PINNacle: A Comprehensive Benchmark of Physics-Informed Neural Networks for Solving PDEs

Zhongkai Hao[1]*, **Jiachen Yao**[1*], **Chang Su**[1*], **Hang Su**[1], **Ziao Wang**[1],

Fanzhi Lu[1], **Zeyu Xia**[1], **Yichi Zhang**[1], **Songming Liu**[1], **Lu Lu**[2], **Jun Zhu**[1†]

[1]Dept. of Comp. Sci. and Tech., Institute for AI, BNRist Center, THBI Lab,
Tsinghua-Bosch Joint ML Center, Tsinghua University
[2] Department of Statistics and Data Science, Yale University
hzj21@mails.tsinghua.edu.cn; dcszj@tsinghua.edu.cn

Abstract

While significant progress has been made on Physics-Informed Neural Networks (PINNs), a comprehensive comparison of these methods across a wide range of Partial Differential Equations (PDEs) is still lacking. This study introduces PINNacle, a benchmarking tool designed to fill this gap. PINNacle provides a diverse dataset, comprising over 20 distinct PDEs from various domains, including heat conduction, fluid dynamics, biology, and electromagnetics. These PDEs encapsulate key challenges inherent to real-world problems, such as complex geometry, multi-scale phenomena, nonlinearity, and high dimensionality. PINNacle also offers a user-friendly toolbox, incorporating about 10 state-of-the-art PINN methods for systematic evaluation and comparison. We have conducted extensive experiments with these methods, offering insights into their strengths and weaknesses. In addition to providing a standardized means of assessing performance, PINNacle also offers an in-depth analysis to guide future research such as domain decomposition methods and loss reweighting for handling multi-scale problems. To the best of our knowledge, it is the largest benchmark with a diverse and comprehensive evaluation that will undoubtedly foster further research in PINNs.

1 Introduction

Partial Differential Equations (PDEs) are of paramount importance in science and engineering, as they often underpin our understanding of intricate physical systems such as fluid flow, heat transfer, and stress distribution [37]. The computational simulation of PDE systems has been a focal point of research for an extensive period, leading to the development of numerical methods such as finite difference [6], finite element [45], and finite volume methods [14].

Recent advancements have led to the use of deep neural networks to solve forward and inverse problems involving PDEs [44, 62, 11, 54]. Among these, Physics-Informed Neural Networks (PINNs) have emerged as a promising alternative to traditional numerical methods in solving such problems [44, 25]. PINNs leverage the underlying physical laws and available data to effectively handle various scientific and engineering applications. The growing interest in this field has spurred the development of numerous PINN variants, each tailored to overcome specific challenges or to enhance the performance of the original framework.

*Equal contribution. †The corresponding author.

38th Conference on Neural Information Processing Systems (NeurIPS 2024) Track on Datasets and Benchmarks.

Figure 1: Architecture of PINNacle. It contains a dataset covering more than 20 PDEs, a toolbox that implements about 10 SOTA methods, and an evaluation module. These methods have a wide range of application scenarios like uid mechanics, electromagnetism, heat conduction, geophysics, and so on.

While PINN methods have achieved remarkable progress, a comprehensive comparison of these methods across diverse types of PDEs is currently lacking. Establishing such a benchmark is crucial as it could enable researchers to more thoroughly understand existing methods and pinpoint potential challenges. Despite the availability of several studies comparing sampling methods [30] and reweighting methods [21], there has been no concerted effort to develop a rigorous benchmark using challenging datasets from real-world problems. The sheer variety and inherent complexity of PDEs make it dif cult to conduct a comprehensive analysis. Moreover, different mathematical properties and application scenarios further complicate the task, requiring the benchmark to be adaptable and exhaustive.

To resolve these challenges, we propose PINNacle, a comprehensive benchmark for evaluating and understanding the performance of PINNs. As shown in Fig. 1, PINNacle consists of three major components — a diverse dataset, a toolbox, and evaluation modules. The dataset comprises tasks from over 20 different PDEs from various domains, including heat conduction, uid dynamics, biology, and electromagnetics. Each task brings its own set of challenges, such as complex geometry, multi-scale phenomena, nonlinearity, and high dimensionality, thus providing a rich testing ground for PINNs. The toolbox incorporates more than 10 state-of-the-art (SOTA) PINN methods, enabling a systematic comparison of different strategies, including loss reweighting, variational formulation, adaptive activations, and domain decomposition. These methods can be exibly applied to the tasks in the dataset, offering researchers a convenient way to evaluate the performance of PINNs which is also user-friendly for secondary development. The evaluation modules provide a standardized means of assessing the performance of different PINN methods across all tasks, ensuring consistency in comparison and facilitating the identi cation of strengths and weaknesses in various methods.

PINNacle provides a robust, diverse, and comprehensive benchmark suite for PINNs, contributing signi cantly to the eld's understanding and application. It represents a major step forward in the evolution of PINNs which could foster more innovative research and development in this exciting eld. Code and data are publicly available at https://github.com/i207M/PINNacle .

In a nutshell, our contributions can be summarized as follows:

- We design a dataset encompassing over 20 challenging PDE problems. These problems encapsulate several critical challenges faced by PINNs, including handling complex geometries, multi-scale phenomena, nonlinearity, and high-dimensional problems.
- We systematically evaluate more than 10 carefully selected representative variants of PINNs. We conducted thorough experiments and ablation studies to evaluate their performance. To the best of our knowledge, this is the largest benchmark comparing different PINN variants.

- We provide an in-depth analysis to guide future research. We show using loss reweighting and domain decomposition methods could improve the performance on multi-scale and complex geometry problems. Variational formulation achieves better performance on inverse problems. However, few methods can adequately address nonlinear problems, indicating a future direction for exploration and advancement.

2 Related Work

2.1 Benchmarks and datasets in scientific machine learning

The growing trend of AI in scientific research has stimulated the development of various benchmarks and datasets, which differ greatly in data formats, sizes, and governing principles. For instance, [33] presents a benchmark for comparing neural operators, while [3, 41] benchmarks methods for learning latent Newtonian mechanics. Furthermore, domain-specific datasets and benchmarks exist in fluid mechanics [20], climate science [42, 5], quantum chemistry [1], and biology [4].

Beyond these domain-specific datasets and benchmarks, physics-informed machine learning has received considerable attention [18, 9] since the advent of Physics-Informed Neural Networks (PINNs) [44]. These methods successfully incorporate physical laws into model training, demonstrating immense potential across a variety of scientific and engineering domains. Various papers have compared different components within the PINN framework; for instance, [10] and [60] investigate the sampling methods of collocation points in PINNs, and [21] compare reweighting techniques for different loss components. PDEBench [52] and PDEArena [17] design multiple tasks to compare different methods in scientific machine learning such as PINNs, FNO, and U-Net. Nevertheless, a comprehensive comparison of various PINN approaches remains absent in the literature.

2.2 Softwares and Toolboxes

A plethora of software solutions have been developed for solving PDEs with neural networks. These include SimNet [19], NeuralPDE [43], TorchDiffEq [8], and PyDEns [29]. More recently, DeepXDE [34] has been introduced as a fundamental library for implementing PINNs across different backends. However, there remains a void for a toolbox that provides a unified implementation for advanced PINN variants. Our PINNacle fills this gap by offering a flexible interface that facilitates the implementation and evaluation of diverse PINN variants. We furnish clear and concise code for researchers to execute benchmarks across all problems and methods.

