

PirateNets: Physics-informed Deep Learning with Residual Adaptive Networks

Sifan Wang

SIFAN.WANG@YALE.EDU

*Institution for Foundation of Data Science
Yale University
New Haven, CT 06520*

Bowen Li

BOWEN.LI@CITYU.EDU.HK

*Department of Mathematics
City University of Hong Kong
Hong Kong SAR*

Yuhan Chen

YCHEN239@NCSU.EDU

*Department of Electrical and Computer Engineering
North Carolina State University
Raleigh, NC 27695*

Paris Perdikaris

PGP@SEAS.UPENN.EDU

*Department of Mechanical Engineering and Applied Mechanics
University of Pennsylvania
Philadelphia, PA 19104*

Editor: Animashree Anandkumar

Abstract

While physics-informed neural networks (PINNs) have become a popular deep learning framework for tackling forward and inverse problems governed by partial differential equations (PDEs), their performance is known to degrade when larger and deeper neural network architectures are employed. Our study identifies that the root of this counter-intuitive behavior lies in the use of multi-layer perceptron (MLP) architectures with non-suitable initialization schemes, which result in poor trainability for the network derivatives, and ultimately lead to an unstable minimization of the PDE residual loss. To address this, we introduce Physics-Informed Residual Adaptive Networks (PirateNets), a novel architecture that is designed to facilitate stable and efficient training of deep PINN models. PirateNets leverage a novel adaptive residual connection, which allows the networks to be initialized as shallow networks that progressively deepen during training. We also show that the proposed initialization scheme allows us to encode appropriate inductive biases corresponding to a given PDE system into the network architecture. We provide comprehensive empirical evidence showing that PirateNets are easier to optimize and can gain accuracy from considerably increased depth, ultimately achieving state-of-the-art results across various benchmarks. All code and data accompanying this manuscript will be made publicly available at <https://github.com/PredictiveIntelligenceLab/jaxpi/tree/pirate>.

Keywords: Deep learning; Physics-informed neural networks; Partial differential equations; Computational physics; Neural solvers

1 Introduction

Machine learning (ML) is making a continual impact on the fields of science and engineering, providing advanced tools for analyzing complex data, uncovering nonlinear relationships, and developing predictive models. Notable examples include AlphaFold (Jumper et al., 2021) in protein structure prediction, Deep Potentials (Zhang et al., 2018) for large-scale molecular dynamics, and GraphCast (Lam et al., 2022) for medium-range weather forecasting. The integration of physical laws and constraints within machine learning models has given rise to physics-informed machine learning (PIML). This emerging field opens new frontiers for traditional scientific research and contributes to addressing persistent challenges in machine learning, such as robustness, interpretability, and generalization.

The fundamental question that PIML aims to address is how to incorporate physical prior knowledge into machine learning models. This can be achieved by modifying key components of the machine learning pipeline, which includes data processing, model architecture, loss functions, optimization algorithms, and fine-tuning and inference. For example, neural operators (Lu et al., 2021; Li et al., 2020) are designed to encode physics from training data that reflects underlying physical laws. Another powerful approach is to tweak network architectures to strictly enforce physical constraints, thereby enhancing the generalizability of the model. One method involves embedding general physical principles such as symmetries, invariances and equivariances into ML models. As highlighted by Cohen *et al.* (Cohen et al., 2019) and Maron *et al.* (Maron et al., 2019, 2018), this strategy leads to simpler models that require less data while achieving higher prediction accuracy. Alternatively, embedding specific physical laws can be equally effective. This is demonstrated by Mohan *et al.* (Mohan et al., 2020), who ensured continuity and mass conservation in neural networks for coarse-graining three-dimensional turbulence by integrating the curl operator. Similarly, Meng *et al.* (Darbon et al., 2020) illustrates that certain neural network architectures are inherently aligned with the physics of specific Hamilton-Jacobi (HJ) partial differential equations (PDEs).

One of the most popular and flexible methods for embedding physical principles in machine learning is through the formulation of tailored loss functions. These loss functions serve as soft constraints that bias ML models towards respecting the underlying physics during training, giving rise to the emergence of physics-informed neural networks (PINNs). Thanks to their flexibility and ease of implementation, PINNs have been extensively used to solve forward and inverse problems involving PDEs by seamlessly integrating noisy experimental data and physical laws into the learning process. In recent years, PINNs have yielded a series of promising results across various domains in computational science, including applications in fluid mechanics (Raissi et al., 2020; Sun et al., 2020; Mathews et al., 2021), bio-engineering (Sahli Costabal et al., 2020; Kissas et al., 2020), and material science (Fang and Zhan, 2019; Chen et al., 2020; Zhang et al., 2022). Furthermore, PINNs have been effectively applied in molecular dynamics (Islam et al., 2021), electromagnetics (Kovacs et al., 2022; Fang, 2021), geosciences (Haghighat and Juanes, 2021; Smith et al., 2022), and in designing thermal systems (Hennigh et al., 2020; Cai et al., 2021).

Although recent studies have showcased some empirical success with PINNs, they also highlight several training pathologies. These include spectral bias (Rahaman et al., 2019; Wang et al., 2021b), unbalanced back-propagated gradients (Wang et al., 2021a, 2022c), and causality violation (Wang et al., 2022a), all of which represent open areas for research

and methodological development. To address these issues, numerous studies have focused on enhancing PINNs’ performance by refining neural network architectures and training algorithms. Notable efforts are loss re-weighting schemes (Wang et al., 2021a, 2022c; McClenny and Braga-Neto, 2020; Maddu et al., 2022) and adaptive resampling of collocation points, such as importance sampling (Nabian et al., 2021), evolutionary sampling (Daw et al., 2022), and residual-based adaptive sampling (Wu et al., 2023). Simultaneously, significant strides have been made in developing new neural network architectures to improve PINNs’ representational capacity, including adaptive activation functions (Jagtap et al., 2020), positional embeddings (Liu et al., 2020; Wang et al., 2021b), and innovative architectures (Wang et al., 2021a; Sitzmann et al., 2020; Gao et al., 2021; Fathony et al., 2021; Moseley et al., 2021; Kang et al., 2022). Further explorations include alternative objective functions, such as those employing numerical differentiation techniques (Chiu et al., 2022; Huang and Alkhalifah, 2024) and variational formulations inspired by Finite Element Methods (FEM) (Kharazmi et al., 2021; Patel et al., 2022), along with additional regularization terms to accelerate PINNs convergence (Yu et al., 2022; Son et al., 2021). The evolution of training strategies also forms an active area of research, with sequential training (Wight and Zhao, 2020; Krishnapriyan et al., 2021) and transfer learning (Desai et al., 2021; Goswami et al., 2020; Chakraborty, 2021) showing promise in accelerating learning and improving accuracy.

Despite significant progress in recent years, most existing works on PINNs tend to employ small, and shallow network architectures, leaving the vast potential of deep networks largely untapped. To bridge this gap, we put forth a novel class of architectures coined Physics-Informed residual adaptive networks (PirateNets). Our main contributions are summarized as follows:

- We argue that the capacity of PINNs to minimize PDE residuals is determined by the capacity of network derivatives.
- We support this argument by proving that, for second-order linear elliptic and parabolic PDEs, the convergence in training error leads to the convergence in the solution and its derivatives.
- We empirically and theoretically reveal that conventional initialization schemes (e.g., Glorot or He initialization) result in problematic initialization of MLP derivatives and thus worse trainability.
- We introduce PirateNets to address the issue of pathological initialization, enabling stable and efficient scaling of PINNs to utilize deep networks. The proposed physics-informed initialization in PirateNets also serves as a novel method for integrating physical priors at the model initialization stage.
- We conduct comprehensive numerical experiments demonstrating that PirateNets achieve consistent improvements in accuracy, robustness, and scalability across various benchmarks.

The rest of this paper is organized as follows. We first give a concise overview of physics-informed neural networks in Section 2, following the original formulation of Raissi *et al.* (Raissi et al., 2019). In Section 3, we take the Allen-Cahn equation as a motivating example to illustrate the growing instability in training PINNs with deep neural networks and then

propose a potential rationale (mechanism): *if a PINN model is trained sufficiently, then both its predictions and derivatives should well approximate the solution of the underlying PDE and its derivatives, respectively.* We theoretically justify this claim for linear elliptic and parabolic PDEs and, by investigating the capabilities of MLP derivatives, find that deep MLP derivatives often suffer from ill-posed initializations, which could result in an unstable training process. To address the aforementioned issues, in Section 4 we propose a novel framework of PirateNets. In Section 5, comprehensive numerical experiments are performed to validate the proposed architecture. In Section 6, we conclude this work with a discussion of potential future research directions.

2 Physics-informed neural network

In this section, we will briefly overview the standard formulation of PINNs, following the one in (Raissi et al., 2019). Without loss of generality, we consider an abstract PDE of the parabolic form:

$$\mathbf{u}_t + \mathcal{D}[\mathbf{u}] = \mathbf{f}, \quad (1)$$

defined on a spatial-temporal domain $[0, T] \times \Omega \subset \mathbb{R}^{1+d}$, where Ω is a bounded domain in \mathbb{R}^d with regular enough boundary $\partial\Omega$, $\mathcal{D}[\cdot]$ is a linear or nonlinear differential operator, and $\mathbf{u}(t, \mathbf{x})$ denotes a unknown solution.

The general initial and boundary conditions can be then formulated as:

$$\mathbf{u}(0, \mathbf{x}) = \mathbf{g}(\mathbf{x}), \quad \mathbf{x} \in \Omega, \quad (2)$$

$$\mathcal{B}[\mathbf{u}] = 0, \quad t \in [0, T], \quad \mathbf{x} \in \partial\Omega. \quad (3)$$

Here, $\mathbf{f}(t, \mathbf{x})$ on $[0, T] \times \Omega$ and $\mathbf{g}(\mathbf{x})$ on Ω are given functions with certain regularity such that PDE (1) is well-posed; $\mathcal{B}[\cdot]$ denotes an abstract boundary operator, representing various boundary conditions such as Dirichlet, Neumann, Robin, and periodic conditions.

We aim at approximating the unknown solution $\mathbf{u}(t, \mathbf{x})$ by a deep neural network $\mathbf{u}_\theta(t, \mathbf{x})$, where θ denotes the set of all trainable parameters of the network (e.g., weights and biases). If a smooth activation function is employed, \mathbf{u}_θ provides a smooth representation that can be queried for any (t, \mathbf{x}) . Moreover, all required gradients with respect to input variables or network parameters θ can be computed via automatic differentiation (Griewank and Walther, 2008). This allows us to define the interior PDE residuals as

$$\mathcal{R}_{\text{int}}[\mathbf{u}_\theta](t, \mathbf{x}) = \frac{\partial \mathbf{u}_\theta}{\partial t}(t, \mathbf{x}) + \mathcal{D}[\mathbf{u}_\theta](t, \mathbf{x}) - \mathbf{f}(\mathbf{x}), \quad (t, \mathbf{x}) \in [0, T] \times \Omega, \quad (4)$$

and spatial and temporal boundary residuals, respectively, by

$$\mathcal{R}_{\text{bc}}[\mathbf{u}_\theta](t, \mathbf{x}) = \mathcal{B}[\mathbf{u}_\theta](t, \mathbf{x}), \quad (t, \mathbf{x}) \in [0, T] \times \partial\Omega, \quad (5)$$

and

$$\mathcal{R}_{\text{ic}}[\mathbf{u}_\theta](\mathbf{x}) = \mathbf{u}_\theta(0, \mathbf{x}) - \mathbf{g}(\mathbf{x}), \quad \mathbf{x} \in \Omega. \quad (6)$$

Then, we train a physics-informed model by minimizing the following composite *empirical loss*:

$$\mathcal{L}(\theta) := \underbrace{\frac{1}{N_{ic}} \sum_{i=1}^{N_{ic}} |\mathcal{R}_{ic}[\mathbf{u}_\theta](\mathbf{x}_{ic}^i)|^2}_{\mathcal{L}_{ic}(\theta)} + \underbrace{\frac{1}{N_{bc}} \sum_{i=1}^{N_{bc}} |\mathcal{R}_{bc}[\mathbf{u}_\theta](t_{bc}^i, \mathbf{x}_{bc}^i)|^2}_{\mathcal{L}_{bc}(\theta)} + \underbrace{\frac{1}{N_r} \sum_{i=1}^{N_r} |\mathcal{R}_{int}[\mathbf{u}_\theta](t_r^i, \mathbf{x}_r^i)|^2}_{\mathcal{L}_r(\theta)}, \quad (7)$$

which aims to enforce the neural network function \mathbf{u}_θ to satisfy the PDEs (1) with initial and spatial boundary conditions (2)–(3). The training data points $\{\mathbf{x}_{ic}^i\}_{i=1}^{N_{ic}}$, $\{t_{bc}^i, \mathbf{x}_{bc}^i\}_{i=1}^{N_{bc}}$ and $\{t_r^i, \mathbf{x}_r^i\}_{i=1}^{N_r}$ can be the vertices of a fixed mesh or points randomly sampled at each iteration of a gradient descent algorithm. We finally remark that although we only review the case of parabolic-type equations here, similar discussions can also be applied to general elliptic and hyperbolic (linear or nonlinear) equations.

3 Initialization pathologies in PINNs

3.1 Challenges in Training PINNs

Training Physics-informed neural networks (PINNs) has proven more challenging than conventional neural networks, with several key issues identified in recent research. Wang *et al.* (Wang et al., 2021b) demonstrated that PINNs suffer from spectral bias, struggling to learn complex, high-frequency functions. One simple and effective way to resolve this is to use random Fourier feature embeddings (Tancik et al., 2020). In addition, PINNs face unbalanced losses, resulting in bias towards minimizing certain loss terms while ignoring others. To address this, various self-adaptive weighting schemes (Wang et al., 2021a, 2022c) have been proposed to balance the convergence rate of individual loss terms during training. Furthermore, Wang *et al.* (Wang et al., 2022a) revealed an implicit bias suggesting that continuous-time PINNs models can violate causality, and hence are susceptible to converge towards erroneous solutions. To mitigate this issue, several causal and curriculum training algorithms (Wang et al., 2022a; Wight and Zhao, 2020; Krishnapriyan et al., 2021) have been developed to encourage the PINNs model to learn PDE solutions in accordance with how the information propagates in time, as the dynamics evolve throughout the spatio-temporal domain.

However, despite these advances, the practical deployment of PINNs has been largely limited to employing small and shallow networks, typically with 5 layers or less. This is in sharp contrast to the exceptional capabilities showcased by deeper networks in the broader field of deep learning (Devlin et al., 2018; Radford et al., 2021; Rombach et al., 2022).

In the following, we will both empirically and theoretically demonstrate that the trainability of PINNs degrades as the network depth increases, particularly when employing multi-layer perceptrons (MLPs) as the backbone architecture. To this end, let us first focus on the Allen-Cahn equation, a challenging benchmark for conventional PINN models that has been extensively studied in recent literature (Wight and Zhao, 2020; Wang et al., 2022a; Es’ kin et al., 2023; Daw et al., 2022; Anagnostopoulos et al., 2023). For simplicity, we




Figure 1: *Allen-Cahn equation*: Relative L^2 error of training PINNs using MLP, ResNet, and PirateNet backbones of varying depths, averaged over 5 random seeds for each architecture.

consider the one-dimensional case with a periodic boundary condition:

$$\begin{aligned} u_t - 0.0001u_{xx} + 5u^3 - 5u &= 0, \quad t \in [0, 1], \quad x \in [-1, 1], \\ u(0, x) &= x^2 \cos(\pi x), \\ u(t, -1) &= u(t, 1), \quad u_x(t, -1) = u_x(t, 1). \end{aligned}$$

In our experimental setup, we adhere closely to the training pipeline described in Wang *et al.* (Wang et al., 2023a). Specifically, we employ random Fourier feature networks with 256 neurons in each hidden layer and hyperbolic tangent (Tanh) activation functions. The models are trained with a batch size of 1,024 over 10^5 steps of gradient descent using the Adam optimizer (Kingma and Ba, 2014). We set the initial learning rate to 10^{-3} , and use an exponential decay rate of 0.9 every 2,000 steps. In particular, the training of our models incorporates learning rate annealing (Wang et al., 2021a, 2022a) and causal training algorithms (Wang et al., 2022a, 2023a) to enhance the model performance and robustness.

Figure 1 shows the resulting relative L^2 errors when training PINNs with MLPs of varying depths. The result is averaged over 5 different random seeds. It can be observed that the prediction error becomes larger as the network depth increases. Eventually, this error almost reaches 100% for PINN models with 18 hidden layers, indicating a complete failure of the model to yield a reasonable solution.

One may argue that the challenges in training deep neural networks can be effectively resolved by incorporating skip connections (He et al., 2016a). To test this hypothesis, our study compares the performance of MLP networks with residual connections (ResNets) under identical hyper-parameter settings. As depicted in the same figure, the overall trend remains consistent: the error proportionally increases with network depth, albeit at a slightly slower rate. This implies that, contrary to expectations, the conventional skip connection fails to address the fundamental issues in training PINNs using deeper networks. These findings highlight a critical limitation in PINN performance when the network depth exceeds a certain threshold. To address this challenge, we propose a novel network architecture in Section 4, referred to as Physics-Informed Residual AdapTive Networks (PirateNets), designed for efficient and stable training of deep PINN models. As can be seen in Figure 1, PirateNets consistently reduces the predictive error as the network depth increases.

