Lagrangian PINNs:
A causality–conforming solution to failure modes
of physics-informed neural networks

Rambod Mojgani[a], Maciej Balajewicz[b], Pedram Hassanzadeh[a,c]

[a]*Department of Mechanical Engineering, Rice University, Houston, TX*
[b]*Realtor.com, Santa Clara, CA*
[c]*Department of Earth, Environmental and Planetary Sciences, Rice University, Houston, TX*

Abstract

Physics–informed neural networks (PINNs) leverage neural–networks to find the solutions of partial differential equation (PDE)–constrained optimization problems with initial conditions and boundary conditions as soft constraints. These soft constraints are often considered to be the sources of the complexity in the training phase of PINNs. Here, we demonstrate that the challenge of training (i) persists even when the boundary conditions are strictly enforced, and (ii) is closely related to the Kolmogorov n–width associated with problems demonstrating transport, convection, traveling waves, or moving fronts. Given this realization, we describe the mechanism underlying the training schemes such as those used in eXtended PINNs (XPINN), curriculum regularization, and sequence–to–sequence learning. For an important category of PDEs, i.e., governed by non–linear convection–diffusion equation, we propose reformulating PINNs on a Lagrangian frame of reference, i.e., LPINNs, as a PDE–informed solution. A parallel architecture with two branches is proposed. One branch solves for the state variables on the characteristics, and the second branch solves for the low–dimensional characteristics curves. The proposed architecture conforms to the causality innate to the convection, and leverages the direction of travel of the information in the domain. Finally, we demonstrate that the loss landscapes of LPINNs are less sensitive to the so–called "complexity" of the problems, compared to those in the traditional PINNs in the Eulerian framework.

Keywords: Deep learning, Kolmogorov n–width, Partial differential equations, Method of characteristics, Lagrangian frame of reference, Physics-informed neural network

1. Introduction

The evolution of many physical phenomena and engineering systems can be derived from first principles leading to governing equations in the form of partial differential equations (PDEs). Although analytical solutions of many of non–linear PDEs are seldom known, development of numerical methods has made approximation of the solution possible. One of the paradigms of solving for a state of the system is through optimization. A solution satisfies the governing PDE,

$$R\left(w\left(x,t\right)\right) := \frac{\partial w\left(x,t\right)}{\partial t} - \mathcal{N}\left(w\left(x,t\right)\right) = 0, \tag{1}$$

where $w\left(x,t\right)$ is the state parameter on a spatial domain of $x \in \Omega$ and $t \in [0,T]$ with appropriate boundary and initial conditions. Therefore, finding a solution of a PDE, w^*, is equivalent to finding a minimizer of the residual equation, i.e.,

$$w^* = \underset{w}{arg\min}\, R\left(w\left(x,t\right)\right), \tag{2}$$

subject to boundary and initial conditions as constraints. Iterative methods are traditionally used to find minimizers of high–dimensional non–linear residual equations.

On the other hand, in the absence of the governing equations, where the phenomena/task cannot be described using first principles, machine learning (ML) methods such as artificial neural networks (ANNs) are recognized to be compelling. Therefore, the application of ANNs to solve the systems with known governing equations seem to be superfluous. Nevertheless, in recent years, these paradigms have merged in physics–informed neural networks (PINNs) [1], such that ANNs are trained to find the minimizer of PDE–constrained optimization.

Despite the maturity of the developed numerical methods of solving PDEs, the flexibility in implementation, readily available adjoints via automatic differentiation, and the low inference cost of PINNs have made it an appealing tool [2], especially in inverse problems and inverse design [3], ill–posed/conditioned problems [4], and control [5].

However, the training phase of PINNs, equivalent to solving the PDEs, faces some challenges [6–14]. The innovations and attempts to improve the accuracy of the PINNs can be classified into two categories. In the first category, the ANN architecture and loss are targeted to improve the training behavior. In [8], it is shown that the eigenvalues of the Neural Tangent Kernels (NTKs) of different loss components explains the training behavior. Accordingly, penalty weights in the loss function are adaptability determined at each iteration of the training. Similarly in [9], the unbalanced gradients of the components of the loss is associated with training failure, and annealing the learning rate is proposed. It is also demonstrated that the architecture of the network can meaningfully change the stiffness of the gradients in the learning phase, and therefore it is suggested that a specialized architecture can be beneficial to specific problems [9]. Subsequently, reformulating the constraints using Augmented Lagrangian method (ALM) [15] demonstrates a flatter/smoother loss landscape, and therefore leads to a more favorable training behavior [13], compared to the originally proposed penalty terms [1]. Notably, the loss landscape of PINNs are less smooth compared to purely data–driven ANNs [10, 13]. The raggedness of loss landscape explains why PINNs are more prone to converge to an unfavorable local minima. Such challenge depends on both the governing equations, and system parameters [6, 10, 13].

In the second category, prior knowledge/property of the system is enforced. For instance, in hyperbolic systems, e.g., inviscid Burgers' equation, total variation diminishing (TVD), and entropy inequalities can be imposed in addition to artificial viscosity [16]. In [6], the flux term of non–linear convection diffusion had to be modified to help with the accuracy of the solution. However, such solutions are PDE specific, and are not applicable to all the challenging test cases. Approaches such as curriculum learning (training on a simple problem and transferring the learned weights to the harder problems) [10], adaptive sampling (in both space and time) [14], or a very similar sequence–to–sequence learning (decomposing the temporal domain) [10] can be applied to a wider range of problems without any prior assumptions. Note that sequence–to–sequence learning [10] or adaptive sampling [14] are different than of the parallel–in–time decomposition of the temporal domain, where different networks, each corresponding to one of the temporal subdomains, are trained [11]. In many of the aforementioned studies, the same network is shown to be capable of expressing a more accurate solution, given additional considerations in the loss and training. Such experiments demonstrate that the used architectures are expressive enough, and therefore, the challenge lies in the training regime. It is argued that all such challenges can be more naturally overcome by respecting the underlying spatio–temporal *causality* [12]. One natural approach to impose such causality is by prioritizing the earlier time steps in the training phase [12].

