Explicit Context-Driven Neural Acoustic Modeling for High-Fidelity RIR Generation

Chen Si[1], **Qianyi Wu**[2], **Chaitanya Amballa**[3], **Romit Roy Choudhury**[3]

[1]University of California, San Diego
[2]Monash University
[3]University of Illinois, Urbana-Champaign

Abstract

Realistic sound simulation plays a critical role in many applications. A key element in sound simulation is the room impulse response (RIR), which characterizes how sound propagates from a source to a listener within a given space. Recent studies have applied neural implicit methods to learn RIR using context information collected from the environment, such as scene images. However, these approaches do not effectively leverage explicit geometric information from the environment. To further exploit the potential of neural implicit models with direct geometric features, we present Mesh-infused Neural Acoustic Field (`MiNAF`), which queries a rough room mesh at given locations and extracts distance distributions as an explicit representation of local context. Our approach demonstrates that incorporating explicit local geometric features can better guide the neural network in generating more accurate RIR predictions. Through comparisons with conventional and state-of-the-art baseline methods, we show that MiNAF performs competitively across various evaluation metrics. Furthermore, we verify the robustness of MiNAF in datasets with limited training samples, demonstrating an advance in high-fidelity sound simulation.

Figure 1: **Task Overview**. To record audio, a loudspeaker (Tx) and several microphones (Rx) in red are placed at different known locations within the room. The room's rough layout is available from images, 3D mesh, or LiDAR scans. `MiNAF` learns to reconstruct the Room Impulse Response (RIR) from sparse audio recordings at red microphones and explicit geometric features from probing the environment. The model can then predict the RIR at any unseen receiver location (green microphone), effectively simulating the binaural acoustic experience anywhere in the room.

Introduction

Realistic sound simulation is an integral part of augmented and virtual reality systems. Research has shown that increasing the fidelity of environmental sound, combined with spatial audio, significantly enhances immersion (Kurucz et al. 2023). The key to estimating a room's acoustic field is the Room Impulse Response (RIR), which models how an impulse emitted from a transmitter (Tx) arrives at a given receiver (Rx) location. The RIR embeds the room's geometry, materials, and other factors influencing sound propagation. Using the RIR from a specific $\langle Tx, Rx \rangle$ pair allows us to recreate the sound heard at the receiver by convolving the original sound with the RIR, enabling simulations such as "What if a speaker's position changes in a living room?" or "What if a piano is moved to the left of a stage?" to assess the acoustic impact of each scenario.

Recent studies on RIR generation have increasingly focused on using neural implicit models to estimate the acoustic field for a given scenario. Inspired by the success of neural implicit models in 3D reconstruction, as demonstrated by NeRF (Mildenhall et al. 2021), which disregards physical interactions such as reflection and reverberation when light propagates through a room, Luo *et al.* (Luo et al. 2022) applied an MLP to learn the acoustic field while maintaining an implicit feature grid for each Tx and Rx location. To better inform the model about environmental information, Liang *et al.*(Liang et al. 2023b) proposed NACF, which incorporates a series of RGB and depth images as additional global context. Furthering this approach, Brunetto *et al.*(Brunetto, Hornauer, and Moutarde 2024) trained a NeRF model on environmental images and queried it to provide color and density information as an enhanced environmental context. However, these methods are based on image-based context extraction, which provides only indirect and implicit environmental information. Therefore, they may not fully exploit the potential of neural implicit models for RIR generation.

To enable practical, accurate, and fast RIR generation that incorporates direct local context, we introduce Mesh-infused Neural Acoustic Field (`MiNAF`). It leverages a rough room mesh along with a few sparsely collected RIR measurements to capture the room's overall acoustic characteristics, enabling it to predict RIRs for any new $\langle Tx, Rx \rangle$ pair, as illustrated in Fig. 1. Unlike previous models that may rely on

global features or learn hidden features, MiNAF gathers precise local context at a given Tx or Rx location by actively probing the surrounding mesh with multiple rays cast in a series of uniformly distributed directions. This ray-based approach captures detailed distance measures from the center to the point of first hit (PoFH) on each ray and their statistical distribution, which includes mean and standard deviation for distances on neighboring rays and a global histogram count of all distances. This provides a clear and direct representation of the local geometry around each queried point. By obtaining these local features, which reflect both immediate spatial configuration and obstacle proximity, MiNAF achieves a deeper understanding of the acoustic environment. Equipped with this robust feature set, a simple MLP is then used to estimate the RIR spectrum effectively.

In this work, we (1) propose a method to better capture local context explicitly and directly by querying the mesh as input to a neural implicit model that learns the acoustic field in a given scenario. (2) We evaluate MiNAF across multiple cases against prior SOTA models and show that direct local context leads to better results. (3) We demonstrate that MiNAF remains robust under limited training data, outperforming previous SOTA methods in data-scarce conditions.

Related Work

Learning Acoustic Fields

Ray tracing has long been a foundational approach to RIR estimation and generation. Allen and Berkley (1979) introduced the image source method, which treats surfaces as mirrors for sound waves to compute RIRs in the time domain. Subsequent advances in ray tracing and parallel computing have significantly improved its speed and accuracy (Tsingos et al. 2001; Funkhouser et al. 2004; Schissler and Manocha 2016). However, ray tracing remains computationally expensive and struggles with modeling diffraction and fine spatial resolution (Rungta et al. 2018). Beyond ray tracing, wave equation-based approaches simulate sound propagation in grid-based representations of space (Gumerov and Duraiswami 2009; Savioja and Valimaki 2003). Acoustic field encoding methods reconstruct spatial sound energy distributions from sparse data using explicit parametric models (Antonello et al. 2017; Jin and Kleijn 2015). While these classical methods laid a strong foundation, they often face scalability and adaptability issues in complex environments.

Introducing visual contexts as an additional helper for RIR generation, several recent studies have proposed a two-step approach that leverages visual inputs—such as images or room models—to first estimate acoustic coefficients from measured RIRs, and then reconstruct the RIR accordingly (Schissler, Loftin, and Manocha 2017; Tang et al. 2020; Ratnarajah, Tang, and Manocha 2020). Alternatively, some researchers have explored direct RIR generation from visual inputs, bypassing intermediate acoustic coefficient estimation and yielding promising results (Liang et al. 2023a; Singh et al. 2021; Majumder et al. 2022). These visual-based methods represent a significant shift toward incorporating environmental context, enabling more adaptable and data-driven approaches to RIR generation.

Implicit Acoustic Field Modeling

Recent advances in implicit scene modeling suggest a promising future for encoding local environmental features with neural networks. Traditional 3D scene modeling relies on explicit representations—such as point clouds, meshes, or voxels—which store scenes as collections of points, surfaces, or volumetric data, requiring substantial storage and explicit iteration to render new views (Bai et al. 2022; Zhang et al. 2023). In contrast, Mildenhall et al. (2021) introduced Neural Radiance Fields (NeRF), which encode color and density directly within a neural network. Building on this idea, later NeRF-based models demonstrated the ability of neural implicit representations to model few-shot scenes (Yu et al. 2021; Niemeyer et al. 2022), enable fast rendering (Barron et al. 2021; Garbin et al. 2021), scale to large environments (Reiser et al. 2021; Turki, Ramanan, and Satyanarayanan 2022), and reconstruct scenes from sparse data (Ni et al. 2024; Lee et al. 2024; Huang et al. 2024). A key insight across these works is to bypass intermediate ray-based physical processes like reflection and reverberation, focusing instead on modeling interactions with spatial environments at the transmitter and receiver.

Inspired by recent progress in neural implicit modeling, several works have explored its potential for learning room acoustic fields. Luo et al. (2022) proposed NAF, which maintains a 2D grid of learnable hidden features to represent the local environment around a given Tx or Rx. These features are queried by a simple MLP to predict the RIR spectrum. As an extension, He et al. (2024) introduced DeepNeRAP, which also models the acoustic field as 2D feature maps but adopts a pyramid-based multi-resolution sampling strategy on the grid map to encode spatial context at different scales. More directly, Richard, Dodds, and Ithapu (2022) introduced IR-MLP, which directly regresses RIRs from spatial-temporal coordinates. To model scenes more explicitly, Su, Chen, and Shlizerman (2022) proposed INRAS, which separately encodes Tx, Rx, and geometry through dedicated neural modules applied sequentially. Using more visual context, NACF (Liang et al. 2023b) uses RGB and depth images to form a global scene representation, combined with relative $\langle Tx, Rx \rangle$ positions as model inputs. NeRAF (Brunetto et al. 2024) fine-tunes a NeRF model to extract color and density features at Tx and Rx locations from multiple viewing perspectives, and then compresses them into local feature embeddings for neural RIR prediction. Complementing these scene-driven approaches, Li et al. (2025) explored which function classes benefit neural audio fields and propose Fourier-KANs to improve representational expressiveness without positional encoding.

Unlike light, sound propagation is significantly affected by reflections and reverberations and is closely linked to the local geometry along the propagation path. Although recent implicit methods have shown promising results in learning acoustic fields by incorporating context, they largely overlook the direct influence of explicit local geometry on sound behavior. Therefore, it is crucial to emphasize the role of explicit local geometry in RIR generation, as it can provide a more precise, physically grounded understanding of how sound interacts with surfaces and obstacles in a space.

