

**APPLICATION OF DENOISING DIFFUSION
PROBABILISTIC MODEL IN SOLVING THE CAUCHY
PROBLEM FOR THE LAPLACE EQUATION
ON A UNIT DISK**

R. Solopatykh, Y. Muzychuk

*Ivan Franko National University of Lviv,
1, Universytetska str., 79007, Lviv, Ukraine,*

e-mail: rostyslav.solopatykh@lnu.edu.ua, yuriy.muzychuk@lnu.edu.ua

In this paper we study the numerical solution of the Cauchy problem for the Laplace equation in a unit disk using Denoising Diffusion Probabilistic Model (DDPM). The Dirichlet and Neumann data are prescribed on the top half of the boundary while the bottom half and the interior are considered unknown. The inverse problem is formulated as conditional generation of a harmonic field represented as a multi-channel image encoding of the domain, boundary geometry, observed boundary, observed Dirichlet and Neumann values. The model is trained to generate a single-channel image that represents the numerical solution within the interior of the domain, discretized on a grid at a prescribed pixel resolution. A synthetic training dataset is generated by sampling the Fourier coefficients associated with the closed-form analytical solution of the well-posed problem, using as Neumann boundary data a truncated Fourier series up to mode K . Empirical evaluations on datasets comprising up to 50,000 samples at resolution of 64×64 indicate that the proposed model degrades gracefully under the tested Dirichlet noise levels. These findings suggest that conditional diffusion models may serve as learned surrogates for certain PDE inverse problems, though broader validation is still needed.

Key words: denoising Diffusion Probabilistic Model, inverse problems, Cauchy problem, Laplace equation, noise-robustness, conditional generative modeling, scientific machine learning.

1. INTRODUCTION

Generative AI models such as DDPM - Denoising Diffusion Probabilistic Models [1] are rapidly advancing as tools for approximating complex data distributions in science and engineering. We observe interest in employing these models for scientific machine learning applications, in which the generated objects are required to satisfy explicit governing equations. In particular, it is important to emphasize the growing interest in the study of inverse problems in partial differential equations (PDEs), as a portion of recent research efforts has been devoted to these topics, as illustrated, for example, in [2] and [3].

Among classical and extensively studied inverse problems, a special role is played by the Cauchy problem for elliptic equations, which is ill-posed in the sense of Hadamard. Therefore, the Cauchy problem for the Laplace equation, as a specific instance of such problems, already benefits from a well-established framework of regularization techniques and numerical solvers, as described, for example, in [4]. Among the available numerical methods for solving such problems, it is worth highlighting the integral equation approach described in [5].

In this paper, we propose a methodology for solving the Cauchy problem for the Laplace equation by employing a Denoising Diffusion Probabilistic Model (DDPM), trained on ground-truth solutions in the unit disk domain, as an efficient and noise-robust

numerical solver. Our goal is to demonstrate the potential of generative AI models for addressing inverse problems. Similar approaches were suggested for electrical impedance tomography image reconstruction in [6] and [7]. Nevertheless, the aforementioned studies primarily concentrate on EIT-based imaging and on models such as the Complete Electrode Model, and their methodological frameworks are explicitly tailored to this specific class of inverse problems. More generally, a unified and systematic framework for deploying such models across the broad class of inverse problems has not yet been established, and their full potential within wider scientific and engineering applications remains largely unexplored.

1.1. PROBLEM STATEMENT

We consider a Cauchy problem for the Laplace equation on a unit disk represented in polar coordinates $\Omega = \{(r, \theta) : r \in [0, 1], \theta \in [0, 2\pi]\}$ as finding a function $u \in H^1(\Omega)$:

$$\Delta u = 0 \quad \text{in } \Omega \quad (1)$$

with partial boundary data $g \in H^1(\partial\Omega)$, $f \in L^2(\partial\Omega)$ on the upper half-circle of the unit disk $\Gamma_1 = \{(r, \theta) : r = 1, \theta \in [0, \pi]\}$:

$$\begin{aligned} u &= g(\theta) \quad \text{on } \Gamma_1 \quad (\text{Dirichlet}) \\ \frac{\partial u}{\partial n} &= f(\theta) \quad \text{on } \Gamma_1 \quad (\text{Neumann}) \end{aligned}$$

and no data on the remaining arc $\Gamma_2 = \{(r, \theta) : r = 1, \theta \in [\pi, 2\pi]\}$.