3.2 Interior Error Estimates of PINNs for Linear PDEs

To gain deeper insights into the reasons behind the poor performance of MLPs and ResNets, we next examine the initialization process of PINNs. Due to the variety and nonlinearity in PDE expressions, a direct analysis of the initialization of the PDE residual network (i.e., $\mathcal{R}_{\text{int}}[\mathbf{u}_\theta]$ in (4)) presents significant challenges. Here, in view of the role of the PDE residual loss in (7) (i.e., enforce the network to satisfy the underlying PDEs), it is reasonable to expect that if the training loss is small, then the derivatives of the network should also closely approximate those of the actual solution. This motivates us to analyze the derivatives of the network to understand the initialization process, instead of examining the full residual network graph $\mathcal{R}_{\text{int}}[\mathbf{u}_\theta]$. To be specific, we put forth the following informal claim, which can be rigorously justified in certain scenarios.

Claim 1 (informal) *Let $\mathbf{u}(t, \mathbf{x})$ be the solution to a given PDE of n -th order in space and m -th order in time, and $\mathbf{u}_\theta(t, \mathbf{x})$ denote a approximate solution by a PINN model with parameters θ , where $\mathbf{x} = (x_1, x_2, \dots, x_d)$. We refer to a k -th order spatial or temporal derivative of \mathbf{u}_θ as the MLP derivative network. Then, if the training loss $\mathcal{L}(\theta)$ converges to 0, then the MLP derivative networks will converge to corresponding solution derivatives:*

$$\frac{\partial^k \mathbf{u}_\theta}{\partial x_i^k}(t, \mathbf{x}) \longrightarrow \frac{\partial^k u}{\partial x_i^k}(t, \mathbf{x}), \quad \frac{\partial^l \mathbf{u}_\theta}{\partial t^l}(t, \mathbf{x}) \longrightarrow \frac{\partial^l u}{\partial t^l}(t, \mathbf{x}), \quad (8)$$

in a certain norm, for $k = 1, 2, \dots, n$, $l = 1, 2, \dots, m$, and $i = 1, 2, \dots, d$.

In fact, we show in Propositions 1 and 2 below that the above convergence (8) indeed holds locally for linear elliptic and parabolic PDEs in some Sobolev spaces. We remark that the global error estimates for PINNs in Sobolev norms have been explored very recently for various PDEs (Zeinhofer et al., 2023). Nevertheless, our results mainly focus on the interior estimates for PINNs, which are not covered by (Zeinhofer et al., 2023) and better support our claim and further analysis. This is because interior estimates typically yield higher regularity (Hörmander, 1958), facilitating the examination of exact network derivatives, while global estimates would give lower regularity due to the L^2 boundary loss.

For ease of expositions, we first introduce some notations. Let $\Omega \subset \mathbb{R}^d$ be a bounded open set, and define the partial differential operator:

$$\mathcal{D}[u] := - \sum_{i,j=1}^n (a^{ij}(\mathbf{x})u_{x_i})_{x_j} + \sum_{i=1}^n b^i(\mathbf{x})u_{x_i} + c(\mathbf{x})u, \quad (9)$$

with smooth coefficients $a^{ij}(\mathbf{x})$, $b^i(\mathbf{x})$, and $c(\mathbf{x})$. We assume that \mathcal{D} is uniformly elliptic, i.e., for some constant $c > 0$, $\sum_{i,j=1}^n a^{ij}(\mathbf{x})\xi_i\xi_j \geq c|\xi|^2$ for a.e. $x \in \Omega$ and all $\xi \in \mathbb{R}^n$.

We also define $B_r(\mathbf{x})$ by the open ball in \mathbb{R}^d of radius r centered at \mathbf{x} and $Q_r(t, \mathbf{x}) := (t, t + r^2) \times B_r(\mathbf{x})$. For a time–space domain $U \subset \mathbb{R}^{d+1}$, we denote the set $C^{r,k+2r}(\bar{U})$ of bounded continuous functions $u(t, x)$ such that their derivatives $\partial_t^\rho \partial_{\mathbf{x}}^\alpha u$ for $|\alpha| + 2\rho \leq k + 2r$ and $\rho \leq r$ are bounded and continuous in U , and can be extended to \bar{U} . For $u \in C^{r,k+2r}(\bar{U})$, we define

$$\|u\|_{W_2^{r,k+2r}(U)} := \sum_{\substack{|\alpha|+2\rho \leq k+2r \\ \rho \leq r}} \|\partial_t^\rho \partial_{\mathbf{x}}^\alpha u\|_{L^2(U)}.$$

Then, the Sobolev space $W_2^{r,k+2r}(U)$ is given by the completion of $C^{r,k+2r}(\bar{U})$ with respect to $\|\cdot\|_{W_2^{r,k+2r}(U)}$.

Proposition 1 *Consider the second-order elliptic Dirichlet problem:*

$$\mathcal{D}[u] = f,$$

with Dirichlet boundary condition and $f \in L^2(\Omega)$. Let $u : \Omega \rightarrow \mathbb{R}$ be its solution and u_θ be a smooth approximation by PINNs. Define the expected loss function:

$$\hat{\mathcal{L}}(\theta) = \int_{\partial\Omega} |u_\theta|^2 \, d\mathbf{x} + \int_{\Omega} |\mathcal{D}[u_\theta] - f|^2 \, d\mathbf{x}. \quad (10)$$

Then, for any compact $V \subset\subset \Omega$, there exist a constant C such that

$$\|u - u_\theta\|_{H^2(V)}^2 \leq C \hat{\mathcal{L}}(\theta). \quad (11)$$

We emphasize that the above proposition is based on the expected loss (10), while the practical training process employs the empirical loss as in (7):

$$\mathcal{L}(\theta) = \frac{1}{N_{bc}} \sum_{i=1}^{N_{bc}} |u_\theta(\mathbf{x}_{bc}^i)|^2 + \frac{1}{N_r} \sum_{i=1}^{N_r} |\mathcal{D}[u_\theta](t_r^i, \mathbf{x}_r^i) - f(t_r^i, \mathbf{x}_r^i)|^2.$$

In general, we have that when the numbers of randomly sampled training points N_{bc} and N_r tend to infinity, $\mathcal{L}(\theta)$ converges to the corresponding $\hat{\mathcal{L}}(\theta)$. We refer the interested readers to (Mishra and Molinaro, 2022, Section 2.4.1) for the detailed analysis. This means that if the network is sufficiently trained, i.e., the loss $\mathcal{L}(\theta)$ is small, the PINN solution u_θ can locally approximate u in the H^2 norm, which, by Sobolev embedding theorem, further implies that in the case of $d = 1$, the derivative $\frac{d}{dx}u_\theta(x)$ approximates $\frac{d}{dx}u(x)$ pointwisely. A similar result can be found in (Zeinhofer et al., 2023, Lemma 4.1). We next state the analogous result for parabolic PDEs, which complements the discussions in (Zeinhofer et al., 2023, Section 4.5).

Proposition 2 *Consider a second-order parabolic equation with Dirichlet boundary condition:*

$$\begin{aligned} u_t + \mathcal{D}[u] &= f, & [0, T] \times \Omega, \\ u &= 0, & [0, T] \times \partial\Omega, \\ u &= g, & \{t = 0\} \times \Omega, \end{aligned}$$

with $f \in L^2([0, T]; L^2(\Omega))$ and $g \in L^2(\Omega)$. Let $u : [0, T] \times \Omega \rightarrow \mathbb{R}$ be its solution and u_θ be a smooth approximation by PINNs. Define the expected loss:

$$\hat{\mathcal{L}}(\theta) = \int_{\Omega} |\mathcal{R}_{ic}[u_\theta](\mathbf{x})|^2 \, d\mathbf{x} + \int_0^T \int_{\partial\Omega} |u_\theta(t, \mathbf{x})|^2 \, dt d\mathbf{x} + \int_0^T \int_{\Omega} |\mathcal{R}_{int}[u_\theta](t, \mathbf{x})|^2 \, dt d\mathbf{x},$$

with $\mathcal{R}_{int}[u_\theta]$ and $\mathcal{R}_{bc}[u_\theta]$ given in (4) and (6), respectively. Then, for any set $Q_r(t, \mathbf{x}) \subset\subset [0, T] \times \Omega$, there exist a constant C such that

$$\|u - u_\theta\|_{W_2^{1,2}(Q_r)}^2 \leq C \hat{\mathcal{L}}(\theta).$$

The proofs of Propositions 1 and 2 can be found in Appendix A, which are essentially the applications of interior estimates of PDEs. Recalling from (Hörmander, 1958) that the solutions of a formally hypoelliptic equation $P(x, D)[u] = f$ are infinitely differentiable if the coefficients in the differential operator $P(x, D)$ and f are infinitely differentiable, one may expect similar results hold for a fairly large class of equations. However, for some evolution equations, e.g., wave equation and Burgers equation, the singularity of the initial data can propagate into the interior of the domain so that the L^2 residual loss would be insufficient to control the higher-order derivatives of the solution, and hence some Sobolev loss may need to be used (Yu et al., 2022; Son et al., 2021). The following proposition discusses the case of the second-order hyperbolic equation, from which we see that in order to obtain the convergence of the first-order derivative of the solution, the H^1 temporal boundary residual is employed. Its proof follows from standard regularity result of hyperbolic equations (Evans, 2022, Section 7.2, Theorem 5); see also (Lasiecka et al., 1986, Theorem 4.1).

Corollary 3 *Consider the hyperbolic initial/boundary-value problem:*

$$\begin{aligned} u_{tt} + \mathcal{D}[u] &= f, & [0, T] \times \Omega, \\ u &= 0, & [0, T] \times \partial\Omega, \\ u &= g, \quad u_t = h, & \{t = 0\} \times \Omega, \end{aligned}$$

with $f \in L^2([0, T]; L^2(\Omega))$, $g \in H_0^1(\Omega)$, and $h \in L^2(\Omega)$. Let u be its solution and u_θ be a smooth approximation by PINNs. Define the expected loss:

$$\begin{aligned} \hat{\mathcal{L}}(\theta) &= \|(u_\theta)_{tt} + \mathcal{D}[u_\theta] - f\|_{L^2([0, T]; L^2(\Omega))}^2 + \|(u_\theta)_t - h\|_{L^2(\Omega)}^2 \\ &\quad + \|u_\theta(0, \cdot) - g\|_{H^1(\Omega)}^2 + \|u_\theta\|_{L^2([0, T]; H^1(\partial\Omega))}^2 + \|u_\theta\|_{H^1([0, T]; L^2(\partial\Omega))}^2. \end{aligned}$$

Then, we have for some constant $C > 0$,

$$\sup_{0 \leq t \leq T} (\|u_\theta - u\|_{H^1(\Omega)}^2 + \|u_t\|_{L^2(\Omega)}^2) \leq C \hat{\mathcal{L}}(\theta).$$

3.3 Analysis of MLP Derivatives

With the results in Section 3.2, we narrow our focus on the analysis of the network derivatives. Without loss of generality, we consider an MLP with scalar inputs and outputs. Specifically, let $x \in \mathbb{R}$ be the input coordinate, $\mathbf{g}^{(0)}(x) = x$ and $d_0 = d_{L+1} = 1$. An MLP $u_\theta(x)$ is recursively defined as follows

$$\mathbf{u}_\theta^{(l)}(x) = \mathbf{W}^{(l)} \cdot \mathbf{g}^{(l-1)}(x) + \mathbf{b}^{(l)}, \quad \mathbf{g}^{(l)}(x) = \sigma(\mathbf{u}_\theta^{(l)}(x)), \quad l = 1, 2, \dots, L, \quad (12)$$

with a final linear layer defined by

$$u_\theta(x) = \mathbf{W}^{(L+1)} \cdot \mathbf{g}^{(L)}(x) + \mathbf{b}^{(L+1)}, \quad (13)$$

where $\mathbf{W}^{(l)} \in \mathbb{R}^{d_l \times d_{l-1}}$ is the weight matrix in l -th layer and σ is an element-wise activation function. Here, $\theta = (\mathbf{W}^{(1)}, \mathbf{b}^{(1)}, \dots, \mathbf{W}^{(L+1)}, \mathbf{b}^{(L+1)})$ represents all trainable parameters of the network. In particular, we suppose that the network is equipped with Tanh activation

functions and all weight matrices are initialized by the Glorot initialization scheme (Glorot and Bengio, 2010), i.e., $\mathbf{W}^{(l)} \sim \mathcal{N}(0, \frac{2}{d_l+d_{l-1}})$ for $l = 1, 2, \dots, L + 1$. Moreover, all bias parameters are initialized to zeros. These are standard practices in the vast majority of existing PINN implementations. In this setting, we have the following proposition with the proof given in Appendix A.

Proposition 4 *Given the above setup, the following holds at initialization of an MLP $u_\theta(x)$.*

(a) *The first-order derivative of u_θ is given by*

$$\frac{\partial u_\theta}{\partial x}(x) = \prod_{i=L+1}^2 \left(\mathbf{W}^{(i)} \cdot \text{diag} \left(\dot{\sigma}(\mathbf{u}_\theta^{(i-1)}(x)) \right) \right) \cdot \mathbf{W}^{(1)}. \quad (14)$$

(b) *If all hidden layers of the MLP have the same number of neurons, i.e., $d_1 = d_2 = \dots = d_L = d$, then*

$$\text{Var} \left(\frac{\partial u_\theta}{\partial x}(x) \right) \lesssim \frac{1}{d}. \quad (15)$$

Consequently, for any $\epsilon > 0$,

$$\mathbb{P} \left(\left| \frac{\partial u_\theta}{\partial x}(x) \right| > \epsilon \right) \lesssim \frac{1}{d\epsilon^2}. \quad (16)$$

From this proposition, it becomes evident that, in contrast to the conventional forward pass of MLPs where the weight matrices are hidden inside the activation functions, all weight matrices directly contribute to the final output of $\frac{\partial u_\theta}{\partial x}$.

To gain further insights, we may assume that MLP lies in a linear regime i.e., $\sigma(x) \approx x$ and $\dot{\sigma}(\cdot) \approx 1$. This is a good start point for analyzing network behaviors at initialization (see (Glorot and Bengio, 2010)), since the inputs to each activation function are relatively small and most activations (e.g., Sin, Tanh, ReLU, Sigmoid) have a slope of 1 at 0. Under this assumption, it follows that:

$$\frac{\partial u_\theta}{\partial x}(x) \approx \mathbf{W}^{(L+1)} \cdot \mathbf{W}^{(L)} \dots \mathbf{W}^{(1)}. \quad (17)$$

Consequently, $\frac{\partial u_\theta}{\partial x}$ behaves as a deep linear network at initialization. Such networks have limited expressivity and are more vulnerable to vanishing and exploding gradients. More critically, this expression is independent of the input coordinates at initialization, therefore preventing $\frac{\partial u_\theta}{\partial x}$ from effectively approximating the solution derivatives in the early stage of training.

Also note that the variance of $\frac{\partial u_\theta}{\partial x}$ is bounded by the network’s width, which is independent of the network depth. As a consequence, $\frac{\partial u_\theta}{\partial x}$ tends to be closer to a constant in wider MLPs. This implies that the network is closer to a trivial zero solution at the initialization, which is particularly unfavorable for solving homogeneous PDEs. Although our theoretical analysis mainly focuses on first-order derivatives and the Tanh activation function, subsequent numerical evidence suggests that these insights extend to higher-order derivatives and other common activation functions.

To illustrate these findings through a concrete numerical test, we consider a simple regression problem for approximating a sinusoidal function $y(x) = \sin(2\pi x)$ for $x \in [0, 1]$. We




Figure 2: *Regression*: *Left*: Variance of the network derivative equipped with different activations across various network widths at initialization. *Middle*: Variance of MLP derivatives of different orders across various network width at initialization. *Right*: Relative L^2 error in approximating $y(x) = \sin(2\pi x)$ with MLP derivatives of different orders. All statistics are averaged over 5 random seeds.

generate a training dataset $\{(x_i, y_i)\}_{i=1}^N$ by evaluating the target function y over a uniform grid of size $N = 256$. Our goal is to investigate the behavior of MLPs and their associated derivative networks at initialization and during training. First, we fix the depth of the MLP at 3 hidden layers and examine $\text{Var}(\frac{u_\theta}{\partial x})$ for various widths and activations at initialization. As shown on the left panel of Figure 2, the variance decays exponentially with increasing width, not only for Tanh but also for GeLU and sine activations. Moreover, in the middle panel of Figure 2, we show the variance of different order derivatives of the MLP at initialization. It can be seen that the second and fourth-order derivatives of the MLP exhibit a similar trend as those of the first-order derivatives. These observations are consistent with Proposition 4.

