

RADIOLUNADIFF: ESTIMATION OF WIRELESS NETWORK SIGNAL STRENGTH IN LUNAR TERRAIN

Paolo Torrado¹, Anders Pearson¹, Jason Klein², Alexander Moscibroda³ and Joshua Smith¹

¹University of Washington, ²Cornell University, ³The Bear Creek School

{patorrad, amp206}@uw.edu, jak532@cornell.edu, alexmosci@outlook.com, jrs@cs.washington.edu

ABSTRACT

In this paper, we propose a novel physics-informed deep learning architecture for predicting radio maps over lunar terrain. Our approach integrates a physics-based lunar terrain generator, which produces realistic topography informed by publicly available NASA data, with a ray-tracing engine to create a high-fidelity dataset of radio propagation scenarios. Building on this dataset, we introduce a triplet-UNet architecture, consisting of two standard UNets and a diffusion network, to model complex propagation effects. Experimental results demonstrate that our method outperforms existing deep learning approaches on our terrain dataset across various metrics. The project website is available at: <https://radiolunadiff.github.io/>.

Index Terms—Radio maps, diffusion models, Helmholtz, lunar communications

1. INTRODUCTION

The NASA proposed lunar exploration communication framework, LunaNet [1], will require prediction of wireless network properties to facilitate reliable and robust communication. In a lunar environment, LunaNet is designed as a cooperative network infrastructure composed of multiple interoperable elements, where nodes communicate with neighbors and relay data at the appropriate link or network layers. Of particular importance are Radio Maps (RMs), which describe the spatial distribution of wireless channel attenuation for a given frequency, environment, and transmitter location.

For example, RMs can help mission planners create attenuation-aware trajectories or determine the shortest path to reconnection in cases of signal loss or degradation. Relaying rovers can also leverage RMs to maintain adequate signal quality when transmitting large data payloads, reducing energy consumption and communication latency.

Navigation is a key pillar of LunaNet, enabling the reliable fulfillment of safety, situational awareness, communication, and scientific objectives. Similar to GPS-RTK systems on Earth—where base stations provide correction data for centimeter-level accuracy—lunar rovers could use RMs to maintain robust data links with landers or beacons acting




Fig. 1: (a) Predicted k^2 probability map over lunar surface, (b) Corresponding model prediction of radio map at 5.8 GHz.

as reference nodes. In addition, RMs can support network reliability for critical services such as space weather monitoring and timely dissemination of solar eruption alerts. Such information is essential for protecting both equipment and astronauts.

Extensive prior research on wireless signal propagation has focused on predicting radio maps (RMs) in urban environments using deep neural networks. Leveraging simulation-based datasets, these studies have demonstrated that neural architectures can accurately estimate signal strength in complex scenes. More recently, researchers have employed physics-informed neural networks (PINNs) based on the Helmholtz equation to better model complex electromagnetic wave propagation, such as diffraction and interference, in dense environments.

In this work, we propose a novel triplet UNet PINN architecture that integrates the Helmholtz equation by modeling wave propagation over the terrain surface. Using a custom lunar terrain generator, we construct a simulation dataset and train a UNet to predict spatial variations in the square wave number, k^2 , enabling the identification of electromagnetic field discontinuities induced by terrain geometry. This predicted k^2 -map is then used as input to a second UNet chained to a diffusion model, both tasked with predicting the RM at either a low or high communication frequency, effectively incorporating terrain-induced propagation effects into the signal strength estimation process. We make the following contributions:

