RECONSTRUCTION OF SOUND FIELD THROUGH DIFFUSION MODELS

Federico Miotello, Luca Comanducci, Mirco Pezzoli,
Alberto Bernardini, Fabio Antonacci, Augusto Sarti

Dipartimento di Elettronica, Informazione e Bioingegneria - Politecnico di Milano, Milan, Italy

ABSTRACT

Reconstructing the sound field in a room is an important task for several applications, such as sound control and augmented (AR) or virtual reality (VR). In this paper, we propose a data-driven generative model for reconstructing the magnitude of acoustic fields in rooms with a focus on the modal frequency range. We introduce, for the first time, the use of a conditional Denoising Diffusion Probabilistic Model (DDPM) trained in order to reconstruct the sound field (SF-Diff) over an extended domain. The architecture is devised in order to be conditioned on a set of limited available measurements at different frequencies and generate the sound field in target, unknown, locations. The results show that SF-Diff is able to provide accurate reconstructions. We conduct a comparative analysis with two state-of-the-art baseline methods, one relying on kernel interpolation and the other on deep learning.

Index Terms— sound field reconstruction, diffusion neural network, space-time processing

1. INTRODUCTION

Sound field reconstruction is a relevant problem in the field of acoustic signal processing, especially when considering the modal frequency range, due to its importance in applications such as sound field control and room compensation [1]. Applications include, for example, the navigation of acoustic scenes [2]. The goal is to estimate the acoustic field over an extended area starting from the information provided by a limited set of sensors. In the literature, several techniques tackling the reconstruction of sound field can be found. Most solutions rely on either a parametric description [3–6] of the acoustic scene, or on models based on the solutions of the wave equations [7, 8] including plane waves [9], spherical waves [10–12] or equivalent sources [13, 14]. In [7, 8] and variations [15, 16], the reconstruction of the acoustic field is achieved using a kernel-interpolation based approach which exploit the solution of the Helmholtz equation as a physically motivated kernel.

A different class of techniques is based on deep learning, which has been widely applied in the field of acoustics [17, 18]. The main advantage of deep learning solutions is to adopt more sparse and irregular microphone array setups for the sound field reconstruction. The first learning-based approach was proposed in [19] and consisted of a U-Net architecture, which was applied in order to reconstruct the magnitude of the sound field with an approach similar to image inpainting. Similarly, in [20], the authors proposed a deep-prior approach to RIR reconstruction following the deep prior paradigm introduced for image inpainting [21]. This approach assumes that the structure of the CNN introduces an implicit prior regularizing the estimation of RIRs. Other solutions instead [22, 23], rely on the physical equation governing the sound propagation i.e., the wave equation as an alternative approach for improving the reconstruction. In [22], the authors introduce a physics-informed neural newtwork for the reconstruction of acoustic fields in the time domain. The model employed a SIREN architecture [24] trained with a physics-informed loss function including the wave equation. The model has been tested on time-domain sound fields captured through a uniform linear microphone array. Although effective, this techniques require to perform per-element training.

More recently, also generative models such as Generative Adversarial Networks (GANs) have been applied to sound field reconstruction problems [25]. In particular, in [25] three different generative models based on GANs are considered: a compressive sensing model, a conditional-GAN and a High-fidelity-GAN. The results in [25] proved that generative neural networks provide an effective approach for the reconstruction of acoustic field. Nonetheless, among generative models, Denoising Diffusion Probabilistic Models (DDPMs) [26] have recently gained interest, due to their enhanced synthesis capabilities and more stable training process with respect to GANs in different tasks [27,28]. As such, they have been applied to a variety of sound-related problems such as speech enhancement [29], speech super-resolution [30], vocoders [31] and audio inverse problems in general [28].

