

Generative Models for Helmholtz Equation Solutions: A Dataset of Acoustic Materials

Riccardo Fosco Gramaccioni^{1*}, Christian Marinoni^{1*}, Fabrizio Frezza¹,
Aurelio Uncini¹, and Danilo Comminiello¹

¹Department of Information Engineering, Electronics and Telecommunications (DIET),
Sapienza University of Rome, Rome, Italy,
{riccardofosco.gramaccioni; christian.marinoni}@uniroma1.it

Abstract—Accurate simulation of wave propagation in complex acoustic materials is crucial for applications in sound design, noise control, and material engineering. Traditional numerical solvers, such as finite element methods, are computationally expensive, especially when dealing with large-scale or real-time scenarios. In this work, we introduce a dataset of 31,000 acoustic materials, named HA30K, designed and simulated solving the Helmholtz equations. For each material, we provide the geometric configuration and the corresponding pressure field solution, enabling data-driven approaches to learn Helmholtz equation solutions. As a baseline, we explore a deep learning approach based on Stable Diffusion with ControlNet, a state-of-the-art model for image generation. Unlike classical solvers, our approach leverages GPU parallelization to process multiple simulations simultaneously, drastically reducing computation time. By representing solutions as images, we bypass the need for complex simulation software and explicit equation-solving. Additionally, the number of diffusion steps can be adjusted at inference time, balancing speed and quality. We aim to demonstrate that deep learning-based methods are particularly useful in early-stage research, where rapid exploration is more critical than absolute accuracy.

Index Terms—acoustic material, helmholtz equation solver, data-driven, diffusion.

I. INTRODUCTION

The Helmholtz equation provides a fundamental mathematical framework for modeling steady-state wave behavior in complex acoustic materials, enabling the prediction of sound pressure fields under different configurations. However, solving the Helmholtz equation numerically is computationally expensive, particularly for high-resolution simulations or real-time applications [1].

Traditional numerical methods, such as the finite element method (FEM) and finite difference method (FDM), provide accurate solutions but require significant computational resources, making them impractical for large-scale or interactive applications. Recent advancements in deep learning have shown promise in accelerating acoustic [2], [3] and physics-based simulations by learning data-driven mappings between material configurations and their corresponding physical responses. Such approaches have been explored in fluid dynamics [4], heat transfer [5], and electromagnetism [6],

*Equal contribution.




Fig. 1. Samples from the HA30K dataset. The red dot represents the sound source location. Black squares represent the obstacles.

but their application to acoustic wave propagation remains relatively underexplored.

In this work, we introduce the *Helmholtz Acoustic 30K* (HA30K) dataset, designed to bridge the gap between computational acoustics and deep learning. The dataset consists of 31,000 2D acoustic material configurations generated using FreeFEM, a finite element solver for partial differential equations. Each sample includes a square domain with one to six square obstacles, each made of different materials such as rubber, wood, or metal. The acoustic source is modeled as

a Gaussian source, randomly placed within the non-occupied regions of the domain. For each configuration, we provide both the material distribution image and the corresponding pressure field p solution obtained by solving the Helmholtz equation. Some examples are depicted in Figure 1. The dataset can be downloaded from Zenodo¹.

The Helmholtz equation, which governs the steady-state behavior of acoustic waves, depends on key parameters such as the wave number k , which encapsulates frequency and material properties, and appropriate boundary conditions to account for wave reflections and absorptions. In our dataset, we enforce conditions that simulate real-world wave interactions, including absorbing boundaries to prevent artificial reflections and Dirichlet conditions on obstacles to model rigid surfaces. The resulting pressure fields provide a rich representation of wave propagation patterns across different material configurations.

To evaluate the potential of deep learning for acoustic simulation, we use as a baseline Stable Diffusion with ControlNet, since diffusion models have demonstrated excellent capabilities in generating high quality images [7]. We trained the model to generate accurate pressure field predictions from material configuration images, offering a fast and data-driven alternative to conventional numerical solvers.

This paper aims to provide a benchmark for future research in applying deep learning to acoustic material simulation. By releasing our dataset and baseline models, we hope to inspire further exploration of deep learning-driven solutions for computational acoustics, enabling faster and more scalable simulations for real-world applications.

