

Data Augmentation Using Neural Acoustic Fields With Retrieval-Augmented Pre-training

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

TR2025-045 April 04, 2025

Abstract

This report details MERL's system for room impulse response (RIR) estimation submitted to the Generative Data Augmentation Workshop at ICASSP 2025 for Augmenting RIR Data (Task 1) and Improving Speaker Distance Estimation (Task 2). We first pre-train a neural acoustic field conditioned by room geometry on an external large-scale dataset in which pairs of RIRs and the geometries are provided. The neural acoustic field is then adapted to each target room by using the enrollment data, where we leverage either the provided room geometries or geometries retrieved from the external dataset, depending on availability. Lastly, we predict the RIRs for each pair of source and receiver locations specified by Task 1, and use these RIRs to train the speaker distance estimation model in Task 2.

*IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)
Satellite Workshop on Generative Data Augmentation for Real-World Signal Processing
Applications (GenDA) 2025*

Data Augmentation Using Neural Acoustic Fields With Retrieval-Augmented Pre-training

Christopher Ick^{*†}, Gordon Wichern^{*}, Yoshiki Masuyama^{*}, François G. Germain^{*}, and Jonathan Le Roux^{*}

^{*}Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA, USA

[†]Music and Audio Research Laboratory, New York University, Brooklyn, NY, USA

Abstract—This report details MERL’s system for room impulse response (RIR) estimation submitted to the Generative Data Augmentation Workshop at ICASSP 2025 for Augmenting RIR Data (Task 1) and Improving Speaker Distance Estimation (Task 2). We first pre-train a neural acoustic field conditioned by room geometry on an external large-scale dataset in which pairs of RIRs and the geometries are provided. The neural acoustic field is then adapted to each target room by using the enrollment data, where we leverage either the provided room geometries or geometries retrieved from the external dataset, depending on availability. Lastly, we predict the RIRs for each pair of source and receiver locations specified by Task 1, and use these RIRs to train the speaker distance estimation model in Task 2.

Index Terms—room impulse response, neural acoustic fields, retrieval-augmented pre-training

I. INTRODUCTION

The goal of the Room Acoustics and Speaker Distance Estimation Challenge, which is part of the Generative Data Augmentation Workshop at ICASSP 2025, is to design systems for room impulse response (RIR) estimation at unseen locations from a small set of training examples [1]. The challenge is split into two tasks, Augmenting RIR Data (Task 1), and Improving Speaker Distance Estimation (Task 2).

Interest in RIR estimation has surged in recent years, with the successful application of room acoustics estimation in immersive experiences for virtual/extended reality, adaptive sound reproduction systems, acoustic scene understanding, and more [2]. Typical settings require a system to predict an RIR from an existing set of acoustic measurements. With modern approaches utilizing deep neural networks, the use of generated acoustic data for augmentation has been essential for training machine learning-based systems.

In our approach for Task 1, we use neural acoustic fields to build spatially continuous representations of acoustic spaces. We use an acoustic similarity-based retrieval method to build a room-specific pre-training dataset, and then fine-tune our model on limited training examples using low-rank adaptation (LoRA) [3]. After pre-training and fine-tuning, the model can be used to estimate the RIR at unseen source/receiver locations within the specified room.

II. PRELIMINARIES

In Task 1, we are asked to generate RIRs for twenty different rooms. For each room, between 5 and 10 RIRs at different

This work was performed while C. Ick was an intern at MERL.

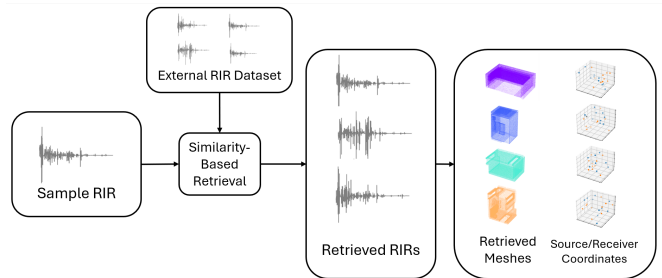


Fig. 1. Retrieval strategy for pre-training dataset selection from a sample RIR.

