PHYSICS-INFORMED ADAPTIVE TRAINING FOR 3D ACOUSTIC WAVE PROPAGATION

Leonardo Mendonça* **& Márcio Marques*** **& Pavel Petrov & Lucas Nissenbaum**
Instituto de Matemática Pura e Aplicada (IMPA)
Rio de Janeiro, Brazil
{leonardo.mendonca, marcio.marques, pavel.petrov, nissenbaum}@impa.br

ABSTRACT

Physics-informed neural networks (PINNs) provide an implicit, meshless approach to solving high-dimensional partial differential equations while avoiding the curse of dimensionality. Despite recent progress, their application to three-dimensional wave propagation problems in time domain remains limited. In this study, a novel stable training scheme that combines adaptive sampling, absorbing boundary conditions, and neural tangent kernel (NTK) dynamical loss balancing is proposed. Exploiting strong space–time localization of wavefields allows accurate propagation modeling, while avoiding the exponential growth of the number of collocation points with the dimension that is inherent to classical numerical methods.

1 INTRODUCTION

Wave propagation in heterogeneous media is a central problem in seismology, radiophysics, optics, underwater acoustics, and exploration geophysics (Virieux & Operto, 2009). While classical numerical methods provide accurate solutions, their computational cost grows rapidly with the problem dimension and required resolution. This is especially true when simulating an acoustic wave in time over three-dimensional space. The use of neural networks to solve these equations is a promising way to overcome these challenges.

Physics-Informed Neural Networks (PINNs) (Raissi et al., 2019) offer a meshless, implicit representation of the solutions of partial differential equations (PDEs) as continuous functions of the spatial coordinates and the time, alleviating some aspects of the curse of dimensionality (Wu et al., 2023). This formulation has motivated increasing interest in applying PINNs to wave propagation problems in heterogeneous media, both in time- and frequency-domain formulations (Moseley et al., 2023; Sethi et al., 2023; Zhang et al., 2023; Song et al., 2021; Alkhalifah & Huang, 2024). Nevertheless, applications of PINNs to the simulation of wave phenomena in three-dimensional environments remains limited, mostly focusing on reconstruction (Lü & Duan, 2025) and on the frequency-domain approaches (Schoder, 2025), or relying on mesh-dependent techniques (Zou et al., 2025; Bizzi et al., 2025).

Despite some promising results, time-domain wave propagation still presents intrinsic challenges for PINNs. For instance, solutions of the wave equation are usually highly localized in space and time due to the concentration of their energy is near propagating wavefronts, that leads to severe gradient imbalance in residual-based training. Moreover, standard neural networks suffer from spectral bias and lack built-in causality (Wang et al., 2022; Xu et al., 2019; Wang et al., 2024; Yu et al., 2022). As a result, training may often converge to incorrect and physically meaningless solutions (Mishra & Shekar, 2025; Krishnapriyan et al., 2021).

Several architectural and training strategies have been proposed to mitigate these effects, including Fourier feature embeddings (Tancik et al., 2020), absorbing boundary conditions and hard enforcement of initial conditions (Ding et al., 2025), causal architectures (Zhao et al., 2024; Bizzi et al., 2025), and hybrid approaches with finite difference methods (Zou et al., 2025). Recent studies further show that adaptive sampling combined with the neural tangent kernel (NTK) establishes a

*Equal contribution.

stable training regime in two-dimensional problems by concentrating collocation points in dynamically important regions, such as wavefronts and media interfaces, where the reflected waves are generated (Marques et al., 2025). This framework mitigates gradient imbalance and spectral bias, leading to improved training stability and accuracy, particularly in two-dimensional environments.

In the case of non-stationary wave equation in three-dimensional spatial domain, adaptive sampling becomes even more critical as the complexity of the problem scales. We propose using an **adaptive sampling** strategy on the meshless PINN framework, whereby we can leverage the sparse structure of wavefields so that the training can focus mostly in a relatively small spacetime region of interest. The use of such strategies makes training substantially more stable, leading to higher accuracy of the resulting solutions.

