

Understanding the Difficulty of Solving Cauchy Problems with PINNs

Tao Wang

TAW003@UCSD.EDU

Bo Zhao

BOZHAO@UCSD.EDU

Sicun Gao

SICUNG@UCSD.EDU

Rose Yu

ROSEYU@UCSD.EDU

University of California, San Diego

Editors: A. Abate, M. Cannon, K. Margellos, A. Papachristodoulou

Abstract

Physics-Informed Neural Networks (PINNs) have gained popularity in scientific computing in recent years. However, they often fail to achieve the same level of accuracy as classical methods in solving differential equations. In this paper, we identify two sources of this issue in the case of Cauchy problems: the use of L^2 residuals as objective functions and the approximation gap of neural networks. We show that minimizing the sum of L^2 residual and initial condition error is not sufficient to guarantee the true solution, as this loss function does not capture the underlying dynamics. Additionally, neural networks are not capable of capturing singularities in the solutions due to the non-compactness of their image sets. This, in turn, influences the existence of global minima and the regularity of the network. We demonstrate that when the global minimum does not exist, machine precision becomes the predominant source of achievable error in practice. We also present numerical experiments in support of our theoretical claims.

1. Introduction

Neural networks have attracted significant attention in the learning of dynamical systems (Wang et al., 2021; Djeumou et al., 2022). One example is solving partial differential equations (PDEs) with Physics-Informed Neural Networks (PINNs) (Sirignano and Spiliopoulos, 2018; Raissi et al., 2019), which are traditionally handled by classical methods such as finite element method (FEM) (Ern and Guermond, 2004). Despite the promise, several works have identified failure modes of PINNs. It has been observed that when solving one-dimensional Burgers equations, classical FEM can achieve the L^2 error¹ on the magnitude of 10^{-7} (Khater et al., 2008), in contrast to the error of magnitude 10^{-2} attained by PINNs (Krishnapriyan et al., 2021) or DeepONet (Lu et al., 2022).

In this paper, we develop a fundamental understanding of the failure mode of PINNs. We focus on Cauchy problems, a class of important PDEs for evolution equations in optimal control (Fleming and Rishel, 1975) and fluid dynamics (Sell and You, 2002). Our analysis reveals two main aspects of the failure mode: the use of L^2 residuals as the loss function and neural networks as function approximators. On the theoretical side, achieving zero loss on L^2 residual and initial/boundary error does not guarantee the true solution (Courant and Hilbert, 1962; Evans, 2010). PINNs solve a PDE over a given compact domain. In Section 3, we demonstrate that such a setting may conflict with the underlying evolutionary dynamics, such as the propagation of characteristics in first-order equations and the non-local behavior of parabolic equations. For instance, when solving the Burgers' equation, minimizing L^2 residual alone produces a smooth solution, whereas the true solution

1. In some work, the mean-squared error (MSE), or the square of the standard L^2 error, is used as the measure of approximation error, usually resulting in lower values. Here we convert it to the corresponding L^2 error.

exhibits high-frequency behavior as in Figure 1. In contrast, traditional methods do not encounter this issue because they "respect" the underlying dynamics and leverage it to solve equations through the method of characteristics and/or Fourier transform.

Regarding the approximation power of neural networks, the universal approximation theorem (UAT) states that a two-layer neural network with sigmoid activation can approximate *any* continuous functions if it is wide enough (Funahashi, 1989; Hornik et al., 1989; Barron, 1993). However, the achievable error in practice is usually much higher than the one UAT suggests. In Section 4, we show

that the intersection of the image set of a neural network and the closed unit ball is *not* compact in function space, which implies that having a global minimum at infinity is possible, especially when the target function has discontinuity. The influence of machine precision is significant in this case due to the exponential decay of gradients in activation functions such as *sigmoid* and *tanh*. Our new lower bound on actual error explains why neural networks cannot achieve an approximation error below certain thresholds in solving Cauchy problems that have discontinuous solutions.

Our main contribution in this paper can be summarized as follows:

- We investigate the difficulty of solving Cauchy problems with PINNs from two perspectives: the loss function of L^2 residuals minimization and the approximation gap of neural networks.
- For the learning objective, we demonstrate that solving Cauchy problems over compact domain can be ill-posed and poorly formulated due to its incapability of capturing the underlying dynamics. Error estimates are obtained for both first-order and second-order equations.
- Regarding the approximation power of neural networks, we show that there might be no global minimum when approximating discontinuous functions. It will, in turn, result in a higher error in solutions due to machine precision.
- Our results suggest two things for PINNs: first, loss functions should be designed in accordance with the evolutionary dynamics within PDEs, for instance, including regularity conditions and proper boundary values; second, when the true solution has discontinuity, neural networks are not capable of representing it accurately.

