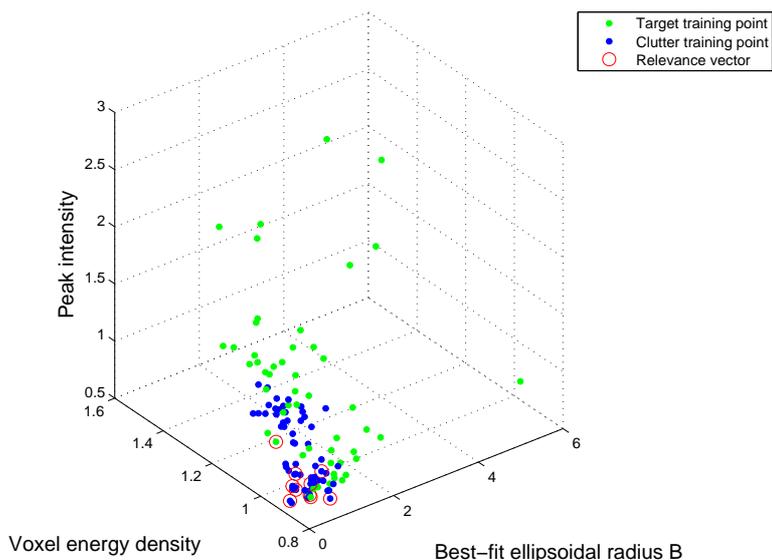


Figure 120: 3D scatter plot showing training points and optimal relevance vectors for the feature triplet B (ellipsoid radii), *Voxel energy density* and *Peak intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles.

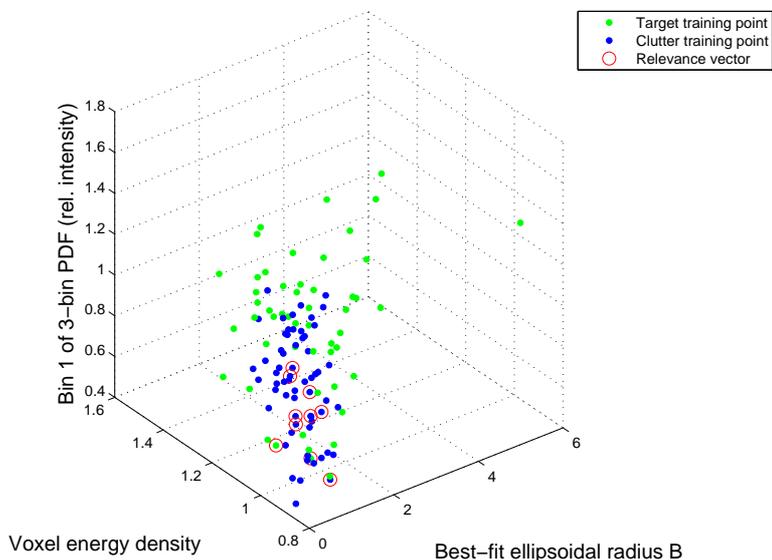


Figure 121: 3D scatter plot showing training points and optimal relevance vectors for the feature triplet B (ellipsoid radii), *Voxel energy density* and *first bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles.

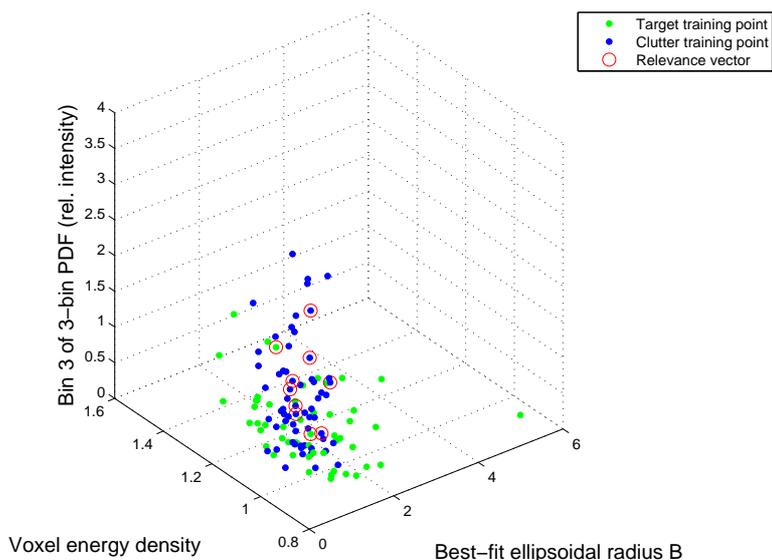


Figure 122: 3D scatter plot showing training points and optimal relevance vectors for the feature triplet B (ellipsoid radii), *Voxel energy density* and *last bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles.

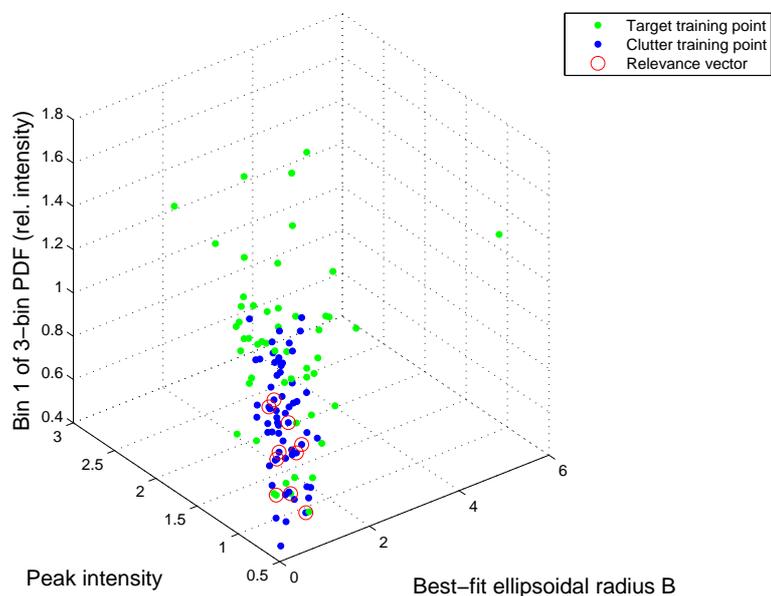


Figure 123: 3D scatter plot showing training points and optimal relevance vectors for the feature triplet B (ellipsoid radii), and *Peak intensity* and *first bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles.

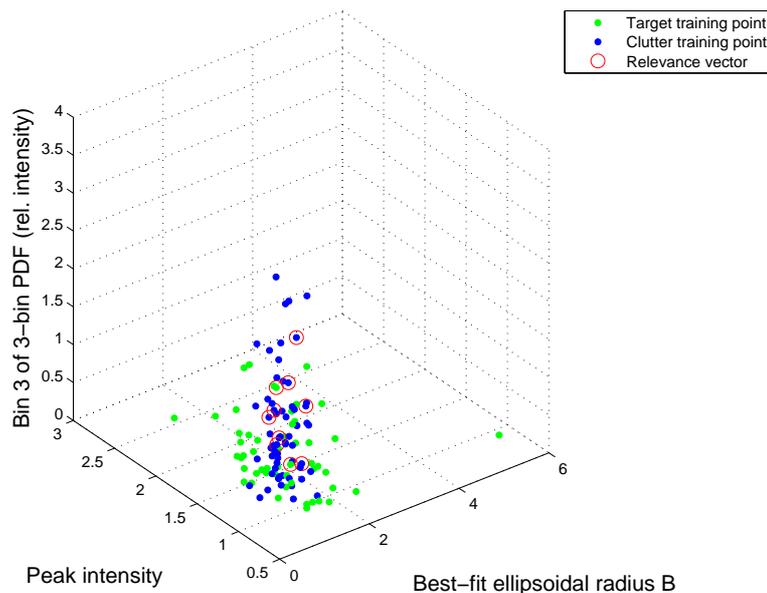


Figure 124: 3D scatter plot showing training points and optimal relevance vectors for the feature triplet B (ellipsoid radii), and Peak intensity and last bin in 3-bin probability distribution of relative intensity. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles.

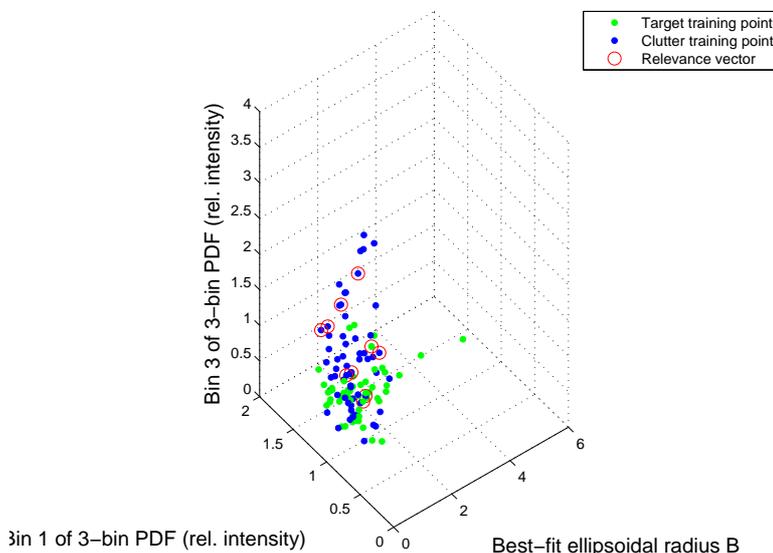


Figure 125: 3D scatter plot showing training points and optimal relevance vectors for the feature triplet B (ellipsoid radii), first bin in 3-bin probability distribution of relative intensity and last bin in 3-bin probability distribution of relative intensity. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles.

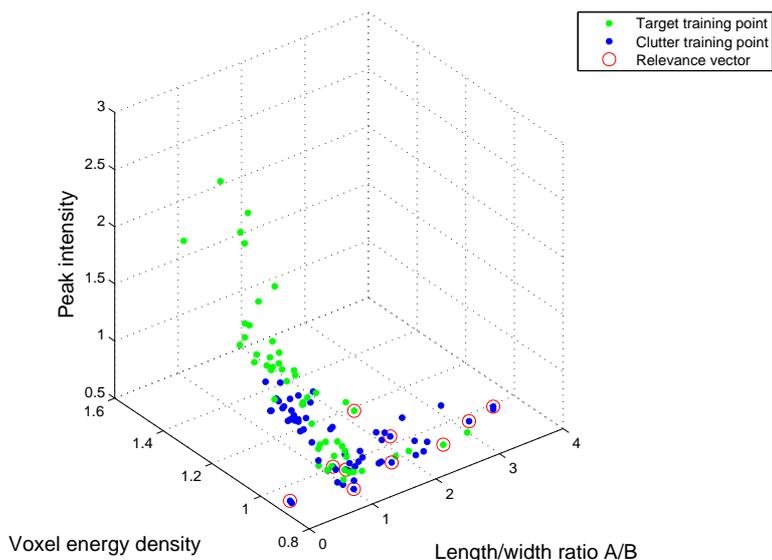


Figure 126: 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *Length to width ratio A/B*, *Voxel energy density* and *Peak intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles.

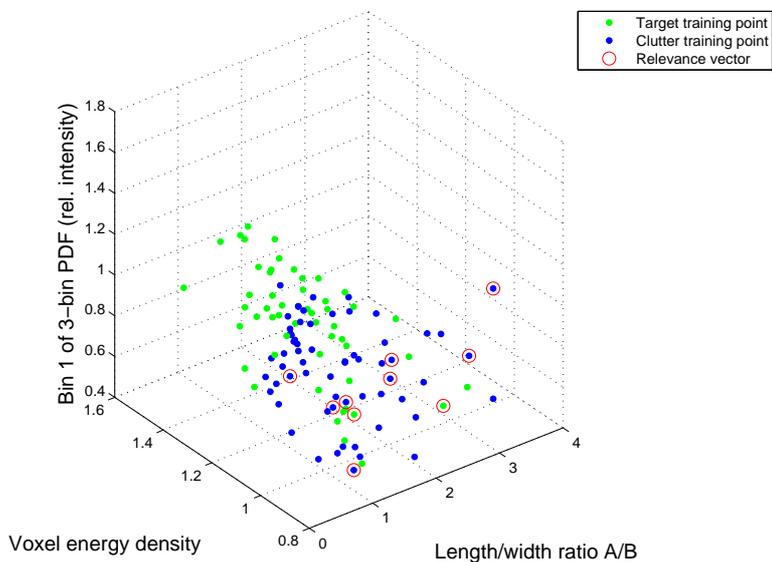


Figure 127: 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *Length to width ratio A/B*, *Voxel energy density* and *first bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles.

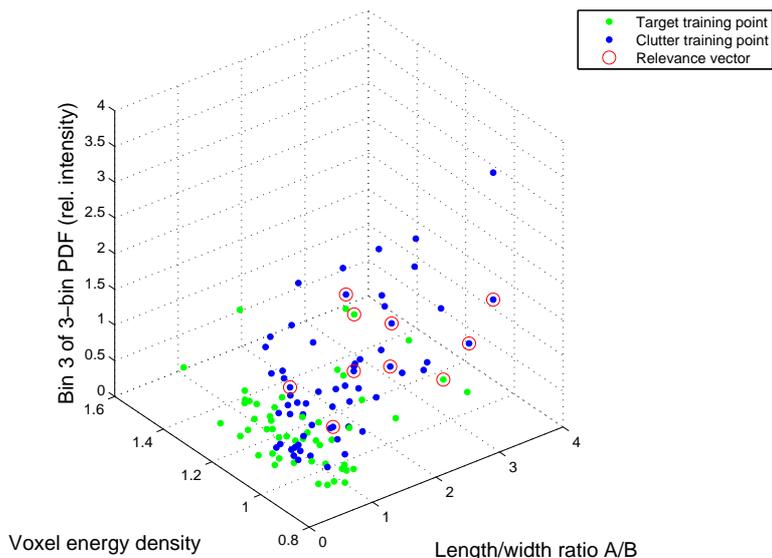


Figure 128: 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *Length to width ratio A/B*, *Voxel energy density* and *last bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles.

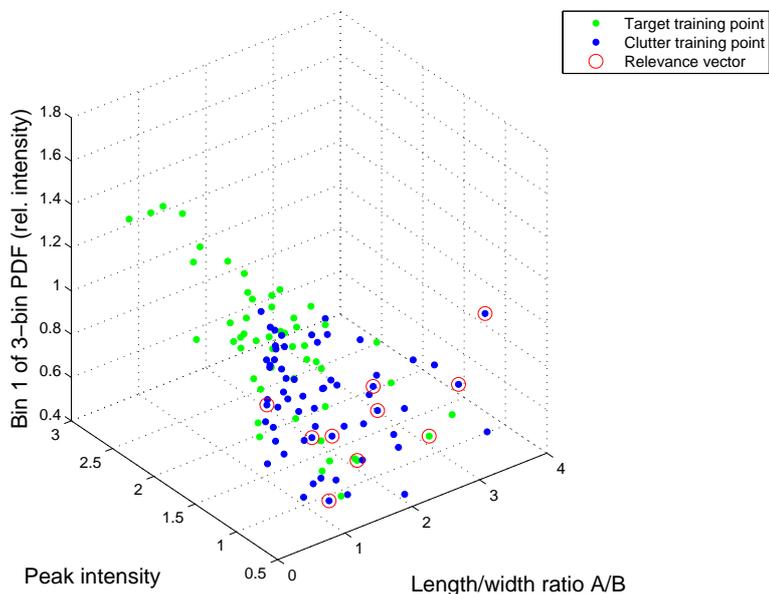


Figure 129: 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *Length to width ratio A/B*, and *Peak intensity*. and *first bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles.

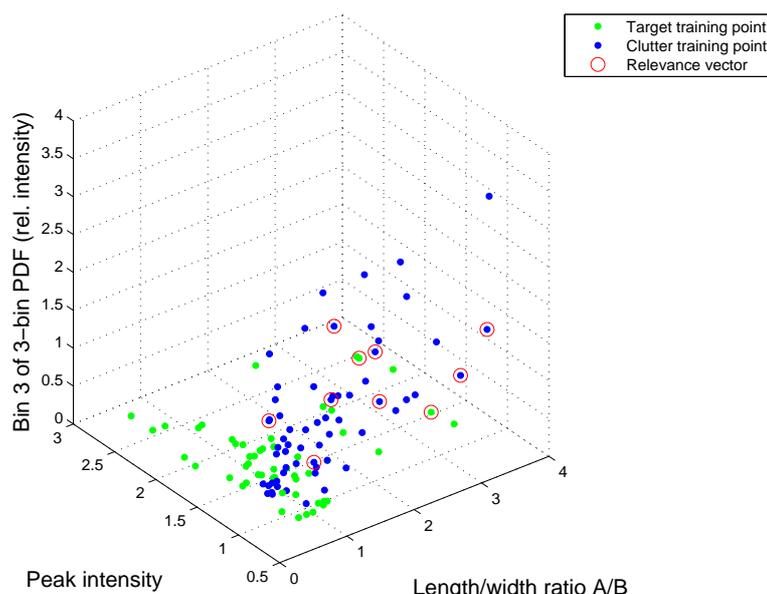


Figure 130: 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *Length to width ratio A/B*, and *Peak intensity*, and *last bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles.

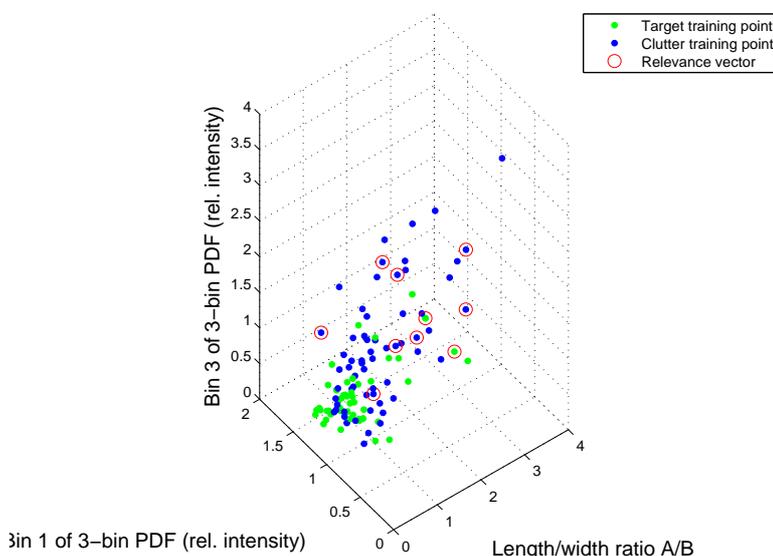


Figure 131: 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *Length to width ratio A/B*, *first bin in 3-bin probability distribution of relative intensity* and *last bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles.