Next, we study the capability of MLP derivative networks of different order to approximate the target function. The computational graph of the k -th derivative network $\frac{\partial^k u_\theta}{\partial x^k}$ can be readily obtained via automatic differentiation, allowing to easily train such networks by minimizing a standard mean squared error loss given by

$$\mathcal{L}(\theta) = \frac{1}{N} \sum_{i=1}^N \left| \frac{\partial^k u_\theta}{\partial x^k}(x_i) - y_i \right|^2, \quad k = 1, 2, 4. \quad (18)$$

For our experiments, we fix the base MLP width at 128, employ Tanh activation functions, and study the effect of various network depths on predictive accuracy. All models are trained for 10^4 steps of full-batch gradient descent using the Adam optimizer. The initial learning rate is 10^{-3} , following an exponential decay with a decay rate of 0.9 every 1,000 steps. The resulting relative L^2 errors are summarized on the right panel of Figure 2. We can see that, in comparison to the MLP output u itself, higher-order derivatives of deeper MLPs tend to yield larger errors. Consequently, we may conclude that the trainability of MLP derivative networks is worse than the base MLP representing u_θ itself. This result serves as a proxy to indicate that evaluating and minimizing the outputs of a PDE residual network $\mathcal{R}_{\text{int}}[\mathbf{u}_\theta]$ is considerably more difficult and unstable than minimizing a conventional supervised loss for fitting \mathbf{u}_θ to data. This also emphasizes a critical challenge in designing and optimizing deep and wide PINNs, particularly in terms of their effectiveness in approximating solution derivatives.

4 Physics-informed residual adaptive networks (PirateNets)

In this section, we introduce Physics-Informed Residual AdapTive Networks (PirateNets), a novel architecture designed to address the initialization issue elaborated in Section 3. Figure 3 below illustrates the main modules of the forward pass of a PirateNet. In detail, the input coordinates \mathbf{x} are first mapped into a high dimensional feature space by an embedding function $\Phi(\mathbf{x})$. In this work, we employ random Fourier features (Tancik et al., 2020):

$$\Phi(\mathbf{x}) = \begin{bmatrix} \cos(\mathbf{B}\mathbf{x}) \\ \sin(\mathbf{B}\mathbf{x}) \end{bmatrix},$$

where each entry in $\mathbf{B} \in \mathbb{R}^{m \times d}$ is i.i.d. sampled from a Gaussian $\mathcal{N}(0, s^2)$ with the standard deviation $s > 0$ being a user-specified hyper-parameter. Such an embedding has been extensively validated for reducing spectral bias, a common pathology in PINNs where the network struggles to learn high-frequency components of the solution. This bias often leads to poor convergence of PINNs, particularly when the underlying solution exhibits multiscale features or sharp transients. By transforming the input space into higher-frequency signals, Fourier embeddings effectively improve the eigenfunction frequency of the resulting Neural Tangent Kernel (NTK), thus enabling a more efficient learning of complex solutions (Wang et al., 2021b).

Then, the embedded coordinates $\Phi(\mathbf{x})$ are sent into two dense layers:

$$\mathbf{U} = \sigma(\mathbf{W}_1\Phi(\mathbf{x}) + \mathbf{b}_1), \quad \mathbf{V} = \sigma(\mathbf{W}_2\Phi(\mathbf{x}) + \mathbf{b}_2),$$

where σ denotes a point-wise activation function. These two encoding maps act as gates in each residual block of the architecture. This step is motivated by (Wang et al., 2021a) and has been widely used to enhance the trainability of MLPs and improve the convergence of PINNs (Wang et al., 2023a; Anagnostopoulos et al., 2023).

Let $x^{(1)} = \Phi(x)$ and $\mathbf{x}^{(l)}$ denote the input of the l -th block for $1 \leq l \leq L$. The forward pass in each PirateNet block is defined by the following iterations:

$$\mathbf{f}^{(l)} = \sigma(\mathbf{W}_1^{(l)}\mathbf{x}^{(l)} + \mathbf{b}_1^{(l)}), \tag{19}$$

$$\mathbf{z}_1^{(l)} = \mathbf{f}^{(l)} \odot \mathbf{U} + (1 - \mathbf{f}^{(l)}) \odot \mathbf{V}, \tag{20}$$

$$\mathbf{g}^{(l)} = \sigma(\mathbf{W}_2^{(l)}\mathbf{z}_1^{(l)} + \mathbf{b}_2^{(l)}), \tag{21}$$

$$\mathbf{z}_2^{(l)} = \mathbf{g}^{(l)} \odot \mathbf{U} + (1 - \mathbf{g}^{(l)}) \odot \mathbf{V}, \tag{22}$$

$$\mathbf{h}^{(l)} = \sigma(\mathbf{W}_3^{(l)}\mathbf{z}_2^{(l)} + \mathbf{b}_3^{(l)}), \tag{23}$$

$$\mathbf{x}^{(l+1)} = \alpha^{(l)}\mathbf{h}^{(l)} + (1 - \alpha^{(l)})\mathbf{x}^{(l)}, \tag{24}$$

where \odot denotes a point-wise multiplication and $\alpha^{(l)} \in \mathbb{R}$ is a trainable parameter. All the weights are initialized by the Glorot scheme (Glorot and Bengio, 2010), while biases are initialized to zero.

The final output of a PirateNet of L residual blocks is given by

$$\mathbf{u}_\theta = \mathbf{W}^{(L+1)}\mathbf{x}^{(L)}. \tag{25}$$

It is worth noting from (19)–(24) that each residual block consists of three dense layers with two gating operations, followed by an adaptive residual connection across the stacked layers. Thus, a PirateNet of L residual blocks will have a depth of $3L$ and its total number of trainable parameters is comparable to that of an MLP of the same depth.

A key aspect of PirateNets is the trainable parameters $\alpha^{(l)}$ in the skip connections, which determine the nonlinearity of the l -th block. Specifically, when $\alpha^{(l)} = 0$, $\mathbf{x}^{(l+1)} = \mathbf{x}^{(l)}$ by (24), meaning that in this case, the l -th block is an identity map. In contrast, when $\alpha^{(l)} = 1$, the mapping becomes fully nonlinear without any shortcuts. Throughout all experiments in this work, we initialize $\alpha^{(l)}$ to zero for all blocks, leading to the final output of a PirateNet \mathbf{u}_θ being a linear combination of the first layer embeddings at initialization, i.e.,

$$\mathbf{u}_\theta(\mathbf{x}) = \mathbf{W}^{(L+1)}\Phi(\mathbf{x}). \quad (26)$$

By initializing α to zero, we effectively begin with a one-layer network, mitigating the challenges associated with training deep residual networks at initialization. This approach circumvents the poor initialization of deep PINNs discussed in Section 3. As training progresses, the network gradually increases in depth through the learned α parameter, allowing for a smooth transition from shallow to deep architectures. This strategy enhances training stability and convergence, while retaining the capacity to capture complex PDE solutions.

In fact, the PDE solution can be simply approximated by a small shallow network or even a linear combination of some basis, similar to spectral and finite element methods. The rationale behind employing deep neural networks lies in leveraging additional nonlinearities to minimize the PDE residuals, thereby enabling the network to learn the solution and its derivatives accordingly.

From equation (26), another key observation is that PirateNets can be viewed as a linear combination of basis functions at initialization. This not only allows one to control the inductive bias of the network by appropriate choice of basis, but also enables the integration of various types of existing data into the initialization phase of the network. Specifically, given a set of solution measurements, denoted as $\mathbf{Y} = \{y_i\}_{i=1}^n$, then one can readily initialize the last linear layer of the model by the following least square problem:

$$\min_{\mathbf{W}} \|\mathbf{W}\Phi - \mathbf{Y}\|_2^2. \quad (27)$$

As a result, PirateNets offers an optimal initial guess, based on the available data, in the L^2 sense. It is important to note that the data for this initialization could come from a variety of sources, including experimental measurements, initial and boundary conditions, solutions derived from surrogate models, or by approximating the solution of the linearized PDE. Indeed, we can apply the same initialization procedure to any network architecture with a linear final layer (see Equation (13)), while a major consideration is the potentially limited expressivity of randomly initialized basis to accurately fit the data. Taken together, our proposed approach paves a new way to incorporate physical priors into machine learning pipelines through appropriate network initialization.

Connections to Prior Work. It is important to acknowledge that the adaptive residual connection proposed in our work is not entirely novel. Similar concepts have been previously

introduced by Bresson *et al.* (Savarese and Figueiredo, 2017), who proposed the so-called *Residual Gates* as follows:

$$\begin{aligned} \mathbf{y} &= \text{ReLU}(\alpha) \cdot (\mathcal{F}(\mathbf{x}) + \mathbf{x}) + (1 - \text{ReLU}(\alpha)) \cdot \mathbf{x}, \\ &= \text{ReLU}(\alpha) \cdot \mathcal{F}(\mathbf{x}) + \mathbf{x}, \end{aligned} \tag{28}$$

where \mathcal{F} represents a residual function and \mathbf{x} and \mathbf{y} denote the input and output of a layer, respectively. While our approach is similar, a subtle yet fundamental difference lies in our initialization $\alpha = 0$ in Equation (24), as opposed to their initialization $\alpha = 1$ in Equation (28). Consequently, Bresson *et al.*'s model is initialized as an original ResNet, while ours is initialized as an identity mapping. Our strategy serves to initialize the model as a shallow neural network for early-phase training, thus avoiding the direct training of deep PINN models which suffer from initialization pathologies, as discussed in Section 3. The significant impact of this initialization choice is clearly illustrated in the ablation studies presented in Figures 10 and 13, which confirm the necessity of initializing α as zero for achieving optimal performance in PINNs. Additionally, He *et al.* (He et al., 2016a,b) has extensively explored various residual connection variants. These typically place the activation after the addition. In contrast, PirateNets requires placing the activation before the addition to ensure an exact identity mapping throughout the entire forward pass by setting $\alpha = 0$ at initialization.

Our proposed initialization technique also shares similarities with the Physics-informed Extreme Learning Machine (PIELM) (Dwivedi and Srinivasan, 2020; Dong and Li, 2021). Both methods focus on initializing the output layer weights using the least squares method. However, PIELM is designed to solve the PDE in a single step without a training process, leading to three important limitations. First, PIELM typically employs shallow randomly initialized networks, potentially limiting their capacity to represent complex solutions effectively. Second, tackling nonlinear problems becomes more challenging due to the nonlinearity of the associated least squares problem, which is notably sensitive to the initial guess and susceptible to local minima. Third, the application of PIELM is limited as it cannot be used for inverse problems. In contrast, PirateNets do not attempt to solve the PDE in one shot. Instead, our goal is to use a linear least squares approach to initialize the model as the best possible approximation based on our available knowledge and/or data. We expect PirateNets to correct and learn the solution during the training phase. An advantage of our method is its simplicity and efficiency in solving the linear least squares problem, which guarantees a global minimum and requires minimal computational resources. Moreover, the flexibility in choosing an appropriate embedding mapping ensures accurate data reconstruction, which greatly benefits from the dedicated design of the adaptive skip connections in PirateNets.

Remark. In PirateNets, random Fourier features are our primary choice for constructing coordinate embeddings. However, it is worth mentioning that there are other feasible options, such as Chebyshev polynomials, radial basis functions, etc. Another alternative is the spectral representation, where frequencies are uniformly sampled on a grid. Despite not being detailed in this paper, we have explored this type of embedding as well. Interestingly, our experiments reveal that compared to random Fourier features, this spectral representation can reconstruct the data with greater accuracy. Yet, it also showed a greater tendency towards overfitting, characterized by a decrease in loss but an increase in the test error. Therefore, we choose to use random Fourier features for the coordinate embedding in PirateNets.




Figure 3: *Physics-informed residual adaptive networks (PirateNets)*: In our model, input coordinates are first projected into a high-dimensional feature space using random Fourier features, then followed by passing through N adaptive residual blocks. Each block consists of three dense layers, augmented with two gating operations that incorporate shallow latent features. The key module of the architecture is the adaptive skip connection with a trainable parameter α initialized at 0, so that at the initialization phase, each block reduces to an identity mapping, and the model can be viewed as a linear combination of the coordinate embeddings. This approach enables the training of a shallow network initially, thereby circumventing the issue of pathological initialization in deep PDE residual networks. We propose a physics-informed initialization for the final layer by solving a least squares problem to fit the available data, while all other weights are initialized following the Glorot scheme, and biases are set to zero. As training progresses, the depth of the model gradually increases as the nonlinearities become activated, enabling the model to progressively recover its approximation capacity.

5 Results

In this section, we demonstrate the effectiveness of the proposed PirateNet architecture across a diverse collection of benchmarks. To establish a strong baseline, we adopt the training pipeline and recommended hyper-parameters from Wang *et al.* (Wang et al., 2023a). Specifically, we employ random Fourier features (Tancik et al., 2020) as coordinate embedding and Modified MLP (Wang et al., 2021a, 2023a) of width 256 as the baseline architecture backbone. Additionally, all weight matrices are enhanced using random weight factorization (RWF) (Wang et al., 2022b). The Tanh activation function is employed by default, unless specified otherwise. Exact periodic boundary conditions are applied as required (Dong and Ni, 2021).

For model training, we use mini-batch gradient descent with the Adam optimizer (Kingma and Ba, 2014), where collocation points are randomly selected in each iteration. Our learning rate schedule includes an initial linear warm-up phase of 5,000 steps, starting from zero and gradually increasing to 10^{-3} , followed by an exponential decay at a rate of 0.9.

Following the recommended training procedure (Wang et al., 2023a), we employ several techniques to address common pathologies in PINNs:

1. Learning rate annealing algorithm (Wang et al., 2021a, 2022c, 2023a): This technique helps mitigate the problem of imbalanced gradients, a pathology where different terms in the loss function converge at vastly different rates. By dynamically adjusting weights

Benchmark	PirateNet	JAX-PI (Wang et al., 2023a)
Allen-Cahn	2.24×10^{-5}	5.37×10^{-5}
Korteweg–De Vries	4.27×10^{-4}	1.96×10^{-3}
Grey-Scott	3.61×10^{-3}	6.13
Ginzburg-Landau	1.49×10^{-2}	3.20×10^{-2}
Lid-driven cavity flow (Re=3200)	4.21×10^{-2}	1.58×10^{-1}

Table 1: State-of-the-art relative L^2 test error for various PDE benchmarks using the PirateNet architecture.

for different components of the loss function, we ensure a more balanced optimization process, leading to improved convergence and stability.

2. Causal Training (Wang et al., 2022a, 2023a): This approach addresses the causality violation pathology often encountered in using PINNs to solve time-dependent PDEs. In standard training, PINNs may inevitably use future information to predict past states, violating physical causality and thereby steering towards erroneous solutions. Causal training enforces PINNs to learn the solution progressively along the time dimension by introducing additional temporal weights. This explicitly respects the causal structure inherent in the solutions of general nonlinear PDEs, ensuring physically consistent results.

For all examples, we compare the proposed PirateNet against our baseline under exactly the same hyper-parameter settings. Our main results are summarized in Table 1. For ease of replication, we detail all hyper-parameters used in our experiments in Tables 4, 5, 6, 7, and 8. For the physics-informed initialization, we employ the `jax.numpy.linalg.lstsq` routine to solve the associated least square problem, which returns the solution with the smallest 2-norm. In cases where the data comes from an initial condition $u_0(x)$, we initialize the weights of the last layer to fit $u_\theta(t, x)$ to $u_0(x)$ for all t . We find that this approach leads to a more stable training process and better results than only fitting the initial conditions at $t = 0$.

Moreover, we conduct comprehensive ablation studies to validate the proposed network architecture and its components. Specifically, these studies focus on: (a) the scalability of PirateNet, (b) the functionality of adaptive residual connection, (c) the efficiency of gating operators, and (d) the effectiveness of physics-informed initialization. For each ablation study, it is worth emphasizing that all compared models are trained under exactly the same hyper-parameter settings and the results are averaged over five random seeds. It is also noteworthy that the computational cost of training PirateNet is comparable to that of training a Modified MLP of equivalent depth. For a more detailed quantitative analysis, readers are referred to Wang *et al.* (Wang et al., 2023a). The code and data for this study will be made publicly available at <https://github.com/PredictiveIntelligenceLab/jaxpi/tree/pirate>.

5.1 Allen-Cahn equation

To demonstrate the effectiveness of PirateNets, let us revisit the Allen-Cahn benchmark we introduced in Section 3. Here we employ a PirateNet with 3 residual blocks (9 hidden layers) and follow the optimal hyper-parameter settings used in JAX-PI (Wang et al., 2023a) for training. The results are shown in Figure 4. As shown at the top panel, the model prediction is in excellent agreement with the reference solution, achieving a relative L^2 error of 2.24×10^{-5} . Moreover, Table 2 summarizes the performance of various PINNs variants on this benchmark, indicating that our method achieves the best result ever reported in the PINNs literature for this example. Additionally, the bottom panel visualizes the training loss, test error, and the nonlinearity α of each residual block during training. We observe that the PirateNet yields faster convergence compared to a modified MLP backbone. Interestingly, after a rapid change in the early training phase, the learned nonlinearities α stabilize around $\mathcal{O}(10^{-2})$ and exhibit minimal variation thereafter. This may imply that for this example, our model efficiently approximates the PDE solution and minimizes the PDE residual with a relatively modest degree of nonlinearity.