2. Related Work

Limitation of ML methods in solving PDEs. Solving partial differential equations (PDEs) is one of the core areas in scientific computing. A number of deep learning algorithms have been developed to learn PDE solutions, such as the deep Galerkin method (Sirignano and Spiliopoulos, 2018), physics-informed neural networks (PINN) (Raissi et al., 2019), Fourier neural operator (Kovachki et al., 2021b), and DeepONet (Lu et al., 2019, 2022). Although some of these methods are




Figure 1: An example where PINNs fail to solve (left), compared to the true solution (right).

proven to have the universal approximation property (Lu et al., 2019; Kovachki et al., 2021a), their performances are often not compared to traditional methods such as FEM, whose error provably approaches zero as the resolution increases. Additionally, PINN has worse accuracy compared to traditional computational fluid dynamics methods (Cai et al., 2021) and can fail to produce a physical solution in certain problems (Chuang and Barba, 2022; Wang et al., 2022a). This limitation may be attributed to spectral bias and convergence rates of different loss components (Wang et al., 2022b). It has been observed that PINN can fail to obtain the true solution even when the neural network possesses enough expressiveness to represent the objective function, which motivates analyzing the loss landscape (Fuks and Tchelepi, 2020; McClenny and Braga-Neto., 2020; Krishnapriyan et al., 2021; Zhao et al., 2022). Our work seeks to understand this limitation by focusing on the conflict between the evolutionary structure of Cauchy problems and the formulation of PINNs.

The approximation power of neural networks. Since the space of square-integrable functions over D , say $L^2(D)$, is separable (Folland, 1999), every basis can perfectly represent any function in $L^2(D)$ if it is allowed to have infinitely many parameters. Thus, the asymptotic guarantee of neural network UAT cannot distinguish neural networks from existing function approximators. Inspired by this fact, many attempts have been made to understand the approximation power of neural networks by evaluating their performance under a fixed number of parameters (Berner et al., 2021; DeVore et al., 2021; Petersen et al., 2021). In particular, the approximation power of neural networks is demonstrated in the way that they can achieve the same level of error as those by classical basis functions, such as polynomials and trigonometric functions, but with fewer parameters (Liang and Srikant, 2016; Elbrächter et al., 2019; Kim et al., 2023). The error analysis is usually done under the assumption that the target function is Lipschitz continuous. However, most of these frameworks do not apply to non-smooth and discontinuous solutions which is common in scientific computing.

3. The Difficulty of Solving Cauchy Problems with PINN Loss

In this section, we investigate why it is difficult for PINNs to obtain accurate solutions in Cauchy problems from the loss function perspective. We consider the Cauchy problem for first-order and second-order parabolic equations separately, and analyze how PINNs lead to failures in these cases.

3.1. Problem formulation

Consider the general form of Cauchy problems:

$$\begin{cases} u_t + \mathcal{F}(x, t, u, \nabla u, \Delta u) = 0, & (x, t) \in \mathbb{R}^d \times [0, T]; \\ u(x, 0) = \phi(x), & x \in \mathbb{R}^d; \end{cases} \quad (1)$$

where $\mathcal{F}(\cdot)$ is the differential operator representing the PDE model, $u(x, t) \in \mathbb{R}^d \times [0, T]$ is the solution, (x, t) is the space-time pair, $[0, T]$ is the time domain and $\phi(\cdot)$ is the initial condition. Under the framework of PINNs, the solution $u(x, t) = u(x, t; \theta)$ is represented by a neural network parameterized by $\theta \in \mathbb{R}^N$. The neural network is trained to minimize the following loss function, which combines the L^2 residual of the PDE model and the fitting error on the initial condition:

$$J(\theta) = \|u_t(\cdot; \theta) + \mathcal{F}(u(\cdot; \theta))\|_{L^2(U \times [0, T])}^2 + \|u(\cdot, 0; \theta) - \phi(\cdot)\|_{L^2(U)}^2. \quad (2)$$

In practice, PINNs are only able to solve PDEs over compact domains. Therefore, U in (2) is usually a compact set instead of \mathbb{R}^d . As we will show next, this limitation may prevent PINNs from obtaining the true solutions.

For the remainder of this section, our focus is on the L^2 error between the optimal solution $v \in C^2(U \times [0, T])$ obtained by PINNs, which achieves zero loss in (2), and the true solution $u(\cdot)$. This assumption implies that $v_t + \mathcal{F}(x, t, v, \nabla v, \Delta v) = 0$ for all $(x, t) \in U \times [0, T]$ and $v(x, 0) = \phi(x)$ for all $x \in U$.

3.2. First-order PDEs

Generally, the solution of a Cauchy problem is determined by three factors: the differential operator, initial conditions and regularity conditions. While many PDEs do not have solutions in the classical sense, others possess infinitely many qualified solutions, among which only one represents the real physical solution. Undesired solutions are eliminated by introducing regularity conditions. Therefore, minimizing the PINN objective in (2) itself is insufficient to determine the true solution as it only reflects the differential operator and initial condition without any regularity conditions.

