Classification of underwater targets from autonomous underwater vehicle sampled bistatic acoustic scattered fields

Erin M. Fischell and Henrik SchmidtAMT

Articles you may be interested in

Environmental effects on seabed object bistatic scattering classification
The Journal of the Acoustical Society of America **141**, (2017); 10.1121/1.4972273

Acoustic identification of buried underwater unexploded ordnance using a numerically trained classifier (L)
The Journal of the Acoustical Society of America **132**, (2012); 10.1121/1.4763997

Time-frequency analysis of the bistatic acoustic scattering from a spherical elastic shell
The Journal of the Acoustical Society of America **131**, (2012); 10.1121/1.3669995

Classification of underwater targets from autonomous underwater vehicle sampled bistatic acoustic scattered fields

Erin M. Fischell[a)] and Henrik Schmidt

Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA

One of the long term goals of autonomous underwater vehicle (AUV) minehunting is to have multiple inexpensive AUVs in a harbor autonomously classify hazards. Existing acoustic methods for target classification using AUV-based sensing, such as sidescan and synthetic aperture sonar, require an expensive payload on each outfitted vehicle and post-processing and/or image interpretation. A vehicle payload and machine learning classification methodology using bistatic angle dependence of target scattering amplitudes between a fixed acoustic source and target has been developed for onboard, fully autonomous classification with lower cost-per-vehicle. To achieve the high-quality, densely sampled three-dimensional (3D) bistatic scattering data required by this research, vehicle sampling behaviors and an acoustic payload for precision timed data acquisition with a 16 element nose array were demonstrated. 3D bistatic scattered field data were collected by an AUV around spherical and cylindrical targets insonified by a 7–9 kHz fixed source. The collected data were compared to simulated scattering models. Classification and confidence estimation were shown for the sphere versus cylinder case on the resulting real and simulated bistatic amplitude data. The final models were used for classification of simulated targets in real time in the LAMSS MOOS-IvP simulation package [M. Benjamin, H. Schmidt, P. Newman, and J. Leonard, J. Field Rob. **27**, 834–875 (2010)]. © *2015 Acoustical Society of America.*
[http://dx.doi.org/10.1121/1.4938017]

I. INTRODUCTION

A growing application for autonomous underwater vehicle (AUV) technology is the localization, classification, and mitigation of underwater hazards in shallow harbor environments. The classification problem has attracted particular attention in recent years with the development of visual and acoustic AUV-based sensors for remote data collection. Because visual inspection of targets can be difficult or impossible in murky harbors and requires precise target localization, acoustic sensors such as sidescan sonar and synthetic aperture sonar (SAS) are used more extensively for AUV-based mine countermeasures missions. While these techniques can provide rich images of targets and the environment, they are difficult to use for real-time target classification and require expensive precision sensors. Sidescan sonar images require expert interpretation and are not useful for locating and classifying buried targets. SAS images are usually computed in post-processing so that navigation corrections may be applied.[1] Both generally rely on image processing for classification assistance, which is difficult with the limited computational power on many AUVs. In addition to these challenges to fully autonomous real-time classification of data from these systems, the sensors themselves are too expensive to be practical in multi-vehicle operations.

To achieve plausible, real-time AUV-based target classification that is expandable to distributed vehicle networks, two key advancements are required: an inexpensive sensing payload and a classification method that can be run in real time on an AUV computer using onboard processing of sensor and navigation data. The advantage of such a sensing system would be the ability to deploy multiple AUVs to carry out the target localization and classification missions with immediate classification and confidence estimates to inform prosecution decisions without having to recover and redeploy vehicles. This paper presents a bistatic acoustic sensing approach to this problem. In this scenario, shown in Fig. 1, a fixed acoustic source insonifies underwater targets and scattering data for target classification are collected using AUV-based linear hydrophone nose arrays cut for sensing at 1–15 kHz. The bistatic scattering data are processed onboard each vehicle for target localization and classification. The bistatic configuration and hydrophone array were selected to limit sensing system cost per vehicle.

It was critical to identify features within this bistatic scattering data that were robust to several meters of error in vehicle location, source location, and target location. The combined navigational uncertainty, plus the computations limitations for data processing on an AUV, made using sensitive time and phase information for target classification impractical. While these features are frequently used for SAS imaging, they would be difficult to use in real time on a bistatic AUV system.

When a target on the ocean bottom is acoustically insonified, the target re-radiates the signal (Fig. 2). This re-radiation consists of multiple time delayed echoes that interfere in the
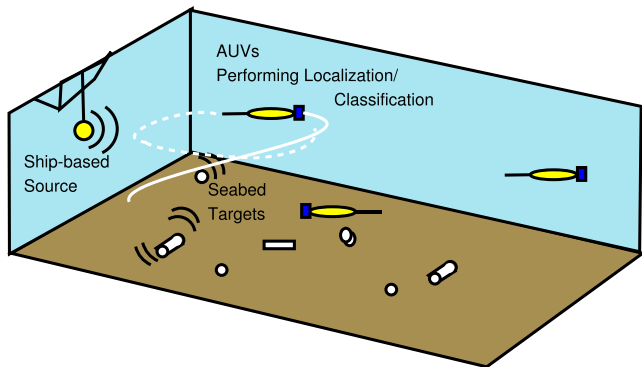
a)Electronic mail: emf43@mit.edu

FIG. 1. (Color online) Multi-vehicle operation mission, where a fixed source insonifies a target field while multiple AUVs sample the bistatic scattering fields around various targets.

frequency domain. The interference of the time-delayed echoes from target scattering result in frequency-dependent minima and maxima in the bistatic radiation pattern from the target. These scattering radiation patterns are distinct for different target types and are mostly dependent on azimuth, showing range and depth independence over meters to tens of meters. The concept for the classification techniques discussed in this paper is that these interference patterns in a given frequency band are stable and can be used to characterize seabed targets. This bistatic radiation pattern can be determined using an AUV with a linear hydrophone nose array, a data acquisition system, and signal processing software to calculate target scattering amplitude as acoustic data are collected around a target. Imaging techniques are not required, as the angular dependence of bistatic scattering amplitude can be analysed directly.

A scattering simulation package was used for modeling target scattering fields expected from bistatic scattering experiments. Figures 3(a) and 3(b) show the simulated intensity-averaged radiation pattern for spherical and cylindrical targets versus azimuthal angle relative to the source for sampling locations 10–60 m from the target. Represented in this fashion, the difference between the two target types is very clear, suggesting a good basis for AUV-based target classification.

Most of the experimental work on bistatic target scattering has been conducted in water tanks and test ponds. For
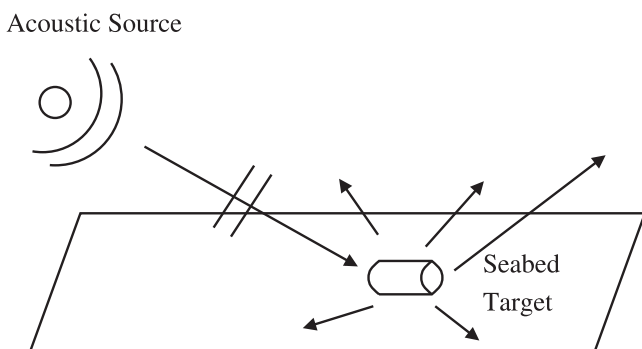

FIG. 2. Insonification of a target results in acoustic scattering, as the target re-radiates the signal in multiple echoes that interfere to form the radiation pattern exploited by the characterization techniques discussed in this paper.

example, Baik, Dudley, and Marston conducted an experiment where they looked at the bistatic response of different cylinders in a test tank for the purposes of holographic imaging.[2] Haumesser et al.[3] took bistatic measurements of scattering from an air-filled elastic cylinder in a test tank to study vibration modes. Another experiment, looking at bistatic scattering from realistic unexploded ordinance geometries, is described in Waters et al.[4] Kargl et al.[5] looked at the bistatic scattering response of aspect-dependent targets in a test pond as a part of the PondEx10 experiment. Other cases of small and full-scale bistatic tank experiments are reviewed in Blondel and Pace.[6]

There have been few seabed target bistatic scattering experiments in real-world harbor environments, and no examples were located describing the target classification approach using bistatic angle mapping of target scattering amplitude for classification. An example of a bistatic scattering experiment in a real-world environment was part of the SITAR project, conducted in Sweden in 2003.[6,7] During this experiment, scattering data were collected around targets located in a dump site using a fixed receiver and a remotely operated vehicle (ROV) based transmitter. The frequencies were higher than those discussed in this paper and the ROV allowed more precise navigation than that possible with the AUV used here, but the concept of having a mobile transmitter and a fixed receiver is very similar.