Furthermore, we perform a series of ablation studies to further investigate the performance of PirateNets. First, we explore the impact of various data sources used for the physics-informed initialization of the PirateNet. We examine two different approaches: (a) fitting the initial conditions, and (b) fitting the solution of the linearized PDE:

$$\begin{aligned} u_t - 0.0001u_{xx} &= 0, \quad t \in (0, 1), x \in (-1, 1), \\ u(0, x) &= x^2 \cos(\pi x), \end{aligned}$$

under periodic boundary conditions. We evaluate the test error resulting from using these two approaches for the initialization of the last linear layer of the PirateNet. The findings, presented on the left panel of Figure 5, indicate comparable performance of both cases. This result aligns with our expectations, given the negligible discrepancy in the solution of the aforementioned linearized PDE compared to the initial condition across the time domain, due to the small diffusivity of 0.0001.

Second, we study the influence of the physics-informed initialization across different network architectures. As mentioned in Section 4, any network with a linear final layer can be viewed as a linear combination of basis functions, and is therefore amenable to our proposed physics-informed initialization technique. In this context, we compared the Modified MLP and PirateNet and present the results in the middle panel of Figure 5. It can be observed that the proposed physics-informed initialization can improve the performance of both architectures and notably, the PirateNet outperforms the Modified MLP baseline, irrespective of the initialization method.

Third, we evaluate the effect of network depth on predictive accuracy for different architectures. The results are presented on the right panel of Figure 5, revealing a clear difference between Modified MLP and PirateNet. Contrary to the increasing error observed in deeper Modified MLP networks, PirateNets demonstrate a consistent improvement in accuracy as the network depth increases. This trend not only underscores PirateNet’s superior performance but also highlights its enhanced scalability and robustness.

Method	Relative L^2 error
Original formulation of Raissi <i>et al.</i> (Raissi et al., 2019)	4.98×10^{-1}
Adaptive time sampling (Wight and Zhao, 2020)	2.33×10^{-2}
Self-attention (McClenny and Braga-Neto, 2020)	2.10×10^{-2}
Time marching (Mattey and Ghosh, 2022)	1.68×10^{-2}
Causal training (Wang et al., 2022a)	1.39×10^{-4}
Dirac delta function causal training (Es' kin et al., 2023)	6.29×10^{-5}
JAX-PI (Wang et al., 2023a)	5.37×10^{-5}
RBA-PINNs (Anagnostopoulos et al., 2023)	4.55×10^{-5}
PirateNet (Ours)	2.24×10^{-5}

Table 2: *Allen-Cahn equation*: Relative L^2 test errors obtained by different PINNs variants.

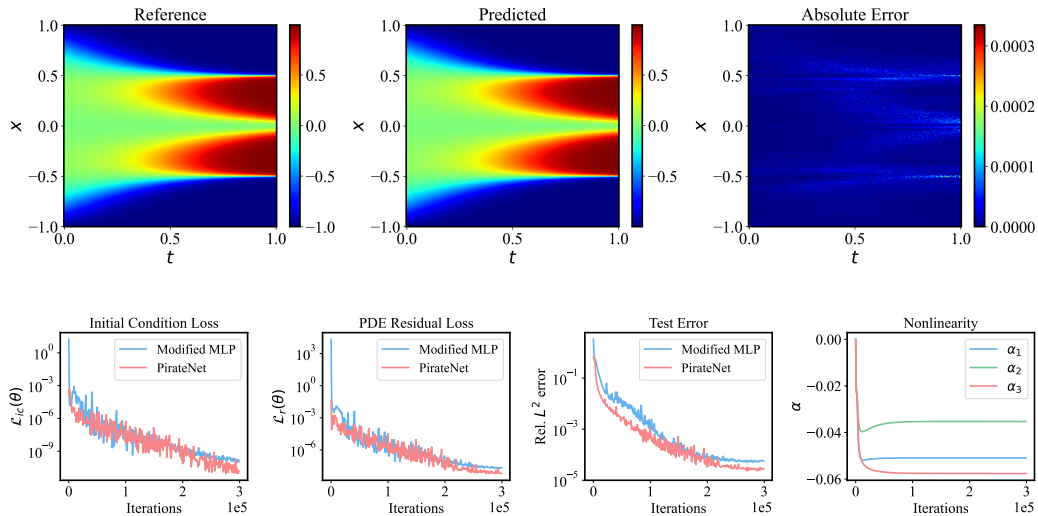


Figure 4: *Allen-Cahn equation*: *Top*: Comparison between the solution predicted by a trained PirateNet and the reference solution. The detailed hyper-parameter settings are presented in Table 4. *Bottom*: Convergence of the initial condition loss, the PDE residual loss, and the relative L^2 test error during the training of a PirateNet and a Modified MLP backbone, alongside the evolution of nonlinearities in each of the PirateNet residual blocks.

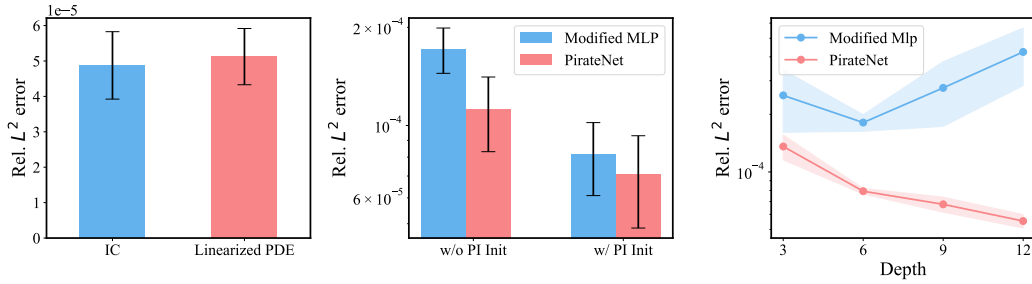


Figure 5: *Allen-Cahn equation*: *Left*: Relative L^2 test errors obtained by a PirateNet with the last layer initialized by the least square solution for fitting the initial condition and the linearized PDE solution, respectively. *Middle*: Relative L^2 test errors of training a Modified MLP and a PirateNet backbone with or without the physics-informed initialization. Without physics-informed initialization, the final layer defaults to a standard dense layer with weights initialized using the Glorot scheme and biases set to zero. *Right*: Relative L^2 errors of training a Modified MLP and a PirateNet backbone of different depth. Each ablation study is performed under the same hyper-parameter settings, with results averaged over 5 random seeds.

5.2 Korteweg–De Vries equation

Next, we explore the one-dimensional Korteweg–De Vries (KdV) equation, a fundamental model used to describe the dynamics of solitary waves, or solitons. The KdV equation is expressed as follows:

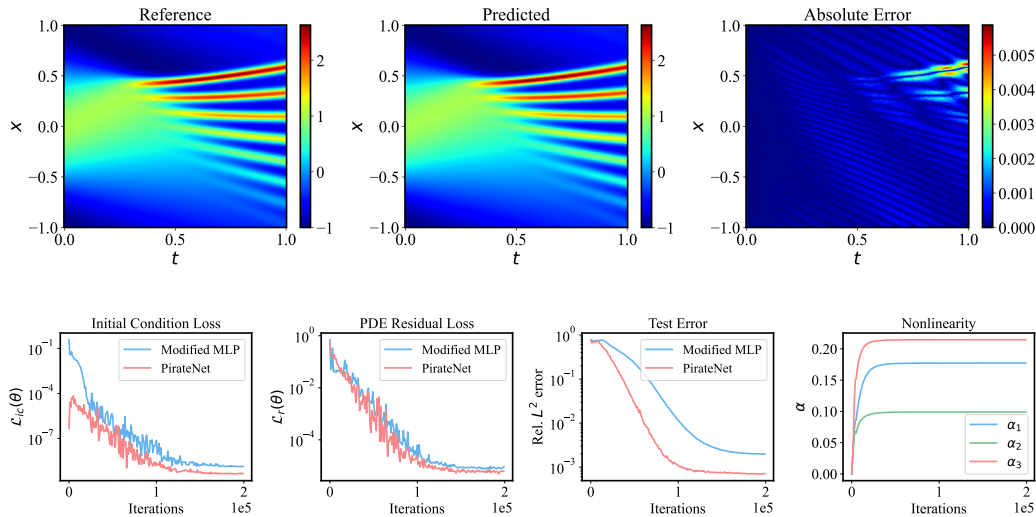
$$\begin{aligned} u_t + \eta uu_x + \mu^2 u_{xxx} &= 0, \quad t \in (0, 1), \quad x \in (-1, 1), \\ u(x, 0) &= \cos(\pi x), \\ u(t, -1) &= u(t, 1), \end{aligned}$$

where η governs the strength of the nonlinearity, while μ controls the dispersion level. Under the KdV dynamics, this initial wave evolves into a series of solitary-type waves. For our study, we adopt the classical parameters of the KdV equation, setting $\eta = 1$ and $\mu = 0.022$ (Zabusky and Kruskal, 1965).

We employ a PirateNet with 3 residual blocks (9 hidden layers) and train the model following the training pipeline and recommended hyper-parameter settings used in JAX-PI (Wang et al., 2023a). The top panel of Figure 6 showcases the predicted solution from the trained model which exhibits an excellent agreement with the numerical solution, resulting in a relative L^2 test error of 4.27×10^{-4} . For further context, in Table 3 we provide a comparison of test errors obtained from competing PINNs approaches on the same benchmark. Impressively, our method outperforms the current state-of-the-art (Es’kin et al., 2023) by approximately an order of magnitude.

Moreover, in the bottom panel of Figure 6 we record the loss and test error obtained by a PirateNet and a Modified MLP backbone, as well as the nonlinearity parameter (α) attained in each PirateNet residual block during training. Consistent with the trends seen in our first example, the PirateNet yields a faster convergence in terms of both loss and test errors compared to modified MLPs. Intriguingly, in this case, the learned nonlinearities (α)

Method	Relative L^2 error
Dirac delta function causal training (Es' kin et al., 2023)	2.45×10^{-3}
JAX-PI (Wang et al., 2023a)	1.96×10^{-3}
PirateNet (Ours)	4.27×10^{-4}

Table 3: *Korteweg–De Vries equation*: Relative L^2 errors obtained by different approaches.Figure 6: *Korteweg–De Vries equation*: *Top*: Comparison between the solution predicted by a trained PirateNet and the reference solution. The detailed hyper-parameter settings are presented in Table 5. *Bottom*: Convergence of the initial condition loss, the PDE residual loss, and the relative L^2 error during the training of a PirateNet and a Modified MLP backbone, alongside the evolution of nonlinearities in each PirateNet residual block.

reach an order of magnitude around 0.1, in contrast to only 0.01 in the first example. This indicates that the problem is inherently more nonlinear and requires a correspondingly more complex nonlinear mapping for accurate approximation.

To further validate our conclusions, we conduct the same ablation studies as in our first example. The results are summarized in Figure 7. The left and the middle panel of the figure reinforce our earlier conclusion that using data from initial conditions or linearized PDEs leads to similar predictive accuracy and the physics-informed initialization can benefit both PirateNet and Modified MLP. However, it is important to note that PirateNet significantly outperforms Modified MLP for this example, even in the absence of physics-informed initialization. Finally, the right panel of the figure further affirms the scalability of the PirateNet architecture, as observed in the first example. This evidence highlights the adaptability and efficiency of the PirateNet architecture in tackling complex, nonlinear problems.

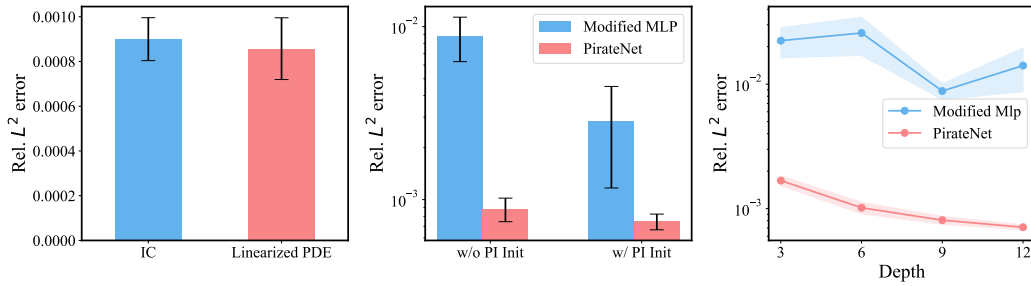


Figure 7: *Korteweg-De Vries equation*: *Left*: Relative L^2 errors of training PirateNet with the last layer initialized by the least square solution to fit the initial condition and the linearized PDE, respectively. *Middle*: Relative L^2 errors of training a Modified MLP and a PirateNet backbone with or without the physics-informed initialization. Without physics-informed initialization, the final layer defaults to a standard dense layer with weights initialized using the Xavier method and biases set to zero. *Right*: Relative L^2 test errors obtained by a Modified MLP and a PirateNet backbone of different depths. Each ablation study is performed under the same hyper-parameter settings, with results averaged over 5 random seeds.

5.3 Grey-Scott equation

In this example, we solve the 2D Grey-Scott equation, a reaction-diffusion system that describes the interaction of two chemical species. The form of this PDE is given as follows

$$\begin{aligned} u_t &= \epsilon_1 \Delta u + b_1(1 - u) - c_1 uv^2, & t \in (0, 2), (x, y) \in (-1, 1)^2, \\ v_t &= \epsilon_2 \Delta v - b_2 v + c_2 uv^2, & t \in (0, 2), (x, y) \in (-1, 1)^2, \end{aligned}$$

subject to the periodic boundary conditions and the initial conditions

$$\begin{aligned} u_0(x, y) &= 1 - \exp(-10((x + 0.05)^2 + (y + 0.02)^2)), \\ v_0(x, y) &= 1 - \exp(-10((x - 0.05)^2 + (y - 0.02)^2)). \end{aligned}$$

Here u and v represent the concentrations of the two species. The system can generate a wide range of patterns, including spots, stripes, and more complex forms, depending on the parameters chosen. For this example, we set $\epsilon_1 = 0.2$, $\epsilon_2 = 0.1$, $b_1 = 40$, $b_2 = 100$, $c_1 = c_2 = 1,000$, resulting in a beautiful random-seeming “rolls”.

We employ a standard time-marching strategy to train our PINN models, as outlined in various studies (Wight and Zhao, 2020; Krishnapriyan et al., 2021; Wang et al., 2023a). Specifically, we divide the temporal domain $[0, 1]$ into 10 equal intervals, employ the PirateNet architecture as a backbone and train a separate PINN model for each time window. The initial condition of each segment is given by the predicted solution at the end of its preceding interval. We initialize the last layer of each PirateNet by the least squares fitting of the associated initial conditions.

Figure 8 displays the predicted solutions of the PirateNet model at the final time $T = 1$. We can see a good alignment between our predictions and the corresponding numerical estimations. The resulting relative L^2 errors of u and v are 3.61×10^{-3} and 9.39×10^{-3} ,

respectively. Moreover, we compare the performance of PirateNet and Modified MLP backbones using different activation functions under the same hyper-parameter settings.

To further evaluate PirateNet’s performance, we conducted a comparative analysis with Modified MLP backbones, using different activation functions under identical hyper-parameter settings. Figure 9 visualizes the loss and error convergence during training for both PirateNets and Modified MLPs with various activation functions. The periodic peaks observed every 10^5 iterations implies transitions between time chunks in our time-marching approach, where a new PINN model is trained to learn the solution for each specific time interval. It can be observed that PirateNets consistently achieve stable loss convergence with both Tanh and swish activations, while the swish activation yields better accuracy for this example. In contrast, the Modified MLP exhibits significant instability in training loss, characterized by frequent spikes and ultimately failing to approximate the PDE solution. This underscores PirateNet’s robustness in terms of activation function choice. The remarkable training stability of PirateNets can be attributed to its initialization as a shallow network, which is inherently easier to train and more resilient to larger learning rates in the early stages of training.

To more thoroughly validate the efficacy of individual components in the PirateNet architecture, we carried out two ablation studies. The goal of the first study is to assess the impact of initializing the nonlinearity parameter $\alpha = 0$ of each residual block. For this purpose, we train PirateNet models of varying depths with initialized α at 0 or 1, respectively, to solve the PDE system within a time window $[0, 0.1]$. The results are summarized on the left panel of Figure 10, which clearly demonstrate that initializing $\alpha = 0$ significantly enhances accuracy by one or two orders of magnitude. Conversely, initializing $\alpha = 1$ results in an increased error, particularly when training deeper networks.

The second ablation study aims to examine the role of gating operations, as described in Equations (20) and (22). To this end, we fix a network depth of 9 and train the PirateNet with and without gating operations over 5 random seeds. As presented on the right panel of Figure 10, we can see that the incorporation of gating operations further reduces predictive errors. These findings strongly support the chosen initialization of $\alpha = 0$ and the design of the gating operation as key factors in the successful training of deep PirateNet models.