Non-uniqueness of the solution. Although regularity conditions assume some level of smoothness on the true solution, smoothness alone does not guarantee uniqueness. According to the Rademacher's theorem, a locally Lipschitz continuous function is differentiable *almost everywhere*. However, the following theorem shows that Lipschitz continuity is not sufficient to determine the uniqueness of the solution:

Theorem 1 *There exists a first-order PDE with proper initial/boundary conditions that has infinitely many Lipschitz continuous solutions.*

The overflow of characteristics. One of the most important properties of first-order PDEs is that their solutions are determined by the propagation of characteristics, a principle that also applies to the solution obtained by minimizing the PINN loss. For instance, consider the one-dimensional transport equation:

$$\begin{cases} u_t + b(x) u_x = c(x), & (x, t) \in \mathbb{R} \times (0, \infty); \\ u(x, 0) = \phi(x), & x \in \mathbb{R}. \end{cases} \quad (3)$$

where $\phi \in C^2(\mathbb{R})$ is the initial condition, $b(\cdot), c(\cdot)$ are sufficiently smooth. Let $D_x \subset \mathbb{R}^d$ be a compact set and $D = D_x \times [0, T]$, then we have the following estimation of the L^2 error:

Theorem 2 *Let $v \in C^1(D)$ be the solution that achieves zero L^2 residual and initial condition error over D , and $u \in C^1(\mathbb{R} \times [0, T])$ be the true solution of (3) over the entire strip $\mathbb{R} \times [0, T]$, then the L^2 error $\|v - u\|_{L^2(D)}$ is determined by the error along the boundary ∂D .*

Proof sketch: To see how it is derived from characteristics, consider the corresponding ODE $\dot{y} = b(y)$, with initial state $y(0) = y_0 \in \mathbb{R}$. Let $V(t) = u(y(t), t)$ and $u \in C^1(\mathbb{R} \times [0, T])$ be a function that satisfies (3) for some $T > 0$. Then, we have

$$\dot{V}(t) = u_t(y(t), t) + u_x(y(t), t) \cdot \dot{y} = u_t(y(t), t) + b(y(t)) u_x(y(t), t) = c(y(t))$$

for all $t \geq 0$. Integrating from 0 to t and applying the initial condition $V(0) = \phi(y_0)$ yields $u(y(t), t) = \phi(y_0) + \int_0^t c(y(s)) ds$. Therefore, the solution is uniquely determined along the

characteristic curve $y(\cdot)$. However, the PINN loss considers a compact domain $D_x \subset \mathbb{R}$ and an arbitrary point $(x, t) \in D_x \times [0, T]$, the characteristics passing through (x, t) is not guaranteed to stay in D before reaching the axis $\mathbb{R} \times \{t = 0\}$ when going backward in time. In this case, we have $u(y(t), t) = u(y_{t_0}, t_0) + \int_{t_0}^t c(y(s)) ds$, where $y_{t_0} \in \partial D_x$ is on the boundary with $0 < t_0 \leq t$ (illustrated in Figure 2). Therefore, the error between u and v at any interior point (x, t) is equal to the error at the initial point of the corresponding characteristics, which is

$$|u(y(t), t) - v(y(t), t)| = |u(y_{t_0}, t_0) - v(y_{t_0}, t_0)|$$

where $y(t) = x$ and $y_{t_0} \in \partial D_x$ if $t_0 > 0$. Note that $u(y(t), t) = v(y(t), t)$ when $t_0 = 0$, since $v(x, 0) = u(x, 0)$ is assumed.

To better illustrate, let's consider a simple example of solving (3) over the compact domain $(x, t) \in D = [-1, 1] \times [0, 1]$ in the case of $b(x) \equiv 1$, $c(x) \equiv 0$ and $T = 1$. In this case, for any point inside the region $\{(x, t) : -1 \leq x \leq 0, x + 1 \leq t \leq 1\}$, the backward characteristics that passes through it will hit the boundary $\{-1\} \times [0, 1]$ first as in Figure 2, which leads to the following analytical formula for the L^2 error:




Figure 2: The shaded region is not covered by valid characteristics.

Corollary 3 *Let $v \in C^1(D)$ be the solution that achieves zero L^2 residual and initial error over D , and $u \in C^1(\mathbb{R} \times [0, T])$ be the true solution of (3) over the entire strip $\mathbb{R} \times [0, T]$ with $b \equiv 1$ and $c \equiv 0$, then the L^2 error is given by*

$$\|v - u\|_{L^2(D)} = \sqrt{\int_0^1 \int_0^t |v(-1, t-x) - u(-1, t-x)|^2 dx dt}. \quad (4)$$

where the term on the right-hand side depends only on the boundary values of u and v .