The GOATS'98 experiment is a rare example of a successful AUV-based bistatic scattering experiment: it included an AUV with a nose array and produced data on the bistatic scattered fields off of fully buried, partially buried, and proud spheres. LePage and Schmidt[8] and Edwards et al.[9] described the AUV experiment and used the array data for SAS imaging. Bistatic scattering of seabed targets with simple geometries was also studied in simulation by Schmidt and Lee,[10] who described the distinguishing characteristics of bistatic scattered fields of different targets and environments. The application of bistatic target modeling data to target identification was discussed by Zampolli et al.,[11] including the possiblity of using modeled monostatic or bistatic target strengths for the design of acoustic waterside security systems.

Machine learning target classification, using acoustic data that include temporal or phase information from monostatic sensing, has been described by Kaminsky and Barbu (buried cylindrical targets such as cables)[12] and Malarkodi et al.[13] (neural networks for classification of target type). A probabilistic method for monostatic target classification is discussed in Ref. 14, which attempts to classify targets using multi-aspect backscatter, wave-based signal processing, and hidden Markov models (HMMs). These techniques differ from those described in this paper in that they use features that include temporal or phase information and only look at backscatter.

This paper describes the AUV payload used to perform bistatic acoustic data collection, the real-world bistatic acoustic data sets collected around spherical and cylindrical seabed targets with that payload during a bistatic scattering experiment, and a machine-learning methodology that utilizes bistatic-angle-dependent amplitude features

Erin M. Fischell and Henrik Schmidt
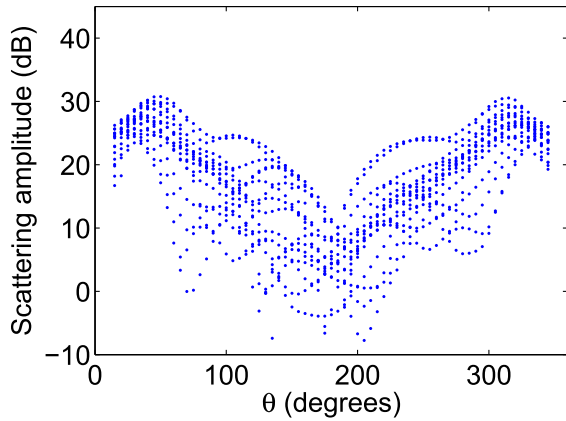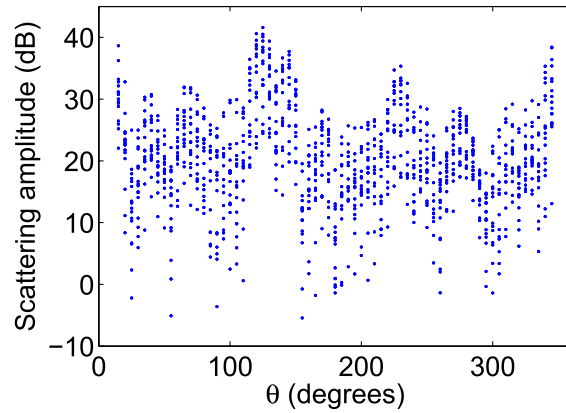
(a) Spherical target.

(b) Cylindrical target.

FIG. 3. (Color online) Simulated scattering amplitude dependence on angle θ for spherical and cylindrical targets. θ is calculated by setting the target at $(0, 0)$ and the source at $(-60, 0)$ such that the source is at $180°$. Amplitudes shown here are for multiple receiver depths and ranges to the target. (a) Spherical target. (b) Cylindrical target.

from the scattered field to classify target shape. First, the AUV experiment used to collect bistatic scattering data is presented, including the experiment setup, vehicle, payload, and signal processing methodology. The machine learning target classification methodology, using bistatic angle dependence of target scattering amplitude, is then described. The novel experimental data set is then presented and compared to simulation results. Finally, the classification results on these data are described and conclusions discussed.

II. BISTATIC SCATTERING DATA COLLECTION

To determine the feasibility of target classification using bistatic angle dependence of target scattering amplitude, it was necessary to collect a real-world bistatic scattering data set. This data set was used to demonstrate signal processing to calculate target scattering amplitude, machine learning model generation, and actual target classification.

A. Experiment parameters

The bistatic scattering data collection was conducted in St. Andrews Bay near Panama City, FL on May 21, 2014 as a part of the BayEx'14 scattering experiment. The configuration of the experiment is shown in Fig. 4. Two targets, a 0.6 m diameter empty steel spherical shell with shell thickness of 0.0168 m and a 3-to-1 solid aluminum cylinder with 0.3048 m diameter, were deployed about 60 m from the ends of the source rail. The source was set at the north end of the rail for the morning to insonify the sphere and the south end of the rail for the afternoon to insonify the cylinder. The source fired a 10 ms 7–9 kHz LFM chirp on a 1 Hz schedule synchronized to global positioning system (GPS) pulse per second (PPS). The water depth was 6–7 m with a mud bottom over sand.

B. Vehicle

The Bluefin 21-inch AUV *Unicorn* was used for data collection (Fig. 5). *Unicorn* is a 3 meter long, 21 in. (0.5334 m) diameter AUV that was outfitted for this experiment with an acoustics and autonomy payload that included a 16 element nose array with 0.05 m element spacing, precision timing/data acquisition hardware, and a computer for autonomy and signal processing. The vehicle also carried a Sea-Bird Electronics model SBE 37-SI conductivity temperature (Bellevue, WA) sensor and a pressure transducer used for depth measurements. The vehicle ran under a front-seat/back-seat control architecture with basic navigation and sensor fusion handled by the front-seat computer and vehicle autonomy, acoustic communications, and acoustic processing handled on the back-seat computer with processes in MOOS, IvP Helm[15] and Goby.[16]
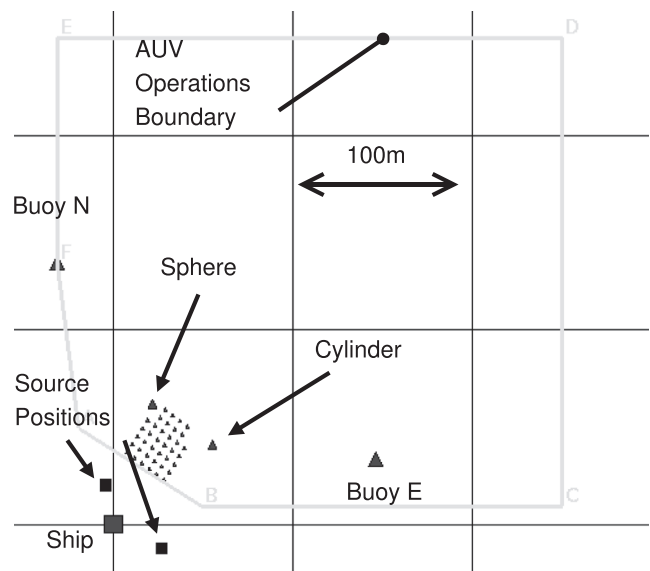

FIG. 4. Experimental setup with source positions, target positions, and AUV operational box.

J. Acoust. Soc. Am. **138** (6), December 2015

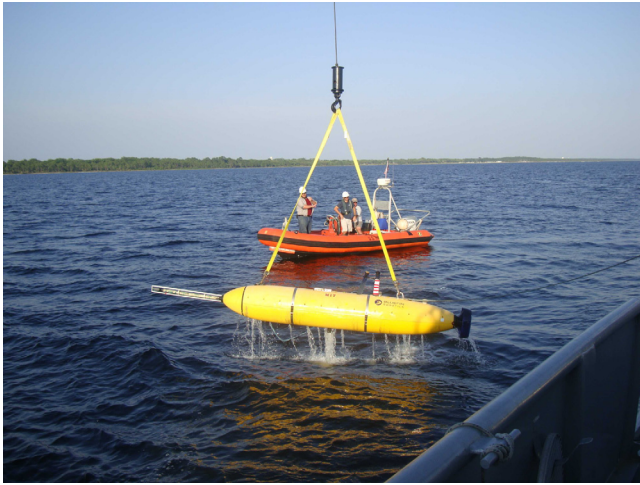Erin M. Fischell and Henrik Schmidt   3775

FIG. 5. (Color online) The AUV *Unicorn* being lifted from the water by the crane of the PCS-12 during the BayEx'14 experiment.

