

Seabed and range estimation of impulsive time series using a convolutional neural network

David F. Van Komen, Tracianne B. Neilsen, Kira Howarth, David P. Knobles, and Peter H. Dahl

Citation: *The Journal of the Acoustical Society of America* **147**, EL403 (2020); doi: 10.1121/10.0001216

View online: <https://doi.org/10.1121/10.0001216>

View Table of Contents: <https://asa.scitation.org/toc/jas/147/5>

Published by the [Acoustical Society of America](#)

ARTICLES YOU MAY BE INTERESTED IN

[Nonlinear time-warping made simple: A step-by-step tutorial on underwater acoustic modal separation with a single hydrophone](#)

The Journal of the Acoustical Society of America **147**, 1897 (2020); <https://doi.org/10.1121/10.0000937>

[Source localization in the deep ocean using a convolutional neural network](#)

The Journal of the Acoustical Society of America **147**, EL314 (2020); <https://doi.org/10.1121/10.0001020>

[Estimation of seabed properties and range from vector acoustic observations of underwater ship noise](#)

The Journal of the Acoustical Society of America **147**, EL345 (2020); <https://doi.org/10.1121/10.0001089>

[Ship source level estimation and uncertainty quantification in shallow water via Bayesian marginalization](#)

The Journal of the Acoustical Society of America **147**, EL339 (2020); <https://doi.org/10.1121/10.0001096>

[A denoising representation framework for underwater acoustic signal recognition](#)

The Journal of the Acoustical Society of America **147**, EL377 (2020); <https://doi.org/10.1121/10.0001130>

[Tracking of multiple surface vessels based on passive acoustic underwater arrays](#)

The Journal of the Acoustical Society of America **147**, EL87 (2020); <https://doi.org/10.1121/10.0000598>

READ NOW!

JASA
THE JOURNAL OF THE
ACOUSTICAL SOCIETY OF AMERICA

Special Issue:
Acoustic Localization

Seabed and range estimation of impulsive time series using a convolutional neural network

.....
David F. Van Komen,^{1,a)} Tracianne B. Neilsen,¹ Kira Howarth,¹ David P. Knobles,²
and Peter H. Dahl³

¹Physics and Astronomy, Brigham Young University, Provo, Utah, 84604, USA

²Knobles Scientific and Analysis, Austin, Texas 73301, USA

³University of Washington, Seattle, Washington 98195, USA

david.vankomen@gmail.com, tbn@byu.edu, howarthke@gmail.com, dpknobles@kphysics.org,
dahl@apl.washington.edu

Abstract: In ocean acoustics, many types of optimizations have been employed to locate acoustic sources and estimate the properties of the seabed. How these tasks can take advantage of recent advances in deep learning remains as open questions, especially due to the lack of labeled field data. In this work, a Convolutional Neural Network (CNN) is used to find seabed type and source range simultaneously from 1 s pressure time series from impulsive sounds. Simulated data are used to train the CNN before application to signals from a single hydrophone signal during the 2017 Seabed Characterization Experiment. The training data includes four seabeds representing deep mud, mud over sand, sandy silt, and sand, and a wide range of source parameters. When applied to measured data, the trained CNN predicts expected seabed types and obtains ranges within 0.5 km when the source-receiver range is greater than 5 km, showing the potential for such algorithms to address these problems. © 2020 Acoustical Society of America

[Editor: Matthew A. Dzieciuch]

Pages: EL403–EL408

Received: 27 December 2019 Accepted: 15 April 2020 Published Online: 7 May 2020

1. Introduction

In ocean acoustics research, knowledge of the ocean environment is important for localizing acoustic sources. However, localizing acoustic sources, using an algorithm such as matched-field processing (MFP),¹ is difficult in an uncertain ocean environment. Matched-field inversions² have been used in many cases to estimate environmental parameters.^{3,4} Simultaneous estimation of source location and environmental characteristics of an unknown ocean environment presents multiple challenges.^{4–6} Some of these challenges include nonlinear relationships between the unknowns, high-dimensional search spaces, and large uncertainty due to ill-conditioning of the inferred parameter values.