5.4 Ginzburg-Landau equation

In this example, we aim to solve a scalar PDE involving a complex variable. To this end, we consider the complex Ginzburg-Landau equation in 2D of the form

$$\frac{\partial A}{\partial t} = \epsilon \Delta A + \mu A - \gamma A |A|^2, \quad t \in (0, 1), \quad (x, y) \in (-1, 1)^2,$$

with an initial condition

$$A_0(x, y) = (10y + 10ix) \exp(-0.01(2500x^2 + 2500y^2)),$$

where A denotes a complex solution and we set $\epsilon = 0.004$, $\mu = 10$ and $\gamma = 10 + 15i$. It describes a vast variety of phenomena from nonlinear waves to second-order phase transitions, from superconductivity, superfluidity, and Bose-Einstein condensation, to liquid crystals and strings in field theory.




Figure 8: *Grey-Scott equation*: Comparisons between the solutions predicted by a trained PirateNet and the reference solutions. The detailed hyper-parameter settings are presented in Table 6.

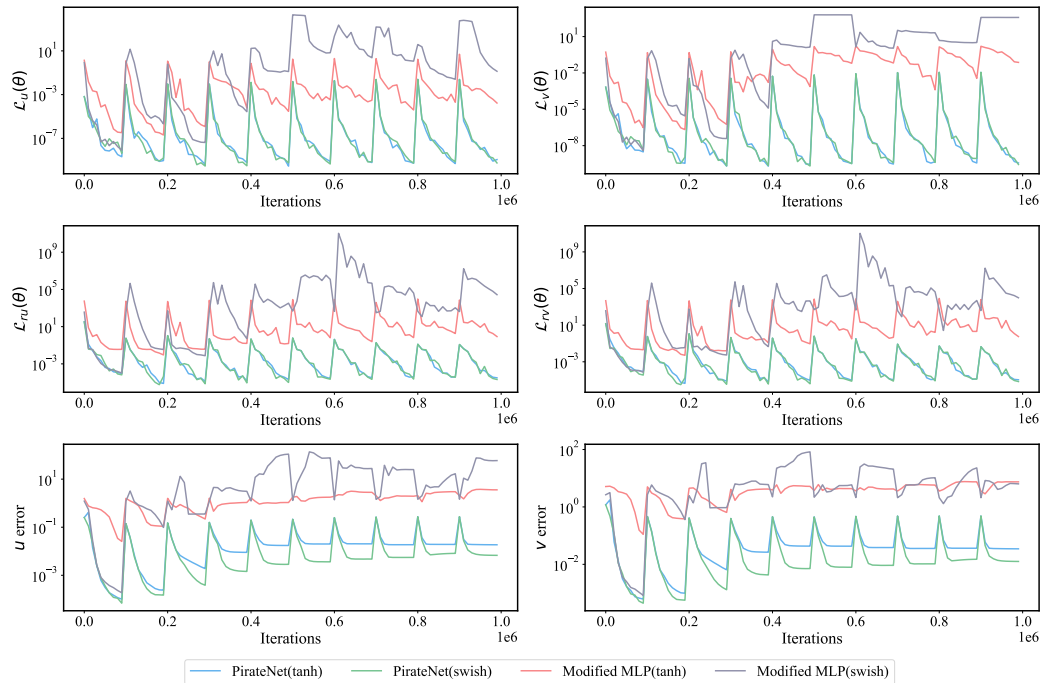


Figure 9: *Grey-Scott equation*: Convergence of the initial condition losses, the PDE residual losses and the relative L^2 errors during the training of PirateNet and Modified MLP backbones with different activation functions.




Figure 10: *Grey-Scott equation*: *Left*: Relative L^2 errors of training PirateNet of different depth, with the nonlinearity parameter α of each residual block initialized to 0 or 1. *Right*: Relative L^2 errors of training PirateNet with and without gating operations, respectively. Each ablation study is performed under the same hyper-parameter settings. All models are trained for a single time window, with the results averaged over five random seeds.

By denoting $A = u + iv$, we can decompose the equation into real and imaginary components, resulting in the following system of PDEs,

$$\begin{aligned}\frac{\partial u}{\partial t} &= \epsilon \Delta u + \mu(u - (u - 1.5v)(u^2 + v^2)), \\ \frac{\partial v}{\partial t} &= \epsilon \Delta v + \mu(v - (v + 1.5u)(u^2 + v^2)).\end{aligned}$$

In our experiments, we solve the decoupled equations using PINNs with a time-marching strategy, analogous to the approach detailed in Section 5.3. Specifically, we partition the time domain into five equal intervals, each of length 0.2, and train a PINN for each interval in sequence. The initial condition of each interval is given by the predicted solution at the end of the previous interval. For PirateNets, we initialize its final layer to align with these initial conditions via least squares.

Figure 11 displays the predicted real and imaginary parts of the solution obtained by using the PirateNet backbone. The predictions show a good agreement with the corresponding ground truth. The computed L^2 errors for the real part u and the imaginary part v are 1.49×10^{-2} and 1.90×10^{-2} , respectively. These results outperform the accuracy achieved by the Modified MLP backbone, which recorded errors of 3.20×10^{-2} for u and 1.94×10^{-2} for v . Figure 12 presents the corresponding training loss and test error trajectories, providing insights into the model’s convergence behavior.

Furthermore, we perform the same ablation studies as in Section 5.3 to further affirm the performance contributions of each element in the PirateNet architecture. As illustrated in the left panel of Figure 13, setting $\alpha = 1$ at initialization results in complete model collapse for deeper networks, while setting $\alpha = 0$ ensures stable training and consistently improved accuracy as network depth increases. In addition, the right panel of the figure supports our earlier finding that the use of gating operations enhances the model’s accuracy.

5.5 Lid-driven Cavity flow

In this example, we focus on a classical benchmark problem in computational fluid dynamics that involves simulating the movement of an incompressible fluid within a two-dimensional




Figure 11: *Ginzburg-Landau equation*: Comparisons between the solutions predicted by a trained PirateNet and the reference solutions. The detailed hyper-parameter settings are presented in Table 7.

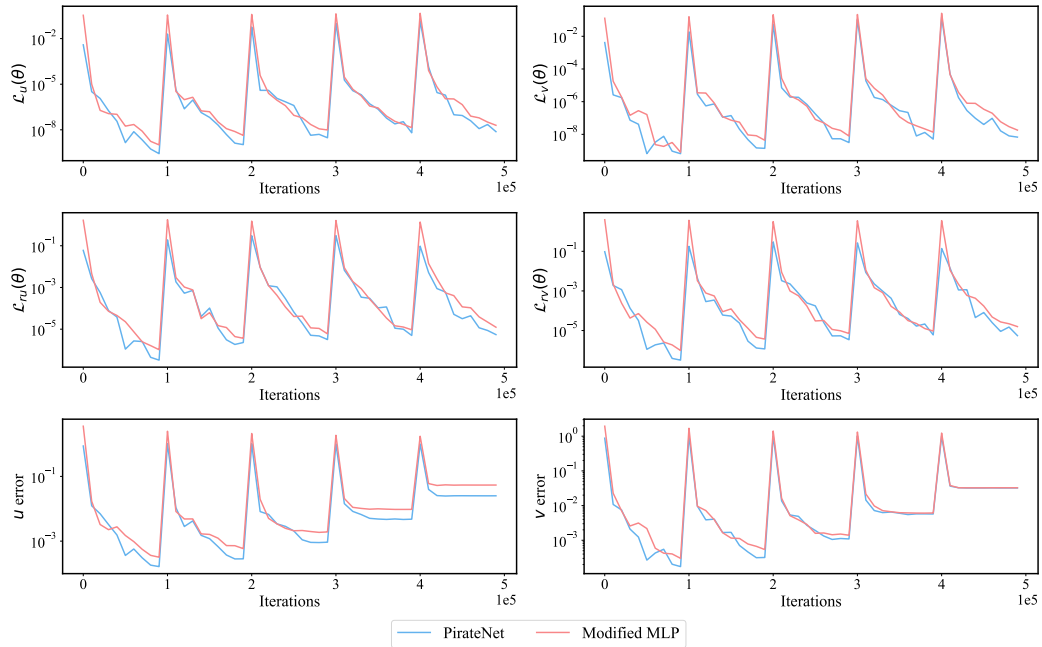


Figure 12: *Ginzburg-Landau equation*: Convergence of the initial condition losses, the PDE residual losses, and the relative L^2 errors during the training of a PirateNet and a Modified MLP backbone.




Figure 13: *Ginzburg-Landau equation*: *Left*: Relative L^2 errors of training PirateNet of different depth, with the nonlinearity parameter α of each residual block initialized to 0 or 1. *Right*: Relative L^2 errors of training PirateNet with and without gating operations, respectively. Each ablation study is performed under the same hyper-parameter setting. All models are trained within the first time window, with the results averaged over five random seeds.

square cavity. The system is governed by the incompressible Navier–Stokes equations written in a non-dimensional form

$$\begin{aligned} \mathbf{u} \cdot \nabla \mathbf{u} + \nabla p - \frac{1}{Re} \Delta \mathbf{u} &= 0, & (x, y) \in (0, 1)^2, \\ \nabla \cdot \mathbf{u} &= 0, & (x, y) \in (0, 1)^2, \end{aligned}$$

where $\mathbf{u} = (u, v)$ denotes the velocity in x and y directions, respectively, and p is the scalar pressure field. To avoid discontinuities of the top lid boundary conditions at two corners, we reformulate the boundary condition as follows

$$u(x, y) = 1 - \frac{\cosh(C_0(x - 0.5))}{\cosh(0.5C_0)}, \quad v(x, y) = 0, \quad (29)$$

where $x \in [0, 1], y = 1, C_0 = 50$. For the other three walls, we enforce a no-slip boundary condition. Our goal is to obtain the velocity and pressure field corresponding to a Reynolds number of 3200.

As illustrated by recent literature (Wang et al., 2023a,b; Cao and Zhang, 2023), training PINNs directly at high Reynolds numbers often leads to instability and a tendency to converge to erroneous solutions. To address this, we employ a curriculum training strategy (Krishnapriyan et al., 2021; Wang et al., 2023a), initially training PINNs at a lower Reynolds number, and gradually increasing the Reynolds numbers during training. This approach allows the model parameters optimized at lower Reynolds numbers to serve as a favorable initialization for training at higher Reynolds numbers. Specifically, we train PINNs across a progressively increasing sequence of Reynolds numbers: [100, 400, 1000, 1600, 3200]. For each Re , we conduct training the model for $10^4, 2 \times 10^4, 5 \times 10^4, 5 \times 10^4, \text{ and } 5 \times 10^5$ iterations, respectively.

Figure 14 plots the predicted velocity field for a Reynolds number of $Re = 3200$. The predictions align closely with the reference results from Ghia *et al.* (Ghia et al., 1982). This is evidenced by a relatively low L^2 error of 4.21×10^{-2} , a significant improvement over the




Figure 14: *Lid-driven cavity* ($Re=3200$): *Left*: Predicted velocity of the trained PirateNet. *Right*: Comparison of the predicted velocity profiles on the vertical and horizontal center-lines against Ghia *et al.* (Ghia *et al.*, 1982). The resulting relative L^2 error against the reference solution is 4.21×10^{-2} . The detailed hyper-parameter settings are presented in Table 8.




Figure 15: *Lid-driven cavity* ($Re=3200$): *Left*: Convergence of the relative L^2 errors during the training of PirateNets of different depths for the final training stage at $Re = 3200$. All experiments are performed under the same hyper-parameter setting. *Right*: Evolution of nonlinearities in each PirateNet residual block of training a PirateNet with 6 residual blocks (18 hidden layers).

1.58×10^{-1} error reported by JAXPI (Wang *et al.*, 2023a). Our ablation study on the model’s depth is visualized on the right panel of Figure 15, which reveals consistent performance enhancements as the depth increases within a moderately large range. However, the model seems to be saturated at depths greater than 20, yet the training remains stable and produces results similar to the one at a depth of 18. This observation highlights PirateNet’s impressive scalability and stability in handling complex, highly nonlinear problems and motivates future advancements in the model architecture. Interestingly, we also examine the evolution of the nonlinearities α in each residual block of the network. A noticeable increase in α is observed when the model transitions to training at higher Reynolds numbers. This trend suggests that the model requires a higher degree of nonlinearity in order to minimize the associated PDE residuals at higher Reynolds numbers.

6 Discussion

We introduce PirateNets, a new backbone for efficient training of deep PINN models. We provide theoretical arguments and illustrative numerical experiments that demonstrate the

instability of training PINNs when using deep MLPs. To explore the underlying reasons behind it, we propose a fundamental hypothesis that a good convergence in the PINNs training loss implies a good convergence of the PDE solution in C^k , depending on the order of the PDE system. Furthermore, under mild conditions, we rigorously prove that this hypothesis holds for linear elliptic and parabolic PDEs. This understanding allows us to shift our focus from analyzing the network’s PDE residuals to examining the MLP derivatives, both at initialization and during training. Our theoretical and empirical analysis reveals that standard initialization schemes result in pathological initialization of MLP derivatives, adversely affecting the trainability and robustness of PINNs.

Motivated by these insights, we introduce PirateNets; a novel architecture with adaptive residual connections. This design choice ensures that the network is initially a linear combination of a chosen basis, effectively addressing the deep MLP derivatives’ initialization issues. It is noteworthy to observe that PirateNet’s approximation capacity is gradually recovered during training, thanks to the trainable parameter α in each residual block. Additionally, the initial linearity of the model offers a unique advantage: it can be tailored to any kind of existing dataset by initializing the final network layer using solutions derived from corresponding least squares problems. Taking all together, our developments provide new insights into training PINNs and yield the state-of-art accuracy across a wide range of benchmarks in computational physics.

While this study marks a substantial advancement in understanding the training of deep physics-informed neural networks, we have to admit that training PINNs still requires significant computational resources and tends to be time-consuming. Future efforts should concentrate on accelerating this process. One promising direction could involve optimizing PirateNet’s coordinate embeddings and tailoring them to specific PDEs. Besides, the physics-informed initialization introduced here opens up new avenues for integrating physical principles into machine learning models. It would be interesting to extend this idea further to the field of operator learning. This may motivate the more efficient design of neural operators in solving parametric PDEs. We believe that pursuing these lines of exploration is crucial, as they are likely to form a key component in evolving physics-informed machine learning into a powerful and reliable tool for computational science and engineering.

Acknowledgments and Disclosure of Funding

We would like to acknowledge support from the US Department of Energy under the Advanced Scientific Computing Research program (grant DE-SC0024563). We also thank the developers of the software that enabled our research, including JAX (Bradbury et al., 2018), Matplotlib (Hunter, 2007), and NumPy (Harris et al., 2020).

