Detection and Classification Capabilities of Two Multibeam Sonars

Running head: CAPABILITIES OF TWO MULTIBEAM SONARS

Emma Cotter (ecotter@whoi.edu) [1, 2]

Brian Polagye[3]

Keywords: Multibeam sonar, marine renewable energy, environmental monitoring, target tracking, target classification, machine learning

[1]Woods Hole Oceanographic Institution, Department of Applied Ocean Physics and Engineering, Falmouth, MA, USA

[2]The author was affiliated with the University of Washington Department of Mechanical Engineering at the time that the work described in this manuscript was conducted.

[3]University of Washington Department of Mechanical Engineering, Seattle, WA, USA

Abstract

Multibeam sonars can be used to detect and classify marine fauna under conditions when optical sensors are ineffective. In this work, we compare the detection and classification capabilities of two multibeam sonars with different operating frequencies, the Tritech Gemini 720is (720 kHz) and the Teledyne BlueView M900-2250 (2250 kHz). The two sonars were deployed with overlapping swaths in a narrow tidal channel with peak currents of approximately 2 m/s where seals, schools of fish, and diving birds were intermittently present. In comparing data concurrently acquired by both sonars, we observe differences in the appearance of detected targets and detection capabilities. The detected target tracks are classified using a random forest model as either individual biological targets (flora and fauna, including fish, diving birds, seals, and plant matter), fish schools, or non-biological targets (including entrained air and sonar artifacts). Despite the observed differences in detection capabilities for the two sonars, automatic classification distinguishes between target classes with precision and recall above 0.7 and 0.9, respectively. These results suggest that while similar methodologies can be used for data analysis, some outcomes from environmental studies using multibeam sonars may be instrument-specific.

Introduction

Multibeam sonars were first applied to bathymetric mapping and hydrographic surveys because their relatively large fields-of-view allowed for more accurate and rapid data collection than single-beam sonars [Kostylev et al., 2001, Colbo et al., 2014]. In recent years, increases in the angular resolution, range, and data acquisition rates of multibeam sonars have fostered new applications [Colbo et al., 2014], including detection of fauna [Melvin, 2016, Parsons et al., 2017], inspection of bubble seeps [Urban et al., 2017], and classification of benthic habitats [Lundblad et al., 2006, Li et al., 2017, Lacharité et al.,

2018].

Due to their operational range (tens of meters) and field-of-view (120+ degrees), multibeam sonars are appealing for observations of fauna at marine energy sites [Copping et al., 2016], and a wide range of sonars have been used for this application (see Table 1). However, data collection at marine energy sites poses several challenges. First, these are characteristically "high-energy" environments, and entrained air from waves or tidal currents or advected debris can create high-intensity returns in the field-of-view, creating challenges for target classification [Melvin and Cochrane, 2015]. Second, interactions between fauna and marine energy converters are likely to be rare, meaning that detection requires continuous monitoring over relatively long time windows (i.e., months) [Copping et al., 2016]. However, data collection from multibeam sonars can quickly accrue large volumes of data that are challenging to store and process. As an example, the BlueView M900-2250 sonar produces over 6 GB of raw data per hour using the manufacturer-provided software, or 4 TB over a one-month period. This motivates the use of automatic target detection algorithms to restrict human data review to periods of interest [Cotter et al., 2017, Lieber et al., 2017, Williamson et al., 2017, Hastie et al., 2019]. Our recent work, Cotter and Polagye [2020a], used machine learning to automatically separate targets associated with flora and fauna from those resulting from sonar artifacts and entrained air in data collected using a BlueView M900-2250 multibeam sonar. In Hastie et al. [2019], a similar approach was used to automatically distinguish between seal and non-seal targets in data collected using a Tritech Gemini 720id sonar, and Hastie et al. note that the classification approach would benefit from a comparative evaluation between sonars.

Previous developments in multibeam sonar target detection and classification have only tested algorithms on one sensor, and we are not aware of any side-by-side comparison of the capabilities of multiple multibeam sonars to observe marine fauna. Here, we compare the detection and classification capabilities of two concurrently deployed multibeam sonars

with overlapping fields-of-view. Our objective is to explore the potential consequences of multibeam sonar selection on environmental monitoring data for marine energy projects and to investigate whether target classification methods are transferable between sonars. To this end, we first compare the appearance of targets (e.g., marine mammals, diving birds, fish, and entrained air) concurrently detected in the two sonar data streams. Following this, we compare the performance of automatic classification of targets detected by the two sonars. Finally, we discuss the implications of these findings for the use of multibeam sonars for environmental monitoring at marine energy sites.