2.3 Variants of Physics-informed neural networks

The PINNs have received much attention due to their remarkable performance in solving both forward and inverse PDE problems. However, vanilla PINNs have many limitations. Researchers have proposed numerous PINN variants to address challenges associated with high-dimensionality, non-linearity, multi-scale issues, and complex geometries [18, 9, 25, 30]. Broadly speaking, these variants can be categorized into: loss reweighting/resampling [57, 58, 53, 60, 40], innovative optimizers [61], novel loss functions such as variational formulations [62, 26, 27, 28] or regularization terms [63, 50], and novel architectures like domain decomposition [21, 31, 38, 24] and adaptive activations [23, 22]. These variants have enhanced PINN's performance across various problems. Here we select representative methods from each category and conduct a comprehensive analysis using our benchmark dataset to evaluate these variants.

3 PINNacle: A Hierarchical Benchmark for PINNs

In this section, we first introduce the preliminaries of PINNs. Then we introduce the details of datasets (tasks), PINN methods, the toolbox framework, and the evaluation metrics.

3.1 Preliminaries of Physics-informed Neural Networks

Physics-informed neural networks are neural network-based methods for solving PDEs as well as inverse problems of PDEs, which have received much attention recently. Specifically, let's consider a

general Partial Differential Equation (PDE) system de ned on , which can be represented as:

$$F(u(x); x) = 0; \quad x \in \quad ; \tag{1}$$

$$B(u(x); x) = 0; \quad x \in @ : \tag{2}$$

where F is a differential operator and B is the boundary/initial condition. PINN uses a neural network $u(x)$ with parameters to approximate $u(x)$. The objective of PINN is to minimize the following loss function:

$$L( ) = \frac{w_c}{N_c} \sum_{i=1}^{N_c} \| F(u(x_c^i); x_c^i)\|^2 + \frac{w_b}{N_b} \sum_{i=1}^{N_b} \| B(u(x_b^i); x_b^i)\|^2 + \frac{w_d}{N_d} \sum_{i=1}^{N_d} \| u(x_d^i) \quad u(x_d^i)\|^2 : \tag{3}$$

where $w_c; w_b; w_d$ are weights. The rst two terms enforce the PDE constraints on $\{x_c^i\}_{1:::N_c}$ and boundary conditions on $\{x_b^i\}_{1:::N_b}$. The last term is data loss, which is optional when there is data available. However, PINNs have several inherent drawbacks. First, PINNs optimize a mixture of imbalance loss terms which might hinder its convergence as illustrated in [57]. Second, nonlinear or stiff PDEs might lead to unstable optimization [58]. Third, the vanilla MLPs might have dif culty in representing multi-scale or high-dimensional functions. For example, [30] shows that vanilla PINNs only work for a small parameter range, even in a simple convection problem. To resolve these challenges, numerous variants of PINNs are proposed. However, a comprehensive comparison of these methods is lacking, and thus it is imperative to develop a benchmark.

3.2 Datasets

To effectively compare PINN variants, we've curated a set of PDE problems (datasets) representing a wide range of challenges. We chose PDEs from diverse domains, re ecting their importance in science and engineering. Our dataset includes 22 unique cases, with further details in Appendix B.

- The Burgers' Equation, fundamental to uid mechanics, considering both one and two-dimensional problems.
- The Poisson's Equation, widely used in math and physics, with four different cases.
- The Heat Equation, a time-dependent PDE that describes diffusion or heat conduction, demonstrated in four unique cases.
- The Navier-Stokes Equation, describing the motion of viscous uid substances, showcased in three scenarios: a lid-driven ow (NS2d-C), a geometrically complex backward step ow (NS2d-CG), and a time-dependent problem (NS2d-LT).
- The Wave Equation, modeling wave behavior, exhibited in three cases.
- Chaotic PDEs, featuring two popular examples: the Gray-Scott (GS) and Kuramoto-Sivashinsky (KS) equations.
- High Dimensional PDEs, including the high-dimensional Poisson equation (PNd) and the high-dimensional diffusion or heat equation (HNd).
- Inverse Problems, focusing on the reconstruction of the coef cient eld from noisy data for the Poisson equation (PInv) and the diffusion equation (HInv).

It is important to note that we have chosen PDEs encompassing a wide range of mathematical properties. This ensures that the benchmarks do not favor a speci c type of PDE. The selected PDE problems introduce several core challenges, which include:

- Complex Geometry: Many PDE problems involve complex or irregular geometry, such as heat conduction or wave propagation around obstacles. These complexities pose signi cant challenges for PINNs in terms of accurate boundary behavior representation.
- Multi-Scale Phenomena: Multi-scale phenomena, where the solution varies signi cantly over different scales, are prevalent in situations such as turbulent uid ow. Achieving a balanced representation across all scales is a challenge for PINNs in multi-scale scenarios.
- Nonlinear Behavior: Many PDEs exhibit nonlinear or even chaotic behavior, where minor variations in initial conditions can lead to substantial divergence in outcomes. The optimization of PINNs becomes intriguing on nonlinear PDEs.

- High Dimensionality: High-dimensional PDE problems, frequently encountered in quantum mechanics, present signi cant challenges for PINNs due to the “curse of dimensionality”. This term refers to the increase in computational complexity with the addition of each dimension, accompanied by statistical issues like data sparsity in high-dimensional space.

These challenges are selected due to their frequent occurrence in numerous real-world applications. As such, a method's performance in addressing these challenges serves as a reliable indicator of its overall practical utility. Table 1 presents a detailed overview of the dataset, the PDEs, and the challenges associated with these problems. We generate data using FEM solver provided by COMSOL 6.0 [39] for problems with complex geometry and spectral method provided by Chebfun [12] for chaotic problems. More details can be found in Appendix B.

Dataset	Complex geometry	Multi-scale	Nonlinearity	High dim
Burgers[1, 2]			p	
Poisson[3, 6]	p	p		
Heat[7, 10]	p	p		
NS[11, 13]	p	p	p	
Wave[14, 16]	p	p		
Chaotic[17, 18]		p	p	
High dim[19, 20]				p
Inverse[21, 22]			p	

Table 1: Overview of our datasets along with their challenges. We chose 22 cases in total to evaluate the methods of PINNs. The left picture shows the visualization of cases with these four challenges, i.e., complex geometry, multi-scale, nonlinearity, and high dimension.

3.3 Methods and Toolbox

After conducting an extensive literature review, we present an overview of diverse PINNs approaches for comparison. Then we present the high-level structure of our PINNacle.

3.3.1 Methods

As mentioned above, variants of PINNs are mainly based on loss functions, architecture, and optimizer [18]. The modi cations to loss functions can be divided into reweighting existing losses and developing novel loss functions like regularization and variational formulation. Variants of architectures include using domain decomposition and adaptive activations.

The methods discussed are directly correlated with the challenges highlighted in Table 1. For example, domain decomposition methods are particularly effective for problems involving complex geometries and multi-scale phenomena. Meanwhile, loss reweighting strategies are adept at addressing imbalances in problems with multiple losses. We have chosen variants from these categories based on their signi cant contributions to the eld.

Here, we list the primary categories and representative methods as summarized in Table 2:

- Loss reweighting/Resampling (2 4): PINNs are trained with a mixed loss of PDE residuals, boundary conditions, and available data losses shown in Eq 3. Various methods [57, 59, 2, 35, 47] propose different strategies to adjust these weights w_c, w_b and w_d at different epochs or resample collocation points $\{x_c^i\}$ and $\{x_b^i\}$ in Eq 3, which indirectly adjust the weights [60, 40]. We choose three famous examples, i.e., reweighting using gradient norms (PINN-LRA) [57], using neural tangent kernel (PINN-NTK) [59], and residual-based resampling (RAR)[34, 60].
- Novel optimizer (5): To handle the problem of multi-scale objectives, some new optimizers [32, 61] are proposed. We chose MultiAdam, which is resistant to domain scale changes.
- Novel loss functions (6 7): Some works introduce novel loss functions like variational formulation [49, 28, 27] and regularization terms to improve training. We choose hp-VPINN [27] and gPINN [63, 50], which are representative examples from these two categories.