Since we do not have access to the correct value on the boundary $\{-1\} \times [0, 1]$ beforehand, the error estimate (4) heavily relies on the initialization of network parameters and hence cannot be controlled. In contrast, traditional methods solve (3) along the propagation of characteristics instead of over a pre-defined compact domain, thereby providing accurate solutions.

3.3. Second-order parabolic equations

For parabolic equations, the L^2 error between a solution that achieves zero loss in (2) and the true solution can be arbitrarily large due to the non-local behavior of the solution. To make it precise, consider the d -dimensional parabolic equation:

$$\begin{cases} u_t + \mathcal{F}(x, t, u, \nabla u, \Delta u) = 0, & (x, t) \in \mathbb{R}^d \times (0, T]; \\ u(x, 0) = \phi(x), & x \in \mathbb{R}^d; \end{cases} \quad (5)$$

where $\mathcal{F}(x, t, u, \nabla u, \Delta u) = -\sum_{i,j=1}^d a_{ij}(x, t)u_{x_i x_j} + \sum_{i=1}^d b_i(x, t)u_{x_i} + c(x, t)u$ and the coefficients $a_{ij}(\cdot, \cdot), b_i(\cdot, \cdot), c(\cdot, \cdot)$ are sufficiently smooth and $\phi \in C^2(\mathbb{R})$ is the initial condition. Also, we assume that there exists $\gamma > 0$ such that $\sum_{i,j=1}^d a_{ij}(x, t)\xi_i \xi_j \geq \gamma|\xi|^2$ for all $x \in U$

and $\xi = [\xi_1, \dots, \xi_d] \in \mathbb{R}^d$ so that the matrix $(a_{ij}(x, t))_{i,j}$ is always positive definite. Under these assumptions, the solution of (5) can be written in the following form:

$$u(x, t) = \int_{\mathbb{R}^d} \Gamma(x, t; \xi, \tau) \phi(\xi) \, d\xi, \quad (x, t) \in \mathbb{R}^d \times [0, T], \quad (6)$$

where $\Gamma(x, t; \xi, \tau)$ is called the fundamental solution of the parabolic equation whose explicit form can be found in the Lemma 2 of Dressel (1940). Note that when (5) is the one-dimensional heat equation $u_t = u_{xx}$, the fundamental solution $\Gamma(x, t; \xi, \tau) = \frac{1}{2\sqrt{\pi t}} e^{-\frac{|x-\xi|^2}{4t}}$ is exactly the heat kernel.

Let $D = D_x \times [0, T]$, with compact subset $D_x \subset \mathbb{R}^d$, denote the compact domain over which we want to solve (5). For each solution candidate $\eta \in C^2(D)$, the objective functional is given by

$$P(\eta) = \|\eta_t + \mathcal{F}(x, t, \eta, \nabla\eta, \Delta\eta)\|_{L^2(D)}^2 + \|\eta - \phi\|_{L^2(D_x)}^2.$$

Suppose that $\phi_2 \in C^2(\mathbb{R}^d)$ is another initial condition that has $\phi_2(x) = \phi(x)$ for all $x \in D_x$, and let $v(x, t) = \int_{\mathbb{R}^d} \Gamma(x, t; \xi, \tau) \phi_2(\xi) \, d\xi$. It can be verified that v solves (5) over D , i.e., $P(v) = 0$. However, ϕ_2 is arbitrary outside D_x and the kernel $\Gamma(x, t; \xi, \tau)$ is positive for all $x \in \mathbb{R}^d$, which means that we can choose ϕ_2 such that the resulting error is arbitrarily large. Therefore, we have proved the following result:

Theorem 4 *For any $K > 0$ and compact set $D_x \subset \mathbb{R}^d$, there exists a function $v \in C^2(D_x \times [0, T])$ that achieves zero L^2 residual and initial condition error over $D = D_x \times [0, T]$, such that $\|v - u\|_{L^2(D)} > K$ where $u(\cdot)$ is the true solution of (5).*

In summary, the non-local property of parabolic equations, induced by the kernel function $\Gamma(x, t; \xi, \tau)$, makes the PINN formulation of the Cauchy problems ill-posed. This leads to the difficulty in learning the true solution over any compact domains, since even achieving zero training loss does not theoretically guarantee small L^2 error.

4. The Approximation Gap of Neural Networks

Aside from the limitations of the PINN loss, the structure of neural networks can also be a problem when representing the solution, especially when it is discontinuous. In this section, we focus on the special structure of their image sets and how it affects the accuracy through machine precision. Let $f : \mathbb{R}^N \rightarrow L^2(D)$ denote a neural network, $D \subset \mathbb{R}^d$ be a compact set and \mathbb{R}^N the parameter space.