In this paper, motivated by the superior performance of DDPMs in different fields, we propose a DDPMs approach

for the reconstruction of room transfer functions (RTFs). We consider the sound field reconstruction as an image-to-image translation problem, where DDPMs have already been successfully applied [32]. Specifically, we focus on the sound field measured on a 2D plane using irregular microphone arrays. We give as input to the DDPM the magnitude of the computed sound field and inject noise where no microphones are deployed. Moreover, we condition the model on the embedding that encodes a considered frequency. Through a simulation campaign, we compare the performance of the proposed method with a kernel interpolation-based signal processing technique [8] and a learning-based approach [19], for configurations with different numbers of microphones, and demonstrate the benefits of applying DDPMs to the sound field reconstruction task.

The rest of the paper is organized as follows: in Sec. 2 we present the adopted data model and formalize the sound field reconstruction problem. In Sec. 3 we present the conditional DDPM for sound field reconstruction, while in Sec. 4 we present results aimed at demonstrating the effectiveness of the proposed technique. Finally, in Sec. 5, we draw some conclusions.

2. PROBLEM FORMULATION

2.1. Data model

Following the approach proposed in [19], let us consider a three-dimensional rectangular room $\mathcal{R} = [0, l_x] \times [0, l_y] \times [0, l_z]$, where $\{l_x, l_y, l_x\} \in \mathbb{R}^3_{>0}$ denote the length, width, and height of the room, respectively. The complex-valued frequency-domain sound field in position $\mathbf{r} \in \mathcal{R}$, can be computed using the Fourier transform as

$$P(\mathbf{r}, \omega) = \int_{\mathbb{R}} p(\mathbf{r}, t) e^{-j\omega t} \, dt, \tag{1}$$

where $\omega \in \mathbb{R}$ is the angular frequency, and $p(\mathbf{r}, t)$ denotes the acoustic sound field measured at position $\mathbf{r}$ at time t. Moreover, for the purpose of this article, we focus on the magnitude $|P(\mathbf{r}, \omega)|$ of the sound fields. In practice, room $\mathcal{R}$ is sampled using a regular rectangular grid $\mathcal{D}_o$, with fixed height $z_o \in [0, l_z]$ and defined as

$$\mathcal{D}_o = \left\{ \left(i\frac{l_x}{I-1}, j\frac{l_y}{J-1}, z_o \right) \right\}_{i,j}, \tag{2}$$

where $i = 0, 1, \ldots, I - 1$, and $j = 0, 1, \ldots, J - 1$ are the indexes of the sampled points in the grid, with $I, J \geq 2$.

2.2. Problem definition

We assume that a limited subset of room measurements $\mathcal{S}_o \subseteq \mathcal{D}_o$ are available. Sound field magnitude reconstruction can be defined as the problem of recovering the missing data $\{|P(\mathbf{r}, \omega)|\}_{\mathbf{r} \in \mathcal{D}_o \setminus \mathcal{S}_o}$, by exploiting the information in

the available observations $\{|P(\mathbf{r}, \omega)|\}_{\mathbf{r} \in \mathcal{S}_o}$ obtained using a limited number of irregularly deployed microphones.

Various techniques have been proposed in the literature to address sound field reconstruction from an under-sampled measurement set [7, 14, 19, 22, 23]. In general, this task can be interpreted in the framework of inverse problems, and a solution to the following minimization problem

$$\begin{aligned} \boldsymbol{\theta}^* = \underset{\boldsymbol{\theta}}{\operatorname{argmin}} \ J(\boldsymbol{\theta}) = \\ E\left(f_{\boldsymbol{\theta}}(\{|P(\mathbf{r}, \omega)|\}_{\mathbf{r} \in \mathcal{S}_o}), \{|P(\mathbf{r}, \omega)|\}_{\mathbf{r} \in \mathcal{D}_o \setminus \mathcal{S}_o}\right), \end{aligned} \tag{3}$$