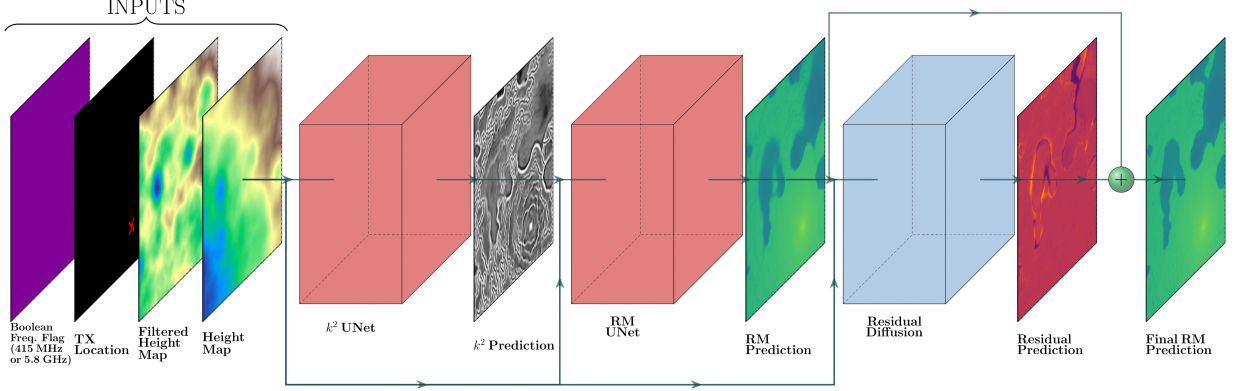


Fig. 2: Proposed system architecture.

Note: RMs shown in figure represent 5.8 GHz.

- A lunar terrain generator that synthesizes realistic topography based on publicly available lunar elevation data.
- A novel physics-informed, diffusion-based model architecture for predicting radio maps over varying terrain at multiple frequencies.

2. RELATED WORKS

Traditional wireless propagation models, such as Rayleigh and TWDP Fading, estimate RMs in a statistical sense, sacrificing knowledge of the location of wireless hotspots and regions of attenuation for a simplified model that can only output a probability. While these models are effective in communications system design, applications such as localization and navigation require more detailed location information, leading to the development of methods that spatially sample and interpolate wireless channel attenuation [2]. Unfortunately, this dense spatial sampling is often infeasible, leading recent research to focus on sampling-free methods like simulation.

Deterministic models such as ray tracers for radio propagation modeling, based on approximations of the Maxwell equations, have become increasingly precise and are widely used both commercially and in research [3]. However, even with advances in efficiency and GPU acceleration, ray tracers remain computationally complex, leading researchers to explore deep learning for RM approximation.

One of the first successful approaches to deep learning-based RM estimation was *RadioUNet*, which applied the UNet architecture to predict RMs using a map of building structures and the location of the transmitter [4]. Another notable adaptation, *PMNet*, modified the traditional convolutional blocks of the UNet to incorporate atrous convolutions for improved feature extraction [5]. More recently, following trends in image generation, highly precise RM construction models have leveraged diffusion architecture [6] [7] [8].

Most existing deep learning approaches for radio map (RM) construction have been developed for planar environments, typically using simplified geometric descriptions of urban layouts represented by binary building maps. In con-

trast, this work focuses on non-planar RMs defined over complex lunar terrain, where elevation and surface irregularities play a dominant role in wave propagation. Prior efforts in this domain have been limited, often relying on simplified models such as the two-ray approximation for RM construction [9], which cannot fully capture the diffraction, refraction, shadowing, and scattering effects inherent to rugged lunar topography.

3. METHODS

3.1. Helmholtz Equation and Laplace-Beltrami Operator

In this work we apply the Helmholtz Equation applied over a surface

$$\Delta_{\mathcal{M}} E(\mathbf{x}) + k^2(\mathbf{x})E(\mathbf{x}) = -f(\mathbf{x}), \quad (1)$$

where $E(\mathbf{x})$ is the complex electric field, $k^2(\mathbf{x})$ is the wave number, $f(\mathbf{x})$ is the source term, and $\Delta_{\mathcal{M}}$ is the Laplace-Beltrami operator defined over the Riemannian manifold \mathcal{M} . The Laplace-Beltrami operator is a generalization of the Laplace operator that works on arbitrary surfaces, as opposed to the Laplace operator which can only be used in Euclidean spaces [10] [11].

3.2. System Overview

We propose a pipeline of three cascaded models—two UNets followed by a diffusion network (Fig. 2)—to predict a radio map $I_{RM} \in \mathbb{R}^{H \times W}$. The input to the pipeline is denoted by \mathbf{X} and consists of: (i) a height-map image $I_{HM} \in \mathbb{R}^{H \times W}$, (ii) a high-pass filtered height map $I_{FM} \in \mathbb{R}^{H \times W}$, (iii) a one-hot encoded image indicating the transmitter location $I_{Tx} \in \{0, 1\}^{H \times W}$, and (iv) a boolean image flag $I_{Hz} \in \{0, 1\}^{H \times W}$ which is all 0 to indicate a 415 MHz target frequency and all 1 for 5.8 GHz.

