

DIRECT LOCALIZATION IN UNDERWATER ACOUSTICS VIA CONVOLUTIONAL NEURAL NETWORKS: A DATA-DRIVEN APPROACH

Amir Weiss, Toros Arikan and Gregory W. Wornell

Massachusetts Institute of Technology

ABSTRACT

Direct localization (DLOC) methods, which use the observed data to localize a source at an unknown position in a one-step procedure, generally outperform their indirect two-step counterparts (e.g., using time-difference of arrivals). However, underwater acoustic DLOC methods require prior knowledge of the environment, and are computationally costly, hence slow. We propose, what is to the best of our knowledge, the first data-driven DLOC method. Inspired by classical and contemporary optimal model-based DLOC solutions, and leveraging the capabilities of convolutional neural networks (CNNs), we devise a holistic CNN-based solution. Our method includes a specifically-tailored input structure, architecture, loss function, and a progressive training procedure, which are of independent interest in the broader context of machine learning. We demonstrate that our method outperforms attractive alternatives, and asymptotically matches the performance of an oracle optimal model-based solution.

Index Terms— Localization, underwater acoustics, deep neural networks, supervised learning, mean cyclic error.

1. INTRODUCTION

Underwater acoustic localization (UAL) is an important and challenging problem that arises in a wide range of emerging applications [1]. As such, it has been extensively addressed, and a host of methods—for the most part model-based—have been proposed for different operational environments [2].

Due to the unprecedented success of deep neural networks (DNNs) in various domains in recent years, the UAL problem had also been approached in a data-driven manner [3]. Motivated by the sensational achievements in classification using CNNs in the visual domain, some methods have been proposed that use such a classification-based methodology for UAL [4–7]. However, this type of framework may not be a well-suited approach to working with 1-dimensional (non-speech) acoustic signals, and a CNN trained for regression may be a more natural fit. Other regression-type solutions [8–10] were not necessarily tailored to the UAL problem, in

the sense that some architectural choices, such as the loss function or training procedure, were made according to standard, “off-the-shelf” ones, which are not necessarily suitable to this challenging domain.

A particularly attractive approach for UAL is *direct* localization (DLOC) [11, 12], which is theoretically superior (in terms of accuracy) to indirect methods, wherein first, some statistics of the data (e.g., time-differences of arrivals (TDOAs)) are estimated, and only then is localization performed based on the estimated statistics. DLOC methods require model characterization of the environment of operation, and typically have a high computational cost. To the best of our knowledge, a data-driven DLOC method, with a specifically-crafted DNN, has not been proposed to date.

In this work, we propose such a DLOC solution. By analyzing DLOC optimal model-based methods, we identify key ingredients in their structures, based on which we devise the input structure, architecture, loss function, and a progressive training procedure. Our main contributions, which are potentially of independent interest beyond the UAL context, are: (i) A novel CNN-based DLOC solution, which, to the best of our knowledge, is the first of its kind; (ii) An extended, model-based semi-blind localization (SBL, [13]) solution for UAL; (iii) A progressive training procedure for a model aiming for the joint estimation of several parameters; and (iv) A suitable loss function for learning to estimate a periodic parameter.

2. PROBLEM FORMULATION

Consider L spatially-diverse, time-synchronized receivers at known locations, each consisting of a single omni-directional hydrophone. Furthermore, consider the presence of an unknown signal, emitted from a source whose unknown position is denoted by the vector of coordinates $\mathbf{p} \in \mathbb{R}^{3 \times 1}$. We assume that the source is static during the observation time interval, and is located sufficiently far from all L receivers to permit a planar wavefront (far-field) approximation, which becomes reasonably accurate in shallow waters at high frequencies.

In such an environment, the underwater acoustic medium gives rise to a rich multipath channel, which can be approximately described using ray propagation [14]. In this model, the acoustic wave that is emitted from the source and measured at the the receiver is represented as a sum of (possibly

Code available at <https://www.weissamir.com/project/DLOC>.

This work was supported, in part, by ONR under Grant No. N00014-19-1-2665, and NSF under Grant No. CCF-1816209.

infinitely many) rays, each propagating according to physical laws governed by the environment's characteristics (e.g., salinity, temperature, bathymetry, etc.). We assume the existence of a parametric representation of the environment, and denote by \mathcal{E} the set of the environmental parameters. Importantly, each ray travels a certain distance from the source to the receiver, and arrives after a time-delay $\tau(\mathbf{p}, \mathcal{E})$, which depends on the source's location \mathbf{p} and the environment \mathcal{E} . Note that in the underwater acoustic environment, the relation between $\tau(\mathbf{p}, \mathcal{E})$ and the distance traveled by the corresponding ray is not necessarily simple, and in particular, linear, as is the case for isovelocity environments.