Moreover, the type of the differential operators, i.e., parabolic, hyperbolic, or elliptic, is believed to describe the difficulty in discovery of the minima. In the case of elliptic and parabolic PDEs, the generalization error, i.e., the difference between a global minimizer of the loss and the solution to the PDEs, converges to zero under certain conditions, given enough number of data points [17]. However, similar results are lacking for hyperbolic equations. More importantly, the optimization error, i.e., the difference between the global and local minima given some data points, is poorly understood.

Table 1 summarizes challenging test cases and the proposed remedies. While chaotic systems, e.g., turbulent flow, and Lorenz 63, and higher order derivatives in governing equations, are known to be challenging since the earliest formulation of PINNs [12], they merely cannot explain the difficulty in training of non–chaotic systems, e.g., advection/convection, reaction, reaction-diffusion, Poisson's, or wave equations.

Although the previous studies have described some of the difficulties and dynamics of the training phase, there is no universal theory on convergence rate or *a priori* measure of success of the training, and the present error estimators need to be tightened for practical non–linear systems [20]. In this paper, we provide some evidence that connects

Table 1: Identified challenging problems in the training of PINNs.

Study	PDEs	Domain	Test case	Proposed remedy
[6]	Convection–diffusion	1D	Buckley–Leverett	Introducing viscosity
[14]	Reaction–diffusion Adsorption/ desorption – surface diffusion	1D 2D 3D	Allen-Cahn Cahn-Hilliard	Time adaptive (sampling/marching), and mini–batching, and regularizing the loss components
[7]	Euler	1D	Sod's shock tube	Characteristic form, Oversampling the shock
[18]	Convection–diffusion	1D	Buckley–Leverett	Tuning the flux term
[9]	Helmholtz Klein–Gordon Navier– Stokes	2D 1D 2D	N/A N/A Lid–driven cavity ($Re = 100$)	Annealing the learning rate
[11]	Advection Shallow-water	Spherical	Traveling feature	Parallel in time decomposition
[10]	Convection Reaction Reaction-diffusion	1D	Challenging in specific regimes	Curriculum regularization, or Sequence–to–sequence learning
[19]	Inviscid Burgers' Convection–diffusion	1D	Shock Rarefaction and shock	Enriching the data–set, and artificial viscosity
[8]	Wave	1D	Traveling wave	Adaptive penalty
[12]	Lorenz 63 Kuramoto–Sivashinsky Navier–Stokes	$\mathbb{R}^3$ 1D 2D	N/A Regular and chaotic Decaying turbulence ($Re = 100$)	Causal training of loss, and/or transformer architecture
[13]	Poisson's	1D 2D	High wave–number forcing	Enforcing the constraints by ALM
[16]	Euler Convection–diffusion	1D	Sod's shock tube Buckley–Leverett	TVD, and entropy inequalities
Ours	Convection–diffusion Viscous Burgers'	1D	Traveling feature, Traveling shock	Lagrangian PINN

the dimensionality of solution, in the Kolmogorov n–width sense, to the difficulties in the training phase. Moreover, we explain how the previous remedies connects to the presented description of complexity. Subsequently, we add an unrecognized complex test case, to the existing list of table 1, i.e., the Burgers' equation in the presence of a shock sweeping the domain (traveling shock). Specifically, we demonstrate that in the cases where the shock sweeps long distances, the network fails to train, a similar behavior as seen for convection equation [10]. We emphasize that PINN is successfully applied to viscous Burgers' equation where a stationary shock forms [1].

Finally, we focus on an important category of the recognized challenging cases, i.e., convection–diffusion problems. We propose Lagrangian physics–informed neural networks (LPINNs) to conform to the "causality" in the system. LPINNs' architecture is informed by the direction of travel of information in the domain, i.e., along the characteristic curves. In this architecture the solution is to be learned with an inherently reduced dimensionality, a simpler task for any of ML architectures.

The paper is organized as follows. In section 2, some of the challenging cases of training of PINNs are discussed. In section 3, we focus on an important canonical set of problems governed by non–linear convection–diffusion equation, and review the traditional architecture of PINNs, and propose LPINNs in section 3.1 and section 3.2, respectively. These architecture are compared with numerical experiments in section 4, and the conclusions is followed in section 5.

2. Kolmogorov n–width of the Failure Modes of PINNs

In this section, we connect the dimensionality of the problem to the training difficulties in PINNs. Specifically, we provide numerical evidence to connect the decay of singular values of the snapshots to the difficulties in the training phase. This measure corresponds to transport phenomena, convection, traveling waves, and moving fronts. Although limited attempts were made to quantify connection of dimensionality and slow convergence of some specific classes of ML architectures [21, 22], many of the questions regarding the choice of activation functions, norms, and architecture remains open [21]. Formal investigations of such questions are out of the scope of the present paper.

In approximation theory, Kolmogorov n–width is a measure of how close n–dimensional subspaces can approximate the solution manifold, $\mathcal{M}$ [23]. The following definitions briefly explains this measure [24].