Recently, machine and deep learning approaches to the problems of source localization have gained traction. These efforts include using neural networks to localize a point source in a homogeneous medium,⁷ supervised learning for real-time range estimation,⁸ classifying underwater targets,⁹ and other machine learning techniques to localize sources in the ocean.^{10,11} Some studies, such as those done by Lefort *et al.*¹² and Niu *et al.*^{13,14} have found that machine learning classifiers outperform MFP, further suggesting the relevance of using machine and deep learning in the underwater acoustics field.

Machine learning has also been used to make predictions about ocean environmental parameters such as attenuation and sound speed. Early efforts included the use of neural and statistical classifiers by Michalopoulou *et al.*,¹⁵ artificial networks on transmission loss by Benson *et al.*,¹⁶ and global and hierarchical approaches by Stephan *et al.*¹⁷ Recently, Piccolo *et al.*¹⁸ used generalized additive models for predicting compressional sound speed and attenuation from time-domain features. A more detailed review of machine learning applications in various fields of acoustics can be found in Bianco *et al.*¹⁹

In this paper, a machine learning model known as a convolutional neural network (CNN) is used to simultaneously predict source range and seabed type. This CNN is trained on synthetic pressure time waveforms and, as a proof of concept, applied to waveforms measured on a single pressure sensor²⁰ from an impulsive signals underwater sound (SUS) source^{21,22} in the New England Mudpatch during the 2017 Seabed Characterization Experiment (SBCEX2017).

^{a)} Author to whom correspondence should be addressed.

2. Background and methodology

Several steps are needed before a CNN can simultaneously predict source range and environment type on a real-world measurement. First, the CNN training data need to reflect the real-world testing data and incorporate the variability likely to be found. The real-world testing data used in this work is described in Sec. 2.1. Second, due to the lack of real-world labeled data, simulated training data were generated using a verified source spectrum and a propagation model, as described in Sec. 2.2. Finally, the training data are used to tune the hyperparameters of the CNN. The CNN model architecture and tools used to build and train the model are presented in Sec. 2.3.

2.1 Intensity vector autonomous recorder in SBCEX2017

SBCEX2017 (Ref. 23) was conducted on the New England Mud Patch centered at [40°28' N, 70°35' W] in the Spring of 2017. The Intensity Vector Autonomous Recorder (IVAR) system deployed by the Applied Physics Laboratory, University of Washington, recorded the waveforms of SUS Mk64 charges deployed at ranges of 2–13 km. Three to five SUSs were deployed at each location. More details about the system, the recordings, locations of the SUS stations, and the experiment conducted can be found in Ref. 20.

While IVAR has multiple sensors for measuring both pressure and particle velocity, the data used here are restricted to the pressure sensor located 1.32 m above the seabed. Examples are displayed in Figs. 1(a) and 1(b). Because the signals were downsampled to 5000 Hz and isolated to 1 s, the simulated training data are each 1 s long and have a sampling frequency of 5000 Hz. To facilitate the generation of trained data for this study and to provide proof of concept, only the data collected from a single pressure sensor IVAR was used; additional SBCEX2017 sensors can be used in future studies.

The environment at SBCEX2017 informed the simulation of the training data. The SUS charges were deployed in an area with an average ocean depth of 74 m and a nearly isospeed sound speed profile in the water [as seen in Fig. 2(b) in Ref. 20]: varying from 1467.8 to 1468.6 m/s. The New England Mud Patch has an uppermost sediment layer consisting of fine-grained mud-like material.²⁴ Thus, training data are simulated for a variety of linear sound speed profiles and four different seabed types, one of which was inferred from the analysis of the SUS waveforms on a different receiving array.²⁵

2.2 Synthetic data generation

Machine learning models require a significant amount of training data to identify and learn patterns. Because of the lack of labeled field data, synthetic training data are used. The range-independent, normal-mode model ORCA (Ref. 26) model was used to generate the ocean's impulse response at frequencies up to 2500 Hz. This impulse response is convolved with a simulated SUS Mk64 signal spanning a bandwidth of 5–2500 Hz to produce the simulated time-series