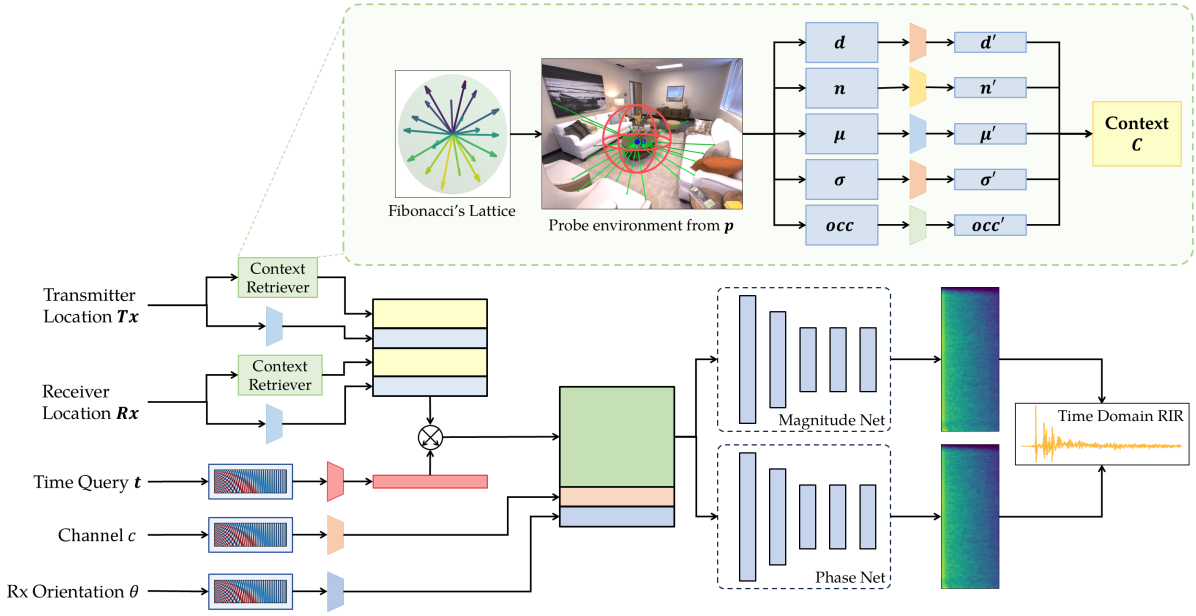

Figure 2: **Workflow Overview**. A context retriever first extracts physical features at the transmitter and receiver locations, generating context vectors $\mathbf{C_{Tx}}$ and $\mathbf{C_{Rx}}$. These are concatenated with their respective positions to form a comprehensive representation of the environment. Next, positional encoding is applied to the time t, channel c, and orientation θ. The encoded time is then element-wise embedded into the combined context, and this enriched context, along with channel and orientation embeddings, is passed to the core neural network. The network outputs the log-magnitude and IF spectra, which are used to reconstruct the Room Impulse Response (RIR).

`MiNAF`: Mesh-infused Neural Acoustic Field

Fig. 2 illustrates the workflow of `MiNAF`. This section outlines the process of collecting local environmental features and constructing the `MiNAF` model. We begin by defining the problem setup in Task Definition. Geometry Context Collection describes how raw explicit geometric context is obtained by probing the mesh. Context Fusion details the composition of input features. Finally, Model Training presents the training procedure for the `MiNAF` model.

Task Definition

Given an arbitrary room with a $\langle Tx, Rx \rangle$ pair inside it, the objective is to estimate the RIR at Rx's position $\mathbf{p_{Rx}} \in \mathbb{R}^3$, using context information $\mathbf{C}$, Tx position $\mathbf{p_{Tx}} \in \mathbb{R}^3$, and Rx orientation θ as inputs. Additionally, since binaural RIR is considered in practice, the (left/right) channel index $c \in [1, 2]$ is also incorporated. Given that the RIR is a time-domain signal and different inputs yield RIRs of varying lengths, the model is queried frame by frame across time t, and the outputs are concatenated to produce the final RIR.

Time-domain RIRs, $\mathbf{w}[t], t \in [0, n]$, are difficult to predict directly due to their highly non-smooth fluctuations. To address this, we apply the Short Time Fourier Transform (STFT), which converts $\mathbf{w}[t]$ into a time-frequency representation where each column corresponds to the Fourier transform of a windowed segment. The final time-domain RIR is recovered by applying the inverse STFT to the predicted spectrum. However, since STFT bins are complex-valued, they are typically decomposed into magnitude and phase components for learning. The phase, however, wraps around at 2π, introducing discontinuities that hinder learning. To mitigate this, we instead predict the *instantaneous frequency* (IF), defined as the time derivative of the phase (Luo et al. 2022; Engel et al. 2019). This representation avoids phase wrapping and provides a smoother, more learnable signal.

In `MiNAF`, two neural networks with identical architecture are used to predict the log-magnitude and IF components: $\Phi_{Mag}(\mathbf{C}, \mathbf{p_{Tx}}, \mathbf{p_{Rx}}, \theta, c, t) \rightarrow M_t$ and $\Phi_{IF}(\mathbf{C}, \mathbf{p_{Tx}}, \mathbf{p_{Rx}}, \theta, c, t) \rightarrow \text{IF}_t$, where M_t and IF_t represent the t^{th} columns of the log-magnitude and IF components on the STFT spectrum.

Geometry Context Collection

Recent advances in SLAM and photogrammetry have greatly streamlined room layout creation, enabling automatic conversion of data from phone cameras, LiDAR, or video into 3D point clouds or meshes (Schönberger et al. 2016; Schönberger and Frahm 2016; Wang et al. 2024; Chen et al. 2024; Yu, Sattler, and Geiger 2024). We use the ground truth mesh from (Chen et al. 2020) to evaluate our components and verify robustness using the reconstructed mesh.

For a given Tx or Rx position $\mathbf{p} \in \mathbb{R}^3$, we employ Fibonacci's lattice to generate N uniformly distributed points on a unit sphere centered at $\mathbf{p}$. These points define N rays, denoted $r_1, \ldots, r_N$, that emanate from $\mathbf{p}$ into the environment. By computing their interactions with the surrounding

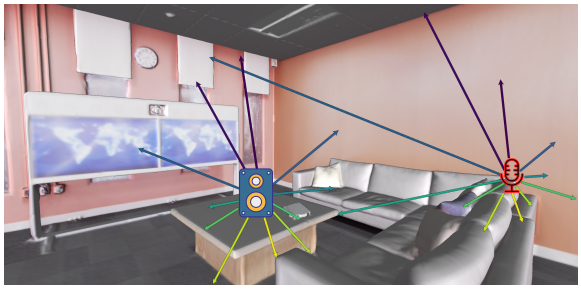

Figure 3: **Demonstration of Context Collection**. For each transmitter and receiver location $\mathbf{p}$, we uniformly sample N rays with Fibonacci's lattice. Then, we measure the distance for a ray from the center $\mathbf{p}$ to its point of first hit (PoFH). A general description of the distribution of neighboring rays' distances is also included in this measurement.

mesh, we derive a feature vector characterizing the spatial properties of the environment, as shown in Fig. 3. We collect the following features from these ray-mesh interactions as a representation of context around $\mathbf{p}$:

1. **Distances:** For each ray r_i, we define the point of first hit (PoFH) as the point where the ray first intersects a mesh surface. We denote this point as $\mathbf{q_i}$ and measure the distance from center $\mathbf{p}$ to $\mathbf{q_i}$. This forms a distance vector $\mathbf{d} \in \mathbb{R}^N$ by querying the mesh for all rays.
2. **Normals:** At each $\mathbf{q_i}$, we record the unit normal vector of the mesh (*i.e.*, the surface orientation). The set of unit normal vectors at all PoFHs is denoted by $\mathbf{n} \in \mathbb{R}^{N\times 3}$.
3. **Proximity statistics:** For each ray r_i, we identify the nearest N_θ rays based on their angles with r_i. We compute the set of mean and standard deviation of these neighboring distances for all rays, denoted as $\boldsymbol{\mu} \in \mathbb{R}^N$ and $\boldsymbol{\sigma} \in \mathbb{R}^N$, respectively.
4. **General distribution:** For all the rays $r_i, i \in [1, N]$, we define N_τ distance thresholds $\delta_1, \ldots, \delta_{N_\tau}$. We count the number of rays with distances within each threshold, forming the final feature vector $\mathbf{occ} \in \mathbb{R}^{N_\tau}$. To be clear, the k-th element of this vector is given by, $\mathbf{occ}_k = |i, |\overline{\mathbf{pq_i}}| \leq \delta_k|$.

To intuitively capture environmental information for robust RIR prediction, we begin with the most direct approach: measuring distances $\mathbf{d}$ to the closest obstacles in all directions. With sufficient rays probed, the distances and their distribution should entail a sparse but accurate description of the local environment. However, directly inputting a collection of raw distances may not allow a neural network to effectively capture critical characteristics in their distribution. To address this, we introduce mean and standard deviation vectors, $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$, along with the occupancy vector $\mathbf{occ}$, to summarize the distribution of distances. These vectors represent a precomputed statistical understanding of distance variability, allowing the model to leverage this distributional context without relying on network nonlinearity to infer it implicitly. In addition, we incorporate normal vectors $\mathbf{n}$ at PoFHs to capture surface orientations, which is crucial for understanding the directions in which sound may get reflected. These normals are particularly valuable when the measurement point $\mathbf{p}$ is close to surfaces, such as in cornered or angled regions, where geometric features like wall intersections significantly influence sound reflection patterns.