Definition: Let $\mathcal{M}$ be a normed linear space and $\widetilde{\mathcal{M}}_n$ any n-dimensional subspace of $\mathcal{M}$. For each $x \in \mathcal{M}$, $\delta\left(x, \widetilde{\mathcal{M}}_n\right)$ shall denote the distance of the n-dimensional subspace $\widetilde{\mathcal{M}}_n$ from x, defined by

$$\delta\left(x; \widetilde{\mathcal{M}}_n\right) = \inf\left\{\|x - y\|_X : y \in \widetilde{\mathcal{M}}_n\right\}. \tag{3}$$

If there exist a $y^* \in \widetilde{\mathcal{M}}_n$ for which $\delta\left(x, \widetilde{\mathcal{M}}_n\right) = \|x - y^*\|$, then y^* is the best approximation of x from $\widetilde{\mathcal{M}}_n$. Extending the concept from a single element of x to $\mathcal{S}$, a given subset of $\mathcal{M}$, the deviation of $\mathcal{S}$ from $\widetilde{\mathcal{M}}_n$ is defined as

$$\delta\left(\mathcal{S}; \widetilde{\mathcal{M}}_n\right) = \sup_{x \in \mathcal{S}} \inf_{y \in \widetilde{\mathcal{M}}_n} \|x - y\|, \tag{4}$$

representing the worst element of $x \in \mathcal{S}$ approximated in $\widetilde{\mathcal{M}}_n$.

Definition: Kolmogorov n–width of $\mathcal{M}$, $d_n(\mathcal{M})$, is defined as

$$d_n(\mathcal{S}; \mathcal{M}) := \inf_{\widetilde{\mathcal{M}}_n} \delta\left(\mathcal{S}; \widetilde{\mathcal{M}}_n\right), \tag{5}$$

where the infimum is taken over all n–dimensional subspaces ($\widetilde{\mathcal{M}}_n$) of the state space, $\mathcal{M}$.

In the context of Petrov–Galerkin projection schemes, n–width correlates with the best achievable rate of convergence for a given snapshots [25]. The connection between singular value decomposition (SVD) of the Hankel operator and the Kolmogorov n–width is rigorously established [26, 27]. It is shown that Kolmogorov n–width is equivalent to the $(n+1)^{st}$ Hankel singular value [28]. The aforementioned studies strictly connected the singular values of the snapshots to convergence of linear subspaces. The extension of such results to non–linear manifolds, such as those discovered through training of PINNs, has not been established yet. However, we numerically evaluate whether the rate of decay of singular values could be used as a simple *a priori* guideline to the complexity of the training phase of PINNs.

In figs. 1 to 4, the snapshots of the solution and the corresponding singular value decays are plotted for given snapshots of convection equation,

$$\frac{\partial w}{\partial t} - c\frac{w}{\partial x} = 0, \tag{6}$$

reaction equation,

$$\frac{\partial w}{\partial t} - \rho w(1 - w) = 0, \tag{7}$$

reaction–diffusion equation,

$$\frac{\partial w}{\partial t} - \nu\frac{\partial^2 w}{\partial x^2} - \rho w(1 - w) = 0, \tag{8}$$

and Burgers' equation,

$$\frac{\partial w}{\partial t} - w\frac{w}{\partial x} = \nu\frac{\partial^2 w}{\partial x^2}. \tag{9}$$

where the state variable $w = w(x, t)$ exists in the domain $(x, t) \in [x_a, x_b] \times [0, T]$, and the PDEs are equipped with initial conditions $w(x, 0) = w_0(x)$, and appropriate boundary conditions at x_a, and x_b. Convection speed, viscosity, and reaction coefficient are denoted by c, ν, and ρ, respectively.

(a) $c = 0$	(b) $c = 50$	(c) Singular value spectra
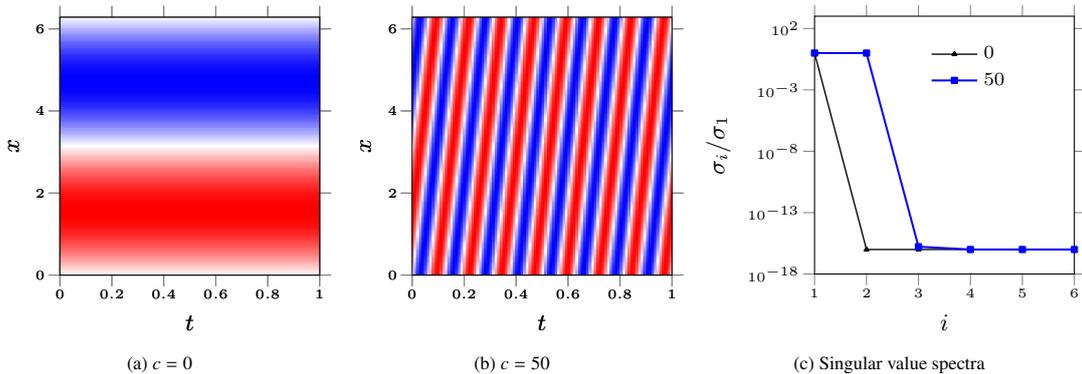
Figure 1: Convection equation. In 1a and 1b, dark red and dark blue represent $+1$ and -1, respectively. In 1c, the black triangle and blue rectangle markers represent $c = 0$ (success mode) and $c = 50$ (failure mode), respectively.

The successful and failing regimes of eqs. (6) to (8) were investigated in [10]. We identify Burgers' equation eq. (9), in certain regimes, as an additional challenging test case for PINNs. In each of these experiments, the solution of the governing equation is depicted in the space–time domain, and the corresponding singular value spectra are also compared, where σ_i is the i^{th} singular value of the snapshots of the solution.