Good navigation and adaptive autonomy were critical for vehicle safety in this experiment because the region for acoustic sampling was near to the source rail and less than $40\,\text{m} \times 50\,\text{m}$. The vehicle's navigation sensors included a Teledyne RD Instruments Doppler velocity log (DVL) (Poway, CA), GPS, a Leica DMC-SX (Norcross, GA) magnetic compass, and a Honeywell HG1700 inertial measurement unit (IMU) (Morris Plains, NJ). The Honeywell IMU was recently installed to improve the navigation of the vehicle while submerged: the previous system resulted in a navigational drift of 1%–5% of distance traveled.[17] The navigational drift with the improved instrumentation was between 0.3% and 0.5% of the distance traveled between GPS fixes. The vehicle surfaced for GPS every 10 min to prevent drift from accumulating significantly.

C. Precision timing

To collect high-quality acoustic bistatic data in this experiment, *Unicorn's* payload was updated and calibrated to ensure timing error of $<70\,\mu\text{s}$. Precision timing is required for bistatic data collection because the source and vehicle are not co-located. The source was triggered directly by a GPS PPS signal, but the GPS signal is not available underwater so *Unicorn* required a separate precise and accurate on-board time source for hardware-triggered data acquisition. A PPS signal indicates the start of a second with the rising edge of a duty-cycled square wave and is used as a trigger for clock synchronization.

The timing and data acquisition hardware included a Microsemi SA.45 (Aliso Viejo, CA) chip scale atomic clock (CSAC),[18] two General Standards Corporation 24DSI12-PLL (Huntsville, AL) analog-to-digital data acquisition boards,[19] and a Garmin 15LxW GPS (Canton of Schaffhausen, Switzerland). The CSAC PPS signal was synchronized to GPS PPS while the vehicle was on the surface and was used to trigger data acquisition at the start of each second. The clock on the autonomy computer, an Advantech 3363[20] (Taipei, Taiwan) with dual core Intel Atom processor (Santa Clara, CA), was synchronized to the

CSAC PPS using network time protocol (NTP) and Generic NMEA GPS Receiver.[21,22] Custom daemons on the computer controlled interfacing and synchronization of the hardware and precision recording of data collected on the data acquisition boards.

Binary files, starting exactly at the start of each second as triggered by CSAC PPS, were recorded on the vehicle computer from the data acquisition boards. To further improve accuracy, the analog and digital delays in the system were characterized in a series of experiments and used to calibrate the system. After correction, the final calibrated system had an arrival time error of $<21.3\,\mu\text{s}$ with 95% confidence, and a phase error of $<8.07\,\text{ns}$ with 99% confidence. To put these numbers in context, a timing error of $70\,\mu\text{s}$ will result in a range localization error of $0.1\,\text{m}$. A phase recording error of $1\,\mu\text{s}$ will result in a phase error of $<1\%$ of a wavelength at $9\,\text{kHz}$.

D. Signal processing

The recorded data files on the computer were passed into MOOS-IvP,[15] which provided a convenient framework for signal processing in real time on the vehicle, and for processing in simulation using navigation and acoustic data logged during an experiment. A MOOS process, pActiveTargetProcess, was used to extract the amplitudes for targets at specified locations from a recorded data files. The operations of this process are shown in Fig. 6. The locations of the target and vehicle for each recorded data file are first used to identify a time window for processing. The time window is centered around the expected arrival time, with a length to either side determined by replica length, navigation uncertainty, and uncertainty in soundspeed estimate. The data from all 16 channels are windowed, and then a fast Fourier transform (FFT) is taken using FFT length *NFFT*. Matched filtering with the source replica is used to identify contacts, and beamforming is used to determine the bearing to the targets and the error in that estimate. The process outputs the band-averaged amplitude from the contact with the location (based on bearing and arrival time)
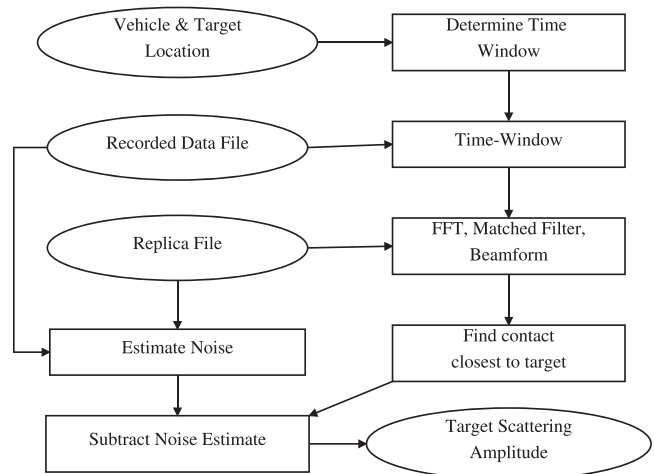

FIG. 6. Processing in pActiveTargetProcess used to extract target amplitudes from the array time series. The recorded data file, vehicle/target location information, and replica are used to estimate the target scattering amplitude.

Erin M. Fischell and Henrik Schmidt

that is closest to the expected target. If the vehicle is in the target's forward scatter region, the target contact cannot be distinguished from the source's direct blast so the process does not produce an amplitude.

Ambient noise was estimated as the band-averaged amplitude from the first *NFFT* samples in the 7–9 kHz band. These samples never included the direct blast or target contacts as the vehicle was always at least 30 m from the source. The estimated noise amplitude was subtracted from the estimated target scattering amplitude associated with each recording to get the estimated target scattering amplitudes used for analysis.

Where decibel levels are shown in this paper based on experimental data, source levels have been estimated and used to calculate scattering strength. While this approach allows for direct comparison with simulation results, it is not practical for real-time classification. Instead, the relative amplitudes are used. A mean and standard deviation for data used by the machine learning model are stored. After two circles of the target, an experimental mean and standard deviation are calculated. The collected experimental data are then normalized using the model mean and standard deviation. Any subsequent data are similarly adjusted. This process allows classification to be performed without perfect knowledge of source level, transmission loss, environmental topography, and data acquisition system gain.

E. AUV deployment

The AUV *Unicorn* was deployed off of the Panama City ship PCS-12, which was anchored on the north end of the operational area shown in Fig. 4. It was then commanded in a sequence of acoustic sampling behaviors around each target. Conventional AUV behaviors, such as lawn mower patterns, are poorly suited for acoustic data acquisition around targets,
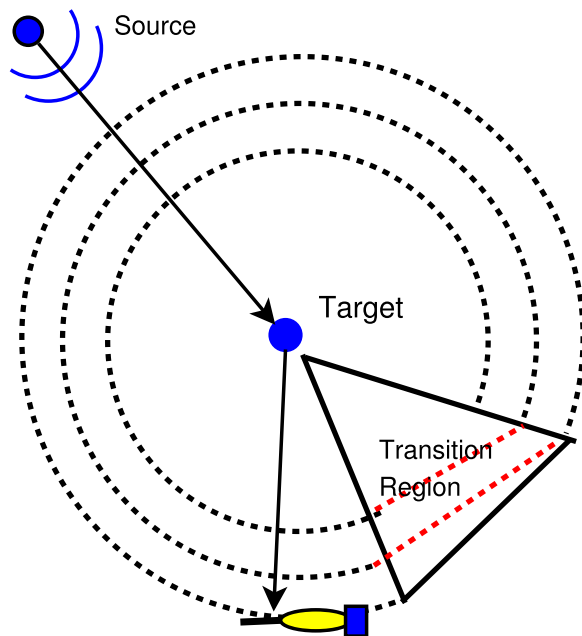

FIG. 7. (Color online) Full field sampling behavior used with the vehicle *Unicorn* for collecting target bistatic data sets. The vehicle circles the target, changing radius in the direct forward-scatter direction.

as the target's contact moves from endfire to broadside and back to endfire. This results in a non-uniform data set. To correct this, a behavior was written to collect a full grid of bistatic amplitudes around a target in depth, range, and azimuth. Figure 7 shows how a single layer of the vehicle path is constructed for this behavior. The vehicle completes a sequence of concentric circles with decreasing radii. By transitioning in radius only in the forward scatter direction, the vehicle goes out of broadside in the region where the target contact cannot be distinguished from the direct blast from the source. This sampling layer is repeated at multiple depths to complete data collection on a target. The behavior is configured using the number of layers sampled, the number of radii sampled, the minimum and maximum depths, the minimum and maximum radii, the minimum permitted distance to the operational boundary, the target location, and the source location.

