

Direction-Aware Neural Acoustic Fields for Few-Shot Interpolation of Ambisonic Impulse Responses

Ick, Christopher; Wichern, Gordon; Masuyama, Yoshiki; Germain, François G; Le Roux, Jonathan

TR2025-120 August 20, 2025

Abstract

The characteristics of a sound field are intrinsically linked to the geometric and spatial properties of the environment surrounding a sound source and a listener. The physics of sound propagation is captured in a time-domain signal known as a room impulse response (RIR). Prior work using neural fields (NFs) has allowed learning spatially-continuous representations of RIRs from finite RIR measurements. However, previous NF-based methods have focused on monaural omnidirectional or at most binaural listeners, which does not precisely capture the directional characteristics of a real sound field at a single point. We propose a direction-aware neural field (DANF) that more explicitly incorporates the directional information by Ambisonic-format RIRs. While DANF inherently captures spatial relations between sources and listeners, we further propose a direction-aware loss. In addition, we investigate DANF's ability to adapt to new rooms in various ways including low-rank adaptation. Index Terms: spatial audio, neural acoustic field, room impulse response, Ambisonics

Interspeech 2025

Direction-Aware Neural Acoustic Fields for Few-Shot Interpolation of Ambisonic Impulse Responses

Christopher Ick^{1,2}, Gordon Wichern¹, Yoshiki Masuyama¹, François Germain¹, Jonathan Le Roux¹

¹Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA, USA

²Music and Audio Research Laboratory, New York University, Brooklyn, NY, USA

chris.ick@nyu.edu {wichern,masuyama,germain,leroux}@merl.com

Abstract

The characteristics of a sound field are intrinsically linked to the geometric and spatial properties of the environment surrounding a sound source and a listener. The physics of sound propagation is captured in a time-domain signal known as a room impulse response (RIR). Prior work using neural fields (NFs) has allowed learning spatially-continuous representations of RIRs from finite RIR measurements. However, previous NF-based methods have focused on monaural omnidirectional or at most binaural listeners, which does not precisely capture the directional characteristics of a real sound field at a single point. We propose a direction-aware neural field (DANF) that more explicitly incorporates the directional information by Ambisonic-format RIRs. While DANF inherently captures spatial relations between sources and listeners, we further propose a direction-aware loss. In addition, we investigate DANF’s ability to adapt to new rooms in various ways including low-rank adaptation.

Index Terms: spatial audio, neural acoustic field, room impulse response, Ambisonics

1. Introduction

Inferring how an acoustic signal emitted from a sound source is received by a listener involves modeling the physically complex propagation of that signal through their shared environment [1]. Even when the source and listener are in direct line-of-sight, reflections of the sound signal at various surfaces and objects add significant contributions to what the listener hears. Precisely capturing the acoustic properties of this process, which is typically done in the form of a room impulse response (RIR) to be convolved with a dry signal at the source, is crucial for accurate sound rendering, which in turn enables immersive experiences in media, particularly in virtual and augmented reality [2, 3].

Because the process of recording RIRs can be time-consuming and impractical in certain settings, there exists a broad set of methods for simulating RIRs, typically either by geometric acoustics [4–6] or by wave-based methods [7–10]. Each method has different benefits depending on the application, but both require an accurate geometric and material specification of a given scene to produce accurate RIRs.

The problem of interpolating RIRs from a finite number of measurements has been widely studied in a variety of settings, including wave-based models [11, 12] and geometric models [13, 14]. In the particular case of learning spatially continuous RIRs for a scene with varying source/listener positions, inspiration has been taken from novel-view synthesis from the vision community [15] to develop continuous representations of acoustic scenes using neural fields (NFs) [16–20] and physics-informed neural networks [21–23]. These methods draw from prior work in omnidirectional single-channel IR estimation, in

which room acoustic metrics (such as T60, C50, and EDT) play a large role in perceptual accuracy, and have mainly evaluated binaural RIRs by averaging metrics across each channel. However, these approaches fail to account for inter-channel directional information, which are crucial not only for delivering immersive audio experiences in mixed reality applications [2, 3] but also for enhancing downstream tasks such as event localization [24], especially when generalizing to multi-channel formats beyond binaural audio.