In [10], it was shown that PINNs for convection become hard to train for $c > 10$ (failure modes). Subsequently, the cases with $c = 0$ and $c = 50$ are compared in fig. 1. It is clear that σ_2/σ_1 increases quickly as c increases. For the reaction equation, the decay of singular values becomes slower as ρ increases, in a similar trend, the PINNs are more difficult to train as ρ increases. For the cases of reaction–diffusion in [10], all cases show similar slow rate of decay of singular values, and in a similar trend, training in all the cases encounters difficulties. Finally, we consider Burgers' equation, a case successfully solved in early studies of PINNs [1]. However, in that case, a viscous shock forms and collapses on its place, without *traveling* in the domain. In this paper, to impose the shock moving through the domain, the initial condition is offset from zero. Similar to the convection case, the rate of decay of singular values decreases as the speed of the shock is increased. Subsequently, similar challenges in training of PINNs for Burgers' equation in convection–dominated regimes is observed. In such cases, the shock travels in the domain before collapsing/diffusing (see section 4.3).

These experiments show that although the neural networks (NNs) can express a non–linear manifold, the dimensionality of the solution on the linear optimal subspace of singular vectors can still inform the convergence behavior of PINNs, for comparable governing equations. However, the singular value spectra between two different governing equations do not explain difficulty of the training, e.g., collapsing shock has a slower rate of decay of singular values compared to convection with high speed and yet the training phase is well–behaved. Nevertheless, in all cases the network is hard to train at the presence of traveling features, such as shock, fronts, and gradients.

Some of the successful remedies in [7, 10, 11] can be explained by breaking the Kolmogorov n–width, a recognized paradigm in finite element method of solving PDEs [29, 30], data assimilation [31], NNs–based reduced order models (ROMs) [32–34], projection-based ROMs [30, 35–43], flexDeepONet [44], and projection–based ROMs on NN–based manifolds [45, 46]. To demonstrate the effect of the proposed remedies on the rate of decay of (normalized) singular values, we consider the synthetic data of fig. 5a, representing a traveling shock. The snapshot is constructed on $N_x = 256$ spatial grid points and $N_t = 500$ time steps.

It is shown that in PINNs, accuracy can be increased by over–sampling the traveling shock, in contrast to the uniformly sampled data [7]. Here, at each time level 25 data points are sampled (i) randomly and uniformly, (ii) randomly but weighted where the probability of sampling is proportional to the absolute value of the gradient of the data. As in fig. 5b, while sampled data has a lower rank compared to the full data, the first singular values of the data remains unchanged when sampled uniformly (a property used in randomized singular value decomposition). However, the rate of decay of the singular values increases by over–sampling the shock.

Parallel–in–time decomposition [11] and sequence-to-sequence learning [10] decompose the temporal domain into short time intervals. This strategy is similar to principal interval decomposition (PID) (in linear subspace), applied to ROMs [40] and Long short–term memory (LSTM) networks [33], and effectively reduces the Kolmogorov

(a) $\rho = 1$ (b) $\rho = 10$ (c) Singular value spectra
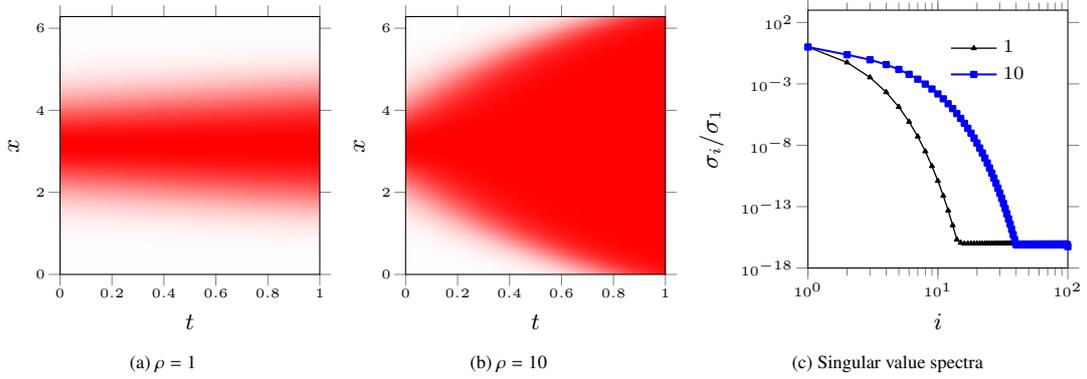
Figure 2: Reaction equation. In 2a and 2b, dark red and white represent +1 and 0, respectively. In 2c, the black triangle and blue rectangle markers represent $\rho = 1$ (success mode) and $\rho = 10$ (failure mode), respectively.

(a) $(v, \rho) = (2, 5)$ (b) $(v, \rho) = (6, 5)$ (c) Singular value spectra
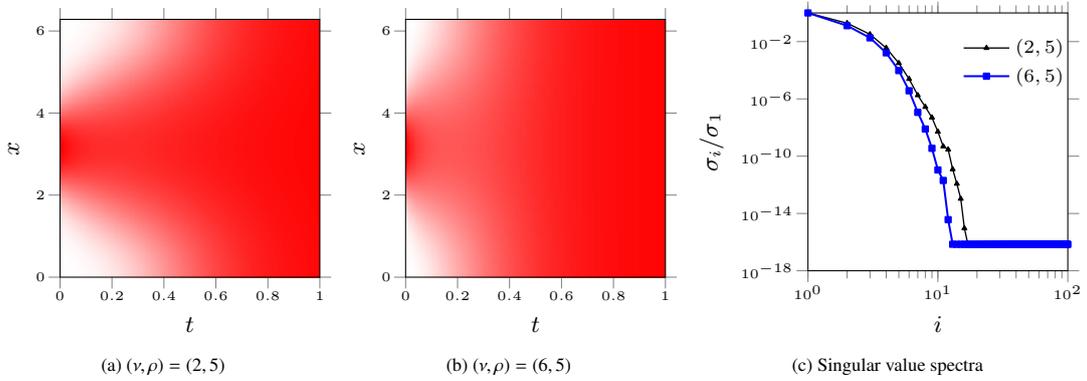
Figure 3: Reaction–diffusion equation. In 3a and 3b, dark red and white represent +1 and 0, respectively. In 3c, the black triangle and blue rectangle markers represent $(v, \rho) = (5, 2)$ (failure mode) and $(v, \rho) = (6, 5)$ (failure mode), respectively.

n–width of the data. Considering the synthetic data, the temporal domain is decomposed to 25 consecutive time steps, significantly increasing the rate of decay of singular values as in fig. 5b. Subsequently, increasing the temporal domain, decreases the rate of decay. Decomposition of the computational domain, in both space and time, is possible using extended PINNs (XPINNs), originally developed to tackle the scalibility of PINNs [47].