Appendix A. Proofs

A.1 Proof of Proposition 1

Consider the second-order elliptic Dirichlet problem:

$$\mathcal{D}[u] = f,$$

with Dirichlet boundary condition and $f \in L^2(\Omega)$. Let $u : \Omega \rightarrow \mathbb{R}$ be its solution and u_θ be a smooth approximation by PINNs. Define the *expected loss* function:

$$\hat{\mathcal{L}}(\theta) = \int_{\partial\Omega} |u_\theta|^2 \, d\mathbf{x} + \int_{\Omega} |\mathcal{D}[u_\theta] - f|^2 \, d\mathbf{x}.$$

Then, for any compact $V \subset\subset \Omega$, there exist a constant C such that

$$\|u - u_\theta\|_{H^2(V)}^2 \leq C \hat{\mathcal{L}}(\theta).$$

Proof

Note that the PINN solution $u_\theta \in C^\infty(\Omega)$ if the activation function is smooth. We define the interior residual $\mathcal{R}_{\text{int},\theta} = \mathcal{D}[u_\theta] - f$ and the spatial boundary residual $\mathcal{R}_{\text{bc},\theta} = u_\theta$. With this notation, we see

$$\hat{\mathcal{L}}(\theta) = \|\mathcal{R}_{\text{bc},\theta}\|_{L^2(\partial\Omega)}^2 + \|\mathcal{R}_{\text{int},\theta}\|_{L^2(\Omega)}^2.$$

Let $u^* = u_{\theta^*}$ be a PINN solution generated by the algorithm. It is easy to see that the error $\hat{u} = u - u^*$ satisfies $\mathcal{D}[\hat{u}] = -\mathcal{R}_{\text{int},\theta^*}$ with $\hat{u} = \mathcal{R}_{\text{bc},\theta^*}$ on $\partial\Omega$. By interior regularity (Evans, 2022, Section 6, Theorem 2), we have for any $m \geq 0$, there holds $\hat{u} \in H_{\text{loc}}^{m+2}(\Omega)$ with for $V \subset\subset \Omega$,

$$\|\hat{u}\|_{H^{m+2}(V)} \leq C(\|\mathcal{R}_{\text{int},\theta^*}\|_{H^m(\Omega)} + \|\hat{u}\|_{L^2(\Omega)}).$$

Let \mathcal{M} be the Dirichlet map defined by $u := \mathcal{M}[g]$ if and only if $\mathcal{D}[u] = 0$ with $u = g$ on $\partial\Omega$. It is known that \mathcal{M} is a bounded linear map from $H^s(\partial\Omega)$ to $H^{s+1/2}(\Omega)$ for any real s (Bramble and Schatz, 1970). Then, since $\mathcal{R}_{\text{bc},\theta^*}$ is smooth, we can find a smooth solution such that $\mathcal{D}[v] = 0$ with $v = \mathcal{R}_{\text{bc},\theta^*}$ on $\partial\Omega$ and there holds

$$\|v\|_{L^2(\Omega)} \leq C\|\mathcal{R}_{\text{bc},\theta^*}\|_{L^2(\partial\Omega)},$$

which gives $\mathcal{D}[\hat{u} - v] = \mathcal{R}_{\text{int},\theta^*}$ with Dirichlet boundary. By the well-posedness, it follows that

$$\|\hat{u} - v\|_{L^2(\Omega)} \leq C\|\mathcal{R}_{\text{int},\theta^*}\|_{L^2(\Omega)}.$$

Hence, we have

$$\|\hat{u}\|_{H^2(V)} \leq C(\|\mathcal{R}_{\text{bc},\theta^*}\|_{L^2(\partial\Omega)} + \|\mathcal{R}_{\text{int},\theta^*}\|_{L^2(\Omega)}) \leq C\mathcal{L}(\theta)^{1/2}.$$

■

A.2 Proof of Proposition 2

Consider a second-order parabolic equation with Dirichlet boundary condition:

$$\begin{aligned} u_t + \mathcal{D}[u] &= f, & [0, T] \times \Omega, \\ u &= 0, & [0, T] \times \partial\Omega, \\ u &= g, & \{t = 0\} \times \Omega, \end{aligned}$$

with $f \in L^2([0, T]; L^2(\Omega))$ and $g \in L^2(\Omega)$. Let $u : [0, T] \times \Omega \rightarrow \mathbb{R}$ be its solution and u_θ be a smooth approximation by PINNs. Define the *expected loss*:

$$\hat{\mathcal{L}}(\theta) = \int_{\Omega} |\mathcal{R}_{\text{int}}[u_\theta](\mathbf{x})|^2 dx + \int_0^T \int_{\partial\Omega} |u_\theta(t, \mathbf{x})|^2 dt d\mathbf{x} + \int_0^T \int_{\Omega} |\mathcal{R}_{\text{int}}[u_\theta](t, \mathbf{x})|^2 dt d\mathbf{x},$$

with $\mathcal{R}_{\text{int}}[u_\theta]$ and $\mathcal{R}_{\text{bc}}[u_\theta]$ given in (4) and (6), respectively. Then, for any set $Q_r(t, \mathbf{x}) \subset \subset [0, T] \times \Omega$, there exist a constant C such that

$$\|u - u_\theta\|_{W_2^{1,2}(Q_r)}^2 \leq C \hat{\mathcal{L}}(\theta).$$

Proof Similarly, we denote the PINN solution by $u^* = u_{\theta^*}$. Then, $\hat{u} = u - u^*$ satisfies

$$\hat{u}_t + \mathcal{D}[\hat{u}] =: \mathcal{R}_{\text{int}, \theta^*},$$

with $\hat{u} =: \mathcal{R}_{\text{bc}, \theta^*}$ on $[0, T] \times \partial\Omega$ and $\hat{u} =: \mathcal{R}_{\text{ic}, \theta^*}$ on $\{t = 0\} \times \Omega$. For a set $Q_r(t, \mathbf{x}) \subset [0, T] \times \Omega$, the interior estimate of parabolic equation gives (Krylov, 1996, Theorem 2.4.7)

$$\|\hat{u}\|_{W_2^{1,k+2}(Q_r)} \leq C(\|\mathcal{R}_{\text{int}, \theta^*}\|_{L^2([0, T]; L^2(\Omega))} + \|\hat{u}\|_{L^2([0, T]; L^2(\Omega))}).$$

Let v be the solution to $v_t + \mathcal{D}[v] = 0$ with $v = \mathcal{R}_{\text{bc}, \theta^*}$ on $[0, T] \times \partial\Omega$ and $v = \mathcal{R}_{\text{ic}, \theta^*}$ on $\{t = 0\} \times \Omega$. The regularity estimate for v can be given by the standard semigroup method, which we recall below for the reader's convenience. We denote by $S(t)$ the strongly continuous analytic semigroup generated by $-\mathcal{D}$ and recall the Dirichlet map \mathcal{M} introduced above. Then, we have the following representation for the solution v (Lasiecka, 1986, Section 2)

$$v = S(t)[\mathcal{R}_{\text{ic}, \theta^*}] + A[\mathcal{R}_{\text{bc}, \theta^*}],$$

where

$$A[u] = \mathcal{D} \int_0^t S(t-z) \mathcal{M}[u](z) dz : L^2([0, T]; L^2(\partial\Omega)) \longrightarrow L^2([0, T]; L^2(\Omega)).$$

It follows that $v \in L^2([0, T]; L^2(\Omega))$ and

$$\|v\|_{L^2([0, T]; L^2(\Omega))} \leq C(\|\mathcal{R}_{\text{ic}, \theta^*}\|_{L^2(\Omega)} + \|\mathcal{R}_{\text{bc}, \theta^*}(t)\|_{L^2([0, T]; L^2(\partial\Omega))}).$$

Moreover, we have (Evans, 2022, Section 7, Theorem 5)

$$\|\hat{u} - v\|_{L^2([0, T]; H^2(\Omega))} \leq C \|\mathcal{R}_{\text{int}, \theta^*}\|_{L^2([0, T]; L^2(\Omega))}.$$

Therefore, we have

$$\|\hat{u}\|_{W_2^{1,k+2}(Q_r)} \leq C \left(\|\mathcal{R}_{\text{int}, \theta^*}\|_{L^2([0, T]; L^2(\Omega))} + \|\mathcal{R}_{\text{ic}, \theta^*}\|_{L^2(\Omega)} + \|\mathcal{R}_{\text{bc}, \theta^*}\|_{L^2([0, T]; L^2(\partial\Omega))} \right) \leq C \mathcal{L}(\theta)^{1/2}.$$

■

A.3 Proof of Proposition 4

The following holds at initialization of an MLP $u_\theta(x)$.

- (a) The first-order derivative of u_θ is given by

$$\frac{\partial u_\theta}{\partial x}(x) = \prod_{i=L+1}^2 \left(\mathbf{W}^{(i)} \cdot \text{diag} \left(\dot{\sigma}(\mathbf{u}_\theta^{(i-1)}(x)) \right) \right) \cdot \mathbf{W}^{(1)}.$$

Furthermore, assuming that the MLP lies in a linear regime, i.e., $\sigma(x) \approx x$ and $\dot{\sigma}(\cdot) \approx 1$, it follows that

$$\frac{\partial u_\theta}{\partial x}(x) \approx \mathbf{W}^{(L+1)} \cdot \mathbf{W}^{(L)} \dots \mathbf{W}^{(1)}.$$

- (b) If all hidden layers of the MLP have the same number of neurons, i.e., $d_1 = d_2 = \dots = d_L = d$, then

$$\text{Var} \left(\frac{\partial u_\theta}{\partial x}(x) \right) \lesssim \frac{1}{d}.$$

Consequently, for any $\epsilon > 0$,

$$\mathbb{P} \left(\left| \frac{\partial u_\theta}{\partial x}(x) \right| > \epsilon \right) \lesssim \frac{1}{d\epsilon^2}.$$

Proof (a) The first statement is a simple consequence of the chain rule. If we further assume the linear regime that $\sigma(\cdot) \approx 1$, then $\text{diag}(\dot{\sigma}(\mathbf{u}_\theta^{(k)}(x))) \approx I$ for all k . Then we obtain

$$\frac{\partial u_\theta}{\partial x}(x) \approx \mathbf{W}^{(L+1)} \cdot \mathbf{W}^{(L)} \dots \mathbf{W}^{(1)}. \quad (30)$$

(b) We prove this result by induction on the number of network layers L . For $L = 1$, the network derivative is given by

$$\frac{\partial u_\theta}{\partial x}(x) = \sum_{k=1}^d W_k^{(2)} \dot{\sigma} \left(W_k^{(1)} x + b^{(1)} \right) W_k^{(1)}. \quad (31)$$

Since $W_k^{(l)}$ is independent from each other for all $l = 1, 2$ and $k = 1, 2, \dots, d$, we have

$$\begin{aligned} \text{Var} \left(\frac{\partial u_\theta}{\partial x}(x) \right) &= \sum_{k=1}^d \text{Var} \left(W_k^{(2)} \dot{\sigma} \left(W_k^{(1)} x + b^{(1)} \right) W_k^{(1)} \right) \\ &= \sum_{k=1}^d \left[\mathbb{E} \left(W_k^{(2)} \dot{\sigma} \left(W_k^{(1)} x + b^{(1)} \right) W_k^{(1)} \right)^2 - \left(\mathbb{E} \left(W_k^{(2)} \dot{\sigma} \left(W_k^{(1)} x + b^{(1)} \right) W_k^{(1)} \right) \right)^2 \right] \\ &\leq \sum_{k=1}^d \left[\mathbb{E} \left(W_k^{(2)} W_k^{(1)} \right)^2 - \left(\mathbb{E} \left(W_k^{(2)} \right) \cdot \mathbb{E} \left(\dot{\sigma} \left(W_k^{(1)} x + b^{(1)} \right) W_k^{(1)} \right) \right)^2 \right] \\ &= \sum_{k=1}^d \text{Var} \left(W_k^{(2)} \right) \text{Var} \left(W_k^{(1)} \right) \\ &\lesssim d \cdot \frac{1}{d} \cdot \frac{1}{d} = \frac{1}{d}. \end{aligned} \quad (32)$$

Now, assume that this holds for a scalar network of L layers. Then, by (a), for $L + 1$, we have

$$\frac{\partial u_\theta}{\partial x}(x) = \prod_{i=L+1}^2 \left(\mathbf{W}^{(i)} \cdot \text{diag} \left(\dot{\sigma} \left(\mathbf{u}_\theta^{(i-1)}(x) \right) \right) \right) \cdot \mathbf{W}^{(1)}. \quad (33)$$

Let $X = \prod_{i=L}^2 \left(\mathbf{W}^{(i)} \cdot \text{diag} \left(\dot{\sigma} \left(\mathbf{u}_\theta^{(i-1)}(x) \right) \right) \right) \cdot \mathbf{W}^{(1)}$. The key observation is that each element X_k is a scalar network of L layers. By induction, we have $\text{Var}(X_k) \lesssim \frac{1}{d}$ for $k = 1, 2, \dots, d$. Therefore,

$$\frac{\partial u_\theta}{\partial x}(x) = \sum_{k=1}^d W_k^{(L+1)} \dot{\sigma} \left(\mathbf{u}_\theta^{(L)}(x) \right) X_k. \quad (34)$$

Similar to (32), we can prove that

$$\frac{\partial u_\theta}{\partial x}(x) \lesssim \frac{1}{d}. \quad (35)$$

Finally, applying Chebyshev's inequality, for any $\epsilon > 0$, we obtain

$$\text{Var} \left(\frac{\partial u_\theta}{\partial x}(x) \right) \lesssim \frac{1}{d\epsilon^2}. \quad (36)$$

■

Appendix B. Allen-Cahn equation

Data generation. We solve the Allen-Cahn equation using conventional spectral methods. Specifically, assuming periodic boundary conditions, we start from the initial condition $u_0(x) = x^2 \cos(\pi x)$ and integrate the system up to the final time $T = 1$. Synthetic validation data is generated using the Chebfun package (Driscoll et al., 2014) with a spectral Fourier discretization with 512 modes and a fourth-order stiff time-stepping scheme (ETDRK4) (Cox and Matthews, 2002) with the time-step size of 10^{-5} . We record the solution at intervals of $\Delta t = 0.005$, yielding a validation dataset with a resolution of 200×512 .

Table 4: *Allen-Cahn equation*: Hyper-parameter configuration for reproducing results in Figure 4.

Parameter	Value
Architecture	
Number of layers	9
Number of channels	256
Activation	Tanh
Fourier feature scale	2.0
Random weight factorization	$\mu = 1.0, \sigma = 0.1$
Learning rate schedule	
Initial learning rate	10^{-3}
Decay rate	0.9
Decay steps	5×10^3
Warmup steps	5×10^3
Training	
Training steps	3×10^5
Batch size	8,192
Weighting	
Weighting scheme	NTK (Wang et al., 2022c, 2023a)
Causal tolerance	1.0
Number of chunks	32

Appendix C. Korteweg–De Vries equation

Data generation. We solve the Korteweg–De Vries equation using conventional spectral methods. Specifically, assuming periodic boundary conditions, we start from the initial condition $u_0(x) = \cos(\pi x)$ and integrate the system up to the final time $T = 1$. Synthetic validation data is generated using the Chebfun package (Driscoll et al., 2014) with a spectral Fourier discretization with 512 modes and a fourth-order stiff time-stepping scheme (ETDRK4) (Cox and Matthews, 2002) with the time-step size of 10^{-5} . We record the solution at intervals of $\Delta t = 0.005$, yielding a validation dataset with a resolution of 200×512 in temporal spatial domain..

Table 5: *Korteweg–De Vries equation*: Hyper-parameter configuration for reproducing results in Figure 6.

Parameter	Value
Architecture	
Number of layers	9
Number of channels	256
Activation	Tanh
Fourier feature scale	1.0
Random weight factorization	$\mu = 1.0, \sigma = 0.1$
Learning rate schedule	
Initial learning rate	10^{-3}
Decay rate	0.9
Decay steps	2×10^3
Warmup steps	5×10^3
Training	
Training steps	2×10^5
Batch size	4,096
Weighting	
Weighting scheme	Grad norm (Wang et al., 2021a)
Causal tolerance	1.0
Number of chunks	16

Appendix D. Grey-Scott equation

Data generation. We solve the Grey-Scott equation using conventional spectral methods. Specifically, assuming periodic boundary conditions, we start from the initial condition and integrate the system up to the final time $T = 2$. Synthetic validation data is generated using the Chebfun package (Driscoll et al., 2014) with a spectral Fourier discretization with 200×200 modes in 2D and a fourth-order stiff time-stepping scheme (ETDRK4) (Cox and Matthews, 2002) with the time-step size of 10^{-3} . We record the solution at intervals of $\Delta t = 0.02$, yielding a validation dataset with a resolution of $100 \times 200 \times 200$ in temporal spatial domain.

Table 6: *Grey-Scott equation*: Hyper-parameter configuration for reproducing results in Figure 8.

Parameter	Value
Architecture	
Number of layers	9
Number of channels	256
Activation	Swish
Fourier feature scale	1.0
Random weight factorization	$\mu = 0.5, \sigma = 0.1$
Learning rate schedule	
Initial learning rate	10^{-3}
Decay rate	0.9
Decay steps	2×10^3
Warmup steps	5×10^3
Training	
Number of time windows	10
Training steps per time window	10^5
Batch size	4,096
Weighting	
Scheme	Grad norm (Wang et al., 2021a)
Causal tolerance	1.0
Number of chunks	32

Appendix E. Ginzburg-Laudau equation

Data generation. We solve the Ginzburg-Laudau equation using conventional spectral methods. Specifically, assuming periodic boundary conditions, we start from the initial condition and integrate the system up to the final time $T = 1$. Synthetic validation data is generated using the Chebfun package (Driscoll et al., 2014) with a spectral Fourier discretization with 200×200 modes in 2D and a fourth-order stiff time-stepping scheme (ETDRK4) (Cox and Matthews, 2002) with time-step size of 10^{-3} . We record the solution at intervals of $\Delta t = 0.01$, yielding a validation dataset with a resolution of $100 \times 200 \times 200$ in temporal spatial domain..

Table 7: *Ginzburg-Laudau equation*: Hyper-parameter configuration for reproducing results in Figure 11.

Parameter	Value
Architecture	
Number of layers	9
Number of channels	256
Activation	Swish
Fourier feature scale	1.0
Random weight factorization	$\mu = 0.5, \sigma = 0.1$
Learning rate schedule	
Initial learning rate	10^{-3}
Decay rate	0.9
Decay steps	2×10^3
Warmup steps	5×10^3
Training	
Number of time windows	5
Training steps per time window	10^5
Batch size	8,192
Weighting	
Scheme	Grad norm (Wang et al., 2021a)
Causal tolerance	5.0
Number of chunks	16

Appendix F. Lid-driven Cavity flow

Data generation. We compare our results against (Ghia et al., 1982).

Table 8: *Lid-driven Cavity flow*: Hyper-parameter configuration for reproducing results in Figure 14.