Materials and procedures

Sensors and data collection

The two multibeam sonars used in this study (Teledyne BlueView M900-2250 and Gemini 720is, hereafter referred to as BlueView and Gemini, respectively) are described in Table 2. The sonars were deployed using an Adaptable Monitoring Package (AMP), a cabled sensor platform that facilitates deployment of active acoustic, passive acoustic, and optical sensors [Polagye et al., 2020]. The Gemini was mounted with a 10-degree upward tilt to compensate for an internal 10-degree downward tilt of the transducer, such that the sonar swaths overlapped to the greatest extent possible. During this deployment, the AMP was also instrumented with a Nortek Signature 500 acoustic Doppler current profiler (ADCP) to provide contextual information about tidal currents and Allied Vision Manta G507b optical cameras for opportunistic target identification. The sonars and ADCP were synchronized to minimize active acoustic crosstalk. We deployed the AMP in the tidal channel at the entrance to Sequim Bay, WA, USA at a location with a water depth of approximately 7 m and peak tidal currents of approximately 2 m/s. Both sonars were approximately 1.1 m

above the seafloor, and the angle of the sonar swaths with the seafloor was varied intermittently during deployment with a tilt motor (15-20° above horizontal). Harbor seals (*Phoca vitulina*), schools of fish (including Pacific herring, *Clupea pallasii*), and diving birds (including pigeon guillemots, *Cepphus columba*) were observed at the test site. BlueView data were stored as 814 x 444 8-bit Cartesian images in PNG format. The BlueView can render larger images, but image size was limited by CPU resources on the acquisition computer. Gemini data were stored as 512 x 1229 8-bit polar images in PNG format, where each row of the image represented a beam and each column represented a range bin.

We consider two data sets in this paper: the first consists of 8 hours of continuously acquired data collected between 11 AM and 3 PM on 29 March and 3 April, 2019 and the second consists of 265 one-minute segments of sonar data collected on 21-25 March and 16-19 April 2019 when a target of interest was detected in the BlueView data using the real-time methods described in Cotter and Polagye [2020a]. The first data set provides an unbiased comparison of sonar imagery, while the second contains a sufficient number of targets to train a machine learning classification model.

Target detection and tracking

Except where otherwise mentioned, target detection and tracking methods follow those outlined in Cotter and Polagye [2020a], and all processing was implemented in MATLAB R2019b. The process is summarized here, and was the same for both sonars excepting sonar-specific thresholds. For target detection, first, the foreground of each image was identified by subtracting the median of the previous 20 seconds of data (background). While the background could be updated every ping, we limited the update rate to every 10 seconds to reduce the computational burden for real-time processing. Second, a sonar-specific, empirically-selected intensity threshold was applied to the image foreground.

For the BlueView data, a range-varying threshold with the highest intensity closest to the transducer was used to offset higher observed backscatter intensities in this region, while a static threshold was used for the Gemini data. Third, a two-dimensional median filter was applied to each thresholded image and only targets appearing larger than 0.015 and 0.02 m^2, for the continuously acquired data and the one-minute sequences, respectively, were retained. More permissive thresholds were used for the continuously acquired data than for the one-minute sequences to increase the number of candidate targets available for comparison. In BlueView images, targets in the outermost 10 degrees were rejected for direct comparison with the Gemini swath, and targets in the first 0.75 meters of range were rejected due to frequent high-intensity noise and poor representation of targets in this region. Gemini data were transformed to Cartesian coordinates before target detection and targets were rejected in a 2.2 m^2 region in the center of the image (3% of the total image) due to persistent acoustic or electrical crosstalk from another AMP sensor. Finally, a morphological dilation operation was used to aggregate targets that might be part of the same object (e.g., fish schools), a method adapted from Williamson et al. [2017]. This is the only portion of the target detection and tracking that was substantially modified from Cotter and Polagye [2020a]. In the BlueView data, targets within 1.5 m Cartesian distance were aggregated, and in the Gemini data, targets within a range of 1.5 m and an azimuthal distance of two degrees were aggregated.

A Kalman filter [Maybeck, 1979] was then used to track targets through the sonar swath, with the tidal current at the time of detection taken as an initial velocity assumption. If tidal current data was not available within 15 minutes of the target detection (e.g., interruption in ADCP data stream), an initial velocity of zero was assumed, since the Kalman filter performance was relatively insensitive to the choice of initial velocity. If any targets were detected within a specified distance of the predicted location in the subsequent frame, the closest target to the predicted location was associated with the track. This association distance was 0.75 m for the continuously acquired data and 0.5 m for the

one-minute sequences. The Kalman filter was then updated with the new target position and velocity. When no target was associated with a track for one second, the track was considered complete and, if consisting of more than five target detections, considered valid.