Given the numerical evidence in this section, one paradigm to tame the training in PINNs is to reformulate the problem on a manifold such that Kolmogorov n–width is decreased, or equivalently, the rate of decay of the singular values is increased. For non–linear convection–diffusion flows, where convection dominates diffusion, such goal is achievable by reformulating the governing equations on the characteristics curves [48, 49].

3. PINNs for Non–linear Convection–Diffusion

Consider the following scalar, one–dimensional convection–diffusion equation

$$R := \frac{\partial w(x,t)}{\partial t} + f_1(x,t,w)\frac{\partial w(x,t)}{\partial x} - f_2(x,t,w)\frac{\partial^2 w(x,t)}{\partial x^2} = 0, \tag{10}$$

in the domain $(x,t) \in [x_a, x_b] \times [0, T]$, with initial conditions $w(x, 0) = w_0(x)$, and appropriate boundary conditions at x_a and x_b.

To address the complexity of the training, the governing equation eq. (10), is reformulated in the Lagrangian frame

6
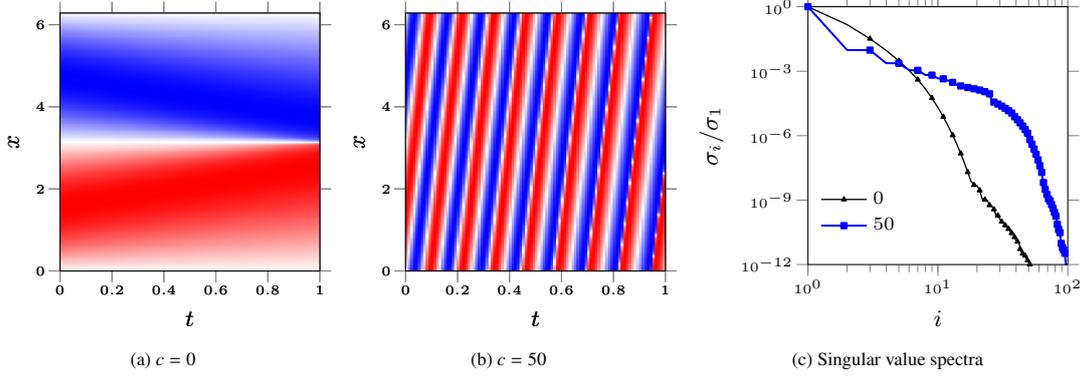
Figure 4: Burgers' equation. In 4a and 4b, dark red and dark blue represent $c + 1$ and $c - 1$, respectively. In 1c, the black triangle and blue rectangle markers represent $c = 0$ (success mode) and $c = 50$ (failure mode), respectively.
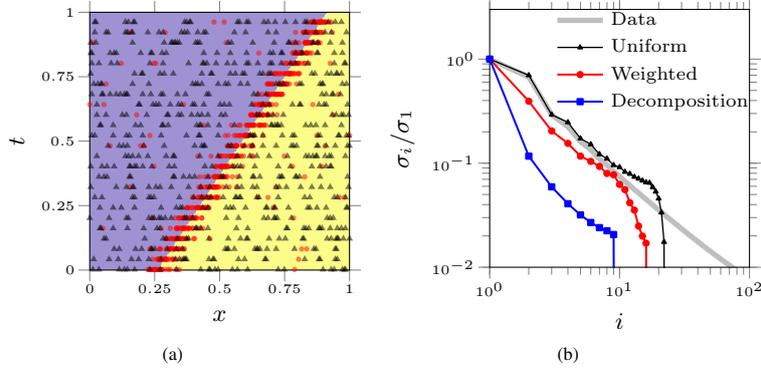
Figure 5: a. Synthetic traveling shock and the location of the samples, (i) uniform sampling (black triangles), and (ii) weighted sampling (red circles), b. The corresponding singular value spectra of the uniformly sampled data (black triangles), the weighted sampled data (red circles), and decomposed data in a limited temporal domain (blue squares).

of reference

$$R_x := \frac{dx}{dt} - f_1(x, t, w) = 0, \tag{11a}$$

$$R_w := \frac{\partial w}{\partial t} - f_2(x, t, w)\frac{\partial^2 w}{\partial x^2} = 0. \tag{11b}$$

where x is the characteristic curves and w is the state variable on the characteristic curves. For the sake of simplicity of the notations, we do not differentiate between the state variable on the Eulerian formulation (stationary grid, eq. (10)) and the Lagrangian formulation (moving grid, eq. (11)).