Formally, the sampled, baseband-converted signal from the ℓ -th receiver is given by

$$x_\ell[n] = \sum_{r=1}^R b_{r\ell} s_{r\ell}[n] + v_\ell[n] \triangleq \mathbf{s}_\ell^T[n] \mathbf{b}_\ell + v_\ell[n] \in \mathbb{C}, \quad (1)$$

$$n = 1, \dots, N, \quad \forall \ell \in \{1, \dots, L\},$$

where we have defined $\mathbf{s}_\ell[n] = [s_{1\ell}[n] \cdots s_{R\ell}[n]]^T \in \mathbb{C}^{R \times 1}$ and $\mathbf{b}_\ell = [b_{1\ell} \cdots b_{R\ell}]^T \in \mathbb{C}^{R \times 1}$, using the notations:

1. $b_{r\ell} \in \mathbb{C}$ as the unknown attenuation coefficient from the source to the ℓ -th sensor associated with the r -th signal component (line-of-sight (LOS) or other non-LOS (NLOS) reflections);
2. $s_{r\ell}[n] \triangleq s(t - \tau_{r\ell}(\mathbf{p}, \mathcal{E}))|_{t=nT_s} \in \mathbb{C}$ as the sampled r -th component of the unknown emitted signal's waveform at the ℓ -th sensor, where $s(t - \tau_{r\ell}(\mathbf{p}, \mathcal{E}))$ is the analog, continuous-time waveform delayed by $\tau_{r\ell}(\mathbf{p}, \mathcal{E})$, and T_s is the (known) sampling period; and
3. $v_\ell[n] \in \mathbb{C}$ as the additive noise at the ℓ -th receiver, representing the overall contributions of internal receiver noise and ocean ambient noise, modeled as a zero-mean random process with an unknown variance $\sigma_{v_\ell}^2$.

Applying the normalized DFT¹ to (1) yields the equivalent frequency-domain representation for all $\ell \in \{1, \dots, L\}$,

$$\bar{x}_\ell[k] = \sum_{r=1}^R b_{r\ell} \bar{s}[k] e^{-j\omega_k \tau_{r\ell}(\mathbf{p}, \mathcal{E})} + \bar{v}_\ell[k]$$

$$= \bar{s}[k] \cdot \underbrace{\mathbf{d}_\ell^H[k] \mathbf{b}_\ell}_{\bar{h}_\ell[k]} + \bar{v}_\ell[k] = \bar{s}[k] \cdot \bar{h}_\ell[k] + \bar{v}_\ell[k] \in \mathbb{C}, \quad (2)$$

where we have defined $\bar{h}_\ell[k] \triangleq \mathbf{d}_\ell^H[k] \mathbf{b}_\ell \in \mathbb{C}$, and

$$\mathbf{d}_\ell[k] \triangleq [e^{-j\omega_k \tau_{1\ell}(\mathbf{p}, \mathcal{E})} \cdots e^{-j\omega_k \tau_{R\ell}(\mathbf{p}, \mathcal{E})}]^H \in \mathbb{C}^{R \times 1},$$

$$\omega_k \triangleq \frac{2\pi(k-1)}{NT_s} \in \mathbb{R}_+, \quad k = 1, \dots, N.$$

As shorthand notation, we further define

$$\bar{\mathbf{x}}_\ell \triangleq [\bar{x}_\ell[1] \cdots \bar{x}_\ell[N]]^T, \quad \bar{\mathbf{s}} \triangleq [\bar{s}[1] \cdots \bar{s}[N]]^T,$$

$$\bar{\mathbf{v}}_\ell \triangleq [\bar{v}_\ell[1] \cdots \bar{v}_\ell[N]]^T, \quad \mathbf{D}_\ell \triangleq [\mathbf{d}_\ell[1] \cdots \mathbf{d}_\ell[N]]^T \in \mathbb{C}^{N \times R},$$

$$\bar{\mathbf{X}}_\ell \triangleq \text{Diag}(\bar{\mathbf{x}}_\ell), \quad \bar{\mathbf{S}} \triangleq \text{Diag}(\bar{\mathbf{s}}), \quad \mathbf{H}_\ell(\mathbf{p}, \mathcal{E}) \triangleq \text{Diag}(\mathbf{D}_\ell \mathbf{b}_\ell),$$

¹ \bar{z} denotes the normalized discrete Fourier transform (DFT) of z .

where $\text{Diag}(\cdot)$ forms a diagonal matrix from its vector argument. Note that $\mathbf{H}_\ell(\mathbf{p}, \mathcal{E})$ is generally a nonlinear function of the unknown source position \mathbf{p} and the environmental parameters \mathcal{E} , e.g., as seen from by the definition of $\mathbf{d}_\ell[k]$. With this notation, we may now write (2) compactly as

$$\bar{\mathbf{x}}_\ell = \mathbf{H}_\ell(\mathbf{p}, \mathcal{E}) \bar{\mathbf{s}} + \bar{\mathbf{v}}_\ell \in \mathbb{C}^{N \times 1}, \quad \forall \ell \in \{1, \dots, L\}. \quad (3)$$

Thus, the localization problem is formulated as follows:

Given the observations $\{\bar{\mathbf{x}}_\ell \in \mathbb{C}^{N \times 1}\}_{\ell=1}^L$ of the signal as per model (3), localize the source. Specifically, design an estimator $\hat{\mathbf{p}}$ of \mathbf{p} , so as to minimize $\mathbb{E} [\|\hat{\mathbf{p}} - \mathbf{p}\|_2^2]$.