2 PHYSICS-INFORMED NEURAL NETWORKS

Physics-Informed Neural Networks (PINNs) (Raissi et al., 2019) approximate the solution of partial differential equations (PDE) through a meshless, implicit neural representation over the domain. Instead of relying on structured meshes, the PINN framework allows a more flexible approach to choosing collocation points, including the use of adaptive sampling that focuses on dynamically important regions, such as wavefronts and media interfaces. This flexibility enables accurate resolution of critical solution features without uniformly increasing the sampling density.

In this study, we consider the application of PINNs to wave propagation modeling in the context of underwater acoustics (for the sake of concreteness). The propagation of acoustic waves in a three-dimensional heterogeneous medium is described, in the time domain, by the wave equation

$$\frac{\partial^2 p(\mathbf{x},t)}{\partial t^2} - c^2(\mathbf{x})\Delta p(\mathbf{x},t) = S(\mathbf{x},t,\mathbf{x}_s)\,, \tag{1}$$

where $p(\mathbf{x},t)$ is unknown acoustic pressure, while $c(\mathbf{x})$ is known sound velocity distribution. We adopt a common setup in seismic simulations, where a point source located at $\mathbf{x}_s$ is approximated by a separable space-time source term S, consisting of a Gaussian spatial kernel, and a source waveform $s(t)$, given by a Ricker wavelet. We assume that the system is initially at rest, i.e, $p(\mathbf{x},0) = 0$ and $\partial_t p(\mathbf{x},0) = 0$, and that the propagation domain is infinite, therefore we employ high-order Higdon's absorbing boundary conditions (BCs) (Higdon, 1986). A similar approach was taken in Marques et al. (2025) and Ding et al. (2025). For a more precise description, see Appendix A.

In this context, we represent the acoustic wavefield by a multilayer perceptron (MLP) enhanced by Fourier feature mapping (Tancik et al., 2020). More precisely, we adopt a hard-constrained formulation, following the architectures used by Ding et al. (2025) and Marques et al. (2025). The output of the neural network is here multiplied by a t^2 factor, which enforces both a vanishing wavefield and a zero time derivative at $t = 0$. As a result, the initial conditions are satisfied exactly by construction. This contrasts with the widely used strategy in the PINN literature of enforcing initial conditions as soft constraints (Raissi et al., 2019; Hu et al., 2024), which may lead to unstable training dynamics due to unbalanced or competing gradients and does not guarantee a solution satisfying the initial conditions (Wang et al., 2021).

In the resulting framework, the governing physics is enforced through a weighted sum of residuals of the wave equation and absorbing boundary conditions (Higdon, 1986) evaluated at collocation points $\mathbf{v}_i = (\mathbf{x}_i, t_i)$,

$$\mathcal{R}_{\text{PDE}}(\theta,\mathbf{v}_i) = \left(\frac{\partial^2}{\partial t^2} - c^2\Delta\right) p(\mathbf{v}_i) - S(\mathbf{v}_i)\,,\ \mathcal{R}_{\text{ABC}}(\theta,\mathbf{v}_i) = \prod_{j=1}^{N}\left(\frac{\partial}{\partial t} + C_j\frac{\partial}{\partial \mathbf{n}}\right) p(\mathbf{v}_i)\,, \tag{2}$$

where $C_j = c(\mathbf{v}_i)$ for all j. The associated loss terms are defined as empirical mean-squared residuals over the collocation sets $\mathcal{X}_{\text{PDE}}$ and $\mathcal{X}_{\text{ABC}}$, containing N_{PDE} and N_{ABC} points, respectively,

$$\mathcal{L}_{\text{PDE}}(\theta,\mathcal{X}_{\text{PDE}}) = \frac{1}{N_{\text{PDE}}}\sum_{\mathcal{X}_{\text{PDE}}} |\mathcal{R}_{\text{PDE}}(\theta,\mathbf{v}_i)|^2\,, \quad \mathcal{L}_{\text{ABC}}(\theta,\mathcal{X}_{\text{ABC}}) = \frac{1}{N_{\text{ABC}}}\sum_{\mathcal{X}_{\text{ABC}}} |\mathcal{R}_{\text{ABC}}(\theta,\mathbf{v}_i)|^2\,. \tag{3}$$

3 Sampling Strategies and Balancing

The traditional uniform strategy used in PINNs samples collocation points using a regular grid, drawing from a uniform distribution or Latin hypercubes (Raissi et al., 2019; Ding et al., 2025). An interesting alternative is to use adaptive sampling, where the points distribution density is defined by the residual Wu et al. (2023); Daw et al. (2022); Bonfanti et al. (2025).