	Complex Geometry	Multi-scale	Nonlinearity	High dim
Vanilla PINN [1]	p	p		
Reweighting/Resampling [2–4]		p		
Novel Optimizer [7]				
Novel Loss Functions [5, 6]	p	p		
Novel Architecture [8–10]		p		

Table 2: Overview of methods in our PINNacle. denotes the method is potentially designed to solve or show empirical improvements for problems encountering the challenge and vice versa.

- Novel activation architectures (8–10): Some works propose various network architectures, such as using CNN and LSTM [14, 15, 46], custom activation functions [22, 23], and domain decomposition [21, 48, 24, 38]. Among adaptive activations for PINNs, we choose LAAF [22] and GAAF [23]. Domain decomposition is a method that divides the whole domain into multiple subdomains and trains subnetworks on these subdomains. It is helpful for solving multi-scale problems, but multiple subnetworks increase the difficulty of training. XPINNs, cPINNs, and FBPINNs [21, 24, 38] are three representative examples. We choose FBPINNs which is the state-of-the-art domain decomposition that applies domain-specific normalization to stabilize training.

3.3.2 Structure of Toolbox

We provide a user-friendly and concise toolbox for implementing, training, and evaluating diverse PINN variants. Specifically, our codebase is based on DeepXDE and provides a series of encapsulated classes and functions to facilitate high-level training and custom PDEs. These utilities allow for a standardized and streamlined approach to the implementation of various PINN variants and PDEs. Moreover, we provided many auxiliary functions, including computing different metrics, visualizing predictions, and recording results.

Despite the unified implementation of diverse PINNs, we also design an adaptive multi-GPU parallel training framework to enhance the efficiency of systematic evaluations of PINN methods. It addresses the parallelization phase of training on multiple tasks, effectively balancing the computational loads of multiple GPUs. It allows for the execution of larger and more complex tasks. In a nutshell, we provide an example code for training and evaluating PINNs on two Poisson equations using our PINNacle framework in Appendix D.

3.4 Evaluation

To comprehensively analyze the discrepancy between the PINN solutions and the true solutions, we adopt multiple metrics to evaluate the performance of the PINN variants. Generally, we choose several metrics that are commonly used in literature that apply to all methods and problems. We suppose that $\mathbf{y} = (y_i)_{i=1}^n$ is the prediction and $\mathbf{y}^0 = (y_i^0)_{i=1}^n$ to is ground truth, where n is the number of testing examples. Specifically, we use ℓ_2 relative error (L2RE), and ℓ_1 relative error (L1RE) which are two most commonly used metrics to measure the global quality of the solution,

$$\text{L2RE} = \sqrt{\frac{\sum_{i=1}^n (y_i - y_i^0)^2}{\sum_{i=1}^n y_i^{0\,2}}}; \quad \text{L1RE} = \frac{\sum_{i=1}^n |y_i - y_i^0|}{\sum_{i=1}^n |y_i^0|}. \tag{4}$$

We also compute max error (mERR in short), mean square error (MSE), and Fourier error (fMSE) for a detailed analysis of the prediction. These three metrics are computed as follows:

$$\text{MSE} = \frac{1}{n}\sum_{i=1}^n (y_i - y_i^0)^2; \quad \text{mERR} = \max_i |y_i - y_i^0|; \quad \text{fMSE} = \sqrt{\frac{\sum_{k_{\min}}^{k_{\max}} |F(y) - F(y^0)|^2}{k_{\max} - k_{\min} + 1}}; \tag{5}$$

where F denotes Fourier transform of y and $k_{\min}$, $k_{\max}$ are chosen similar to PDEBench [52]. Besides, for time-dependent problems, investigating the quality of the solution with time is important. Therefore we compute the L2RE error varying with time in Appendix E.2.

We assess the performance of PINNs against the reference from numerical solvers. Experimental results utilizing the ℓ_2 relative error (L2RE) metric are incorporated within the main text, while a more exhaustive set of results, based on the aforementioned metrics, is available in the Appendix E.1.

L2RE	Name	Vanilla			Loss Reweighting/Sampling			Optimizer	Loss functions		Architecture		
–		PINN	PINN-w	LBFGS	LRA	NTK	RAR	MultiAdam	gPINN	vPINN	LAAF	GAAF	FBPINN
Burgers	1d-C	1.45E-2	2.63E-2	1.33E-2	2.61E-2	1.84E-2	3.32E-2	4.85E-2	2.16E-1	3.47E	1.43E-2	5.20E-2	2.32E-1
	2d-C	3.24E-1	2.70E-1	4.65E-1	2.60E-1	2.75E-1	3.45E-1	3.33E-1	3.27E-1	6.38E-1	2.77E-1	2.95E-1	–
Poisson	2d-C	6.94E-1	3.49E-2	NaN	1.17E-	1.23E-2	6.99E-1	2.63E-2	6.87E-1	4.91E-1	7.68E-1	6.04E-1	4.49E-2
	2d-CG	6.36E-1	6.08E-2	2.96E-1	4.34E-	1.43E-2	6.48E-1	2.76E-1	7.92E-1	2.86E-1	4.80E-1	8.71E	2.90E-2
	3d-CG	5.60E-1	3.74E-1	7.05E-1	1.02E-1	9.47E-1	5.76E-1	3.63E-1	4.85E-1	7.38E-1	5.79E-1	5.02E-1	7.39E-1
	2d-MS	6.30E-1	7.60E-1	1.45E+0	7.94E-1	7.48E-1	6.44E	5.90E-1	6.16E-1	9.72E-1	5.93E-1	9.31E-1	1.04E+0
Heat	2d-VC	1.01E+0	2.35E-1	2.32E-1	2.12E-1	2.14E-1	9.66E-1	4.75E-1	2.12E+0	9.40E-1	6.42E-1	8.49E-1	9.52E-1
	2d-MS	6.21E-2	2.42E-1	1.73E-2	8.79E-2	4.40E-2	7.49E-2	2.18E-1	1.13E-1	9.30E-1	7.40E-2	9.85E-1	8.20E-2
	2d-CG	3.64E-2	1.45E-1	8.57E-1	1.25E-1	1.16E	2.72E-2	7.12E-2	9.38E-2	–	2.39E-2	4.61E-1	9.16E-2
	2d-LT	9.99E-1	9.99E-1	1.00E+0	9.99E-1	1.00E+0	9.99E-1	1.00E+0	1.00E+0	1.00E+0	9.99E-1	9.99E-1	1.01E+0
NS	2d-C	4.70E-2	1.45E-1	2.14E-1	NaN	1.98E-1	4.69E-1	7.27E-1	7.70E-2	2.92	3.60E-2	3.79E-2	8.45E-2
	2d-CG	1.19E-1	3.26E-1	NaN	3.32E-1	2.93E-1	3.34E-1	4.31E-1	1.54E-1	9.94E-1	3.24E-2	1.74E-1	8.27E+0
	2d-LT	9.96E-1	1.00E+0	9.70E-1	1.00E+0	9.99E-1	1.00E+0	1.00E+0	9.95E-1	1.73E+0	9.98E-1	9.99E-1	1.00E+0
Wave	1d-C	5.88E-1	2.85E-1	NaN	3.61E-	9.79E-2	5.39E-1	1.21E-1	5.56E-1	8.39E-1	4.54E-1	6.77E-1	5.91E-1
	2d-CG	1.84E+0	1.66E+0	1.33E+0	1.48E+0	2.16E+0	1.15E+0	1.09E+0	8.14	7.99E-1	8.19E-1	7.94E-1	1.06E+0
	2d-MS	1.34E+0	1.02E+0	1.37E+0	1.02E+0	1.04E+0	1.35E+0	1.01E+0	1.02E+0	9.82E-1	1.06E+0	1.06E+0	1.03E+0
Chaotic	GS	3.19E-1	1.58E-1	NaN	9.37E-2	2.16E-1	9.46E	9.37E-2	2.48E-1	1.16E+0	9.47E-2	9.46E	7.99E-2
	KS	1.01E+0	9.86E-1	NaN	9.57E-1	9.64E-1	1.01E+0	9.61E-1	9.94E-1	9.72E-1	1.01E+0	1.00E+0	1.02E+0
High dim	PNd	3.04E-3	2.58E-3	4.67E-4	4.58E-4	4.64E-3	3.59E-3	3.98E-3	5.05E-3	–	4.14E-3	7.75E-3	–
	HNd	3.61E-1	4.59E-1	1.19E-4	3.94E-1	3.97E-1	3.57E-1	3.02E-1	3.17E-1	–	5.22E-1	5.21E-1	–
Inverse	PInv	9.42E-2	1.66E-1	NaN	1.54E-1	1.93E-1	9.35E-2	1.30E-	8.03E-2	2.45E-2	1.30E-1	2.54E-1	8.44E-1
	HInv	1.57E+0	5.26E-2	NaN	5.09E-2	7.52E-2	1.52E+0	8.04E-2	4.84E+0	4.56E-1	5.59E-1	2.12E-1	9.27E-1