Parameter	Value
Architecture	
Number of layers	18
Number of channels	256
Activation	Tanh
Fourier feature scale	15.0
Random weight factorization	$\mu = 1.0, \sigma = 0.1$
Learning rate schedule	
Initial learning rate	10^{-3}
Decay rate	0.9
Decay steps	10^4
Warmup steps	5×10^3
Training	
Re	[100, 400, 1,000, 1,600, 3,200]
Training steps	$[10^4, 2 \times 10^4, 5 \times 10^4, 5 \times 10^4, 5 \times 10^5]$
Batch size	4,096
Weighting	
Scheme	Grad norm (Wang et al., 2021a)

Appendix G. Additional results on PINNacle (Hao et al., 2023)

This section presents additional results of PirateNet on benchmark problems from PINNacle (Hao et al., 2023). We focus on three representative benchmarks: Burgers' 1D-C, Wave 1D-C, and the inverse problem of Poisson equations. Our experiments strictly adhere to the problem settings outlined in the PINNacle paper, which are detailed below for each benchmark.

1D Burgers Equation (Burgers 1D-C). The 1D Burgers equation is defined as:

$$u_t + uu_x = \nu u_{xx}, \quad (37)$$

where $(x, t) \in \Omega = [-1, 1] \times [0, 1]$, with initial and boundary conditions:

$$u(x, 0) = -\sin(\pi x), \quad (38)$$

$$u(-1, t) = u(1, t) = 0, \quad (39)$$

and viscosity parameter $\nu = 0.01/\pi$.

1D Wave Equation (Wave1d-C). The governing PDE for the 1D wave equation is:

$$u_{tt} - 4u_{xx} = 0, \quad (40)$$

on the domain $\Omega \times T = [0, 1] \times [0, 1]$, with boundary conditions:

$$u(0, t) = u(1, t) = 0, \quad (41)$$

and initial conditions:

$$u(x, 0) = \sin(\pi x) + \frac{1}{2} \sin(4\pi x), \quad (42)$$

$$u_t(x, 0) = 0. \quad (43)$$

The analytical solution is given by:

$$u(x, t) = \sin(\pi x) \cos(2\pi t) + \frac{1}{2} \sin(4\pi x) \cos(8\pi t). \quad (44)$$

2D Navier-Stokes equation (NS 2D-C). Consider the incompressible Navier-Stokes equations in two dimensions:

$$\mathbf{u} \cdot \nabla \mathbf{u} + \nabla_p - \frac{1}{Re} \Delta \mathbf{u} = 0, \quad x \in \Omega, \quad (45)$$

$$\nabla \cdot \mathbf{u} = 0, \quad x \in \Omega, \quad (46)$$

where \mathbf{u} represents the velocity field, p denotes pressure, and $Re = 100$ is the Reynolds number. The equations are solved on the unit square domain $\Omega = [0, 1]^2$. Let Γ_1 denote the top boundary and Γ_2 represent the remaining boundaries (left, right, and bottom). The boundary conditions are:

$$\mathbf{u}(\mathbf{x}) = (4x(1-x), 0), \quad x \in \Gamma_1, \quad (47)$$

$$\mathbf{u}(\mathbf{x}) = (0, 0), \quad x \in \Gamma_2, \quad (48)$$

$$p = 0, \quad x = (0, 0). \quad (49)$$

2D Navier-Stokes equation (NS 2D-CG). The governing equations remain the same as above:

$$\mathbf{u} \cdot \nabla \mathbf{u} + \nabla p - \frac{1}{Re} \Delta \mathbf{u} = 0, \quad (50)$$

$$\nabla \cdot \mathbf{u} = 0. \quad (51)$$

However, the domain Ω is now defined as $\Omega = [0, 4] \times [0, 2] \setminus ([0, 2] \times [1, 2] \cup R_i)$, representing a rectangular region with its top-left quarter removed. The inlet velocity is given by $u_{\text{in}} = 4y(1 - y)$, the outlet pressure is $p = 0$, and the boundary condition is no-slip: $\mathbf{u} = 0$. The Reynolds number is set to $Re = 100$.

Poisson inverse problem. The governing PDE is

$$-\nabla(a\nabla u) = f. \quad (52)$$

The geometric domain is $\Omega = [0, 1]^2$, and

$$u = \sin(\pi x) \sin(\pi y), \quad (53)$$

in the domain $\Omega = [0, 1]^2$, where $u = \sin(\pi x) \sin(\pi y)$ and the source term f is:

$$f = \frac{2\pi^2 \sin(\pi x) \sin(\pi y)}{1 + x^2 + y^2 + (x - 1)^2 + (y - 1)^2} + \frac{2\pi [(2x - 1) \cos(\pi x) \sin(\pi y) + (2y - 1) \sin(\pi x) \cos(\pi y)]}{[1 + x^2 + y^2 + (x - 1)^2 + (y - 1)^2]^2}. \quad (54)$$

To ensure uniqueness of the solution, we impose a boundary condition of $a(x, y)$, i.e.,

$$a(x, y) = \frac{1}{1 + x^2 + y^2 + (x - 1)^2 + (y - 1)^2}, \quad x \in \partial\Omega. \quad (55)$$

We sample data of $u(x, y)$ with 2,500 uniformly distributed 50×50 points and add Gaussian noise $\mathcal{N}(0, 0.1)$ to it. The objective is to reconstruct the diffusion coefficient $a(x, y)$, whose ground truth is:

$$a(x, y) = \frac{1}{1 + x^2 + y^2 + (x - 1)^2 + (y - 1)^2}, \quad x \in \Omega. \quad (56)$$

The performance of PirateNet and several baseline methods on these benchmarks is summarized in Table 9. A notable observation is that our baseline implementation (JAX-PI) demonstrates remarkably strong performance, achieving accuracy improvements of about two orders of magnitude for forward problems compared to the results reported in PINNacle. More significantly, PirateNet consistently outperforms all baselines across the three problems, showcasing its superior effectiveness and versatility in handling various types of PDEs. To provide a more comprehensive view of PirateNet’s performance, we have included visualizations of the predicted solutions, along with loss and test error convergence plots during training. These can be found in Figure 16 - 20. These visualizations offer insights into the predicted solution quality and convergence behavior of PirateNet across different problem types, further illustrating its robust performance.

Benchmark	PINNacle	PIELM	JAX-PI	PirateNet
Burger 1D-C	1.33×10^{-2}	1.63×10^{-1}	3.58×10^{-4}	8.20×10^{-5}
Wave 1D-C	9.79×10^{-2}	8.58×10^{-5}	9.80×10^{-5}	5.15×10^{-5}
NS 2D-C	3.60×10^{-2}	N/A	2.39×10^{-3}	9.52×10^{-4}
NS 2D-CG	8.24×10^{-2}	N/A	4.10×10^{-2}	2.10×10^{-2}
Poisson Inverse	1.96×10^{-2}	N/A	1.99×10^{-2}	5.34×10^{-3}

Table 9: Performance of PirateNet on PINNacle benchmarks. For PINNacle results, We directly report the results from its original paper. The metric is relative L^2 error between the predicted and reference solutions.

In particular, we also compare PirateNet against PIELM (Physics-Informed Extreme Learning Machine). Our results indicate that PIELM performs exceptionally well for linear PDEs, because of its approach of solving an underlying linear optimization problem with provable convergence. However, for non-linear problems, PIELM exhibits significant sensitivity to optimizer choice, network width, and selection of collocation points. In these cases, it tends to converge easily to poor local minima. Moreover, PIELM is not applicable to inverse problems, further limiting its versatility compared to PirateNet.

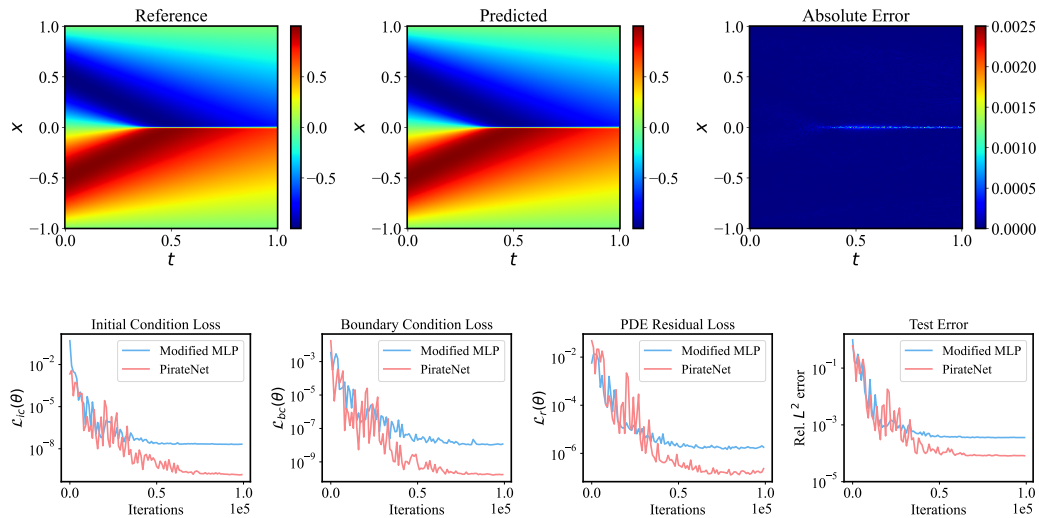


Figure 16: *Burgers equation*: *Top*: Comparison between the solution predicted by a trained PirateNet and the reference solution. *Bottom*: Convergence of the initial condition loss, the PDE residual loss, and the relative L^2 error during the training of a PirateNet and a Modified MLP backbone.

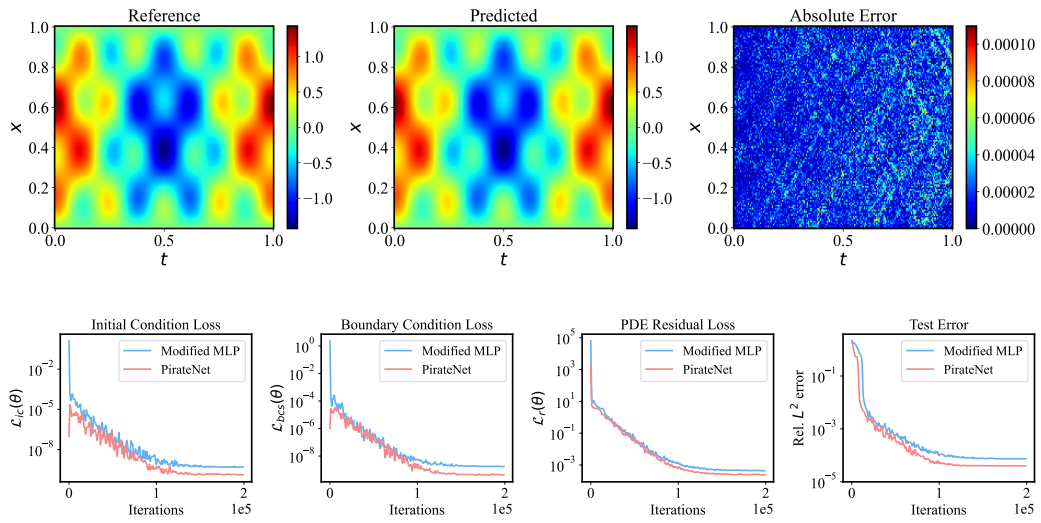


Figure 17: *Wave equation*: *Top*: Comparison between the solution predicted by a trained PirateNet and the reference solution. *Bottom*: Convergence of the initial condition loss, the PDE residual loss, and the relative L^2 error during the training of a PirateNet and a Modified MLP backbone.

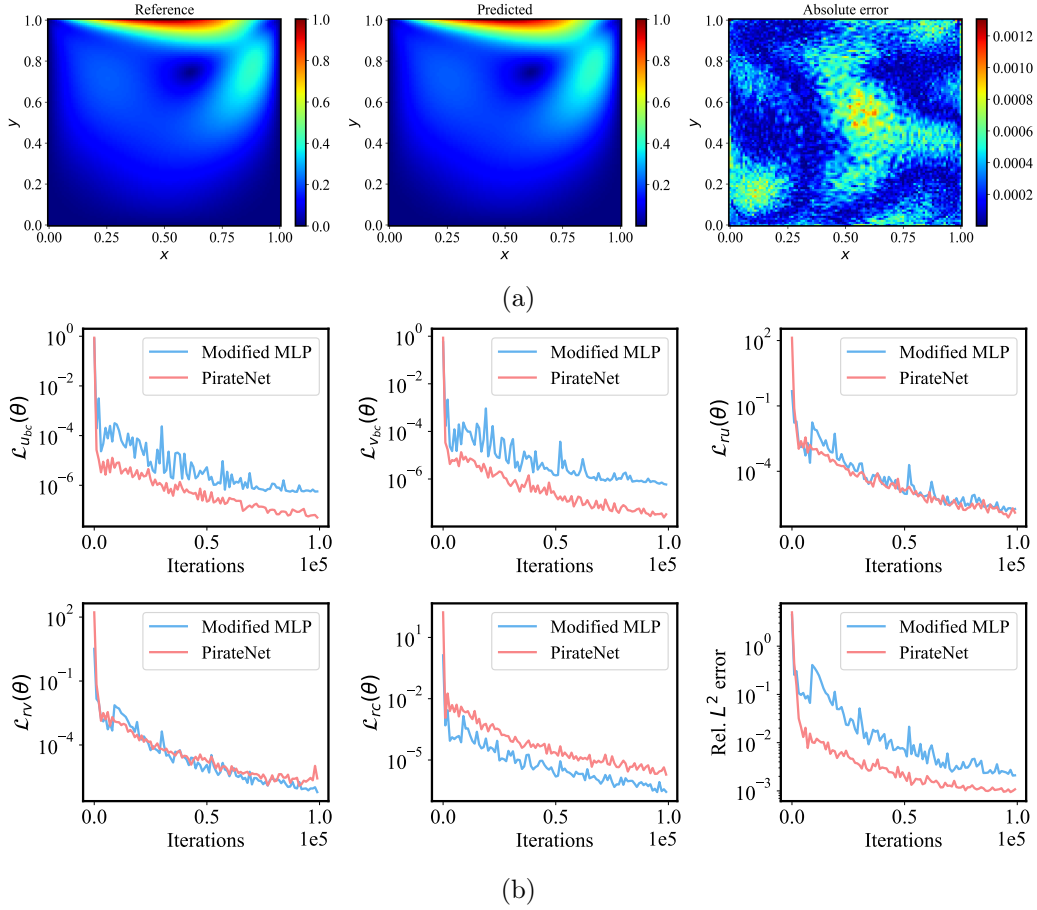


Figure 18: *NS 2D-C*: (a): Comparison between the solution predicted by a trained PirateNet and the reference solution. (b): Convergence of the boundary condition losses, the PDE residual losses, and the relative L^2 error during the training of a PirateNet and a Modified MLP backbone.

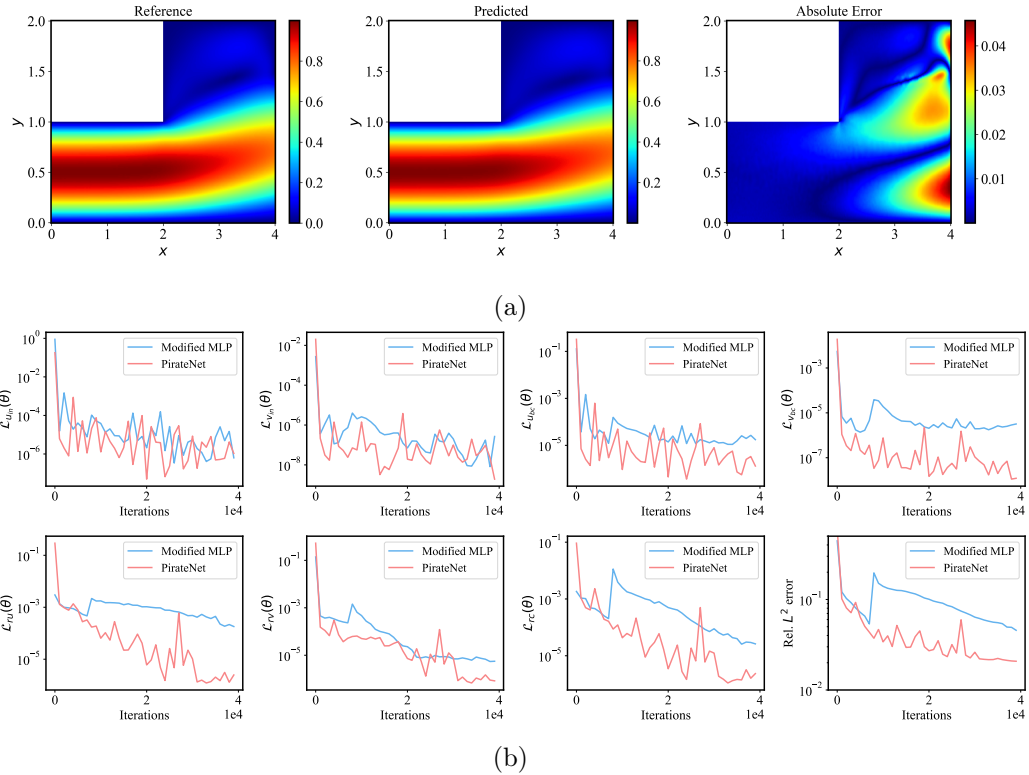


Figure 19: *NS 2D-CG*: (a): Comparison between the solution predicted by a trained PirateNet and the reference solution. (b): Convergence of the boundary condition losses, the PDE residual losses, and the relative L^2 error during the training of a PirateNet and a Modified MLP backbone.

