Physics-informed Data-driven Communication Performance Prediction for Underwater Vehicles

Mandar Chitre and Li Kexin

***Abstract*—Underwater vehicles usually rely on acoustics for communication and navigation. Reliable communication and accurate navigation require the vehicle to plan a path through areas with good acoustic coverage from communication gateways and beacons. Planning such a path can be challenging in areas with complex acoustic propagation, especially when the signal strength does not monotonically reduce as a function of distance from a transmitter. When the environmental parameters are not fully known, traditional acoustic propagation models are unable to provide accurate predictions. We develop an online physics-informed data-driven method to predict acoustic signal quality in a region ahead of the underwater vehicle to inform the vehicle's path-planning algorithm.**

I. Introduction

While most autonomous underwater vehicles (AUVs) today use acoustics for communication and navigation, operators find that the performance of acoustic systems can be highly variable and often cannot be relied upon. Especially in littoral waters, multipath arrivals interfere constructively and destructively, creating zones of high and low signal strength and highly variable delay spread. Acoustic performance in such channels does not monotonically reduce with range, but is a complicated function of source and receiver location. If an AUV can predict acoustic performance as a function of its own location, it may be able to plan a path to provide guarantees on navigation and communication performance [1], [2].

If the acoustic environment is well understood, physics-based acoustic propagation models such as BELLHOP [3] may be used to predict performance. However, in practice, sufficiently detailed environmental knowledge is usually not available a priori to plan an acoustically performant path. While signal quality measurements may be available as the vehicle executes its mission, physics-based acoustic propagation models cannot readily incorporate these measurements to improve their predictions. In this paper, we demonstrate how an online data-driven method may be used to predict acoustic signal quality in a region ahead of the vehicle to inform the vehicle's path-planning algorithm. The method is informed by acoustic propagation physics to enable it to learn about the acoustic environment with only a few measurements.

II. Problem Formulation

In a small region in the far field of a source, the acoustic pressure amplitude measured at frequency f and location $\mathbf{r}$ can be expressed in terms of a sum of pressure amplitudes of M plane waves (each represented by a ray):

$$p(\mathbf{r}, f) = \left| \sum_{m=1}^{M} A_m e^{ik\hat{\mathbf{k}}_m \cdot \mathbf{r}} \right|, \tag{1}$$

where A_m is the complex amplitude of the m^{th} ray, $k = 2\pi f/c$ is the wavenumber, c is the speed of sound, and $\hat{\mathbf{k}}_m$ is a unit vector along the direction of travel of the ray. If we make measurements $\hat{P}_{jn}$ of acoustic field amplitude at various frequencies f_n and locations $\mathbf{r}_j$, we can estimate the unknown pressure amplitudes A_{mn} as a solution to a non-linear optimization problem:

$$\arg\min_{\{A_{mn}\}} \sum_{j,n} \left(\hat{P}_{jn} - \left| \sum_m D_{jnm} A_{mn} \right| \right)^2, \tag{2}$$

where $D_{jnm} \equiv e^{ik_n \hat{\mathbf{k}}_m \cdot \mathbf{r}_j}$, and A_{mn} is the complex amplitude of the m^{th} ray for frequency f_n.

In general, we do not know the number of rays M or the direction vectors $\hat{\mathbf{k}}_m$. We can, however, take M to be large and spread the $\hat{\mathbf{k}}_m$ uniformly over the range of possible directions, and allow the optimization to set A_{mn} for directions without a ray to be 0. Since we know a priori that M is small, we desire solutions to (2) that exhibit row sparsity in matrix $[A_{mn}]$. Once we have the estimated values of A_{mn}, we can determine the pressure amplitude (and hence acoustic signal strength[1]) at any frequency and location using (1).

III. Optimization

Row-sparse solutions to the non-linear optimization problem in (2) can be obtained adding a 1-norm regularization term $\mathcal{R}$ to the original optimization problem:

$$\arg\min_{\{A_{mn}\}} \left[\sum_{j,n} \left(\hat{P}_{jn} - \left| \sum_m D_{jnm} A_{mn} \right| \right)^2 + \mathcal{R} \right], \tag{3}$$

where

$$\mathcal{R} \equiv \xi \sum_m \sqrt{\sum_n |A_{mn}|^2},$$

is the 1-norm of the vector formed from the 2-norms of each row of matrix $[A_{mn}]$ and ξ is a hyper-parameter to control the degree of sparsity.