The features are projected into a latent space of dimension h using separate nonlinear projection networks that act as dimension compressors; exact formulations are provided in the supplementary material. We then concatenate the resulting feature vectors to form the final latent feature matrix at position $\mathbf{p}$, denoted as $\mathbf{C_p} = [\mathbf{d}'\ \mathbf{n}'\ \boldsymbol{\mu}'\ \boldsymbol{\sigma}'\ \mathbf{occ}'] \in \mathbb{R}^{h\times 5}$.

Context Fusion

As stated in Task Definition, the neural acoustic field Φ takes the locations of Tx and Rx: $\mathbf{p_{Tx}} \in \mathbb{R}^3$ and $\mathbf{p_{Rx}} \in \mathbb{R}^3$. We apply sinusoidal position encoding to them, projecting to a higher-dimensional space: $\gamma(\mathbf{p_{Tx}}), \gamma(\mathbf{p_{Rx}}) \in \mathbb{R}^{2\times L}$, where L is the number of frequency bands. This encoding, widely used in transformer models, effectively maps scalar inputs to high-dimensional representations (Dufter, Schmitt, and Schütze 2022). The encoded positions are then projected into the h-dimensional latent space to align with remaining context features. In parallel, we collect Tx and Rx geometric context features using Geometry Context Collection, yielding $\mathbf{C_{Tx}}, \mathbf{C_{Rx}} \in \mathbb{R}^{h\times 5}$. We concatenate them with the encoded positions to form the full context: $\mathbf{C_p} = [\mathbf{C_{Tx}}\ \mathbf{p}'_{\mathbf{Tx}}\ \mathbf{C_{Rx}}\ \mathbf{p}'_{\mathbf{Rx}}] \in \mathbb{R}^{h\times 12}$. Detailed formulations are in Section A.2 of the supplementary material.

In most prior works, the time index t is simply concatenated with input features. However, this often causes the model to ignore temporal distinctions across spectral columns, leading to over-smoothed outputs where adjacent time steps lack distinct spectral characteristics. To mitigate this, we follow Liang et al. (2023b) and apply sinusoidal position encoding to the time index using L frequencies, followed by projection into the h-dimensional latent space. The detailed formulation is provided in Section A.2 of the supplementary material. We then element-wise multiply each column of the context $\mathbf{C}$ with the projected time vector t', i.e., $\mathbf{C'_t} = \mathbf{C} \odot t'$, where $\mathbf{C'_{ij}} = \mathbf{C_{ij}} \times t'_j$. This operation preserves the shape of the context matrix, yielding $\mathbf{C'_t} \in \mathbb{R}^{h\times 12}$ with temporal information effectively embedded.

Model Training

After we obtain the time-embedded context information, we pass it into the core neural network along with the fixed receiver orientation θ and channel c. `MiNAF` uses a simple MLP to map between the input position-context information and the output RIR in the form of STFT spectrum.

To train this core neural network, we employ a loss function composed of an L1 loss between the spectrum and an additional term based on the difference between the ground-truth and the predicted audio's Schroeder curve, which reflects the decay of RIR energy. In general, the loss can be expressed as: $\mathcal{L} = \mathcal{L}_1(\text{spectrum}) + \alpha \times \mathcal{L}_{\text{Schroeder}}(\text{waveform})$, where α is a scaling factor that balances the two loss terms. It's observed that the L1 loss primarily ensures the predicted spectrum resembles the overall shape of the ground truth, while the Schroeder curve-based loss term focuses on refin-

ing details in the energy decay profile by highlighting and capturing cross-column relationships within the spectrum.

Experiments

Dataset. We utilize the SoundSpaces dataset (Chen et al. 2020), which derives its 3D room environments from the Replica dataset (Straub et al. 2019). SoundSpaces includes 18 high-quality indoor scenes of varying sizes, each with a grid of positions and simulated RIRs for every $\langle Tx, Rx \rangle$ pair. The transmitter is omnidirectional, while the receiver has a directional gain pattern with 4 possible orientations $\theta \in 0°, 90°, 180°, 270°$. For fair comparison with prior work (e.g., (Luo et al. 2022)), we use the same 6 rooms: 2 single rooms with rectangular walls (*office 4*, *room2*), 2 with non-rectangular walls (*frl apartment 2*, *frl apartment 5*), and 2 multi-room layouts (*apartment 1*, *apartment 2*). To train MiNAF, we use the Replica mesh and SoundSpaces RIRs for each room, splitting the data into 80% training, 5% validation, and 15% test sets.

Metrics. For fair comparison, we adopt the same evaluation metrics as prior works (e.g., (Luo et al. 2022)):

- **T60**: Time for the waveform energy to decay by 60dB after an initial 5dB drop, describing the reverberation tail.
- **C50**: Ratio of early (within 50ms) to late sound energy. Higher C50 values indicate better speech clarity.
- **EDT**: Time for energy to decay by 10dB from the initial level, providing a fast reverberation assessment.

Baselines. We compare five neural implicit models and two traditional encoding methods.[1]

Comparison. Predicting the time domain RIR requires predicting both magnitude and phase of the STFT. Note that 3 of the 5 baselines use different phase estimation methods. NAF (Luo et al. 2022) proposes to predict the magnitude and phase but reports improved results by reconstructing the waveform with a randomly initialized phase. Their results have aligned with findings in (Singh et al. 2021). AV-NeRF (Liang et al. 2023a) reconstructs the waveform using the ground-truth phase. NeRAF (Brunetto et al. 2024) applies the Griffin-Lim algorithm (Griffin and Lim 1984), which iteratively refines the phase estimate starting from a random initial guess. To enable a comprehensive comparison with all baselines, we implemented four variants of MiNAF: (1) MiNAF(GTP), which reconstructs with ground-truth phase; (2) MiNAF(PreP), which reconstructs with predicted phase; (3) MiNAF(RanP), using random phase for reconstruction; and (4) MiNAF(GLim), which employs the Griffin-Lim algorithm. Additionally, to assess the gap between predicted and true phase, we report (5) MiNAF(GTM+PreP), which reconstructs the RIR using ground-truth magnitude and predicted phase.

[1]We follow prior SOTA baselines: INRAS (Su, Chen, and Shlizerman 2022), NAF (Luo et al. 2022), NACF (Liang et al. 2023b), AV-NeRF (Liang et al. 2023a), and NeRAF (Brunetto, Hornauer, and Moutarde 2024); and traditional codecs AAC (International Organization for Standardization 2006) and Opus (Xiph.Org Foundation 2012). A detailed description of each baseline is provided in the supplementary material.

Method	T60 (%) ↓	C50 (dB) ↓	EDT (sec) ↓
Opus-nearest	10.10	3.58	0.115
Opus-linear	8.64	3.13	0.097
AAC-nearest	9.35	1.67	0.059
AAC-linear	7.88	1.68	0.057
INRAS 2022 $^{\%}$	3.14	0.60	0.019
NAF 2022 *	3.18	1.06	0.031
NACF 2023 $^{\%}$	2.36	0.50	0.014
AV-NeRF 2023 $^{+}$	2.47	0.57	0.016
NeRAF 2024 $^{\#}$	2.04	**0.39**	0.011
MiNAF(GTP) $^{+}$	**1.49**	**0.42**	**0.002_3**
MiNAF(PreP) $^{\%}$	2.35	0.58	0.004_2
MiNAF(RanP) *	**1.42**	**0.41**	**0.002_2**
MiNAF(GLim) $^{\#}$	**1.59**	**0.41**	**0.002_2**
MiNAF(GTM+PreP)	2.17	0.50	0.003_4

Table 1: Comparison of MiNAF with baselines on SoundSpaces in terms of T60, C50, and EDT metrics, averaged over all scenes. For all metrics, lower values indicate better performance. Top 3 results for each metric are **bolded**.

Method	Large	Medium	Small
NAF	0.404	0.384	0.350
MiNAF (log-mag. prediction)	0.389	0.371	0.312

Table 2: Spec. loss comparison against NAF and MiNAF over different-sized scenarios. Lower value is better.

Results

Quality of RIR Reconstruction

A shared superscript between a baseline and MiNAF indicates the same method for reconstructing time-domain RIR from the predicted spectrum. ■ When we use the same reconstruction method, MiNAF consistently outperforms all corresponding baselines in terms of T60 and EDT, and is either better or comparable in C50. This marks a good ability of MiNAF to capture the RIR's pattern in energy decay. ■ For example, with Griffin–Lim, MiNAF(GLim) achieves a 22% lower T60 than NeRAF, with nearly the same C50. With the ground-truth phase, MiNAF(GTP) also outperforms AV-NeRF across all metrics, with a notable 40% and 26% improvement in T60 and C50. ■ Even when relying on predicted phase, MiNAF(PreP) has a slight improvement in T60 and a comparable C50 than NACF. Using random phase, MiNAF(RanP) performs almost a $2\times$ as its baseline NAF over T60 and C50. ■ Finally, MiNAF and all neural implicit baselines show large improvements over conventional audio coding methods reported in the top 4 rows of Table 1.