Next, we derive the square wave number map $k^2(\mathbf{x})$, represented in discrete form as $I_{KM} \in \mathbb{R}^{H \times W}$, by adapting *RadioDiff- k^2* [8] for surface data by rearranging (1). In practice, the source term $f(\mathbf{x})$ is determined by the transmitter location, which is not on the surface the electric field is defined on and therefore it is ignored.

Let $I_E \in \mathbb{R}^{H \times W}$ denote the complex electric field on the terrain surface. Since the radio map I_{RM} is a scalar field proportional to the squared field magnitude, $I_{RM}(i, j) \propto |I_E(i, j)|^2$, we approximate k^2 as:

$$I_{KM}(i, j) = \frac{-\Delta_{\mathcal{M}} I_E(i, j)}{I_E(i, j)}. \quad (2)$$

In practice, we utilize PyTorch and its auto-differentiable capabilities to derive the I_{KM} [12]. The final step is to convert the KMs into binary values with pixels with values < 0 being set to 1, and all non-negative pixels being set to 0.

Using I_{HM} , I_{FM} , I_{Tx} , and I_{Hz} as inputs, the first UNet, f_{θ_1} , estimates \hat{I}_{KM} according to the formulation in (2). The second UNet, f_{θ_2} , takes as input these three maps together with the predicted \hat{I}_{KM} from the first UNet to estimate \hat{I}_{RM} . Finally, the third model, a Denoising Diffusion Probabilistic Model (DDPM) denoted f_{θ_3} estimates the residual map \hat{r} . The final prediction is reconstructed as:

$$\hat{I}_{RM}^{\text{final}} = \hat{I}_{RM, \theta_2} + \hat{r}_{\theta_3}. \quad (3)$$

Stage 1: Physics-informed diffusion for wave number prediction. The first UNet, f_{θ_1} , estimates the binary square wave number map \hat{I}_{KM} from \mathbf{X} and is trained with Binary Cross-Entropy (BCE) loss against the ground-truth I_{KM} ,

$$\mathcal{L}_1 = -\frac{1}{HW} \sum_{i=1}^H \sum_{j=1}^W \left[\text{BCE}(I_{KM, ij}, \hat{I}_{KM, ij}) \right]. \quad (4)$$

which compels the model to output a probability map \hat{I}_{KM} . Each pixel in this map represents the model's confidence that a k^2 discontinuity is present. While this map could be binarized using a 0.5 threshold, we retain the continuous probabilities as a "soft" input for the subsequent stage, thereby preserving the model's uncertainty.

Stage 2: Initial Radio Map Prediction. The second UNet, f_{θ_2} , produces an initial estimate of the radio map \hat{I}_{RM, θ_2} from the inputs \mathbf{X} and the predicted square wave number map \hat{I}_{KM, θ_1} . The model is trained by minimizing the Mean Squared Error (MSE) loss:

$$\mathcal{L}_2 = \frac{1}{HW} \sum_{i=1}^H \sum_{j=1}^W (I_{RM, ij} - f_{\theta_2}(\mathbf{X}, \hat{I}_{KM, \theta_1})_{ij})^2. \quad (5)$$

Stage 3: Diffusion for Residual Refinement. In the final state, we utilize the architecture DDPM as proposed by Ho et al. [13] and used in [6] to act as a refinement model. The model, f_{θ_3} , which refines the initial radio map by predicting the residual $r = I_{RM} - \hat{I}_{RM, \theta_2}$. Its UNet core U_{θ_3} takes a noisy residual r_t at timestep t together with conditioning inputs $c = (\mathbf{X}, \hat{I}_{KM, \theta_1}, \hat{I}_{RM, \theta_2})$, and predicts the velocity v_{pred} . The model is trained with a hybrid objective that combines velocity prediction [14] and residual reconstruction:

$$\mathcal{L}_3 = \mathcal{L}_v + \lambda \mathcal{L}_{\text{recon}}, \quad (6)$$

where the MSE between the true velocity v and the network's prediction, v_{pred} , is

$$\mathcal{L}_v = \mathbb{E}_{r, \epsilon, t, c} [\|v - U_{\theta_3}(r_t, t, c)\|^2], \quad (7)$$

and the MSE between the true residual r and the predicted residual, \hat{r}_{θ_3} , is

$$\mathcal{L}_{\text{recon}} = \mathbb{E}_{r, \epsilon, t, c} [\|r(\mathbf{x}) - \hat{r}_{\theta_3}\|^2]. \quad (8)$$