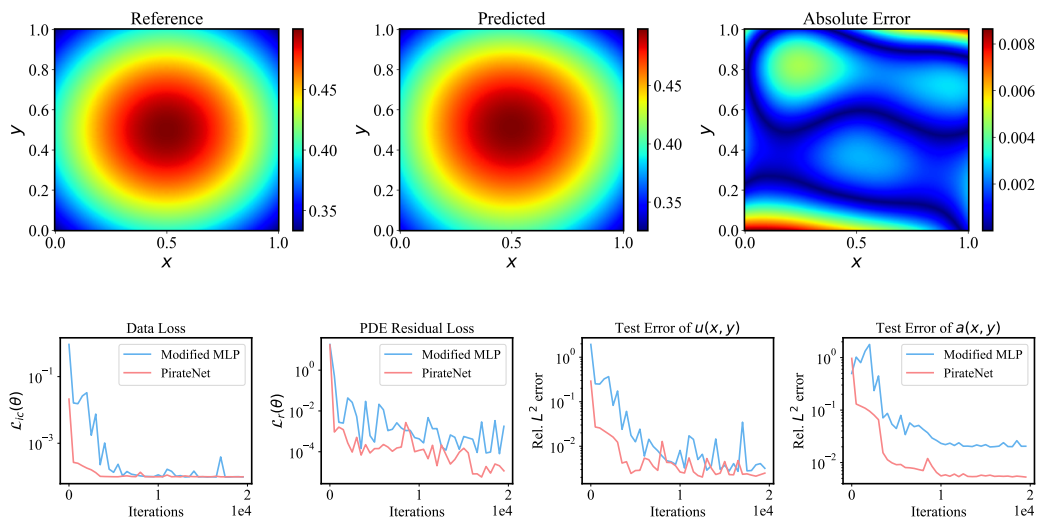


Figure 20: *Poisson inverse*: *Top*: Comparison between the coefficients predicted by a trained PirateNet and the reference solution. *Bottom*: Convergence of the initial condition loss, the PDE residual loss, and the relative L^2 error during the training of a PirateNet and a Modified MLP backbone.

References

- Sokratis J Anagnostopoulos, Juan Diego Toscano, Nikolaos Stergiopoulos, and George Em Karniadakis. Residual-based attention and connection to information bottleneck theory in pinns. *arXiv preprint arXiv:2307.00379*, 2023.
- James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/google/jax>.
- James H Bramble and Alfred H Schatz. Rayleigh-ritz-galerkin methods for dirichlet’s problem using subspaces without boundary conditions. *Communications on Pure and Applied Mathematics*, 23(4):653–675, 1970.
- Shengze Cai, Zhicheng Wang, Sifan Wang, Paris Perdikaris, and George Em Karniadakis. Physics-informed neural networks for heat transfer problems. *Journal of Heat Transfer*, 143(6), 2021.
- Wenbo Cao and Weiwei Zhang. Tsnn: Time-stepping-oriented neural network for solving partial differential equations. *arXiv preprint arXiv:2310.16491*, 2023.
- Souvik Chakraborty. Transfer learning based multi-fidelity physics informed deep neural network. *Journal of Computational Physics*, 426:109942, 2021.
- Yuyao Chen, Lu Lu, George Em Karniadakis, and Luca Dal Negro. Physics-informed neural networks for inverse problems in nano-optics and metamaterials. *Optics express*, 28(8): 11618–11633, 2020.
- Pao-Hsiung Chiu, Jian Cheng Wong, Chinchun Ooi, My Ha Dao, and Yew-Soon Ong. Canpinn: A fast physics-informed neural network based on coupled-automatic–numerical differentiation method. *Computer Methods in Applied Mechanics and Engineering*, 395: 114909, 2022.
- Taco Cohen, Maurice Weiler, Berkay Kicanaoglu, and Max Welling. Gauge equivariant convolutional networks and the icosahedral cnn. In *International conference on Machine learning*, pages 1321–1330. PMLR, 2019.
- Steven M Cox and Paul C Matthews. Exponential time differencing for stiff systems. *Journal of Computational Physics*, 176(2):430–455, 2002.
- Jérôme Darbon, Gabriel P Langlois, and Tingwei Meng. Overcoming the curse of dimensionality for some hamilton–jacobi partial differential equations via neural network architectures. *Research in the Mathematical Sciences*, 7:1–50, 2020.
- Arka Daw, Jie Bu, Sifan Wang, Paris Perdikaris, and Anuj Karpatne. Rethinking the importance of sampling in physics-informed neural networks. *arXiv preprint arXiv:2207.02338*, 2022.

- Shaan Desai, Marios Mattheakis, Hayden Joy, Pavlos Protopapas, and Stephen Roberts. One-shot transfer learning of physics-informed neural networks. *arXiv preprint arXiv:2110.11286*, 2021.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Suchuan Dong and Zongwei Li. Local extreme learning machines and domain decomposition for solving linear and nonlinear partial differential equations. *Computer Methods in Applied Mechanics and Engineering*, 387:114129, 2021.
- Suchuan Dong and Naxian Ni. A method for representing periodic functions and enforcing exactly periodic boundary conditions with deep neural networks. *Journal of Computational Physics*, 435:110242, 2021.
- Tobin A Driscoll, Nicholas Hale, and Lloyd N Trefethen. Chebfun guide, 2014.
- Vikas Dwivedi and Balaji Srinivasan. Physics informed extreme learning machine (pielm)—a rapid method for the numerical solution of partial differential equations. *Neurocomputing*, 391:96–118, 2020.
- Vasiliy A Es' kin, Danil V Davydov, Ekaterina D Egorova, Alexey O Malkhanov, Mikhail A Akhukov, and Mikhail E Smorkalov. About optimal loss function for training physics-informed neural networks under respecting causality. *arXiv preprint arXiv:2304.02282*, 2023.
- Lawrence C Evans. *Partial differential equations*, volume 19. American Mathematical Society, 2022.
- Zhiwei Fang. A high-efficient hybrid physics-informed neural networks based on convolutional neural network. *IEEE Transactions on Neural Networks and Learning Systems*, 33(10): 5514–5526, 2021.
- Zhiwei Fang and Justin Zhan. Deep physical informed neural networks for metamaterial design. *IEEE Access*, 8:24506–24513, 2019.
- Rizal Fathony, Anit Kumar Sahu, Devin Willmott, and J Zico Kolter. Multiplicative filter networks. In *International Conference on Learning Representations*, 2021.
- 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.
- U Ghia, K.N Ghia, and C.T Shin. High-re solutions for incompressible flow using the navier-stokes equations and a multigrid method. *Journal of Computational Physics*, 48(3): 387–411, 1982. ISSN 0021-9991. doi: [https://doi.org/10.1016/0021-9991\(82\)90058-4](https://doi.org/10.1016/0021-9991(82)90058-4).
- Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256, 2010.

- Somdatta Goswami, Cosmin Anitescu, Souvik Chakraborty, and Timon Rabczuk. Transfer learning enhanced physics informed neural network for phase-field modeling of fracture. *Theoretical and Applied Fracture Mechanics*, 106:102447, 2020.
- Andreas Griewank and Andrea Walther. *Evaluating derivatives: principles and techniques of algorithmic differentiation*. SIAM, 2008.
- Ehsan Haghighat and Ruben Juanes. Sciann: A keras/tensorflow wrapper for scientific computations and physics-informed deep learning using artificial neural networks. *Computer Methods in Applied Mechanics and Engineering*, 373:113552, 2021.
- Zhongkai Hao, 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. *arXiv preprint arXiv:2306.08827*, 2023.
- Charles R Harris, K Jarrod Millman, Stéfan J van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J Smith, et al. Array programming with numpy. *Nature*, 585(7825):357–362, 2020.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016a.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14*, pages 630–645. Springer, 2016b.
- Oliver Hennigh, Susheela Narasimhan, Mohammad Amin Nabian, Akshay Subramaniam, Kaustubh Tangsali, Max Rietmann, Jose del Aguila Ferrandis, Wonmin Byeon, Zhiwei Fang, and Sanjay Choudhry. Nvidia simnetTM: an ai-accelerated multi-physics simulation framework. *arXiv preprint arXiv:2012.07938*, 2020.
- Lars Hörmander. On interior regularity of the solutions of partial differential equations. *Communications on Pure and Applied Mathematics*, 11(2):197–218, 1958.
- Xinquan Huang and Tariq Alkhalifah. Efficient physics-informed neural networks using hash encoding. *Journal of Computational Physics*, page 112760, 2024.
- John D Hunter. Matplotlib: A 2D graphics environment. *IEEE Annals of the History of Computing*, 9(03):90–95, 2007.
- Mahmudul Islam, Md Shajedul Hoque Thakur, Satyajit Mojumder, and Mohammad Nasim Hasan. Extraction of material properties through multi-fidelity deep learning from molecular dynamics simulation. *Computational Materials Science*, 188:110187, 2021.
- 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.

- John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, 2021.
- Namgyu Kang, Byeonghyeon Lee, Youngjoon Hong, Seok-Bae Yun, and Eunbyung Park. Pixel: Physics-informed cell representations for fast and accurate pde solvers. *arXiv preprint arXiv:2207.12800*, 2022.
- 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.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Georgios Kissas, Yibo Yang, Eileen Hwuang, Walter R Witschey, John A Detre, and Paris Perdikaris. Machine learning in cardiovascular flows modeling: Predicting arterial blood pressure from non-invasive 4D flow MRI data using physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering*, 358:112623, 2020.
- Alexander Kovacs, Lukas Exl, Alexander Kornell, Johann Fischbacher, Markus Hovorka, Markus Gusenbauer, Leoni Breth, Harald Oezelt, Masao Yano, Noritsugu Sakuma, et al. Conditional physics informed neural networks. *Communications in Nonlinear Science and Numerical Simulation*, 104:106041, 2022.
- Aditi S Krishnapriyan, Amir Gholami, Shandian Zhe, Robert M Kirby, and Michael W Mahoney. Characterizing possible failure modes in physics-informed neural networks. *arXiv preprint arXiv:2109.01050*, 2021.
- Nikolai Vladimirovich Krylov. *Lectures on elliptic and parabolic equations in Holder spaces*. Number 12 in Graduate studies in mathematics. American Mathematical Soc., 1996.
- Remi Lam, Alvaro Sanchez-Gonzalez, Matthew Willson, Peter Wirnsberger, Meire Fortunato, Alexander Pritzel, Suman Ravuri, Timo Ewalds, Ferran Alet, Zach Eaton-Rosen, et al. Graphcast: Learning skillful medium-range global weather forecasting. *arXiv preprint arXiv:2212.12794*, 2022.
- I Lasićka. Galerkin approximations of abstract parabolic boundary value problems with rough boundary data— l_p theory. *mathematics of computation*, 47(175):55–75, 1986.
- Irena Lasićka, J-L Lions, and Roberto Triggiani. Non homogeneous boundary value problems for second order hyperbolic operators. *Journal de Mathématiques pures et Appliquées*, 65(2):149–192, 1986.
- Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*, 2020.

- Ziqi Liu, Wei Cai, and Zhi-Qin John Xu. Multi-scale deep neural network (MscaleDNN) for solving Poisson-Boltzmann equation in complex domains. *arXiv preprint arXiv:2007.11207*, 2020.
- Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3(3):218–229, 2021.
- 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.
- Haggai Maron, Heli Ben-Hamu, Nadav Shamir, and Yaron Lipman. Invariant and equivariant graph networks. *arXiv preprint arXiv:1812.09902*, 2018.
- Haggai Maron, Ethan Fetaya, Nimrod Segol, and Yaron Lipman. On the universality of invariant networks. In *International conference on machine learning*, pages 4363–4371. PMLR, 2019.
- Abhilash Mathews, Manaure Francisquez, Jerry W Hughes, David R Hatch, Ben Zhu, and Barrett N Rogers. Uncovering turbulent plasma dynamics via deep learning from partial observations. *Physical Review E*, 104(2):025205, 2021.
- Revanth Matthey and Susanta Ghosh. A novel sequential method to train physics informed neural networks for allen cahn and cahn hilliard equations. *Computer Methods in Applied Mechanics and Engineering*, 390:114474, 2022.
- Levi McClenny and Ulisses Braga-Neto. Self-adaptive physics-informed neural networks using a soft attention mechanism. *arXiv preprint arXiv:2009.04544*, 2020.
- Siddhartha Mishra and Roberto Molinaro. Estimates on the generalization error of physics-informed neural networks for approximating a class of inverse problems for pdes. *IMA Journal of Numerical Analysis*, 42(2):981–1022, 2022.
- Arvind T Mohan, Nicholas Lubbers, Daniel Livescu, and Michael Chertkov. Embedding hard physical constraints in neural network coarse-graining of 3d turbulence. *arXiv preprint arXiv:2002.00021*, 2020.
- 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.
- 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*, 2021.
- Ravi G Patel, Indu Manickam, Nathaniel A Trask, Mitchell A Wood, Myoungkyu Lee, Ignacio Tomas, and Eric C Cyr. Thermodynamically consistent physics-informed neural networks for hyperbolic systems. *Journal of Computational Physics*, 449:110754, 2022.

- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua Bengio, and Aaron Courville. On the spectral bias of neural networks. In *International Conference on Machine Learning*, pages 5301–5310, 2019.
- 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.
- Maziar Raissi, Alireza Yazdani, and George Em Karniadakis. Hidden fluid mechanics: Learning velocity and pressure fields from flow visualizations. *Science*, 367(6481):1026–1030, 2020.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022.
- Francisco Sahli Costabal, Yibo Yang, Paris Perdikaris, Daniel E Hurtado, and Ellen Kuhl. Physics-informed neural networks for cardiac activation mapping. *Frontiers in Physics*, 8: 42, 2020.
- Pedro Savarese and Daniel Figueiredo. Residual gates: A simple mechanism for improved network optimization. In *Proc. Int. Conf. Learn. Representations*, 2017.
- Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *Advances in Neural Information Processing Systems*, 33:7462–7473, 2020.
- Jonathan D Smith, Zachary E Ross, Kamyar Azizzadenesheli, and Jack B Muir. Hyposvi: Hypocentre inversion with stein variational inference and physics informed neural networks. *Geophysical Journal International*, 228(1):698–710, 2022.
- Hwijae Son, Jin Woo Jang, Woo Jin Han, and Hyung Ju Hwang. Sobolev training for physics informed neural networks. *arXiv preprint arXiv:2101.08932*, 2021.
- Luning Sun, Han Gao, Shaowu Pan, and Jian-Xun Wang. Surrogate modeling for fluid flows based on physics-constrained deep learning without simulation data. *Computer Methods in Applied Mechanics and Engineering*, 361:112732, 2020.
- Matthew Tancik, Pratul P Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan T Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *arXiv preprint arXiv:2006.10739*, 2020.

- 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, 2021a.
- 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, 2021b.
- Sifan Wang, Shyam Sankaran, and Paris Perdikaris. Respecting causality is all you need for training physics-informed neural networks. *arXiv preprint arXiv:2203.07404*, 2022a.
- Sifan Wang, Hanwen Wang, Jacob H Seidman, and Paris Perdikaris. Random weight factorization improves the training of continuous neural representations. *arXiv preprint arXiv:2210.01274*, 2022b.
- 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, 2022c.
- Sifan Wang, Shyam Sankaran, Hanwen Wang, and Paris Perdikaris. An expert’s guide to training physics-informed neural networks. *arXiv preprint arXiv:2308.08468*, 2023a.
- Zhicheng Wang, Xuhui Meng, Xiaomo Jiang, Hui Xiang, and George Em Karniadakis. Solution multiplicity and effects of data and eddy viscosity on navier-stokes solutions inferred by physics-informed neural networks. *arXiv preprint arXiv:2309.06010*, 2023b.
- Colby L Wight and Jia Zhao. Solving Allen-Cahn and Cahn-Hilliard equations using the adaptive physics informed neural networks. *arXiv preprint arXiv:2007.04542*, 2020.
- 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.
- 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.
- Norman J Zabusky and Martin D Kruskal. Interaction of "solitons" in a collisionless plasma and the recurrence of initial states. *Physical review letters*, 15(6):240, 1965.
- Marius Zeinhofer, Rami Masri, and Kent-André Mardal. A unified framework for the error analysis of physics-informed neural networks. *arXiv preprint arXiv:2311.00529*, 2023.
- Enrui Zhang, Ming Dao, George Em Karniadakis, and Subra Suresh. Analyses of internal structures and defects in materials using physics-informed neural networks. *Science advances*, 8(7):eabk0644, 2022.
- Linfeng Zhang, Jiequn Han, Han Wang, Roberto Car, and EJPRL Weinan. Deep potential molecular dynamics: a scalable model with the accuracy of quantum mechanics. *Physical review letters*, 120(14):143001, 2018.