In this experiment, we simulate the imposition of boundary alternating currents on the unit disk Ω as Neumann boundary data, with the objective of reconstructing the corresponding harmonic interior electric potential. This setting is closely related to the problem formulated in [7], where the Complete Electrode Model (CEM) is employed on a circular domain, and it is also consistent with the experimentally acquired data reported in [8]. In particular, the Neumann boundary conditions for the input data are prescribed as in (2):

$$f(\theta) = \sum_{k=1}^K (a_k \cos(k\theta) + b_k \sin(k\theta)) \quad (2)$$

$$a_k, b_k \sim \mathcal{N}(0, \sigma_k^2), k = 1, \dots, K$$

In this setting, the closed-form analytical solution of the well-posed boundary value problem with Neumann boundary conditions is straightforward to obtain in polar coordinates (3):

$$u(r, \theta) = \sum_{k=1}^K \frac{1}{k} r^k (a_k \cos(k\theta) + b_k \sin(k\theta)) + C \quad (3)$$

Subsequently, we employ the analytical solution (3) to construct a synthetic dataset for training the DDPM model on ground-truth solutions, and then we apply the trained model to inverse problems within the same domain for evaluation.

1.2. MOTIVATION

Diffusion probabilistic model is a parameterized Markov chain trained using variational inference to produce samples matching the data after finite time [1]. These models have been widely adopted within the deep learning community and have been applied across a broad range of tasks. Of particular relevance to this work is their application in conditional generative modeling, such as text-to-image generation, which is described in detail in [9]. The central objective of this work is to reformulate the solution of the two-dimensional Cauchy problem as an image generation task. In this framework, conditional generative modeling is employed, with the model conditioned on Dirichlet and Neumann boundary data, in a manner analogous to text-conditioned image generation.

2. METHODS

To solve the Cauchy problem using a diffusion model and to enable conditional generation, the solution must first be represented as an image-like object. This image-based representation must be constructed so as to explicitly encode the geometry of the domain, the precise location of the boundary, and an unambiguous distinction between boundary values and interior values. In addition, it is essential to ensure that the prescribed Neumann data are strictly consistent with the corresponding Dirichlet data. This compatibility condition is necessary to guarantee the mathematical correctness of the Cauchy problem formulation.

To satisfy the aforementioned requirements, we adopt a fundamental concept from computer vision-image multichanneling. In this framework, the solution is represented as a 64×64 single-channel image, while the domain geometry, boundary locations, and boundary values are likewise encoded as single-channel images. These channels are stacked to form a multichannel representation, all defined on a square grid $[-1 - \epsilon, 1 + \epsilon]$ that fully contains the unit disk. Consequently, the *U-Net* architecture underlying the diffusion model is trained on samples, each consisting of a single ground-truth solution channel and five conditioning channels:

1. Interior mask – equal to 1 inside the unit disk and 0 outside.
2. Full boundary mask – equal to 1 on pixels within a half-pixel diagonal of $r = 1$.
3. Observed boundary mask – boundary pixels with $\theta \in [0, \pi]$.
4. Dirichlet values – $g(\theta)$ at observed boundary pixels and 0 elsewhere.
5. Neumann values – $f(\theta)$ at observed boundary pixels and 0 elsewhere.

The proposed conditioning strategy allows us to generate solutions to the Cauchy problem by employing the partial Dirichlet and Neumann boundary data as conditioning channels 4 and 5, respectively. A visualization of this interpretation is provided in Fig. 1. In the example, we make use of a ground-truth solution corresponding to simulated input currents $f(\theta)$, which is exactly a truncated Fourier series up to mode $K = 5$.

Denoising Diffusion Probabilistic Models (DDPMs) are characterized by relatively large model architectures. In the original work by [1], the model was trained on the CIFAR-10 dataset, which contains approximately 60,000 samples. This suggests that, for effective application of DDPMs we need a training dataset on the order of 50,000 samples. Therefore, in this study, we deliberately restrict our attention to a simplified unit disk domain, enabling the use of a closed-form analytical solution expressed in polar coordinates (3). We exploit this polar-coordinate solution to sample its coefficients (2),




Fig. 1. Multichannel solution representation in 64x64 pixel resolution

thereby generating tens of thousands of training samples and constructing a dataset of sufficient size for DDPM training. It is worth noting that by sampling these coefficients, we obtain not only the ground-truth solution of the underlying problem but also its associated Neumann boundary data (2), represented as simulated alternating currents $f(\theta)$.