In this work, we propose an NF that represents spatial RIRs, allowing for direction-aware modeling of the room acoustics over continuous source and listener locations. The proposed NF, dubbed direction-aware neural field (DANF), incorporates Ambisonic RIRs and captures the directivity of sound as illustrated in Fig. 1. Furthermore, to enhance its direction awareness, we propose a loss function defined on an intensity vector [24, 25]. Our experiments show that the intensity-vector loss can improve not only the sound source direction accuracy but also other single-channel RIR metrics. In addition, we explore the adaptation of DANF to novel rooms from few measurements and demonstrate its few-shot capability.

2. Ambisonics Room Impulse Responses

The transfer function characterizing the physical process of sound propagation is commonly referred to as RIR. It is typically defined for given locations of the source and listener in a particular acoustic environment [1]. An RIR can be treated as a linear filter: by convolving it with an anechoic sound source, we can mimic the acoustic properties of the source-listener-environment configuration (referred to as a “scene”) in the rendered sound. RIRs are most commonly recorded with omnidirectional microphones, which, by averaging out directional information with equal weighting, remove any spatial cues in their measurements.

A popular approach to capturing a sound field at a point is to encode it into an Ambisonics format [1]. Ambisonics captures the local pressure gradient vector by projecting it onto a spherical harmonic basis up to a given order n , where higher orders correspond to harmonics with increased spatial resolution. The number of harmonics, and corresponding encoding channels, needed to go up to order n is $(n + 1)^2$, meaning the popular first-order Ambisonics (FOA) encoding features 4 channels: one omnidirectional channel w and 3 figure-of-8 channels aligned with any 3 Cartesian directions, typically referred to as x , y , and z , as shown in top parts of Fig. 1.

Each component is treated as a separate channel in a 4-channel FOA signal. For a given omnidirectional source, a listener records FOA signal $u(t) = [w(t), x(t), y(t), z(t)]^T$, where $t = 1, \dots, T$ is the time sample index, and T is the

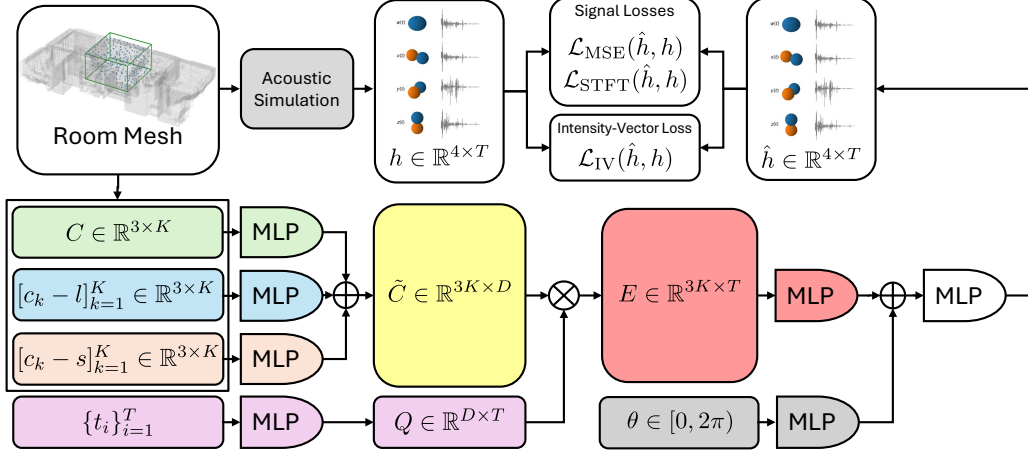


Figure 1: Overview of the proposed DANF framework. The input is the environmental context in the form of K bounce points $c_k \in \mathbb{R}^3$, as well as the source and listener locations $s, l \in \mathbb{R}^3$, all of which are concatenated into spatial feature \tilde{C} (denoted by operator \oplus) used to create a spatial-temporal encoding $E \in \mathbb{R}^{3K \times T}$. For the impulse generation module, the spatio-temporal embedding E and the orientation of the listener $\theta \in [0, 2\pi)$ are decoded into an Ambisonics impulse response $h \in \mathbb{R}^{4 \times T}$.