where $f_{\boldsymbol{\theta}}(\{|P(\mathbf{r}, \omega)|\}_{\mathbf{r} \in \mathcal{S}_o})$ is a function that generates the estimated sound field using parameters $\boldsymbol{\theta}$ having access to available measurements, and $E(\cdot)$ is a data-fidelity term, e.g., the mean squared error (MSE), between the estimated data and the observations. It is worth noting that in (3), the evaluation of the reconstruction error is performed in the observed locations $\{\mathbf{r}\} \in \mathcal{S}_o$. However, f must be able to provide a meaningful estimate also in locations different from the available ones, i.e., $\{\mathbf{r}\} \in \mathcal{D}_o \setminus \mathcal{S}_o$. Therefore, the solution to the ill-posed problem (3) is constrained using regularization strategies. Typical techniques include compressed sensing frameworks based on assumptions about the signal model [14], such as plane and spherical wave expansions [13] or the RIRs structure [33], as well as deep learning approaches [19, 23].

3. PROPOSED SOLUTION

3.1. Diffusion Model for RTF reconstruction

In this work, we aim at solving the sound field reconstruction problem (3) in order to provide an estimate of the magnitude of a sound field as

$$|\hat{P}(\mathbf{r}, \omega)| = f_{\boldsymbol{\theta}}\left(\{|P(\mathbf{r}, \omega)|\}_{\mathbf{r} \in \mathcal{S}_o}\right), \tag{4}$$

where the function $f(\cdot)$ represents a neural network. In particular, we exploit the power of diffusion models, which have recently emerged as the cutting-edge technology in the field of deep learning-based generation, becoming the new state-of-the-art [26]. In particular, these are quickly replacing Generative Adversarial Networks (GANs) and Variational Autoencoders across various tasks and domains, thanks to their straightforward training process and the higher generation accuracy. Diffusion models employ an iterative denoising process to transform samples generated from a standard Gaussian distribution into samples that align with an empirical data distribution. In particular, we employ Palette [32], a conditional denoising diffusion model of the form $p(\boldsymbol{y}|\boldsymbol{x})$, in which the denoising process is conditioned by an input signal. For example, in the case of the proposed method, $\boldsymbol{x}$ would be the under-sampled sound field, while $\boldsymbol{y}$ the reconstructed one. Palette proved to be effective in many image-to-image translation tasks (e.g., inpainting, colorization, uncropping, etc.), outperforming already existing state-of-the-art methods.
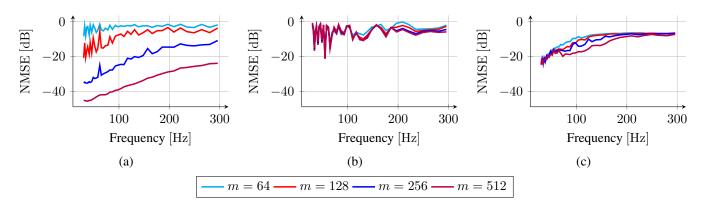
Fig. 1. Normalized Mean Squared Error (NMSE) for different number of microphones m measured over the reconstructed magnitude using the proposed method (a), Ueno et al. [8] (b) and Lluis et al. [19] (c).

The training process of Palette is carried out by optimizing an objective function in the form

$$\mathbb{E}_{(\boldsymbol{x},\boldsymbol{y})}\mathbb{E}_{\boldsymbol{\epsilon}\sim\mathcal{N}(0,I)}\mathbb{E}_{\gamma}\|f_{\theta}(\boldsymbol{x},\underbrace{\sqrt{\gamma}\boldsymbol{y}+\sqrt{1-\gamma}\boldsymbol{\epsilon}}_{\tilde{\boldsymbol{y}}},\gamma)-\boldsymbol{\epsilon}\|_2^2, \quad (5)$$

where $\boldsymbol{y}$ is a training output sample, $\tilde{\boldsymbol{y}}$ is its noisy version and $\boldsymbol{\epsilon}$ a noise vector. We train a neural network f with parameters $\boldsymbol{\theta}$ to denoise $\tilde{\boldsymbol{y}}$ given $\boldsymbol{y}$ and a noise level indicator γ [32].