Our contribution can be summarized as follows:

- We introduce a novel dataset of 31,000 2D acoustic material configurations, including pressure field solutions obtained using FreeFEM.
- We test our dataset on a baseline model derived from Stable Diffusion with ControlNet to predict Helmholtz equation solutions from material configurations.
- We perform an evaluation based on some of the most common objective metrics for image quality estimation and a comparative analysis between traditional numerical solvers and deep learning-based approaches, highlighting their advantages and limitations for acoustic wave simulation.

The rest of the paper is organized as follows. Section II presents the related works, Section III the proposed dataset structure, while in Section IV we discuss the baseline model and we validate the obtained experimental results. Finally, conclusions are drawn in Section V.

II. RELATED WORKS

Deep learning has emerged as a powerful tool for advancing the design and analysis of complex materials, including acoustic materials. Recent studies demonstrate the potential of generative models and deep learning-based inverse design techniques in discovering novel material properties and structures.

Diffusion models, in particular, have gained popularity as physics-informed approaches capable of generating physically plausible designs while maintaining computational efficiency [8], [9]. These methods have been successfully applied to inverse-design tasks for nonlinear mechanical metamaterials and multi-modal resonant structures [10].

For acoustic materials, deep learning has proven effective in optimizing surfaces for sound absorption. Different studies have used machine learning techniques to design materials with high degrees of freedom [11], [12], enabling the automated discovery of novel acoustic structures. The application of inverse design has also extended to phononic crystals and multishape metamaterials, allowing for tailored wave propagation characteristics [13], [14].

In wider terms, advancements in data-driven methods have demonstrated the feasibility of using deep learning to optimize materials for specific mechanical, optical, and electromagnetic properties [15]–[18]. Several works have explored deep learning as a tool for accelerating the discovery of new material configurations, facilitating rapid generation and evaluation of candidate structures [19], [20]. Notably, generative models have been employed for mechanistic-based learning and on-demand inverse design of metamaterials, highlighting the increasing role of deep learning in material science [21], [22].

Building on this foundation, our work presents a dataset of simulated acoustic materials designed using the Helmholtz equation, providing a benchmark for evaluating deep learning-based methods for acoustic simulation. Unlike previous studies focusing on specific material categories or optimization tasks, our dataset aims to facilitate the development of generalizable models capable of predicting acoustic wave behavior in diverse material configurations.

III. HA30K DATASET

In this section, we introduce the HA30K dataset—a collection of 31,000 numerical samples derived from solutions of the Helmholtz equation. These samples capture a wide range of acoustic scenarios, making the dataset a valuable resource for research in wave propagation, numerical simulation, and machine learning applications in acoustics. All numerical solutions were computed using FreeFEM, an open-source high-level multiphysics finite element solver.

A. Theoretical Background

The propagation of acoustic waves is modeled by the Helmholtz equation. In its strong form, the equation is given by

$$\Delta p + k^2 p = f \quad \text{in } \Omega, \quad (1)$$

where p denotes the pressure field, k is the wave number, and f is a given source term. To ensure a well-posed problem, we impose boundary conditions, such as the Sommerfeld radiation condition on the boundary of the domain Γ_1 , representing the boundary of the domain Ω :

$$\frac{\partial p}{\partial n} - ikp = 0 \quad \text{on } \Gamma_1, \quad (2)$$

¹<https://zenodo.org/records/15683385>

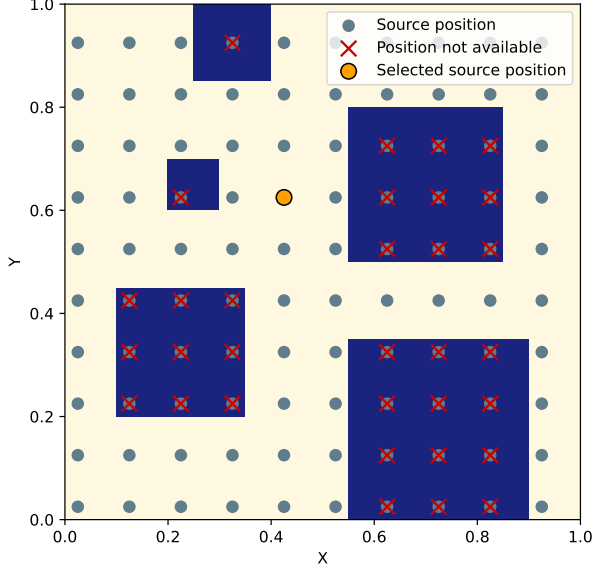


Fig. 2. An example showing the possible locations of the sound source given the position of obstacles (in blue).