In particular, for three-dimensional wave problems the solution tends to be localized in space and time and to contain highly varying degrees of complexity throughout the domain. This suggests that sampling in a way that attributes equal importance to all regions of spacetime is inefficient at best, and unfeasible at worst. We adopt the adaptive sampling strategy of Marques et al. (2025), which we outline below.

After every 100 training steps, the wave equation residue is evaluated on a regular grid of collocation points. An adaptive distribution for each grid cube is constructed, where the probability of a cube is proportional to the residue at its center point. We then sample from a uniform distribution the box centered on this point. This approach leads to sampling primarily on regions where the PINN residue is high, thus locating complex parts of the solution, such as media interfaces or complex waveforms, that the PINN is currently in the process of learning, as shown in Figures 1b and 1d.

The adaptive distribution is used to sample a set $\mathcal{X}_A$ of adaptive points. We also sample a set $\mathcal{X}_U$ of uniform points, so as to keep the model from forgetting parts of the solution where the residue is currently low. The loss is then obtained as a combination of losses on these two sets:

$$\mathcal{L}_{PDE}(\theta) = \mathcal{L}_{PDE}(\theta, \mathcal{X}_U) + \beta \mathcal{L}_{PDE}(\theta, \mathcal{X}_A), \tag{4}$$

where β is a weight parameter. In order to keep adaptive points' influence stable throughout training, we dynamically adjust the value of β using the Neural Tangent Kernel (Jacot et al., 2018).

The Neural Tangent Kernel of PINNs (Wang et al., 2022) describes the evolution of separate loss terms throughout training as coupled dynamical systems, which leads to an algorithm for adjusting the weights of different losses and ensuring that no single loss term dominates the training. This algorithm has been applied in many PINN problems to weigh different loss terms (Wu et al., 2024; Zhongkai et al., 2024), and is adapted by Marques et al. (2025) to the context of adaptive sampling.

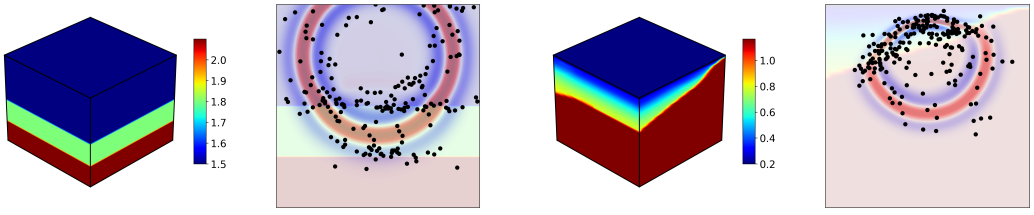

(a) Three layers (b) Three layers points (c) Utah Forge (d) Utah Forge points

Figure 1: (a) and (c): Sound velocity $c(\mathbf{x})$ distribution in the three-dimensional domains used in the experiments. (b) and (d): A vertical slice of the solution $p(\mathbf{x}, t)$ at $t = 0.35$ showing adaptive points after 10000 training steps. In both domains, the adaptively sampled points follow the wavefront and focus on the interfaces between media.

4 Experiments and Results

As the objective of this paper is to extend the PINN techniques from Marques et al. (2025) to three-dimensional problems, we consider two propagation domains. **Three layers** environment consists of parallel horizontal layers with sharp interfaces, leading to transmitted and reflected waves that present a nontrivial benchmark for PINNs, even in two dimensions (Bizzi et al., 2025). **Utah Forge** case is a real-world seismic field described by Vasco & Chan (2022), leading to a solution with a more pronounced three-dimensional effects. In order to normalize the problem dimensions, we divide the original velocities and lengths in Utah Forge by a factor of 5. Due to the linearity of the