In addition to the shared metrics mentioned above, MiNAF is also evaluated on the unique metrics used in each prior work. As shown in Table 2, MiNAF achieves lower spectrum loss than NAF across all scenario sizes, indicating a better spectrum learning and generalizing ability. As shown in Fig. 4, compared with INRAS, MiNAF performs better on SNR (GTP) but falls short on cases using predicted

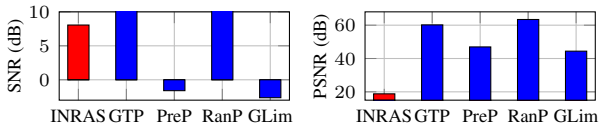

Figure 4: Comparisons: (1) Average SNR and (2) Average PSNR of MiNAF against INRAS.

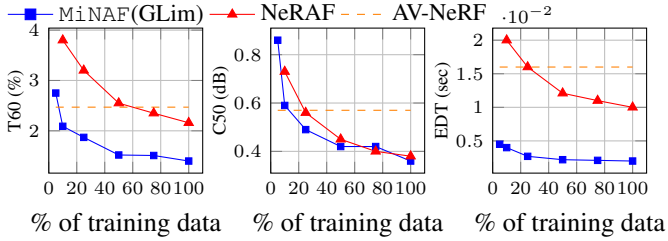

Figure 5: **Few-shot Experiment**. MiNAF trained with 10% of data outperforms AV-NeRF trained with the entire dataset. MiNAF also outperforms NeRAF in its ability to quickly adapt to few-shot datasets.

phases. In terms of PSNR, MiNAF produces better results than INRAS on synthesized RIRs overall methods.

Ablation Studies

Following Brunetto et al. (2024), we experiment with *room 2* to evaluate the effectiveness of each module in MiNAF.

Impact of Less Training Data. Collecting a large number of RIR measurements from an indoor mesh can be computationally intensive. Therefore, we assess the impact of reduced amounts of training data on RIR quality. Fig. 5 presents results for all three metrics, comparing the top-performing baseline, NeRAF, with MiNAF(GLim); for reference, the second-best baseline, AV-NeRF, trained on the entire dataset, is also included.

In terms of T60 and EDT, MiNAF(GLim) demonstrates less performance degradation than NeRAF across nearly all reduced data conditions. Notably, MiNAF surpasses AV-NeRF on T60 with just 10% of the data, showing a 15.4% improvement, and achieves a 71.8% gain in EDT with only 5% of the data. The slope of quality degradation for MiNAF is also slower than NeRAF when training data is limited. These results highlight MiNAF's strong resilience and efficiency under data-scarce scenarios.

Impact of Noisy Meshes. Surface reconstruction from multiview images has seen significant advancements in recent years (Wang et al. 2024; Li et al. 2023), providing a solid foundation for our approach, where we use explicit geometry for accurate acoustic modeling. However, in real-world applications, reconstructed meshes may still exhibit noticeable geometric errors. To assess the robustness of MiNAF under such conditions, we conducted two sets of experiments: one with controlled synthetic noise and another with a practically reconstructed mesh. First, we introduced Gaussian noise with increasing variances to the ground-truth mesh vertices. As shown in Fig. 7 and Fig. 6, MiNAF's per-

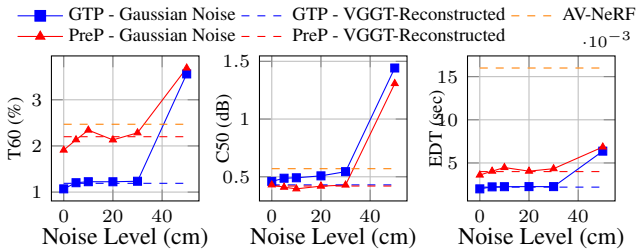

Figure 6: **Noisy Mesh Experiment**. MiNAF trained with meshes contaminated with Gaussian noise of 5cm, 10cm, and 20cm all perform at the same level with the original one, despite some minor performance degradation. When noise added reaches 50cm, the model failed to converge.

Method	T60 (%) ↓	C50 (dB) ↓	EDT (sec) ↓
MiNAF(PreP)	**1.91**	0.4331	**0.0036**
- w/o **C**	2.51	**0.4314**	0.0047
- w/o **n**	2.46	0.4587	0.0047
- w/o μ and σ	**2.45**	**0.4171**	**0.0042**
- w/o **occ**	**2.34**	**0.3895**	**0.0044**
MiNAF(RanP)	**1.11**	**0.4159**	**0.0021**
- w/o **C**	1.42	**0.3839**	0.0026
- w/o **n**	**1.23**	0.4989	**0.0023**
- w/o μ and σ	1.26	**0.4531**	**0.0023**
- w/o **occ**	**1.24**	0.5308	**0.0023**
MiNAF(GLim)	**1.40**	**0.4038**	**0.0020**
- w/o **C**	**1.27**	**0.3847**	**0.0023**
- w/o **n**	1.49	0.7010	0.0027
- w/o μ and σ	**1.44**	**0.6317**	**0.0026**
- w/o **occ**	1.59	0.7563	0.0029

Table 3: Comparison of MiNAF with entire or partial context features removed during training. When removing one or more context items, the model performance generally drops, indicating the utility of local context to RIR reconstruction. Top-1, top-2, and top-3 results are highlighted in **bold**.

Method	T60 (%) ↓	C50 (dB) ↓	EDT (sec) ↓
MiNAF(PreP)	**1.91**	**0.4331**	**0.0036**
- w/o t-EMB	2.88	0.5568	0.0055
MiNAF(RanP)	**1.11**	**0.4159**	**0.0021**
- w/o t-EMB	1.56	0.3885	0.0029
MiNAF(GLim)	1.40	0.4038	**0.0020**
- w/o t-EMB	**1.37**	**0.4097**	0.0026

Table 4: Comparison of MiNAF(PreP), MiNAF(RanP), and MiNAF(GLim) with time embedding removed during training versus the baseline model. It can be observed that when concatenating time as a direct input, the model performance drops compared to the results from embedding time into the context, suggesting the influence of using this time embedding in context. Top results are highlighted in **bold**.

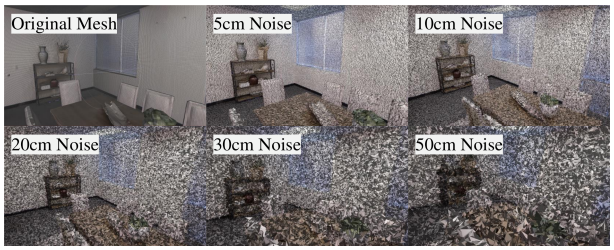

Figure 7: Preview of mesh with/without noise.

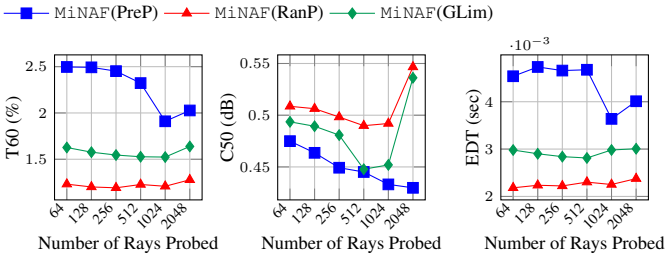

Figure 8: **Probe Density Experiment**: Performance of MiNAF tends to decrease with fewer rays, etc.

formance degrades gracefully under moderate noise, with minimal impact observed up to 30cm variance. Performance only deteriorates sharply when the mesh becomes severely distorted at 50cm, where it becomes largely unrecognizable.

To further test robustness under realistic reconstruction errors, we used VGGT (Wang et al. 2025)—a recent multiview surface reconstruction method—to generate a mesh from 45 casually captured RGB images (matching the image budget of prior works such as NACF and NeRAF). Mesh generation completed within 30s using VGGT's pretrained model. We then re-evaluated MiNAF on this reconstructed mesh. As shown in Fig. 6, performance remains comparable to that on slightly noisy synthetic meshes and close to that on the ground-truth mesh. This indicates that MiNAF is robust not only to artificial noise but also to realistic reconstruction inaccuracies. Moreover, since mesh generation is a one-time, offline step, MiNAF retains the same runtime efficiency as image-based baselines while offering improved prediction accuracy even in the presence of reconstruction noise.

Context Information. Table 3 reports the impact of adding or removing context features in MiNAF. In almost all methods analyzed, removing the entire context from the input notably affects performance, resulting in a reduction of 31% and 28% in T60 for PreP and RanP, respectively. This performance pattern holds consistently across C50 and EDT metrics, underscoring the value of incorporating context in elevating RIR quality. Examining individual components, we observe that removing $\mathbf{n}$, which captures the local geometry around PoFHs, has the most substantial impact on all three metrics. Since surface normals provide an approximation of the local geometric structure where the last reflection occurs before a ray is received, this result highlights the importance of explicitly integrating geometric information for accurate RIR generation. Likewise, removing $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$, which describe the distribution of neighboring distances, also affects performance, although their impact is somewhat less pronounced than that of surface normals. Similarly, excluding occlusion counts, which reflect the distribution of all distances, results in performance degradation.