We construct the dataset by sampling the coefficients of the analytical solution (3) as independent and identically distributed random variables drawn from normal distributions with mode-dependent variances given by $\sigma_k^2 = \frac{1}{k^2}$ for $k = 1, \dots, K$. This modeling choice suppresses higher-mode contributions, produces smoother sampled boundary data, and remains compatible with the weak formulation of (1).

In this study, the DDPM model architecture is configured to exactly replicate the settings of the original DDPM framework proposed by [1]. Specifically, the following design choices are adopted:

1. Linear noise-variance schedule with $\beta \in [10^{-4}, 2 \times 10^{-2}]$ as in [1];
2. Total of $T = 1000$ diffusion time steps [1];
3. Training batch size of 128, matching the configuration used for CIFAR-10 in the original work [1];
4. Learning rate of $2 \cdot 10^{-4}$, matching the original configuration [1];
5. U-Net with 3 downsampling blocks and 3 corresponding upsampling blocks, where each block comprises 2 ResNet blocks followed by a Downsampler or Upsampler block, respectively [1];
6. Self-attention block applied at a feature map resolution of 16×16 , as in the original work [1].

The objective of this study was not only to demonstrate the capability of DDPM to solve the inverse problem, but also to assess its robustness with respect to noise. We evaluated the model's performance on Cauchy problems with noisy input data representing perturbed "voltage measurements," i.e., noisy Dirichlet boundary data, while the Neumann boundary data were kept unchanged. We introduced multiple noise levels, specifically 0 %, 1 %, 2 %, 5 %, and 10 %, and assessed performance using several quantitative metrics, namely the mean squared error (MSE) and its root (RMSE), the relative L_2 error (4), and the relative L_∞ error (5).

$$\text{rel. } L_2 \text{ error} = \frac{\|u_{\text{pred}} - u_{\text{true}}\|_2}{\|u_{\text{true}}\|_2} \quad (4)$$

$$\text{rel. } L_\infty \text{ error} = \frac{\max_{\Omega} |u_{\text{pred}} - u_{\text{true}}|}{\max_{\Omega} u_{\text{true}} - \min_{\Omega} u_{\text{true}}} \quad (5)$$

where u_{pred} denotes the solution generated by the DDPM model, and u_{true} represents the reference (ground-truth) solution to the Cauchy problem.

3. RESULTS

For all experiments conducted in this study, we fixed the model architecture and training hyperparameters as specified in the previous section. A single training dataset comprising 50,000 samples was generated (using coefficient sampling, Fig. 2), and a single instance of the model was trained and subsequently used across all experiments. Both the training and validation sets used clean boundary input data, as the analysis of training with noisy inputs is postponed to future work.

Model training was conducted for 200 epochs using the AdamW optimizer with a learning rate of $2 \cdot 10^{-4}$ on a training set comprising 50,000 samples and a validation set comprising 5,000 samples. The effective number of training epochs was determined by analyzing the evolution of the loss function, as depicted in Fig. 4. We stopped at epoch 160 because training loss had plateaued and later validation-loss gains were modest.




Fig. 2. Sampled coefficients of $f(\theta)$ corresponding to modes $k = 1, 2, 3, 4, 5$




Fig. 3. Histogram of amplitude values for the Neumann boundary data $f(\theta)$




Fig. 4. Evolution of training and validation loss over DDPM training epochs

We first generated few solutions and visually compared them with the ground truth (see Fig. 5). Subsequently, we performed a quantitative evaluation of the model on a test set comprising 500 previously unseen samples, using several metrics and considering different noise levels added to the observed Dirichlet input data (0 %, 1 %, 2 %, 5 %, and 10 % noise levels).

The results of model evaluation across different noise levels are summarized in Tabl. 1 and visualized in Fig. 6.