To ensure *Unicorn* did not hit anything while operating in the tight region between the targets and the source rail, an operation area and obstacle avoidance points were selected to keep it away from the buoys and other collision dangers. In addition, adaptive autonomy built into the acoustic sampling behavior kept the vehicle moving perpendicular to intersection with and at least 10 m away from the operational boundary. Altitude safeties prevented the vehicle from nosing into the bottom by aborting the mission when *Unicorn* got within 2 m of the bottom. The MOOS-IvP[15] infrastructure with Goby[16] interface to acoustic communications meant new commands could be sent without recovering the vehicle to improve sampling or increase vehicle safety in the tight operational area.

III. CLASSIFICATION METHODOLOGY

The goal of this research was to investigate the plausibility of classifying underwater using bistatic scattered amplitudes calculated by an AUV from acoustic data collected between the source and target. We used two data sets for this demonstration, the real bistatic scattering data set collected around spherical and cylindrical targets during the experiment, and a simulation data set matched as closely as possible to the experiment conditions.

A. Machine learning approach

Our approach to the challenge of classifying targets using amplitude-only bistatic acoustic data was to apply supervised machine learning. In a machine learning approach, data are represented using example vectors in a particular feature space and used to train a model that can be used to classify subsequent data. This approach has drawbacks and benefits. Because the method is dependent on well-represented data instead of a physics-based model, it can be more susceptible to "garbage in, garbage out," and poor independent testing can lead to misleadingly good results. However, with sufficient care in problem construction and validation, machine learning can be more powerful than the model-based approach, as it accounts for effects that show up in real data but are neglected in conventional

J. Acoust. Soc. Am. **138** (6), December 2015

Erin M. Fischell and Henrik Schmidt   3777

models. The complexity of bistatic scattering physics makes a data-based approach appealing.

For this problem, we selected a type of supervised machine learning called support vector machines (SVMs). The SVM algorithm was selected for this problem for several reasons. SVMs handle large feature spaces easily, adapt well to different kernels, and have well-implemented off-the-shelf optimization packages. Perhaps most importantly for this real time application, while an SVM model can take significant time and memory to train classification using an existing model is fast. Classification of an example vector using this algorithm also results in a margin, which is an indication of the strength of the classification.

SVM classification works by maximizing the minimum Euclidean distance from a separating hyperplane to the set of training vectors. The soft-margin SVM classification formulation, originally derived by Vapnik,[23] is used

$$\underset{\mathbf{w}, \xi, b}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{w}\|^2$$
$$+ C \sum_{i=1}^{l} \xi_i \text{ subject to} \begin{cases} y_i(\langle \mathbf{w}, \mathbf{x_i} \rangle + b) \geq \xi_i \\ \xi_i \geq 0, \end{cases} \quad (1)$$

where $\mathbf{w}$ is the normal vector to the separating hyperplane that defines the binary classification, ξ is the slack variable that allows the optimization to deviate from perfect classification in the selection of a solution, C is used to adjust the trade-off between the size of $\|\mathbf{w}\|$ and the tolerance for misclassification, and b is the offset from the origin of the classification solution. $\mathbf{x_i}$ is the ith example vector and y_i its label (1 or −1).

A training data set, $\mathbf{X_t}$, for the SVM is represented as

$$\mathbf{X_t} = (\mathbf{x}_1, y_1), ..., (\mathbf{x_l}, y_l) \subset \chi x \mathbb{R}, \quad (2)$$

where χ represents the space of the input such that $\chi = \mathbb{R}^d$ if there are d features.

This optimization selects a separating hyperplane that maximizes the minimum distance, or margin, from the nearest training data points to the hyperplane, subject to the set of conditions. The SVM-Light software package[24] was used for this optimization. The trained SVM model can be represented by $\mathbf{w}^*$, which is the normal vector to the separating hyperplane selected by training. A new data set, $\mathbf{x_i}$, is classified by comparing it to this separating hyperplane. This results in a margin, a_i, which is the Euclidean distance from the test example to the separating hyperplane, and is calculated as the dot product of $\mathbf{w}^*$ and the new example $\mathbf{x_i}$.

$$a_i = \langle \mathbf{w}^*, \mathbf{x_i} \rangle + b. \quad (3)$$

If $a_i > 0$, the class is positive (sphere) and if $a_i < 0$, the class is negative (cylinder). A larger margin indicates that the model ascribes greater confidence to an example vector. For the purposes of this paper, $\mathbf{a}_+$ will be used to represent margins from examples that come from the positive class (spheres), and $\mathbf{a}_-$ will be used to represent margins from examples that come from the negative class (cylinders).

Assessing the validity of a given model is very important to the success of this methodology. Two metrics are used in this paper: test accuracy and test minimum margin ratio. Test accuracy is simply the accuracy of classification of the examples in the test set. The test set is independent of the training set and the validation set used in selecting SVM model parameters. Positive margin ratio is the ratio between the largest true positive margin and the largest false positive margin, i.e., the ratio between the strongest true sphere classification and worst false sphere classification. The minimum margin ratio is the minimum of positive and negative margin ratios

$$\beta_{\min} = \min\left\{ \frac{\max(\mathbf{a}_+ > 0)}{\max(\mathbf{a}_- > 0)}, \frac{\max(\mathbf{a}_- < 0)}{\max(\mathbf{a}_+ < 0)} \right\}. \quad (4)$$

If $\beta_{\min} < 1$, the confidence in classification fails to approach 100% as the margin increases. A larger value of $\beta_{\min}$ results in greater utility of classification and a better confidence model.

B. Training and analysis

Real and simulated bistatic data were used to train, validate, and test SVM models in the training and analysis procedure, shown in Fig. 8. In the real-world experiment, a grid of amplitude data was collected around each target using the AUV *Unicorn*. Each second, the vehicle's calibrated data acquisition system began recording exactly on the second. The onboard signal processing chain, shown in Fig. 6, then extracted the amplitude for each target of interest from the
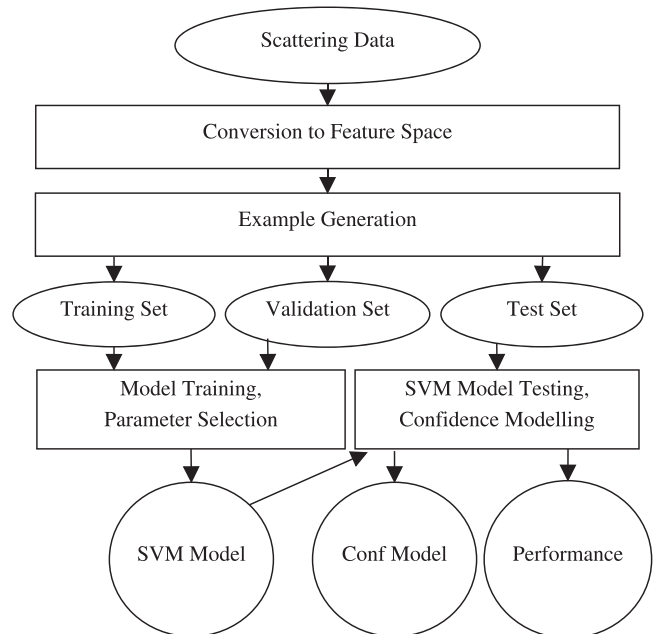

FIG. 8. Training and analysis process for machine learning methodology. Acoustic scattering amplitude data are converted to a feature space and used to construct example vectors. Independent example vectors form training, validation, and test data sets. Classification model training is conducted on the training set, and the validation set is used in the selection of model parameters. The test set is then used to determine the model's generalization performance and construct a confidence model, used to estimate the probability of correct classification given the number of samples and the classification margin.

Erin M. Fischell and Henrik Schmidt

recorded acoustic data file, subtracting the ambient noise. The result was a grid of bistatic scattering amplitudes in range, depth, and azimuth around each target.

In simulation, three-dimensional (3D) data sets were generated using the scattering simulation package for environment, source, and target parameters similar to those in the real experiment. The sphere parameters were matched very closely using an elastic fluid-filled shell model. Because the simulation package does not currently include an elastic cylinder model, the cylinder was modeled as rigid with the same dimensions and orientation as the real solid aluminum cylinder. The outputs of the simulator were files containing the azimuthal Fourier orders for the sphere and cylinder scattered fields. These data were converted into a grid of amplitude values in range, depth, and azimuth. It is this grid that was used in SVM example generation.