The predicted residual is obtained from the velocity estimate as

$$\hat{r}_{\theta_3} = \sqrt{\alpha_t} r_t - \sqrt{1 - \alpha_t} v_{\text{pred}}. \quad (9)$$

Overall objective. The final loss is:

$$\mathcal{L}_{\text{total}} = \underbrace{\mathcal{L}_1}_{\text{Stage 1: } k^2 \text{ map}} + \underbrace{\mathcal{L}_2}_{\text{Stage 2: radio map}} + \underbrace{\mathcal{L}_3}_{\text{Stage 3: residual}}. \quad (10)$$

UNET Architecture, Training and Inference. f_{θ_1} is implemented using a custom UNet, f_{θ_2} is implemented by modifying the architecture from *PMNet* [5], and f_{θ_3} is implemented via architecture provided by the Hugging Face Diffusers python library [15]. Each stage is trained separately, with gradients from \mathcal{L}_2 blocked from flowing into θ_1 by the stop-gradient on \hat{I}_{KM} , and gradients from \mathcal{L}_3 blocked from flowing into θ_2 and θ_1 by the stop-gradient on \hat{I}_{RM} . At inference, we first generate \hat{I}_{KM} using f_{θ_1} and then condition f_{θ_2} on \hat{I}_{KM} to produce \hat{I}_{RM} , and finally condition f_{θ_3} on \hat{I}_{KM} and \hat{I}_{RM} to produce \hat{r} , finally reconstructing $\hat{I}_{RM}^{\text{final}}$.

4. EXPERIMENTAL SETUP AND DATASETS

4.1. Lunar Terrain Height Map Generation

Lunar terrain is dominated by impact craters and the geomorphic consequences of cratering. We generate synthetic lunar terrains following Appendix C of Cai and Fa [16] with numerical details following [17] [18] [19]. Specifically, we alternate between (i) simulating a subset of cratering events corresponding to a fraction of the target surface age and (ii) applying diffusion over that same interval. Iterating these two steps until the total age is reached yields complex terrains containing both fresh, sharply defined craters and older, highly degraded craters.

4.2. Sionna Simulation Dataset

We built our dataset using Sionna-RT, an open source ray-tracer for radio propagation modeling [3]. Based on LunaNet specifications we choose frequencies 5.8GHz and 415MHz [20]. Ray-tracing methods are inherently limited by the finite number of discretized rays, which often results in radio maps with null pixels. To mitigate this, we randomize the seed and recompute each map instance multiple times. Bilinear interpolation was used to complete 415 MHz maps, which exhibited scattered data gaps, while a static fill was applied to




Fig. 3: Qualitative comparison of generated radiomaps at 415 MHz.

5.8 GHz maps, which contained large contiguous voids and few scattered points.

The electromagnetic properties of lunar regolith—relative permittivity ϵ'_{reg} (real part) and electrical conductivity σ —are taken from NASA specifications [3].

5. EVALUATION AND RESULTS

5.1. Metrics

We evaluate pixel-wise performance using the root mean squared error (RMSE), normalized mean squared error (NMSE), structural similarity index (SSIM) and peak signal-to-noise ratio (PSNR), as used in [21] [6] [8].

5.2. Quantitative Results

To evaluate our model, we compared it against neural network-based RM prediction methods that could be easily re-trained with our dataset. Specifically, we selected PMNet [5] and RadioUNet [4]. PMNet achieved the best performance in the First Pathloss Radio Map Prediction Challenge [21], while RadioUNet was among the first and most effective NN-based RM prediction models. More recent architectures such as RMDM [7], RadioDiff [6], and RadioDiff- k^2 [8] leverage diffusion models but require substantial modification for our terrain setting or do not yet have publicly available code.