Time Embedding. Table 4 compares MiNAF with and without time embedded into the input context. In this configuration, rather than using element-wise embedding, the time vector after positional encoding is concatenated as an additional input to the neural network. The results indicate that incorporating time indices into the context yields substantial performance improvements, with $\approx$ 30% increase in T60 observed for both PreP and RanP. This enhancement supports the effectiveness of fusing time embeddings into context features in enabling the model to capture the scene's spatiotemporal characteristics more accurately, reinforcing their importance in improving RIR prediction quality.

Probe Density. We evaluate the effect of varying the number of emanating rays sampled at each Tx or Rx location. Since a non-linear projection is used to compress input features to a fixed dimension h, we adjust the widths of these projection layers to match the number of rays. This ensures that the projections serve strictly as dimension compressors, avoiding interference with the feature-to-RIR mapping task. As shown in Fig. 8, the performance of MiNAF generally declines as the ray count decreases. Reducing it from 1024 to 512 results in an 18% drop in T60 for PreP. This trend continues with fewer rays, with the most pronounced impact observed in PreP, where phase predictions are used. These results indicate that probing density plays a critical role in phase estimation, as phases are intricately linked to propagation paths, heavily influenced by local geometry.

Discussion and Conclusion

Discussion. Since local features vary by scene, MiNAF exhibits limitations in generalizability. A promising direction for future work is to develop a model that directly predicts RIRs on unseen scenes with few-shot fine-tuning. Another challenge is the complexity of feature collection when gathering RIRs at many unique locations. While SoundSpaces provides fixed locations for all $\langle Tx, Rx \rangle$ pairs, querying ray-mesh interactions for varied positions is computationally expensive. This could be mitigated by high-efficiency parallelism in context retrieval. Our method also demonstrates that explicit geometric guidance benefits the learning of acoustic fields. Incorporating more interpretable features into RIR models may further support real-world deployment.

Conclusion. We introduce MiNAF as a novel, fast, and accurate approach for RIR generation, which uses neural networks to represent a room's acoustic field from rough room meshes and sparse RIR measurements. By incorporating direct and explicit local geometry information as input features, MiNAF enhances the interpretability and reliability of the contextual data. Extensive experimental results show that MiNAF outperforms traditional encoding-based methods and neural implicit baselines in RIR fidelity, even with limited data. This makes MiNAF a promising tool for high-fidelity sound simulations in AR/XR and scientific research.

References

Allen, J. B.; and Berkley, D. A. 1979. Image method for efficiently simulating small-room acoustics. *The Journal of the Acoustical Society of America*, 65(4): 943–950.

Antonello, N.; De Sena, E.; Moonen, M.; Naylor, P. A.; and Van Waterschoot, T. 2017. Room impulse response interpolation using a sparse spatio-temporal representation of the sound field. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 25(10): 1929–1941.

Bai, X.; Zhou, J.; Ning, X.; and Wang, C. 2022. 3D data computation and visualization.

Barron, J. T.; Mildenhall, B.; Tancik, M.; Hedman, P.; Martin-Brualla, R.; and Srinivasan, P. P. 2021. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF international conference on computer vision*, 5855–5864.

Brunetto, A.; Hornauer, S.; and Moutarde, F. 2024. NeRAF: 3D Scene Infused Neural Radiance and Acoustic Fields. *CoRR*, abs/2405.18213.

Chen, C.; Jain, U.; Schissler, C.; Gari, S. V. A.; Al-Halah, Z.; Ithapu, V. K.; Robinson, P.; and Grauman, K. 2020. Soundspaces: Audio-visual navigation in 3d environments. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VI 16*, 17–36. Springer.

Chen, D.; Li, H.; Ye, W.; Wang, Y.; Xie, W.; Zhai, S.; Wang, N.; Liu, H.; Bao, H.; and Zhang, G. 2024. PGSR: Planar-based Gaussian Splatting for Efficient and High-Fidelity Surface Reconstruction. *arXiv preprint arXiv:2406.06521*.

Dufter, P.; Schmitt, M.; and Schütze, H. 2022. Position information in transformers: An overview. *Computational Linguistics*, 48(3): 733–763.

Engel, J.; Agrawal, K. K.; Chen, S.; Gulrajani, I.; Donahue, C.; and Roberts, A. 2019. Gansynth: Adversarial neural audio synthesis. *arXiv preprint arXiv:1902.08710*.

Funkhouser, T.; Tsingos, N.; Carlbom, I.; Elko, G.; Sondhi, M.; West, J. E.; Pingali, G.; Min, P.; and Ngan, A. 2004. A beam tracing method for interactive architectural acoustics. *The Journal of the acoustical society of America*, 115(2): 739–756.

Garbin, S. J.; Kowalski, M.; Johnson, M.; Shotton, J.; and Valentin, J. 2021. Fastnerf: High-fidelity neural rendering at 200fps. In *Proceedings of the IEEE/CVF international conference on computer vision*, 14346–14355.

Griffin, D.; and Lim, J. 1984. Signal estimation from modified short-time Fourier transform. *IEEE Transactions on acoustics, speech, and signal processing*, 32(2): 236–243.

Gumerov, N. A.; and Duraiswami, R. 2009. A broadband fast multipole accelerated boundary element method for the three dimensional Helmholtz equation. *The Journal of the Acoustical Society of America*, 125(1): 191–205.

He, Y.; Cherian, A.; Wichern, G.; and Markham, A. 2024. Deep Neural Room Acoustics Primitive. In Salakhutdinov, R.; Kolter, Z.; Heller, K.; Weller, A.; Oliver, N.; Scarlett, J.; and Berkenkamp, F., eds., *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, 17842–17857. PMLR.

Huang, X.; Li, H.; Yang, Z.; Wang, Z.; and Liang, X. 2024. 3D visibility-aware generalizable neural radiance fields for interacting hands. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 2400–2408.

International Organization for Standardization. 2006. Advanced Audio Coding (AAC). ISO/IEC 13818-7:2006.

Jin, W.; and Kleijn, W. B. 2015. Theory and design of multizone soundfield reproduction using sparse methods. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 23(12): 2343–2355.

Kingma, D. P. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Kurucz, P.; Baghaei, N.; Serafin, S.; and Klein, E. 2023. Enhancing Auditory Immersion in Interactive Virtual Reality Environments. In *2023 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*, 789–792. IEEE.

Lee, S.; Choi, J.; Kim, S.; Kim, I.-J.; and Cho, J. 2024. Few-shot neural radiance fields under unconstrained illumination. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 2938–2946.

Li, L.; Zhang, L.; Wang, Z.; Zhang, F.; Li, Z.; and Shen, Y. 2025. Representing Sounds as Neural Amplitude Fields: A Benchmark of Coordinate-MLPs and a Fourier Kolmogorov-Arnold Framework. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 24458–24466. AAAI Press. Available at: AAAI 2025.

Li, Z.; Müller, T.; Evans, A.; Taylor, R. H.; Unberath, M.; Liu, M.-Y.; and Lin, C.-H. 2023. Neuralangelo: High-fidelity neural surface reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8456–8465.

Liang, S.; Huang, C.; Tian, Y.; Kumar, A.; and Xu, C. 2023a. Av-nerf: Learning neural fields for real-world audio-visual scene synthesis. *Advances in Neural Information Processing Systems*, 36: 37472–37490.

Liang, S.; Huang, C.; Tian, Y.; Kumar, A.; and Xu, C. 2023b. Neural acoustic context field: Rendering realistic room impulse response with neural fields. *arXiv preprint arXiv:2309.15977*.

Luo, A.; Du, Y.; Tarr, M.; Tenenbaum, J.; Torralba, A.; and Gan, C. 2022. Learning neural acoustic fields. *Advances in Neural Information Processing Systems*, 35: 3165–3177.

Majumder, S.; Chen, C.; Al-Halah, Z.; and Grauman, K. 2022. Few-shot audio-visual learning of environment acoustics. *Advances in Neural Information Processing Systems*, 35: 2522–2536.

Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2021. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1): 99–106.

Ni, Z.; Yang, P.; Yang, W.; Ma, L.; and Kwong, S. 2024. ColNeRF: Collaboration for Generalizable Sparse Input Neural Radiance Field. 38(5): 4325–4333.

Niemeyer, M.; Barron, J. T.; Mildenhall, B.; Sajjadi, M. S.; Geiger, A.; and Radwan, N. 2022. Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5480–5490.

Ratnarajah, A.; Tang, Z.; and Manocha, D. 2020. IR-GAN: Room impulse response generator for far-field speech recognition. *arXiv preprint arXiv:2010.13219*.

Reiser, C.; Peng, S.; Liao, Y.; and Geiger, A. 2021. Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In *Proceedings of the IEEE/CVF international conference on computer vision*, 14335–14345.

Richard, A.; Dodds, P.; and Ithapu, V. K. 2022. Deep impulse responses: Estimating and parameterizing filters with deep networks. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 3209–3213. IEEE.

Rungta, A.; Schissler, C.; Rewkowski, N.; Mehra, R.; and Manocha, D. 2018. Diffraction kernels for interactive sound propagation in dynamic environments. *IEEE transactions on visualization and computer graphics*, 24(4): 1613–1622.