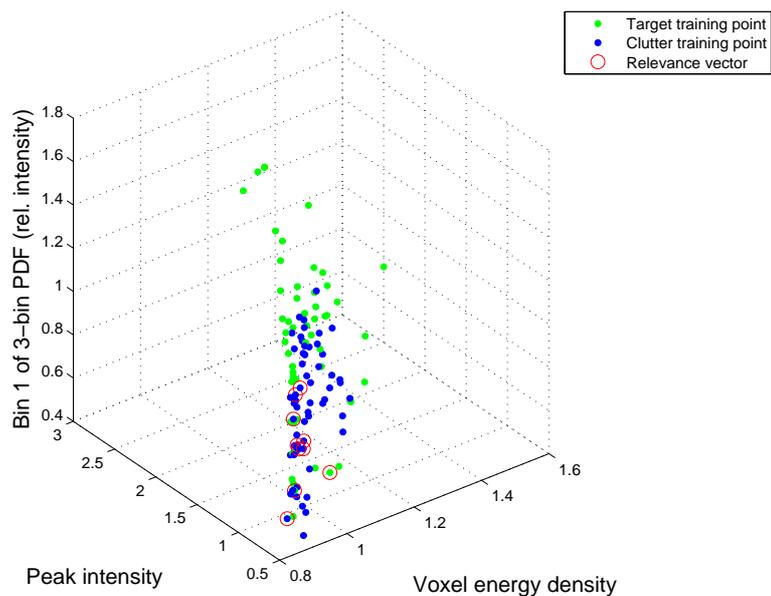


Figure 132: 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *Voxel energy density*, *Peak intensity* and *first bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles.

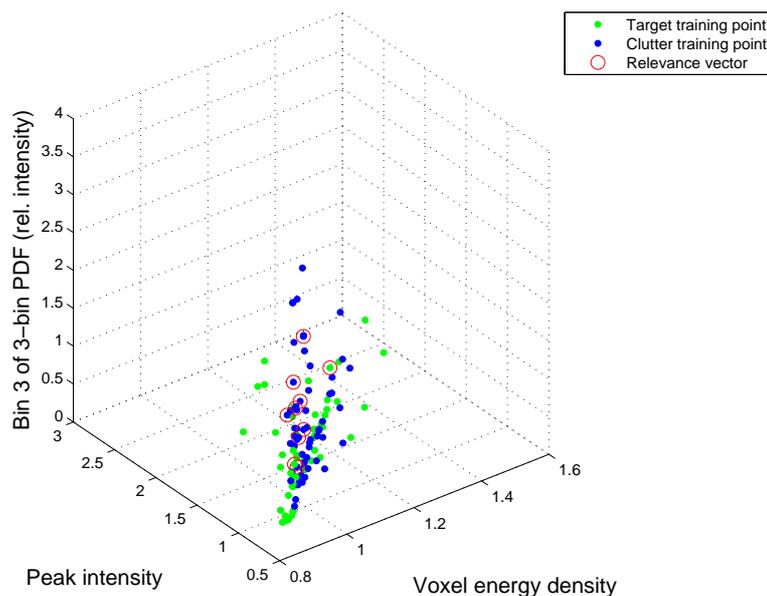


Figure 133: 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *Voxel energy density, Peak intensity* and *last bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles.

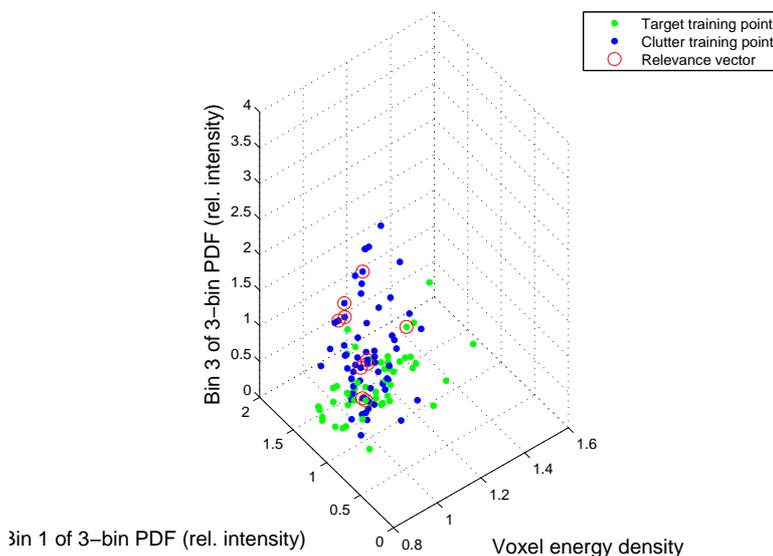


Figure 134: 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *Voxel energy density, first bin in 3-bin probability distribution of relative intensity* and *last bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles.

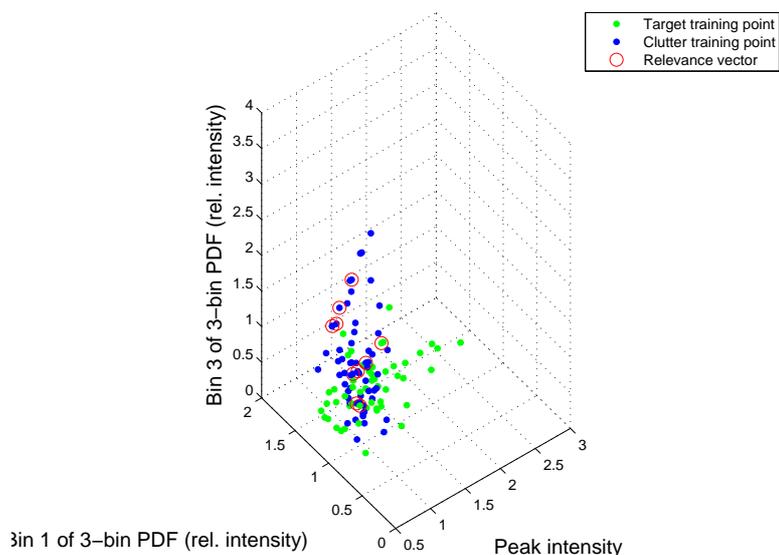


Figure 135: 3D scatter plot showing training points and optimal relevance vectors for the feature triplet *Peak intensity, first bin in 3-bin probability distribution of relative intensity and last bin in 3-bin probability distribution of relative intensity*. The features were derived from the BOSS AUV FEST 2008 data set. The training vectors used for classifier, and corresponding to known target and clutter items are indicated by green and blue discs, respectively. Of these, the optimal relevance vector are those indicated by the red circles.

A.6 Statistical realizations for classifier training

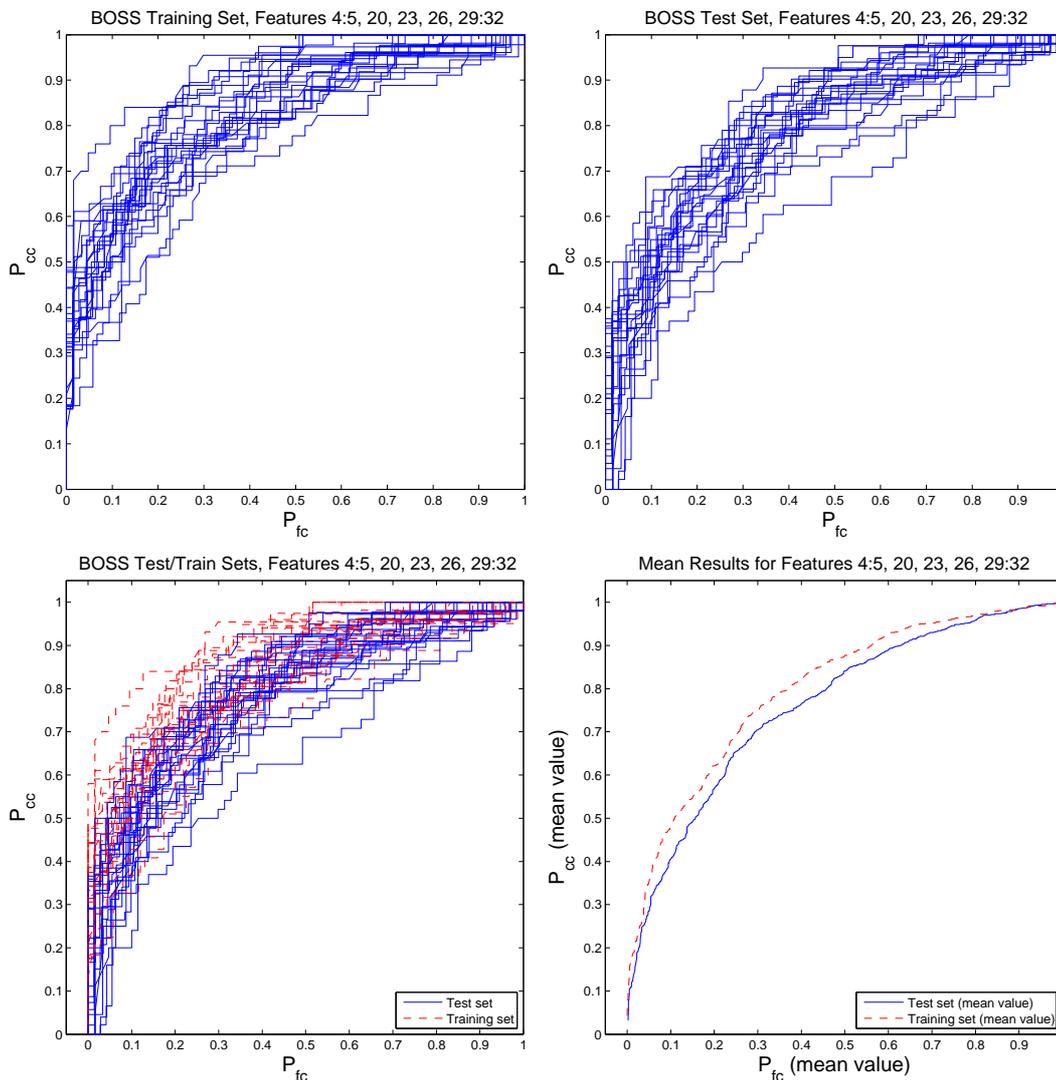


Figure 136: Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets.

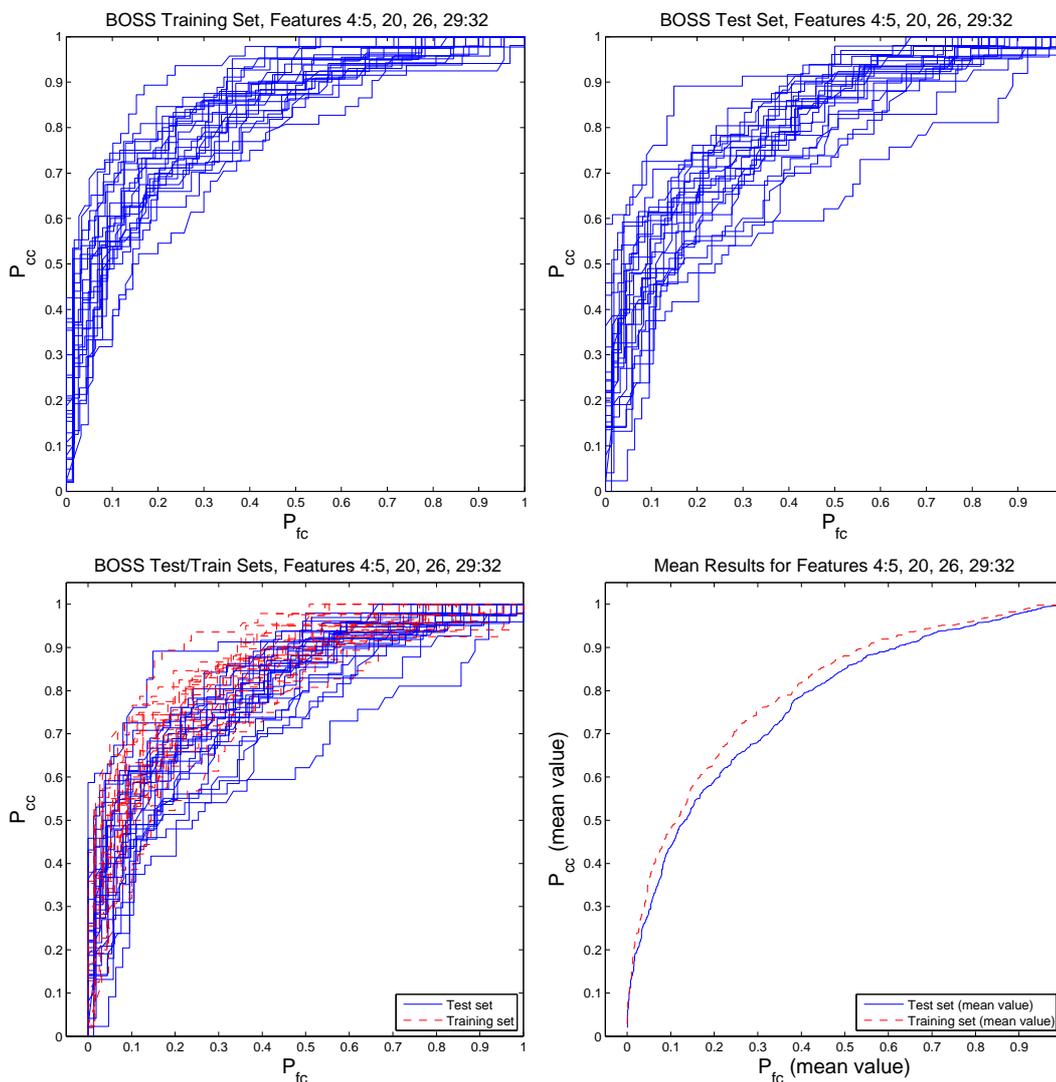


Figure 137: Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets.

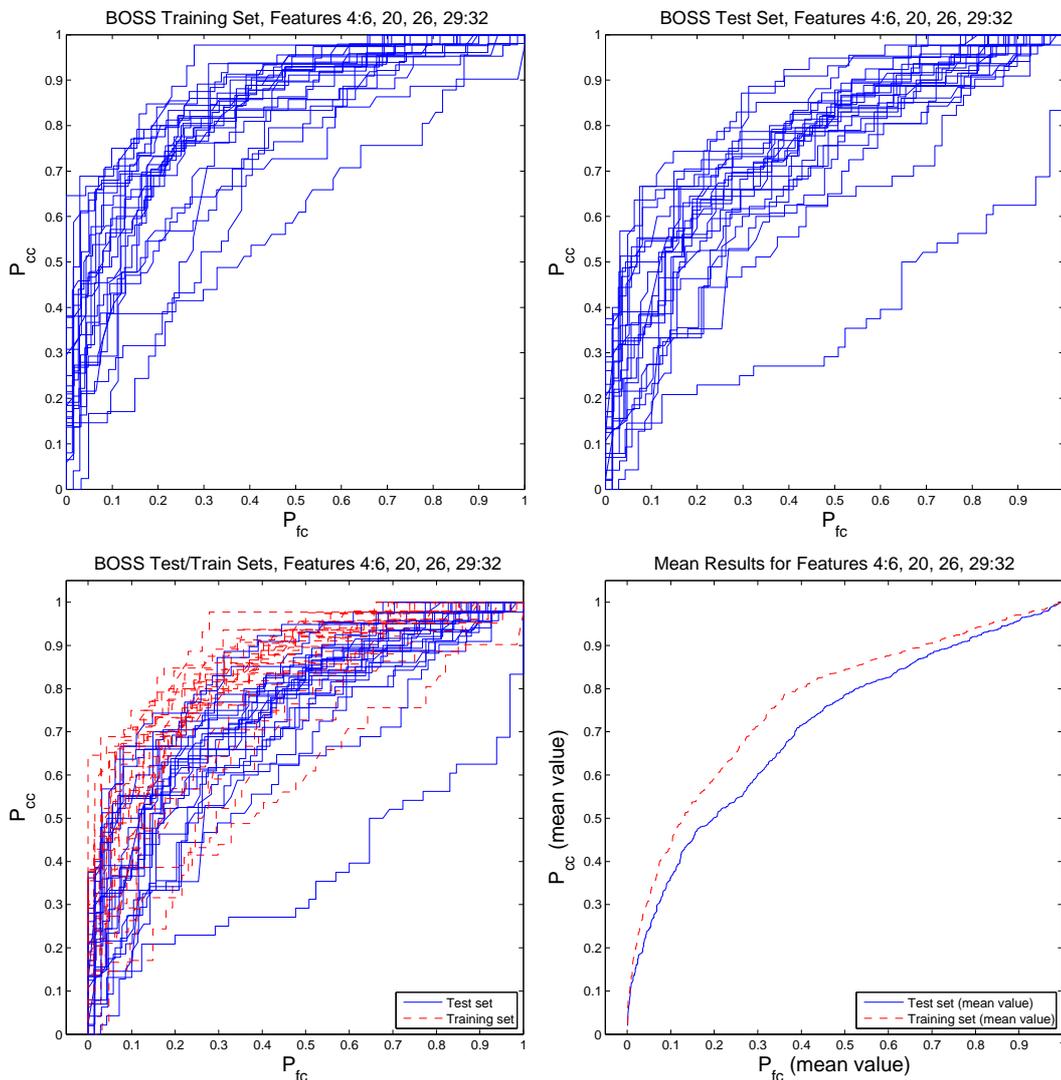


Figure 138: Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets.