Target track annotation

All valid target tracks were annotated as either seals, diving birds, fish schools, "small targets," or non-biological targets by a human reviewer. We used concurrent data from the AMP optical cameras to opportunistically verify classification when the target fell within their field of view and ambient lighting and water clarity allowed. However, such data was infrequently available, which is one of the primary motivations for use of multibeam sonars in this application. The "small targets" class included individual fish and drifting pieces of plant matter, because, absent concurrent optical data, a reviewer cannot confidently differentiate these targets. The other target classes have distinct attributes (e.g., appearance, patterns of motion) that allow reliable classification in the absence of optical data. Non-biological targets included entrained air bubbles and sonar artifacts.

All target tracks from both sonars in the continuously acquired data were manually annotated by the reviewer. In the one-minute sequences, only BlueView target tracks were initially annotated and tracks for which the reviewer was not able to confidently assign a target class were removed from the data set (seven total tracks). These tracks were compared to the set of Gemini target tracks and concurrent targets were assigned the same class. The non-concurrent Gemini target tracks were then manually annotated. As discussed in the next section, some Gemini target tracks were retained in the continuously acquired data despite an ambiguous appearance. The consequences of including this ambiguous training data in a classification model provides an instructive contrast in classification model accuracy and the relative utility of the two sonars.

Comparison of classification capabilities

We separated the annotated target tracks into three classes to compare automatic classification capabilities: individual biological targets, fish schools, and non-biological targets. The individual biological target class included small targets, seals, and diving birds. While diving birds and seals produce distinct acoustic imagery, there were insufficient training data available to train a machine learning classification model for these subclasses (e.g., 5 seal tracks, 16 diving bird tracks). While species or genus-level classification is ideal, coarse classification is still valuable since this can automatically reject the large number of non-biological targets that were detected on both sonars, thereby reducing the required human review effort [Cotter and Polagye, 2020a].

The training data set for automatic classification was limited to the target tracks from the one-minute sequences to isolate differences in classification capability from differences in detection capability. The total number of target tracks associated with each class for each sonar are listed in Table 3. As discussed in the section *Comparison of detection capabilities*, frequent vessel traffic during collection of the continuously acquired data resulted in a relatively large number of non-biological targets in the Gemini data (entrained air from boat wakes) that were not present in the BlueView data. These non-biological targets were difficult to differentiate from small targets (individual target class) drifting near the surface. To explore the effects that ambiguous training data can have on classification model performance, we also evaluated the performance of a Gemini target track classification model trained using the target tracks from both the continuous and one-minute data sets.

We used the same methods as in Cotter and Polagye [2020a] to train a classification model for each sonar. Twenty-nine hand-engineered features (Cotter and Polagye [2020a], Appendix 1) described the shape, motion, and intensity of each target track, as well as environmental covariates at the time of detection (e.g., current speed). Each feature was

normalized such that the 10^{th} and 90^{th} percentiles were equal to zero and one, respectively, and a hill-climbing feature selection algorithm [Dobeck et al., 1997] was used to identify the subset of the 29 features that provided optimal classification performance for each sonar. A random forest algorithm (100 trees, MATLAB default values for other parameters) was used for classification, and the holdout method was used for model cross-validation: random selection of 70% of the available data to train the classification model and the remaining 30% used to validate performance. To address disparity in the number of samples from each class (e.g., 985 non-biological targets compared to 115 fish schools in the Gemini data), target tracks were randomly subsampled [Wallace et al., 2011] such that there was equal representation from each class in both the training and validation data subsets. The predicted classes of the validation data subset were used to evaluate model performance using two metrics:

- the recall, or true positive rate, for each class, defined as the fraction of target tracks belonging to that class that were correctly classified; and
- the precision for each class, defined as the fraction of target tracks predicted to belong to the class that were correctly classified.

This cross-validation process was repeated 100 times to address any variability from the randomized subsampling.

Assessment

Comparison of detection capabilities

Figure 1 shows concurrent imagery of each biological target class on both sonars. In general, concurrently detected targets have similar appearances, though the higher

resolution of the BlueView results in sharper, more defined edges, and, in the case of the seal, the shape of the animal is more clear.

Other than the appearance of targets, two key differences were noted between the detection capabilities of the two sonars. First, air bubbles from surface waves, boat wakes, and diving fauna produced higher intensity returns in the Gemini data. An example of entrained air from a boat wake that was detected in the Gemini data but not in the BlueView data is shown in Figure 2A, and a similar example of entrained air from a diving bird is shown in Figure 2B. These differences may be explained by scattering and absorption of sound from bubbles. Specifically, the lower operating frequency of the Gemini may be closer to the acoustic resonance peaks of small, entrained bubbles [Vagle and Farmer, 1991]. Similarly, the higher frequency BlueView signal attenuates more rapidly than that of the Gemini and the excess attenuation due to bubbles may result in shadow zones beyond such clouds where targets are not detectable. Excess attenuation due to bubbles [Medwin and Clay, 1998] is a factor for both systems, but may simply be a more important factor for the BlueView. This is consistent with a lack of detected targets in the BlueView data within regions that contained bubbles on the Gemini. Figure 2C shows a small target, possibly kelp based on its passive trajectory, detected at $\sim$9 m range at a time when bubbles were observed in the Gemini data and not visible in the BlueView data. Finally, as shown in Figure 2D, the entrained air from a boat wake at closer range to the AMP has a starkly difference appearance in each sonar. Specifically, there is a high intensity return close to the transducer in the BlueView image and more diffuse scattering farther from the transducer in the Gemini image. Unlike the prior example of air bubble masking, we believe this is a consequence of thresholds used in the sonars' internal processing.