Savioja, L.; and Valimaki, V. 2003. Interpolated rectangular 3-D digital waveguide mesh algorithms with frequency warping. *IEEE transactions on speech and audio processing*, 11(6): 783–790.

Schissler, C.; Loftin, C.; and Manocha, D. 2017. Acoustic classification and optimization for multi-modal rendering of real-world scenes. *IEEE transactions on visualization and computer graphics*, 24(3): 1246–1259.

Schissler, C.; and Manocha, D. 2016. Interactive sound propagation and rendering for large multi-source scenes. *ACM Transactions on Graphics (TOG)*, 36(4): 1.

Schönberger, J. L.; and Frahm, J.-M. 2016. Structure-from-Motion Revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.

Schönberger, J. L.; Zheng, E.; Pollefeys, M.; and Frahm, J.-M. 2016. Pixelwise View Selection for Unstructured Multi-View Stereo. In *European Conference on Computer Vision (ECCV)*.

Singh, N.; Mentch, J.; Ng, J.; Beveridge, M.; and Drori, I. 2021. Image2reverb: Cross-modal reverb impulse response synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 286–295.

Straub, J.; Whelan, T.; Ma, L.; Chen, Y.; Wijmans, E.; Green, S.; Engel, J. J.; Mur-Artal, R.; Ren, C.; Verma, S.; et al. 2019. The Replica dataset: A digital replica of indoor spaces. *arXiv preprint arXiv:1906.05797*.

Su, K.; Chen, M.; and Shlizerman, E. 2022. Inras: Implicit neural representation for audio scenes. *Advances in Neural Information Processing Systems*, 35: 8144–8158.

Tang, Z.; Bryan, N. J.; Li, D.; Langlois, T. R.; and Manocha, D. 2020. Scene-aware audio rendering via deep acoustic analysis. *IEEE transactions on visualization and computer graphics*, 26(5): 1991–2001.

Tsingos, N.; Funkhouser, T.; Ngan, A.; and Carlbom, I. 2001. Modeling acoustics in virtual environments using the uniform theory of diffraction. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, 545–552.

Turki, H.; Ramanan, D.; and Satyanarayanan, M. 2022. Mega-nerf: Scalable construction of large-scale nerfs for virtual fly-throughs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12922–12931.

Wang, J.; Chen, M.; Karaev, N.; Vedaldi, A.; Rupprecht, C.; and Novotny, D. 2025. VGGT: Visual Geometry Grounded Transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.

Wang, S.; Leroy, V.; Cabon, Y.; Chidlovskii, B.; and Revaud, J. 2024. Dust3r: Geometric 3d vision made easy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20697–20709.

Xiph.Org Foundation. 2012. Xiph Opus. https://opus-codec.org/.

Yu, A.; Ye, V.; Tancik, M.; and Kanazawa, A. 2021. pixelnerf: Neural radiance fields from one or few images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4578–4587.

Yu, Z.; Sattler, T.; and Geiger, A. 2024. Gaussian Opacity Fields: Efficient Adaptive Surface Reconstruction in Unbounded Scenes. *ACM Transactions on Graphics*.

Zhang, H.; Wang, C.; Tian, S.; Lu, B.; Zhang, L.; Ning, X.; and Bai, X. 2023. Deep learning-based 3D point cloud classification: A systematic survey and outlook. *Displays*, 79: 102456.

Appendix

In this appendix, we provide additional details to support the main findings of our paper. Section A details the mathematical formulations used to project raw geometric and positional features into the model's latent space. Section B outlines key implementation and training settings, as well as expanded descriptions of baseline models used for comparison. In Section C, we report the complete set of T60, C50, and EDT metrics across all evaluated scenes, offering insight into `MiNAF`'s performance across room types and complexities. Section D compares the predicted and ground-truth RIR spectra in the frequency domain, emphasizing the model's strengths in magnitude estimation and challenges in phase reconstruction. Finally, Section E presents visualizations of time-domain RIRs and their associated Schroeder curves, providing further analysis of reverberation characteristics and phase reconstruction methods.

A. Projection Formulations

This section provides the detailed formulas for how each geometric and positional feature is processed as part of the fused context, as described in Section *MiNAF: Mesh-infused Neural Acoustic Field* in the main paper.

A.1 Feature Embedding Functions

After we probe the mesh and obtain the raw explicit geometry features, each raw feature is projected into a latent space of dimension h using a separate nonlinear mapping (MLP):

$$f_d : \mathbb{R}^N \to \mathbb{R}^h, \quad \mathbf{d}' = f_d(\mathbf{d}) \in \mathbb{R}^h \tag{1}$$

$$f_n : \mathbb{R}^{N\times 3} \to \mathbb{R}^h, \quad \mathbf{n}' = f_n(\mathbf{n}) \in \mathbb{R}^h \tag{2}$$

$$f_\mu : \mathbb{R}^N \to \mathbb{R}^h, \quad \boldsymbol{\mu}' = f_\mu(\boldsymbol{\mu}) \in \mathbb{R}^h \tag{3}$$

$$f_\sigma : \mathbb{R}^N \to \mathbb{R}^h, \quad \boldsymbol{\sigma}' = f_\sigma(\boldsymbol{\sigma}) \in \mathbb{R}^h \tag{4}$$

$$f_{occ} : \mathbb{R}^{N_\tau} \to \mathbb{R}^h, \quad \mathbf{occ}' = f_{occ}(\mathbf{occ}) \in \mathbb{R}^h \tag{5}$$

, where each nonlinear projection function is implemented as a two-layer MLP with ReLU activation.

A.2 Positional and Temporal Embedding Functions

As the transmitter and receiver positions are low-dimensional compared to the model's hidden space h, we first apply sinusoidal positional encoding $\gamma(\cdot)$ with L frequency bands to enrich their representations. The encoded vectors are then projected into the latent space via separate nonlinear mappings:

$$f_{p_{Tx}} : \mathbb{R}^{2\times L} \to \mathbb{R}^h, \quad \mathbf{p}'_{\mathbf{Tx}} = f_{p_{Tx}}(\gamma(\mathbf{p_{Tx}})) \in \mathbb{R}^h \tag{6}$$

$$f_{p_{Rx}} : \mathbb{R}^{2\times L} \to \mathbb{R}^h, \quad \mathbf{p}'_{\mathbf{Rx}} = f_{p_{Rx}}(\gamma(\mathbf{p_{Rx}})) \in \mathbb{R}^h \tag{7}$$

We apply the same encoding and projection procedure to the STFT time index t, enabling the model to capture temporal structure in the spectrum:

$$f_t : \mathbb{R}^{2\times L} \to \mathbb{R}^h, \quad t' = f_t(\gamma(t)) \in \mathbb{R}^h \tag{8}$$

The resulting vector t' is later fused with the context matrix via element-wise multiplication, as detailed in Subsection *Context Fusion* of the main paper.

B. Implementation Details

Context Feature Extraction. We set the number of rays emitted from each transmitter or receiver location to $N = 1024$. For proximity statistics, we compute the mean and standard deviation using $N_\theta = 8$ nearest neighboring rays. The global histogram features use $N_\tau = 8$ thresholds, with values $\delta_1, \ldots, \delta_{N_\tau}$ linearly spaced from 0.5 m to 10 m.

Neural Network Architecture. The core prediction module is a 5-layer fully connected MLP with hidden widths that decrease exponentially from 2048 to 512, ending in a 257-dimensional output matching the STFT column size. Notably, this is a much simpler architecture compared to those used in existing state-of-the-art neural implicit baselines.

Orientation and Channel Embeddings. We use 4 learnable vectors of length 64 to represent the hidden feature corresponding to the 4 discrete receiver orientations. Similarly, we assign 2 learnable 64-dimensional vectors for the 2 channels in our binaural setting. These embeddings are concatenated with the MLP input at every layer.

Positional Encoding. We use sinusoidal positional encoding with $L = 10$ frequency bands for all position-based inputs: time index, orientation, and channel index.

Training Setup. We train `MiNAF` using the Adam optimizer (Kingma 2014) with an initial learning rate of 5×10^{-4}, decaying by a factor of 0.9 per epoch. Batch size is set to 256. A fixed random seed (42) is used for all experiments to ensure reproducibility. All models are trained and evaluated on a single NVIDIA RTX 4080 GPU.

Experimental Runs. Each reported result, including ablation studies, comes from a single training run without additional ensembling or averaging.

Baseline Implementations. As stated in Section Experiments, we compare `MiNAF` against five recent neural implicit models and two traditional audio codecs methods. The details of them are below:

- **INRAS (2022)** (Su, Chen, and Shlizerman 2022): Models the acoustic field via separate neural modules dedicated to the transmitter, propagation path, and receiver. These modules encode respective physical and spatial characteristics, which are then synthesized by a listener module to generate the final RIR.
- **NAF (2022)** (Luo et al. 2022): Represents the scene with 2D learnable feature grids placed around each transmitter and receiver location. These grids encode local context, which is queried by a deep MLP to predict the RIR spectrum. The model avoids using explicit geometry, relying instead on end-to-end learned spatial embeddings.
- **NACF (2023)** (Liang et al. 2023b): Leverages RGB images captured along the scene's boundary to construct a global visual representation of the environment. This visual context is combined with $\langle Tx, Rx \rangle$ positional information to predict the RIR spectrum. While it benefits from visual cues, the model relies heavily on camera coverage and does not encode local geometry explicitly.
- **AV-NeRF (2023)** (Liang et al. 2023a): Extends NeRF by learning a joint audio-visual representation of the scene. It synthesizes both novel views and corresponding RIRs for arbitrary listener locations and orientations.