[1] We use acoustic signal strength as a measure for signal quality in this paper. However, other measures such as delay spread, impulse response, etc. may also be obtained from $p(\mathbf{r}, f)$ through an inverse Fourier transform.

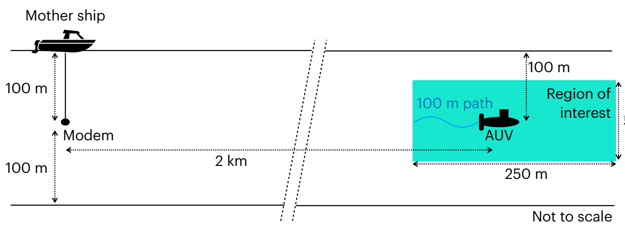

Figure 1: Simulation setup.

The resulting problem is non-convex and is solved iteratively using gradient descent based methods in machine learning, such as the ADAM optimizer [4]. While these methods do not guarantee a global minimum, we find that the minimum they find is usually a very good one, as long as the measurements are scaled appropriately, and the ADAM step size and ξ are chosen carefully[2]. The gradient of the loss function can easily be obtained via automatic differentiation [5].

IV. Simulation Study

A. Simulation setup

Consider an AUV operating in mid-water column, at a range of about 2 km from a mother ship, in a water depth of 200 m and an iso-velocity sound speed profile. The AUV is assumed to know the sound speed in the local area around it, but not the bathymetry or the location of the mother ship. The AUV communicates with the mother ship via an acoustic communication link in the 9–14 kHz band (consistent with the JANUS standard [6]). The mother ship transmits regularly, and the AUV makes acoustic signal strength measurements at every 1 m, over a 100 m distance, and uses this as training data $\{\hat{P}_{jn}\}$. To provide some diversity in the depth axis and to break symmetry, we bootstrap the measurements by asking the AUV to slowly vary its depth by ± 2 m as shown in Figure 1. With only 100 measurements in a very small region (100 m range) along its path, the AUV is then asked to predict the signal strength in a much larger region (250 m range $\times$ 50 m depth) around it. These predictions may then be used for the AUV's path planning.

B. Results

The *ground truth* acoustic signal strength field (Figure 2) is produced over the region of interest using BELLHOP. The AUV path is shown in blue, and the 100 simulated measurement locations are marked with solid dots. The *model* parameters $[A_{mn}]$ are obtained by solving the optimization problem in (2) as described in the previous section. Solving the problem involves discretization of frequency and ray angles. We discretize frequency in steps of 500 Hz, and choose ray angles from -20° to 20° in steps of 0.5°. At a range of 2 km from the source in shallow waters, we do not expect much energy arriving at steeper angles. We perform the optimization over 500 epochs of the training data. The model in (1) is then used to predict the acoustic signal strength field shown in Figure 3. The similarity between the two fields is clearly seen.

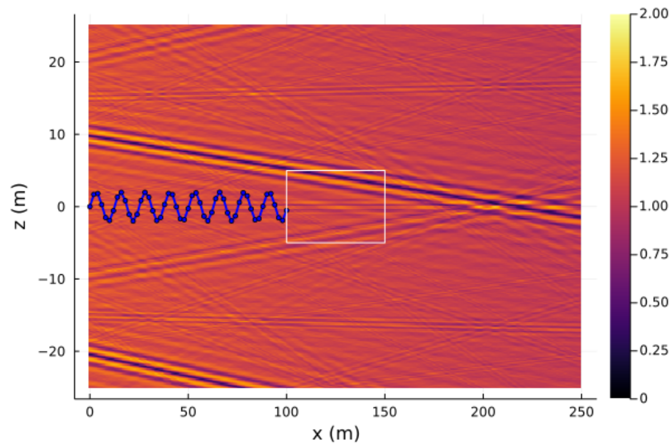

Figure 2: Ground truth signal strength integrated over the 9–14 kHz band, as obtained from BELLHOP over the region of interest. The AUV path is shown in blue, with solid dots showing the measurement locations. The white rectangle shows a region marked for closer inspection.