Second, radial segments in the BlueView image masked the detection of some targets, as shown in Figure 2E. These segments are produced by overlapping sidelobes of the three individual transducers that make up the BlueView (personal communication, Tyler

Whitaker, Teledyne BlueView), and obscure approximately 11% of the BlueView image (13 m^2). The Kalman filter target tracking scheme frequently did not effectively track targets moving across the sonar swath when they passed through in these regions, resulting in multiple, shorter tracks for a single target. These regions masked two of the nine diving birds that were observed in the continuously acquired Gemini data.

Comparison of classification capabilities

Figure 3 shows the precision and recall for automatic classification of each target class for both sonars. In general, performance is comparable between the two sensors, though some minor differences are observed. Recall rates were above 0.9 and similar for both sonars (BlueView: 0.91, 0.94, and 0.93; Gemini: 0.91, 0.96, and 0.90 for non-biological targets, fish schools, and individual targets, respectively), with little variability observed due to subsampling (interquartile ranges below 0.06 in all cases). Individual targets were classified with higher precision in the BlueView data (0.91, compared to 0.77), likely due to their sharper edges and better-defined shapes. Precision was similar between the two sonars for fish schools and non-biological targets, though for these classes, the BlueView had larger interquartile ranges than were observed for the Gemini, indicating that there was more track-to-track variation in that sonar's data.

When all available target tracks (one-minute sequences and continuously acquired data) were used to train the Gemini classification model, recall decreased for individual targets and non-biological targets. Similarly, precision for individual targets decreased by over 0.2. This is likely due to the difficulty in differentiating between drifting targets near the surface (small targets) and entrained air (non-biological targets) by the human reviewer. Classification model performance for fish schools was unchanged because there was no ambiguity in their annotation by the human reviewers.

Discussion

Both sonars that we evaluated had strengths and weaknesses, and, in some situations, it may be advantageous to concurrently deploy multibeam sonars with contrasting capabilities. For example, human annotation of target tracks, particularly for smaller targets, is easier using the higher-frequency BlueView data, but more targets are observable in the Gemini data because of less signal attenuation from entrained air and a lack of sidelobe interference.

Our findings have several implications for environmental monitoring at marine energy sites that may be more broadly applicable. First, air bubbles created high-intensity targets in the lower-frequency Gemini data but were infrequently visible in the higher-frequency BlueView data. However, this did not mean that target detection was superior in the BlueView data since it appeared that bubbles effectively shadowed some biological targets at greater range. What role excess attenuation from bubbles and instrument internal thresholds play in these observations remains unclear and cannot be fully addressed with the available information. However, this should be considered when interpreting data collected from shallow or turbulent environments where entrained air is common. We note that boat wakes were a significant source of entrained air at our test site, but that this would not be a factor when data are collected at greater depth or at sites with less vessel traffic. This could, however, be a factor for vessel-based data collection in any water depth. Second, we found that the target tracking and classification approach we introduced in Cotter and Polagye [2020a] is translatable between sonars, suggesting that it could be readily adapted to other multibeam sonars (Table 1). Third, ease of human classification is not necessarily analogous to accuracy of machine learning classification. Specifically, a human reviewer could unambiguously assign a target class to more detected tracks in the BlueView data, but, once trained, the random forest algorithm was able to distinguish between fish schools, individual biological targets, and non-biological targets with similar

performance for both sonars. The difference in ease of human classification has implications for the accuracy of training data, which is foundational to the classification model. This is demonstrated in the decrease in Gemini classification model performance when all target tracks, including those with ambiguous class annotations, were included in the training data set.

Comments and recommendations

We have qualitatively and quantitatively compared the target detection and classification capabilities of two multibeam sonars. We found that differences in operating frequency, transducer design, and internal processing result in different detection capabilities. Most notably, entrained air created high-intensity targets in the lower-frequency Tritech sonar, but may have shadowed targets in the higher-frequency BlueView sonar. While human classification of targets was more straightforward in the higher-frequency sonar data because of its higher resolution, we found that a random forest algorithm was able to classify targets in both sonar data streams with similar results. Differences in detection capabilities should be considered when selecting multibeam sonars for environmental monitoring, particularly at sites where entrained air is prevalent.