source and receiver locations are provided in the enrollment set. For rooms 1–10, room geometry information is provided in the form of a room mesh object file. For rooms 11–20, only the RIRs and source/listener locations are provided. An additional room (Room 0) is provided with recorded and simulated RIRs at 20 source/listener locations, and an additional 405 simulated RIRs at grid locations within the room. The goal for Task 1 is to design a system to estimate the RIRs in each of the twenty rooms at unseen source/receiver locations specified by the organizers. The generated RIRs are then evaluated on a variety of acoustic quality measures, namely RT60, EDF, and DRR [1].

III. METHODOLOGY

A. Retrieval for Pre-training

Each room $\{R_i\}_{i=1}^{20}$ is provided with N RIRs $\{h_{i,j}\}_{j=1}^N$, and pairs of corresponding source and receiver positions $(s_{i,j}, r_{i,j})$. To find appropriate data for pre-training, we use each room’s provided RIRs to retrieve RIRs from an external dataset. We use the GWA dataset [4], a large-scale RIR dataset that contains over 2 million RIRs that were simulated based on 3D scene meshes from the 3D-FRONT dataset [5], [6] to ensure acoustic diversity. Each scene in the 3D-FRONT dataset contains multiple rooms and objects; by building bounding boxes defined by the boundaries of the scene meshes, we defined a set of rooms $\{\tilde{R}_i\}_{i>20}$ corresponding to different regions within the 3D-front scenes.

To select an appropriate set of rooms from this dataset, for each room, we use a similarity-based retrieval method shown in 1. For each RIR $h_{i,j}$ in the enrollment data for room R_i , we compute a multi-band RT60 with B bands $\text{RT}_{60}(h_{i,j}) \in \mathbb{R}_+^B$. With it, we can query the dataset by using the L_2 distance

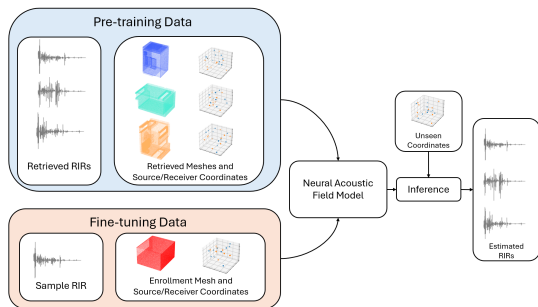


Fig. 2. Pre-training and fine-tuning strategies for neural acoustic field training.

between each measured multi-band RT60 and that of each of the RIR samples \tilde{h}_l in the GWA set $\|RT_{60}(h_{i,j}) - RT_{60}(\tilde{h}_l)\|_2$, selecting M closest RIRs in the dataset for each of the N enrollment RIRs. For each of these $N \times M$ retrieved RIRs, we collect the room(s) containing the source and receiver. We can then rank these rooms based on their frequency in the set of retrieved RIRs, providing us with a distribution of matching rooms, each with several RIR measurements at various source/receiver locations. In our experiments, for training efficiency, we limited these rooms to the 100 most frequently retrieved rooms. This set of rooms can be used to create a room-specific pre-training set for each of our 20 rooms. We compare this retrieval method with randomly selected rooms in IV. In the cases of rooms 11–20, where no room mesh is provided, we further retain the most closely matched room’s geometry for use as the estimated geometry for the queried room.

B. Neural Acoustic Field

For our model, we use a simplified single-channel version of the INRAS architecture [7], excluding orientation information. This model takes a pair of source and listener positions in 3D space, as well as geometric features in the form of sampled points from the surface of the room mesh (referred to as bounce points). Following [8], for each room, we evenly sample the room’s mesh using Poisson disk sampling, such that we are provided with K evenly spaced points characterizing the room’s geometry. These bounce points, as well as the coordinates of the source and receiver, are used as inputs to our model. The model then uses a sinusoidal encoding and a multilayer perceptron to generate representations of the room’s IR measurement at a given pair of source and receiver locations. After training, this latent representation can then be used to generate RIRs at an unseen pair of source and receiver positions.