4.1. The topology of image set $Im(f)$

Now we study the complexity of the image set $Im(f) \subset L^2(D)$. Since unbounded sets are never compact in normed spaces, we will focus on their intersections with the closed unit ball \mathcal{B} in $L^2(D)$, namely $\overline{Im(f)} \cap \mathcal{B}$, to better describe the complexity of the closure of image set $\overline{Im(f)}$. Every closed and bounded set in finite-dimensional Euclidean spaces is compact, but closed and bounded sets may not be compact when the underlying space is infinite-dimensional, such as in L^2 spaces. Therefore, compactness is no longer a trivial implication of boundedness plus closedness. To get a better sense of what compact sets in $L^2(D)$ look like, the following theorem provides a full characterization of when a bounded set has compact closure in function spaces:

Proposition 5 (*Fréchet-Kolmogorov theorem, Brezis (2010)*) Let $1 \leq p < \infty$ and $D \subset \mathbb{R}^d$ be a bounded measurable set. Then, for any bounded set $S \subset L^p(D)$, its closure \bar{S} is compact if and only if for any $\epsilon > 0$, there exists $\delta > 0$ such that $\|f(x+h) - f(x)\|_p < \epsilon$ for all $f \in S$ and all $h \in \mathbb{R}^d$ with $|h| < \delta$.

The above condition is equivalent to the statement that the functions in S have $\lim_{|h| \rightarrow 0} \|f(x+h) - f(x)\|_p = 0$ uniformly, indicating that S cannot contain "too many" discontinuous functions. For instance, consider the set $S = \{g_i\}_{i=1}^\infty \subset L^2([0, 1])$ where $g_i(x) = \sqrt{i}\chi_{[0, \frac{1}{i}]}(x)$, then for any $\delta > 0$, there exists $i_0 > 2\delta^{-1}$ such that $\|g_{i_0}(x + \frac{\delta}{2}) - g_{i_0}(x)\|_2 = \sqrt{2}$, which implies that S is not compact according to Proposition 5.

Next, we consider the image set of a neural network whose activation function is either *sigmoid* or *ReLU*. Neural networks may converge to step functions as the corresponding weights approach infinity as in Figure 3, which leads to the following result:

Theorem 6 Let $f : \mathbb{R}^N \rightarrow L^2(D)$ be a neural network with sigmoid or ReLU activation function, having a width greater than or equal to 2 and at least 2 hidden layers. Then, the set $Im(f) \cap \mathcal{B}$ is not compact.

The non-compactness of neural network image sets can lead to the following two problems:

- *Global minimum at infinity*: Usually, if there exists a global minimum, it has a region of attraction in its neighborhood in which the objective function is strictly convex and guarantees local convergence. When there is no global minimum, however, the minimal achievable error is much higher as discussed in 4.2.
- *Loss of regularity*: While neural networks represent Lipschitz continuous functions with bounded parameters, they may lose smoothness as weights approach infinity. According to Theorem 6, there could be a sequence of parameters $\{\theta_i\} \subset \mathbb{R}^N$ that produces a decreasing sequence $\{J(\theta_i)\}$ but diverges to infinity itself. While each neural network $f(\theta_i)$ is Lipschitz continuous for all $i \in \mathbb{N}$, the continuity of $\lim_{i \rightarrow \infty} f(\theta_i)$ is not always implied (cite one or two papers here.). Therefore, regularity cannot be guaranteed through training, even if the training loss continues to decrease.




Figure 3: Neural networks can converge to step functions.

4.2. Unreachable Infinity

Having the global minimum at infinity would not affect trainability if we could approach it. Unfortunately, this is not the case due to machine precision since the weights in a neural network must tend towards infinity, which is prohibited by exponentially-decaying gradients for activation functions like *sigmoid* or *tanh*. In particular, we have the following estimate:

Theorem 7 (*Lower bound on actual error*) Let f be a neural network with sigmoid or tanh activation functions and ϕ be the target function. Suppose that the target function ϕ has discontinuity, then the smallest attainable L^2 error $e \sim \mathcal{O}(\sqrt{\Delta x} |\log \epsilon|^{-\frac{1}{2}})$ where $\Delta x < 1$ is the grid size and $\epsilon > 0$ is the machine precision.




Figure 5: (a) Solution by Adam; (b) Solution by SGD; (c) Training curves; (d) True solution of (7).

A detailed discussion can be found in Appendix A. Theorem 7 suggests that the unreachable infinity problem is hard to address for two reasons: first, increasing the machine precision does not significantly improve the performance, as the minimum of error is proportional to $|\log \epsilon|^{-\frac{1}{2}}$ which decays very slowly (Figure 4); second, The discretization error is a fundamental limitation present in both PINNs and traditional methods, suffering from the curse of dimensionality (Aubin, 2000). Additional discussions on numerical issues when the global minimum does not exist can also be found in Glorot and Bengio (2010) and Gallon et al. (2022).