3.2. Architecture

Palette exploits the U-Net architecture [34], which was originally proposed for multimodal medical image segmentation. In particular, the network has been enhanced with modifications proposed in recent works and is based on the 256×256 class-conditional U-Net model used in [35]. However, for the purpose of the task we are addressing and following the approach proposed [32], the architecture we are employing does not present class-conditioning but rather an additional conditioning of the input data via concatenation [36].

Even though, in principle, the method is able to reconstruct any arbitrary sound field, we focus on the reconstruction of the magnitude of room transfer functions (RTFs), which correspond to the Fourier transform of the impulse response of a room, measured in positions $\boldsymbol{r} \in \mathcal{D}_o$ and computed as in (1). The input to the network is composed by the concatenation of the RTFs magnitude matrix $\mathbf{P}$ and a frequency embedding $\mathbf{F}$.

Similarly to the approach proposed in [32], in correspondence of the unknown data points in $\mathbf{P}$, the RTF magnitude value is replaced with noise coming from a Gaussian distribution $\mathcal{N}(0,1)$. The role of $\mathbf{F}$ is to provide the extra conditioning, needed for the diffusion model to consistently learn how to reconstruct the magnitude of a room transfer function at a certain frequency, starting from a noisy version of it. During the training phase, we restrict the computation of the loss function, i.e., the MSE, only to the the available measures of the RTFs.

4. EXPERIMENTAL VALIDATION

4.1. Setup

Following the approach proposed in [19], we trained the proposed architecture using a simulated data set of 10000 randomly generated rectangular rooms, considering one frequency out of forty in the range $30 - 300$ Hz for each room. In particularly, we sample each room into a grid of 32×32 uniformly-spaced points (independently of the room dimensions). Thus, the RTFs are in the form $\mathbf{P} \in \mathbb{R}^{32 \times 32}$. For each room, we then randomly generated a binary mask $\mathbf{M} \in \{0,1\}^{32 \times 32}$, in order to select the number of microphones placed in the room. In particular, during training we considered masks selecting a number of microphones in the range $64 - 512$, which corresponds to the $6.25 - 50\%$ of the total sound field. It is worth noting that, as in [32], we do not directly provide the mask $\mathbf{M}$ as input to the network. Instead, we fill the masked measures with random Gaussian noise. Similarly to the training dataset, the test dataset is composed of 250 rooms randomly simulated rooms, considering forty different frequencies in the range $30 - 300$ Hz for each room. RTFs are approximated by using Green's function, which represents the solution as an infinite summation of room modes in the room [19]. The room dimensions are randomly sampled, considering a floor area in the range $20 - 60\,\text{m}^2$, and constant reverberation time T60 $= 0.6\,\text{s}$.

We trained the model for 1000 epochs, which we empirically verified to sufficient for the employed architecture to converge, using the Adam optimizer with a fixed 5×10^{-5} learning rate. During training, a linear noise schedule ranging from 1e^{-6} to 0.01 is applied over 2000 time-steps. Similarly, 2000 refinement steps are employed at inference, considering a linear noise schedule from 1e^{-4} to 0.09.

4.2. Results

We assessed the performance of the proposed sound field reconstruction method in terms of normalized mean squared er-
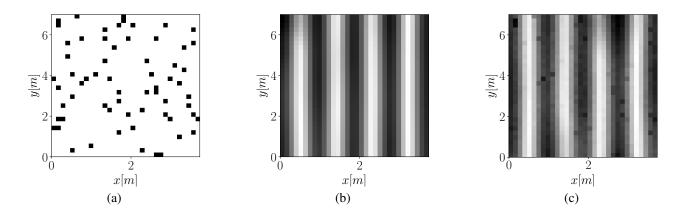
3

Fig. 2. Magnitude of the sound field in a randomly generated [3.7 m × 7 m × 26.1 m] room with a 98 Hz active source **s** positioned at $[0.9 \text{ m}, 0.3 \text{ m}, 2.4 \text{ m}]^T$, obtained using the proposed method (c), using the 64 active microphone configuration depicted in (a). Ground truth magnitude is shown in (b).