C. Fine-Tuning

During pre-training, we expose the model to a wide range of source and receiver location pairs inside a variety of room geometries. After this step, we can then fine-tune the model on the N provided enrollment RIRs. Several approaches for fine-tuning can be taken; recent work in LLM fine-tuning has shown the success of using low-rank adaptation, which adjusts a low-rank matrix of weights to adapt a model to a

TABLE I
PERFORMANCE ON ROOM 0 OF MODELS USING DIFFERENT PRE-TRAINING SETS AND FINE-TUNING METHODS IN TERMS OF THE ERROR IN RT60 [%], EDF [dB], AND DRR [dB].

Pre-training set	Fine-tuning method	RT60	EDF	DRR
Retrieved-GWA	LoRA-1	0.090	0.520	3.009
Retrieved-GWA	All Parameters	0.174	2.736	4.626
Random-GWA	LoRA-1	0.186	0.652	5.155
Random-GWA	All Parameters	0.792	0.562	11.172
None	All Parameters	0.417	4.416	8.917

new domain with a limited amount of data [3]. For a given $d \times d$ weight matrix W of the model, two additional weight matrices $A, B \in \mathbb{R}^{d \times r}$ are introduced to update the weights as $\bar{W} = W + BA^T$, where BA^T defines a rank r weight matrix learned on the enrollment data. We also experiment with fine-tuning the model by simply using the pre-trained model’s weights to initialize a model trained on the provided training examples. After fine-tuning, we perform inference by providing the coordinates of the target source and receiver locations for evaluation, along with the provided (or retrieved, for rooms 11–20) room mesh.

IV. RESULTS

We tested our method on the 5 provided examples from “Room_0”, with the results shown in Table IV, evaluating on the remaining provided examples in that room. We can see in Table IV that our strategy’s ability to reconstruct rooms is effective; we found that Rank-1 LoRA was the most effective method for capturing the EDF of the RIR. We used the trained models for each room to generate data for Task 2 using the coordinates provided in the enrollment dataset, then fine-tuned the provided SDE model on this data and submitted the results.

REFERENCES

- [1] J. Lin, G. Götz, H. S. Llopis, H. Hafsteinsson, S. Guðjónsson, D. G. Nielsen, F. Pind, P. Smaragdīs, D. Manocha, J. Hershey, T. Kristjánsson, and M. Kim, “Generative Data Augmentation Challenge: Synthesis of room acoustics for speaker distance estimation,” in *Proc. ICASSP*, 2025.
- [2] A. Tewari, J. Thies, B. Mildenhall, P. P. Srinivasan, E. Tretschk, W. Yifan, C. Lassner, V. Sitzmann, R. Martin-Brualla, S. Lombardi, T. Simon, C. Theobalt, M. Nießner, J. T. Barron, G. Wetzstein, M. Zollhöfer, and V. Golyanik, “Advances in neural rendering,” *arXiv preprint arXiv:2111.05849*, 2021.
- [3] E. J. Hu, Y. 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.
- [4] Z. Tang, R. Aralikkatti, A. Ratnarajah, and D. Manocha, “GWA: A large geometric-wave acoustic dataset for audio processing,” in *Proc. SIGGRAPH*, 2022.
- [5] H. Fu, B. Cai, L. Gao, L.-X. Zhang, J. Wang, C. Li, Q. Zeng, C. Sun, R. Jia, B. Zhao *et al.*, “3D-FRONT: 3D furnished rooms with layouts and semantics,” in *Proc. ICCV*, 2021, pp. 10 933–10 942.
- [6] H. Fu, R. Jia, L. Gao, M. Gong, B. Zhao, S. Maybank, and D. Tao, “3D-FUTURE: 3D FURNITURE SHAPE WITH TEXTURE,” *IJCV*, pp. 1–25, 2021.
- [7] K. Su, M. Chen, and E. Shlizerman, “INRAS: Implicit neural representation for audio scenes,” in *Proc. NeurIPS*, vol. 35, 2022, pp. 8144–8158.
- [8] 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.