Acknowledgments

The authors wish to acknowledge Zachary Tully for his assistance with data annotation, Christopher Bassett for productive conversations regarding acoustic scattering, Paul Murphy for his work developing data acquisition software, and James Joslin and Paul Gibbs for their hard work that made deployment of the Adaptable Monitoring Package possible. Figures showing sonar data use the “solar” colormap from the cmocean package

[Thyng et al., 2016]. Portions of this research were funded by the US Department of Energy under DE-EE0007827. Emma Cotter was supported by a National Science Foundation Graduate Research Fellowship under grant number DGE-1762114.

References

K. Colbo, T. Ross, C. Brown, and T. Weber. A review of oceanographic applications of water column data from multibeam echosounders. *Estuatrine, Coastal, and Shelf Science*, 145:41–56, 2014. doi: 10.1016/j.ecss.2014.04.002.

A. Copping, N. Sather, L. Hanna, J. Whiting, G. Zydlewski, G. Staines, A. Gill, I. Hutchinson, and Others. State of the Science Report: Environmental Effects of Marine Renewable Energy Development Around the World. Technical report, Report for Ocean Energy Systems (OES), 2016.

E. Cotter and B. Polagye. Automatic Classification of Biological Targets in a Tidal Channel using a Multibeam Sonar. *Journal of Atmospheric and Oceanic Technology*, 2020a. doi: 10.1175/JTECH-D-19-0222.1.

E. Cotter and B. Polagye. Supplemental data for detection and classification capabilities of two multibeam sonars. *University of Washington ResearchWorks Archive*, 2020b. URL `http://hdl.handle.net/1773/46231`.

E. Cotter, P. Murphy, and B. Polagye. Benchmarking sensor fusion capabilities of an integrated instrumentation package. *International Journal of Marine Energy*, 2017. doi: 10.1016/j.ijome.2017.09.003.

G. J. Dobeck, J. C. Hyland, and L. Smedley. Automated Detection/Classification of Sea Mines in Sonar Imagery. In *Proceedings of SPIE*, volume 3079, pages 90–110, Orlando, FL, 1997.

F. Francisco and J. Sundberg. Detection of Visual Signatures of Marine Mammals and Fish within Marine Renewable Energy Farms using Multibeam Imaging Sonar. *Journal of Marine Science and Engineering*, 7(22):1–19, 2019. doi: 10.3390/jmse7020022.

G. D. Hastie, G. Wu, S. Moss, P. Jepp, J. MacAulay, A. Lee, C. E. Sparling, C. Evers, and Others. Automated detection and tracking of marine mammals: a novel sonar tool for monitoring effects of marine industry. *Aquatic Conservation: Marine and Freshwater Ecosystems*, 29(S1):119–130, 2019. doi: 10.1002/aqc.3103.

V. E. Kostylev, B. J. Todd, G. B. J. Fader, R. C. Courtney, G. D. M. Cameron, and R. A. Pickrill. Benthic habitat mapping on the Scotian Shelf based on multibeam bathymetry, surficial geology and sea floor photographs. *Marine Ecology Progress Series*, 219: 121–137, 2001. doi: 10.3354/meps219121.

M. Lacharité, C. J. Brown, and V. Gazzola. Multisource multibeam backscatter data : developing a strategy for the production of benthic habitat maps using semi-automated seafloor classification methods. *Marine Geophysical Research*, pages 307–322, 2018. doi: 10.1007/s11001-017-9331-6.

D. Li, C. Tang, C. Xia, and H. Zhang. Acoustic mapping and classification of benthic habitat using unsupervised learning in artificial reef water. *Estuarine, Coastal and Shelf Science*, 128:11–21, 2017. doi: 10.1016/j.ecss.2016.12.001.

L. Lieber, B. Williamson, C. S. Jones, L. R. Noble, A. Brierley, P. Miller, and B. Scott. Introducing novel uses of multibeam sonar to study basking sharks in the light of marine rnewable energy extraction. In *Proceedings of the 2nd International Conference on Environmental Interactions of Marine Renewable Energy Technologies*, page 3 pp., Stornoway, Isle of Lewis, Outer Hebrides, Scotland, 2014.

L. Lieber, T. Nilsen, C. Zambrano, and L. Kregting. Optimising multiple multibeam sonars

to assess marine life interactions with an underwater kite. In *Proceedings of the 12th European Wave and Tidal Energy Conference*, pages 1–6, Cork, Ireland, 2017.

E. R. Lundblad, D. J. Wright, J. Miller, E. M. Larkin, R. Rinehart, D. F. Naar, B. T. Donahue, S. M. Anderson, and Others. A benthic terrain classification scheme for American Samoa. *Marine Geodesy*, 29(2):89–111, 2006. doi: 10.1080/01490410600738021.