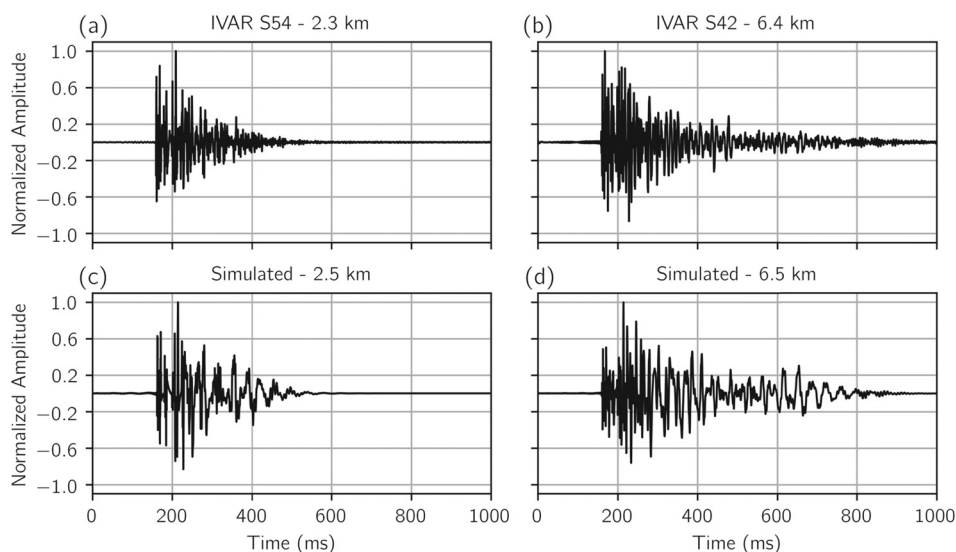


Fig. 1. Normalized pressure waveforms measured on IVAR from SUS station S54 (a), and S42 (b) (the exact positions are shown in Fig. 3 of Ref. 20) and two simulated, normalized waveforms, in a representative environment, at 2.5 (c) and 6.5 km (d) away from the source.

waveform using the model in Wilson *et al.*²² Examples of the simulated samples are shown in Figs. 1(c) and 1(d).

The time series were simulated using four different seabeds representing deep mud, mud over sand, sandy silt, and sand environments. The “deep mud” environment comes from a study done in the gulf of Mexico by Knobles *et al.*²⁷ The “mud over sand” environment comes from maximum entropy statistical inference on data collected during the SUS circle experiments in SBCEX.²⁵ The “sandy silt” environment comes from geoacoustic inversions done in the New England Bight by Potty *et al.*²⁸ The “sand” environment comes from a study on sandy seafloors done by Zhou *et al.*²⁹ A visualization of these seabeds and values for sound speeds, densities, and attenuations is shown in Fig. 1 of Van Komen *et al.*³⁰

A variety of sound speed profiles were also used in the simulation to expand the training dataset. Fifty different linear, downward-refracting sound speed profiles were used with water depths between 73 and 75 m, representative of the measurements taken at SBCEX2017. These profiles (spanning 1465–1470 m/s) were selected to add variability that was present in the measurements. This variability also allows the machine learning model to better generalize the internal feature extraction used for prediction.

Source and receiver parameters were selected to cover the variability in the measured IVAR data. The signals were simulated to be received on a hydrophone located 1.3 m from the ocean floor at ranges between 0.5 and 15 km at 0.5 km intervals. The true IVAR-to-SUS charge ranges (as reported by station locations in Ref. 20) were not used in the simulations so as to test the model’s ability to generalize on the range label. Because the nominal source depth of the SUS charges is 18.3 m, the simulations were performed at eight different source depths spanning 10.3 to 24.3 m from the surface to account for any variability.

With the four seabed types, 50 sound speed profiles, 30 ranges, and eight source depths, a dataset of 48 000 signals was generated. The 1 s signals, normalized and sampled at 5000 Hz, were aligned with the arrival time occurring at approximately the same time in each sample, as illustrated in Fig. 1. Thus, the data contain no information about absolute travel time or absolute amplitude.