1. Feature selection

The presentation of the data is one of the critical aspects for successfully using machine learning for target classification. For SVMs, this takes the form of the feature representation used for example vectors. Amplitudes were mapped to features using the bistatic angle of the samples, allowing the model to exploit the differences between minima and maxima in the bistatic scattering pattern. This feature space was defined in purely spatial terms, meaning that the model does not take into account sampling order. A representation of the angular feature space is shown in Fig. 9.

Each example vector consisted of a sequence of feature-value pairs, where each value is the median scattering amplitude sampled within the angular region defined by the feature number. The feature number, F_n, was calculated as a function of the location of a sample's bistatic angle relative to the source-target line, θ_s, and a bin size in azimuth, $\Delta\theta$.

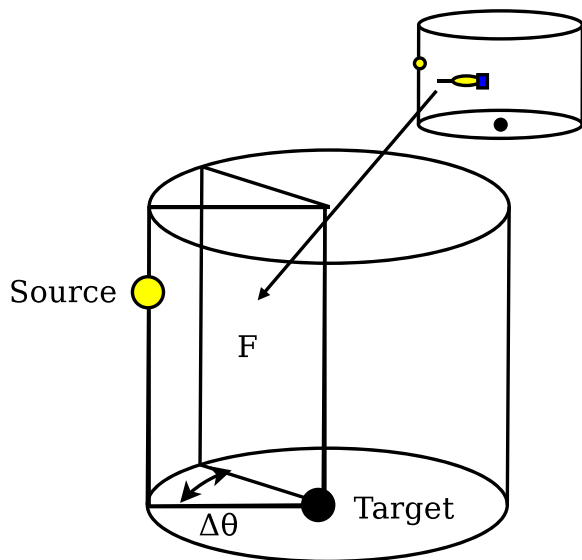$$F_n = \left\lfloor \frac{\theta_s}{\Delta\theta} \right\rfloor. \tag{5}$$

FIG. 9. (Color online) Angularly dependent feature space, configured using parameter $\Delta\theta$.

When multiple samples are collected from the same feature, the median amplitude is taken. Feature mapping is performed for each geometric point the AUV has sampled, such that the feature vector is composed of a number of feature-value pairs and the label y_i is the target's class

$$\mathbf{x_i} = \{[F_1, A_1], ..., [F_N, A_N]\}, \quad y_i = \begin{cases} 1, & \text{if sphere} \\ -1, & \text{if cylinder.} \end{cases} \tag{6}$$

This is a rapid calculation that can easily be performed on an AUV. The value of the parameter that describes the feature space, $\Delta\theta$, was selected using a design of experiment (DOE) reducing grid search.

2. Example generation

Once the grid of scattered field amplitude data is represented in terms of the feature space, training, validation, and test example vectors are constructed. Each example represents the data collected by an AUV approximately circling a target for some time τ. Because the vehicle collects one acoustic file each second, this involves collecting $N = \tau$ samples. To properly simulate this while constructing example vectors from simulation data, each angular feature is sampled either m, $m - 1$, or $m + 1$ times, where $m = N/n_F$ and n_F is the number of features in the feature space defined by $\Delta\theta$. The full set of example vectors was split into three independent data sets such that 50% of examples were used for training set $\mathbf{X_t}$, 25% for validation set $\mathbf{X_v}$, and 25% for testing set $\mathbf{X_x}$ as suggested by Hastie et al.[25] in The Elements of Statistical Learning.

The data sets collected during the real-world experiment for the sphere and cylinder targets were used to directly create example vectors. Each data set, $\mathbf{A}$, was broken into three parts: half for training ($\mathbf{A_t}$), a quarter for validation ($\mathbf{A_v}$), and a quarter for testing ($\mathbf{A_x}$). Examples were then created from each set by selecting N sequential data points at a time. If the set of amplitudes designated for training is represented by $\mathbf{A_t} = [(\theta_0, A_0), (\theta_1, A_1), ..., (\theta_M, A_M)]$ the first training example, $\mathbf{x}_1$, would be created using the data $[(\theta_0, A_0), (\theta_1, A_1), ..., (\theta_N, A_N)]$ and the second example, $\mathbf{x}_2$, would be created using the data $[(\theta_1, A_1), (\theta_2, A_2), ..., (\theta_{N+1}, A_{N+1})]$. This process is repeated until $N + i = M$. The training, validation, and test data are kept entirely independent such that if the full sphere data set consisted of 2000 data points, the first 1000 data points would be used for training, the next 500 for validation, and the final 500 for testing. This ensures that performance is tested realistically on sequential data collected by the AUV that are separate from data used in model training. The value of N was varied to observe the relationship between amount of sampling and classification accuracy.

3. Model training and analysis

An SVM model was trained using the training set. The parameters for the SVM model were selected using a logarithmic grid search in C, using the training set to train a model then classifying the validation set. Training sets were

generated with different numbers of training examples and used to classify the fully independent test set to confirm that the amount of data being used to train the model was appropriate. The performance of the final model was assessed by classifying the test set of example vectors.

4. Confidence estimation

One of the parameters important to target classification is confidence estimation: while an SVM outputs the class (sphere or cylinder) and distance to the separating hyperplane in the classification margin a, that value does not translate directly into a probability of correct classification. The confidence was determined as the probability of correct classification given the classification margin and the amount of data collected by the vehicle $P(\tilde{y}_i = y_i | a \geq a_i, N)$, where a_i is the margin. This probability was calculated empirically by using the final SVM model to classify sets of example vectors that represent different values of N, and the results were converted into a lookup table for rapid confidence calculation in real time.

C. Onboard target classification

Once the SVM model and confidence model are trained, they are used to perform real-time target classification. These processes are run within MOOS-IvP,[15] which allows nearly seamless runtime/simulation trade-off and gives a way to test signal processing on the bench with simulated or logged data. The onboard classification processing chain (Fig. 10) uses the results from the training and analysis process, and includes synchronous and asynchronous components. These processes were demonstrated in simulation and bench tests.
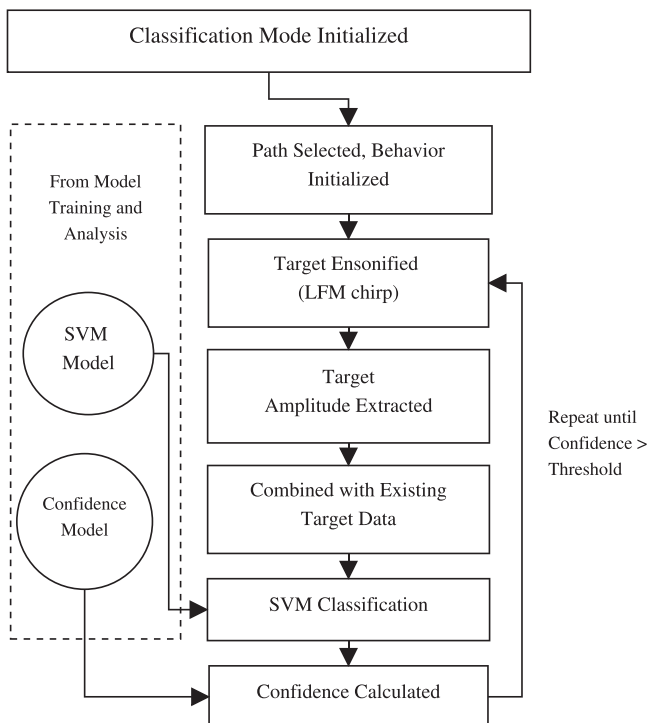
Each second, the data acquisition system records the first 0.21 s of data, which should include the direct blast from the source and target contacts that may be of interest. The target contacts are extracted from the data and tracked. Each target report includes target locations and confidence. Once a target's location has some confidence, it can be prosecuted by initiating classification. To give the best classification result, the AUV is commanded to approximately circle the target. This gives sampling across all bistatic angles.

On the real vehicle, data are recorded each second to a file which is then read by pActiveTargetProcess, which performs the matched filter, beamforming, and selection to choose the contact amplitude from the target of interest. In simulation, an acoustic simulator was developed to emulate multipath bistatic acoustic arrivals on a virtual array. This multipath is combined with simulated scattering data to simulate amplitudes collected by the AUV as it passes through the scattered field.

The scattering amplitudes collected in simulation or runtime are converted to the correct feature space. An SVM model is specified to an SVM interface application, which then runs classification on the amplitude data. The full process runs continuously as data are collected by a real or simulated AUV, constantly updating classification and confidence until a confidence threshold is met for the target.