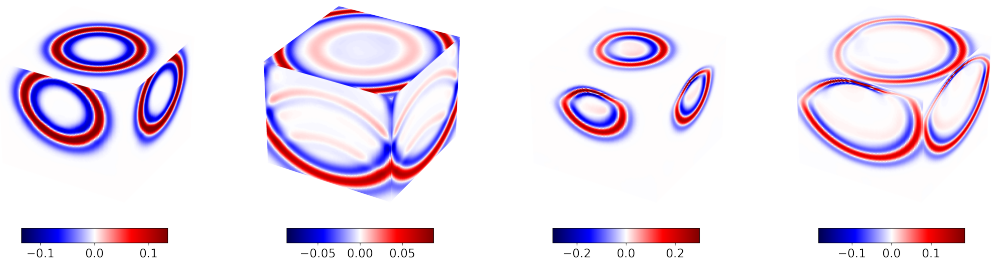

(a) Three layers $t = 0.30$ (b) Three layers $t = 0.50$ (c) Utah Forge $t = 0.30$ (d) Utah Forge $t = 0.50$

Figure 2: PINN solutions obtained for the Three-layered medium (a and b) and Utah Forge (c and d). Using adaptive sampling, NTK dynamic weights and the other techniques described, the PINN is able to accurately learn the correct form of the wave and its reflections.

wave equation, this scaling does not affect the true solution, but PINNs strongly benefit from this normalization (Xu et al., 2025). These domains are depicted in Figure 1 and in Appendix B.

Each domain has coordinates $\mathbf{x} \in [0, 1]^3$ and $t \in [0, 0.5]$. For the forcing term, the Gaussian kernel has width $\sigma = 0.025$ and is centered at $\mathbf{x}_0 = (0, 0, 0.25)$, and the Ricker wavelet $s(t)$ has amplitude $M_0 = 10^4$ and frequency $f_0 = 10$.

We use the same architecture as in Marques et al. (2025), with 180 random Fourier features and 5 hidden layers of each 150 neurons each. At each training step, we calculate the PDE residue on 20000 collocation points and the absorbing boundary conditions on 3000 points per boundary. In the case of adaptive sampling, the PDE points are split into 17500 uniform and 2500 adaptive points. The model is trained for 30000 steps, which is sufficient for the loss to become stable in every experiment.

Table 1: Relative L_1 and L_2 errors and training time (TT, in seconds) for wave propagation in both domains, with their average and standard deviation calculated over 5 different runs.

Sampler	**Three layers**			**Utah Forge**		
	L1RE ↓	**L2RE** ↓	**TT**	**L1RE** ↓	**L2RE** ↓	**TT**
Uniform	1.165 ± 0.515	1.024 ± 0.376	$\mathbf{2578 \pm 62}$	0.959 ± 0.165	0.849 ± 0.187	$\mathbf{2775 \pm 5}$
Adaptive	$\mathbf{0.067 \pm 0.003}$	$\mathbf{0.062 \pm 0.002}$	4358 ± 52	$\mathbf{0.048 \pm 0.012}$	$\mathbf{0.037 \pm 0.005}$	4622 ± 11

For each trained model, we evaluate the relative L_1 and L_2 errors against a reference solution as described in Appendix B, as in Zhongkai et al. (2024). The results, shown in Table 1 show that the adaptive training successfully converges to the correct wavefield, while the uniform strategy fails to converge even after 30000 steps. Despite the greater visual complexity of the Utah forge medium the three layer medium introduces higher frequencies when reflection occurs as it has a sharper discontinuity. This leads both methods to have more ease finding accurate representations of the solution for the Utah Forge dataset. We observed that adaptive training was able to track this physical phenomenon qualitatively as it is able to capture reflection at the media, as shown in Figure 2 above.

5 Conclusion

We introduced a stable Physics-Informed Neural Network framework for three-dimensional time-domain wave propagation in heterogeneous media. By combining adaptive collocation sampling, absorbing boundary conditions, and Neural Tangent Kernel (NTK)-guided dynamic loss balancing, the proposed approach addresses key challenges in training PINNs for wave equations, including gradient imbalance, spectral bias, and the lack of causality in standard neural architectures.