number of time samples. To gain insight on the direction-of-arrival (DoA) of sources in a scene, which tends to be associated with the azimuth and elevation of maximum intensity of the FOA signal $u \in \mathbb{R}^{4 \times T}$, an intensity vector is computed from the complex time-frequency representation of each of the FOA components [24, 25]. For example obtained using the short-time Fourier transform (STFT) $U \in \mathbb{C}^{4 \times M \times F}$, where M and F denote the number of time frames and frequency bins, respectively:

$$\text{IV}(U)(m, f) = \begin{bmatrix} \text{Re}(W(m, f)^* X(m, f)) \\ \text{Re}(W(m, f)^* Y(m, f)) \\ \text{Re}(W(m, f)^* Z(m, f)) \end{bmatrix}, \quad (1)$$

where m denotes the time frame index, f the frequency bin index, $*$ the complex conjugate, and Re the real component of a complex number.

3. Methods

3.1. Neural Acoustic Fields

Prior approaches in neural acoustic fields establish the problem of estimating an RIR $h \in \mathbb{R}^{N \times T}$ [17–19], where N denotes the number of channels. Prior work has focused on the monaural or binaural cases, and while some have allowed for different listener orientations $\theta \in [0, 2\pi)$, their training methods and evaluations did not consider inter-channel directional features that would correspond to changes in listener orientation for fixed source/listener positions. A neural field can be used to estimate the binaural RIR $h \in \mathbb{R}^{2 \times T}$ as a function of the listener position $l \in \mathbb{R}^3$, the source position $s \in \mathbb{R}^3$, and some fixed environmental context C :

$$h = \text{NF}(s, l, \theta, C). \quad (2)$$

In many neural acoustic fields, C is implicit, as a separate NF will be trained for each environment. This environmental context can also be specified via a set of “bounce points”, as in [17]: each environment’s geometry can be characterized by capturing a set of uniformly sampled points along the surface of each room’s mesh, defining the environmental context

as $C = [c_1, \dots, c_K] \in \mathbb{R}^{3 \times K}$, with each bounce point c_k being normalized such that the origin is at the midpoint of the minimum/maximum coordinates as defined by the dataset metadata.

3.2. Direction-Aware Neural Field

We propose to go beyond binaural RIR modeling, and model an Ambisonic RIR $h \in \mathbb{R}^{4 \times T}$ using a direction-aware neural field (DANF) framework:

$$h = \text{DANF}(s, l, \theta, C), \quad (3)$$

where DANF is optimized based on reconstruction losses including a direction-aware component, as described below and shown in Fig 1.

We follow the architecture established in [17], which can be split into a spatial embedding module and an impulse response generation module. Using the principles of acoustic radiance transfer, the spatial embedding module learns features derived from the bounce points $c_1, \dots, c_K \in \mathbb{R}^3$, as well as the relative location vectors from the bounce points to the listener $[c_k - l]_{k=1}^K \in \mathbb{R}^{3 \times K}$ and from the bounce points to the source $[c_k - s]_{k=1}^K \in \mathbb{R}^{3 \times K}$.

Each of these feature sets is passed through a sinusoidal encoding and to a separate network (a 2-layer MLP for the bounce points matrix $C \in \mathbb{R}^{3 \times K}$ and a single layer MLP for the relative location vectors) to obtain three separate spatial features, which are concatenated into a single spatial feature $\tilde{C} \in \mathbb{R}^{3K \times D}$. We then choose a set of time samples $\{t_i\}_{i=1}^T$ which we encode via sinusoidal encoding and an MLP to create a set of time-domain basis functions $Q \in \mathbb{R}^{D \times T}$, which is multiplied with our concatenated spatial features to produce a spatial-temporal feature $E \in \mathbb{R}^{3K \times T}$. This spatial-temporal feature can then be concatenated with a learned encoding of the listener orientation $\theta \in [0, 2\pi)$ at the impulse response generation module, which predicts an N -channel impulse response via an MLP network.