Table 3: Mean L2RE of different PINN variants on PINNacle. Best results are highligh blue and second-places lightblue. We do not bold any result if errors of all methods are about 100% "NaN" means the method does not converge and "–" means the method is not suitable for the problem.

4 Experiments

4.1 Main Results

We now present experimental results. Except for the ablation study in Sec 4.3 and Appendix E.2, we use a learning rate of 0.001 and train all models with 20,000 epochs. We repeat all experiments three times and record the mean and std. We run all experiments on a Linux server with 20 Intel(R) Xeon(R) Silver 4210 CPUs @ 2.20GHz and eight NVIDIA GeForce RTX 2080 Ti each with 12 GB GPU memory. All experiments in Table E.1 require a total about 776 GPU hours, which can be completed in about 4 days on our cluster. Table 3 presents the main results for all methods on our tasks and shows their average ℓ_2 relative errors (with standard deviation results available in Appendix E.1).

PINN. We use PINN-w to denote training PINNs with larger boundary weights. The vanilla PINNs struggle to accurately solve complex physics systems, indicating substantial room for improvement. Using an ℓ_2 relative error (L2RE) of 10% as a threshold for a successful solution, we find that vanilla PINN only solves 10 out of 22 tasks, most of which involve simpler equations (e.g., 1.45% on Burgers-1d-C). They encounter significant difficulties when faced with physics systems characterized by complex geometries, multi-scale phenomena, nonlinearity, and longer time spans. This shows that directly optimizing an average of the PDE losses and initial/boundary condition losses leads to critical issues such as loss imbalance, suboptimal convergence, and limited expressiveness.

PINN variants. PINN variants offer approaches to addressing some of these challenges to varying degrees. Methods involving loss reweighting and resampling have shown improved performance in some cases involving complex geometries and multi-scale phenomena (e.g., 1.43% on Poisson-2d-CG). This is due to the configuration of loss weights and sampled collocation points, which adaptively place more weight on more challenging domains during the training process. However, these methods still struggle with Wave equations, Navier-Stokes equations, and other cases with higher dimensions or longer time spans. MultiAdam, a representative of novel optimizers, solves several simple cases and the chaotic GS equation (8.37%), but does not significantly outperform other methods. The new loss term of variational form demonstrates significant superiority in solving inverse problems (e.g., 1.19% on HInv for vPINN), but no clear improvement in fitting error over standard PINN in forward cases. Changes in architecture can enhance expressiveness and flexibility for cases with complex

L2RE	Name	Vanilla	Loss Reweighting/Sampling			Optimizer	Loss functions		Architecture		
–	–	PINN	LRA	NTK	RAR	MultiAdam	gPINN	vPINN	LAAF	GAAF	FBPINN
Burgers-P	2d-C	4.74E-01	4.36E-01	4.13E-01	4.71E-01	4.93E-01	4.91E-01	2.82E+0	4.37E-	4.34E-01	-
Poisson-P	2d-C	2.24E-01	7.07E-02	1.66E-02	2.33E-01	8.24E-02	4.03E-01	5.51E-1	1.84E-01	2.97E	2.87E-2
Heat-P	2d-MS	1.73E-01	1.23E-01	1.50E-01	1.53E-01	4.00E-01	4.59E-01	5.12	6.27E-02	1.89E-01	2.46E-1
NS-P	2d-C	3.89E-01	-	4.52E-01	3.91E-01	9.33E-01	7.19E-	3.76E-1	3.63E-01	4.85E-01	3.99E-1
Wave-P	1d-C	5.22E-01	3.44E-01	2.69E-01	5.05E-01	6.89E-01	7.66E-01	3.58E-1	4.03E-01	9.00E-01	1.15E+0
High dim-P	HNd	7.66E-03	6.53E-03	9.04E-03	8.07E-0	2.22E-03	7.87E-03	-	6.97E-03	1.94E-01	-

Table 4: Results of different PINN variants on parametric PDEs. We report average L2RE on all examples within a class of PDE. We hold the best results across all methods.

geometries and multi-scale systems. For example, FBPINN achieves the smallest error on the chaotic GS equation (7.99%), while LAAF delivers the best tting result on Heat-2d-CG (2.39%).

Discussion.For challenges related to complex geometries and multi-scale phenomena, some methods can mitigate these issues by implementing mechanisms like loss reweighting, novel optimizers, and better capacity through adaptive activation. This holds true for the 2D cases of Heat and Poisson equations, which are classic linear equations. However, when systems have higher dimensions (Poisson3d-CG) or longer time spans (Heat2d-LT), all methods fail to solve, highlighting the dif cul-ties associated with complex geometries and multi-scale systems.

In contrast, nonlinear, long-time PDEs like 2D Burgers, NS, and KS pose challenges for most methods. These equations are sensitive to initial conditions, resulting in complicated solution spaces and more local minima for PINNs [51]. The Wave equation, featuring a second-order time derivative and periodic behavior, is particularly hard for PINNs, which often become unstable and may violate conservation laws [24, 55]. Although all methods perform well on Poisson-Nd, only PINN with LBFGS solves Heat-Nd, indicating the potential of a second-order optimizer for solving high dimensional PDEs[53].

4.2 Parameterized PDE Experiments

To investigate whether PINNs could handle a class of PDEs, we design this experiment to solve the same PDEs with different parameters. We choose 6 PDEs, i.e., Burgers2d-C, Poisson2d-C, Heat2d-MS, NS-C, Wave1d-C, and Heat-Nd (HNd), with each case containing ve parameterized examples. Details of the parametrized PDEs are shown in Appendix B. Here we report the average L2RE metric on these parameterized PDEs for every case, and results are shown in the following Table 4. First, we see that compared with the corresponding cases in Table E.1, the mean L2RE of parameterized PDEs is usually higher. We suggest that this is because there are some dif cult cases under certain parameters for these PDEs with very high errors. Secondly, we nd that PINN-NTK works well on parameterized PDE tasks which achieve three best results among all six experiments. We speculate that solving PDEs with different parameters requires different weights for loss terms, and PINN-NTK is a powerful method for automatically balancing these weights.

4.3 Hyperparameter Analysis

The performance of PINNs is strongly affected by hyperparameters, with each variant introducing its own unique set. The results are shown in Figure 2. We focus on a set of problems, i.e., Burgers1d, GS, Heat2d-CG, and Poisson2d-C. Detailed results and additional ndings are in Appendix E.2.