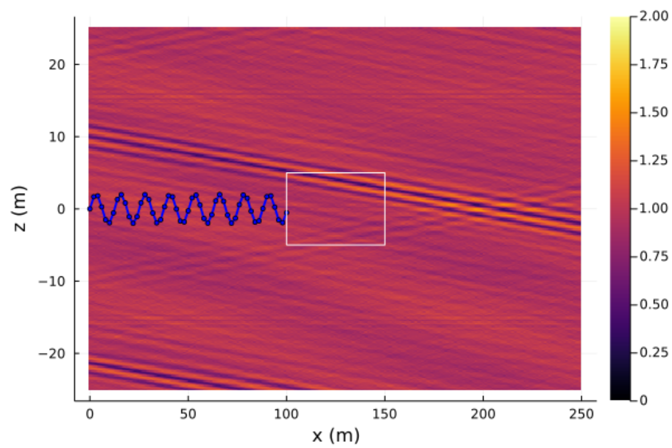

Figure 3: Predicted signal strength integrated over the 9–14 kHz band, over the region of interest. The AUV path is shown in blue, with solid dots showing the measurement locations. The white rectangle shows a region marked for closer inspection.

The white rectangles in Figure 2 and Figure 3 show regions just ahead of the AUV that we choose for closer inspection at specific frequencies. Figure 4 and Figure 5 show the signal strength field in the corresponding zoomed in regions at 9 kHz. Figure 6 and Figure 7 show the signal strength field at 14 kHz. The interference structure is clearly visible in both cases, with deep nulls separating regions of good signal strength. We observe the similarity in field patterns between the ground truth and the predictions. While some of the detailed structure in the ground truth is

[2]For all examples shown in this paper, we scale the signal strength measurements to have unit mean, and use a ADAM step size of 0.5 and sparsity parameter $\xi = 5$.

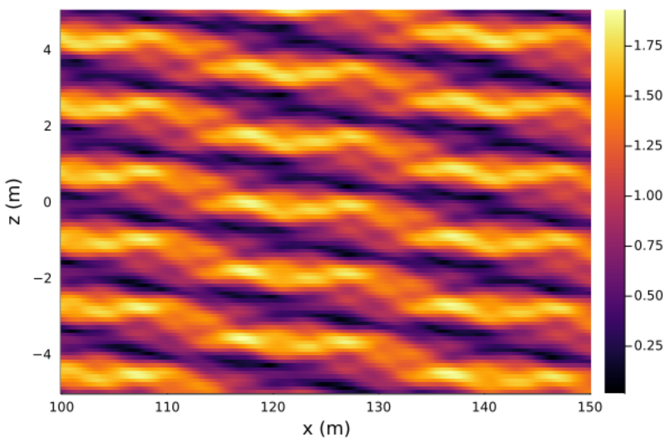

Figure 4: Zoomed in view of the region ahead of the AUV showing ground truth field strength at 9 kHz.

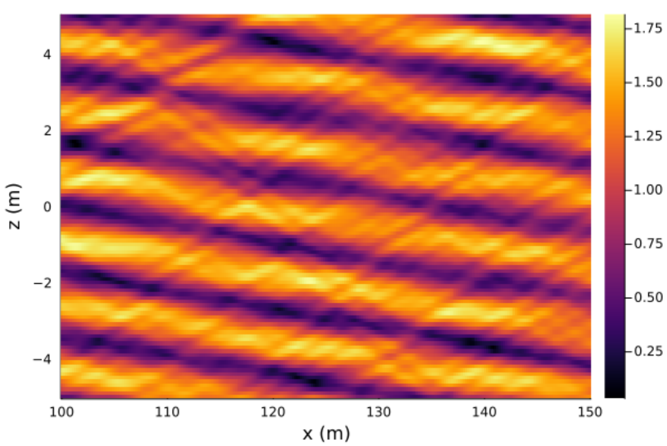

Figure 5: Zoomed in view of the region ahead of the AUV showing predicted field strength at 9 kHz.