Figure 4: The minimal achievable error is determined by $|\log \epsilon|$.

5. Experiments

In this section, we validate our theoretical results in the one-dimensional Burgers' equation (Basdevant et al., 1986):

$$\begin{cases} \frac{\partial u}{\partial t} + \mu u \frac{\partial u}{\partial x} = \nu \frac{\partial^2 u}{\partial x^2}, & x \in [-1, 1], t \in [0, 1]; \\ u(x, 0) = \sin\left(\frac{\pi x}{2}\right), & x \in \mathbb{R}; \end{cases} \quad (7)$$

where $\mu = -1, \nu = 10^{-3}$. The Burgers' equation is a nonlinear partial differential equation that models a combination of advection and diffusion. This system is carefully selected because both of its initial condition and differential operator are smooth, but they finally lead to a discontinuous solution at $t = 1$. We first solve (7) by minimizing its PINN objective. Then, we train on a new loss function that combines the L^2 residual with data from the true solution.

Solving with PINN Loss. The solution is represented by a 2-hidden-layer *sigmoid* network whose size is 256×256 . We minimize the loss function (2) via SGD and Adam separately, and the results are presented in Figure 5. It turns out that both optimizers provide smooth approximations that do not match the true solution.

Approximation gap of neural networks. To examine the limitation of neural network approximation for discontinuous solutions, we minimize the new loss function

$$\tilde{J}(\theta) = \frac{1}{K} \sum_{i=1}^K |u(x_i, t_i; \theta) - u_i|^2,$$

where u_i are the values of true solution at (x_i, t_i) , which are collected by computing the numerical solution using the spectral method implemented in Binder (2021). The time resolution is 0.01 and spatial resolution is 0.001. Figure 5 (d) visualizes the solution. We evaluate the mean squared error of fully connected networks with *sigmoid* activation trained under SGD. In the first set of experiment, we use a 2 layer network with hidden size ranging from 2 to 256. In the second set, we vary the number of layers while keeping all hidden dimension at 4. Figure 6 (b,c) show that a better approximation is obtained than using L^2 residual only. They also show the approximation error of neural networks with different width and depth. Figure 7 visualizes the predicted solution against the numerical solution at different t .




Figure 6: (a) Numerical solution represented by neural network; (b) Neural network approximation error vs width; (c) Neural network approximation error vs number of layers.




Figure 7: Comparison of predicted and analytical solution of Burgers' equation at different time: (a) $t = 0$; (b) $t = 0.25$; (c) $t = 0.5$; (d) $t = 0.75$; (e) $t = 1$. A shock wave is formed at $t = 1$.

Analysis. Figure 6 (a) shows that it is possible to use neural networks to represent the true solution. This verifies that the failure case in Figure 5 is due to the limitation of PINN loss (2) in capturing the emergence of shock waves as suggested in Section 3. On the other hand, even with the real data, the best achieved L^2 error still exceeds the square root of the resolution ($\delta = 10^{-3}$), which further validates the claim on machine precision in Theorem 7.

6. Conclusion

In this work, we discuss the fundamental limitations of physics-informed neural networks (PINNs) in solving Cauchy problems. However, these issues rarely affect classical methods, prompting a reconsideration of using neural networks for scientific computing. While we focus on the aspects of L^2 residual and neural network approximation, it is worth noting that various factors, such as the sampling scheme and optimizer choice, can also influence the numerical performance. Overall,

we advocate for a comprehensive understanding of the foundation of dynamical models for future integration of deep learning in scientific computing.

Appendix A. Details on the Machine Precision Result

Consider a finite dataset $X = \{x_1, \dots, x_N\} \subset \mathbb{R}^d$, and a neuron $f(x; \theta) = \sigma(w^T x + b)$ where σ is the *sigmoid* activation function. Since the precision of a computing machine is always finite, let $\epsilon = 2^{-p}$ be the precision where $p \in \mathbb{N}$. Then for any quantity $\eta \in \mathbb{R}$ with $|\eta| < \epsilon$, it yields $\eta = 0$ on the machine. Let $\phi(\cdot)$ be the target function with $\max_{x \in D} |\phi(x)| \leq M$. For each $x \in D$, the norm of the gradient of approximation error $L(x, \theta) = (f(x; \theta) - \phi(x))^2$ with respect to w is

$$|\nabla_w L(x, \theta)| = 2|f(x; \theta) - \phi(x)| \left| \frac{2xe^{-(w^T x + b)}}{(1 + e^{-(w^T x + b)})^2} \right| \leq 2(M + 1)M'e^{-|w^T x + b|}$$