Batch size and training epochs.Figure 19 presents the effects of varying batch sizes and training epochs. Larger batch sizes generally yield better outcomes due to more accurate gradient estimations, though saturation is observed beyond a batch size of 2048 for the GS and Poisson2d-C problems. Similarly, increasing the number of training epochs reduces the L2RE, indicating an improvement in model accuracy. However, this bene t plateaus around 20k to 80k epochs, where further increases in epochs do not signi cantly reduce the error.

Learning Rates. The performance of standard PINNs under various learning rates and learning rate schedules is shown in Figure 3. We observe that the in uence of the learning rate on performance is intricate, with optimal learning rates varying across problems. Furthermore, PINN training tends to be unstable. High learning rates, such as 10^{-2}, often lead to error spikes, while low learning rates,

Figure 2: Performance of vanilla PINNs under different batch sizes (number of collocation points), which is shown in the left gure; and number of training epochs, which is shown in the right gure.

Figure 3: Convergence curve of PINNs with different learning rate schedules on Burgers1d, Heat2d-CG, and Poisson2d-C.

like 10^{-5}, result in slow convergence. Our ndings suggest that a moderate learning rate, such as 10^{-3} or 10^{-4}, or a step decay learning rate schedule, tends to yield more stable performance.

5 Limitations

First, real-world problems are often more complex, with giant geometric domains or chaotic behaviors. Good performance on PINNacle does not guarantee it solves practical problems. We could explore larger-scale PINN training methods or ef cient domain decomposition methods [7]. Second, the issues of safety in the PINN methods pose potential roadblocks. Developing theoretical convergence for PINN like stability and convergence analysis [13] could help resolve these limitations.

6 Conclusion

In this work, we introduced PINNacle, a comprehensive benchmark offering a user-friendly toolbox that encompasses over 20 PDE problems and 10 PINN methods with extensive experiments and ablation studies. Looking forward, we plan to expand the benchmark by integrating additional state-of-the-art methods and incorporating more practical problem scenarios. Our analysis of the experimental results yields several key insights. First, domain decomposition is bene cial for addressing problems characterized by complex geometries, and PINN-NTK is a strong method for balancing loss weights as experiments show. Second, selecting appropriate hyperparameters is crucial to the performance of PINNs. However, the best hyperparameters usually vary with PDEs. Third, we identify high-dimensional and nonlinear problems as a pressing challenge. The overall performance of PINNs is not yet on par with traditional numerical methods [16]. Fourth, there are only a few attempts exploring of PINNs' loss landscape [30]. Finally, integrating the strengths of neural networks with numerical methods like preconditioning, weak formulation, and multigrid may present a promising avenue toward overcoming the challenges[36].

Acknowledgement

This work was supported by NSFC Projects (Nos. 62350080, 92370124, 92248303, 62276149, U2341228, 62061136001, 62076147), BNRist (BNR2022RC01006), Tsinghua Institute for Guo Qiang, and the High Performance Computing Center, Tsinghua University. J. Zhu was also supported by the XPlorer Prize.

References

[1] Simon Axelrod and Rafael Gomez-Bombarelli. Geom, energy-annotated molecular conformations for property prediction and molecular generation. Scientific Data, 9(1):185, 2022.

[2] Rafael Bischof and Michael Kraus. Multi-objective loss balancing for physics-informed deep learning. arXiv preprint arXiv:2110.09813, 2021.

[3] Aleksandar Botev, Andrew Jaegle, Peter Wirnsberger, Daniel Hennes, and Irina Higgins. Which priors matter? benchmarking models for learning latent dynamics. arXiv preprint arXiv:2111.05458, 2021.

[4] Emmanuel Boutet, Damien Lieberherr, Michael Tognolli, Michel Schneider, Parit Bansal, Alan J Bridge, Sylvain Poux, Lydie Bougueleret, and Ioannis Xenarios. Uniprotkb/swiss-prot, the manually annotated section of the uniprot knowledgebase: how to use the entry view. Plant bioinformatics: methods and protocols, pages 23–54, 2016.

[5] Salva Rühling Cachay, Venkatesh Ramesh, Jason NS Cole, Howard Barker, and David Rolnick. Climart: A benchmark dataset for emulating atmospheric radiative transfer in weather and climate models. arXiv preprint arXiv:2111.14674, 2021.

[6] DM Causon and CG Mingham. Introductory finite difference methods for PDEs. Ventus Publishing ApS, 2010.

[7] Nithin Chalapathi, Yiheng Du, and Aditi Krishnapriyan. Scaling physics-informed hard constraints with mixture-of-experts. arXiv preprint arXiv:2402.13412, 2024.

[8] Feiyu Chen, David Sondak, Pavlos Protopapas, Marios Mattheakis, Shuheng Liu, Devansh Agarwal, and Marco Di Giovanni. Neurodiffeq: A python package for solving differential equations with neural networks. Journal of Open Source Software, 5(46):1931, 2020.

[9] Salvatore Cuomo, Vincenzo Schiano Di Cola, Fabio Giampaolo, Gianluigi Rozza, Maziar Raissi, and Francesco Piccialli. Scientific machine learning through physics–informed neural networks: where we are and what's next. Journal of Scientific Computing, 92(3):88, 2022.

[10] Sourav Das and Solomon Tesfamariam. State-of-the-art review of design of experiments for physics-informed deep learning. arXiv preprint arXiv:2202.06416, 2022.

[11] Arka Daw, Anuj Karpatne, William Watkins, Jordan Read, and Vipin Kumar. Physics-guided neural networks (pgnn): An application in lake temperature modeling. arXiv preprint arXiv:1710.11431, 2017.

[12] Tobin A Driscoll, Nicholas Hale, and Lloyd N Trefethen. Chebfun guide, 2014.

[13] Francisco Eiras, Adel Bibi, Rudy Bunel, Krishnamurthy Dj Dvijotham, Philip Torr, and M Pawan Kumar. Provably correct physics-informed neural networks. arXiv preprint arXiv:2305.10157, 2023.

[14] Robert Eymard, Thierry Gallouët, and Raphaèle Herbin. Finite volume methods. Handbook of numerical analysis, 7:713–1018, 2000.

[15] Han Gao, Luning Sun, and Jian-Xun Wang. Phygeonet: Physics-informed geometry-adaptive convolutional neural networks for solving parameterized steady-state pdes on irregular domain. Journal of Computational Physics, 428:110079, 2021.

[16] Tamara G Grossmann, Urszula Julia Komorowska, Jonas Latz, and Carola-Bibiane Schönlieb. Can physics-informed neural networks beat the finite element method? arXiv preprint arXiv:2302.04107, 2023.

[17] Jayesh K Gupta and Johannes Brandstetter. Towards multi-spatiotemporal-scale generalized pde modeling. arXiv preprint arXiv:2209.15616, 2022.

[18] Zhongkai Hao, Songming Liu, Yichi Zhang, Chengyang Ying, Yao Feng, Hang Su, and Jun Zhu. Physics-informed machine learning: A survey on problems, methods and applications. arXiv preprint arXiv:2211.08064, 2022.

[19] Oliver Hennigh, Susheela Narasimhan, Mohammad Amin Nabian, Akshay Subramaniam, Kaustubh Tangsali, Zhiwei Fang, Max Rietmann, Wonmin Byeon, and Sanjay Choudhry. Nvidia simnet™: An ai-accelerated multi-physics simulation framework. In International Conference on Computational Science, pages 447–461. Springer, 2021.

[20] Zizhou Huang, Teseo Schneider, Minchen Li, Chenfanfu Jiang, Denis Zorin, and Daniele Panozzo. A large-scale benchmark for the incompressible navier-stokes equations. arXiv preprint arXiv:2112.05309, 2021.

[21] Ameya D Jagtap and George E Karniadakis. Extended physics-informed neural networks (xpinns): A generalized space-time domain decomposition based deep learning framework for nonlinear partial differential equations. In AAAI Spring Symposium: MLPS, 2021.