and a Robin boundary condition on obstacle surfaces $\Gamma_{\text{obstacles}}$:

$$\frac{\partial p}{\partial n} + iZp = 0 \quad \text{on } \Gamma_{\text{obstacles}}, \quad (3)$$

where Z is the impedance parameter. For numerical simulations, we derive the weak formulation by multiplying the strong form by a test function q , integrating over the domain Ω , and applying integration by parts. The resulting weak formulation is:

$$\underbrace{\int_{\Omega} (\nabla p \cdot \nabla q - k^2 p q) dx}_{\text{Weak formulation of the Helmholtz eq.}} + \underbrace{\int_{\Omega} f q dx}_{\text{Gaussian source}} - \underbrace{\int_{\Gamma_1} i k p q ds}_{\text{Sommerfeld boundary condition}} + \underbrace{\int_{\Gamma_{\text{obstacles}}} i Z p q ds}_{\text{Robin boundary condition}} = 0.$$

The Sommerfeld condition is imposed to simulate free radiation (i.e., to avoid reflections at the domain boundaries), while the Robin condition on obstacles models impedance effects. By varying the parameter Z , we can simulate different obstacle materials (i.e., $Z = 150$ for foam, $Z = 600$ for rubber, $Z = 1500$ for wood, $Z = 1 \times 10^6$ for metal).

B. Domain definition

The generation of each sample in the HA30K dataset involves the definition of the domain and its characteristics. The domain is a unit square discretized with a 256×256 mesh. Within this domain, the speed of sound c is set according to the material properties of the medium, influencing the wave

TABLE I
SPEEDS OF SOUND c GIVEN DIFFERENT DOMAIN MATERIALS.

Material	Speed of sound (m/s)
Rubber	60
Air at 20°C	343
Lead	1210
Gold	3240
Glass	4540
Aluminum	6320

propagation characteristics. Table I summarizes the values of c given the different materials. Similarly, the frequency of the sound source can have a value between 100 and 4000 Hz (with increments of 100 Hz). Additionally, obstacles are modeled as squares and are randomly placed in non-overlapping configurations. Their dimensions are sampled from the interval $[0.1, 0.4]$ (with a step of 0.05) and their number varies from 1 to 6. Larger obstacles are assigned a higher probability when the total number of obstacles is lower, where on average 32% of the domain surface occupied by obstacles. This flexible domain setup enables the simulation of a wide range of acoustic scenarios.

C. Implementation with FreeFEM

FreeFEM was employed to solve the weak formulation of the Helmholtz equation using the finite element method. Its flexibility allowed us to incorporate complex boundary conditions and heterogeneous material properties, thereby enabling the generation of a large number of simulations under varying physical and geometrical configurations. The solution p to the Helmholtz equation is complex-valued. For each simulation, the results are stored as a set of values $[X, Y, \text{Re}(p), \text{Im}(p)]$. From the distribution of obstacles we derive the source images of the dataset, which represent with black pixels the obstacles and with white pixels the free portion of the domain. For the target images, we convert the real component of the pressure field ($\text{Re}(p)$) into RGB images by applying the viridis colormap from matplotlib. The resulting dataset consists of 31,000 pairs of source and target images, each with dimensions of 256×256 pixels.

IV. EXPERIMENTAL RESULTS

A. Baseline

As a baseline for evaluating the effectiveness of deep learning methods in predicting the Helmholtz equation solutions, we employ Stable Diffusion with ControlNet [23]. We use official repository, available on GitHub² and the provided checkpoints. The model is trained on our dataset, which is split into training (80%), validation (15%), and test (5%) sets.

The input to ControlNet consists of the image representing the material domain, which includes the surface and obstacles, providing spatial information about the acoustic environment.