In this paper, we assume that we do not have access to precise knowledge of the environment \mathcal{E} nor the underlying propagation model, from which $\mathbf{H}_\ell(\mathbf{p}, \mathcal{E})$ can be computed for any \mathbf{p} . Rather, we assume that a dataset of signals with labeled source positions $\mathcal{D}_{\text{train}}^{(J)} \triangleq \{(\bar{\mathbf{x}}_1^{(i)}, \dots, \bar{\mathbf{x}}_L^{(i)}, \mathbf{p}^{(i)})\}_{i=1}^J$ of size J is available. Such a dataset can be obtained by taking measurements using a deployed system, or by generating synthetic signals, possibly using simulators such as Bellhop [15], and combining them with recorded ocean ambient noise, which is readily accessible, e.g., [16, 17].

3. OPTIMAL MODEL-BASED SOLUTIONS

To motivate our proposed approach and gain theoretically-justified intuition into the implementation and architectural choices, it is instructive to first present two model-based solutions, which are optimal, each for its respective model, under slightly different assumptions.

3.1. Matched Field Processing

The key assumption of Matched Field Processing (MFP) [18] approaches is that, for a given hypothesized source location \mathbf{p} , the channel response $\mathbf{H}_\ell(\mathbf{p}, \mathcal{E})$ is *fully* predictable, known,² and computable. In many practical cases, this assumption is not a realistic one. For example, the attenuation coefficients associated with the bottom reflections generally depend on the angle of incidence, the sediment type at the area of operation (e.g., sand, silt, clay, etc.) and various other factors.

Nevertheless, assuming that \mathcal{E} is perfectly known, the MFP solution serves as a theoretical benchmark, since it is the maximum likelihood estimator (MLE) under the additive white Gaussian noise model, and hence is asymptotically efficient (in the statistical sense). It can be shown that after simplification (see, e.g., [13]), the MFP solution is given by

$$\hat{\mathbf{p}}_{\text{MFP}} = \underset{\mathbf{p} \in \mathbb{R}^{3 \times 1}}{\text{argmax}} \sum_{k=1}^N \frac{|\bar{\mathbf{x}}[k]^H \bar{\mathbf{h}}_k(\mathbf{p}, \mathcal{E})|^2}{\|\bar{\mathbf{h}}_k(\mathbf{p}, \mathcal{E})\|_2^2}, \quad (4)$$

where we have defined, for every k -th DFT component,

$$\bar{\mathbf{x}}[k] \triangleq [\bar{x}_1[k] \cdots \bar{x}_L[k]]^T \in \mathbb{C}^{L \times 1},$$

$$\bar{\mathbf{h}}_k(\mathbf{p}, \mathcal{E}) \triangleq [\bar{h}_1[k] \cdots \bar{h}_L[k]]^T \in \mathbb{C}^{L \times 1}.$$

²Otherwise, infeasible high-dimensional optimization is required.

As mentioned above, MFP requires the precise knowledge of the channel impulse response. When this knowledge is not available, (4) can be thought of as an oracle MLE.

3.2. Semi-Blind Localization

Unlike the MFP framework, the formulation of the semi-blind localization (SBL) problem [13] considers the channel attenuation coefficients $\{\mathbf{b}_\ell\}$ to be arbitrary (deterministic) unknown parameters, rather than as functions of \mathbf{p} and/or \mathcal{E} . Thus, by intentionally ignoring the functional structure of $\{\mathbf{b}_\ell\}$, additional (possibly redundant) nuisance parameters are introduced, and must be (implicitly) estimated in order to estimate \mathbf{p} . It turns out that this redundancy pays off in terms of robustness, and the SBL estimator is more resilient than other methods to some unmodeled environmental structures, such as unknown occluders [13, Sec. VI].

We now present an extended version of the SBL method, which had been proposed for a three-ray model that takes into account the LOS component, and the primary NLOS surface and bottom reflections. This extension is generally of independent interest for UAL, beyond the main contributions of this paper, to be presented in the next section.

For the general ray propagation model (1), possibly with more than three rays and not necessarily for an isovelocity environment, we define the (extended) SBL estimator as

$$\hat{\mathbf{p}}_{\text{SBL}} \triangleq \underset{\mathbf{p} \in \mathbb{R}^{3 \times 1}}{\operatorname{argmin}} \left\{ \min_{\substack{\bar{\mathbf{s}} \in \mathcal{S}_N \\ \mathbf{B} \in \mathbb{C}^{R \times L}}} \sum_{\ell=1}^L \|\bar{\mathbf{x}}_\ell - \mathbf{H}_\ell(\mathbf{p}, \mathcal{E}) \bar{\mathbf{s}}\|_2^2 \right\}, \quad (5)$$

where $\bar{\mathbf{s}} \in \mathcal{S}_N \triangleq \{\mathbf{z} \in \mathbb{C}^{N \times 1} : \|\mathbf{z}\|_2 = 1\}$ is used in (5) without loss of generality, since both $\bar{\mathbf{s}}$ and \mathbf{b}_ℓ are unknown,³ and $\mathbf{B} \triangleq [\mathbf{b}_1 \cdots \mathbf{b}_L] \in \mathbb{C}^{R \times L}$. We have the following result:

Proposition 1. *Assume $\{\mathbf{D}_\ell\}_{\ell=1}^L$ are full column rank. For any potential position \mathbf{p} , define the data-dependent matrix,*

$$\mathbf{Q}(\mathbf{p}, \mathcal{E}) \triangleq \sum_{\ell=1}^L \bar{\mathbf{X}}_\ell \mathbf{D}_\ell^* \left(\mathbf{D}_\ell^T \mathbf{D}_\ell^* \right)^{-1} \left(\bar{\mathbf{X}}_\ell \mathbf{D}_\ell^* \right)^H \in \mathbb{C}^{N \times N}. \quad (6)$$

Then, for any spectrally flat waveform, i.e., $|\bar{s}[k]| = \rho$, for all $k \in \{1, \dots, N\}$ and some $\rho > 0$, the SBL estimator is

$$\hat{\mathbf{p}}_{\text{SBL}} = \underset{\mathbf{p} \in \mathbb{R}^{3 \times 1}}{\operatorname{argmax}} \lambda_{\max}(\mathbf{Q}(\mathbf{p}, \mathcal{E})), \quad (7)$$

where $\lambda_{\max}(\mathbf{A})$ denotes the largest eigenvalue of $\mathbf{A} \in \mathbb{C}^{N \times N}$.

Due to space limitations, the proof is given in [19].

Although (7) holds exactly for spectrally flat signals, it is also a reliable solution for general waveforms [13, Prop. 3]. However, the complexity of (7) is $\mathcal{O}(RN)$, and it can be implemented only when the time-delays $\{\tau_{r,\ell}(\mathbf{p}, \mathcal{E})\}$ can be computed for any \mathbf{p} . In practice, this is not always possible, since in some cases \mathcal{E} is not fully known. Moreover, even

³Giving rise to an inherent scaling ambiguity: $\bar{\mathbf{s}} \cdot \mathbf{b}_\ell = (\alpha \bar{\mathbf{s}}) \cdot (\frac{1}{\alpha} \mathbf{b}_\ell)$.

when \mathcal{E} is known, the computation can be nontrivial, and has to be executed for *any* tested hypothesized position \mathbf{p} .

Before we present our proposed solution, we point to the following observations regarding the MFP and SBL solutions:

- *Sufficient statistic:* Although MFP and SBL differ in their simplified forms (4) and (7), respectively, both require *only* the (same) correlation functions of the observed data. This can be observed by rewriting the numerator of (4), and writing explicitly the statistically-equivalent form of (6)— $\tilde{\mathbf{Q}}(\mathbf{p}, \mathcal{E})$ (given in [19]), due to $\lambda_{\max}(\mathbf{Q}(\mathbf{p}, \mathcal{E})) = \lambda_{\max}(\tilde{\mathbf{Q}}(\mathbf{p}, \mathcal{E}))$ —as

$$\begin{aligned} & \left| \bar{\mathbf{x}}[k]^H \bar{\mathbf{h}}_k(\mathbf{p}, \mathcal{E}) \right|^2 = \bar{\mathbf{h}}_k(\mathbf{p}, \mathcal{E})^H \underbrace{\bar{\mathbf{x}}[k] \bar{\mathbf{x}}[k]^H}_{\tilde{\mathbf{Q}}(\mathbf{p}, \mathcal{E})} \bar{\mathbf{h}}_k(\mathbf{p}, \mathcal{E}), \\ \tilde{\mathbf{Q}}(\mathbf{p}, \mathcal{E}) &= \begin{bmatrix} \mathbf{G}_1^H \underbrace{\bar{\mathbf{X}}_1^H \bar{\mathbf{X}}_1}_{\tilde{\mathbf{Q}}_1} \mathbf{G}_1 & \cdots & \mathbf{G}_1^H \underbrace{\bar{\mathbf{X}}_1^H \bar{\mathbf{X}}_L}_{\tilde{\mathbf{Q}}_{1L}} \mathbf{G}_L \\ \vdots & \ddots & \vdots \\ \mathbf{G}_L^H \underbrace{\bar{\mathbf{X}}_L^H \bar{\mathbf{X}}_1}_{\tilde{\mathbf{Q}}_{L1}} \mathbf{G}_1 & \cdots & \mathbf{G}_L^H \underbrace{\bar{\mathbf{X}}_L^H \bar{\mathbf{X}}_L}_{\tilde{\mathbf{Q}}_L} \mathbf{G}_L \end{bmatrix}, \\ & \implies \mathbf{e}_{\ell_1}^T \bar{\mathbf{x}}[k] \bar{\mathbf{x}}[k]^H \mathbf{e}_{\ell_2} = \mathbf{e}_k^T \bar{\mathbf{X}}_{\ell_1}^H \bar{\mathbf{X}}_{\ell_2} \mathbf{e}_k, \end{aligned}$$

where $\{\mathbf{G}_\ell \in \mathbb{C}^{N \times R}\}_{\ell=1}^L$, defined in [19], are data-independent, and \mathbf{e}_n is the n -th standard basis vector. It follows that MFP and SBL share *the same* sufficient statistic, only arranged differently. Recall that an element-wise product of a DFT vector with the conjugate of another is equivalent to correlating the corresponding time-domain vectors. Hence, the sufficient statistics for the optimal model-based solutions are the second-order statistics (SOS), i.e., the empirical correlation functions of all L observed signals $\{x_\ell[n]\}$.