P. Maybeck. Introduction. In *Stocastic models, estimation, and control*, pages 1–19. Academic Press, 1979.

H. Medwin and C. Clay. *Fundamentals of Acoustical Oceanography*. Academic Press, 1998. ISBN 012487570X.

G. D. Melvin. Observations of in situ Atlantic bluefin tuna with 500-kHz multibeam sonar. *ICES Journal of Marine Science*, 73(8):1975–1986, 2016. doi: 10.1093/icesjms/fst048.

G. D. Melvin and N. A. Cochrane. Multibeam Acoustic Detection of Fish and Water Column Targets at High-Flow Sites. *Estuaries and Coasts*, 38(1):227–240, 2015. doi: 10.1007/s12237-014-9828-z.

M. J.G. Parsons, E. Fenny, K. Lucke, S. Osterrieder, G. Jenkins, B. J. Saunders, P. Jepp, and I. M. Parnum. Imaging Marine Fauna with a Tritech Gemini 720i Sonar. *Acoustics Australia*, 45(1):41–49, 2017. doi: 10.1007/s40857-016-0076-1.

Brian Polagye, James Joslin, Paul Murphy, Emma Cotter, Mitchell Scott, Paul Gibbs, Christopher Bassett, and Andrew Stewart. Adaptable monitoring package development and deployment: Lessons learned for integrated instrumentation at marine energy sites. *Journal of Marine Science and Engineering*, 8(8):553, 2020. doi: 10.3390/jmse8080553.

K. M. Thyng, C. A. Greene, R. D. Hetland, H. M. Zimmerle, and S. F. DiMarco. True colors of oceanography: Guidelines for effective and accurate colormap selection. *Oceanography*, 29(3):9–13, 2016. doi: 10.5670/oceanog.2016.66.

P. Urban, K. Köser, and J. Greinert. Processing of multibeam water column image data for automated bubble/seep detection and repeated mapping. *Limnology and Oceanography: Methods*, 15:1–21, 2017. doi: 10.1002/lom3.10138.

S. Vagle and D. M. Farmer. The Measurement of Bubble-Size Distributions by Acoustical Backscatter. *Journal of Atmospheric and Oceanic Technology*, 9:630–644, 1991. doi: 10.1175/1520-0426(1992)009⟨0630:TMOBSD⟩2.0.CO;2.

H. A. Viehman and G. B. Zydlewski. Fish Interactions with a Commercial-Scale Tidal Energy Device in the Natural Environment. *Estuaries and Coasts*, 38:1–12, 2014. doi: 10.1007/s12237-014-9767-8.

B. C. Wallace, K. Small, C. E. Brodley, and T. A. Trikalinos. Class Imbalance, Redux. In *2011 IEEE 11th International Conference on Data Mining*, pages 754–763. IEEE Computer Society, 2011. doi: 10.1109/ICDM.2011.33.

B. Williamson, S. Fraser, P. Blondel, P. S. Bell, J. J. Waggit, and B. E. Scott. Multisensor Acoustic Tracking of Fish and Seabird Behavior Around Tidal Turbine Structures in Scotland. *IEEE Journal of Ocean Engineering*, 42:1–18, 2017. doi: 10.1109/JOE.2016.2637179.

List of Figures

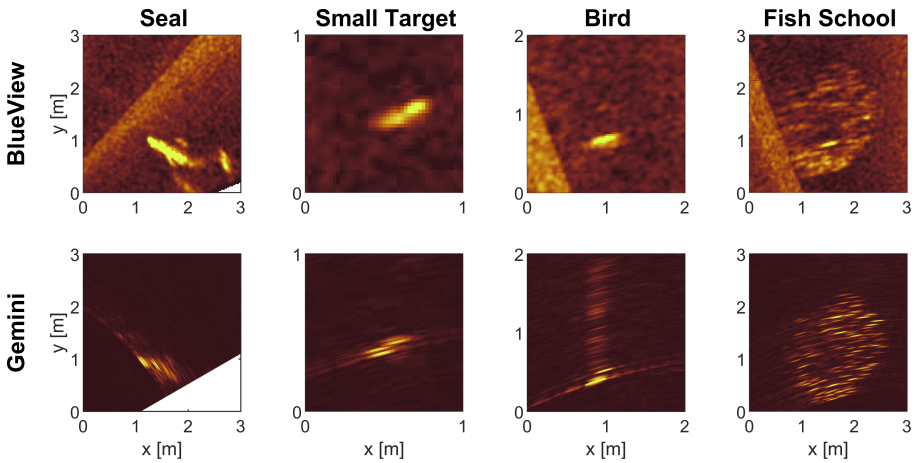

Figure 1: Concurrent detection of each biological target class on the BlueView and Gemini sonars. Note that the spatial scale varies between classes. The same intensity scale is used for both sonars.