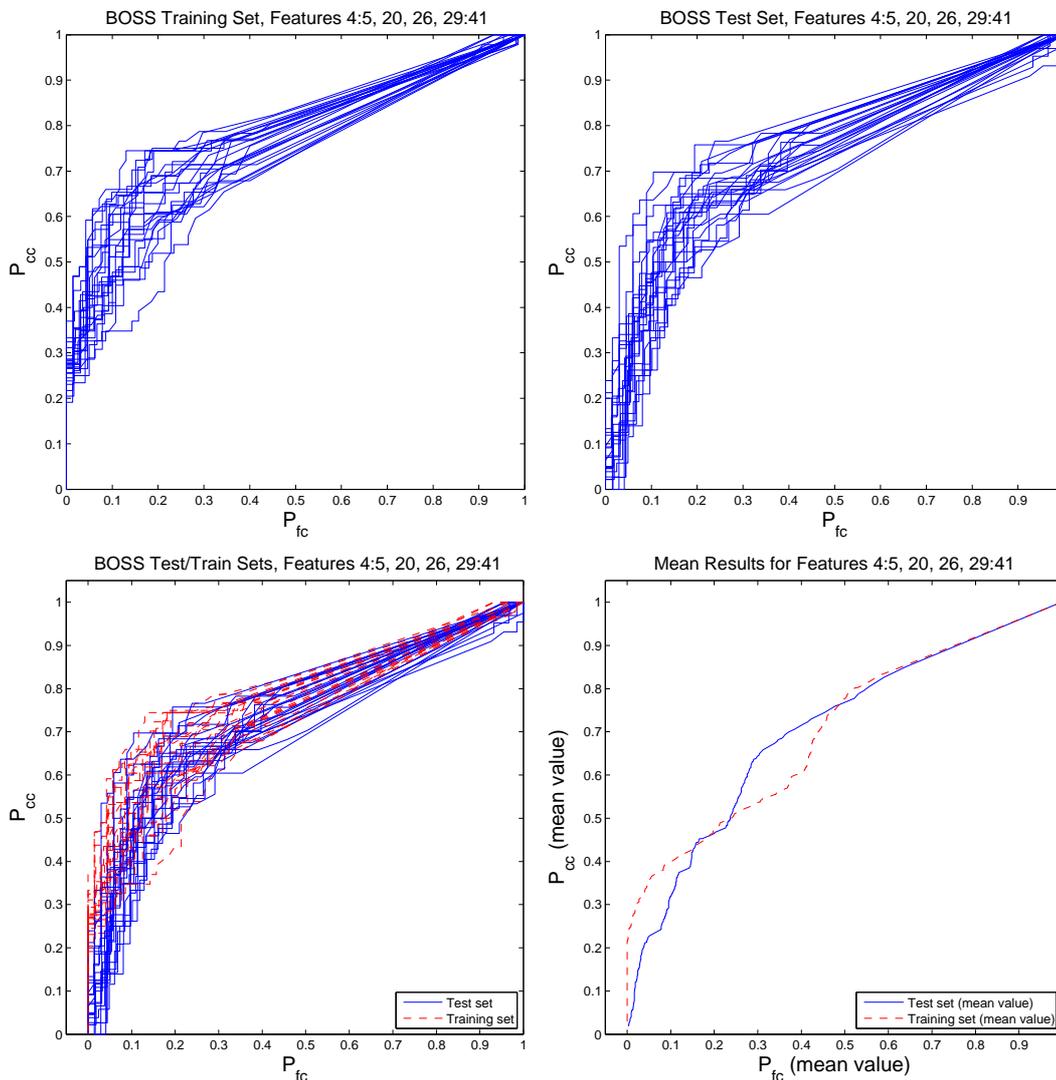


Figure 139: Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets.

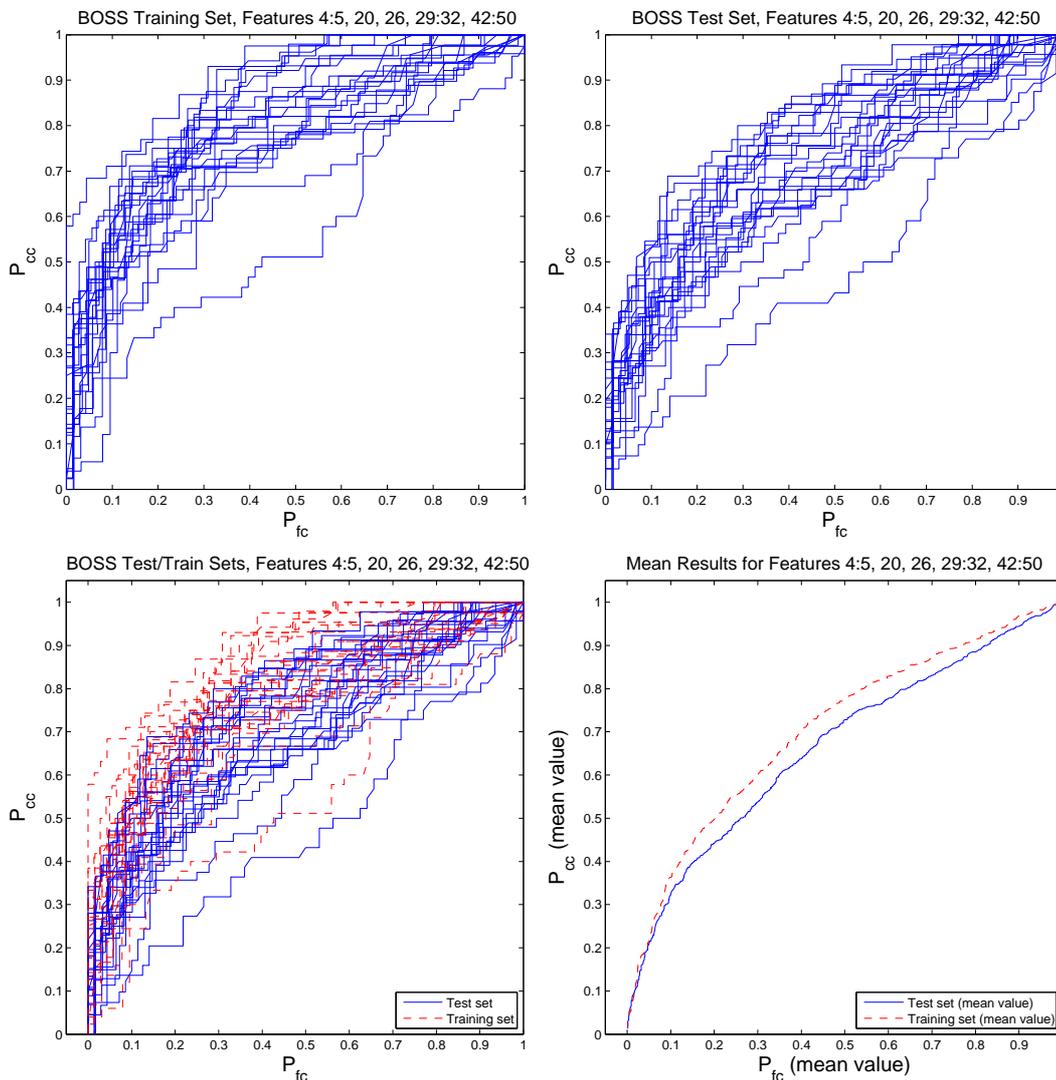


Figure 140: Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets.

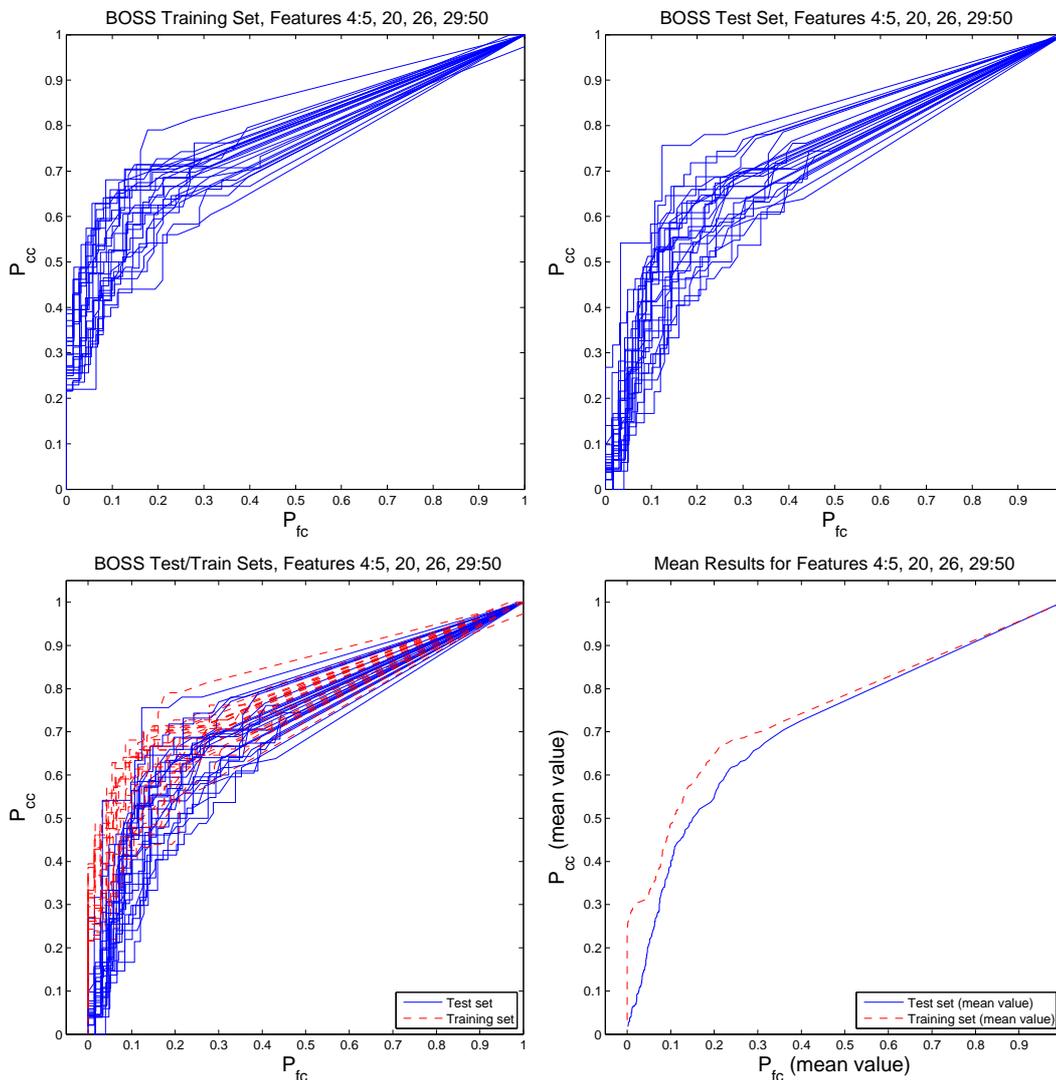


Figure 141: Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets.

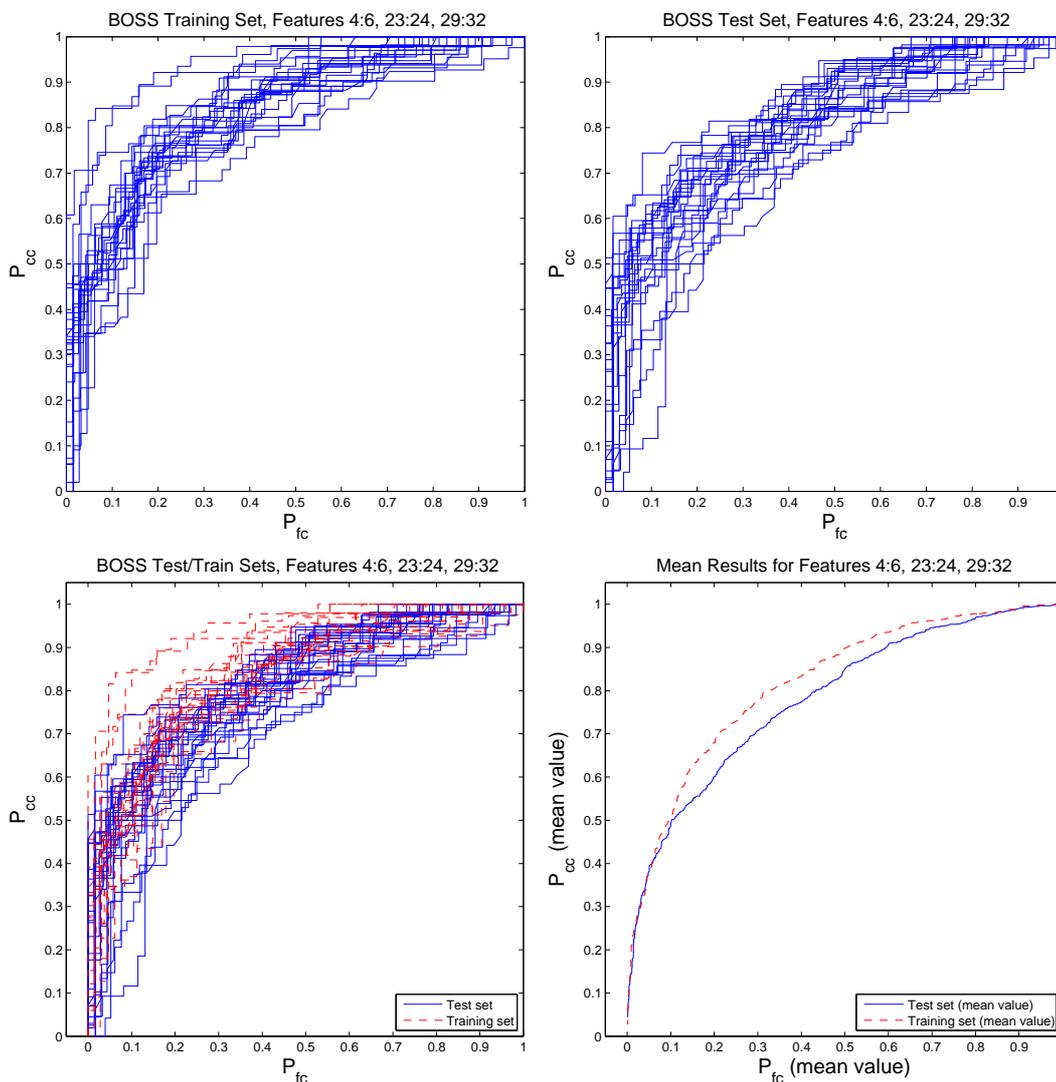


Figure 142: Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets.

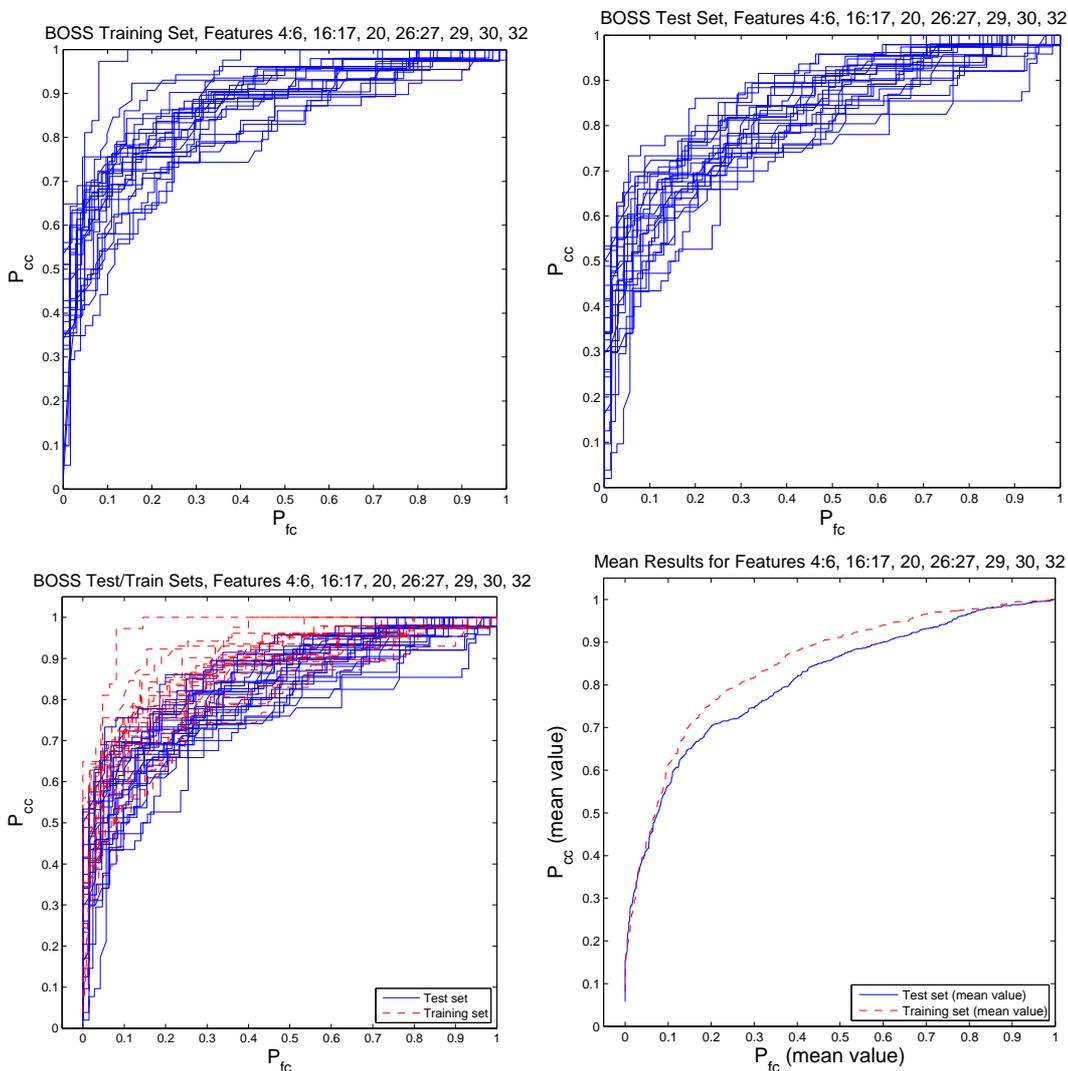


Figure 143: Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets.