IV. RESULTS

A. Acoustic data

During the experiment, *Unicorn* was successfully commanded to 15 target sampling missions over the course of the day, 5 for sphere sampling in the morning and 10 for cylinder sampling in the afternoon. Three depth levels were used in sampling and five radii, and the full data collection was repeated at least twice around each target. Commanding data sampling deeper than 3.5 m resulted in a depth abort because of the shallow water depth, so the commanded depths were between 1.5 m and 3.5 m. Figure 11 shows the locations of all acoustic data files collected around both
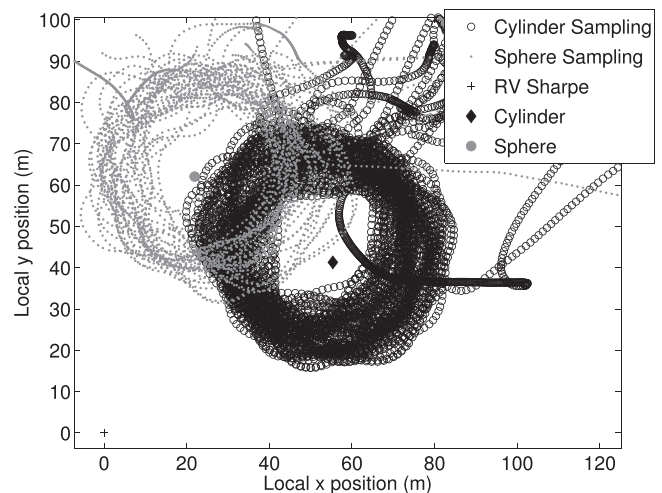

FIG. 10. Classification processing chain run onboard an AUV.

FIG. 11. Locations of collected acoustic data files in *x* and *y* relative to the position of the Research Vessel Sharpe.

Erin M. Fischell and Henrik Schmidt

targets in the original coordinate system. Each of the collected acoustic data files starts at the beginning of the second and, therefore, includes the direct blast. The sampling rate on the data acquisition boards was set to 37500 Hz, and each data file contained 8000 samples for the 16 channels. During the experiment, 15 of the 16 channels worked properly. The data from the broken element were ignored during processing.

1. Scattering amplitude extraction

The final data sets included 2162 usable scattering amplitude points around the sphere and 4784 usable points around the cylinder. This excludes data in the forward scatter direction indistinguishable from the direct blast, data made noisy by surface transport of the vehicle by the rib boat, data when the vehicle is far from the target, data when the source is off, and data where no target contact is identified by the signal processing chain. Approximately one third of the sphere data set contained high levels of noise caused by a rope wrapped in the vehicle's tail cone. These data were not used in analysis.

Each scattering amplitude represents the processing of an acoustic file consisting of 16 channels of data. For processing, *NFFT* was set to 1024, 30 beams were used, and the matched filter operation was performed with a 90% overlap for high time resolution. The band-average amplitude over the 7–9 kHz frequency band was reported. Modeling for this frequency range and environment showed low reverberation compared to ambient noise. This result was consistent with the experimental findings, which showed similar scattering strengths to those found in target-only simulation data.

Figure 12 shows the comparison of median radiation patterns calculated from real scattering field data for the sphere and cylindrical targets. These polar radiation pattern plots are valuable for comparing the overall patterns in the scattering fields of the targets. Most importantly, the sphere and cylinder are easily distinguished based on these bistatic scattering patterns, which was the purpose of collecting this bistatic data set.

The sphere data show a nearly symmetric pattern, with overall intensity lower in the backscatter direction than the forward scatter direction. Maxima are present at 130/230 degrees

and 150/210 degrees, with strong minima at 180 degrees in the direct backscatter direction. The strong minima in the backscatter direction is likely the same minima observed by Simpson *et al.*[26] between 7 and 8 kHz when measuring backscatter from this exact spherical shell in a 2003 laboratory experiment. For the thin-shelled sphere, this phenomenon has been shown to be a result of out-of-phase acoustic radiation resonance and specular reflection in the target.[27]

The cylinder shows a glint around 45 degrees, a strong lobe at 240 degrees, and stronger backscatter than the sphere. The glint and the source signal should have equal angles from broadside to the main cylinder axis. The cylinder axis is at 24 ± 5 degrees in the experiment, which puts broadside at 114 ± 5 degrees. This means that the predicted glint should be at 48 ± 5 degrees, where it is observed in the data.

2. Comparison with simulation

The densest part of the grid of scattering field amplitudes around the spherical and cylindrical targets was compared directly with scattering simulations that matched source and environment parameters as closely as possible given the limitations of the scattering simulator. The simulated scattering amplitudes at the locations that *Unicorn* sampled during the experiment around the spherical targets were used to create simulated target scattering data sets.

Figure 13 shows a comparison of scattering amplitudes in the 3 m depth bin in simulation and real data. The normalized amplitudes are similar for real and simulated data for both the sphere and cylinder. The range of scattering amplitudes is larger for the real sphere than the simulated, and the opposite is true for the cylinder.

The sphere real data set shows nearly identical locations of maxima and minima to the simulation. Important features and symmetry appear in common to both simulation and real models. The cylinder simulation is less similar to the real data, although general location of minima and maxima are consistent between the model and the real data. For example, the glint around 48 degrees is present in both real and simulated scattered fields, as are relative maxima at 120, 180, and 210 degrees. The most obvious difference
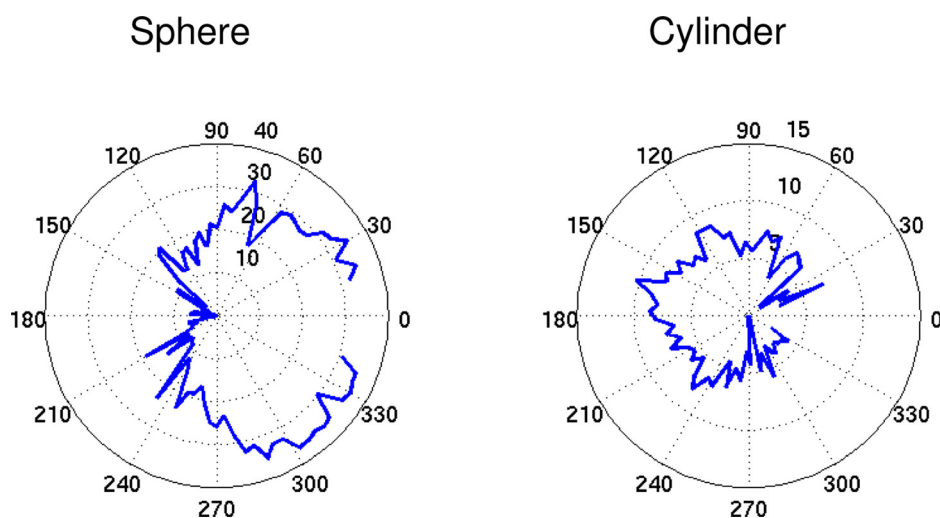
Sphere
Cylinder

FIG. 12. (Color online) Polar plot showing angle dependence of median target scattering amplitude for spherical and cylindrical targets. Difference between intensity-averaged amplitude and minimum amplitude is plotted on the r axis and angle in the source-target coordinate system on the θ axis.