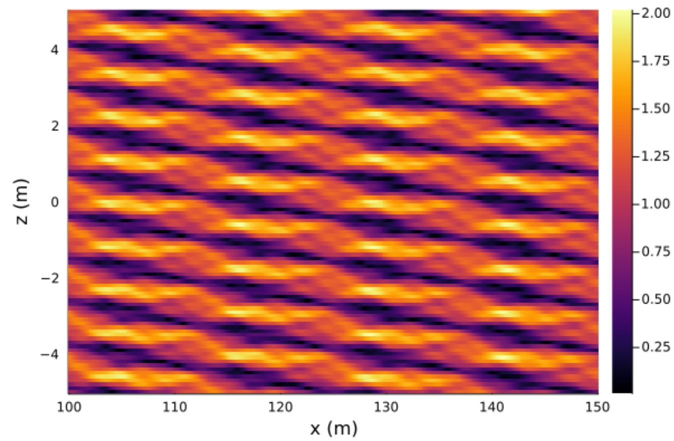

Figure 6: Zoomed in view of the region ahead of the AUV showing ground truth field strength at 14 kHz.

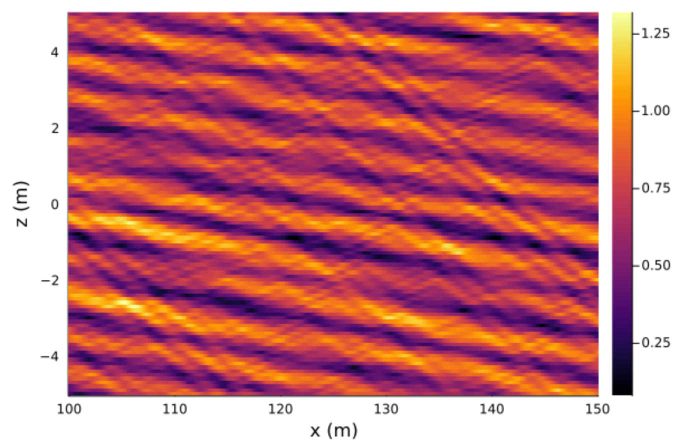

Figure 7: Zoomed in view of the region ahead of the AUV showing predicted field strength at 14 kHz.

not reproduced in the prediction, the overall pattern is accurately predicted and can be used effectively for AUV path planning or adaptation.

We visualize the recovered model parameters (matrix $[A_{mn}]$) in Figure 8. The row-spare structure of the matrix is clearly seen in terms of the horizontal bands at specific ray angles. The angles are consistent with the expected arrival angles based on our knowledge of the channel geometry.

V. Discussion

The problem at hand requires us to predict signal strength in a region around the AUV by observing the signal strength at a few locations. This is essentially a *regression* problem, but traditional machine learning algorithms such as deep neural networks (DNN) or Gaussian process regression (GPR) are unable to tackle it effectively. DNNs and GPR are extremely effective at interpolating data, but do not extrapolate well. Our problem requires extensive extrapolation. By incorporating knowledge of the physics of acoustic propagation in the structure of the algorithm, we are able to not only extrapolate well, but also learn from very little data [7]. DNNs are universal approximators, and are able to learn to approximate any function. Unlike DNNs, our algorithm is only able to learn to approximate a class of functions that are solutions to the acoustic wave equation. This aids in extrapolation, as the solution discovered is constrained by the acoustic wave equation in regions where no data points are available. The algorithm presented in this paper is closely related to that in [7], but uses a fixed set of direction of arrivals and combines data from multiple frequency bands to enable learning from data from a very small spatial region (as discussed next).

We observe the signal strength at only a few locations, but we collect data at each location at many frequencies. This is critically important for the algorithm to work. In most cases, the problem of determining the ray directions with only a few observations at a single frequency is vastly under-determined and therefore not uniquely solvable. However, by combining information over a range of frequencies and imposing a row-sparse structure on the

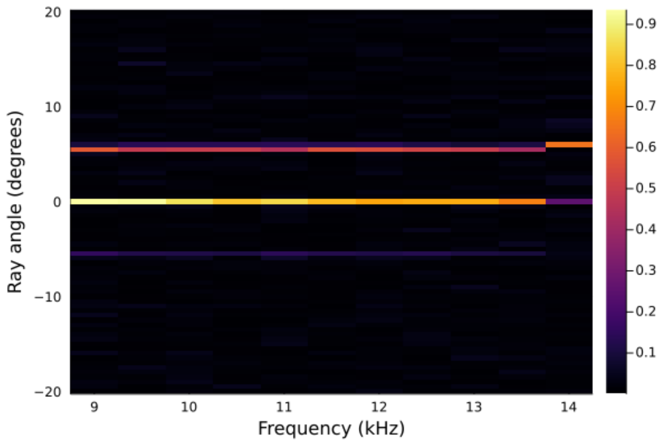

Figure 8: Learnt model parameters (matrix $[A_{mn}]$) showing discrete ray arrivals at specific angles. The row-sparse structure of the recovered matrix is clearly seen.

model parameters, we provide additional constraints that enable the problem to be solved.