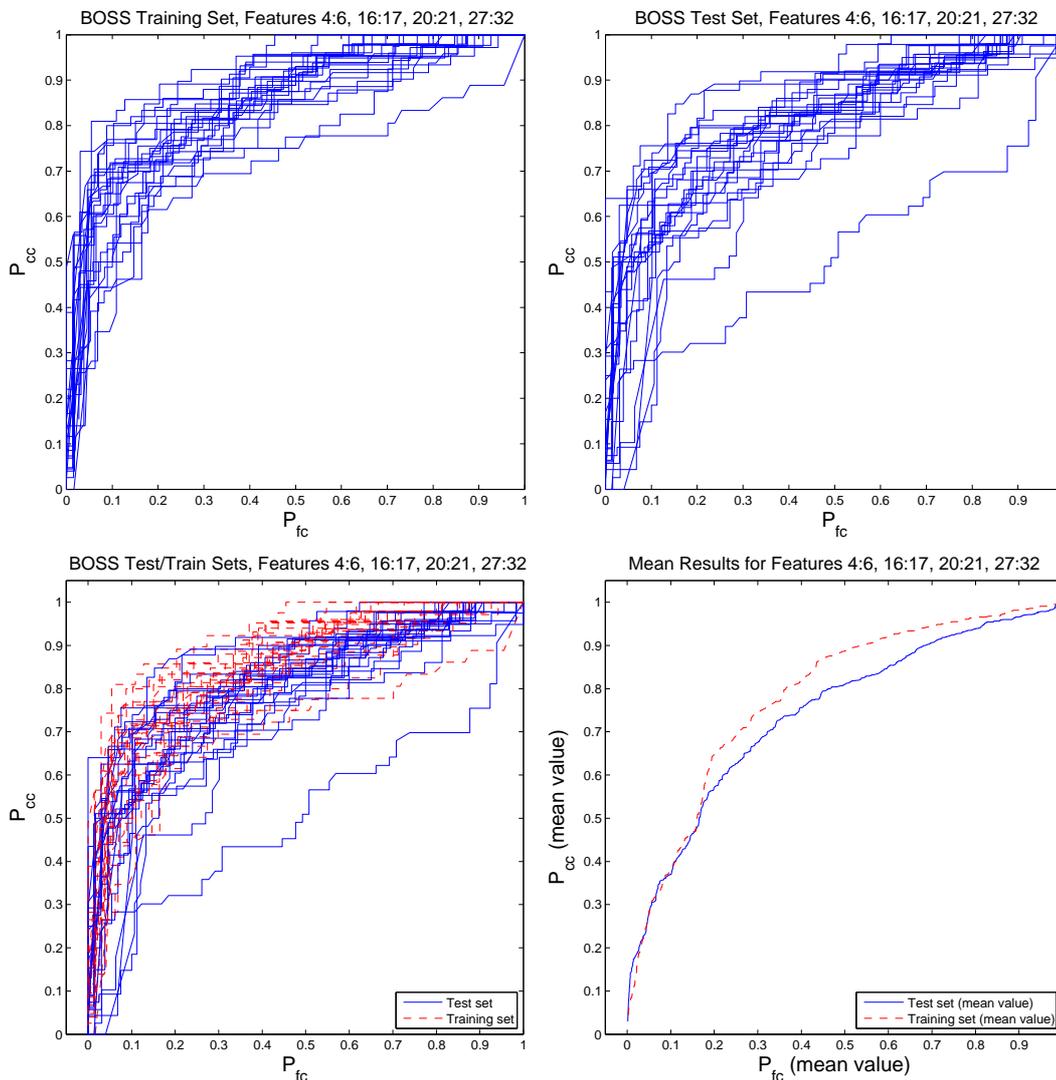


Figure 144: Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets.

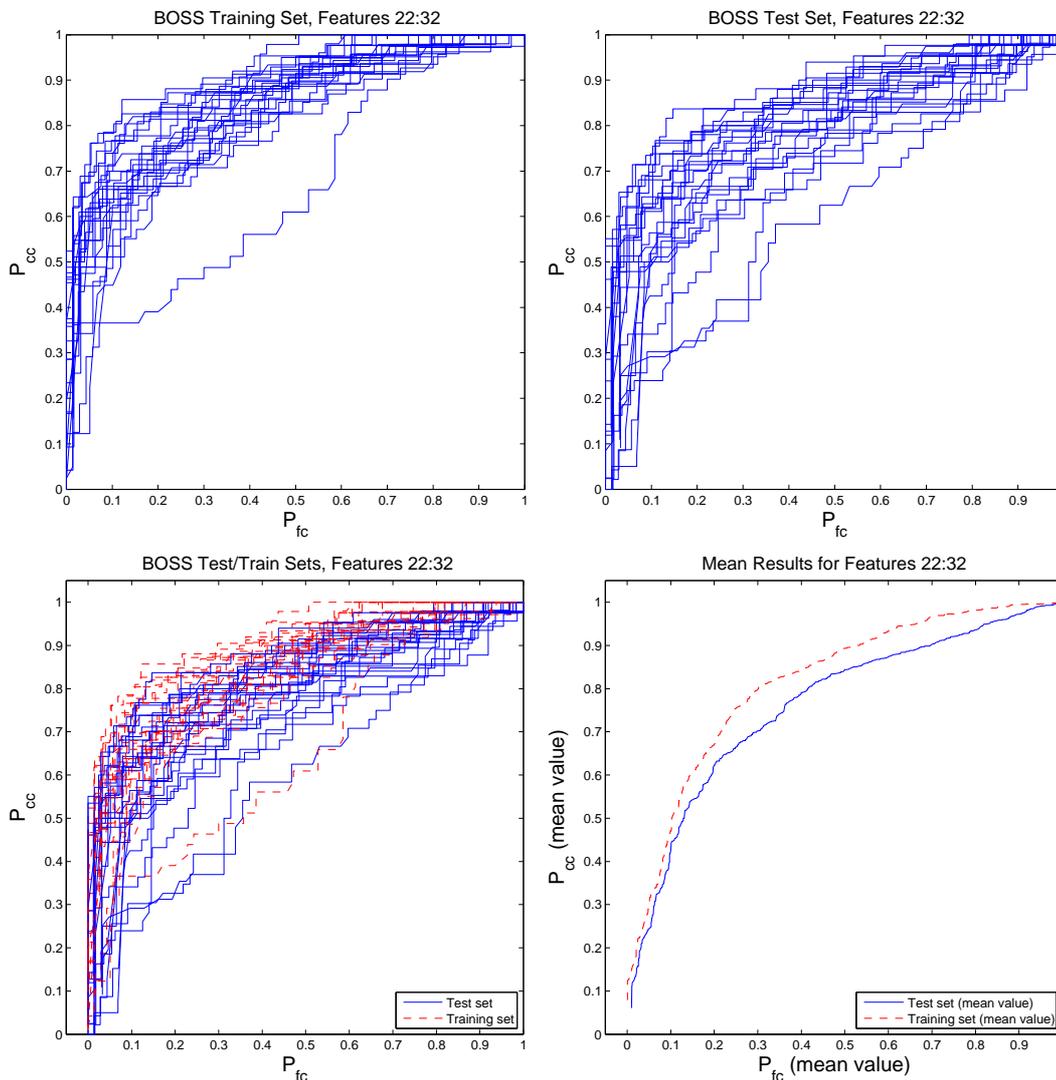


Figure 145: Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets.

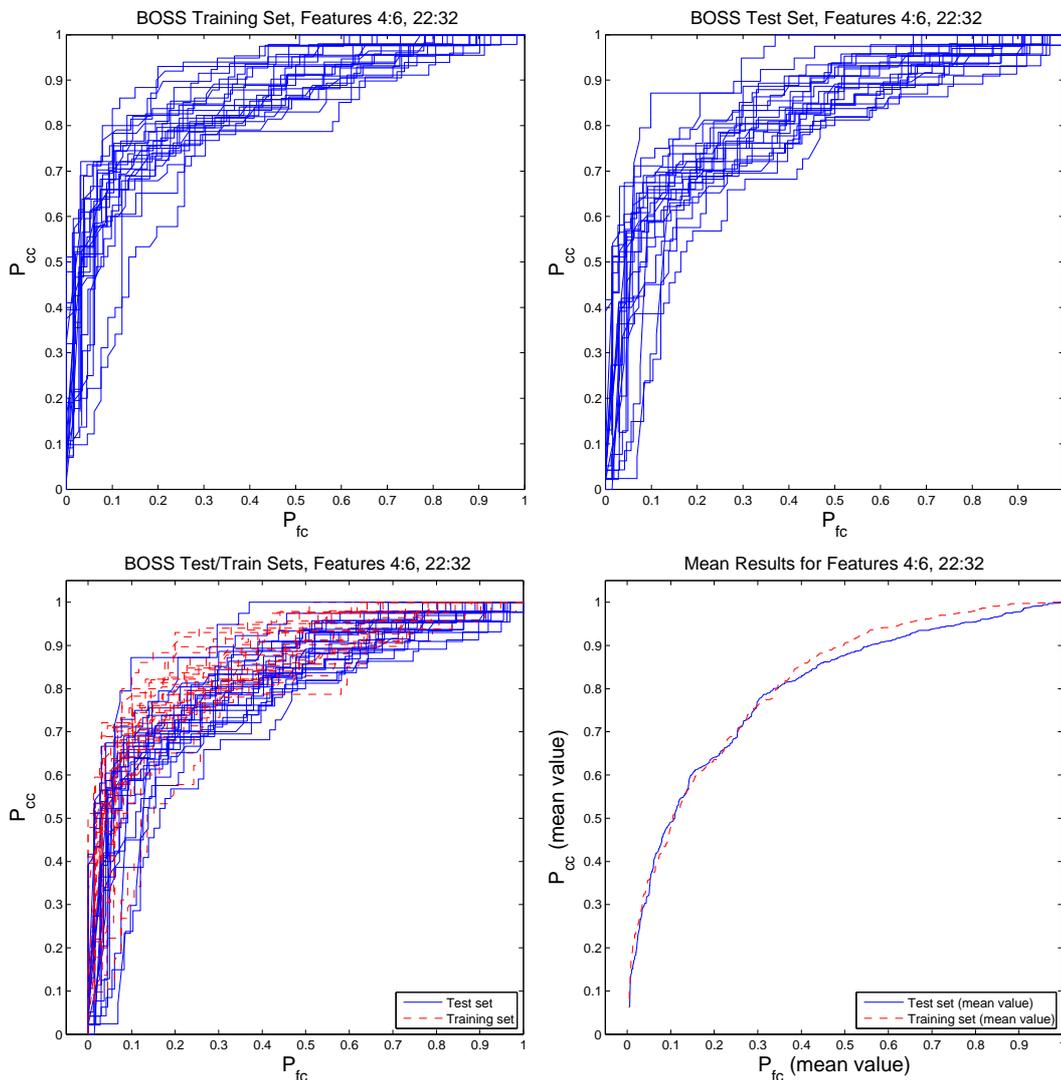


Figure 146: Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets.

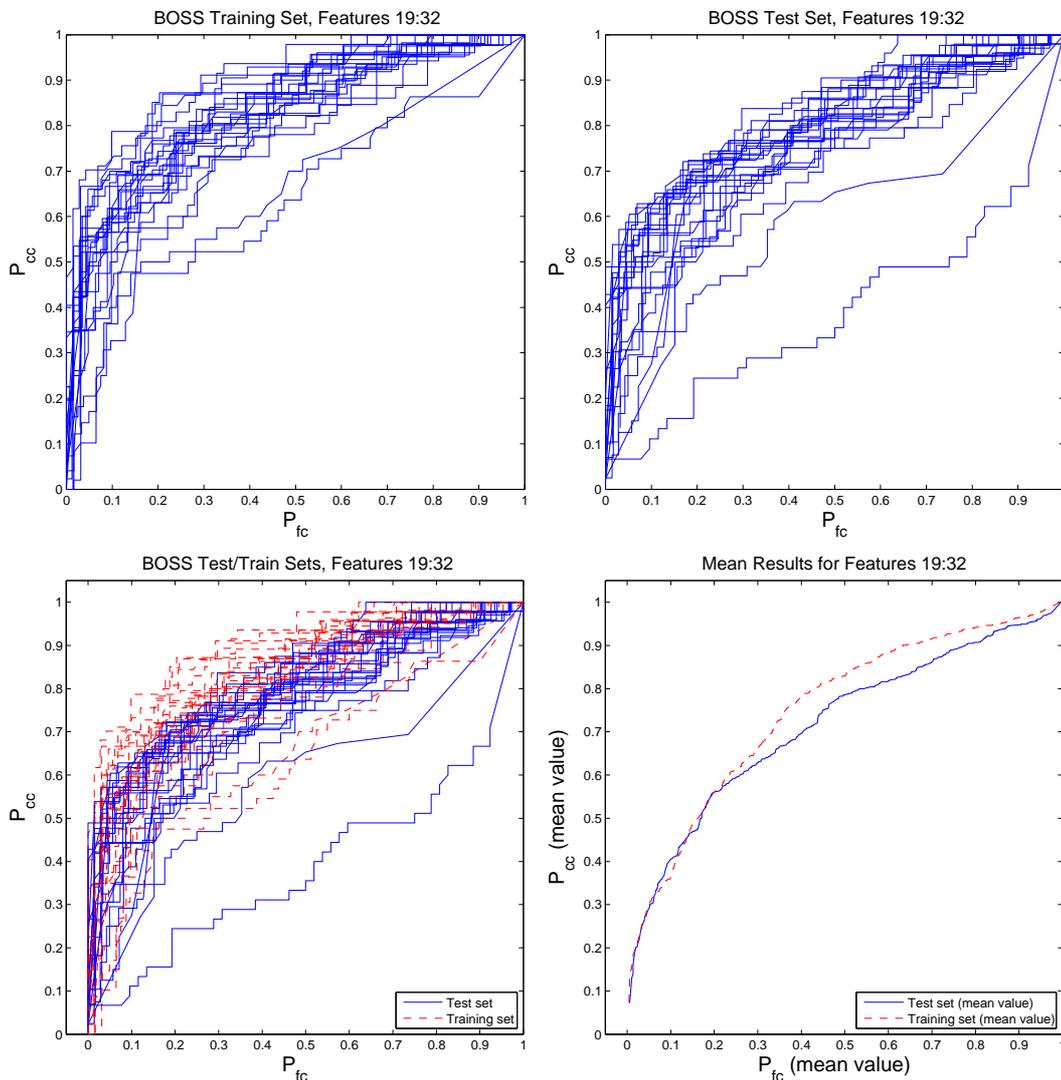


Figure 147: Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets.

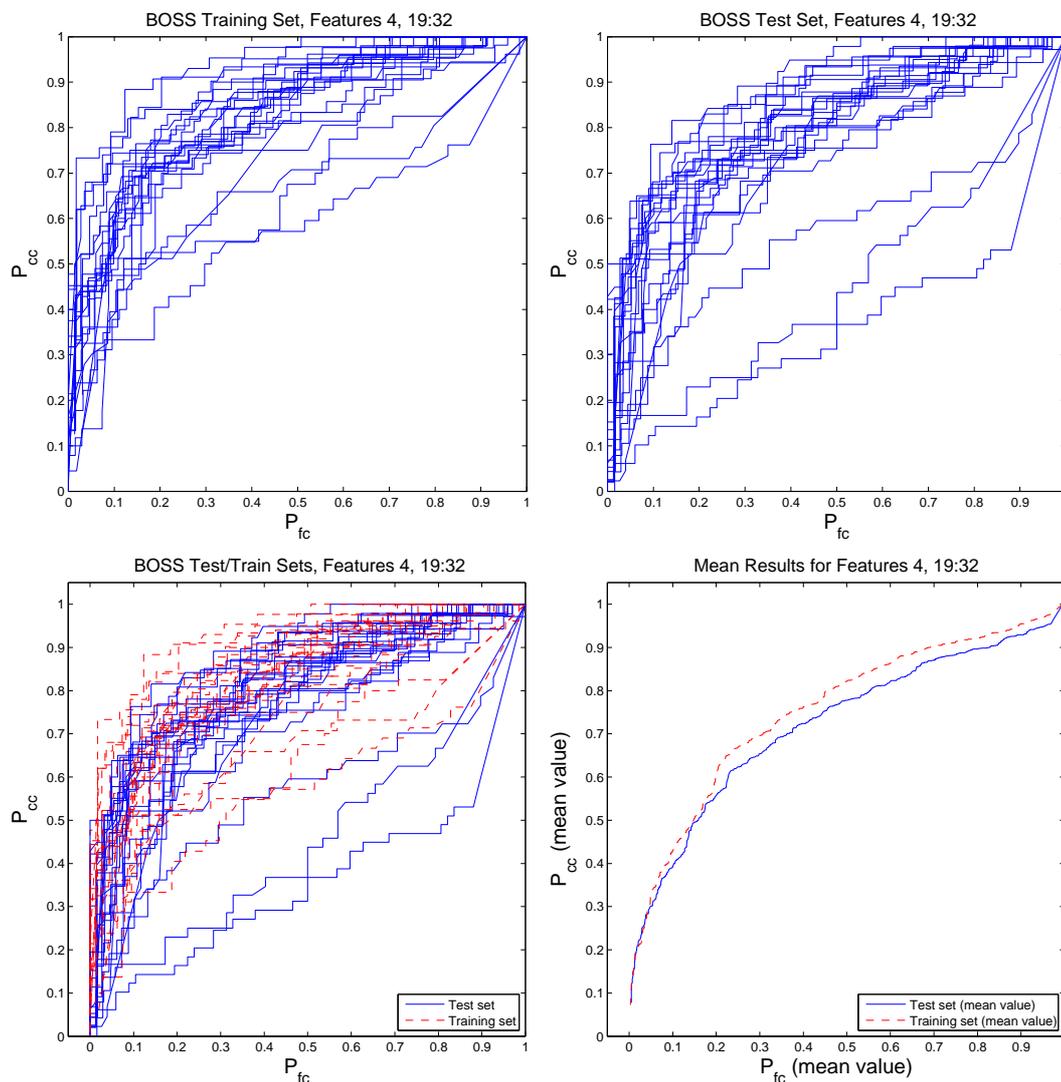


Figure 148: Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets.