2.3 Machine learning model architecture and training

The machine learning model employed in this study is a CNN. A CNN is chosen due to its ability to find patterns in gridded data, which makes it an ideal candidate for analysis on time-series waveforms. Using a CNN also eliminates the need for data preprocessing because a CNN learns to extract features necessary for accurate predictions. More information on CNNs can be found in Goodfellow *et al.*³¹ The CNN architecture used in this study is a simple network with four convolutional layers and two fully connected layers and is the same network architecture found and explained in more detail in Van Komen *et al.*³⁰

To implement and train the machine learning models, the Python package PyTorch (Ref. 32) is used. PyTorch automates many of the functions used in machine learning for developing and training models and provides an interface for learning on a GPU, a necessity for training quickly on 48 000 training samples. A Tesla T4 GPU by NVIDIA was used to accelerate computations by a factor of 5 (approximately 15 minutes versus 75 minutes on an Intel Xeon Gold 6134 CPU) for this study. PyTorch also provides the algorithms needed for the models to learn, such as the Adam³³ optimizer used to train the network’s weights.

The dataset of 48 000 simulated signals were labeled with the source-receiver range (in km) and a number representing the seabed type. The dataset was randomly divided into a training-validation split of 95%/5%, leading to 45 600 training and 2400 validation samples. This split was chosen because the final testing dataset is the IVAR signals. Validation errors for this network are reported in Van Komen *et al.*³⁰

For a single training session, the CNN was trained over 200 epochs. The learning rate began at 0.001 and was annealed via a cosine function to allow the algorithms to make large adjustments in the early epochs and then make smaller adjustments in the later epochs. These hyperparameters were selected over several trial runs and led to a sufficiently low validation error; the results on the IVAR SBCEX2017 data (Sec. 3) also confirm the selections. Unfortunately, there is no special formula or method for selecting hyperparameters, so there is a possibility that other hyperparameters could give more precise results, but for the sake of consistency, these parameters were chosen and used across all tests in this study.

The network was trained multiple times to get a broader picture of the potential the CNN has to make predictions on the IVAR SBCEX2017 data. This process of training multiple networks was chosen to account for the random initialization of weights and the random training-validation split. Ten instances of this training-validation split led to ten different networks that were then applied to the IVAR SBCEX2017 data. The predictions from these ten different networks are shown and discussed in Sec. 3.

3. Results and discussion

This section has been divided into two sections: results from a single trained network applied to the IVAR SBCEX2017 data and results from multiple trained networks. The division of the results is to illustrate how different random initializations and splits of training data can lead to different predictions. The network, in all cases, attempts to make predictions via regression and outputs two values for each sample: the range (in km) and a number representing the seabed type.

3.1 Individual network results

After training the CNN on the simulated data, the 37 IVAR SBCEX2017 data samples were passed through the network to obtain predictions of range and seabed type. The predictions of the CNN from two separate networks are shown in Fig. 2. The predicted ranges are compared to the measured IVAR-to-SUS ranges spanning 2–12 km. If the network predictions matched the true ranges, the points would lie exactly on the diagonal line. In both cases, the signals of range 7 km and closer are overpredicted by the network, and those farther away are closer to the truth or slightly underpredicted. In these cases, the root mean squared error for all range predictions is 0.84 km.

For each data sample, the CNN also yields a number corresponding to a prediction of the seabed type. The predicted seabed type is a decimal number, which is rounded to the nearest integer to easily display the results. In Fig. 2, the different symbols/colors indicate which seabed type is predicted. For the two cases shown, the CNN predicts that the signals farther than 5 km from IVAR are identified as being from the mud over sand seabed (green squares), which most closely resembles the New England Mudpatch area. However, the network is predicting environments that are more reflective: the sandy silt environment (blue triangles), the “sandy” environment (red pluses), and even a “5th” environment (black crosses) which does not exist for the closer ranges. This variability likely ties to the physics of the sound propagation. One potential explanation is the variation in propagation paths. Another potential explanation is that the longer ranges sense primarily the upper portion of the seafloor, and, when the range is shorter, deeper features of the seafloor have more impact on the propagation. Thus, the mud over sand seafloor represents what influences propagation to ranges greater than 6 km, while the tendency toward a seafloor with less bottom loss at short ranges indicates that the deeper features of the mud over sand seafloor are likely too absorptive.