3.1. Traditional PINNs

A PINN architecture is composed of a densely connected ANN that minimizes a loss comprised of the residual equation eq. (10) evaluated at the collocation points, data, and the initial and boundary points [1]. The output of the network, is the state parameter, w,

$$w = \mathcal{N}(u) = \phi_m\left(W_m\phi_{m-1}\left(W_{m-1}\cdots\phi_1\left(W_1u + b_1\right)\cdots + b_{m-1}\right) + b_m\right), \tag{12}$$

where the input vector is the concatenation of the spatial and temporal location, i.e., $u = [x_i, t_i]^\top \in \mathbb{R}^{d+1}$, where d is the dimension of physical space, $\phi_i(.)$ is the activation function at the i^{th}-layer, $W_1 \in \mathbb{R}^{w\times(d+1)}$, $W_i \in \mathbb{R}^{w\times w}$, $\forall i \in \{2, \cdots, m-1\}$, and $W_w \in \mathbb{R}^{d\times w}$ are the weights, and $b_1 \in \mathbb{R}^w$, $\forall i \in \{1, \cdots, m-1\}$, and $b_m \in \mathbb{R}^d$ are biases. The

7
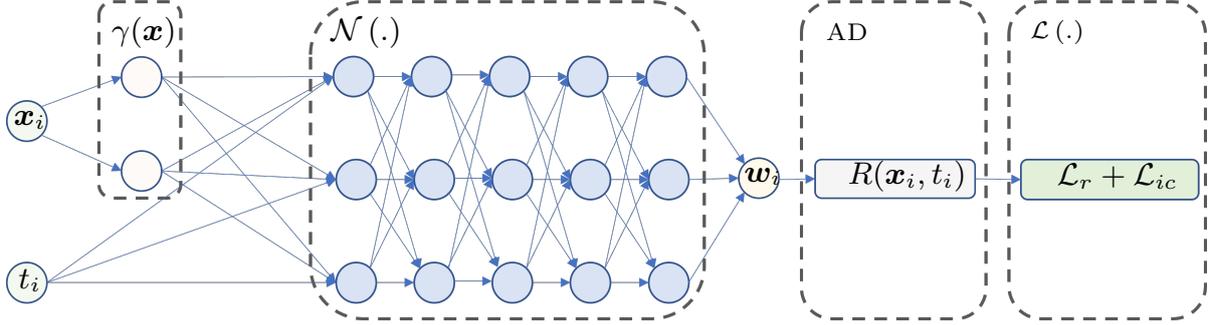
Figure 6: The traditional PINNs architecture with periodic boundary condition.

weights and biases are learned to minimize the so–called physics informed loss, minimizing the residual equation and the appropriate boundary and initial conditions, i.e.,

$$\mathcal{L} = \mathcal{L}_r + \mathcal{L}_{bc} + \mathcal{L}_{ic}, \tag{13}$$

where

$$\mathcal{L}_r = \lambda_r \frac{1}{N_r} \sum_{i=1}^{N_r} |R(\boldsymbol{x}_i, t_i)|_2, \tag{14a}$$

$$\mathcal{L}_{bc} = \lambda_{bc} \frac{1}{N_{bc}} \sum_{i=1}^{N_{bc}} \left| \mathcal{B}[w]\left(\boldsymbol{x}_{bc}^i, t_b c^i\right) \right|_2, \tag{14b}$$

$$\mathcal{L}_{ic} = \lambda_{ic} \frac{1}{N_{ic}} \sum_{i=1}^{N_{ic}} \left| \boldsymbol{w}(\boldsymbol{x}, 0) - g\left(\boldsymbol{x}_{ic}^i\right) \right|_2, \tag{14c}$$

and $\{t_r^i, x_{ic}^i\}_{i=1}^{N_r}$ is the set of temporal and spatial coordinates of the collocation points where the residual is evaluated, and $\{x_{ic}^i\}_{i=1}^{N_{ic}}$, is the set of coordinates where the initial condition is known ($g\left(\boldsymbol{x}_{ic}^i\right)$ at $t = 0$), and $\{t_{bc}^i, x_{ic}^i\}_{i=1}^{N_{bc}}$ is a set of temporal and spatial coordinates of the boundary points ($\mathcal{B}[w]\left(\boldsymbol{x}_{bc}^i, t_b c^i\right)$). The hyperparameters λ_r, λ_{bc}, and λ_{ic} are scalars tuned to enhance the convergence. Augmenting the loss using more data points leads to faster convergence, especially in convection–dominated problems [19]. However, we intentionally refrain from using any data points to evaluate the convergence of PINNs as a solver (no–data regime). To solve convection–diffusion equation, eq. (10) is used as the residual term in eq. (14a).

In this paper, without loss in generality, we limit the spatial domain to one–dimensional (1D) space. The periodic boundary condition is strictly enforced using a custom layer [11], and therefore $\lambda_{bc} = 0$. The custom layer in $x \in [0, 2\pi]$ transforms the domain to a polar coordinate, i.e.,

$$\gamma(x) = [\cos(x), \sin(x)]^\top, \tag{15}$$

and out–put of this layer is fed into the traditional NN as described in eq. (12) with appropriate adjustment of the dimension of the weight of the first layer, i.e., $\boldsymbol{W}_1 \in \mathbb{R}^{w \times (2d+1)}$, increasing the number of network variables by only $w \times d$. Further discussion of strictly enforcing the boundary conditions can be found in [50].