3.3. Losses

A common way to assess the accuracy of an estimated signal \hat{h} is via direct mean squared error (MSE) loss [16, 17]:

$$\mathcal{L}_{\text{MSE}}(\hat{h}, h) = \|\hat{h} - h\|_2^2. \quad (4)$$

Other typical approaches approximate human perception and utilize a time-frequency representation of the signal, such as the STFT H . The magnitude STFT $|H|$ and STFT phase $\angle H$ are also often considered. Prior approaches in NFs utilize some combination of time domain and time-frequency domain signal reconstruction losses, such as the STFT loss as defined in [17], which is a combination of the magnitude loss, $\mathcal{L}_{\text{mag}} = \|\hat{H} - |H|\|_1$, phase loss $\mathcal{L}_{\text{phase}} = \|\angle \hat{H} - \angle H\|_2$, and spectral convergence $\mathcal{L}_{\text{sc}} = \frac{\|\hat{H} - |H|\|_2}{\|H\|_2}$:

$$\mathcal{L}_{\text{STFT}} = \mathcal{L}_{\text{mag}} + \mathcal{L}_{\text{phase}} + \mathcal{L}_{\text{sc}}. \quad (5)$$

To account for interchannel information that indicates directivity, we apply a novel loss utilizing the intensity vectors of the original and reconstructed RIRs. We integrate the intensity vector $\text{IV}(U)(m, f)$ in (1) along the frequency axis and the time axis, leading to a vector $\mathbf{I}(h) \in \mathbb{R}^3$. This averaged vector contains information relating to the DoA of a signal received from a single source to a listener [26]. We compute the IV loss based on the cosine distance between these averaged intensity vectors for the reference and estimated RIRs:

$$\mathcal{L}_{\text{IV}}(\hat{h}, h) = \frac{1}{2} \left(1 - \frac{\langle \mathbf{I}(\hat{h}); \mathbf{I}(h) \rangle}{\|\mathbf{I}(\hat{h})\| \|\mathbf{I}(h)\|} \right). \quad (6)$$

4. Experiments

4.1. Dataset

To define our acoustic spaces, we use the Matterport3D-RGB dataset, which contains high quality RGB-D images of multi-room interior home environments [27]. We sample 5 environments segmented into individual rooms, resulting in 100 unique rooms at varying volumes and acoustic properties. Within each room, we uniformly sample a grid of points at a fixed height of 1.5 m to avoid overlapping with furniture or other objects in the room. From this grid, for each room, we uniformly sample 1000 source-listener pairs, with a uniformly sampled listener orientation in $[0, 2\pi)$. We use the SoundSpaces simulation platform [28] to compute RIRs in FOA for each source-listener location pair.

4.2. Single-Room Performance

To evaluate the model’s performance in accurately reconstructing spatial cues, we separately examined 10 distinct rooms at varying volumes. For each room, we trained and evaluated a separate DANF model using a set of 1000 source-listener pairs, with 800 randomly assigned for training and 200 for testing.

Each model was trained on a weighted sum of the MSE signal loss, the STFT loss, and the intensity vector loss:

$$\mathcal{L}_{\text{train}} = \mathcal{L}_{\text{MSE}} + \mathcal{L}_{\text{STFT}} + \lambda \mathcal{L}_{\text{IV}}, \quad (7)$$

where we set $\lambda = 0$ as a baseline with no intensity vector loss, or $\lambda \in \{10^\gamma \mid \gamma = -4, \dots, 4\}$ to explore the effect of the intensity vector loss on the model’s ability to accurately reconstruct the directivity of a given RIR.

4.3. Room Adaptation

Prior work [17, 20] has shown NFs’ ability to be trained on multiple scenes and still provide reasonable performance for novel scenes by incorporating features from room geometry. To demonstrate DANF’s ability to adapt to new rooms, we initially pre-train a model with RIRs from a wide range of rooms

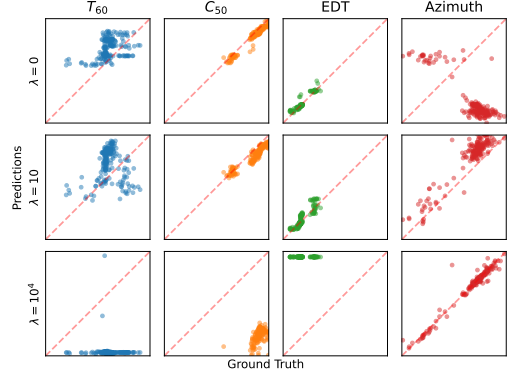


Figure 2: Single-room performance for different values of λ (row) and metrics (columns). Each metric is normalized to $[0, 1]$ for ease of visualization.