ror (NMSE) between ground truth and reconstructed RTFs. In particular, we computed the NMSE as

$$\text{NMSE} = 10 \log_{10} \left(\frac{1}{N} \sum_{i=0}^{N} \frac{\|\hat{P}_i(\mathbf{r}, \omega) - P_i(\mathbf{r}, \omega)\|_2^2}{\|P_i(\mathbf{r}, \omega)\|_2^2} \right) \quad (6)$$

where N is the number of samples in the testing dataset. The performances of the method are computed with respect to the number of microphones placed in each room. Thus, for each room we consider 64, 128, 256, 512 microphones, randomly arranged in space.

Figure 1(a) shows the NMSE value with respect to frequency, for reconstructions performed using the proposed SF-Diff method and following the setup described in Section 4.1. As expected, the number of available measurements highly affects the reconstruction error, which ranges from a minimum value of −8.35 dB in the 64 mics configuration, to a minimum value of −45.39 dB in the 512 mics configuration. Having access to more information about the RTF in a certain room, the model is able to better reconstruct the pressure values in the unknown positions. Also, in all configurations, the error increases with the frequency value. This is due to the fact that RTFs at higher frequencies present complicated magnitude patterns, graphically, which are more difficult to reconstruct. In fact, the spatial distribution of room modes becomes more intricate, resulting in peaks and nulls in the RTF at specific locations. At lower frequencies, instead, the room response is more uniform and thus easier to reconstruct.

Figure 2(c) shows an example of RTF reconstruction performed using SF-Diff, considering the measurement setup represented in Figure 2(a) and ground-truth RTF represented in Figure 2(b). As it can be seen, the method is able to provide a coherent reconstruction, in which most characteristics of the ground-truth image are present, leading to an NMSE value of −11.72 dB.

As a comparison, we performed RTF reconstruction us-

ing the methods proposed in [8] and in [19], leveraging the same dataset employed in our testing procedure. Specifically, the approach proposed in [8] addresses the problem by exploiting the kernel ridge regression with the constraint of the Helmholtz equation. On the other hand, [19] exploits a super-resolution approach based on deep learning.

Results are shown in Figure 1(b) and Figure 1(c), respectively. In scenarios where 64 measurements are available, both baseline methods demonstrate superior reconstruction performance, with respect to SF-Diff. In particular, [8] achieves a minimum NMSE of −15.78 dB at 34 Hz, while [19] obtains a minimum NMSE of −22.13 dB at the same frequency. However, when evaluating setups with an increased number of available measurements (128, 256 and 512 microphones), the reconstruction error does not diminish as prominently as observed in the case of SF-Diff, indicating a limited performance characteristic under varying measurement conditions.

5. CONCLUSION

In this paper we have, to the best of our knowledge, proposed the first application of Denoising Diffusion Probabilistic Models to the problem of sound field reconstruction. Specifically, we consider a pre-existing architecture used for image-inpainting and adapt accordingly. We consider sound fields on a grid positioned on a two-dimensional plane, where measurements are available only for a limited number of arbitrarily positioned microphones. Following the DDPM procedure, we inject noise in the positions for which no measurements available and through the proposed model, we learn to denoise them in order to reconstruct the ground truth sound field. Through an experimental simulation campaign we compare the performance of the proposed technique with a signal processing-based kernel interpolation method and a learning-based method, demonstrating the effectiveness of the DDPM-based approach. The obtained results encourage us to further develop the application of DDPMs to the problem of sound field reconstruction by adapting the technique in order to perform in more challenging scenarios.

6. REFERENCES

[1] J. M. Schmid, E. Fernandez-Grande, M. Hahmann, C. Gurbuz, M. Eser, and S. Marburg, "Spatial reconstruction of the sound field in a room in the modal frequency range using bayesian inference," *The Journal of the Acoustical Society of America*, vol. 150, no. 6, pp. 4385–4394, 2021.