By exploiting the strong space-time localization of wavefields, the method focuses computational effort on dynamically relevant regions of the domain, enabling accurate wave propagation without the exponential growth in collocation points typical of mesh-based discretizations. These results suggest that implicit, meshless PINN formulations equipped with suitable stabilization and sampling strategies provide a promising foundation for extensions to more complex geometries, higher-frequency regimes, and data-informed wave modeling problems. As future work directions, we believe that applying adaptive sampling training strategies and weight balancing schemes for other complex partial differential equations and setups is merited.

ACKNOWLEDGMENTS

This work was supported by Petrobras.

REFERENCES

Tariq Alkhalifah and Xinquan Huang. Physics-informed neural wavefields with gabor basis functions. *Neural Networks*, 175:106286, 2024.

Arthur Bizzi, Leonardo M. Moreira, Márcio Marques, Leonardo Mendonça, Christian Júnior de Oliveira, Vitor Balestro, Lucas dos Santos Fernandez, Daniel Yukimura, Pavel Petrov, João M. Pereira, Tiago Novello, and Lucas Nissenbaum. Neuro-spectral architectures for causal physics-informed networks. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025.

Andrea Bonfanti, Ismael Medina, Roman List, Björn Staeves, Roberto Santana, and Marco Ellero. Pinn balls: Scaling second-order methods for pinns with domain decomposition and adaptive sampling. *arXiv preprint arXiv:2510.21262*, 2025.

Arka Daw, Jie Bu, Sifan Wang, Paris Perdikaris, and Anuj Karpatne. Mitigating propagation failures in physics-informed neural networks using retain-resample-release (r3) sampling. *arXiv preprint arXiv:2207.02338*, 2022.

Yi Ding, Su Chen, Hiroe Miyake, and Xiaojun Li. Physics-informed neural networks with fourier features for seismic wavefield simulation in time-domain nonsmooth complex media. *IEEE Transactions on Geoscience and Remote Sensing*, 63:1–13, 2025. doi: 10.1109/TGRS.2025.3581638.

Robert L Higdon. Absorbing boundary conditions for difference approximations to the multidimensional wave equation. *Math. comput.*, 47(176):437–459, 1986.

Zheyuan Hu, Khemraj Shukla, George Em Karniadakis, and Kenji Kawaguchi. Tackling the curse of dimensionality with physics-informed neural networks. *Neural Networks*, 176:106369, 2024. ISSN 0893-6080.

Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *Advances in neural information processing systems*, 31, 2018.

Aditi Krishnapriyan, Amir Gholami, Shandian Zhe, Robert Kirby, and Michael W Mahoney. Characterizing possible failure modes in physics-informed neural networks. *Advances in neural information processing systems*, 34:26548–26560, 2021.

Zhuosheng Lü and Lixia Duan. Reconstructing and predicting nonlinear waves from locally dense sparse 3d data using deep learning. *Nonlinear Dynamics*, pp. 1–16, 2025.

Márcio Marques, Leonardo Mendonça, Arthur Bizzi, Leonardo Moreira, Christian Oliveira, Deborah Oliveira, Lucas Fernandez, Vitor Balestro, João Pereira, Daniel Yukimura, et al. Stable adaptive training for physics-informed neural networks in acoustic wave propagation. *JASA Express Letters*, 5(11), 2025.

S Mishra and B Shekar. Wave equation and neural solvers: What works and what needs to change? In *86th EAGE Annual Conference & Exhibition*, volume 2025, pp. 1–5. European Association of Geoscientists & Engineers, 2025.

Ben Moseley, Andrew Markham, and Tarje Nissen-Meyer. Finite basis physics-informed neural networks (fbpinns): a scalable domain decomposition approach for solving differential equations. *Advances in Computational Mathematics*, 49(4):62, 2023.

Maziar Raissi, Paris Perdikaris, and George 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*, 378:686–707, 2019.

Stefan Schoder. Physics-informed neural networks for modal wave field predictions in 3d room acoustics. *Applied Sciences*, 15(2):939, 2025.