The problem formulation and optimization approach described in Section II and Section III are quite general, and can be customized for the problem at hand. In the next few paragraphs, we illustrate how the approach can be modified to relax some of the assumptions made in this paper.

We assumed the region of interest to be small, and far away from the source. This enabled us to approximate the acoustic propagation in the region as plane waves, and to ignore spreading loss and absorption over the region of interest. The effects of these simplifying approximations are clearly seen by comparing Figure 2 with Figure 3. Firstly, we see a gradual decrease in signal strength in the ground truth field with increasing x, while we do not see this in the predicted field. Secondly, the wavefront curvature causes the angle of arrival of rays at the near end of the region of interest to differ from the ones at the far end by about 0.6°. This model mismatch leads to some errors in prediction of detailed structure. Should prediction in a larger region of interest, or operations at shorter ranges be desired, a more sophisticated model that includes wavefront curvature and spreading loss has to be adopted. This is easily done by modifying (1) and following the steps outlined in Section II and Section III, also at the cost of an increased number of model parameters.

We explicitly assumed an iso-velocity profile in our simulation study. By allowing sound speed c to be a function of depth and/or range, we can model channels with non iso-velocity sound speed profiles. This has to be coupled with adoption of a ray tracer that determines the refraction of the wavefront through the medium, instead of the plane wave model we used in (1). Since our optimization technique uses automatic differentiation to obtain gradient estimates, one must use a differentiable ray tracer (e.g. `RaySolver` [8], [9]) during training.

For low frequency long range acoustic applications, it is common to use the method of normal modes for propagation modeling. In modeling low frequency communication performance, one may substitute the ray model in (1) with the equivalent expression from a normal mode model [10]. The rest of the method described in Section II and Section III can then be applied to find the parameters of the modes propagating through the region of interest.

VI. Conclusions

We have shown that it is possible to predict acoustic signal quality in a region of interest around an AUV by simply observing the signal strength at a few locations along the AUV's trajectory, as long as the observations are made at many different frequencies. While we demonstrated this with a simple plane wave propagation model in an iso-velocity channel, the same approach can be adapted for use with spherical wave models, non iso-velocity channels, and even normal mode models. By incorporating knowledge of physics of acoustic propagation into the model, the AUV is able to make predictions with very little data, and with measurements made in a small subset of the region of interest. As new measurements flow in, the model may be updated through online training. The model output may then be used in path-planning algorithms for adaptive missions honoring acoustic navigation and communication guarantees.

References

[1] M. Chitre, "A holistic approach to underwater sensor network design," 2011.

[2] H. Schmidt and T. Schneider, "Acoustic communication and navigation in the new arctic — a model case for environmental adaptation," in *2016 IEEE third underwater communications and networking conference (UComms)*, 2016, pp. 1–4.

[3] M. B. Porter, "The BELLHOP manual and user's guide: Preliminary draft," *Heat, Light, and Sound Research, Inc., La Jolla, CA, USA, Tech. Rep*, vol. 260, 2011.

[4] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.

[5] A. Griewank, "On automatic differentiation," *Mathematical Programming: recent developments and applications*, vol. 6, no. 6, pp. 83–107, 1989.

[6] J. Potter, J. Alves, D. Green, G. Zappa, I. Nissen, and K. McCoy, "The JANUS underwater communications standard," in *2014 underwater communications and networking (UComms)*, 2014, pp. 1–4.

[7] K. Li and M. Chitre, "Data-aided underwater acoustic ray propagation modeling," *arXiv preprint arXiv:2205.06066*, 2022.

[8] M. Chitre, "RaySolver: A differentiable Gaussian beam tracer," *GitHub repository*. https://github.com/org-arl/AcousticRayTracers.jl; GitHub, 2022.

[9] M. Chitre, "Underwater acoustics in the age of differentiable and probabilistic programming," *UComms 2020 webinar*. Dec. 2020.

[10] F. B. Jensen, W. A. Kuperman, M. B. Porter, H. Schmidt, and A. Tolstoy, *Computational ocean acoustics*, vol. 794. Springer, 2011.