3.2. Proposed Lagrangian PINNs

In this section, we describe the additional changes to the architecture of the traditional PINNs to conform with the Lagrangian formulation to satisfy eq. (11). We propose a parallel architecture comprised of two branches. The first branch solves for the characteristics curves and minimizes eq. (11a), i.e.,

$$\boldsymbol{x} = \mathcal{N}_x(\boldsymbol{u}) = \phi_m\left(\boldsymbol{W}_m \phi_{m-1}\left(\boldsymbol{W}_{m-1} \cdots \phi_1\left(\boldsymbol{W}_1 \boldsymbol{u} + \boldsymbol{b}_1\right) \cdots + \boldsymbol{b}_{m-1}\right) + \boldsymbol{b}_m\right), \tag{16}$$
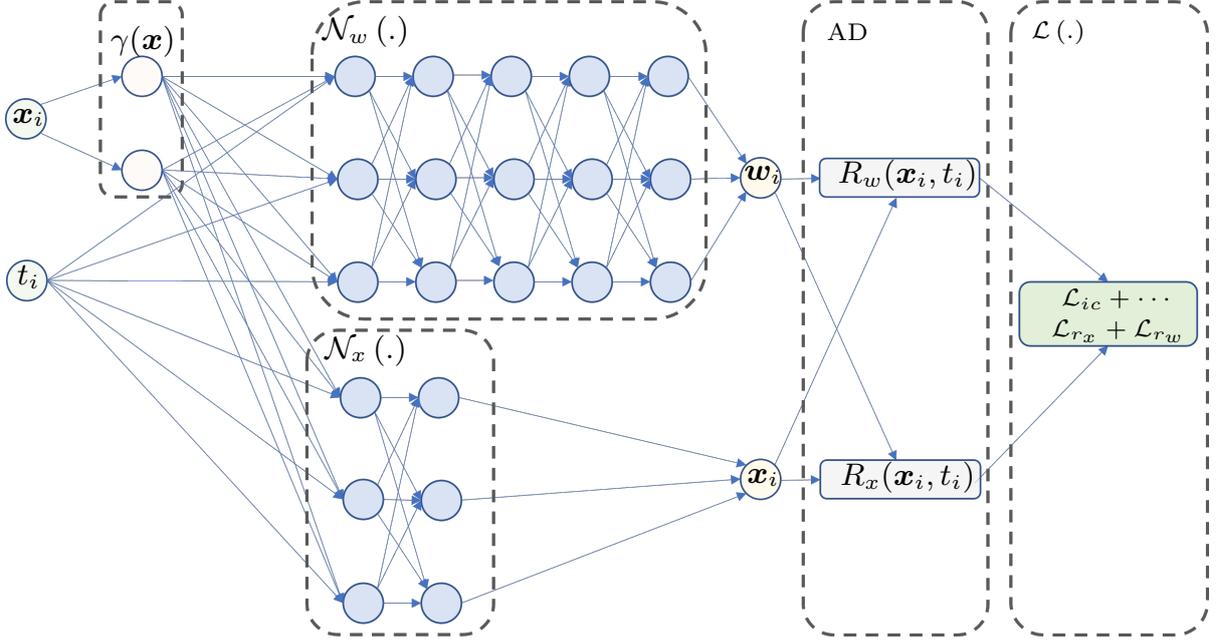
Figure 7: The proposed LPINNs architecture with periodic boundary condition.

and the second branch solves for the state parameter on the characteristics curves, and minimizes eq. (11b), i.e.,

$$w = \mathcal{N}_w(u) = \phi_m\left(W_m\phi_{m-1}\left(W_{m-1}\cdots\phi_1\left(W_1u + b_1\right)\cdots + b_{m-1}\right) + b_m\right), \qquad (17)$$

where all the parameters are defined similar to the network in section 3.1. The two branches can be of different width, and depth. The output of the network is the state parameter on the characteristics curves minimizing the loss, i.e.,

$$\mathcal{L} = \mathcal{L}_{r_x} + \mathcal{L}_{r_w} + \mathcal{L}_{ic}, \qquad (18)$$

where $\mathcal{L}_{r_x}$ and $\mathcal{L}_{r_x}$ are the residual associated with eq. (11), and $\mathcal{L}_{ic}$ is the loss associated with initial condition of both the state and grid. Finally, one can interpolate the states from the Lagrangian to the Eulerian frame of reference, w^*, if necessary. The proposed architecture is depicted in fig. 7.

We recognize the residual equations in eq. (11) can also be minimized in an architecture similar to that of the traditional PINNs. However, the proposed two–branch architecture leverages the inherent low–dimensionality of the characteristics [48], to build a shallow and efficient network to solve eq. (11a).

4. Experimental Results

In this section, the traditional PINN and the proposed LPINN are compared. The spatio–temporal grid is of size $(N_x, N_t) = \left(2^8, 100\right)$ for $(x, t) \in [0, 2\pi] \times [0, 1]$. In the convection and convection–diffusion cases, the initial condition is $w_0 = \sin(2\pi x)$, and the numerical solutions in [10] are considered as the truth. In the the Burgers' case, the initial condition is $w_0 = \sin(2\pi x) + c$, and solution of a Fourier pseudo–spectral solver is considered as the truth.

In all cases, the PINNs have 4 hidden-layers, and the LPINNs have an additional shallow branch with 2 hidden-layers. All activation functions are $\phi(.) = \tanh(.)$, except those of the last layer, where activation is linear (identity). The hyper–parameters in eq. (13) are $\lambda_{bc} = 0$, $\lambda_{ic} = 1000$, and $\lambda_r = 10$. Adam optimizer [51] with 10^5 iterations and learning rate of 0.01 are used for all the cases.

The error is defined as,

$$\text{Error} = \frac{\|w - w^*\|_2}{\|w\|_2}, \qquad (19)$$

where w is the truth, and w^* is the output of the network (interpolated) on the Eulerian grid, both after removing the boundary condition. A quadratic scheme is used to interpolate the output of LPINN to the Eulerian grid.

9
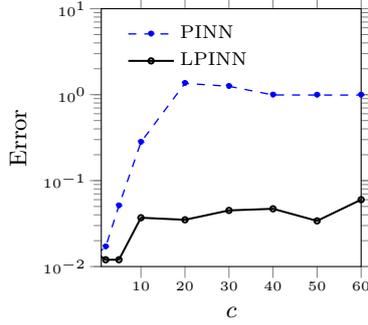
Figure 8: Comparison of the error in PINN (dashed blue) vs. the proposed LPINNs (solid back) for the convection equation eq. (20).