[2] J. G. Tylka and E. Y. Choueiri, "Fundamentals of a parametric method for virtual navigation within an array of ambisonics microphones," *J. Audio Eng. Soc.*, vol. 68, no. 3, pp. 120–137, 2020.

[3] O. Thiergart, G. Del Galdo, M. Taseska, and E. A. P. Habets, "Geometry-based spatial sound acquisition using distributed microphone arrays," *Trans. Acoust., Speech, Signal Process.*, vol. 21, no. 12, pp. 2583–2594, 2013.

[4] M. Pezzoli, F. Borra, F. Antonacci, S. Tubaro, and A. Sarti, "A parametric approach to virtual miking for sources of arbitrary directivity," *IEEE/ACM Trans. Acoust., Speech, Signal Process.*, vol. 28, pp. 2333–2348, 2020.

[5] L. McCormack, A. Politis, R. Gonzalez, T. Lokki, and V. Pulkki, "Parametric ambisonic encoding of arbitrary microphone arrays," *Trans. Acoust., Speech, Signal Process.*, vol. 30, pp. 2062–2075, 2022.

[6] M. Pezzoli, F. Borra, F. Antonacci, A. Sarti, and S. Tubaro, "Reconstruction of the virtual microphone signal based on the distributed ray space transform," in *26th European Signal Processing Conference (EU-SIPCO)*, pp. 1537–1541, IEEE, 2018.

[7] N. Ueno, S. Koyama, and H. Saruwatari, "Sound field recording using distributed microphones based on harmonic analysis of infinite order," *IEEE Signal Processing Letters*, vol. 25, no. 1, pp. 135–139, 2017.

[8] N. Ueno, S. Koyama, and H. Saruwatari, "Kernel ridge regression with constraint of helmholtz equation for sound field interpolation," in *Int. Workshop Acoust. Signal Enhanc.*, pp. 1–440, IEEE, 2018.

[9] W. Jin and W. B. Kleijn, "Theory and design of multizone soundfield reproduction using sparse methods," *Trans. Acoust., Speech, Signal Process.*, vol. 23, no. 12, pp. 2343–2355, 2015.

[10] M. Pezzoli, M. Cobos, F. Antonacci, and A. Sarti, "Sparsity-based sound field separation in the spherical harmonics domain," in *Int. Conf. Acoust. Speech Signal Process*, IEEE, 2022.

[11] F. Borra, I. D. Gebru, and D. Markovic, "Soundfield reconstruction in reverberant environments using higher-order microphones and impulse response measurements," in *Int. Conf. Acoust. Speech Signal Process*, pp. 281–285, IEEE, 2019.

[12] A. Fahim, P. N. Samarasinghe, and T. D. Abhayapala, "Sound field separation in a mixed acoustic environment using a sparse array of higher order spherical microphones," in *2017 Hands-free Speech Communications and Microphone Arrays (HSCMA)*, pp. 151–155, IEEE, 2017.

[13] S. Koyama and L. Daudet, "Sparse representation of a spatial sound field in a reverberant environment," *IEEE Journal of Selected Topics in Signal Processing*, vol. 13, no. 1, pp. 172–184, 2019.

[14] S. Damiano, F. Borra, A. Bernardini, F. Antonacci, and A. Sarti, "Soundfield reconstruction in reverberant rooms based on compressive sensing and image-source models of early reflections," in *2021 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, IEEE, 2021.

[15] J. G. Ribeiro, N. Ueno, S. Koyama, and H. Saruwatari, "Region-to-region kernel interpolation of acoustic transfer functions constrained by physical properties," *Trans. Acoust., Speech, Signal Process.*, vol. 30, pp. 2944–2954, 2022.

[16] J. G. Ribeiro, S. Koyama, and H. Saruwatari, "Kernel interpolation of acoustic transfer functions with adaptive kernel for directed and residual reverberations," *arXiv preprint arXiv:2303.03869*, 2023.