- *Joint optimization of orthogonal coordinates:* Although range, inclination and azimuth are algebraically orthogonal, and therefore in some sense decoupled, they are nevertheless *statistically coupled*. Thus, DLOC, i.e., the estimation of \mathbf{p} directly from the data can (and does) lead to improved accuracy [11].

4. A DATA-DRIVEN CNN-BASED APPROACH

Bearing in mind the practical challenges faced by UAL model-based approaches, namely the assumption that we have prior knowledge of the environment, and taking into account the prohibitive computational cost of a 3-dimensional optimization (search for the emitter) in the volume of interest, we propose a different approach. As we show in Section 5, our method is able to attain the performance of an *oracle* optimal model-based solution, which has access to perfect knowledge of the impulse response. Furthermore, our solution is exceptionally computationally efficient compared to state-of-the-art alternatives, which typically require some grid-search. Thanks to the data-driven nature of our approach, we bypass the need to analytically describe the physical laws that govern the true underlying propagation model, which are manifested in the resulting impulse responses $\{\mathbf{H}_\ell(\mathbf{p}, \mathcal{E})\}$.

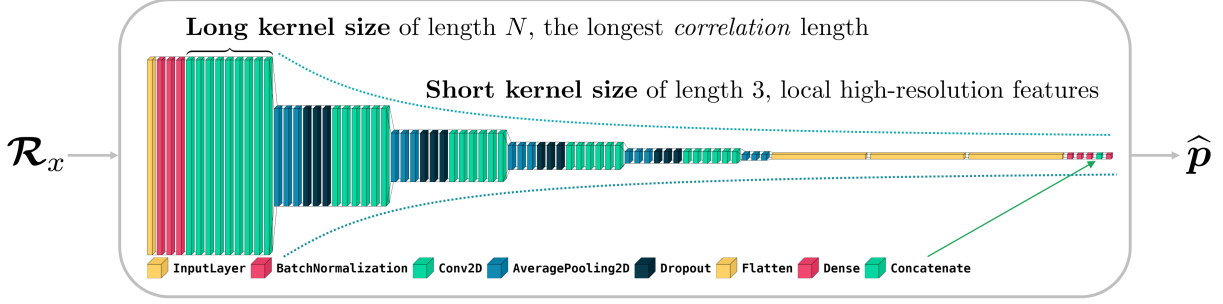


Fig. 1: Architecture of the proposed DNN for data-driven DLOC. The model is comprised of three sub-models, which are initially trained individually to estimate range, azimuth and inclination. Each sub-model contains 5 main building blocks, where each block contains 2D-convolutional layers, followed by 2D average pooling. Dropouts are used in the first 4 blocks. The 3 sub-models are connected with a dense layer. For more implementation details, see [19].

We propose to use a CNN-based architecture, that is constructed and trained progressively. Fig. 1 illustrates the final structure, where some of the implementation details are provided in the caption; see [19] for the full details. The proposed input structure, chosen architecture and loss function, are designed according to properties shared by the optimal model-based solutions (4), (7), as described in detail next.

4.1. Input: The SOS Tensor

As explained in the Section 3, the MFP and SBL optimal model-based solutions use (only) all the empirical auto- and cross-correlation functions of the observed signals from all L receivers. Therefore, in order to provide a more succinct version of the data, which is in agreement with what the optimal model-based solutions require, we define the SOS tensor as⁴

$$\begin{aligned} \mathcal{R}_x[l_{\text{corr}}(\ell_1, \ell_2), m, 1] &\triangleq \Re \left\{ \frac{1}{N} \sum_n x_{\ell_1}[n+m] x_{\ell_2}^*[n] \right\}, \\ \mathcal{R}_x[l_{\text{corr}}(\ell_1, \ell_2), m, 2] &\triangleq \Im \left\{ \frac{1}{N} \sum_n x_{\ell_1}[n+m] x_{\ell_2}^*[n] \right\}, \\ \implies \mathcal{R}_x &\in \mathbb{R}^{L_{\text{corr}} \times (2N-1) \times 2}, \quad L_{\text{corr}} \triangleq \frac{(L+1)L}{2}, \end{aligned}$$

and use it as the input to our CNN. Here, $l_{\text{corr}}(\ell_1, \ell_2)$ is a bijective function that maps all the unique pairs of receiver indices to a unique index in $\{1, \dots, L_{\text{corr}}\}$ (corresponding to all the auto- and cross-correlation functions, without repetition).

With this input, it is guaranteed that the model has the data-dependent (but not necessarily prior) information required to (at least) match the performance of the optimal model-based solutions presented in Section 3. This inspires the choice of 2D-convolutional layers, as with this choice, the CNN can jointly process the correlation functions per-layer.