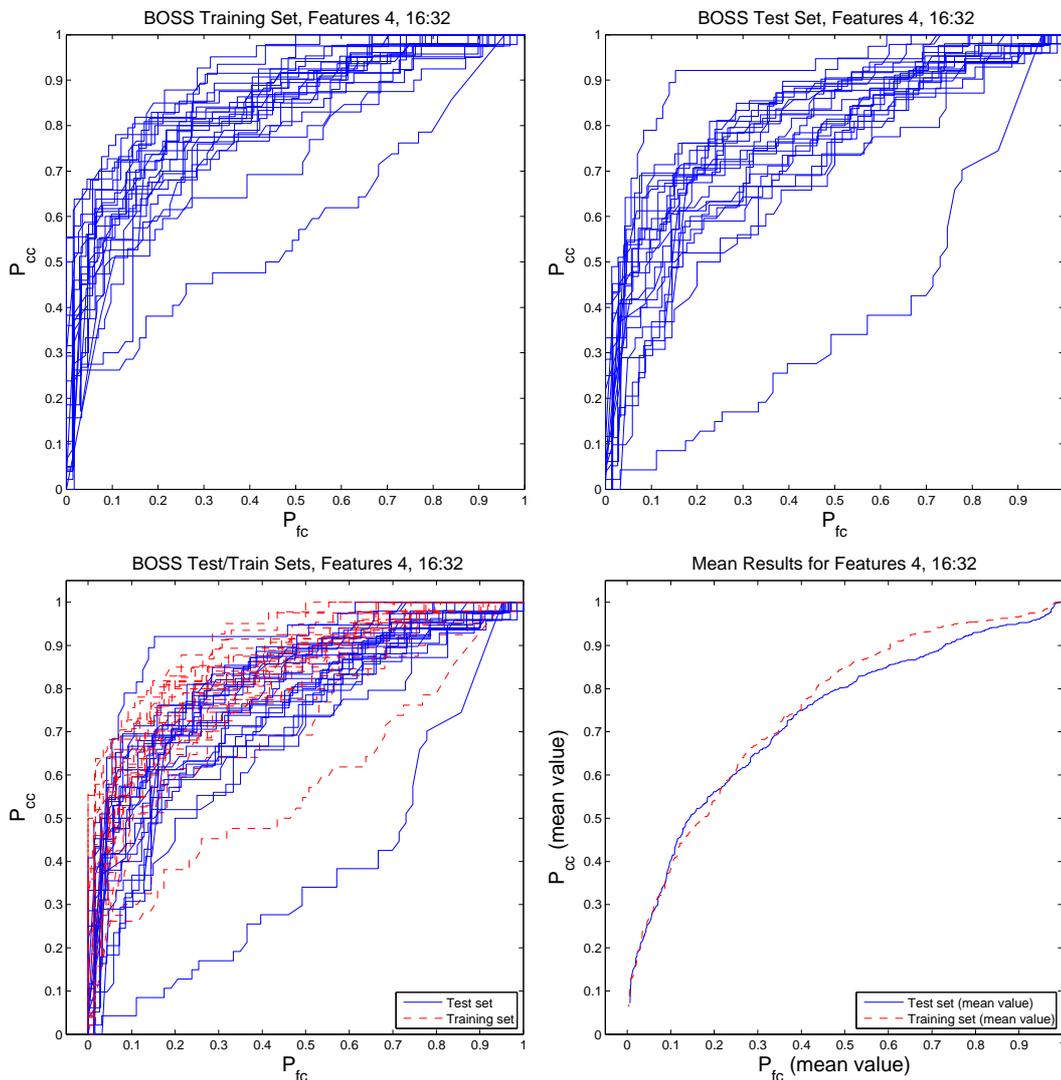


Figure 149: Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets.

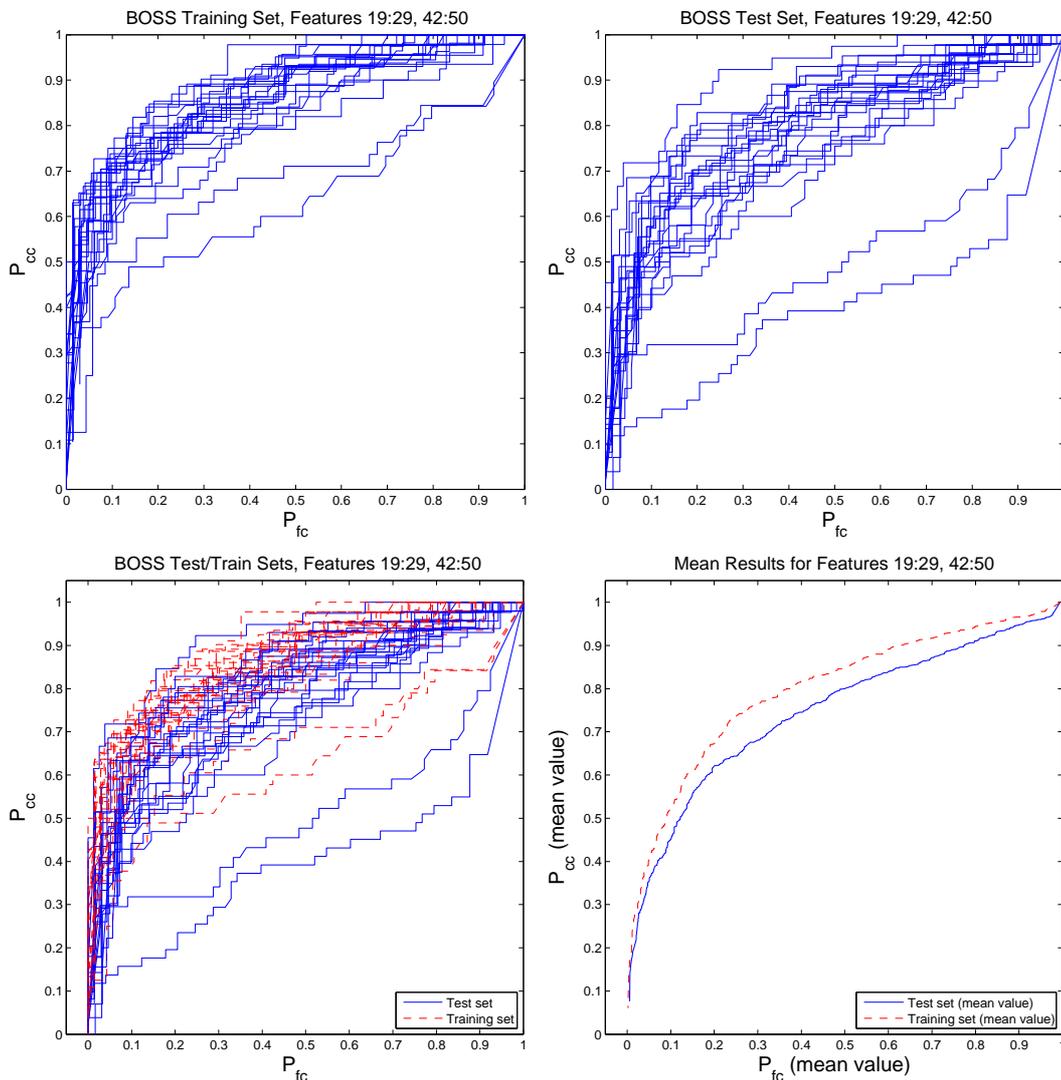


Figure 150: Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets.

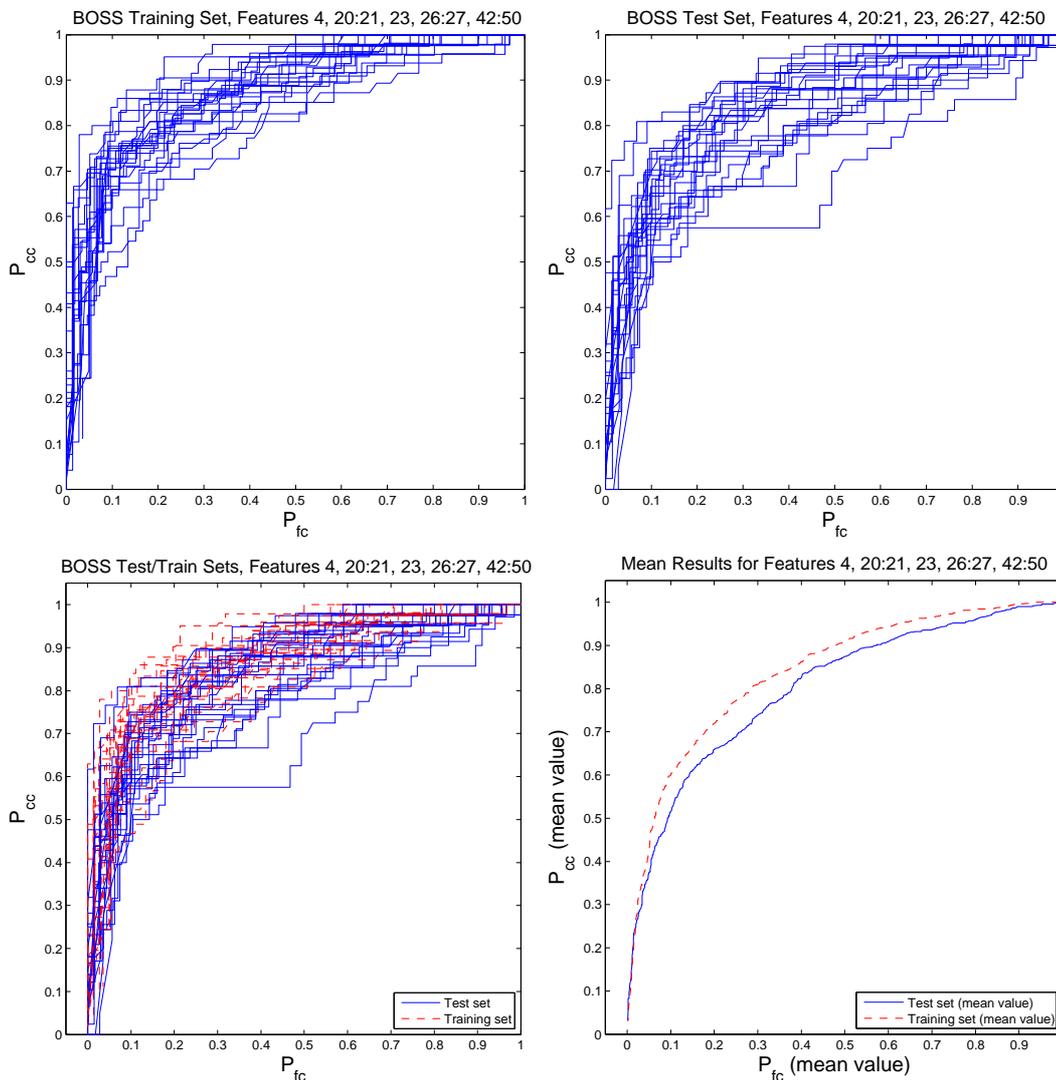


Figure 151: Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets.

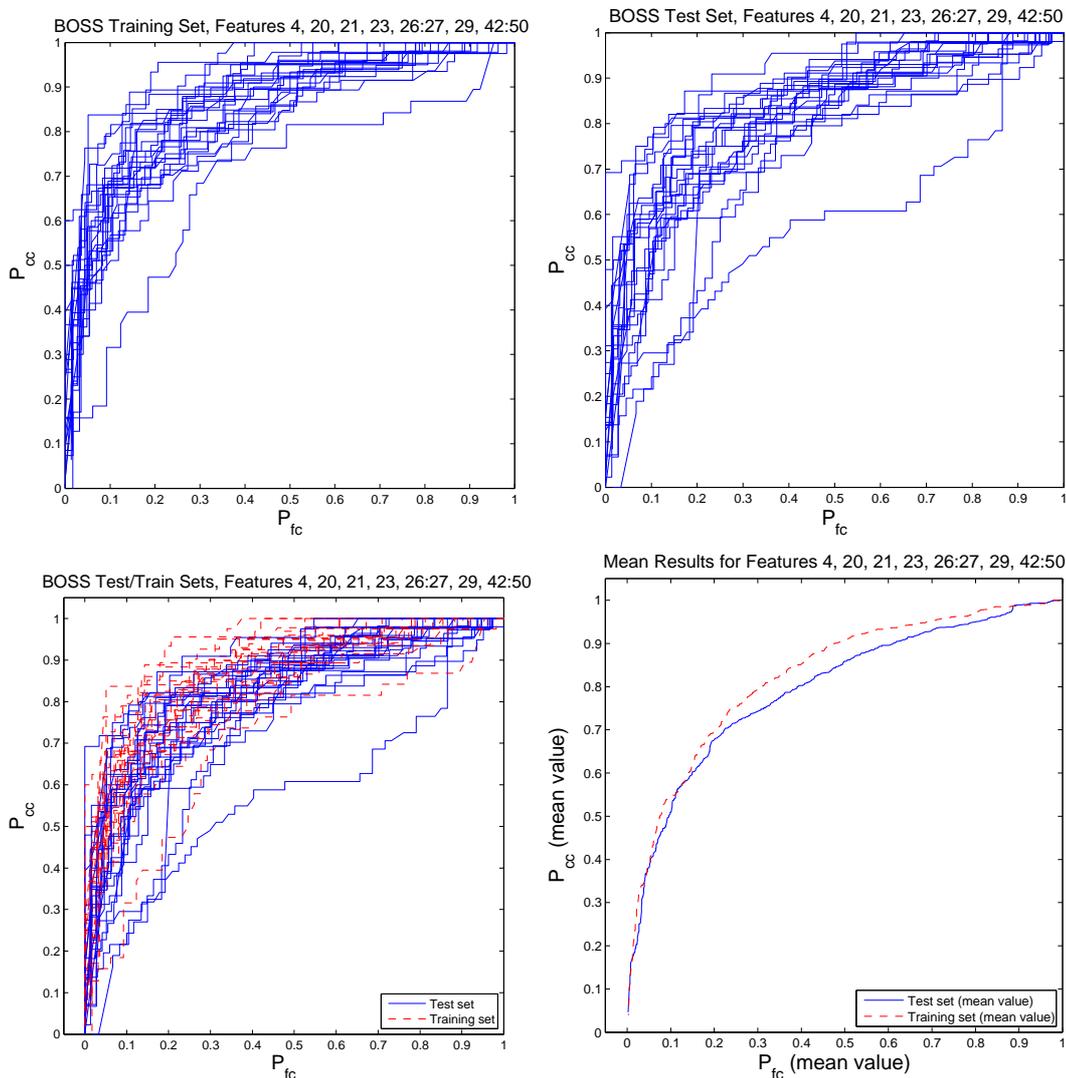


Figure 152: Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets.

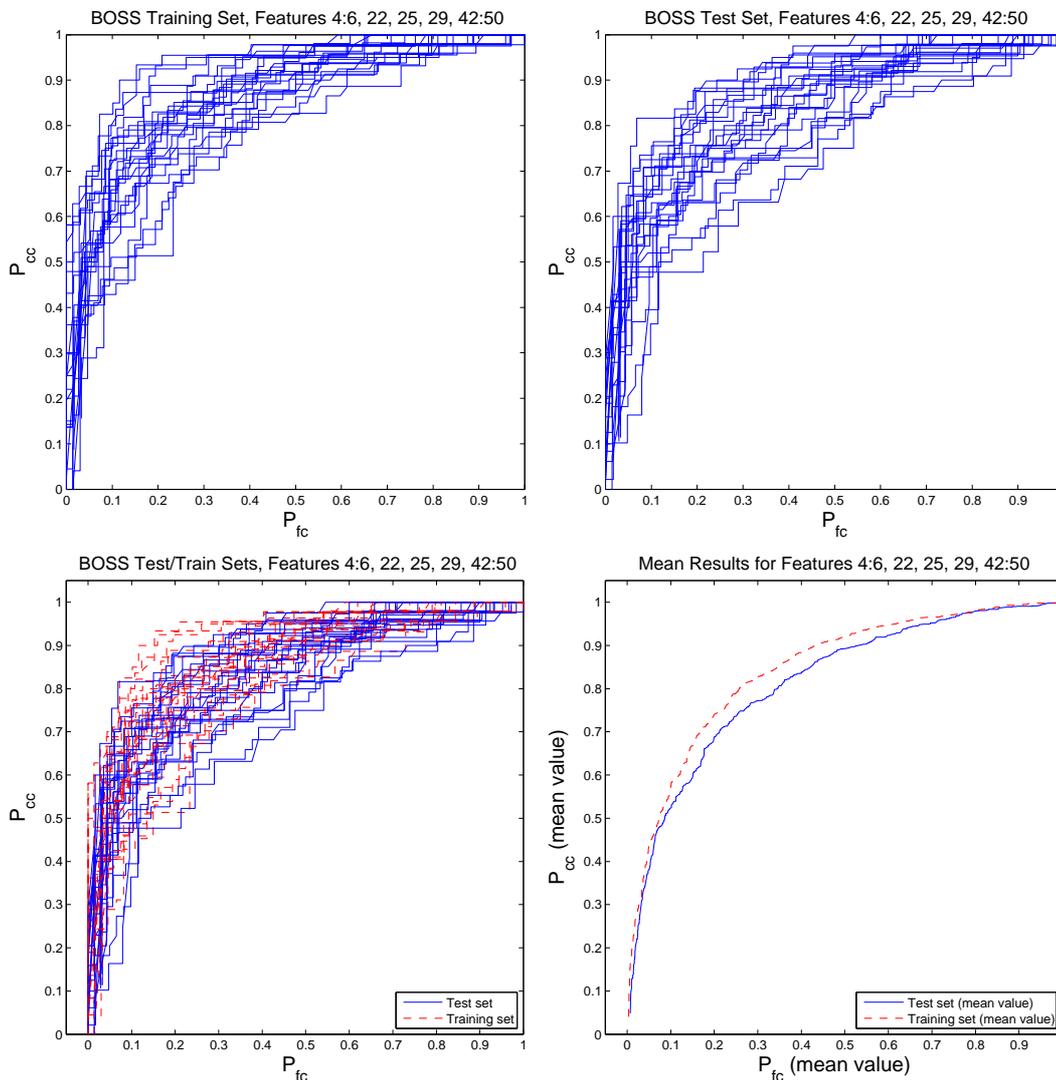


Figure 153: Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets.