[17] M. J. Bianco, P. Gerstoft, J. Traer, E. Ozanich, M. A. Roch, S. Gannot, and C.-A. Deledalle, "Machine learning in acoustics: Theory and applications," *The Journal of the Acoustical Society of America (JASA)*, vol. 146, no. 5, pp. 3590–3628, 2019.

[18] M. Olivieri, M. Pezzoli, F. Antonacci, and A. Sarti, "A physics-informed neural network approach for nearfield acoustic holography," *Sensors*, vol. 21, no. 23, 2021.

[19] F. Lluís, P. Martínez-Nuevo, M. Bo Møller, and S. Ewan Shepstone, "Sound field reconstruction in rooms: Inpainting meets super-resolution," *The Journal of the Acoustical Society of America*, vol. 148, no. 2, pp. 649–659, 2020.

[20] M. Pezzoli, D. Perini, A. Bernardini, F. Borra, F. Antonacci, and A. Sarti, "Deep prior approach for room impulse response reconstruction," *Sensors*, vol. 22, no. 7, p. 2710, 2022.

[21] D. Ulyanov, A. Vedaldi, and V. Lempitsky, "Deep image prior," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 9446–9454, 2018.

[22] 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," in *Forum Acusticum 2023*, EAA, 2023.

[23] X. Karakonstantis and E. Fernandez-Grande, "Generative adversarial networks with physical sound field priors," *The Journal of the Acoustical Society of America*, vol. 154, no. 2, pp. 1226–1238, 2023.

[24] V. Sitzmann, J. Martel, A. Bergman, D. Lindell, and G. Wetzstein, "Implicit neural representations with periodic activation functions," *Advances in Neural Information Processing Systems*, vol. 33, pp. 7462–7473, 2020.

[25] E. Fernandez-Grande, X. Karakonstantis, D. Caviedes-Nozal, and P. Gerstoft, "Generative models for sound field reconstruction," *J. Acou. Soc. Am.*, vol. 153, no. 2, pp. 1179–1190, 2023.

[26] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," *Advances in neural information processing systems*, vol. 33, pp. 6840–6851, 2020.

[27] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-resolution image synthesis with latent diffusion models," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10684–10695, 2022.

[28] E. Moliner, J. Lehtinen, and V. Välimäki, "Solving audio inverse problems with a diffusion model," in *Int. Conf. Acoust. Speech Signal Process*, pp. 1–5, IEEE, 2023.

[29] H. Yen, F. G. Germain, G. Wichern, and J. Le Roux, "Cold diffusion for speech enhancement," in *Int. Conf. Acoust. Speech Signal Process*, pp. 1–5, IEEE, 2023.

[30] C.-Y. Yu, S.-L. Yeh, G. Fazekas, and H. Tang, "Conditioning and sampling in variational diffusion models for speech super-resolution," in *Int. Conf. Acoust. Speech Signal Process*, pp. 1–5, IEEE, 2023.

[31] N. Takahashi, M. Kumar, Y. Mitsufuji, *et al.*, "Hierarchical diffusion models for singing voice neural vocoder," in *Int. Conf. Acoust. Speech Signal Process*, pp. 1–5, IEEE, 2023.

[32] C. Saharia, W. Chan, H. Chang, C. Lee, J. Ho, T. Salimans, D. Fleet, and M. Norouzi, "Palette: Image-to-image diffusion models," in *ACM SIGGRAPH 2022 Conference Proceedings*, pp. 1–10, 2022.

[33] E. Zea, "Compressed sensing of impulse responses in rooms of unknown properties and contents," *Journal of Sound and Vibration*, vol. 459, p. 114871, 2019.

[34] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th Int. Conf., Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pp. 234–241, Springer, 2015.

[35] A. Q. Nichol and P. Dhariwal, "Improved denoising diffusion probabilistic models," in *International Conference on Machine Learning*, pp. 8162–8171, PMLR, 2021.

[36] C. Saharia, J. Ho, W. Chan, T. Salimans, D. J. Fleet, and M. Norouzi, "Image super-resolution via iterative refinement," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 4, pp. 4713–4726, 2022.