C_{pretrain} . We then freeze all trainable parameters of the model and choose a new room $C' \notin C_{\text{pretrain}}$ with a corresponding set of training and test data, with 800 and 200 samples each respectively. To evaluate the model’s ability to adapt with limited data access, we constrained the number of training samples to $[1, 80, 800]$.

We evaluated our model under three baselines settings: the model with no adaptation to examine zero-shot performance (*Zero-Shot*); a model initialized and trained from scratch with limited data (*Cold-Start*); our pre-trained model fine-tuned with the limited training set, updating all parameters (*Warm-Start*).

Additionally, we propose using low-rank adaptation (LoRA(r)) [29] to fine-tune our models. LoRA updates a weight matrix $W \in \mathbb{R}^{i \times j}$ by an r -rank adaptation matrix defined by matrices $B \in \mathbb{R}^{i \times r}$ and $A \in \mathbb{R}^{r \times j}$, such that the updated weights are given by $W' = W + \frac{1}{r}BA$. LoRA was implemented for the weights of the MLPs after the spatial temporal encoding E using the limited training set from C' , allowing for adjustment of a lower number of parameters N_p .

4.4. Evaluation

Following prior NF-based work for RIR reconstruction, we evaluate our reconstructed RIRs on three acoustic parameters: reverberation time (T60), clarity score (C50), and early decay time (EDT) [16, 17, 20]. We evaluate our models on their ability to reconstruct DoA via computing the cosine distance between the intensity vectors of the target and estimated RIRs.

5. Results

5.1. Spatial Improvements

Our results show that the IV loss significantly improves DANF’s ability to accurately reconstruct the directional characteristics of Ambisonic-RIRs.

Figure 2 shows the correlations between the acoustic measures of the oracle and the rendered RIRs. We can see the effect of adding the IV loss as the estimates of the DoA azimuth improve at increasing scales of λ . As λ increases, we see the model overcompensate towards accurate DoA estimates, leading the non-directional acoustic measures to start to fail, most performing worse than a baseline model at $\lambda \geq 10^3$. We can see the failure of the estimates of T60, C50, and EDT in Fig. 2, as we see the model either significantly under-estimating or over-estimating each parameter, despite strong DoA azimuth perfor-

Table 1: Comparison of few-shot learning techniques based on number of parameters N_p , evaluating performance when fine-tuned on 1, 80, or 800 training samples.

| | N_p | 1 training example | | | | 80 training examples | | | | 800 training examples | | | |
|------------|-------------------|--------------------|-------|--------|--------|----------------------|-------|--------|-------|-----------------------|------|-------|-------|
| | | T60 | C50 | EDT | DoA | T60 | C50 | EDT | DoA | T60 | C50 | EDT | DoA |
| Zero-Shot | 0 | 4.78 | 14.07 | 474.24 | 111.10 | - | - | - | - | - | - | - | - |
| Cold-Start | 3.5×10^6 | 22.06 | 26.79 | 945.64 | 67.87 | 22.72 | 26.80 | 944.19 | 59.99 | 0.46 | 2.89 | 10.21 | 32.86 |
| Warm-Start | 3.5×10^6 | 2.68 | 4.68 | 29.34 | 52.56 | 1.22 | 2.73 | 18.29 | 31.94 | 0.49 | 2.39 | 9.28 | 27.08 |
| LoRA(3) | 2.9×10^4 | 2.34 | 6.36 | 39.88 | 67.32 | 1.44 | 3.75 | 21.51 | 33.68 | 1.40 | 3.29 | 20.52 | 34.12 |
| LoRA(1) | 9.6×10^3 | 4.82 | 7.07 | 88.15 | 55.43 | 1.82 | 3.76 | 27.22 | 42.32 | 1.31 | 3.86 | 22.66 | 40.67 |