Table 1

DDPM reconstruction error metrics evaluated on 500 unseen test samples with varying noise levels applied to Dirichlet data

Metric	$\sigma = 0.0$	$\sigma = 0.01$	$\sigma = 0.02$	$\sigma = 0.05$	$\sigma = 0.1$
Rel. L^2 Error	0.0315	0.0325	0.0371	0.0575	0.0998
Rel. L^∞ Error	0.0305	0.0316	0.0365	0.0574	0.0982
MSE	0.0003	0.0004	0.0005	0.0010	0.0030
RMSE	0.0170	0.0176	0.0197	0.0292	0.0498

4. DISCUSSION

We analyzed the generated training dataset to confirm that the sampled coefficients are sufficiently varied, ensuring the model sees a broad range of functions rather than overfitting to a narrow subset of possible “input currents”. This diversity in sampled coefficients is evident in Fig. 2, which shows no clustering along any radius on the $\frac{1}{k}$ circles. Additionally, we verified that the amplitudes of Neumann boundary value ranges are approximately normally distributed, thereby ensuring that the training set includes both high- and low-amplitude functions. As illustrated in Fig. 3, the Neumann value amplitudes indeed closely conform to a normal distribution.

Numerical solutions generated with conditioning on noise-free Dirichlet boundary observations closely matched the ground-truth solutions (see Fig. 5). Therefore, we exam-




Fig. 5. Numerical solutions generated by the DDPM conditioned on noise-free Dirichlet boundary observations

ined how DDPM reconstruction accuracy changes with Dirichlet boundary observation noise level σ . Each panel of Fig. 6 presents one error metric (relative L_2 , relative L_∞ , MSE, RMSE) computed over a set of 500 unseen test samples. The solid line represents the mean value of the error metric, while the error bars and shaded region denote its standard deviation. The model was trained only on noise-free data, as noted previously.

The evaluation results in Tabl.1 and Fig. 6 indicate that the model is reasonably robust to noise: all error metrics increase gradually with increasing noise, without abrupt spikes or exponential growth. At $\sigma = 0$, the relative L_2 and L_∞ errors are approximately 3 %, indicating high accuracy. The variance of each error metric grows with noise level,




Fig. 6. DDPM reconstruction error metrics vs Dirichlet data noise levels

as seen from the widening confidence band. We hypothesize that min-max scaling causes noise to affect low-amplitude samples more than high-amplitude ones, thereby widening the error metric deviation interval. The influence of noise on the generated solutions is illustrated in Fig. 7.

5. CONCLUSIONS

In this study, we demonstrate that a generative AI model, specifically a denoising diffusion probabilistic model (DDPM), can be employed to solve complex inverse problems, including the severely ill-posed Cauchy problem for the Laplace equation. We considered a simplified experimental configuration in the unit disk, in which we simulated the injection of currents along the upper portion of the boundary. This setup is conceptually related to image reconstruction in electrical impedance tomography (EIT), in a manner analogous to [7].

To enforce the ill-posed nature of the Cauchy problem, we prescribed both Dirichlet and Neumann boundary data only on the upper semicircle of the disk. We used a simplified model with a closed-form analytical solution in polar coordinates, which allowed straightforward generation of a synthetic dataset by sampling the corresponding expansion coefficients.

The DDPM was trained on a dataset consisting of 50,000 samples and subsequently evaluated on an independent test set of 500 samples under five noise levels: 0%, 1%, 2%, 5%, and 10%. The evaluation of solution reconstruction error metrics across these noise levels indicates that the model functions as a noise-robust solver for the investigated inverse problem.




Fig. 7. Influence of noise applied to the Dirichlet data (conditioning channel)

In future work, we aim to compare the DDPM-based reconstructions against those obtained with conventional approaches, in particular the finite element method (FEM) combined with Tikhonov regularization. Furthermore, we plan to investigate training the model on noisy input data, in order to more accurately reflect the characteristics of real-world inverse problems.

REFERENCES

1. Ho J. Denoising diffusion probabilistic models / Jonathan Ho, Ajay Jain, Pieter Abbeel // arXiv preprint, arXiv:2006.11239. – 2020. – DOI: <https://doi.org/10.48550/arXiv.2006.11239>.
2. Chung H. Diffusion models for inverse problems / Hyungjin Chung, Jeongsol Kim, Jong Chul Ye // arXiv preprint. – 2025. – URL: <https://arxiv.org/abs/2508.01975>.
3. Huang J. Diffusionpde: Generative pde-solving under partial observation / Jiahe Huang, Guandao Yang, Zichen Wang, Jeong Joon Park // Advances in Neural Information Processing Systems 37 (NeurIPS 2024). – 2024. – URL: <https://arxiv.org/abs/2406.17763>. – DOI: <https://doi.org/10.52202/079017-4140>.