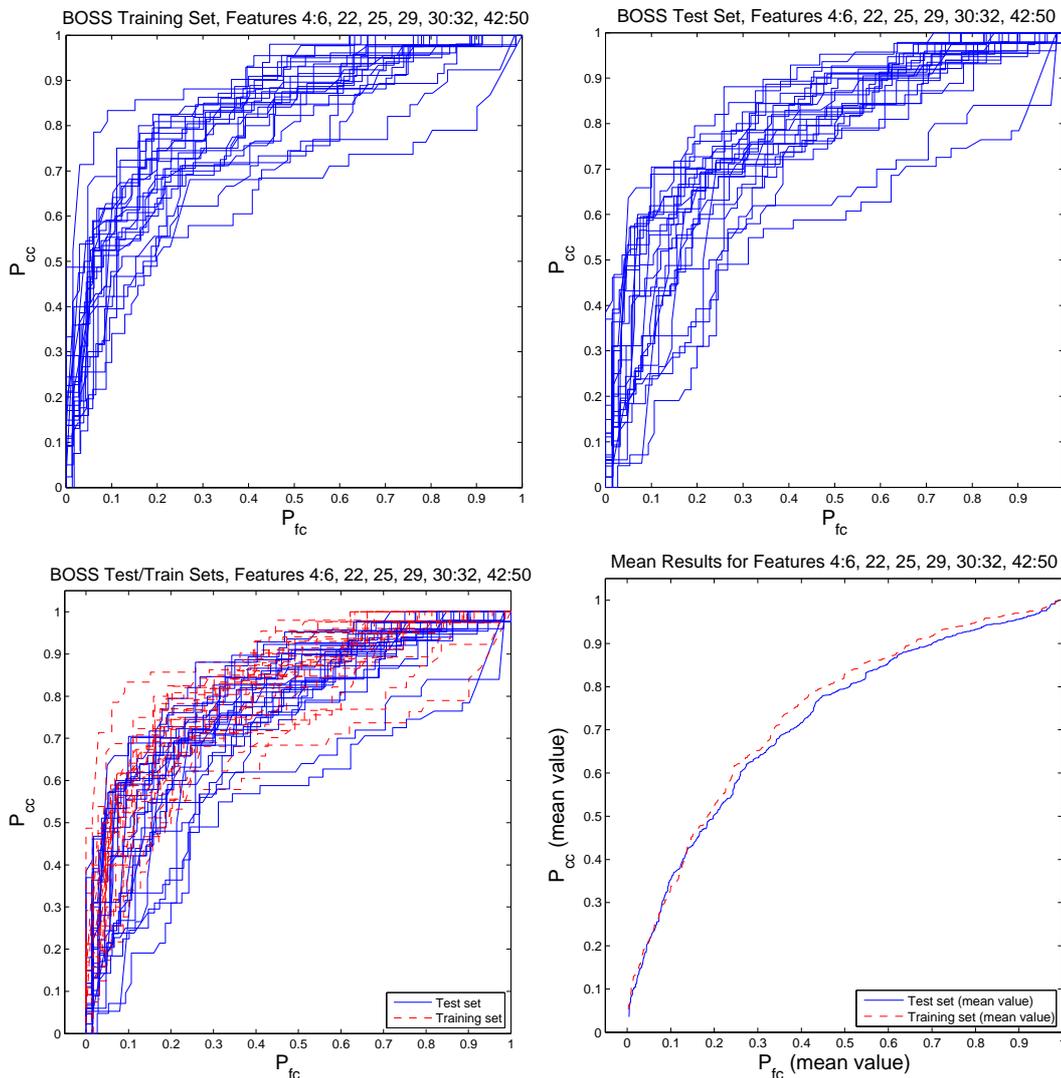


Figure 154: Classification results for AUV Fest 2008 BOSS data using the Relevance Vector Machine. In training, typically half of a data set is randomly chosen as the ‘training set’ and the remaining half is chosen as the ‘test set’. Thus, there are numerous random subsets that could be chosen, and so for statistical robustness in choosing an optimal feature subset, the selection should be made with respect to the ensemble results for various subsets of target/clutter training points. In the results shown here, 25 realizations of training vectors were chosen, and the corresponding ROC curves for each set are shown. **Upper left:** ROC classifier performance for the training set (half of target/clutter exemplars) **Upper right:** ROC classifier performance for the test set (the other half of the target/clutter exemplars). **Lower left:** ROC classifier performance for training *and* test sets.. **Lower right:** Mean value of the ROC curves in each dimension from the training and test sets.

A.7 BOSS processing results

In this section we display our processing results of the data collects described in Sec. 4.3.1 using the signal and image processing steps described in Secs. 3.3.1 – 3.3.4. We display 2D image constructs corresponding to the maximum intensity projection mapping in the horizontal (x,y) plane i.e., $P_{xy}(x,y)$ in equation (114). These results are for both the mud and sand sites and for various platform altitudes above the seabeds. These varying conditions are indicated in the figure notations, and these notations are defined below in Table 9. Figures 155 – 165 show various MIP

Notation	Definition
ss	The 'ss' runs were collected at the 'sand site', just inside St. Andrew's Bay, south of Panama City
boss	The 'boss' runs were at the 'mud site', in the northern part of St. Andrew's Bay, directly West of the airport
low	These are low-altitude runs, typically 2-3 m above the seabed. The total water depth was approximately 10 m
high	These are high-altitude runs, typically 5-6 m above the seabed. The total water depth was approximately 10 m

Table 9: Data notation definitions for the BOSS data collections and processing results.

images for the sand and mud sites. In addition, we show in Figs. 166 – 172 show detailed views of a particular target at the mud site for various depth planes corresponding to B_{x_F} in eq. (111).

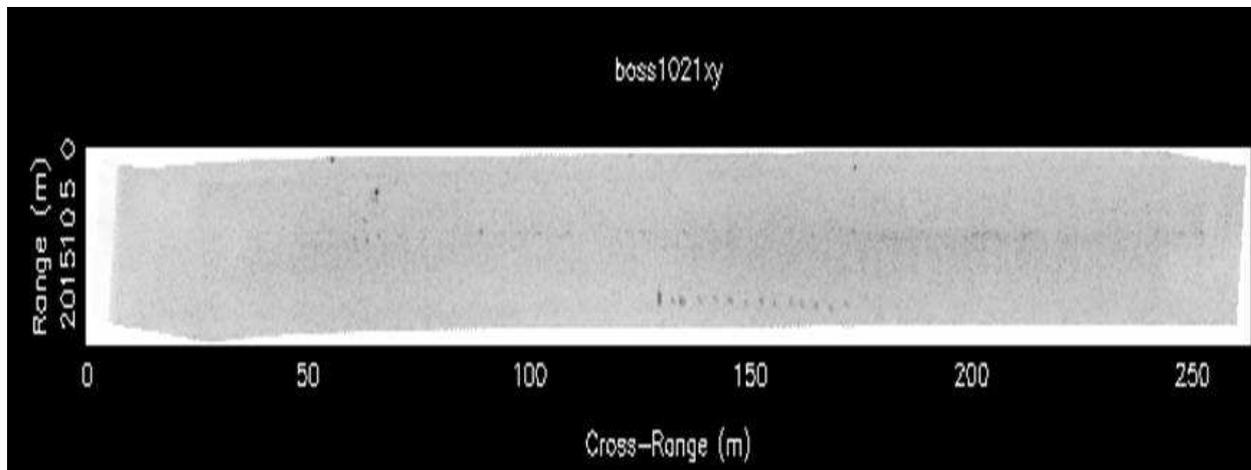


Figure 155: This figure shows P_{xy} (maximum value over all z planes at each (x,y) point) for an entire run at the mud site. Total water depth is approximately 10 m.

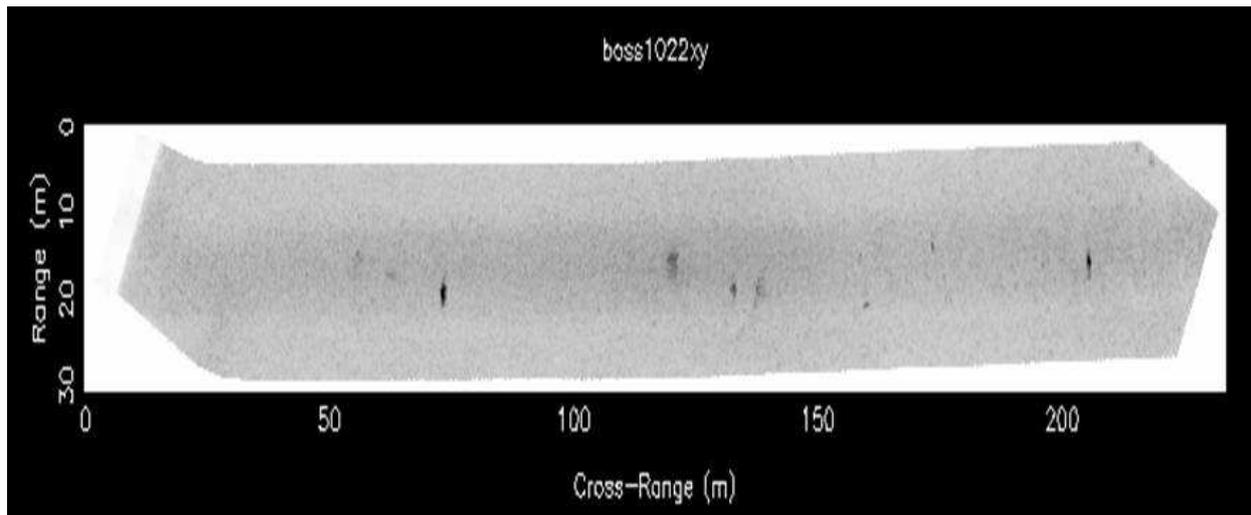


Figure 156: Similar to Fig. 155 but for a different run.

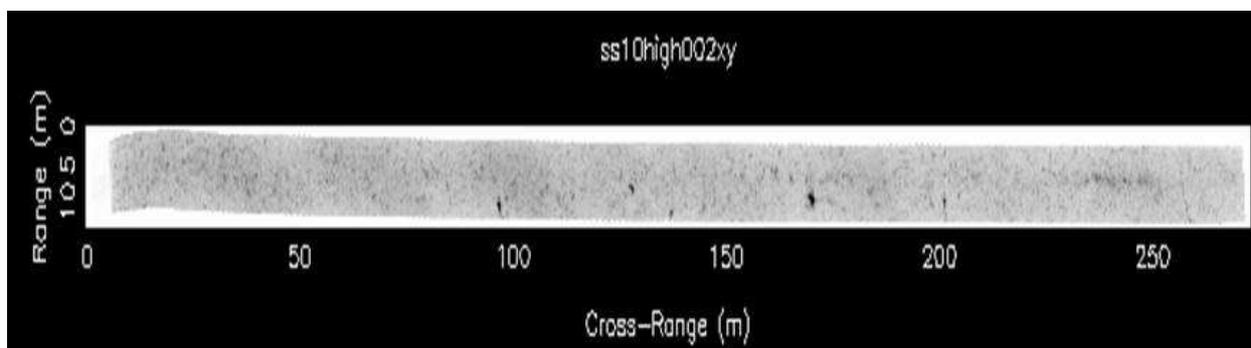


Figure 157: This figure shows P_{xy} (maximum value over all z planes at each (x,y) point) for an entire run at the sand site. The platform was at a ‘high’ altitude above the seabed (typically 5 to 6 m). Total water depth is approximately 10 m. Exploded views of this run are shown in Figs. 158 and 159.

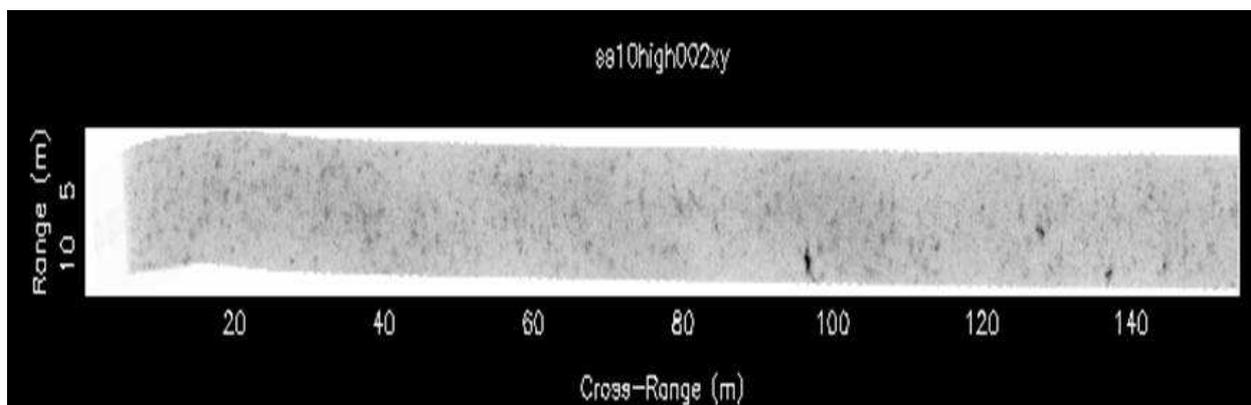


Figure 158: Exploded view of Fig. 157 for the *first half* of the run.

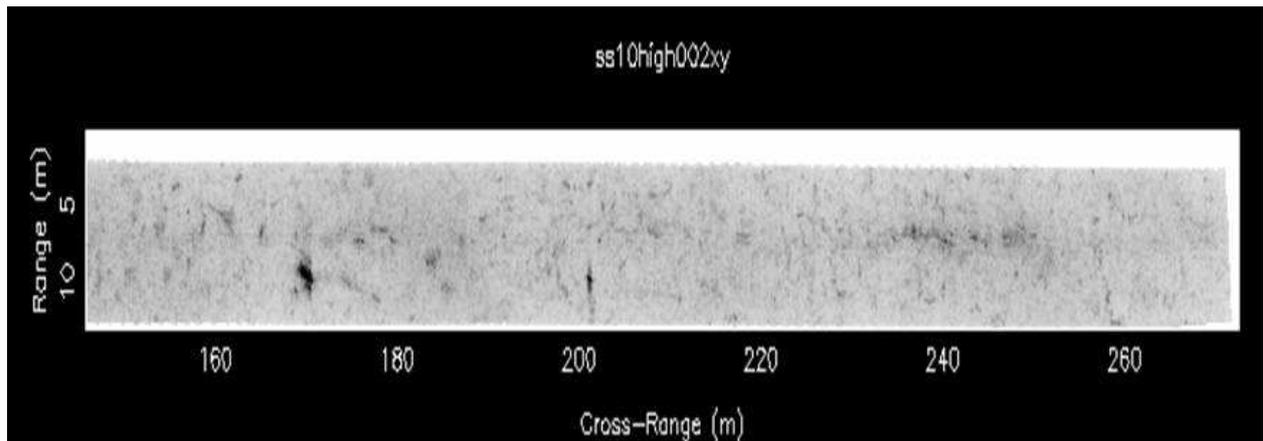


Figure 159: Exploded view of Fig. 157 for the *second half* of the run.

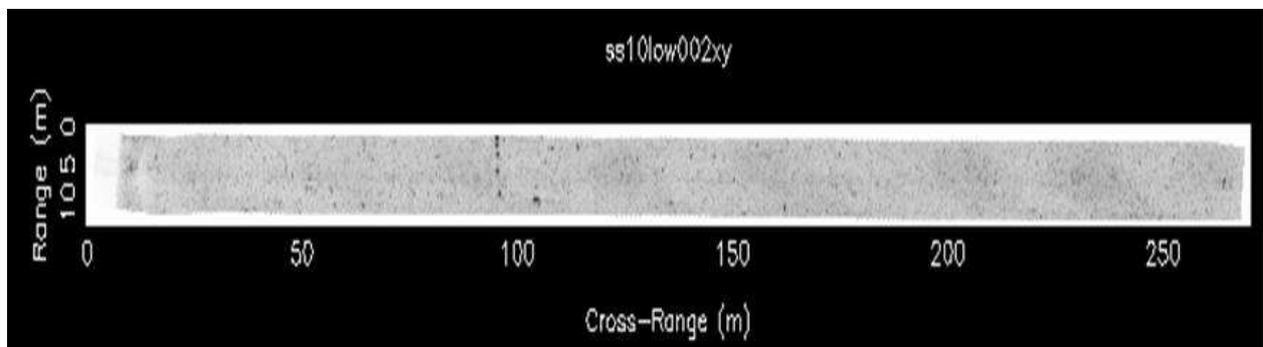


Figure 160: This figure shows P_{xy} (maximum value over all z planes at each (x,y) point) for an entire run at the sand site. The platform was at a ‘low’ altitude above the seabed (typically 2 to 3 m). Total water depth is approximately 10 m. Exploded views of this run are shown in Figs. 161 and 162.

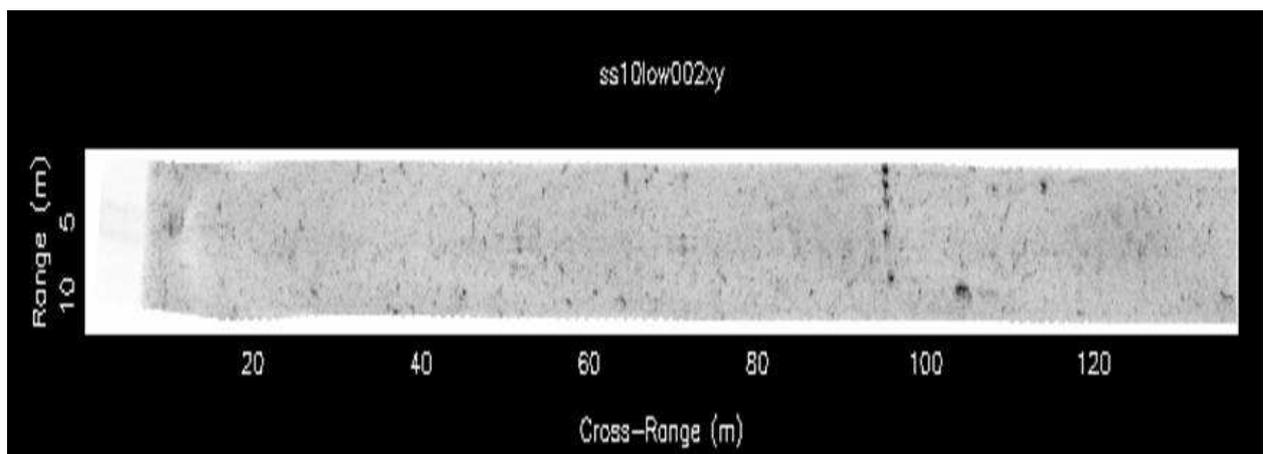


Figure 161: Exploded view of Fig. 160 for the *first half* of the run.

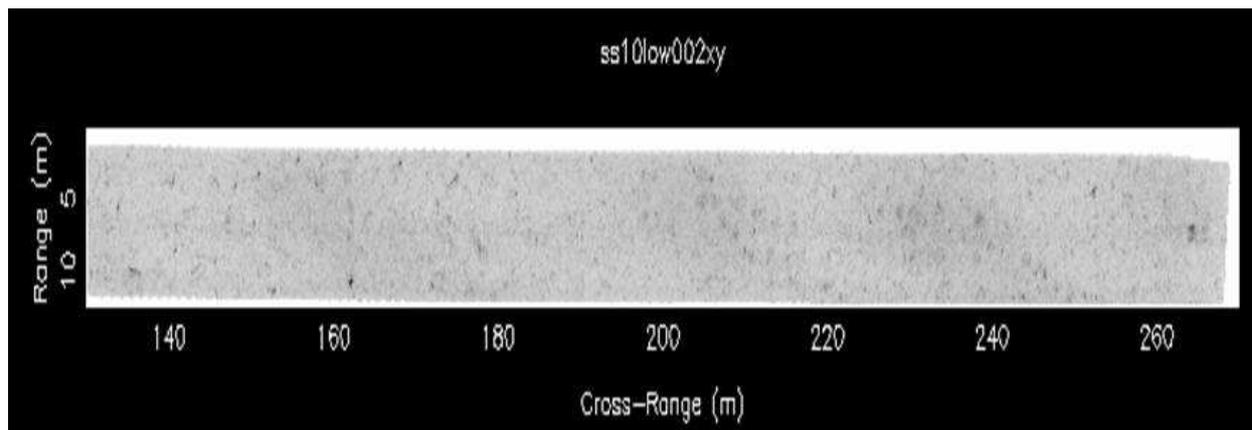


Figure 162: Exploded view of Fig. 160 for the *second half* of the run.