²<https://github.com/Illyasviel/ControlNet>

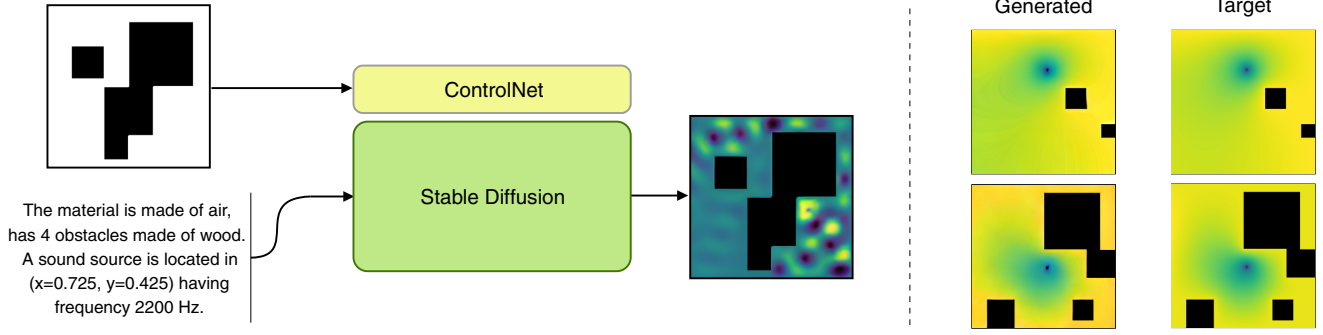


Fig. 3. **Baseline architecture:** The baseline architecture takes the topological structure of the domain along with a textual description of the acoustic properties and outputs an estimate of the pressure field. The target is the pressure field obtained by solving the Helmholtz equation.

Additionally, global parameters such as frequency, material of the surface, material of the obstacles, number of obstacles, frequency, and position of the sound source are encoded as a single text string having structure: “The material is made of $\{material_domain\}$, has $\{num_obstacles\}$ obstacles made of $\{material_obstacles\}$. A sound source is located in ($x=\{s_x\}$, $y=\{s_y\}$) having frequency $\{source_freq\}$ Hz”. The text string is used as cross-attention conditioning for the model. The target is the 256x256 pixel image derived by the real components of the pressure values $\text{Re}(p)$, as resulted from the simulations.

During training, the Stable Diffusion weights remain frozen, and we train only the ControlNet parameters to adapt the diffusion process to our specific dataset. This setup ensures that the model learns to generate pressure field solutions while leveraging the powerful generative capabilities of diffusion models. A block diagram of the baseline architecture is proposed in Figure 3.

The model is trained on a single 48 GB Nvidia RTX A6000. During training, we used a batch size of 4 with a learning rate kept fixed at $1e-5$ for a total of 100k diffusion steps, where each step is defined by 1000 timesteps. We use the Stable Diffusion 2.1³ checkpoint and we keep the model frozen, training only the ControlNet parameters. Consequently, we use the original values for all other hyperparameters of the model.

B. Objective Evaluation

To demonstrate the performance of the baseline used to test our dataset in predicting the solutions of the Helmholtz equation, we employ the following objective evaluation metrics:

- **Mean Squared Error (MSE):** this metric measures the average squared difference between the predicted and ground truth pressure field images. A lower MSE indicates a higher accuracy in the reconstruction of the pressure distribution.
- **Fréchet Inception Distance (FID):** FID compares the statistical distribution of generated and real images in a feature space extracted from a pre-trained neural network.

³<https://huggingface.co/stabilityai/stable-diffusion-2-1-base/tree/main>

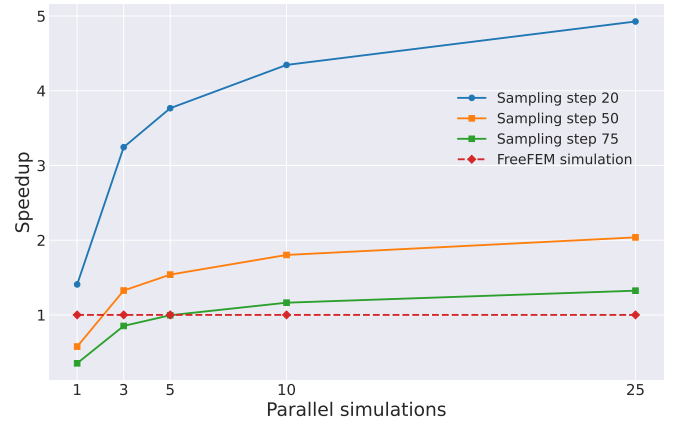


Fig. 4. Speedup obtained using parallel calculation with our baseline compared to sequential simulations with FreeFEM.