Harpreet Sethi, Doris Pan, Pavel Dimitrov, Jeffrey Shragge, Gunter Roth, and Ken Hester. Hard enforcement of physics-informed neural network solutions of acoustic wave propagation. *Computational geosciences*, 27(5):737–751, 2023.

Chao Song, Tariq Alkhalifah, and Umair Bin Waheed. Solving the frequency-domain acoustic vti wave equation using physics-informed neural networks. *Geophysical Journal International*, 225 (2):846–859, 2021.

Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in neural information processing systems*, 33:7537–7547, 2020.

Don Vasco and Coral Chan. Utah forge: Development of a reservoir seismic velocity model and seismic resolution study. Geothermal Data Repository, Array Information Technology, https://doi.org/10.15121/1989942, 2022. Accessed: 2025-11-26.

Jean Virieux and Stéphane Operto. An overview of full-waveform inversion in exploration geophysics. *Geophysics*, 74(6):WCC1–WCC26, 2009.

Sifan Wang, Yujun Teng, and Paris Perdikaris. Understanding and mitigating gradient flow pathologies in physics-informed neural networks. *SIAM Journal on Scientific Computing*, 43(5):A3055–A3081, 2021.

Sifan Wang, Xinling Yu, and Paris Perdikaris. When and why pinns fail to train: A neural tangent kernel perspective. *Journal of Computational Physics*, 449:110768, 2022.

Sifan Wang, Shyam Sankaran, and Paris Perdikaris. Respecting causality for training physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering*, 421: 116813, 2024.

Chenxi Wu, Min Zhu, Qinyang Tan, Yadhu Kartha, and Lu Lu. A comprehensive study of non-adaptive and residual-based adaptive sampling for physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering*, 403:115671, 2023.

Haixu Wu, Huakun Luo, Yuezhou Ma, Jianmin Wang, and Mingsheng Long. Ropinn: Region optimized physics-informed neural networks. *Advances in Neural Information Processing Systems*, 37:110494–110532, 2024.

Shengfeng Xu, Yuanjun Dai, Chang Yan, Zhenxu Sun, Renfang Huang, Dilong Guo, and Guowei Yang. On the preprocessing of physics-informed neural networks: How to better utilize data in fluid mechanics. *Journal of Computational Physics*, 528:113837, 2025.

Zhi-Qin John Xu, Yaoyu Zhang, Tao Luo, Yanyang Xiao, and Zheng Ma. Frequency principle: Fourier analysis sheds light on deep neural networks. *arXiv preprint arXiv:1901.06523*, 2019.

Jeremy Yu, Lu Lu, Xuhui Meng, and George Em Karniadakis. Gradient-enhanced physics-informed neural networks for forward and inverse pde problems. *Computer Methods in Applied Mechanics and Engineering*, 393:114823, 2022.

Yijie Zhang, Xueyu Zhu, and Jinghuai Gao. Seismic inversion based on acoustic wave equations using physics-informed neural network. *IEEE transactions on geoscience and remote sensing*, 61:1–11, 2023.

Zhiyuan Zhao, Xueying Ding, and B. Aditya Prakash. PINNsformer: A transformer-based framework for physics-informed neural networks. In *The Twelfth International Conference on Learning Representations*, 2024.

Hao Zhongkai, Jiachen Yao, Chang Su, Hang Su, Ziao Wang, Fanzhi Lu, Zeyu Xia, Yichi Zhang, Songming Liu, Lu Lu, et al. Pinnacle: A comprehensive benchmark of physics-informed neural networks for solving pdes. *Advances in Neural Information Processing Systems*, 37:76721–76774, 2024.

Jingbo Zou, Cai Liu, Yanghua Wang, Chao Song, Umair bin Waheed, and Pengfei Zhao. Accelerating the convergence of physics-informed neural networks for seismic wave simulation. *Geophysics*, 90(2):T23–T32, 2025.