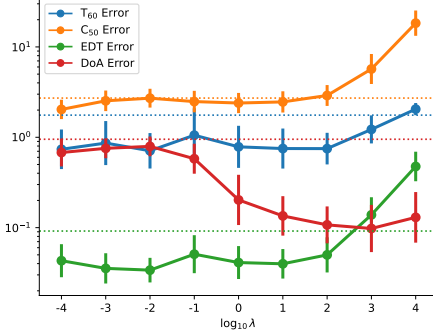


Figure 3: Model performance on single rooms with varying intensity vector loss weight λ , with the dashed line representing no intensity vector loss, or $\lambda = 0$.

mance. On the top row, the model trained without the IV loss resulted in inaccurate estimates of the DoA azimuth. With the IV loss with $\lambda = 10$, the predicted RIRs were enforced to preserve the original IVs and showed a better correlation in the azimuth of the sound source. The larger weight, $\lambda = 10^4$, further improved the DoA performance but deteriorated the distribution of other acoustic measures. This result suggests that the predicted RIRs focus on the direct sound and neglect the shape of late reverberation.

Figure 3 depicts the acoustic measures and the DoA performance with different weights for the IV loss in (6). We can see that the addition of IV loss starts to improve model performance at $\lambda = 10^{-1}$ and continues to significantly improve DoA estimation. The dotted lines correspond to the model without the IV loss, i.e., $\lambda = 0$. Although too large a weight, e.g., $\lambda = 10^4$, degrades the non-directional acoustic measures (T60, C50, EDT), we can see that the intensity vector with a moderate weight, $\lambda = 10$, improves both room acoustic metrics and directional performance. As we add the IV loss in Table 2, we can see the DoA error decrease significantly, showing that without IV loss, the model either does not capture any spatial information, or incorrectly resolves the directivity of the sound field in the opposite direction.

5.2. Fine-tuning

We compare the performance of DANF pre-trained on a wide variety of rooms, and fine-tuned either by warm-starting the model on pre-trained weights, or adapting the model’s pre-trained weights using LoRA. In addition to *Cold-Start*, we demonstrate the model’s zero-shot performance as a baseline. This model still leverages “bounce points” of the target room as its environmental context.

Table 2: Per-room performance in terms of the error in T60 [%], C50 [dB], EDT [ms] and DoA [degrees] for models trained without ($\lambda=0$) and with ($\lambda=10$) IV loss. Lower is better.

| c | T60 | | C50 | | EDT | | DoA | |
|-----|-------------|--------------|-------------|--------------|-------------|--------------|-------------|--------------|
| | $\lambda=0$ | $\lambda=10$ | $\lambda=0$ | $\lambda=10$ | $\lambda=0$ | $\lambda=10$ | $\lambda=0$ | $\lambda=10$ |
| 0 | 4.61 | 5.35 | 3.85 | 4.03 | 58.42 | 46.84 | 158 | 7 |
| 1 | 1.26 | 1.10 | 2.86 | 1.74 | 16.91 | 23.55 | 68 | 35 |
| 2 | 7.42 | 0.87 | 4.34 | 3.97 | 104.94 | 38.77 | 99 | 21 |
| 3 | 0.11 | 0.08 | 0.84 | 0.90 | 23.21 | 18.91 | 57 | 35 |
| 4 | 0.24 | 0.24 | 1.80 | 4.23 | 49.14 | 81.99 | 58 | 46 |
| 5 | 0.37 | 0.39 | 3.02 | 2.29 | 10.08 | 9.75 | 157 | 19 |
| 6 | 1.09 | 1.00 | 1.95 | 2.32 | 27.29 | 27.35 | 83 | 76 |
| 7 | 0.12 | 0.10 | 1.35 | 1.09 | 31.95 | 24.57 | 112 | 41 |
| 8 | 0.19 | 0.21 | 1.11 | 0.90 | 11.46 | 13.83 | 34 | 14 |
| 9 | 0.53 | 0.60 | 4.67 | 6.57 | 234.95 | 310.06 | 85 | 69 |

We can see in Table 1 that with RIRs for a large number of source-receiver pairs, training a model from scratch outperforms most fine-tuning measures. Conversely, in the low-data regime, we see that warm-starting the model and LoRA are able to more efficiently provide better estimates of a new environment’s acoustic properties. We see rank-3 LoRA achieving near equivalent performance with warm-starting the model, despite using less than 1% of the number of parameters as retraining the full model, producing similar performance when fine-tuned on a single measurement.