Lower FID values indicate that the generated pressure fields are perceptually closer to the ground truth.

- **Structural Similarity Index Measure (SSIM):** SSIM assesses the perceptual similarity between predicted and reference images by considering luminance, contrast, and structural information. Higher SSIM values indicate better structural preservation in the generated pressure fields.

These metrics provide a comprehensive evaluation of both numerical accuracy and perceptual quality for the acoustic material simulations.

C. Discussion

The results of our study highlight the potential of deep learning methods for accelerating material simulations, particularly for pressure field estimation using the Helmholtz equation. Traditional numerical solvers, such as FreeFEM, while highly accurate, require high computational costs and have limited scalability, as each simulation must be performed sequentially. In contrast, deep learning models allow for parallelized inference, leveraging batch processing and multi-GPU acceleration to significantly reduce computation time.

TABLE II
OBJECTIVE EVALUATION FOR DIFFERENT SAMPLING STEPS

Sampling steps	FID ↓	SSIM ↑	MSE ↓
20	53.31	0.657	0.1156
50	40.93	0.662	0.1290
75	43.62	0.669	0.1163

Generating high-quality, labeled data is a fundamental step toward making this kind of methods a viable alternative to traditional solvers. By providing a dataset of 31,000 simulated materials, each with its corresponding pressure field solution, we enable the training and evaluation of data-driven models capable of learning complex physical relationships.

To assess the efficiency of our approach, we evaluated the performance of our model using the DDIM sampler with different batch sizes (1, 3, 5, 10, and 25) and various sampling steps (20, 50, 75, and 100). Our analysis demonstrates that the generated pressure fields maintain high fidelity to the ground truth, as indicated by the objective evaluation metrics (**MSE**, **FID**, and **SSIM**) reported in Table II. We observed that increasing the batch size and optimizing the number of sampling steps allows for a significant reduction in inference time while maintaining competitive accuracy, as shown in Figure 4.

By parallelizing inference and tuning sampling parameters, our approach enables rapid and scalable material simulations, making it a viable alternative for large-scale acoustic material analysis. Furthermore, our dataset can be used as a benchmark for future research.

V. CONCLUSIONS

In this work, we introduced the HA30K dataset, a dataset for deep learning-based acoustic material simulation, specifically addressing pressure field estimation via the Helmholtz equation. By leveraging deep generative models, we demonstrated the advantages of using data-driven approaches instead of traditional numerical solvers to significantly accelerate simulation times through parallelized inference. Our experimental results confirm the high fidelity of the generated pressure fields compared to ground truth simulations, while also showcasing the efficiency of batch processing and optimized sampling strategies.

Our findings highlight the potential of deep learning for material simulation, paving the way for more scalable and efficient acoustic analysis methods. Furthermore, the dataset we provide establishes a new base for future research in data-driven material design and inverse modeling, fostering advancements in physics-informed deep learning applications.

ACKNOWLEDGMENT

This work was supported by the European Union under the Italian National Recovery and Resilience Plan (NRRP) of NextGenerationEU, partnership on “National Centre for HPC, Big Data and Quantum Computing” (CN00000013 - Spoke 6: Multiscale Modelling & Engineering Applications).