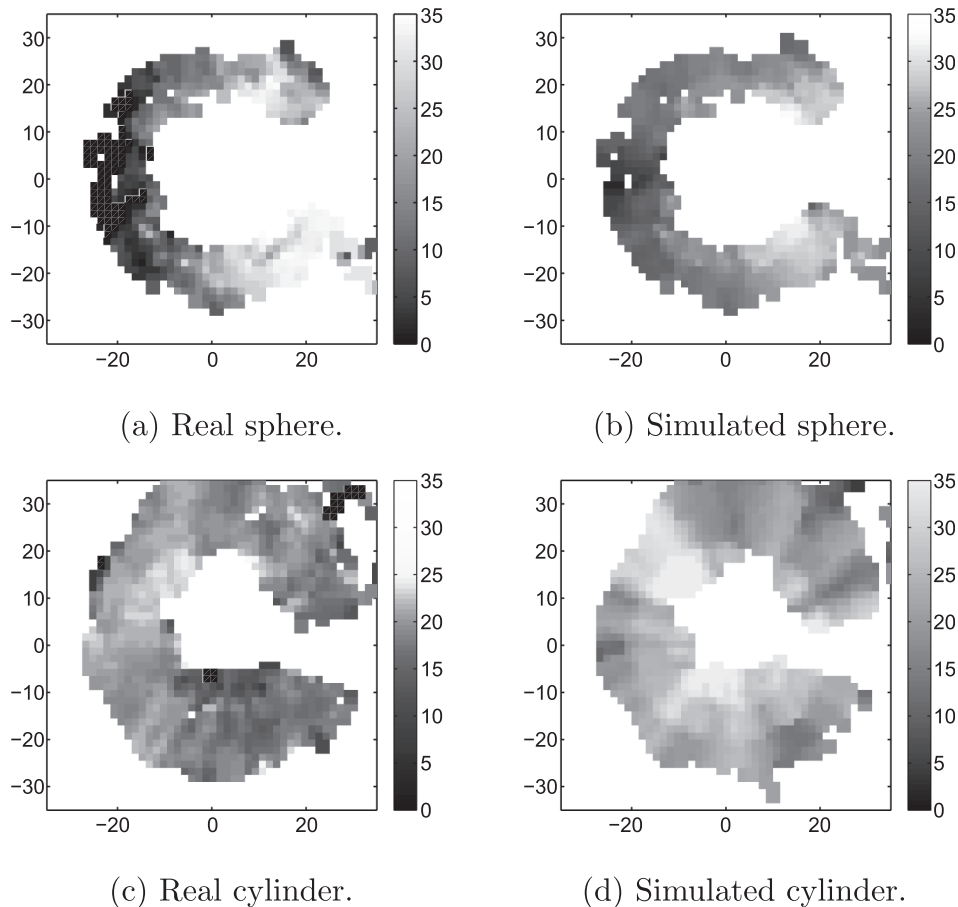(a) Real sphere.

(b) Simulated sphere.

FIG. 13. Comparison of real versus simulated scattered fields between 2.5 m and 3.5 m depth for spherical and cylindrical targets. (a) Real sphere. (b) Simulated sphere. (c) Real cylinder. (d) Simulated cylinder.

(c) Real cylinder.

(d) Simulated cylinder.

between the patterns is the greater backscatter intensity in the real cylinder's scattered field relative to the forward scatter intensity.

B. Feature and SVM parameter selection

The value of $\Delta\theta = 9°$ gave the best performance in terms of minimum margin ratio when a model was trained and validated on real data. This value also gave good performance for a model trained on simulation data and used to classify real data. The set of curves for β_{min} versus $\Delta\theta$ for some of the tested values of N is shown in Fig. 14. Larger values of N are not plotted because as N increases the value of β_{min} approaches infinity as accuracy goes to 100%. The plot clearly shows the best feature space at $\Delta\theta = 9°$.

A linear SVM model performed extremely well with the angularly dependent feature space used for classification. This minimized the complexity of the model and meant that additional parameters did not need to be selected—adjusting the value of C did not affect the model in this case.

C. Training/analysis results

Two models were used in training and analysis. The first was trained based on the real bistatic data, the second on simulation data matched to the environment of the experiment.

1. Training on real data to classify real data

Data from the real experiment were turned into training examples and then a linear SVM classification model. The test set was classified using the resulting trained model. The accuracy of the resulting classification was highly dependent on the value of N. A plot of accuracy versus N is shown in Fig. 15. Overall, the SVM model was very effective for classifying independent test example vectors once the vehicle had completed at least one full circle around the target. With
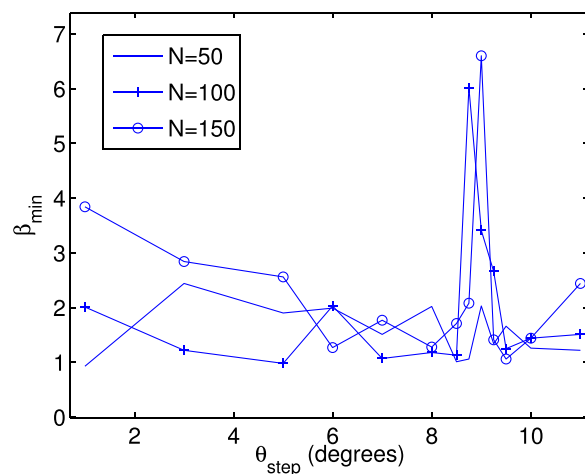

FIG. 14. (Color online) Selection of $\Delta\theta$ based on the minimum margin ratio, β_{min}, at increasing values of N. $\Delta\theta = 9°$ was selected because it converged most quickly to $\beta_{min} = \infty$ as the accuracy reached 100%.

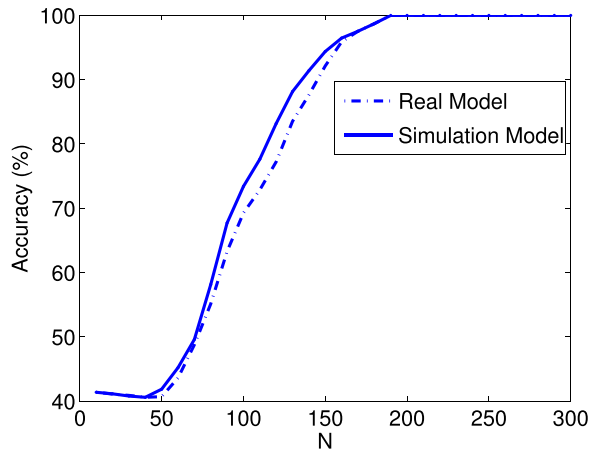Erin M. Fischell and Henrik Schmidt

FIG. 15. (Color online) N versus accuracy for model trained on real and simulated data with feature space where $\Delta\theta = 9°$. As N increases, the accuracy increases until it reaches 100%. This behavior is expected as additional data improve the averaging in each feature. After $N = 190$, the accuracy goes to 100%. When $N = 190$, the vehicle has generally completed two circles of the target.

two complete circles of the target, the accuracy of the classification model in classifying new test examples reached 100%. The performance at different values of N and estimated confidence as real scattering data are collected would be used to inform vehicle behaviors during classification.

2. Training on simulation data to classify real data

The simulated scattering fields are a good match to those calculated from the real-world BayEx'14 scattering data. To show empirically that this was the case, a SVM model was trained using the sphere and cylinder simulation scattering models and used to classify the same sets of example vectors used to test the SVM model trained on real data. The classification results for the real test examples were very similar using the simulated-data-based model and the real-data-based model. Figure 15 shows the plot of accuracy versus N in classifying the test example sets using the real and simulated SVM models. The trend for accuracy versus N is nearly identical in the two cases. The similarity of these results suggests that, at least for simple targets, a simulation approach could be used to augment real data in constructing SVM models used to classify targets in new environments. However, the power of the machine learning approach remains the flexibility to deal with targets and environments that are not easily modeled numerically or analytically by using real data to construct a model.

D. Confidence model

The curves describing the empirically determined confidence in correct classification versus classification margin a for different values of N is shown in Fig. 16. The general behavior shows an approximately logistic relationship between a and confidence. As N increases, the probability of correct classification from a lower output margin also increases. Once N increases past 190, the confidence of correct classification approaches 100% for all margin values, indicating no false classifications in the test set.
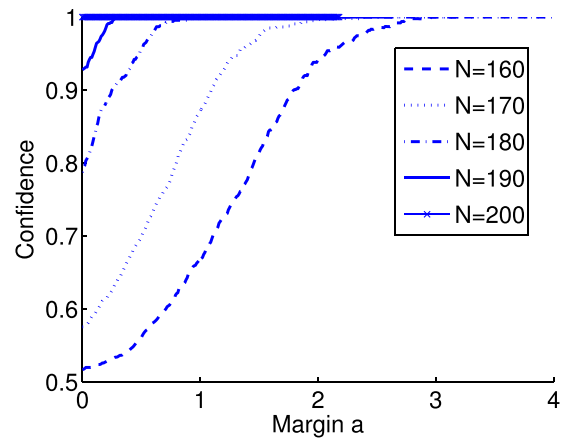

FIG. 16. (Color online) Classification confidence versus margin and N for sphere versus cylinder classification.

E. Real-time classification

The use of real and simulated models for real-time classification was tested in simulation. Simulation studies and bench tests with the vehicle computer show the full processing chain successfully completing each second: it takes ~0.3 s to calculate the target amplitude from an acoustic file, the incorporation of acoustic data into the existing SVM example for classification takes <0.05 s, and the actual classification, which is only run when the vehicle exits a feature (every 5–10 s depending on range to the target), takes <0.5 s. This shows the plausibility of using this method for real-time analysis and classification. These numbers were shown on the bench with the *Unicorn* computer when only the classification processing chain was running. Adding the target localization processing chain increased processing times significantly so that the acoustic data were fully processed only once every 3 s instead of every second. To simultaneously run localization and classification processing on *Unicorn* every second, a more powerful or second computer would be required.