Furthermore, even in the low-data regime, all models do relatively well at estimating the DoA characteristics of a room; the spatial relationship between listener and source is generally unaffected by the acoustic properties of the environment they are in, so these relationships should translate relatively easily from pre-trained models. All models adapting from pre-trained model outperform the zero-shot case, demonstrating the effectiveness of even the low-parameter LoRA approaches.

6. Conclusion

In this work, we presented DANF, a novel neural acoustic field model for estimating direction-aware RIRs. To our knowledge, DANF is the first NF model estimating Ambisonic-format RIRs. By utilizing a directional intensity vector loss, DANF is able to accurately capture not only the acoustic properties of the environment, but also the direction-dependent acoustic properties. This is a first-of-its-kind approach to capturing direction-of-arrival (DoA) with an acoustic NF. Furthermore, we demonstrate the model’s ability to adapt to new unseen environments with limited training data. By utilizing fine-tuning strategies such as LoRA, we can see that this model’s ability to adapt outperforms training-from-scratch and zero-shot approaches.

7. References

- [1] A. Roginska and P. Geluso, *Immersive sound: The Art and Science of Binaural and Multi-channel Audio*, 1st ed. Focal Press, 2018.
- [2] M. Zaunschirm, C. Schörkhuber, and R. Höldrich, “Binaural rendering of ambisonic signals by head-related impulse response time alignment and a diffuseness constraint,” *The Journal of the Acoustical Society of America*, vol. 143, no. 6, pp. 3616–3627, 06 2018.
- [3] A. Rana, C. Ozcinar, and A. Smolic, “Towards generating ambisonics using audio-visual cue for virtual reality,” in *Proc. ICASSP*, 2019, pp. 2012–2016.
- [4] J. B. Allen and D. Berkley, “Image method for efficiently simulating small-room acoustics,” *J. Acoust. Soc. Am.*, vol. 65, no. 4, pp. 943–950, 04 1979.
- [5] R. Scheibler, E. Bezzam, and I. Dokmanic, “Pyroomacoustics: A python package for audio room simulation and array processing algorithms,” in *Proc. ICASSP*, Apr. 2018, p. 351–355.
- [6] C. Chen, U. Jain, C. Schissler, S. V. A. Gari, Z. Al-Halah, V. K. Ithapu, P. Robinson, and K. Grauman, “SoundSpaces: Audio-visual navigation in 3D environments,” in *Proc. ECCV*, 2020.
- [7] L. L. Thompson, “A review of finite-element methods for time-harmonic acoustics,” *J. Acoust. Soc. Am.*, vol. 119, no. 3, pp. 1315–1330, 03 2006.
- [8] R. Mehra, L. Antani, S. Kim, and D. Manocha, “Source and listener directivity for interactive wave-based sound propagation,” *IEEE Trans. Vis. Comput. Graph.*, vol. 20, no. 4, pp. 495–503, 2014.
- [9] R. Mehra, N. Raghuvanshi, L. Savioja, M. C. Lin, and D. Manocha, “An efficient GPU-based time domain solver for the acoustic wave equation,” *Appl. Acoust.*, vol. 73, no. 2, pp. 83–94, 2012.
- [10] J.-T. Chen, Y.-T. Lee, and Y.-J. Lin, “Analysis of multiple-shepers radiation and scattering problems by using a null-field integral equation approach,” *Appl. Acoust.*, vol. 71, no. 8, pp. 690–700, 2010.
- [11] N. Raghuvanshi and J. Snyder, “Parametric wave field coding for precomputed sound propagation,” *ACM Trans. Graph.*, vol. 33, no. 4, Jul. 2014.