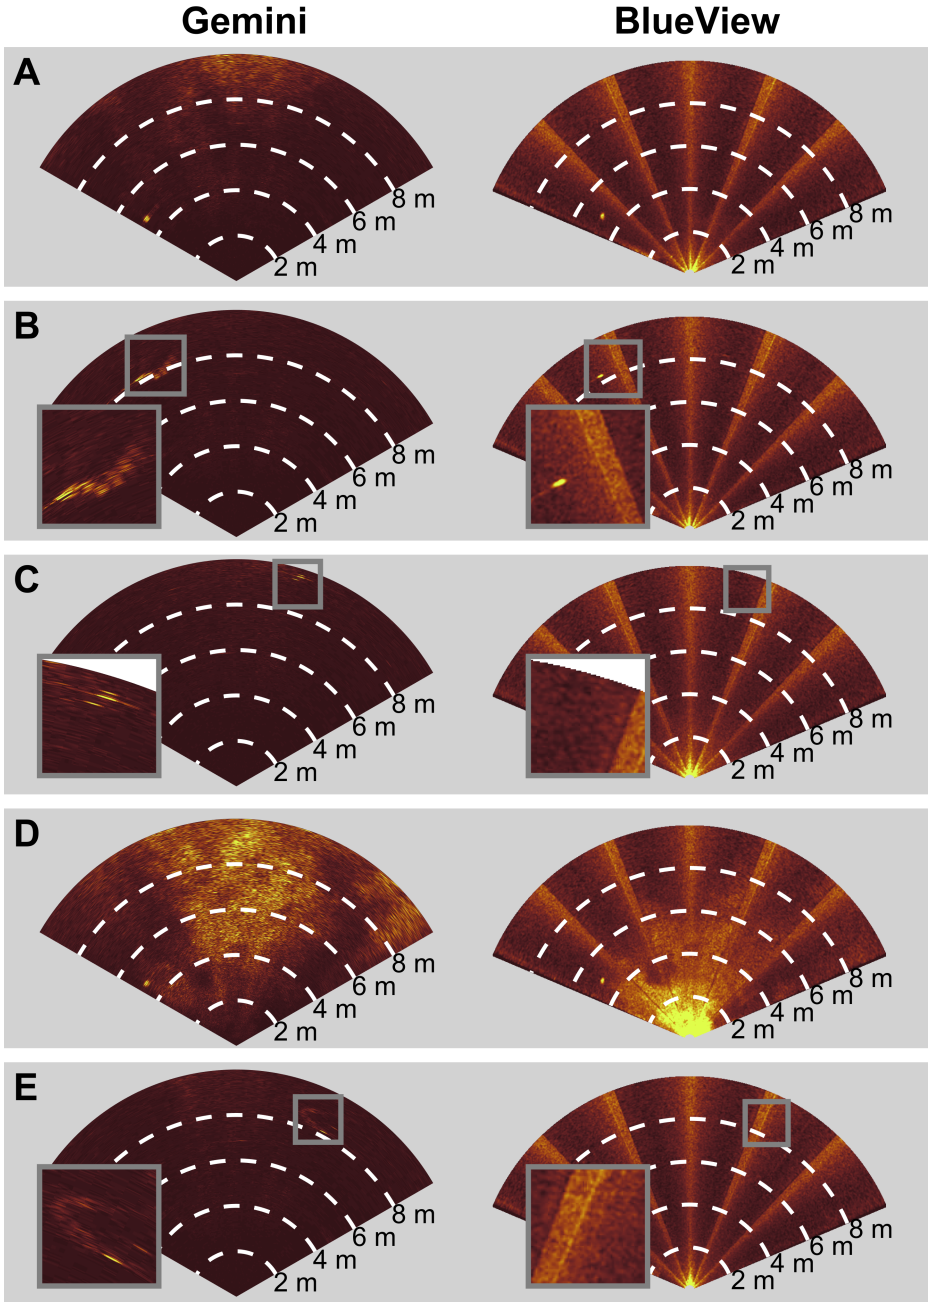

Figure 2: Concurrent sonar imagery from the BlueView and Gemini sonars. A) Surface interference visible in the Gemini data, but not the BlueView data, B) concurrent detection of a diving bird (bubbles entrained in the dive are apparent in the Gemini data, but not the BlueView data), C) detection of a small target (possibly drifting kelp) on the Gemini that is not visible on the BlueView due to attenuation, D) a boat wake visible at close range on the BlueView and long range on the Gemini, and E) detection of a diving bird in the Gemini data that is not visible in the BlueView data due to low signal-to-noise in the beam overlap region. The persistent target at a range of approximately 5 m on the left-hand-side of all images is an instrument platform from another project. The same intensity scale is used for all BlueView and Gemini images, respectively, but different scales are used for each sonar to highlight detected targets. Video sequences containing the images in each subplot can be found in [Cotter and Polagye, 2020b].