[22] Ameya D Jagtap, Kenji Kawaguchi, and George Em Karniadakis. Locally adaptive activation functions with slope recovery for deep and physics-informed neural networks. Proceedings of the Royal Society A, 476(2239):20200334, 2020.

[23] Ameya D Jagtap, Kenji Kawaguchi, and George Em Karniadakis. Adaptive activation functions accelerate convergence in deep and physics-informed neural networks. Journal of Computational Physics, 404:109136, 2020.

[24] Ameya D Jagtap, Ehsan Kharazmi, and George Em Karniadakis. Conservative physics-informed neural networks on discrete domains for conservation laws: Applications to forward and inverse problems. Computer Methods in Applied Mechanics and Engineering, 365:113028, 2020.

[25] George Em Karniadakis, Ioannis G Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. Nature Reviews Physics, 3(6):422–440, 2021.

[26] Ehsan Kharazmi, Zhongqiang Zhang, and George Em Karniadakis. Variational physics-informed neural networks for solving partial differential equations. arXiv preprint arXiv:1912.00873, 2019.

[27] Ehsan Kharazmi, Zhongqiang Zhang, and George Em Karniadakis. hp-vpinns: Variational physics-informed neural networks with domain decomposition. Computer Methods in Applied Mechanics and Engineering, 374:113547, 2021.

[28] Reza Khodayi-Mehr and Michael Zavlanos. Varnet: Variational neural networks for the solution of partial differential equations. In Learning for Dynamics and Control, pages 298–307. PMLR, 2020.

[29] R Khudorozhkov, S Tsimfer, and A Koryagin. Pydens framework for solving differential equations with deep learning. arXiv preprint arXiv:1909.11544, 2019.

[30] 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.

[31] Ke Li, Kejun Tang, Tianfan Wu, and Qifeng Liao. D3m: A deep domain decomposition method for partial differential equations. IEEE Access, 8:5283–5294, 2019.

[32] Yuhao Liu, Xiaojian Li, and Zhengxian Liu. An improved physics-informed neural network based on a new adaptive gradient descent algorithm for solving partial differential equations of nonlinear systems. 2022.

[33] Lu Lu, Xuhui Meng, Shengze Cai, Zhiping Mao, Somdatta Goswami, Zhongqiang Zhang, and George Em Karniadakis. A comprehensive and fair comparison of two neural operators (with practical extensions) based on fair data. Computer Methods in Applied Mechanics and Engineering, 393:114778, 2022.

[34] Lu Lu, Xuhui Meng, Zhiping Mao, and George Em Karniadakis. Deepxde: A deep learning library for solving differential equations. SIAM review, 63(1):208–228, 2021.

[35] Suryanarayana Maddu, Dominik Sturm, Christian L Müller, and Ivo F Sbalzarini. Inverse dirichlet weighting enables reliable training of physics informed neural networks. Machine Learning: Science and Technology, 3(1):015026, 2022.

[36] Stefano Markidis. The old and the new: Can physics-informed deep-learning replace traditional linear solvers? Frontiers in big Data, 4:669097, 2021.

[37] Keith W Morton and David Francis Mayers. Numerical solution of partial differential equations: an introduction. Cambridge university press, 2005.

[38] Ben Moseley, Andrew Markham, and Tarje Nissen-Meyer. Finite basis physics-informed neural networks (fbpinns): a scalable domain decomposition approach for solving differential equations. arXiv preprint arXiv:2107.07871, 2021.

[39] COMSOL Multiphysics. Introduction to comsol multiphysics®. COMSOL Multiphysics, Burlington, MA, accessed Feb, 9:2018, 1998.

[40] Mohammad Amin Nabian, Rini Jasmine Gladstone, and Hadi Meidani. Efficient training of physics-informed neural networks via importance sampling. Computer-Aided Civil and Infrastructure Engineering, 36(8):962–977, 2021.

[41] Karl Otness, Arvi Gjoka, Joan Bruna, Daniele Panozzo, Benjamin Peherstorfer, Teseo Schneider, and Denis Zorin. An extensible benchmark suite for learning to simulate physical systems. arXiv preprint arXiv:2108.07799, 2021.

[42] Evan Racah, Christopher Beckham, Tegan Maharaj, Samira Ebrahimi Kahou, Mr Prabhat, and Chris Pal. Extremeweather: A large-scale climate dataset for semi-supervised detection, localization, and understanding of extreme weather events. Advances in neural information processing systems, 30, 2017.

[43] Christopher Rackauckas and Qing Nie. Differentialequations. jl–a performant and feature-rich ecosystem for solving differential equations in julia. Journal of open research software, 5(1), 2017.

[44] 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.

[45] Junuthula Narasimha Reddy. Introduction to the finite element method. McGraw-Hill Education, 2019.

[46] Pu Ren, Chengping Rao, Yang Liu, Jian-Xun Wang, and Hao Sun. Phycrnet: Physics-informed convolutional-recurrent network for solving spatiotemporal pdes. Computer Methods in Applied Mechanics and Engineering, 389:114399, 2022.

[47] Franz M Rohrhofer, Stefan Posch, and Bernhard C Geiger. On the pareto front of physics-informed neural networks. arXiv preprint arXiv:2105.00862, 2021.

[48] Khemraj Shukla, Ameya D Jagtap, and George Em Karniadakis. Parallel physics-informed neural networks via domain decomposition. Journal of Computational Physics, 447:110683, 2021.

[49] Justin Sirignano and Konstantinos Spiliopoulos. Dgm: A deep learning algorithm for solving partial differential equations. Journal of computational physics, 375:1339–1364, 2018.

[50] Hwijae Son, Jin Woo Jang, Woo Jin Han, and Hyung Ju Hwang. Sobolev training for the neural network solutions of pdes. 2021.

[51] Sophie Steger, Franz M Rohrhofer, and Bernhard C Geiger. How pinns cheat: Predicting chaotic motion of a double pendulum. In The Symbiosis of Deep Learning and Differential Equations II, 2022.

[52] Makoto Takamoto, Timothy Praditia, Raphael Leiteritz, Daniel MacKinlay, Francesco Alesiani, Dirk P üger, and Mathias Niepert. Pdebench: An extensive benchmark for scienti c machine learning. Advances in Neural Information Processing Systems, 35:1596–1611, 2022.

[53] Kejun Tang, Xiaoliang Wan, and Chao Yang. Das: A deep adaptive sampling method for solving partial differential equations. arXiv preprint arXiv:2112.14038, 2021.

[54] Nanzhe Wang, Dongxiao Zhang, Haibin Chang, and Heng Li. Deep learning of subsurface ow via theory-guided neural network. Journal of Hydrology, 584:124700, 2020.

[55] Sifan Wang, Shyam Sankaran, and Paris Perdikaris. Respecting causality is all you need for training physics-informed neural networks. arXiv preprint arXiv:2203.07404, 2022.

[56] Sifan Wang, Shyam Sankaran, Hanwen Wang, and Paris Perdikaris. An expert's guide to training physics-informed neural networks. arXiv preprint arXiv:2308.08468, 2023.

[57] Sifan Wang, Yujun Teng, and Paris Perdikaris. Understanding and mitigating gradient ow pathologies in physics-informed neural networks. SIAM Journal on Scienti c Computing, 43(5):A3055–A3081, 2021.

[58] Sifan Wang, Hanwen Wang, and Paris Perdikaris. On the eigenvector bias of fourier feature networks: From regression to solving multi-scale pdes with physics-informed neural networks. Computer Methods in Applied Mechanics and Engineering, 384:113938, 2021.

[59] 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.

[60] 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.

[61] Jiachen Yao, Chang Su, Zhongkai Hao, Songming Liu, Hang Su, and Jun Zhu. Multiadam: Parameter-wise scale-invariant optimizer for multiscale training of physics-informed neural networks. arXiv preprint arXiv:2306.02810, 2023.