4.2. Loss Function with Spherical Coordinates

As the objective function, we choose to minimize

$$\mathcal{L}_{\text{sph}}(\hat{\mathbf{p}}(\mathbf{w}), \mathbf{p}) \triangleq \|\hat{\mathbf{p}}(\mathbf{w}) - \mathbf{p}\|_2^2 = \hat{r}^2(\mathbf{w}) - r^2 - 2\hat{r}(\mathbf{w})r \cdot \left[\sin(\hat{\varphi}(\mathbf{w})) \sin(\varphi) \cos(\hat{\theta}(\mathbf{w}) - \theta) + \cos(\hat{\varphi}(\mathbf{w})) \cos(\varphi) \right], \quad (8)$$

⁴The signals in the sum are zero-padded wherever necessary.

namely, the squared Euclidean norm of the estimation error in *spherical coordinates*, where r, θ, φ are the range, azimuth and inclination, respectively, associated with \mathbf{p} (similarly for $\hat{r}(\mathbf{w}), \hat{\theta}(\mathbf{w}), \hat{\varphi}(\mathbf{w})$ and $\hat{\mathbf{p}}(\mathbf{w})$), and \mathbf{w} is the vector of all the weights of the DNN to be trained. Several considerations lead to this particular choice.

First, although it is analytically convenient to decouple different coordinates, when training a NN we prefer that the gradient of the loss function with respect to one estimand (e.g., range) would depend on another (e.g., azimuth), so as to *exploit statistical dependencies*. Indeed, it can be seen that the gradient of (8) with respect to $\hat{r}(\mathbf{w})$ is a function of $\hat{\theta}(\mathbf{w})$ and $\hat{\varphi}(\mathbf{w})$. This is not the case for Cartesian coordinates.

Second, in spherical coordinates, two coordinate axes are angles, which are periodic by definition. This way, we create a loss function with *infinitely many globally optimal* points, since the metric (8) is invariant to additions of 2π (and hence $2\pi\mathbb{Z}$) to $\hat{\theta}(\mathbf{w})$ and $\hat{\varphi}(\mathbf{w})$. This is also not true for the common, but not necessarily justified, choice of Cartesian coordinates.

While the choices above are inspired by the optimal model-based solutions' structure, there is still a need to account for the optimization procedure for computing an estimate given the data. Based on our empirical observations, training our DNN to directly output an estimate of \mathbf{p} without an educated initialization does not seem to work. This possibly results from a highly non-smooth "landscape" induced by the objective (8) due to the NLOS components, giving rise to many spurious source locations. We address this difficulty using a progressive training procedure.

4.3. A Two-Step Progressive Training Procedure

Rather than starting with one model that outputs the location of the source \mathbf{p} , we start by training three independent models, for estimating *individually* the range, inclination and azimuth. Although these models may yield sub-optimal estimates of the different coordinates, these are much simpler tasks (known as range and direction-of-arrival estimation). The idea behind progressive training is the following: After each model is trained individually, the models are combined (using the concatenation operator), and then serve as the initial weights for the training process of the DLOC model. Ob-

	x [m]	y [m]	z [m]
Receiver 1	150	-250	10
Receiver 2	50	-250	15
Receiver 3	-50	-250	20
Receiver 4	-150	-250	25

Table 1: Positions of the four receivers in Cartesian coordinates.

serve that in Fig. 1, except for the input and output layers, all layers appear in multiples of three, which is due to the concatenation of the three sub-models trained to individually estimate r , θ , and φ .

For the range model, we choose $\|r - \hat{r}(\mathbf{w}_r)\|_2^2$ as the loss function, where \mathbf{w}_r denotes the vector of weights of the model. As for the azimuth θ , we would like to exploit the inherent periodicity to create multiple globally optimal solutions, as explained in Section 4.2. A suitable loss for the estimation of periodic parameters is the mean cyclic error (MCE),

$$\text{MCE}(\hat{\theta}, \theta) \triangleq 2 - 2\mathbb{E} \left[\cos(\hat{\theta} - \theta) \right] \in \mathbb{R}_+. \quad (9)$$

Therefore, we choose the empirical MCE loss, defined as

$$\text{EMCE}(\hat{\theta}(\mathbf{w}_\theta), \theta) \triangleq 2 - 2 \cos(\hat{\theta}(\mathbf{w}_\theta) - \theta),$$

where \mathbf{w}_θ denotes the vector of weights of the model. We use the same loss for the inclination angle, scaled by a factor of 2, since its range is $[0, \pi]$. To summarize the training process:

Phase 1: Train the *individual* sub-models for r, θ, φ .

Phase 2: Train the model for DLOC, i.e., \mathbf{p} , where the three concatenated “core branches” are initialized by the weights of the previously trained sub-models.

5. SIMULATION RESULTS

We focus on a special case of the general ray model—the isovelocity three-ray model [13, Fig. 1], which considers the three primary, typically most dominant, components of the impulse response: the direct-path LOS and the NLOS surface and bottom reflections. This model yields a method that already allows for successful localization in nontrivial settings (e.g., complete loss of LOS [13]). Additionally, it enables a comparison to the oracle MFP, so as to evaluate how close the proposed solution can get to the oracle’s performance, which constitutes a lower bound. By this, we are able to assess whether our model has the required capacity to learn the environment, the receivers geometry, and the optimal estimation rule (4). We note, however, that our approach can cope with the more complex propagation models of other environments.