3.2 Multiple network results

After training and applying the network that produced the predictions shown in Fig. 2, nine more networks were trained to test how well the network architecture generalizes with multiple initializations. The number of networks (ten) was selected because adding more instances did not significantly change the distribution of the results, which are displayed using violin plots. Violin plots show the median (white dot), inter-quartile range (black rectangle), and the full distribution of the data (in color). The violin plots for results from the ten trained networks are shown in Fig. 3, where Fig. 3(a) shows the distribution of ranges and Fig. 3(b) shows the same for environment number.




Fig. 2. (Color online) Predictions from one training instance of the CNN on the 37 data samples of SUS signals recorded on IVAR. The “True Range” corresponds to the measured distance; the diagonal dashed line indicates where the “ideal” predictions should lie. The symbol indicates the predicted seabed type. The pluses indicate the sandy seabed, the diamonds indicate the sandy-silt seabed, and the squares indicate the mud over sand seabed while the crosses indicate an environment out of scope.

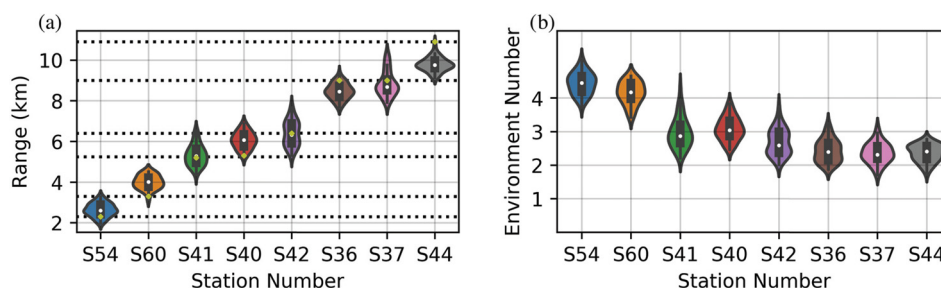


Fig. 3. (Color online) Violin plots (a combination of probability density and box plots) showing the distributions of the predictions a CNN made on IVAR data samples for (a) range and (b) seabed type separated by the SUS station number. (A map of the stations is shown in Fig. 3 of Ref. 20. The yellow diamonds show the expected values, with the horizontal dotted lines highlighting the true range. The seabed types are numbered as 1: deep mud, 2: mud over sand, 3: sandy silt, and 4: sandy.

The violin plots for range [Fig. 3(a)] show how well the network learned to predict IVAR-to-SUS ranges using the simulated training dataset. The median ranges (white dots) follow the same trend as the expected values. For several of the stations, the centers of the distributions are close to the expected values (S54, S41, S42, and S37), while the other centers are off by 1–2 km, with closer ranges (S60 and S40) being overpredicted and longer ranges (S36 and S44) being underpredicted. These results show that the network is learning how to predict ranges, though there is certainly room for improvement.

The violin plots for seabed type [Fig. 3(b)] reflect the trend seen in Fig. 2. For the two closest stations (Stations S54 and S60 at approximately 2.4 and 3.3 km), the network never predicts a seabed below three. The next two stations (S41 and S40 at 5.3 and 5.2 km, as well as some results from S42 at 6.2 km) also predict seabeds closer to three than the expected two. However, the results for the remaining stations show the majority of their distributions approach the mud over sand environment with some extension to the sandy silt environment. In addition, the distribution of the seabed type predictions becomes narrower as the range increases. These plots suggest that the network learns how to predict the seabed type, but some additional information needs to be included in the training dataset to increase the efficacy of the network at close ranges.