V. CONCLUSIONS

The experiment demonstrated the navigation, timing, and vehicle behaviors necessary for high-quality 3D bistatic scattering data collection by an AUV. Navigation problems were addressed by installing a new IMU that improved drift to <0.5% of distance traveled, and by surfacing frequently for GPS. The time synchronization issues were solved using a CSAC as a time source, a PPS-triggered data acquisition system, and characterizing all delays to achieve better than 70 μs accuracy. AUV sampling behaviors were developed to keep the vehicle broadside to the target, resulting in more uniform data quality through the sampling region. The final system was successfully deployed, and in one day of data collection acquired sufficient sphere and cylinder scattering data to compare real data to existing bistatic scattering models.

Classification of spherical versus cylindrical targets using scattering amplitude data collected by an AUV was successfully demonstrated using real and simulated target

J. Acoust. Soc. Am. **138** (6), December 2015

Erin M. Fischell and Henrik Schmidt   3783

scattering data. Furthermore, it was shown in simulation on the bench that all processes required for target classification using this methodology can be run in less than a second, which means AUV-based real-time classification and confidence estimation are plausible. While the sphere versus cylinder classification investigated here is a simplification of the target geometries of interest in mine countermeasures, this research shows the potential of the combination of sensing bistatic scattering fields with a linear array payload and applying machine learning classification of calculated acoustic amplitudes for solving the real-time target classification problem.

Several avenues of future study are suggested based on these results. First, data were acquired on only a single cylinder rotation during the experiment, and collecting scattering data from different cylinder rotations would be valuable for learning more about aspect-dependent bistatic scattering. Data from target types other than spheres and cylinders and on buried targets would be required to extend this work from the basic sphere versus cylinder classification to a broader mine countermeasures application. Looking at the influence of the cylinder aspect ratio and absolute target size on the efficacy of the methodology should be further explored, as should the effect of the target fill and environmental factors such as bottom type. Another exciting possibility is extending the work to incorporate multiple vehicles and multistatics with an acoustic source located on one of the AUVs.

ACKNOWLEDGMENTS

[1]R. E. Hansen, H. J. Callow, T. O. Sabo, and S. A. V. Synnes, "Challenges in seafloor imaging and mapping with synthetic aperture sonar," IEEE Trans. Geosci. Remote Sens. **49**, 3677–3687 (2011).

[2]K. Baik, C. Dudley, and P. L. Marston, "Acoustic quasi-holographic images of scattering by vertical cylinders from one-dimensional bistatic scans," J. Acoust. Soc. Am. **130**, 3838–3851 (2011).

[3]L. Haumesser, A. Baillard, D. Decultot, and G. Maze, "Experimental identification of finite cylindrical shell vibration modes," J. Acoust. Soc. Am. **111**, 2034–2039 (2002).

[4]Z. J. Waters1, H. J. Simpson, A. Sarkissian, S. Dey, B. H. Houston, J. A. Bucaro, and T. J. Yoder, "Bistatic, above-critical angle scattering

[5]measurements of fully buried unexploded ordnance (UXO) and clutter," J. Acoust. Soc. Am. **132**, 3076–3085 (2012).

[5]S. G. Kargl, K. L. Williams, A. L. España, J. L. Kennedy, T. T. Marston, J. L. Lopes, and R. Lim, "Acoustic scattering from underwater munitions near a water-sediment interface," J. Acoust. Soc. Am. **129**, 2685–2697 (2011).

[6]P. Blondel and N. G. Pace, "Bistatic sonars: Sea trials, laboratory experiments and future surveys," Arch. Acoust. **34**, 95–109 (2009).

[7]P. Moren, A. Caiti, M. Zakharia, M. A. Larsen, Ph. Blondel, and J. Dybedal, "Acoustic sea trial in the Möja Söderfjärd dumpsite," in *Buried Waste in the Seabed*, edited by Ph. Blondel and A. Caiti (Springer, Chichester, 2006), pp. 87–102.

[8]K. D. LePage and H. Schmidt, "Bistatic synthetic aperture imaging of proud and buried targets from an AUV," IEEE J. Ocean. Eng. **27**(3), 471–483 (2002).

[9]J. R. Edwards, H. Schmidt, and K. D. LePage, "Bistatic synthetic aperture target detection and imaging with an AUV," IEEE J. Ocean. Eng. **26**(4), 690–699 (2001).

[10]H. Schmidt and J. Lee, "Physics of 3-D scattering from rippled seabeds and buried targets in shallow water," J. Acoust. Soc. Am. **105**, 1605–1617 (1999).

[11]M. Zampolli, F. B. Jensen, and A. Tesei, "Review of target strength information for waterside security application," in *Proc. of WSS2008*, Technical University of Denmark, Aug. 25–28, 2008.

[12]E. J. Kaminsky and M. Barbu, "Classification of cylindrical targets buried in seafloor sediments," in *Proc. of IEEE Region 5 Technical Conference*, Apr. 20–21, 2007, pp. 117–123.

[13]A. Malarkodi, D. Manamalli, G. Kavitha, and G. Latha, "Acoustic scattering of underwater targets," in *Proc. of SYMPOL* (2013), pp. 127–132.

[14]P. Runkle, L. Carin, L. Couchman, J. A. Bucaro, and T. J. Yoder, "Multiaspect identification of submerged elastic targets via wave-based matching pursuits and hidden Markov models," J. Acoust. Soc. Am. **106**, 605–616 (1999).

[15]M. Benjamin, H. Schmidt, P. Newman, and J. Leonard, "Nested autonomy for unmanned marine vehicles with MOOS-IvP," J. Field Rob. **27**, 834–875 (2010).

[16]T. Schneider and H. Schmidt, "Goby-Acomms version 2: Extensible marshalling, queuing, and link layer interfacing for acoustic telemetry," in *9th IFAC Conference on Manoeuvring and Control of Marine Craft'12*, Arenzano, Italy (2012).

[17]S. Petillo and H. Schmidt, "Exploiting adaptive and collaborative AUV autonomy for detection and characterization of internal waves," IEEE J. Ocean. Eng. **39**, 150–164 (2014).

[18]Symmetricom, "Leading edge technology enables a chip scale atomic clock" (2012) available at http://www.ecnmag.com/sites/ecnmag.com/files/legacyimages/WP_LeadingEdgeTechnology_CSAC.pdf (Last viewed March 20, 2014).

[19]General Standards Corporation, "24DSI12, 12-channel 24-bit delta-sigma PMC analog input board" available at http://www.generalstandards.com/specs/24dsi12_spec.pdf (Last viewed April 8, 2014).

[20]Advantech, "PCM-3363," Jan. 2014, available at http://downloadt.advantech.com/ProductFile/PIS/PCM-3363/Product%20-%20Datasheet/PCM-3363_DS%2801.15.14%2920140205144952.pdf (Last viewed April 8, 2014).

[21]"Generic NMEA GPS receiver," available at https://www.eecis.udel.edu/~mills/ntp/html/drivers/driver20.html (Last viewed April 8, 2014).

[22]"LinuxPPS," available at http://linuxpps.org/wiki/index.php/Main_Page (Last viewed April 8, 2014).

[23] Vladimir and N. Vapnik, *The Nature of Statistical Learning Theory* (Springer, Berlin, 1995).

[24]T. Joachims, *Making Large-Scale SVM Learning Practical. Advances in Kernel Methods—Support Vector Learning*, edited by B. Schölkopf, C. Burges, and A. Smola (MIT Press, Cambridge, MA, 1999).

[25]T. Hastie, R. Tibshirani, and J. Friedman, "Model assessment and selection," in *The Elements of Statistical Learning—Data Mining, Inference and Prediction*, 2nd ed. (Springer, Berlin, 2009), Sec. 7.1, p. 222.

[26]H. J. Simpson, B. H. Houston, and R. Lin, "Laboratory measurements of sound scattering from a buried sphere above and below the critical angle," J. Acoust. Soc. Am. **113**, 39–42 (2003).

[27]P. L. Marston and N. H. Sun, "Resonance and interference scattering near the coincidence frequency of a thin spherical shell: An approximate ray synthesis," J. Acoust. Soc. Am. **92**, 3315–3319 (1992).