We evaluated all models at 415 MHz, 5.8 GHz and on combined results across both bands. Table 1 reports the quantitative metrics, where our method consistently outperforms the baselines across all metrics, with the greatest difference in SSIM and PSNR. In particular, our approach yields lower error and better structural fidelity, demonstrating its superiority over PMNet and RadioUNet. We performed an ablation study by either removing the final residual diffusion step (w/o Diff) or replacing the residual target with the RM itself (RM Diff). As shown in Table 1, both modifications degraded performance compared to our full model.

A qualitative comparison further highlights the advantages of our model. As shown in Fig. 3, our predictions preserve fine geometric details and accurately localize electromagnetic singularities, especially in low-connectivity regions such as craters and terrain shadowed by hills. In contrast, the baseline methods tend to blur sharp features in the radio map

or misplace singularities, leading to degraded interpretability and less reliable coverage prediction.

Table 1: Quantitative Comparison of Models.

Note: Results in **bold-underlined** and underlined represent the best and second best scores, respectively.

	Comparative		Ablation		
Methods	Radio-UNet	PMNet	w/o Diff	RM Diff	Ours
415 MHz					
NMSE	.001060	.000869	<u>.000820</u>	.000898	.000795
RMSE	.0243	.0220	<u>.0214</u>	.0224	.0211
SSIM \uparrow	.8810	.8867	<u>0.9035</u>	0.8975	.9069
PSNR \uparrow	33.54	34.26	<u>34.77</u>	34.53	35.40
5.8 GHz					
NMSE	.00339	<u>.00258</u>	.00264	.00270	.00248
RMSE	.0354	<u>.0309</u>	.0312	.0315	.0302
SSIM \uparrow	.8826	0.8898	.9009	<u>.9014</u>	0.9067
PSNR \uparrow	30.35	31.39	31.64	<u>32.10</u>	32.65
Both					
NMSE	.00223	<u>.00172</u>	.00173	.00180	.00164
RMSE	.0299	.0265	<u>.0263</u>	.0270	.0257
SSIM \uparrow	.8818	.8883	<u>.9022</u>	.8995	.9068
PSNR \uparrow	31.95	32.83	33.20	<u>33.32</u>	34.03

6. CONCLUSION

In this work, we presented a physics-informed deep learning framework for predicting radio maps over complex lunar terrain. Our approach integrates a synthetic lunar terrain generator with a ray-tracing pipeline to construct training datasets, and introduces a physics based triplet-UNet architecture composed of two UNets and a diffusion model. Experimental results on simulated lunar data demonstrate that our method improves accuracy over baseline deep learning models. Furthermore, we carried out ablation studies to show the need for the k^2 map. The ability to generate reliable radio maps over lunar terrain supports future communication-aware mission planning for the LunaNet framework.

7. ACKNOWLEDGMENT

This work was supported by NASA Grant 80NSSC22K0258 to the University of Washington, with a subcontract to Astrobotic Technology. Computations were performed on Hyak, the University of Washington's high-performance cluster funded by the student technology fee.