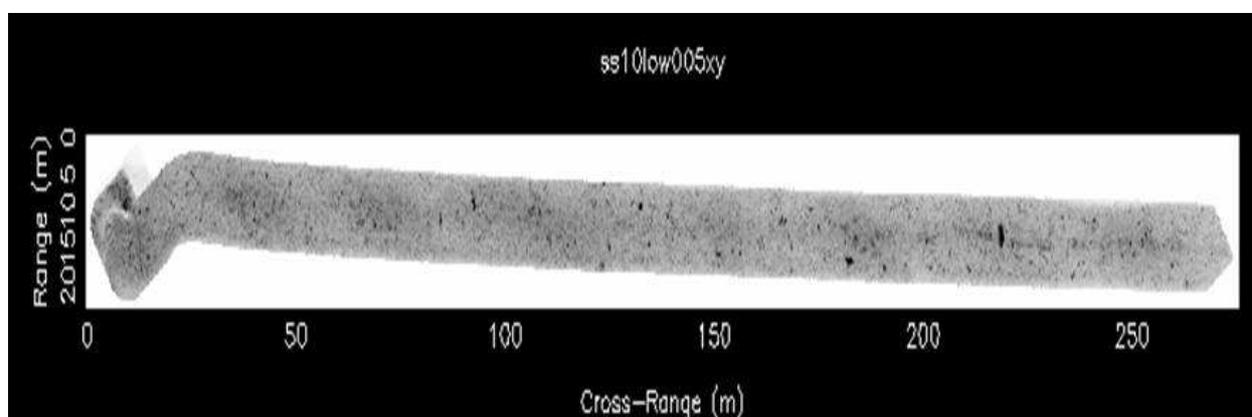


Figure 163: Similar to Fig. 160 but for a second run.

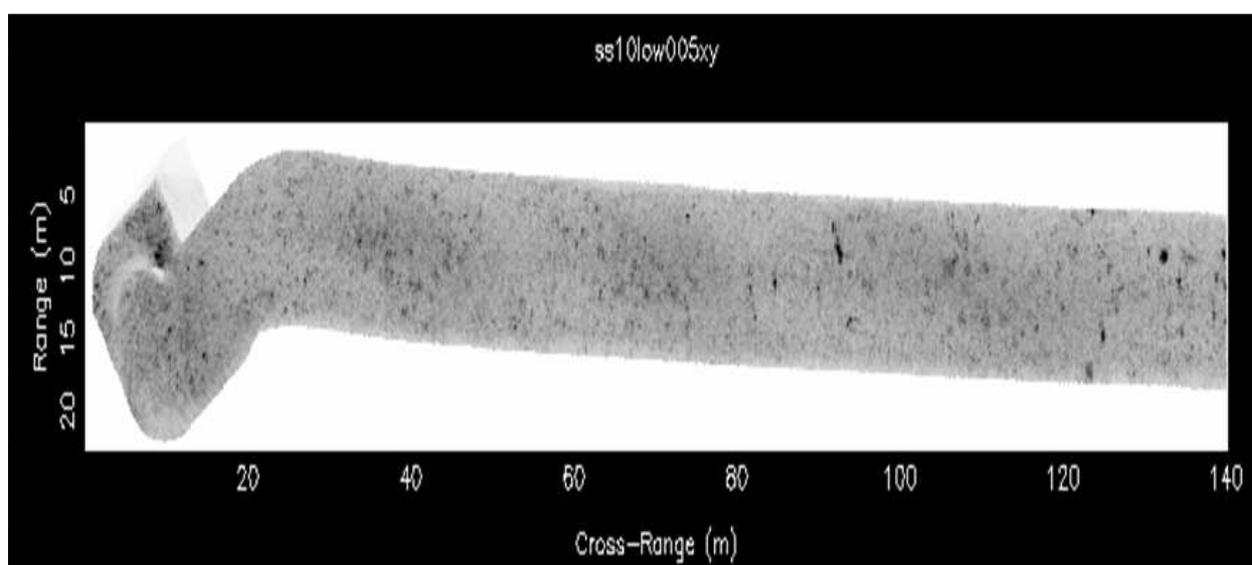


Figure 164: Exploded view of Fig. 163 for the *first half* of the run.

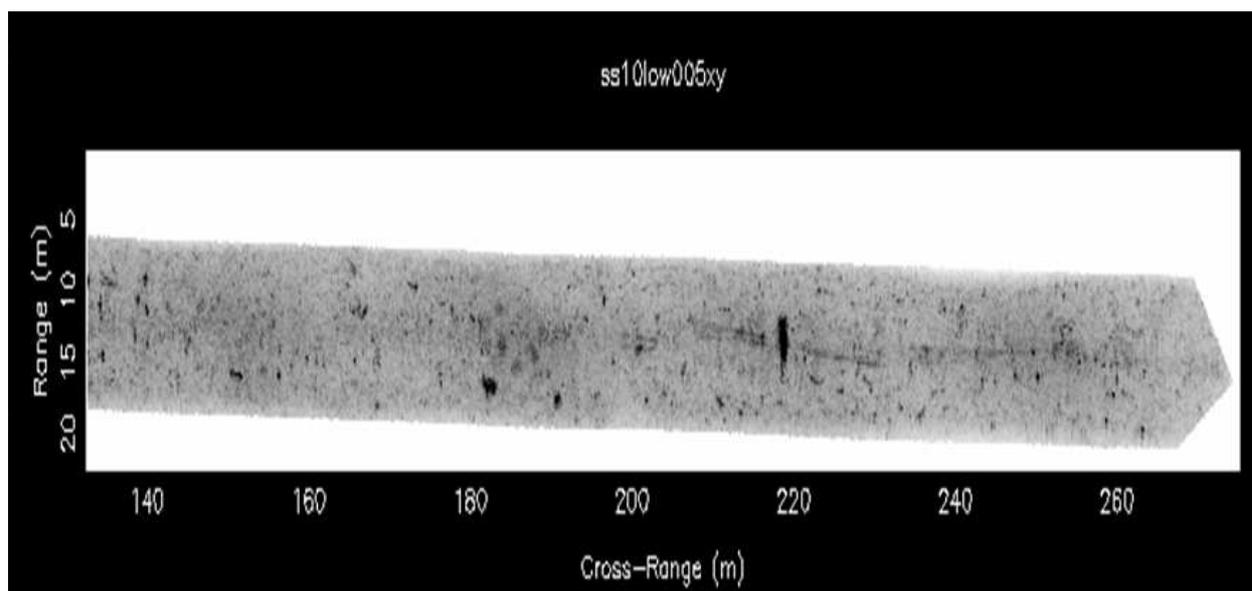


Figure 165: Exploded view of Fig. 163 for the *second half* of the run.

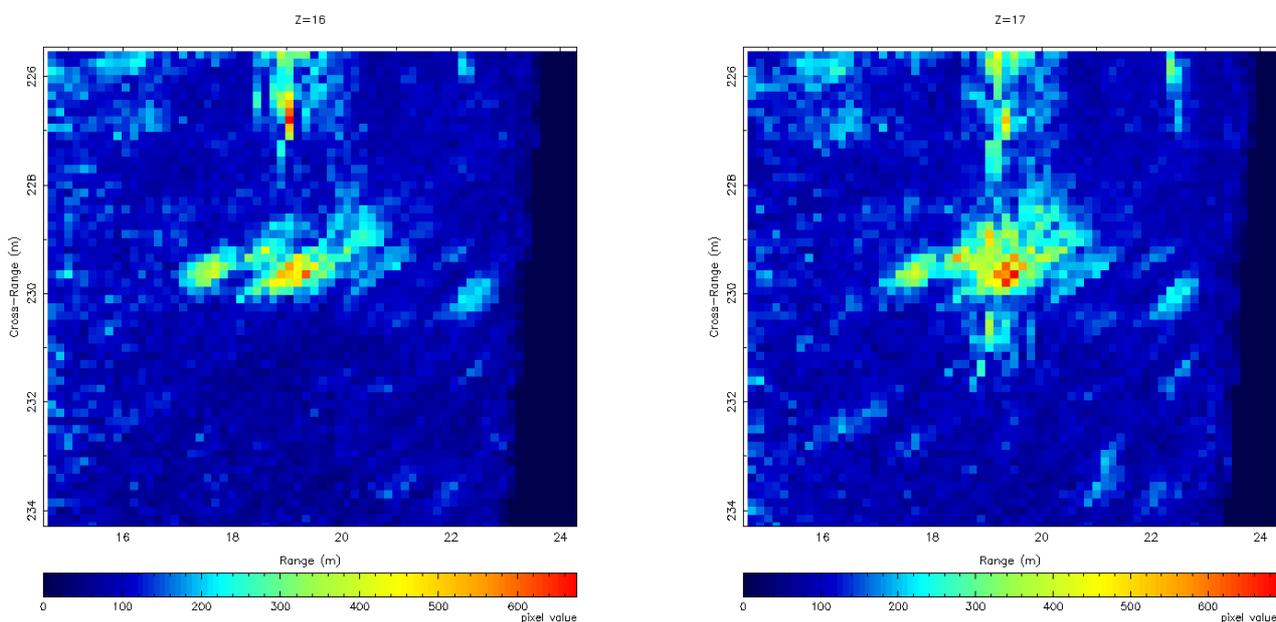


Figure 166: This figure shows a detailed view of $B(\mathbf{x}_F)$ (eq. 111) for a particular target at the mud site. This view shows slices on the z -plane at depth indices 16 and 17.

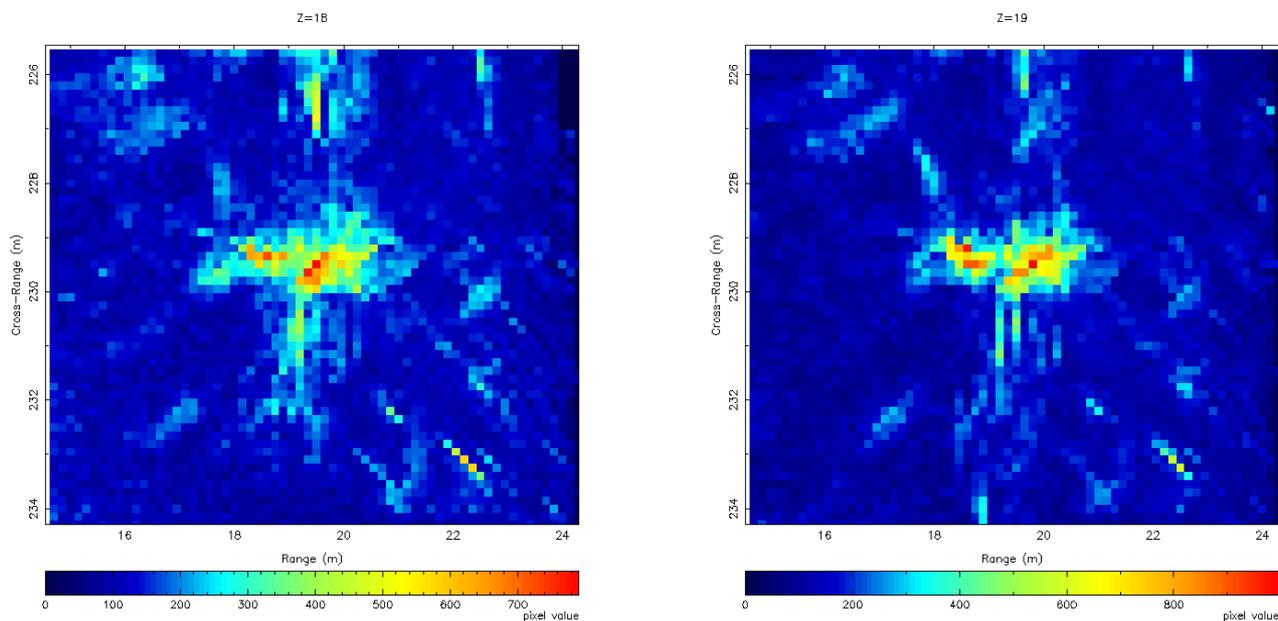


Figure 167: Similar to Fig. 166 but for depth slices 18 and 19.

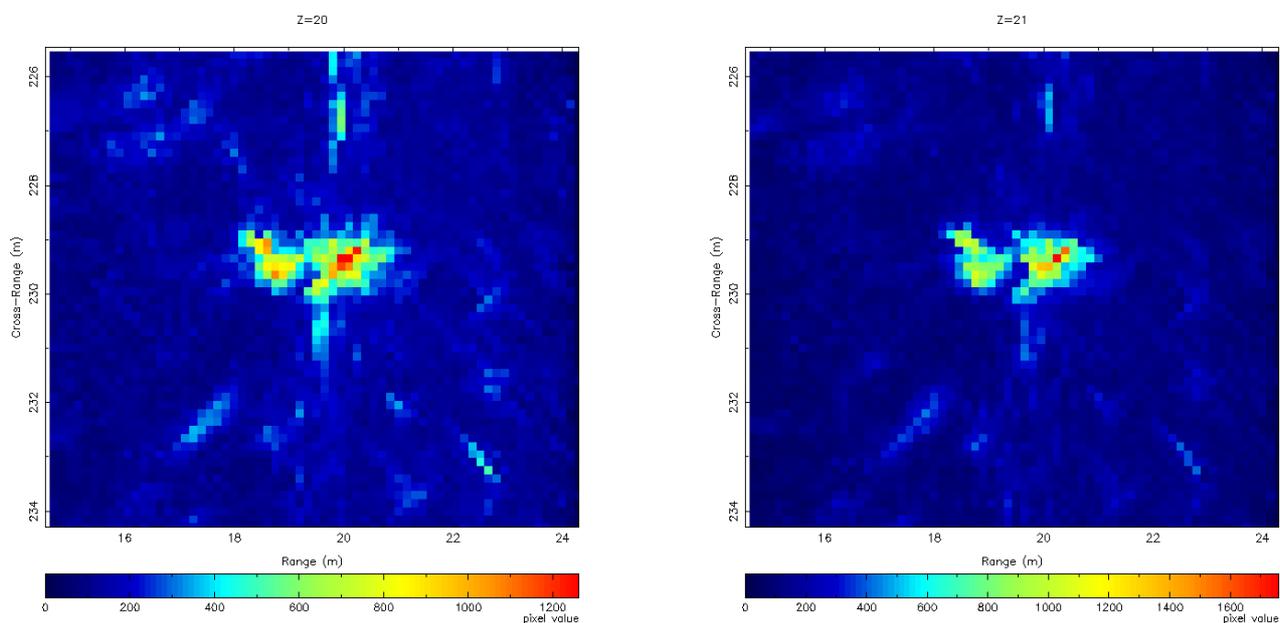


Figure 168: Similar to Fig. 166 but for depth slices 20 and 21.

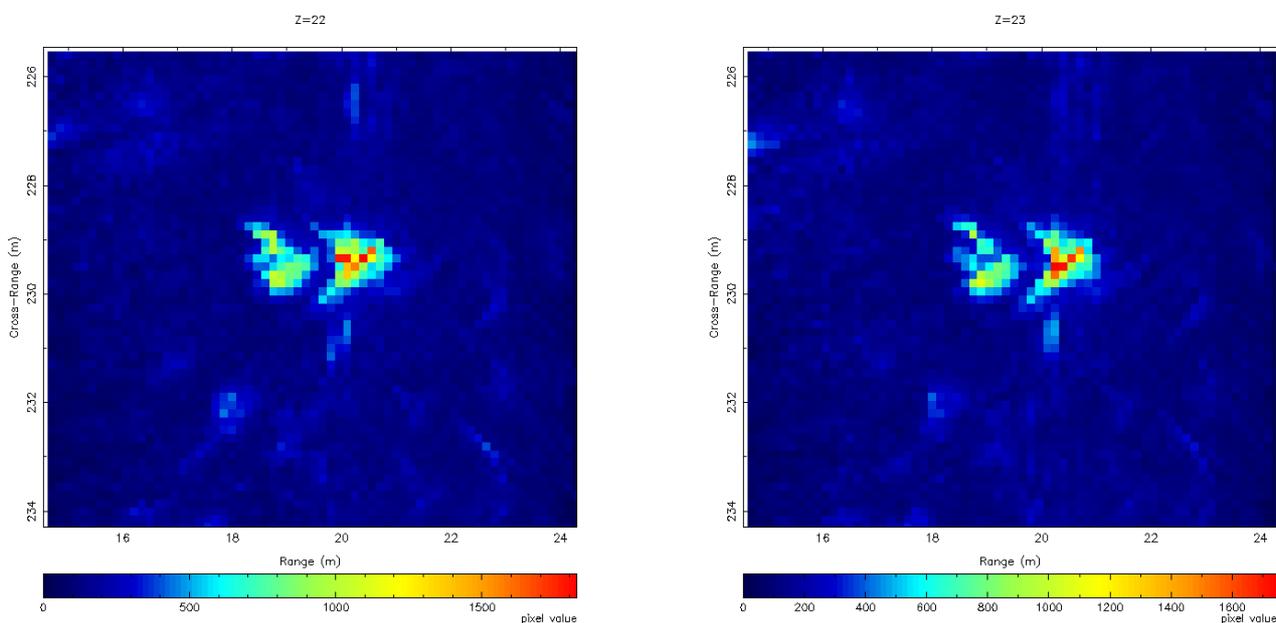


Figure 169: Similar to Fig. 166 but for depth slices 22 and 23.