when $M' = \max_{x \in D} |x|$. Similarly, we have $|\nabla_b L(x, \theta)| \leq 2(M + 1)e^{-|w^T x + b|}$. Therefore, $\nabla_{\theta} L(x, \theta)$ is zero on the machine when both quantities are less than ϵ , i.e., there exists

$$\delta = \max\{(p - 1) \log 2 - \log((M + 1)M'), (p - 1) \log 2 - \log(M + 1), 0\}$$

such that $\nabla_{\theta} L(x_i, \theta) = 0$ for all data point x_i outside the region $\mathcal{E}(w, b) = \{x \in \mathbb{R}^d : |w^T x + b| \geq \delta\}$. In other words, such points have no contribution to the gradient of the loss objective with respect to the parameter of this neuron. Notice that $\mathcal{E}(w, b)$ is exactly the collection of all points whose distance to the hyperplane $w^T x + b = 0$ is $\frac{\delta}{\|w\|}$ where the constant δ depends only on the precision and the target function, the width of effective region $\mathcal{E}(w, b)$ decays as $\|w\|$ grows at the rate of $\mathcal{O}(\|w\|^{-1})$.

Suppose the sample domain D is partitioned into equal grids of length $\delta > 0$, the gradient of the MSE $\mathcal{L}(\theta)$ is actually estimated by

$$\nabla \mathcal{L}(\theta) \simeq \frac{1}{N} \sum_{i=1}^N \frac{\partial L}{\partial \theta}(x_i, \theta)$$




Figure 8: The effective region $\mathcal{E}(w, b)$.

which means that the parameters will stop updating when the effective region is too small to contain any data points.

For instance, consider the approximation of the target function $\phi = \chi_{[0,1]}$ over the domain $D = [-1, 1]$ using a single neuron $f(x; w) = \sigma(wx)$, so that the global minimum is at $w = \infty$ with zero loss. Now consider the partition $x_i = -1 + \frac{i-1}{K}$ for $i = 1, 2, \dots, 2K + 1$, according to the previous results, the largest reachable value of $\|w\|$ by gradient descent cannot exceed $\frac{(p-1) \log 2}{\Delta x}$ where $\Delta x = \frac{1}{K}$ is the grid size.

On the other hand, the actual L^2 error at $w' = \frac{(p-1) \log 2}{\Delta x}$ is

$$\|f(\cdot; w') - \phi\|_2 = \sqrt{\frac{2}{w'} \left(\log \left(\frac{2}{1 + e^{-w'}} \right) + 1 - \frac{2}{1 + e^{w'}} \right)} \sim \mathcal{O}(\sqrt{\Delta x} |\log \epsilon|^{-\frac{1}{2}}) \quad (8)$$

when $|w'| \gg 1$, which then leads to Theorem 7.

Acknowledgments

This work was supported in part by Army-ECASE award W911NF-23-1-0231, the U.S. Department Of Energy, Office of Science under #DE-SC0022255, IARPA HAYSTAC Program, CDC-RFA-FT-23-0069, NSF Grants #2205093, #2146343, and #2134274.

We would like to thank the anonymous reviewers for their valuable suggestions.

References

- J.-P. Aubin. *Applied Functional Analysis*. Wiley, 2000.
- A. R. Barron. Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information Theory*, 39(3):930–945, 1993.
- C. Basdevant, M. Deville, P. Haldenwang, J. M. Lacroix, J. Ouazzani, R. Peyret, P. Orlandi, and A. T. Patera. Spectral and finite difference solutions of the Burgers equation. *Computers & Fluids*, 14(1):23–41, 1986.
- J. Berner, P. Grohs, G. Kutyniok, and P. Petersenz. The modern mathematics of deep learning. *arXiv preprint arXiv:2105.04026*, 2021.
- S. Binder. Étude de l’observation et de la modélisation des ondes de surface en eau peu profonde., 2021. URL https://github.com/sachabinder/Burgers_equation_simulation.
- H. Brezis. *Functional Analysis, Sobolev Spaces and Partial Differential Equations*. Springer, 2010.
- S. Cai, Z. Mao, Z. Wang, M. Yin, and G. E. Karniadakis. Physics-informed neural networks (PINNs) for fluid mechanics: A review. *Acta Mechanica Sinica*, 37(12):1727–1738, 2021.
- P.-Y. Chuang and L. A. Barba. Experience report of physics-informed neural networks in fluid simulations: pitfalls and frustration. *arXiv preprint arXiv:2205.14249*, 2022.
- R. Courant and D. Hilbert. *Methods of Mathematical Physics*. Wiley-Interscience, 1962.
- R. DeVore, B. Hanin, and G. Petrova. Neural network approximation. *Acta Numerica*, pages 327–444, 2021.
- F. Djeumou, C. Neary, E. Goubault, S. Putot, and U. Topcu. Neural networks with physics-informed architectures and constraints for dynamical systems modeling. *Proceedings of The 4th Annual Learning for Dynamics and Control Conference, PMLR*, 168:263–277, 2022.
- F. G. Dressel. Fundamental solution of the parabolic equation, II. *Duke Mathematical Journal*, 7: 186–203, 1940.
- D. Elbrächter, D. Perekrestenko, P. Grohs, and H. Bölcskei. Deep neural network approximation theory. *arXiv preprint arXiv:1901.02220*, 2019.
- A. Ern and J.-L. Guermond. *Theory and Practice of Finite Elements*. Springer, 2004.
- L. C. Evans. *Partial Differential Equations*. American Mathematical Society, 2010.