We consider a scenario with $L = 4$ receivers, in a volume with bottom depth $h = 50$ m. The locations of the receivers are given in Table 1. The attenuation coefficients were drawn independently from the complex normal (CN) distribution, such that $\mathbb{E}[|b_{rl}|^2] = 1$, with variance 0.1^2 . The speed of sound was set to $c = 1500$ m/s, and we used $N = 100$.

We trained our model as described in Section 4 using “hybrid” received signals $\{\bar{\mathbf{x}}_1^{(i)}, \dots, \bar{\mathbf{x}}_L^{(i)}\}$ that were generated

according to (2) with $R = 3$. The source’s baseband-converted waveform was drawn from the standard CN distribution, i.e., $\bar{s}[k] \stackrel{\text{iid}}{\sim} \mathcal{CN}(0, 1)$, and the noise realizations were taken from recordings of real ocean ambient noise from the KAM11 experiment [16]. Thus, the model can potentially learn the underlying statistics of real ocean noise, which is generally not temporally white or Gaussian [20]. The size of the dataset $\mathcal{D}_{\text{train}}^{(J)}$ was $J = 10^4$ per SNR level, with twenty equidistant levels on $[-10, 30]$ dB, where we set $\forall \ell : \sigma_{v_\ell}^2 = \sigma_v^2$, and define $\text{SNR} \triangleq 1/\sigma_v^2$. For each example, the source position was drawn independently in Cartesian coordinates according to

$$\mathbf{p}^{(i)} = [x^{(i)} \ y^{(i)} \ z^{(i)}]^T : \begin{cases} x^{(i)} \sim \text{Unif}(-150, 150) [m], \\ y^{(i)} \sim \text{Unif}(-100, 0) [m], \\ z^{(i)} \sim \text{Unif}(5, 45) [m]. \end{cases}$$

More details regarding the training process (optimizer, learning rate, etc.) are given in the supplementary materials [19].

We first test our model with additive white CN noise, and compare its performance with three methods in the literature: the “generalized cross-correlation with phase transform” (GCC-PHAT) [21] TDOA method, which is recognized as being resilient to multipath; SBL; and an oracle MFP, which uses the *exact* values of $\{\mathbf{b}_\ell\}$ for each realization. Fig. 2 shows the root mean squared error (RMSE), $\sqrt{\mathbb{E}[\|\hat{\mathbf{p}} - \mathbf{p}\|_2^2]}$, vs. the SNR for each method.⁵ As expected, the oracle MFP, which is also the oracle MLE in this case, dominates all methods. However, from 0 dB onward our CNN-based DLOC model is the best non-oracle method, and at high SNR approximately *matches the oracle performance*. We conclude that the model has the capacity to learn the statistics of the channel and the optimal estimation rule at high SNR.

We repeat the experiment, for synthetic CN and real ocean ambient noise from the KAM11 experiment [16], but now for a “dynamic” propagation model, in which the surface height is slightly perturbed, giving rise to perturbations in the surface-reflection associated time-delays, $\{\tau_{2\ell}(\mathbf{p}, \mathcal{E})\}$. To this end, in each realization we add zero-mean Gaussian perturbations to the time-delays $\{\tau_{2\ell}(\mathbf{p}, \mathcal{E})\}$, with a standard deviation of 1% of the average distance from the source to the receivers, divided by c . Figs. 3 and 4 reflect that due to these perturbations, the performances of the model-based methods already deteriorate considerably. In contrast, the proposed data-driven method, which was trained with data reflecting this phenomenon, adjusts and thereby attains the highest accuracy—from 0 dB onwards for (synthetic) CN noise, and uniformly for real ocean ambient noise. These results demonstrate the potential to implement data-driven DLOC methods that go beyond what is offered by model-based methods derived for relatively simple propagation models. Further, such methods can learn the true underlying statistics of ocean ambient noise, and exploit them for higher localization accuracy.

⁵Results are based on 10^3 and 10^4 trials for the model-based methods and CNN-DLOC, respectively, due to their different computational complexity.

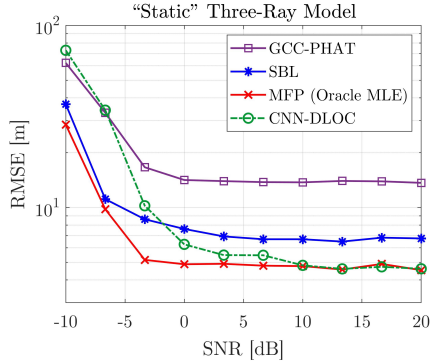


Fig. 2: RMSE vs. SNR for the “static” model, with a flat and stationary surface. Already from 0 dB, CNN-DLOC dominates the non-oracle methods.




Fig. 3: RMSE vs. SNR for the “dynamic” model, with a flat, randomly shifted surface. Here, MFP is not the MLE, hence referred to as “semi-oracle”.

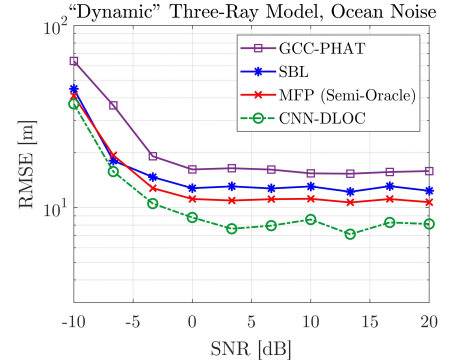


Fig. 4: RMSE vs. SNR for the “dynamic” model with ocean ambient noise. Here, CNN-DLOC dominates all methods, even at low SNRs.