A Acoustic wave propagation in a three-dimensional medium

The propagation of acoustic waves in a three-dimensional heterogeneous medium is described, in the time domain, by the scalar wave equation. The acoustic pressure field $p(\mathbf{x}, t)$ satisfies

$$\frac{\partial^2 p(\mathbf{x},t)}{\partial t^2} - c^2(\mathbf{x})\Delta p(\mathbf{x},t) = S(\mathbf{x},t) := G(\mathbf{x},\mathbf{x}_s)\, s(t), \tag{5}$$

where $c(\mathbf{x})$ denotes the heterogeneous acoustic wave speed. A point source located at $\mathbf{x}_s$ is approximated by a separable space-time source term, consisting of a Gaussian spatial kernel and a source time function,

$$G(\mathbf{x},\mathbf{x}_s) = e^{-\frac{\|\mathbf{x}-\mathbf{x}_s\|^2}{2\alpha^2}} \qquad \text{and} \qquad s(t) = M_0 \left(1 - 2\pi^2 f_0^2 \left(t - \frac{1}{f_0}\right)^2\right) \mathrm{e}^{-\pi^2 f_0^2 \left(t-\frac{1}{f_0}\right)^2}, \tag{6}$$

with $t_0 = 1/f_0$, where f_0 denotes the dominant frequency and M_0 the source amplitude.

We assume that the system is initially at rest, i.e, $p(\mathbf{x}, 0) = 0$ and $\partial_t p(\mathbf{x}, 0) = 0$. Therefore, to simulate an infinite domain, we employ high-order Higdon's absorbing BCs Higdon (1986), which can be written as

$$\prod_{j=1}^{N} \left(\frac{\partial}{\partial t} + C_j \frac{\partial}{\partial \mathbf{n}}\right) p = 0\,, \tag{7}$$

where N is the order of the condition, and $\mathbf{n}$ is the outward-normal vector at the given point of the boundary.

B Experimental Setup

We consider propagating a wave using PINNs in the two media detailed in Figure 3.

The PINN solutions are evaluated by comparing them to a reference solution obtained by a pseudo-spectral method with discretization $\Delta x = \Delta y = \Delta z = 10^{-2}$, and an explicit fourth-order Runge-Kutta solver with $\Delta t = 5 \times 10^{-4}$.

Models are evaluated on the relative L1 and L2 errors as in Zhongkai et al. (2024), defined as:

$$L1RE = \frac{\sum_i |u_{pred}(t_i, \mathbf{x_i}) - u_{ref}(t_i, \mathbf{x_i})|}{\sum_i |u_{ref}(t_i, \mathbf{x_i})|}, \tag{8}$$

$$L2RE = \sqrt{\frac{\sum_i \left(u_{pred}(t_i, \mathbf{x_i}) - u_{ref}(t_i, \mathbf{x_i})\right)^2}{\sum_i u_{ref}(t_i, \mathbf{x_i})^2}}. \tag{9}$$

where the summations are over a regular grid of 100 points in each space and time dimension.

All experiments were run on the same machine, having an i9-13900K CPU, an RTX 4090 GPU and 128GB RAM. In order to mitigate random effects, each experiment was repeated using 5 different random seeds.

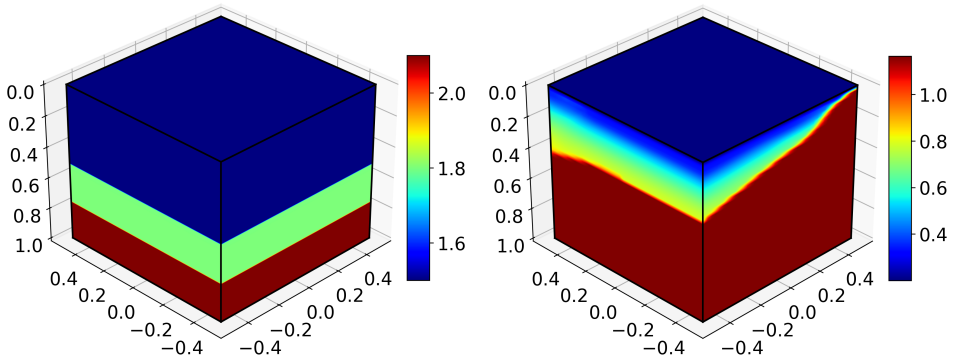

(a) Three layers (b) Utah Forge

Figure 3: Wave propagation velocity in both domains considered in the experiments.