Room	MiNAF(GTP)			MiNAF(GTM+PreP)			MiNAF(PreP)			MiNAF(RanP)			MiNAF(GLim)		
	T60	C50	EDT	T60	C50	EDT	T60	C50	EDT	T60	C50	EDT	T60	C50	EDT
office 3	1.32	0.39	0.002_2	2.19	0.42	0.003_0	2.34	0.46	0.005_0	1.31	0.39	0.002_1	1.49	0.38	0.002_1
office 4	1.32	0.38	0.002_1	2.34	0.35	0.003_0	2.35	0.41	0.004_0	1.35	0.36	0.002_1	1.42	0.35	0.002_3
room 2	1.07	0.47	0.002_0	1.94	0.47	0.003_0	1.91	0.43	0.003_6	1.11	0.42	0.002_1	1.40	0.40	0.002_0
apartment 1	2.06	0.58	0.002_2	2.31	0.62	0.002_7	2.69	0.85	0.003_6	2.03	0.56	0.002_2	2.29	0.59	0.002_3
apartment 2	1.87	0.49	0.002_3	2.27	0.65	0.003_0	2.51	0.78	0.004_0	1.69	0.47	0.002_1	1.71	0.47	0.002_2
frl apartment 2	1.45	0.30	0.002_8	2.00	0.59	0.005_0	2.59	0.59	0.005_5	1.27	0.34	0.003_0	1.50	0.34	0.002_4
frl apartment 5	1.35	0.32	0.002_7	2.07	0.49	0.004_9	2.32	0.52	0.004_1	1.36	0.34	0.002_4	1.35	0.31	0.002_4
Average	1.49	0.42	0.002_3	2.17	0.50	0.003_4	2.35	0.58	0.004_2	1.42	0.41	0.002_2	1.59	0.41	0.002_2

Table 5: **MiNAF Evaluations on All Scenarios.** We show the raw T60, C50, and EDT measurements on different methods of evaluation here. We use *office 3* as a development scene and report the overall evaluation of MiNAF on the rest six scenarios.

The visual backbone is used to encode spatial appearance, which is fused with listener information to generate acoustic output.

- **NeRAF (2024)** (Brunetto, Hornauer, and Moutarde 2024): Extracts RGB and density features from a NeRF model pre-trained on multi-view images of the scene. These features are sampled at the transmitter and receiver positions and assembled into a compact feature grid. The resulting grid is then passed to a neural network to predict the RIR spectrum, integrating photometric and structural cues.
- **AAC (Advanced Audio Coding)** (International Organization for Standardization 2006): A standardized audio codec used to compress RIR waveforms into latent vectors. In our baseline, we encode RIRs using AAC, then perform linear or nearest-neighbor interpolation in the latent space to predict RIRs at new positions. This method provides a non-learning-based baseline for spatial generalization.
- **Opus (Xiph Foundation)** (Xiph.Org Foundation 2012): A real-time audio codec optimized for low-latency communication, adapted here for RIR modeling. We encode each RIR into a latent vector using Opus, and perform latent space interpolation to synthesize unseen RIRs.

C. Measurements of All Scenarios

While *office 3* serves as the primary development scene for MiNAF, we also evaluate its performance on six additional rooms from the SoundSpaces dataset (Chen et al. 2020): (1) two single-room environments with rectangular layouts, *office 4* and *room 2*; (2) two single-room environments with non-rectangular structures, *frl apartment 2* and *frl apartment 5*; and (3) two multi-room environments, *apartment 1* and *apartment 2*. Table 5 reports the T60, C50, and EDT metrics across all scenes using five different evaluation methods.

MiNAF maintains stable performance on T60 and EDT across most single-room scenarios, especially in standard rectangular layouts. Its accuracy remains consistent for rooms of similar size, demonstrating strong generalization in compact environments. This suggests that by accounting for local geometry, MiNAF effectively learns acoustic fields across a variety of room shapes, including standard rectangular configurations and more complex layouts with stairs or irregular corners.

A closer inspection of the two rectangular single-room scenes reveals an interesting pattern: in the smaller *room 2*, the model slightly overfits T60, likely due to the effect of the energy curve loss, which prioritizes reverberation decay at the expense of early reflection metrics like C50. This highlights the need for a better balance between the L1 loss, which emphasizes overall spectral fidelity, and the energy-based loss that drives accurate T60 and EDT estimation. While capturing energy decay is important, the model must also preserve the relative amplitude of early reflections.

Another notable finding from the cross-scenario analysis is that larger rooms tend to exhibit slower convergence on these metrics. The two multi-room settings, in particular, perform worse than the single-room scenarios. This discrepancy can be attributed to rays in larger scenes traversing longer distances or passing through open doorways into adjacent rooms before encountering walls. Such rays may introduce outliers into the local distance distributions, which could influence MiNAF's local context estimation. These outliers deviate from an accurate representation of the local geometry, highlighting a challenge in capturing the spatial intricacies of complex environments.

D. RIR Spectrum Visualizations

Fig. 9 illustrates a comparison of the STFT spectra in terms of log-magnitude and *instantaneous frequency* (IF), respectively, across representatives for each type of room in our dataset, including *room 2*, *office 3*, *frl apartment 2*, and *apartment 1*.

The log-magnitude plots reveal that the predicted spectra closely match the ground truth, suggesting that the model effectively captures the magnitude component of the acoustic response. Although the predicted IF spectrum roughly matches the ground truth, it fails to capture the subtle variations across individual time frames. This misalignment reinforces that MiNAF(PreP) underperforms compared to MiNAF(GTP), highlighting a limitation in the current network architecture's ability to accurately learn phase information. This is consistent with findings reported in NAF

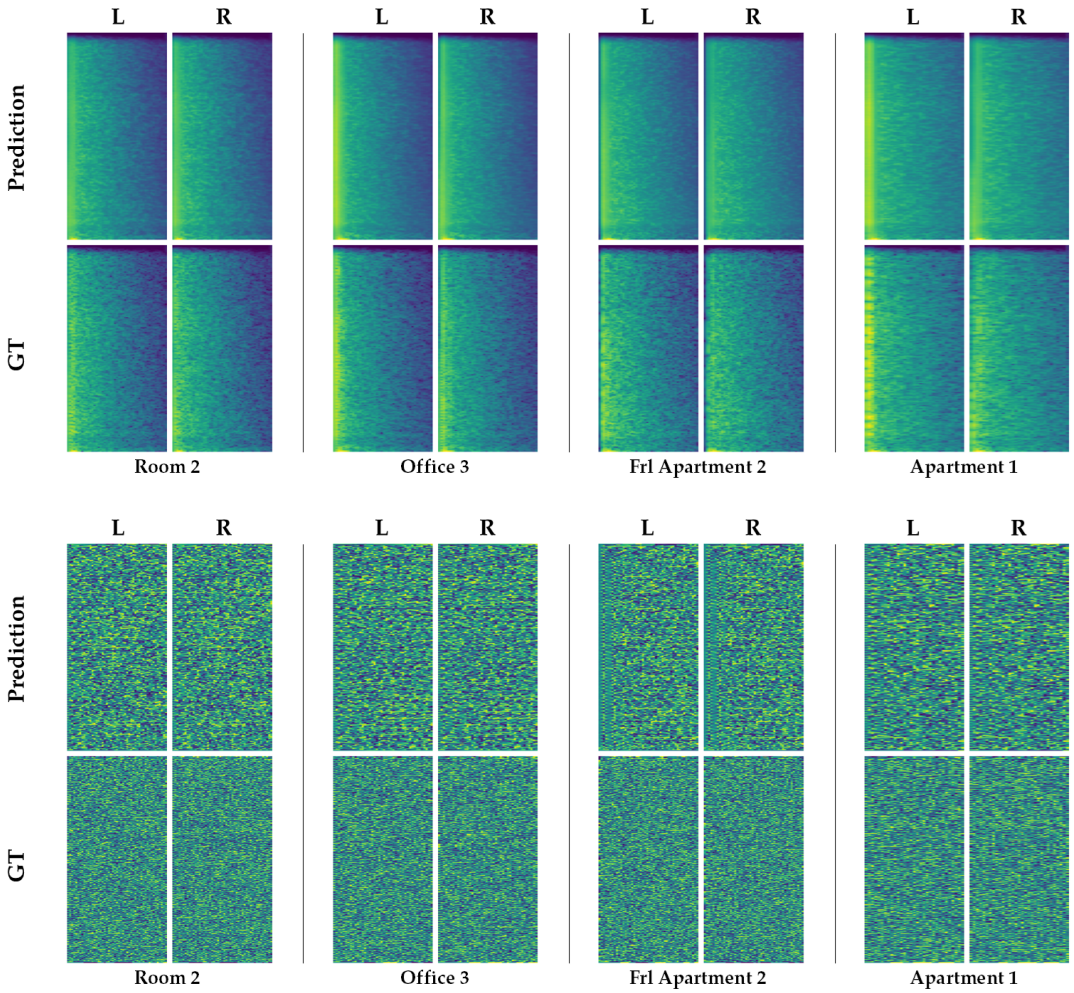

Figure 9: **Log-magnitude and IF Spectra**. Here we show the predicted and ground-truth log-magnitude and *instantaneous frequency* (IF) spectra of samples from four representative rooms: *room 2*, *office 3*, *frl apartment 2*, and *apartment 1*. The visualizations demonstrate that `MiNAF` excels in capturing the magnitude spectrum with high fidelity. However, learning phase information proves to be more challenging. While the predicted IF spectra generally align with the ground truth, they struggle to capture the coarse temporal variations across frames. This observation is consistent with our finding that `MiNAF`(PreP) underperforms compared to `MiNAF`(GTP).