[62] Bing Yu et al. The deep ritz method: a deep learning-based numerical algorithm for solving variational problems. Communications in Mathematics and Statistics, 6(1):1–12, 2018.

[63] 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.

[64] Ruiyang Zhang, Yang Liu, and Hao Sun. Physics-informed multi-lstm networks for meta-modeling of nonlinear structures. Computer Methods in Applied Mechanics and Engineering, 369:113226, 2020.

[65] Zahari Zlatev, Ivan Dimov, István Faragó, and Ágnes Havasi. Richardson extrapolation: Practical aspects and applications, volume 2. Walter de Gruyter GmbH & Co KG, 2017.

Checklist

The checklist follows the references. Please read the checklist guidelines carefully for information on how to answer these questions. For each question, change the default [TODO] to [Yes] , [No] , or [N/A] . You are strongly encouraged to include a justi cation to your answer , either by referencing the appropriate section of your paper or providing a brief inline description. For example:

- Did you include the license to the code and datasets? [Yes] See Section 1.

Please do not modify the questions and only use the provided macros for your answers. Note that the Checklist section does not count towards the page limit. In your paper, please delete this instructions block and only keep the Checklist section heading above along with the questions/answers below.

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately re ect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [Yes]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments (e.g. for benchmarks)...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [Yes]
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [Yes]
   (e) Did you discuss whether the data you are using/curating contains personally identi able information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Overview of Appendices

We provide supplementary details about problems and experiments for the main text in the Appendix. In Appendix B, we provide mathematical descriptions and visualization for all PDEs in this paper. In Appendix C, we list the detailed hyperparameters and training/testing settings. In Appendix D, we provide a high-level overview of the codebase of the toolbox. In Appendix E, the results for the main experiments, i.e., the performance of L2RE, L1RE, MSE, and runtime for all methods on all PDEs are displayed. In Appendix F, we show the visualization results for several methods on some problems.

B Details of PDEs and Methods

Here provide details of PDE tasks used for evaluating different variants of PINNs. Denote u to be the function to solve and $x; t$ to be spatial and temporal variables.

B.1 De nitions for PDEs in main experiments

1. One-dimensional Burgers Equation (Burgers1d)

The Burgers 1D equation is given by

$$u_t + uu_x = u_{xx}: \tag{6}$$

The domain is de ned as

$$(x; t) 2 = [1; 1] [0; 1]: \tag{7}$$

The initial and boundary conditions are

$$u(x; 0) = \sin x; \tag{8}$$

$$u(1; t) = u(1; t) = 0: \tag{9}$$

The parameter is

$$= \frac{0:01}{}: \tag{10}$$

2. 2D Coupled Burgers equation (Burgers 2d)

The 2D Coupled Burgers equation is given by

$$u_t + u \quad r \; u \qquad u = 0; \tag{11}$$

$$u(0; y; t) = u(L; y; t); \quad u(x; 0; t) = u(x; L; t); \tag{12}$$

$$f x; y g 2 [0; L]; \quad t 2 [0; T]; \tag{13}$$

Figure 4: Reference solution of Burgers1d using FEM solver.

Figure 5: Reference solution of Burgers2d at timesteps $t = 0; 0.2, 0.4, 1.0$ using FEM solver.

The domain is defined as

$$(x, y, t) \in \Omega = [0, L]^2 \times [0, 1]. \tag{14}$$

The initial conditions are given by

$$w(x, y) = \sum_{i=-L}^{L} \sum_{j=-L}^{L} a_{ij} \sin(2\pi(ix + jy)) + b_{ij} \cos(2\pi(ix + jy)), \tag{15}$$

$$u(x, y, 0) = 2w(x, y) + c \tag{16}$$

where $a, b, c \sim N(0, 1)$. The parameters are

$$L = 4, \quad T = 1, \quad \nu = 0.001. \tag{17}$$

3. Poisson 2D Classic (Poisson2d-C)

The Poisson 2D equation is given by

$$\nabla u = 0. \tag{18}$$

The domain is a rectangle minus four circles $\Omega = \Omega_{rec} \setminus \bigcup R_i$ where $\Omega_{rec} = [-0.5, 0.5]^2$ is the rectangle and R_i denotes four circle areas:

$$R_1 = \{(x, y) : (x - 0.3)^2 + (y - 0.3)^2 \leq 0.1^2\}, \tag{19}$$

$$R_2 = \{(x, y) : (x + 0.3)^2 + (y - 0.3)^2 \leq 0.1^2\}, \tag{20}$$

$$R_3 = \{(x, y) : (x - 0.3)^2 + (y + 0.3)^2 \leq 0.1^2\}, \tag{21}$$

$$R_4 = \{(x, y) : (x + 0.3)^2 + (y + 0.3)^2 \leq 0.1^2\}. \tag{22}$$

The boundary condition is

$$u = 0, x \in \partial R_i, \tag{23}$$

$$u = 1, x \in \partial \Omega_{rec}. \tag{24}$$

Figure 6: Reference solution of Poisson2d-CG by FEM solver.

4. Poisson-Boltzmann (Helmholtz) 2D Irregular Geometry (Poisson2d-CG)

The Poisson-Boltzmann (Helmholtz) 2D equation is given by

$$u + k^2 u = f(x;y): \tag{25}$$

The function $f(x)$ is de ned as

$$f(x) = A \sum_i \left({}^2_i + x_i^2 \right) \sin({}_1 x_1) \sin({}_2 x_2): \tag{26}$$

The domain is $[1;1]^2$ and the boundary conditions are

$$u = 0:2; \quad x 2 @ {}_{rec}; \tag{27}$$

$$u = 1; \quad x 2 @ {}_{circle}: \tag{28}$$

Parameter references are

$$ {}_1 = 1; \quad {}_2 = 4; \quad k = 8; \quad A = 10: \tag{29}$$

The domain is $[1;1]^2$ with several circles removed. The circles ${}_{circle} = [{}^4_{i=1} R_i$ are

$$R_1 = f(x;y) : (x \quad 0:5)^2 + (y \quad 0:5)^2 \quad 0:2^2 g \tag{30}$$

$$R_2 = f(x;y) : (x \quad 0:4)^2 + (y + 0:4)^2 \quad 0:4^2 g \tag{31}$$

$$R_3 = f(x;y) : (x + 0:2)^2 + (y + 0:7)^2 \quad 0:1^2 g \tag{32}$$

$$R_4 = f(x;y) : (x + 0:6)^2 + (y \quad 0:5)^2 \quad 0:3^2 g \tag{33}$$

5. Poisson 3D Complex Geometry with Two Domains (Poisson3d-CG)

The Poisson 3D equation with two domains is given by

$$ {}_i \quad u + k_i^2 u = f(x;y;z); \quad i = 1;2: \tag{34}$$

The function $f(x;y;z)$ is de ned as

$$\begin{aligned} f(x;y;z) = A_1 \frac{\exp(\sin m_1 x \ + \sin m_2 y \ + \sin m_3 z \ )}{x^2 + y^2 + z^2 + 1} (x^2 + y^2 + z^2 \quad 1) \\ + A_2 \sin(m_1 x \ ) \sin(m_2 y \ ) \sin(m_3 z \ ): \end{aligned} \tag{35}$$

The coef cients are de ned as $\begin{array}{l} = \ {}_1; \quad k = k_1; \quad x 2 \quad {}_1; \\ = \ {}_2; \quad k = k_2; \quad x 2 \quad {}_2: \end{array}$

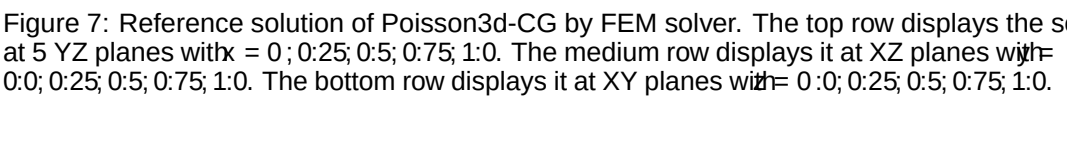

Figure 7: Reference solution of Poisson3d-CG by FEM solver. The top row displays the solution at 5 YZ planes with $x = 0; 0.25, 0.5, 0.75, 1.0$. The medium row displays it at XZ planes with $y = 0.0, 0.25, 0.5, 0.75, 1.0$. The bottom row displays it at XY planes with $z = 0.0, 0.25, 0.5, 0.75, 1.0$.