- [12] R. Mignot, G. Chardon, and L. Daudet, “Low frequency interpolation of room impulse responses using compressed sensing,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 22, no. 1, pp. 205–216, 2014.
- [13] N. Raghuvanshi and J. Snyder, “Parametric directional coding for precomputed sound propagation,” *ACM Trans. Graph.*, vol. 37, no. 4, Jul. 2018.
- [14] I. Tsunokuni, K. Kurokawa, H. Matsushashi, Y. Ikeda, and N. Osaka, “Spatial extrapolation of early room impulse responses in local area using sparse equivalent sources and image source method,” *Appl. Acoust.*, vol. 179, p. 108027, 2021.
- [15] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, “NeRF: Representing scenes as neural radiance fields for view synthesis,” in *Proc. ECCV*, 2020.
- [16] A. Luo, Y. Du, M. Tarr, J. Tenenbaum, A. Torralba, and C. Gan, “Learning neural acoustic fields,” in *Proc. NeurIPS*, vol. 35, 2022, pp. 3165–3177.
- [17] K. Su, M. Chen, and E. Shlizerman, “INRAS: Implicit neural representation for audio scenes,” in *Proc. NeurIPS*, vol. 35, 2022, pp. 8144–8158.
- [18] S. Liang, C. Huang, Y. Tian, A. Kumar, and C. Xu, “AV-NeRF: Learning neural fields for real-world audio-visual scene synthesis,” in *Proc. NeurIPS*, 2023.
- [19] —, “Neural acoustic context field: Rendering realistic room impulse response with neural fields,” in *Proc. ICCV*, 2023.
- [20] Z. Chen, I. D. Gebru, C. Richardt, A. Kumar, W. Laney, A. Owens, and A. Richard, “Real acoustic fields: An audio-visual room acoustics dataset and benchmark,” in *Proc. CVPR*, 2024.
- [21] M. Pezzoli, F. Antonacci, and A. Sarti, “Implicit neural representation with physics-informed neural networks for the reconstruction of the early part of room impulse responses,” *arXiv preprint arXiv:2306.11509*, 2023.
- [22] X. Karakonstantis, D. Caviedes-Nozal, A. Richard, and E. Fernandez-Grande, “Room impulse response reconstruction with physics-informed deep learning,” *J. Acoust. Soc. Am.*, vol. 155, no. 2, pp. 1048–1059, 2024.
- [23] S. Koyama, J. G. C. Ribeiro, T. Nakamura, N. Ueno, and M. Pezzoli, “Physics-informed machine learning for sound field estimation: Fundamentals, state of the art, and challenges,” *IEEE Signal Process. Mag.*, vol. 41, no. 6, pp. 60–71, 2024.
- [24] L. Perotin, R. Serizel, E. Vincent, and A. Guérin, “CRNN-based joint azimuth and elevation localization with the ambisonics intensity vector,” in *Proc. IWAENC*, 2018, pp. 241–245.
- [25] H. Khaddour, J. Schimmel, and M. Trzos, “Three-dimensional sound source localization using B-format signals,” *IJATES²*, vol. 2, Mar. 2013.
- [26] D. Charalampous, K. George, A. Konstantinos, and P. George, “Improved localization of sound sources using multi-band processing of ambisonic components,” *J. Audio Eng. Soc.*, no. 7691, May 2009.
- [27] A. Chang, A. Dai, T. Funkhouser, M. Halber, M. Niessner, M. Savva, S. Song, A. Zeng, and Y. Zhang, “Matterport3D: Learning from RGB-D data in indoor environments,” in *Proc. 3DV*, 2017.
- [28] C. Chen, C. Schissler, S. Garg, P. Kobernik, A. Clegg, P. Calamia, D. Batra, P. W. Robinson, and K. Grauman, “SoundSpaces 2.0: A simulation platform for visual-acoustic learning,” in *Proc. NeurIPS Datasets and Benchmarks Track*, 2022.
- [29] E. J. Hu, Yelong shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, “LoRA: Low-rank adaptation of large language models,” in *Proc. ICLR*, 2022.