4. Kirsch A. An Introduction to the Mathematical Theory of Inverse Problems / Andreas Kirsch // Applied Mathematical Sciences. – 2021. – Vol. 120. – ISBN 978-3-030-63343-1. – Switzerland, Cham: Springer, 2021. – DOI: <https://doi.org/10.1007/978-3-030-63343-1>.
5. Chapko R. On the use of an integral equation approach for the numerical solution of a cauchy problem for laplace equation in a doubly connected planar domain / R. Chapko, B. T. Johansson, Y. Savka // Inverse Problems in Science and Engineering. – 2013. – Vol. 21 (6). – P. 805–828. – DOI: <https://doi.org/10.1080/17415977.2013.829467>.
6. Shi D. Conditional diffusion model for electrical impedance tomography / Duanpeng Shi, Wendong Zheng, Di Guo, Huaping Liu // arXiv preprint arXiv:2501.05769. – 2025. – URL <https://arxiv.org/pdf/2501.05769>. – DOI: <https://doi.org/10.48550/arXiv.2501.05769>.
7. Shi S. A conditional diffusion model for electrical impedance tomography image reconstruction / Shuaikai Shi, Ruiyuan Kang, Panos Liatsis // arXiv preprint arXiv:2412.16979. – 2024. – URL <https://arxiv.org/pdf/2412.16979>. – DOI: <https://doi.org/10.48550/arXiv.2412.16979>.
8. Hauptmann A. Open 2d electrical impedance tomography data archive / Hauptmann Andreas, Kolehmainen Ville, Mach Nguyet Minh, Savolainen Tuomo, Seppänen Aku, Siltanen Samuli // arXiv preprint arXiv:1704.01178. – 2017. – DOI: <https://doi.org/10.48550/arXiv.1704.01178>.
9. Ramesh A. Hierarchical text-conditional image generation with clip latents / Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, Mark Chen // arXiv preprint. – 2022. – DOI: <https://doi.org/10.48550/arXiv.2204.06125>.

Article: received 11.02.2026

revised 11.03.2026

printing adoption 16.03.2026

ЗАСТОСУВАННЯ ЗНЕШУМЛЮВАЛЬНОЇ ДИФУЗІЙНОЇ ЙМОВІРНІСНОЇ МОДЕЛІ ДО РОЗВ'ЯЗАННЯ ЗАДАЧІ КОШІ ДЛЯ РІВНЯННЯ ЛАПЛАСА В ОДИНИЧНОМУ КРУЗІ

Р. Солопатич, Ю. Музичук

*Львівський національний університет імені Івана Франка,
вул. Університетська 1, Львів, 79007, Україна,
e-mail: rostyslav.solopatykh@lnu.edu.ua, yuriy.muzychuk@lnu.edu.ua*

Досліджуємо чисельне розв'язання задачі Коші для рівняння Лапласа в одиничному крузі за допомогою знешумлювальної дифузійної ймовірнісної моделі. Вхідні дані Діріхле та Неймана задані на частині межі, а саме – на верхньому напівколі, в той час як значення розв'язку на нижній частині межі та всередині кола вважаються невідомими. Така обернена задача формулюється як задача умовної генерації гармонічного поля поданого як багатоканальне зображення, що містить інформацію про форму області, розташування межі, частину межі, що спостерігається, а також дані Діріхле та Неймана. Модель тренується генерувати одноканальне зображення, що відповідає чисельному розв'язку задачі в області, яка дискретизована на сітці з заданою роздільною здатністю. Синтетичний набір тренувальних даних згенеровано за допомогою випадкової вибірки коефіцієнтів Фур'є, що відповідають аналітичному розв'язку коректної крайової задачі, в якій дані Неймана задано як часткову сумму ряду Фур'є до члена K . Емпіричні оцінки моделі на наборі з 50,000 екземплярів із роздільною здатністю 64×64 підтверджують, що запропонована модель демонструє певну стійкість до шуму. Отримані результати демонструють потенційні можливості

застосування генеративних моделей штучного інтелекту для розв'язання складних обернених задач.

Ключові слова: знешумлювальна дифузійна ймовірнісна модель, обернені задачі, задача Коші, рівняння Лапласа, стійкість до шуму, умовна генерація, наукове машинне навчання.