4. Conclusion

This paper has shown that a CNN can be trained on simulated data to make seabed type and range predictions simultaneously on real-world data. As the IVAR-to-SUS range increases, the mean distribution of range predictions tends to be underpredicted as the environmental predictions tend toward the correct seabed. At ranges less than 5 km, a more correct answer for range tends to be coupled with a prediction of a more reflective seabed type. Although predictions from this network are not perfect, these results show the potential for machine learning models to make range and seabed predictions simultaneously if the training dataset represents the variability of the real world data. The violin plots shown here also provide one way of tracking uncertainty due to random initialization of the network.

Future work will seek to improve the results by refining the simulated dataset to include more variation in the environment and range to allow for further generalization. For example, only downward refracting sound speeds were used in this study, though some measurements also show that at some locations an isovelocity profile was observed. This variation was not included in these simulations, so a larger dataset that takes varying ocean sound speeds could prove beneficial. As with all applications of machine learning, there is also the possibility of developing deeper and more advanced networks to improve results.

Acknowledgments

This research was supported in part by the Office of Naval Research, SBIR Grant No. N68335-18-C-0806.

References and links

- ¹A. Tolstoy, *Matched Field Processing for Underwater Acoustics* (World Scientific, Singapore, 1993).
- ²M. D. Collins, W. A. Kuperman, and H. Schmidt, “Nonlinear inversion for ocean-bottom properties,” *J. Acoust. Soc. Am.* **92**(5), 2770–2783 (1992).
- ³S. Dosso, M. Jeremy, J. Ozard, and N. Chapman, “Estimation of ocean-bottom properties by matched-field inversion of acoustic field data,” *IEEE J. Oceanic Eng.* **18**(3), 232–239 (1993).
- ⁴Z.-H. Michalopoulou and U. Ghosh-Dastidar, “Tabu for matched-field source localization and geoacoustic inversion,” *J. Acoust. Soc. Am.* **115**(1), 135–145 (2004).