8. REFERENCES

- [1] D. J. Israel, J. W. Mitchell, M. A. Johnson, K. D. Mauldin, A. A. Pulkkinen, S. D. Christe, C. J. Roberts, L. V. D. Cooper, and C. J. Gramling, "LunaNet: A Flexible and Extensible Lunar Exploration Communications and Navigation Infrastructure," in *2020 IEEE Aerospace Conference*. IEEE, 2020, pp. 1–14.
- [2] K. Sato, K. Suto, K. Inage, K. Adachi, and T. Fujii, "Space-frequency-interpolated radio map," *IEEE Transactions on Vehicular Technology*, vol. 70, no. 1, pp. 714–725, 2021.
- [3] J. Hoydis, S. Cammerer, F. Ait Aoudia, M. Nimier-David, L. Maggi, G. Marcus, A. Vem, and A. Keller, "Sionna," 2022, <https://nvlabs.github.io/sionna/>.
- [4] R. Levie, Ç. Yapar, G. Kutyniok, and G. Caire, "Radiounet: Fast radio map estimation with convolutional neural networks," *IEEE Transactions on Wireless Communications*, vol. 20, no. 6, pp. 4001–4015, 2021.
- [5] J.-H. Lee, J. Lee, S.-H. Lee, and A. F. Molisch, "Pmnet: Large-scale channel prediction system for icassp 2023 first pathloss radio map prediction challenge," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–2.
- [6] X. Wang, K. Tao, N. Cheng, Z. Yin, Z. Li, Y. Zhang, and X. Shen, "Radiodiff: An effective generative diffusion model for sampling-free dynamic radio map construction," *IEEE Transactions on Cognitive Communications and Networking*, vol. 11, no. 2, pp. 738–750, 2025.
- [7] H. Jia, W. Chen, Z. Huang, H. Xiao, N. Jia, K. Wu, S. Lai, and Y. Yue, "Rmdm: Radio map diffusion model with physics informed," 2025. [Online]. Available: <https://arxiv.org/abs/2501.19160>
- [8] X. Wang *et al.*, "RadioDiff- k^2 : Helmholtz Equation Informed Generative Diffusion Model for Multi-Path Aware Radio Map Construction," *arXiv preprint arXiv:2504.15623*, apr 2025, arXiv:2504.15623. [Online]. Available: <https://arxiv.org/abs/2504.15623>
- [9] E. Staudinger, R. Giubilato, M. J. Schuster, R. Pöhlmann, S. Zhang, A. Dömel, A. Wedler, and A. Dammann, "Terrain-aware communication coverage prediction for cooperative networked robots in unstructured environments," *Acta Astronautica*, vol. 202, pp. 799–805, 2023. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0094576522006002>
- [10] J. Jost, *Riemannian Geometry and Geometric Analysis*, 8th ed. Springer, 2020. [Online]. Available: <https://doi.org/10.1007/978-3-030-56620-1>
- [11] E. Burman, P. Hansbo, M. G. Larson, and A. Massing, "A stable cut finite element method for partial differential equations on surfaces: The helmholtz-beltrami operator," 2018. [Online]. Available: <https://arxiv.org/abs/1810.04217>
- [12] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic differentiation in pytorch," 2017.
- [13] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," *arXiv preprint arXiv:2006.11239*, 2020. [Online]. Available: <https://arxiv.org/abs/2006.11239>
- [14] T. Salimans and J. Ho, "Progressive distillation for fast sampling of diffusion models," 2022. [Online]. Available: <https://arxiv.org/abs/2202.00512>
- [15] P. von Platen, S. Patil, A. Lozhkov, P. Cuenca, N. Lambert, K. Rasul, M. Davaadorj, D. Nair, S. Paul, W. Berman, Y. Xu, S. Liu, and T. Wolf, "Diffusers: State-of-the-art diffusion models," <https://github.com/huggingface/diffusers>, 2022.
- [16] Y. Cai and W. Fa, "Meter-scale topographic roughness of the Moon: The effect of small impact craters," *Journal of Geophysical Research: Planets*, vol. 125, p. e2020JE006429, 2020. [Online]. Available: <https://doi.org/10.1029/2020JE006429>
- [17] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes 3rd Edition: The Art of Scientific Computing*. New York: Cambridge University Press, 2007.
- [18] G. Neukum, B. A. Ivanov, and W. K. Hartmann, "Cratering records in the inner solar system in relation to the lunar reference system," *Space Science Reviews*, vol. 96, pp. 55–86, 2001.
- [19] C. I. Fassett and B. J. Thomson, "Crater degradation on the lunar maria: Topographic diffusion and the rate of erosion on the Moon," *Journal of Geophysical Research: Planets*, vol. 119, pp. 2255–2271, 2014.
- [20] D. J. Israel and C. J. Gramling, "LunaNet Interoperability Specification Document," National Aeronautics and Space Administration (NASA), Goddard Space Flight Center, Greenbelt, MD, Technical Publication NASA/TP–20210021073/Rev.5, Sep 2023.
- [21] Ç. Yapar, F. Jaensch, R. Levie, G. Kutyniok, and G. Caire, "The first pathloss radio map prediction challenge," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–2.