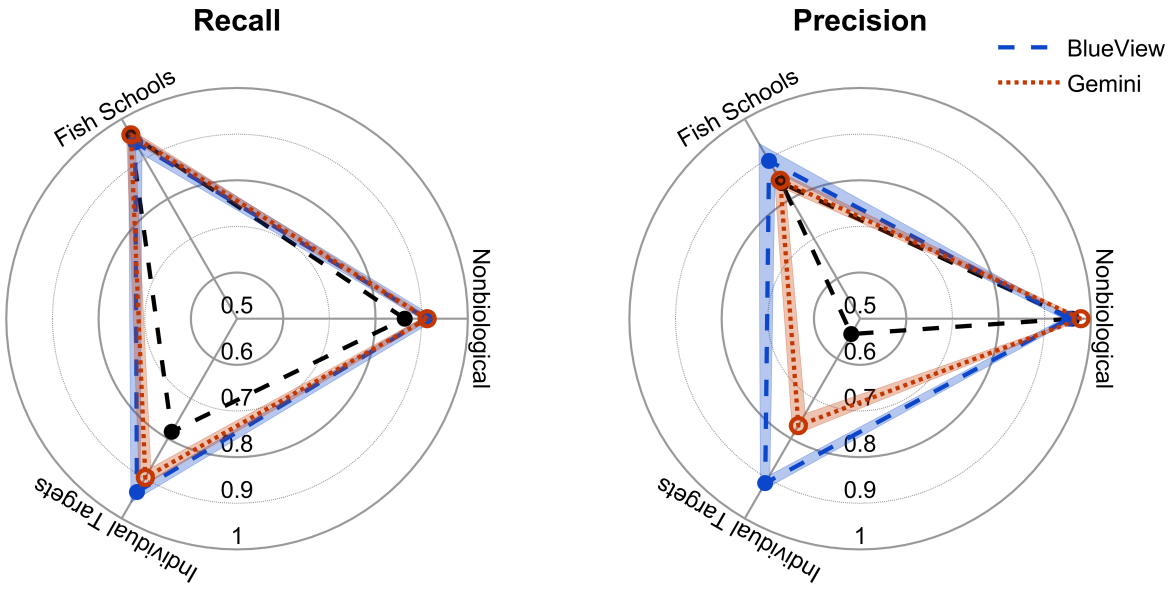

Figure 3: Precision and recall for for each class for both sonars using the target tracks contained in the one-minute sequences as training data. The colored dashed lines indicate the median value for 100 iterations of cross-validation, and the shaded regions indicate the interquartile range. The black dashed line indicates classification performance for the Gemini when all available target tracks are used for classification (one-minute sequences and continuously acquired data). Note that the center of the plot is 0.5, not 0.

List of Tables

Reference	Sonar Manufacturer	Sonar Model)	Operating Frequency (kHz)
Melvin and Cochrane [2015]	Kongsberg	MS 2000	200
Williamson et al. [2017]	Imagenex	837B Delta T	260
Lieber et al. [2014]	Teledyne Reson	Seabat 7128	200/400
Cotter et al. [2017]	Kongsberg	M3	500
Lieber et al. [2017]	Kongsberg	M3	500
Hastie et al. [2019]	Tritech	Gemini 7209id	720
Francisco and Sundberg [2019]	Teledyne BlueView	M900-130-S-MKS-VDSL	900
Viehman and Zydlewski [2014]	Sound Metrics	Didson	1800
Cotter et al. [2017]	Teledyne BlueView	BlueView M900-2250	2250

Table 1: An overview of multibeam sonars tested for environmental monitoring at marine renewable energy sites from 2014-2019

Manufacturer	Tritech	Teledyne
Model	Gemini 720is	BlueView M900-2250
Operating Frequency (kHz)	720	2250
Operating Range (m)	10	10
Ping Rate (Hz)	5	5
Along-Swath Dimension (^{o})	120	130
Across-Swath Dimension (^{o})	20	20

Table 2: Overview of the two multibeam sonars tested. The operating range of the Gemini can be extended beyond 10 m, though resolution decreases with range. The range was limited to 10 m here to provide direct comparison with the BlueView. The BlueView has two transducers that operate at 900 and 2250 kHz, but only the 2250 kHZ transducer was used. Both sonars are capable of acquiring data at ping rates above 5 Hz, but ping rates were limited by synchronization to minimize crosstalk between the sonars.

	Non-biological	Individual Targets	Fish Schools
BlueView (one-minute sequences)	985	261	169
Gemini (one-minute sequences)	299	252	115
Gemini (continuous)	880	3	95

Table 3: Number of target tracks in each class from one-minute sequences used to evaluate classification model performance for each sonar. The tracks from the continuous Gemini data set were added to the one-minute sequences to train the classification model associated with black, dashed line in Figure 3. Due to ambiguity in the sonar data, the target tracks in the continuous data annotated by the human reviewer as "non-biological" likely contains a number of small, individual targets.