The boundary condition is

$$\frac{@u}{@n} = 0; \quad x \; 2 \; @ : \tag{36}$$

The domains and other parameters are de ned as follows:

$$_1 = [0; 1] \quad [0; 1] \quad [0; 0.5] = [\; _{i=1}^{4} R_i; \tag{37}$$

$$_2 = [0; 1] \quad [0; 1] \quad [0.5; 1] = [\; _{i=1}^{4} R_i : \tag{38}$$

The circular regions R_i are

$$R_1 = f(x; y; z) : (x \quad 0.4)^2 + (y \quad 0.3)^2 + (z \quad 0.6)^2 \quad 0.2^2 g \tag{39}$$

$$R_1 = f(x; y; z) : (x \quad 0.6)^2 + (y \quad 0.7)^2 + (z \quad 0.6)^2 \quad 0.2^2 g \tag{40}$$

$$R_1 = f(x; y; z) : (x \quad 0.2)^2 + (y \quad 0.8)^2 + (z \quad 0.7)^2 \quad 0.1^2 g \tag{41}$$

$$R_1 = f(x; y; z) : (x \quad 0.6)^2 + (y \quad 0.2)^2 + (z \quad 0.3)^2 \quad 0.1^2 g \tag{42}$$

$$\tag{43}$$

Other parameters are

$$m_1 = 1; m_2 = 10; m_3 = 5; \quad _1 = 1; \quad _2 = 1; k_1 = 8; k_2 = 10; A_1 = 20; A_2 = 100: \tag{44}$$

6. 2D Poisson equation with many subdomains (Poisson2d-MS)

The PDE and boundary condition is given by

$$r \quad (a(x)r \; u) = f(x; y); x \text{in} \tag{45}$$

$$\frac{@u}{@n} + u = 0; x \; 2 \; @ : \tag{46}$$

Here the domain is $(x; y) \; 2 \; = [\quad 10; 10]^2$. We divide the whole domain into many small squares, and $a(x)$ is a piecewise linear function in each square. We store $a(x)$ in a le in practical implementation.

Figure 8: Reference solution of Poisson2d-MS by FEM solver.

Figure 9: Reference solution of Heat2d-VC by FEM solver at times $t = 0; 0.5; 2.0; 3.5$.

7. 2D Heat with Varying Coef cients (Heat2d-VC)

The 2D heat equation with a varying source is given by

$$\frac{@u}{@t} \quad r \quad (a(x)r \ u) = f(x;t): \tag{47}$$

The domain is $T = [0;1]^2 \quad [0;5]$. The function $a(x)$ is chosen similarly to Darcy ow but with an exponential GRF. The function $f(x;t)$ is de ned as

$$f(x;t) = A \sin(m_1 \ x) \sin(m_2 \ y) \sin(m_3 \ t): \tag{48}$$

with $A = 200; m_1 = 1; m_2 = 5; m_3 = 1$. The initial and boundary conditions are

$$u(x;y;0) = 0; x \ 2 \tag{49}$$
$$u(x;y;t) = 0; x \ 2 \ @: \tag{50}$$

8. 2D Heat Multi-Scale (Heat2d-MS)

The 2D heat multi-scale equation is given by

$$\frac{@u}{@t} \quad \frac{1}{(500 \ )^2} u_{xx} \quad \frac{1}{2} u_{yy} = 0; \tag{51}$$

with domain $T = [0;1]^2 \quad [0;5]$.

The initial and boundary conditions are

$$u(x;y;0) = \sin(20 \ x) \sin(\ y); \quad x \ 2 \quad ; \tag{52}$$
$$u(x;y;t) = 0; \quad x \ 2 \ @: \tag{53}$$

Figure 10: Reference solution of Heat2d-MS by FEM solver.

Figure 11: reference solution of Heat2d-CG by FEM solver at times $t = 0.5; 2.0; 2.5; 3.0$.

9. 2D Heat Complex Geometry (Heat Exchanger, Heat2d-CG)

The 2D heat equation for a complex geometry is given by

$$\frac{@u}{@t} \quad u = 0: \tag{54}$$

The domain is de ned as $T = ([\ 8; 8] \quad [\ 12; 12] \, n \, [\ _i R_i) \quad [0; 3]$.

The boundary condition is

$$n \ (\ cr \ u) = g \quad qu: \tag{55}$$

Here we choose $c = 1$. The positions of large circles are

$$(\ 4; \ 3); \quad (\ 4; \ 9); \quad (0; 0); \quad (0; \ 6); \quad r = 1 \tag{56}$$

with $g = 5$ and $q = 1$. The positions of small circles are

$$(\ 3:2; \ 6); \quad (\ 3:2; 0); \quad r = 0:4 \tag{57}$$

with $g = 1$ and $q = 1$. For the rectangular boundary conditions, $g = 0:1$ and $q = 1$.

10. 2D Heat Long Time (Heat2d-LT)

The governing PDE is

$$\frac{@u}{@t} = 0:001 \quad u + 5 \sin(\ ku^2) \quad 1 + 2 \sin \quad \frac{t}{4} \qquad \sin(m_1 \ x \ ) \sin(m_2 \ y \ ) \tag{58}$$

with domain $T = [0 \ ; 1]^2 \quad [0; 100]$, $m_1 = 4$, $m_2 = 2$, and $k = 1$.

The initial and boundary conditions are given by

$$u(x; y; 0) \ = \ \sin(4 \ x \ ) \sin(3 \ y \ ); x \ 2 \tag{59}$$

$$u(x; y; t) \ = \ 0 \ ; \quad x \ 2 \ @ : \tag{60}$$

11. 2D NS lid-driven ow (NS2d-C).

Figure 12: reference solution of Heat2d-LT by FEM solver at time step $0; 20; 50; 80; 100$.

Figure 13: Reference solution of NS2d-Ld by FEM solver.

The PDE is given by

$$u \quad r\ u + r\ p \quad \frac{1}{Re} \quad u \quad = \quad 0; x \ 2 \tag{61}$$

$$r \quad u \quad = \quad 0; x \ 2 \tag{62}$$

The domain is $= [0 \ ; 1]^2$, the top boundary is $_1$, the left, right and bottom boundary is $_2$.

The boundary conditions are

$$u(x) \quad = \quad (4\ x(1 \quad x); 0); x \ 2 \quad _1 \tag{63}$$

$$u(x) \quad = \quad (0\ ; 0); x \ 2 \quad _2 \tag{64}$$

$$p \quad = \quad 0; x = (0\ ; 0): \tag{65}$$

The Reynolds number $Re = 100$.

12. 2D Back Step Flow (NS-CG)

The equations and boundary conditions are given by

$$u \quad r\ u + r\ p \quad \frac{1}{Re} \quad u \quad = \quad 0; \tag{66}$$

$$r \quad u \quad = \quad 0: \tag{67}$$

The domain is de ned as $= [0 \ ; 4] \quad [0; 2] \ n \ ([0; 2] \quad [1; 2] \quad \bigcup R_i)$ (excluding the top-left quarter).