- W. Fleming and R. W. Rishel. *Deterministic and Stochastic Optimal Control*. Springer-Verlag, 1975.
- G. B. Folland. *Real Analysis: Modern Techniques and Their Applications*. John Wiley, 1999.
- O. Fuks and H. A. Tchelepi. Limitations of physics informed machine learning for nonlinear two-phase transport in porous media. *Journal of Machine Learning for Modeling and Computing*, 2020.
- K. I. Funahashi. On the approximate realization of continuous mappings by neural networks. *Neural Networks*, 2:183–192, 1989.
- D. Gallon, A. Jentzen, and F. Lindner. Blow up phenomena for gradient descent optimization methods in the training of artificial neural networks. *arXiv preprint arXiv:2211.15641*, 2022.
- X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks. *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 249–256, 2010.
- K. Hornik, M. Stinchcombe, and H. White. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2:359–366, 1989.
- A. H. Khater, R. S. Temsah, and M. M. Hassan. A Chebyshev spectral collocation method for solving Burgers-type equations. *Journal of Computational and Applied Math*, 222(2):333–350, 2008.
- N. Kim, C. Min, and S. Park. Minimum width for universal approximation using relu networks on compact domain. *arXiv preprint arXiv:2309.10402*, 2023.
- N. Kovachki, S. Lanthaler, and S. Mishra. On universal approximation and error bounds for fourier neural operators. *The Journal of Machine Learning Research*, 22(1):13237–13312, 2021a.
- N. Kovachki, Z. Li, B. Liu, K. Azizzadenesheli, K. Bhattacharya, A. Stuart, and A. Anandkumar. Neural operator: Learning maps between function spaces. *arXiv preprint arXiv:2108.08481*, 2021b.
- A. Krishnapriyan, A. Gholami, S. Zhe, R. Kirby, and M. W. Mahoney. Characterizing possible failure modes in physics-informed neural networks. *Advances in Neural Information Processing Systems*, 34:26548–26560, 2021.
- S. Liang and R. Srikant. Why deep neural networks for function approximation? *arXiv preprint arXiv:1610.04161*, 2016.
- L. Lu, P. Jin, and G. E. Karniadakis. Deeponet: Learning nonlinear operators for identifying differential equations based on the universal approximation theorem of operators. *arXiv preprint arXiv:1910.03193*, 2019.
- L. Lu, R. Pestourie, S. G. Johnson, and G. Romano. Multifidelity deep neural operators for efficient learning of partial differential equations with application to fast inverse design of nanoscale heat transport. *Physical Review Research*, 4(2):023210, 2022.

- L. McClenny and U. Braga-Neto. Self-adaptive physics-informed neural networks using a soft attention mechanism. *arXiv preprint arXiv:2009.04544*, 2020.
- P. Petersen, M. Raslan, and F. Voigtlaender. Topological properties of the set of functions generated by neural networks of fixed size. *Foundations of Computational Mathematics*, 21:375–444, 2021.
- M. Raissi, P. Perdikaris, and G. 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.
- G. R. Sell and Y. You. *Dynamics of Evolutionary Equations*. Springer, 2002.
- J. Sirignano and K. Spiliopoulos. DGM: A deep learning algorithm for solving partial differential equations. *Journal of computational physics*, 375:1339–1364, 2018.
- C. Wang, S. Li, D. He, and L. Wang. Is L^2 physics-informed loss always suitable for training physics-informed neural network? *arXiv preprint arXiv:2206.02016*, 2022a.
- R. Wang, D. Maddix, C. Faloutsos, Y. Wang, and R. Yu. Bridging physics-based and data-driven modeling for learning dynamical systems. *Proceedings of The 4th Annual Learning for Dynamics and Control Conference, PMLR*, 144:385–398, 2021.
- S. Wang, X. Yu, and P. Perdikaris. When and why PINNs fail to train: A neural tangent kernel perspective. *Journal of Computational Physics*, 449:110768, 2022b.
- Z. Zhao, X. Ding, G. Atulya, A. Davis, and A. Singh. Physics informed machine learning with misspecified priors: an analysis of turning operation in lathe machines. *AAAI 2022 Workshop on AI for Design and Manufacturing (ADAM)*, 2022.