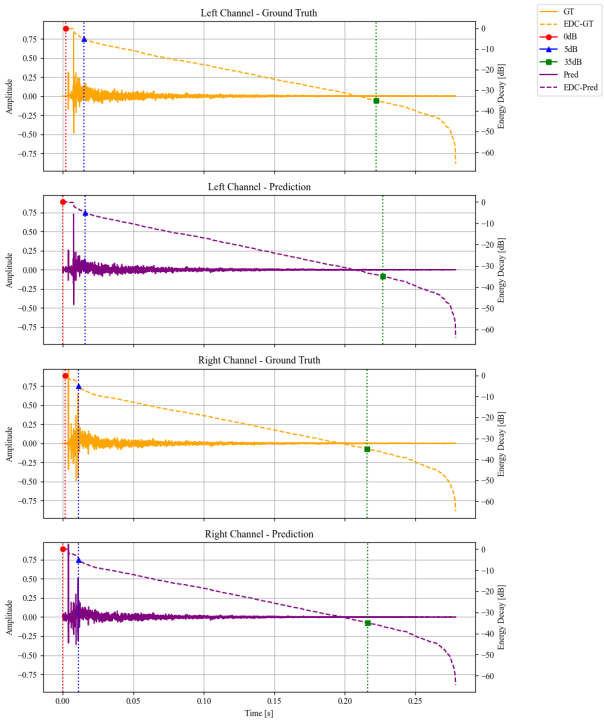

(a) MiNAF(GTP)

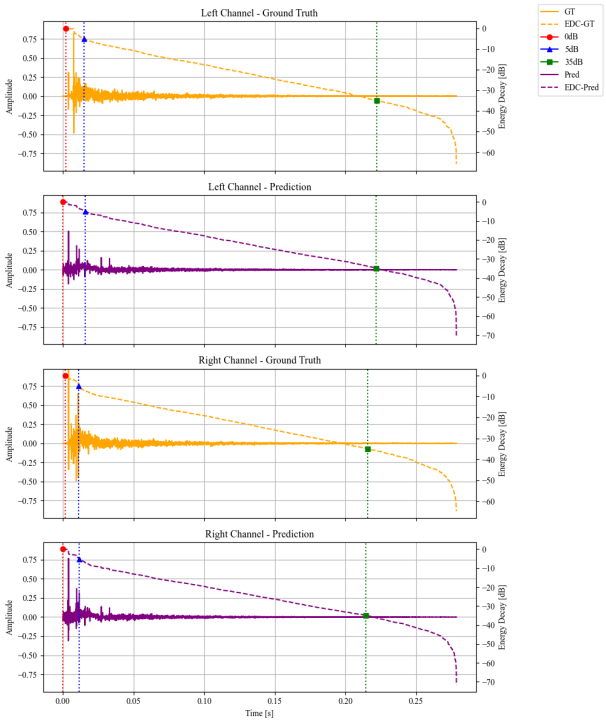

(b) MiNAF(GTM+PreP)

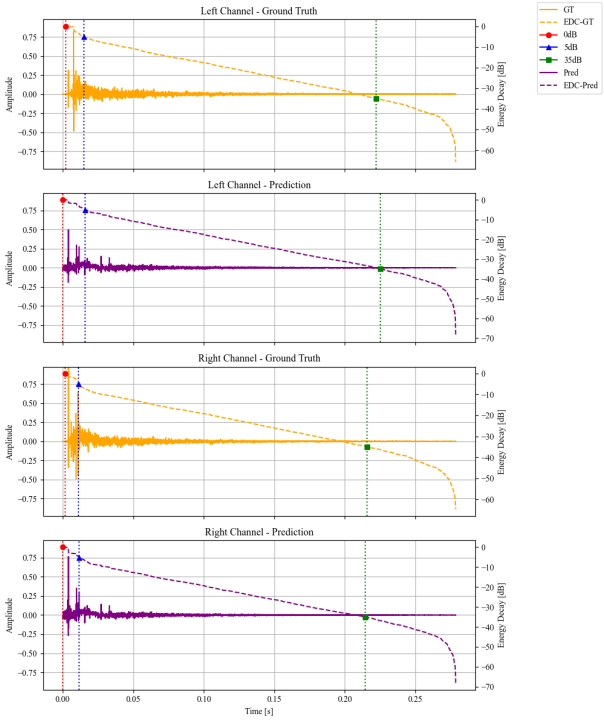

(c) MiNAF(PreP)

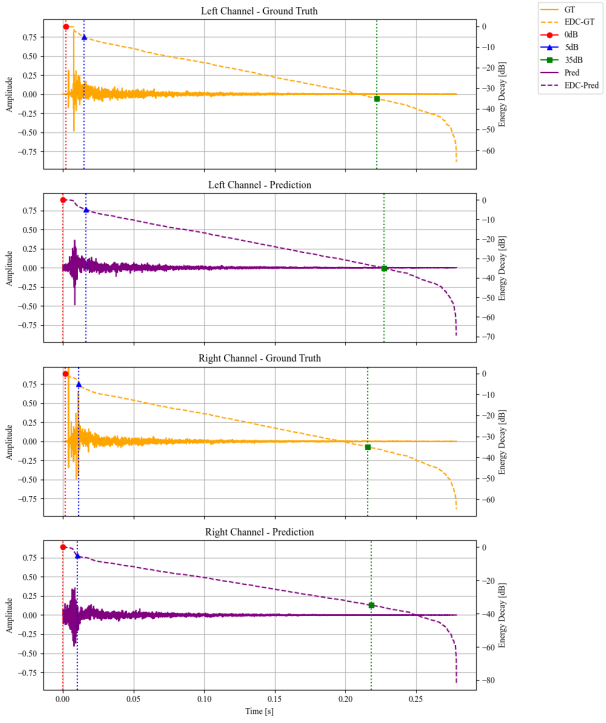

(d) MiNAF(GLim)

Figure 10: **RIR in Time Domain and its Energy Curve**. We show the RIR in the time domain and its energy decay for a random sample in *Room 2*. In each subplot, the dashed line shows the Schroeder Curve, which represents the RIR's energy decay. The red, blue, and green markers represent the point of maximum RIR energy, the start point to measure T60 (when energy decays for 5 dB), and the endpoint to measure T60 (when energy decays for another 30 dB), respectively.

(Luo et al. 2022), which noted the greater difficulty in learning phase components compared to magnitudes.

E. Time-domain RIR and Energy Decay

Fig. 10 visualizes the estimated time-domain RIRs and the corresponding Schroeder Curves for a test sample $\langle Tx, Rx \rangle$ pair in *room 2*, each reconstructed using `MiNAF`(GTP), `MiNAF`(GTM+PreP), `MiNAF`(PreP), and `MiNAF`(GLim), respectively. `MiNAF`(RanP) is not displayed because, while using random phases is adequate for evaluating predicted log magnitude spectra, it does not generate complete RIRs from the predictions.

From Fig. 10a, it is evident that as the log-magnitude predictions closely match the ground truth ones, a high resemblance for both left and right channels in the time-domain RIR constructed with the GTP method is observed. In Fig. 10b and Fig. 10c, although the maximum amplitude values of the predicted time-domain RIRs are similar to the ground truth, the overall amplitude range differs. Notably, the predicted minimum amplitude values are consistently higher than expected. This phenomenon, which can be described as a "hollow bottom" effect, results from inaccurate phase predictions. The predicted IF values do not vary as frequently as those in the ground truth, leading to less accurate phase predictions obtained by integrating these IF values.

In Fig. 10d, the results of RIR reconstruction using the Griffin-Lim algorithm (Griffin and Lim 1984) are presented. This algorithm estimates the RIR by initially guessing a random phase spectrum and iteratively computing the inverse and forward STFTs. After several iterations, the predicted phase tends to converge towards the ground truth. While this method is effective for normal audio pieces and serves as a valuable tool when only magnitude data is available, its application to RIR — a signal dominated by impulsive energy and structured reverberations governed by geometry — is limited. It is observed from Fig. 10d that although the Schroeder Curve of the prediction aligns closely with the ground truth, indicating a comparable T60, the time-domain shape of the predicted RIR significantly deviates from the ground truth. Therefore, while the Griffin-Lim algorithm (GLim) is utilized for comparison against NeRAF (Brunetto, Hornauer, and Moutarde 2024), its use is not recommended for real-world scenarios due to its limitations in accurately reconstructing RIRs.