REFERENCES

- [1] A. A. Oberai and P. M. Pinsky, “A numerical comparison of finite element methods for the helmholtz equation,” *Journal of computational Acoustics*, vol. 8, no. 01, pp. 211–221, 2000.
- [2] R. F. Gramaccioni, C. Marinoni, C. Chen, A. Uncini, and D. Commiello, “L3das23: Learning 3d audio sources for audio-visual extended reality,” *IEEE Open Journal of Signal Processing*, 2024.
- [3] C. Marinoni, R. F. Gramaccioni, C. Chen, A. Uncini, and D. Commiello, “Overview of the l3das23 challenge on audio-visual extended reality,” in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–2, IEEE, 2023.
- [4] J. Tompson, K. Schlachter, P. Sprechmann, and K. Perlin, “Accelerating eulerian fluid simulation with convolutional networks,” in *International conference on machine learning*, pp. 3424–3433, PMLR, 2017.
- [5] S. Cai, Z. Wang, S. Wang, P. Perdikaris, and G. E. Karniadakis, “Physics-informed neural networks for heat transfer problems,” *Journal of Heat Transfer*, vol. 143, no. 6, p. 060801, 2021.
- [6] M. Raissi, P. Perdikaris, and G. E. Karniadakis, “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations,” *Journal of Computational physics*, vol. 378, pp. 686–707, 2019.
- [7] L. Sigillo, R. F. Gramaccioni, A. Nicolosi, and D. Commiello, “Ship in sight: Diffusion models for ship-image super resolution,” in *2024 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8, IEEE, 2024.
- [8] J.-H. Bastek, W. Sun, and D. Kochmann, “Physics-informed diffusion models,” in *The Thirteenth International Conference on Learning Representations*, 2025.
- [9] J.-H. Bastek and D. M. Kochmann, “Inverse-design of nonlinear mechanical metamaterials via video denoising diffusion models,” *Nat. Mac. Intell.*, vol. 5, pp. 1466–1475, 2023.
- [10] S. Dedoncker, C. Donner, L. Taenzer, and B. van Damme, “Generative inverse design of multimodal resonant structures for locally resonant metamaterials,” 2023.
- [11] K. Donda, Y. Zhu, A. Merkel, S.-W. Fan, L. Cao, S. Wan, and B. Assouar, “Ultrathin acoustic absorbing metasurface based on deep learning approach,” *Smart Materials and Structures*, vol. 30, 2021.
- [12] K. Donda, Y. Zhu, A. Merkel, S. Wan, and B. Assouar, “Deep learning approach for designing acoustic absorbing metasurfaces with high degrees of freedom,” *Extreme Mechanics Letters*, 2022.
- [13] D. M. J. Dykstra and C. Coulais, “Inverse design of multishape metamaterials,” 2023.
- [14] C. S. Ha, D. Yao, Z. Xu, C. Liu, H. Liu, D. Elkins, M. Kile, V. S. Deshpande, Z. Kong, M. Bauchy, and X. R. Zheng, “Rapid inverse design of metamaterials based on prescribed mechanical behavior through machine learning,” *Nature Communications*, vol. 14, 2023.
- [15] J. Liu, D. Zhang, D. Yu, M. Ren, and J. Xu, “Machine learning powered ellipsometry,” *Light, Science & Applications*, vol. 10, 2021.
- [16] Y. Liu, G. P. Wang, J. B. Pendry, and S. Zhang, “All-angle reflectionless negative refraction with ideal photonic weyl metamaterials,” *Light, Science & Applications*, vol. 11, 2022.
- [17] R. Gramaccioni, C. Marinoni, F. Frezza, A. Uncini, D. Commiello, et al., “Inverse design of thin-film metamaterials with a transformer-based approach,” in *Proc. WIRN 2024*, 2024.
- [18] Muhammad, J. Kennedy, and C. W. Lim, “Machine learning and deep learning in phononic crystals and metamaterials a review,” *Materials Today Communications*, 2022.
- [19] S. Sarkar, A. Ji, Z. Jermain, R. Lipton, M. L. Brongersma, K. Dayal, and H. Y. Noh, “Physics-informed machine learning for inverse design of optical metamaterials,” *Advanced Photonics Research*, vol. 4, 2023.
- [20] A. Schneuing, Y. Du, C. Harris, A. R. Jamasb, I. Igashov, W. Du, T. L. Blundell, P. Li’o, C. P. Gomes, M. Welling, M. M. Bronstein, and B. E. Correia, “Structure-based drug design with equivariant diffusion models,” *ArXiv*, vol. abs/2210.13695, 2022.
- [21] L. Wang, Y.-C. Chan, F. Ahmed, Z. Liu, P. Zhu, and W. Chen, “Deep generative modeling for mechanistic-based learning and design of metamaterial systems,” *ArXiv*, vol. abs/2006.15274, 2020.
- [22] Z. Wang, A. T. C. Chen, Z.-X. Xu, J. Yang, B. Liang, and J. chun Cheng, “On-demand inverse design of acoustic metamaterials using probabilistic generation network,” *Science China Physics, Mechanics & Astronomy*, vol. 66, pp. 1–8, 2023.
- [23] L. Zhang, A. Rao, and M. Agrawala, “Adding conditional control to text-to-image diffusion models,” 2023.