- ⁵T. B. Neilsen, "An iterative implementation of rotated coordinates for inverse problems," *J. Acoust. Soc. Am.* **113**(5), 2574–2586 (2003).
- ⁶S. E. Dosso and M. J. Wilmut, "Bayesian multiple-source localization in an uncertain ocean environment," *J. Acoust. Soc. Am.* **129**(6), 3577–3589 (2011).
- ⁷B. Z. Steinberg, M. J. Beran, S. H. Chin, and J. H. Howard, Jr., "A neural network approach to source localization," *J. Acoust. Soc. Am.* **90**(4), 2081–2090 (1991).
- ⁸L. Houégnigan, P. Safari, C. Nadeu, M. van der Schaar, and M. André, "A novel approach to real-time range estimation of underwater acoustic sources using supervised machine learning," in *IEEE OCEANS 2017-Aberdeen* (2017), pp. 1–5.
- ⁹E. M. Fischell and H. Schmidt, "Classification of underwater targets from autonomous underwater vehicle sampled bistatic acoustic scattered fields," *J. Acoust. Soc. Am.* **138**(6), 3773–3784 (2015).
- ¹⁰Z. Huang, J. Xu, Z. Gong, H. Wang, and Y. Yan, "Source localization using deep neural networks in a shallow water environment," *J. Acoust. Soc. Am.* **143**(5), 2922–2932 (2018).
- ¹¹Y. Wang and H. Peng, "Underwater acoustic source localization using generalized regression neural network," *J. Acoust. Soc. Am.* **143**(4), 2321–2331 (2018).
- ¹²R. Lefort, G. Real, and A. Drémeau, "Direct regressions for underwater acoustic source localization in fluctuating oceans," *J. Appl. Acoust.* **116**, 303–310 (2017).
- ¹³H. Niu, E. Ozanich, and P. Gerstoft, "Ship localization in Santa Barbara Channel using machine learning classifiers," *J. Acoust. Soc. Am.* **142**(5), EL455–EL460 (2017).
- ¹⁴H. Niu, E. Reeves, and P. Gerstoft, "Source localization in an ocean waveguide using supervised machine learning," *J. Acoust. Soc. Am.* **142**(3), 1176–1188 (2017).
- ¹⁵Z. Michalopoulou, D. Alexandrou, and C. D. Moustier, "Application of neural and statistical classifiers to the problem of seafloor characterization," *IEEE J. Ocean. Eng.* **20**(3), 190–197 (1995).
- ¹⁶J. Benson, N. R. Chapman, and A. Antoniou, "Geoacoustic model inversion using artificial neural networks," *Inverse Problems* **16**(6), 1627 (2000).
- ¹⁷Y. Stephan, X. Demoulin, and O. Sarzeaud, "Neural direct approaches for geoacoustic inversion," *J. Comput. Acoust.* **6**(01n02), 151–166 (1998).
- ¹⁸J. Piccolo, G. Haramuniz, and Z.-H. Michalopoulou, "Geoacoustic inversion with generalized additive models," *J. Acoust. Soc. Am.* **145**(6), EL463–EL468 (2019).
- ¹⁹M. J. Bianco, P. Gerstoft, J. Traer, E. Ozanich, M. A. Roch, S. Gannot, and C.-A. Deledalle, "Machine learning in acoustics: Theory and applications," *J. Acoust. Soc. Am.* **146**(5), 3590–3628 (2019).
- ²⁰P. H. Dahl and D. R. Dall'Osto, "Vector acoustic analysis of time-separated modal arrivals from explosive sound sources during the 2017 seabed characterization experiment," *IEEE J. Ocean. Eng.* **45**, 131–143 (2019).
- ²¹N. R. Chapman, "Source levels of shallow explosive charges," *J. Acoust. Soc. Am.* **84**(2), 697–702 (1988).
- ²²P. S. Wilson, D. P. Knobles, P. H. Dahl, A. R. McNeese, and M. C. Zeh, "Short-range signatures of explosive sounds in shallow water used for seabed characterization," *IEEE J. Ocean. Eng.* **45**, 14–25 (2019).
- ²³P. S. Wilson, D. P. Knobles, and T. B. Neilsen, "Guest editorial an overview of the seabed characterization experiment," *IEEE J. Oceanic Eng.* **45**(1), 1–13 (2020).
- ²⁴J. A. Goff, A. H. Reed, G. Gawarkiewicz, P. S. Wilson, and D. P. Knobles, "Stratigraphic analysis of a sediment pond within the New England Mud Patch: New constraints from high-resolution chirp acoustic reflection data," *Marine Geol.* **412**, 81–94 (2019).
- ²⁵D. P. Knobles, P. S. Wilson, J. A. Goff, L. Wan, M. J. Buckingham, J. D. Chaytor, and M. Badiéy, "Maximum entropy derived statistics of sound-speed structure in a fine-grained sediment inferred from sparse broadband acoustic measurements on the New England Continental Shelf," *IEEE J. Ocean. Eng.* **45**, 1–13 (2019).
- ²⁶E. K. Westwood, C. T. Tindle, and N. R. Chapman, "A normal mode model for acousto-elastic ocean environments," *J. Acoust. Soc. Am.* **100**(6), 3631–3645 (1996).
- ²⁷D. P. Knobles, R. A. Koch, L. A. Thompson, K. C. Focke, and P. E. Eisman, "Broadband sound propagation in shallow water and geoacoustic inversion," *J. Acoust. Soc. Am.* **113**(1), 205–222 (2003).
- ²⁸G. R. Potty, J. H. Miller, and J. F. Lynch, "Inversion for sediment geoacoustic properties at the New England bight," *J. Acoust. Soc. Am.* **114**(4), 1874–1887 (2003).
- ²⁹J.-X. Zhou, X.-Z. Zhang, and D. P. Knobles, "Low-frequency geoacoustic model for the effective properties of sandy seabottoms," *J. Acoust. Soc. Am.* **125**(5), 2847–2866 (2009).
- ³⁰D. F. Van Komen, T. B. Neilsen, D. P. Knobles, and M. Badiéy, "A convolutional neural network for source range and ocean seabed classification using pressure time-series," *Proc. Meet. Acoust.* **36**(1), 070004 (2019).
- ³¹I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, Cambridge, MA, 2016).
- ³²A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, and L. Antiga, "PyTorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems* (Curran Assoc. Inc., Red Hook, NY, 2019), pp. 8024–8035.
- ³³D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014).