6. CONCLUDING REMARKS

We present a computationally efficient, data-driven direct localization approach, that can asymptotically match optimal model-based solutions, which are hard to implement in practice. The proposed holistic solution includes a specifically tailored architecture and a progressive training procedure. Still, many interesting aspects for potential research avenues remain to be explored. Examples include extensions to more complex environments, tracking, coarsely-quantized signals, scaling of the architecture (width and number of layers) with the volume of operation and more.

7. REFERENCES

- [1] J. Waterston, J. Rhea, S. Peterson, L. Bolick, J. Ayers, and J. Ellen, “Ocean of things: Affordable maritime sensors with scalable analysis,” in *Proc. OCEANS Conf.*, 2019, pp. 1–6.
- [2] H.-P. Tan, R. Diamant, W. K. G. Seah, and M. Waldmeyer, “A survey of techniques and challenges in underwater localization,” *Ocean Eng.*, vol. 38, no. 14-15, pp. 1663–1676, 2011.
- [3] M. 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.*, vol. 146, no. 5, pp. 3590–3628, 2019.
- [4] H. Niu, Z. Gong, E. Ozanich, P. Gerstoft, H. Wang, and Z. Li, “Deep-learning source localization using multi-frequency magnitude-only data,” *J. Acoust. Soc. Am.*, vol. 146, no. 1, pp. 211–222, 2019.
- [5] A. Testolin and R. Diamant, “Underwater acoustic detection and localization with a convolutional denoising autoencoder,” in *Proc. of CAMSAP*, 2019, pp. 281–285.
- [6] Z. Gong, C. Li, and F. Jiang, “A machine learning-based approach for auto-detection and localization of targets in underwater acoustic array networks,” *IEEE Trans. Veh. Technol.*, vol. 69, no. 12, pp. 15 857–15 866, 2020.
- [7] R. Chen and H. Schmidt, “Model-based convolutional neural network approach to underwater source-range estimation,” *J. Acoust. Soc. Am.*, vol. 149, no. 1, pp. 405–420, 2021.
- [8] R. Lefort, G. Real, and A. Drémeau, “Direct regressions for underwater acoustic source localization in fluctuating oceans,” *Applied Acoustics*, vol. 116, pp. 303–310, 2017.
- [9] Y. Wang and H. Peng, “Underwater acoustic source localization using generalized regression neural network,” *J. Acoust. Soc. Am.*, vol. 143, no. 4, pp. 2321–2331, 2018.
- [10] D. Qin, J. Tang, and Z. Yan, “Underwater acoustic source localization using LSTM neural network,” in *2020 39th Chinese Control Conference (CCC)*, 2020, pp. 7452–7457.
- [11] A. J. Weiss, “Direct position determination of narrowband radio frequency transmitters,” *IEEE Signal Process. Lett.*, vol. 11, no. 5, pp. 513–516, 2004.
- [12] L. Wang, Y. Yang, and X. Liu, “A direct position determination approach for underwater acoustic sensor networks,” *IEEE Trans. Veh. Technol.*, vol. 69, no. 11, pp. 13 033–13 044, 2020.
- [13] A. Weiss, T. Arikan, H. Vishnu, G. B. Deane, A. C. Singer, and G. W. Wornell, “A semi-blind method for localization of underwater acoustic sources,” *IEEE Trans. Signal Process.*, vol. 70, pp. 3090–3106, 2022.
- [14] P. C. Etter, *Underwater acoustic modeling and simulation*. CRC Press, 2018.
- [15] 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.
- [16] W. Hodgkiss and J. Preisig, “Kauai ACOMMS MURI 2011 (KAM11) experiment,” in *Proc. of Eur. Conf. Underwater Acoust.*, 2012, pp. 993–1000.
- [17] S. Menze, D. P. Zitterbart, I. van Opzeeland, and O. Boebel, “The influence of sea ice, wind speed and marine mammals on southern ocean ambient sound,” *Royal Society Open Science*, vol. 4, no. 1, p. 160370, 2017.
- [18] A. B. Baggeroer, W. A. Kuperman, and P. N. Mikhalevsky, “An overview of matched field methods in ocean acoustics,” *IEEE J. Ocean. Eng.*, vol. 18, no. 4, pp. 401–424, 1993.
- [19] A. Weiss, T. Arikan, and G. W. Wornell, “Blind direct localization via convolutional neural networks: A data-driven approach, supplementary materials,” <https://www.weissamir.com/project/DLOC>, 2022.
- [20] M. Stojanovic and J. Preisig, “Underwater acoustic communication channels: Propagation models and statistical characterization,” *IEEE Commun. Mag.*, vol. 47, no. 1, pp. 84–89, 2009.
- [21] C. Zhang, D. Florêncio, and Z. Zhang, “Why does PHAT work well in low noise, reverberative environments?” in *Proc. of ICASSP*, 2008, pp. 2565–2568.