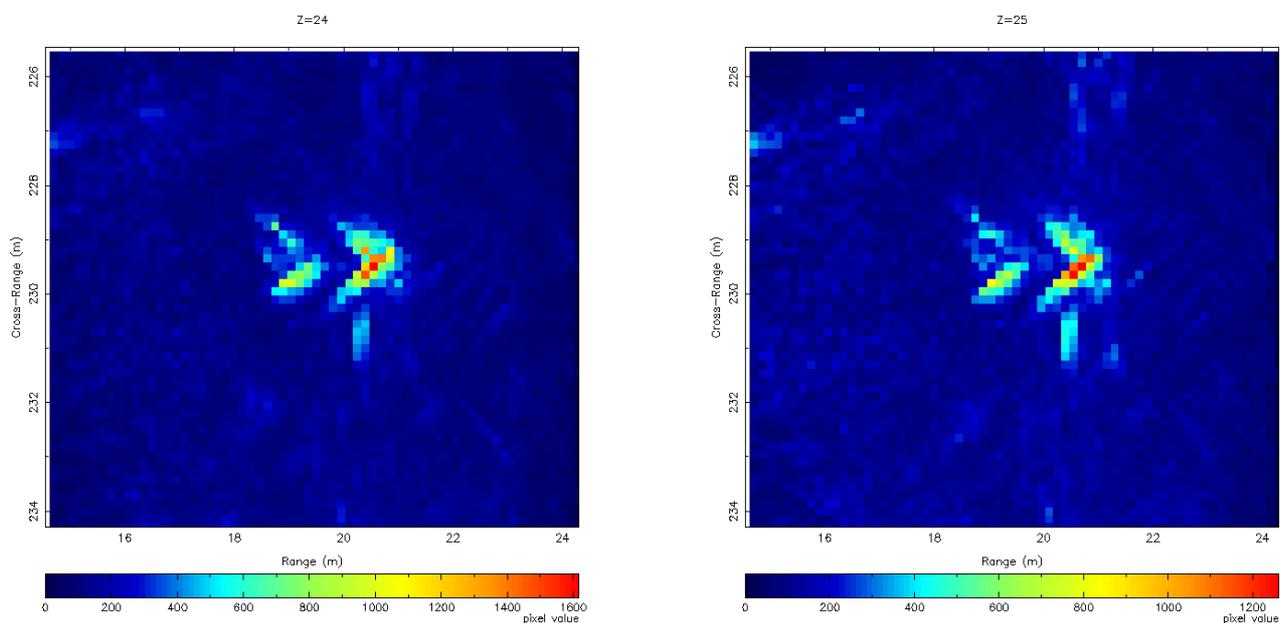


Figure 170: Similar to Fig. 166 but for depth slices 24 and 25.

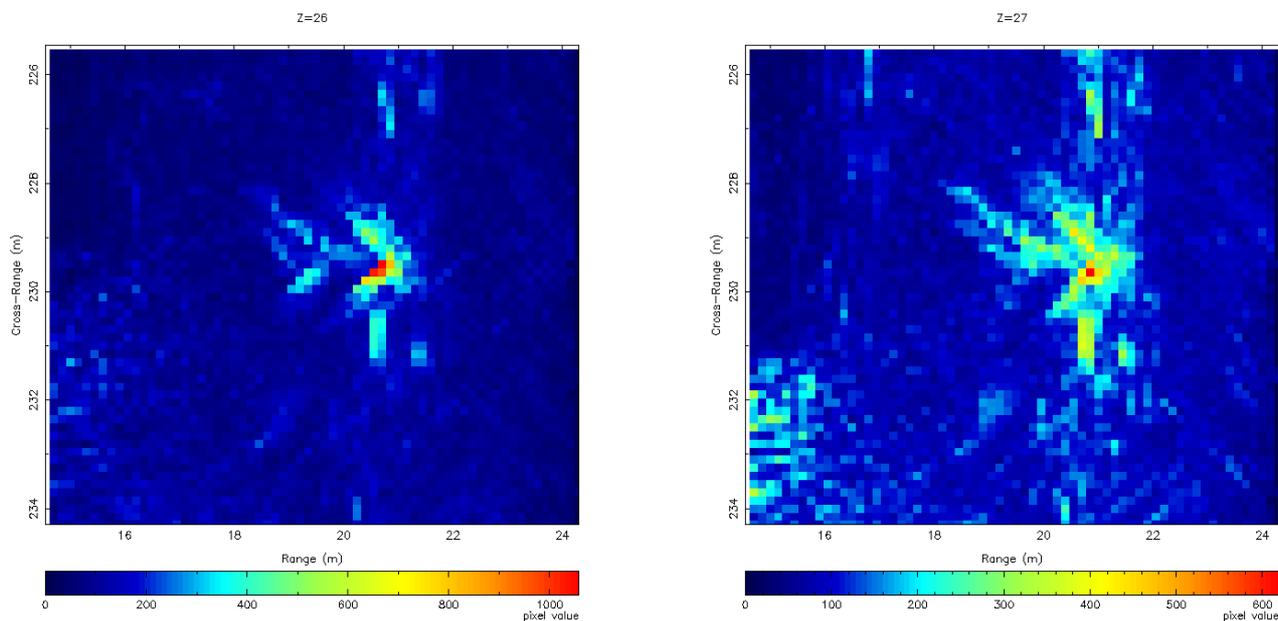


Figure 171: Similar to Fig. 166 but for depth slices 26 and 27.

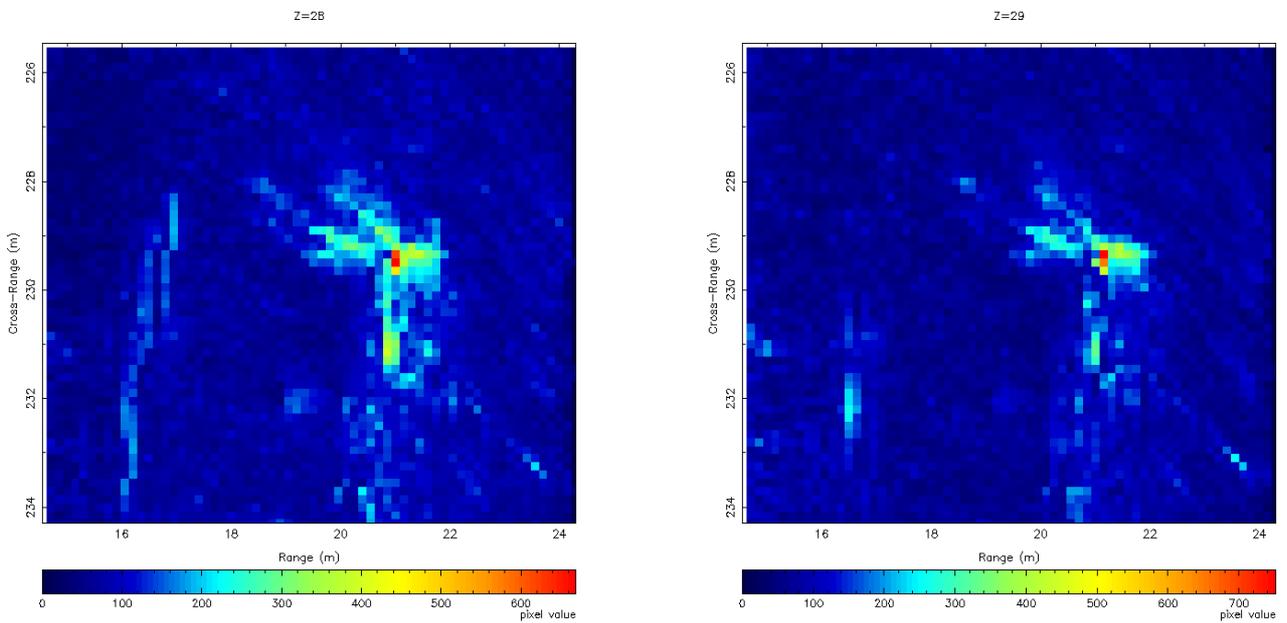


Figure 172: Similar to Fig. 166 but for depth slices 28 and 29.

B List of Scientific/Technical Publications

B.1 Papers

The following publications were supported by this SERDP program:

1. C. Morency and J. Tromp, Spectral-element simulations of wave propagation in porous media, *Geophys. J. Int.*, vol. 175, pp. 301–345 (2008).
2. C. Morency, Y. Luo and Jeroen Tromp, Finite-frequency kernels for wave propagation in porous media based upon adjoint methods, submitted to *Geophys. J. Int.*, accepted 2009 July 10.
3. C. Morency, Y. Luo and Jeroen Tromp, Spectral-Element Simulations of Wave Propagation in Porous Media: Finite-Frequency Sensitivity Kernels based upon Adjoint Methods, in *POROMECHANICS IV (Proceedings of the Fourth Biot Conference on Poromechanics)*, DEStech Publications, Inc., Pennsylvania, USA.

B.2 Technical abstract for SERDP/ESTCP Annual Symposium (2006)

SYNTHETIC APERTURE SONAR FOR DETECTION AND DISCRIMINATION OF UNDERWATER UXO

EUGENE M. LAVELY
BAE Systems AIT
6 New England Executive Park
Burlington, MA 01803
(781) 273-3388 Ext. 294
eugene.lavely@baesystems.com

CO-PERFORMERS: Dr. Angela Putney (Applied Signal Technology, Inc.)
Professor Jeroen Tromp (California Institute of Technology, Department of Geological and Planetary Sciences)

Abstract

We are pursuing a combined simulation and observational study to assess the potential performance of sonar-based discrimination systems for UXO targets in littoral and other shallow water environments. Sonar is a natural candidate for UXO detection due to its wide-area surveillance capability and target sensitivity. Specifically, we are analyzing synthetic aperture sonar (SAS) imagery for both buried and bottom surface targets to determine repeatability of discriminating features, especially given the complex wave propagation characteristics in the water column and sediments. Environment complexity can couple the target response to the detailed variation of the embedding medium (including bulk properties and surface properties such as sediment ripples). To better understand this we intend to apply a high-fidelity wave simulation for modeling of wave propagation in heterogeneous porous media and for targets

embedded in that media. The resulting sonar phase histories will then be processed with a SAS image formation processor for image generation. We are using the Biot formulation of the wave equations for the simulations and we are assigning realistic values of Biot propagation parameters using experimental characterizations derived from previous studies. The Spectral Element Method (SEM) is currently being evaluated as the numerical method for the wave solution, but in principle, other methods are applicable as well (e.g. finite difference, time domain). To be useful, the numerical method must capture arbitrary target/environment complexity and accurately model all wave phenomena (resonant modes, surface waves, diffractions, specular scattering, target-environment coupling) over the low and high frequency regimes of interest. To support this effort we are in the process of building target and environment models. This activity, along with the results data products, target signatures and features from both simulations and real data collections will provide an outstanding resource for evaluation, design and development of target recognition systems. This work is supported under the Munitions Management (MM) focus area of SERDP as Project MM-1533.

B.3 Technical abstract for SERDP/ESTCP Annual Symposium (2007)

SYNTHETIC APERTURE SONAR FOR DETECTION AND DISCRIMINATION OF UNDERWATER UXO

EUGENE M. LAVELY
BAE Systems AIT
6 New England Executive Park
Burlington, MA 01803
(781) 273-3388 Ext. 4294
eugene.lavely@baesystems.com

CO-PERFORMERS: Matthew Nelson (Applied Signal Technology, Inc.)
Professor Jeroen Tromp and Christina Morency (California Institute of Technology, Department of Geological and Planetary Sciences)

Abstract

We are pursuing a combined simulation and observational study to assess the potential performance of sonar-based discrimination systems for UXO targets in littoral and other shallow water environments. Sonar is a natural candidate for UXO detection due to its wide-area surveillance capability and target sensitivity. Specifically, we are analyzing synthetic aperture sonar (SAS) imagery for both buried and bottom surface targets to determine repeatability of discriminating features, especially given the complex wave propagation characteristics in the water column and sediments. Environment complexity can couple the target response to the detailed variation of the embedding medium (including bulk properties and surface properties such as sediment ripples). To better understand this we intend to apply a high-fidelity wave simulation for modeling of acoustic wave propagation, wave propagation in heterogeneous porous media and wave interaction with targets embedded in that media. The resulting sonar phase histories will then be processed with a SAS image formation processor for image generation. We are using the Biot formulation of the wave equations in the porous medium, derived

using the averaging principle in a biphasic system, for the simulations and we are assigning realistic values of Biot propagation parameters using experimental characterizations derived from previous studies. For the acoustic equation of motion, a displacement potential formulation is used, which results in an acoustic-poroelastic coupling based on a non-iterative time scheme. The Spectral Element Method (SEM) is currently being evaluated as the numerical method for the wave solution. In principle, other methods are applicable as well (e.g., finite difference, time domain), but the SEM is a high-order variational method, which allows to compute synthetic seismograms in heterogeneous models with deformed geometry with very high accuracy and low numerical dispersion. To be useful, the numerical method must capture arbitrary target/environment complexity and accurately model all wave phenomena (resonant modes, surface waves, diffractions, specular scattering, target-environment coupling) over the low and high frequency regimes of interest. To support this effort we are in the process of building target and environment models. This activity, along with the results data products, target signatures and features from both simulations and real data collections will provide an outstanding resource for evaluation, design and development of target recognition systems. This work is supported under the Munitions Management (MM) focus area of SERDP as Project MM- 1533.

We are pursuing a combined simulation and observational study to assess the potential performance of sonar-based discrimination systems for UXO targets in littoral and other shallow water environments. Sonar is a natural candidate for UXO detection due to its wide-area surveillance capability and target sensitivity. Specifically, we are analyzing synthetic aperture sonar (SAS) imagery for both buried and bottom surface targets to determine repeatability of discriminating features, especially given the complex wave propagation characteristics in the water column and sediments. Environment complexity can couple the target response to the detailed variation of the embedding medium (including bulk properties and surface properties such as sediment ripples). To better understand this we intend to apply a high-fidelity wave simulation for modeling of wave propagation in heterogeneous porous media and for targets embedded in that media. The resulting sonar phase histories will then be processed with a SAS image formation processor for image generation. We are using the Biot formulation of the wave equations for the simulations and we are assigning realistic values of Biot propagation parameters using experimental characterizations derived from previous studies. The Spectral Element Method (SEM) is currently being evaluated as the numerical method for the wave solution, but in principle, other methods are applicable as well (e.g. finite difference, time domain). To be useful, the numerical method must capture arbitrary target/environment complexity and accurately model all wave phenomena (resonant modes, surface waves, diffractions, specular scattering, target-environment coupling) over the low and high frequency regimes of interest. To support this effort we are in the process of building target and environment models. This activity, along with the results data products, target signatures and features from both simulations and real data collections will provide an outstanding resource for evaluation, design and development of target recognition systems. This work is supported under the Munitions Management (MM) focus area of SERDP as Project MM-1533.

B.4 Technical abstract for SERDP/ESTCP Annual Symposium (2008)

UNDERWATER BURIED MUNITIONS DETECTION AND CLASSIFICATION WITH SONAR-BASED FEATURES

EUGENE M. LAVELY
BAE Systems AIT
6 New England Executive Park
Burlington, MA 01803
(781) 273-3388 Ext. 4294
eugene.lavelly@baesystems.com

CO-PERFORMERS: Matthew Nelson and Kent Harbaugh (Applied Signal Technology, Inc.); Jeroen Tromp and Christina Morency (Princeton University)

Abstract

We develop an underwater munitions classifier algorithm with a special focus on buried objects using bottom-penetrating low-frequency sonar data. Objects in the classifier training set include various munition targets, man-made clutter objects such as concrete pipes, cylinders and spheres, and natural clutter responses. One of our key criteria for feature selection is robustness to natural variability, and to specific details of target-sensor geometry. An additional issue for feature identification is not just the class of feature (e.g., shape-based, intensity histogram, waveform-derived, tomographic estimates of material properties, etc.) but also the tuning parameters used to extract the feature. Effective choice of these parameters can reduce false alarms. We use the Relevance Vector Machine as the basic classification tool since it provides a Bayesian framework, a means for principled estimation of posterior class membership probabilities, and can be used to automatically determine the relevance of each feature to the identification of class membership probability. We report on classifier performance using both real data sets and synthetic data sets generated with a high-fidelity simulation tool based on a recently developed spectral element method. Synthetic responses from both acoustic-elastic and poroelastic simulations are considered. This project is supported under SERDP Project MM-1533.

B.5 Technical abstract for American Geophysical Union Meeting (Winter, 2008)

Spectral-Element Simulations of Wave Propagation in Porous Media: Finite-Frequency Sensitivity Kernels Based Upon Adjoint Methods

Christina Morency, Yang Luo and Jeroen Tromp - Princeton University Department of Geosciences, 114 Guyot Hall, Princeton, NJ 08544-1003, United States

The mathematical formulation of wave propagation in porous media developed by Biot is based upon the principle of virtual work, ignoring processes at the microscopic level, and does not explicitly incorporate gradients in porosity. Based on recent studies focusing on averaging techniques, we derive the macroscopic porous medium equations from the microscale, with a particular emphasis on the effects of gradients in porosity. In doing so, we are able to naturally determine two key terms in the momentum equations and constitutive relationships, directly translating the coupling between the solid and fluid phases, namely a drag force and an interfacial strain tensor. In

both terms, gradients in porosity arise. One remarkable result is that when we rewrite this set of equations in terms of the well known Biot variables (u_s, \mathbf{w}) , terms involving gradients in porosity are naturally accommodated by gradients involving \mathbf{w} , the fluid motion relative to the solid, and Biot's formulation is recovered, i.e., it remains valid in the presence of porosity gradients. We have developed a numerical implementation of the Biot equations for two-dimensional problems based upon the spectral-element method (SEM) in the time domain. The SEM is a high-order variational method, which has the advantage of accommodating complex geometries like a finite-element method, while keeping the exponential convergence rate of (pseudo)spectral methods. As in the elastic and acoustic cases, poroelastic wave propagation based upon the SEM involves a diagonal mass matrix, which leads to explicit time integration schemes that are well-suited to simulations on parallel computers. Effects associated with physical dispersion and attenuation and frequency-dependent viscous resistance are addressed by using a memory variable approach. Various benchmarks involving poroelastic wave propagation in the high- and low-frequency regimes, and acoustic-poroelastic and poroelastic-poroelastic discontinuities have been successfully performed. We present finite-frequency sensitivity kernels for wave propagation in porous media based upon adjoint methods. We first show that the adjoint equations in porous media are similar to the regular Biot equations upon defining an appropriate adjoint source. Then we present finite-frequency kernels for seismic phases in porous media (e.g., fast P, slow P, and S). These kernels illustrate the sensitivity of seismic observables to structural parameters and form the basis of tomographic inversions. Finally, we show an application of this imaging technique related to the detection of buried landmines and unexploded ordnance (UXO) in porous environments.