4.1. Convection

Consider the inviscid convection equation,

$$\frac{\partial w(x,t)}{\partial t} - c\frac{w(x,t)}{\partial x} = 0, \tag{20}$$

and its reformulation in the Lagrangian frame of reference,

$$\frac{dx}{dt} = c, \tag{21a}$$

$$\frac{\partial w}{\partial t} = 0. \tag{21b}$$

The solution to eq. (21) is straightforward. Equation (21b) dictates the grid points to move with the constant convection velocity, c, while the state variable remains constant along the moving points, eq. (21b). The accuracy of PINN and LPINN are compared for different convection velocity in fig. 8. Similar to [10], the error of PINN increases for larger values of c, such that for $c \geq 20$, the PINN cannot be trained. In case of the proposed LPINN, where the problem is simply reformulated on the Lagrangian frame of reference, the error for all cases remains below 5%. Note that the reported error is also comprised of the error originating from interpolating the predicted state from the moving grid of the Lagrangian frame to the stationary grid of the Eulerian frame.

To evaluate the optimality of the trained network, the loss landscape of the network at the end of the training phase is often used as a descriptive measure [6, 10, 13, 52]. To compute the loss landscape, the two dominant eigenvectors of the Hessian of the loss with respect to the trainable parameters of the networks, δ and η, are computed using the code provided in [53]. Subsequently, the network is perturbed along the eigenvectors and its loss, $\mathcal{L}'$, is evaluated, i.e.,

$$\mathcal{L}'(\alpha,\beta) = \mathcal{L}(\theta + \alpha\delta + \beta\eta), \tag{22}$$

where $(\alpha,\beta) \in [-\alpha_0, \alpha_0] \times [-\beta_0, \beta_0]$. Finally, $\log(\mathcal{L}'(\alpha,\beta))$ is visualized in fig. 9, for $c = \{0, 30, 50\}$ and for both PINN and LPINN architectures. In fig. 9a, we recover the saddle shape of the loss landscape for small convection speed as reported for PINN in [10]. Similarly, by increasing c, the landscape becomes less smooth (sharper, or more rugged), implying the trained network is not at a minimizer (figs. 9b to 9c). In the case of LPINN (figs. 9d to 9f), the loss landscapes are significantly smoother compared to their PINN counterparts (figs. 9a to 9c). Moreover, the landscape is smooth (flat), even at high c, increasing the confidence that the obtained minimizer is a global one.

4.2. Convection-diffusion

Consider the viscous convection–diffusion equation,

$$\frac{\partial w(x,t)}{\partial t} - c\frac{w(x,t)}{\partial x} = v\frac{\partial^2 w(x,t)}{\partial x^2}, \tag{23}$$

10
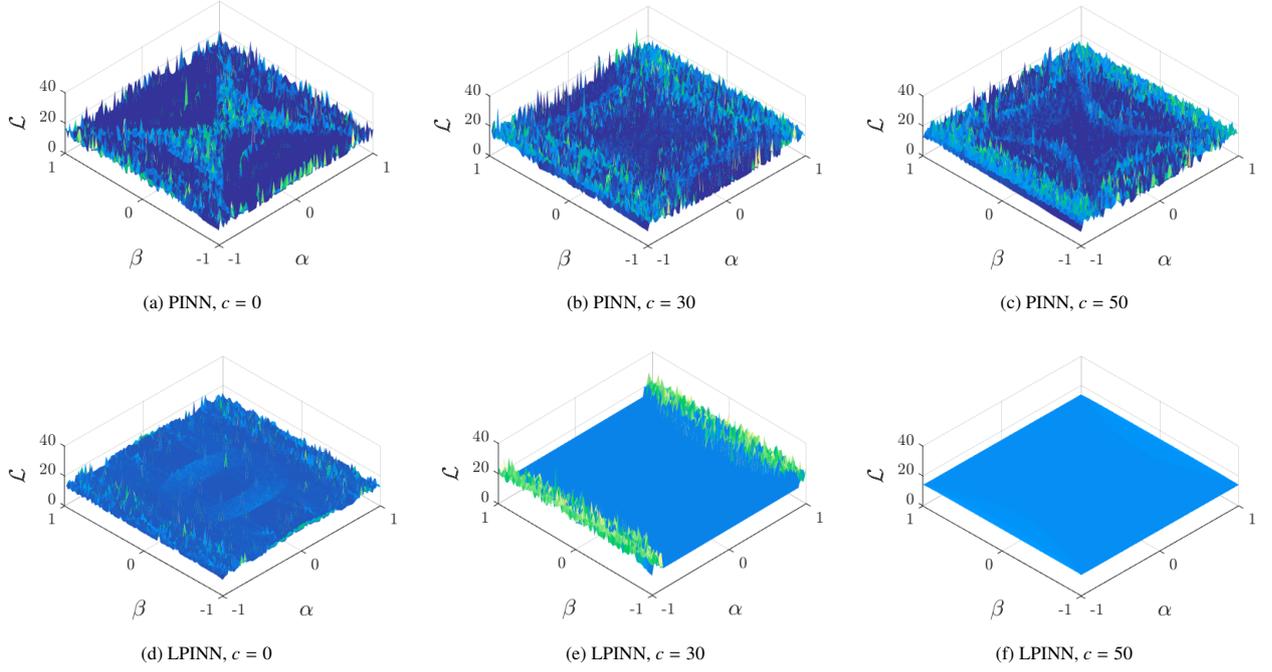
Figure 9: The (log of) loss landscape of convection equation given different convection speeds, $c \in \{0, 30, 50\}$. a-c. PINN, d-e. LPINN.
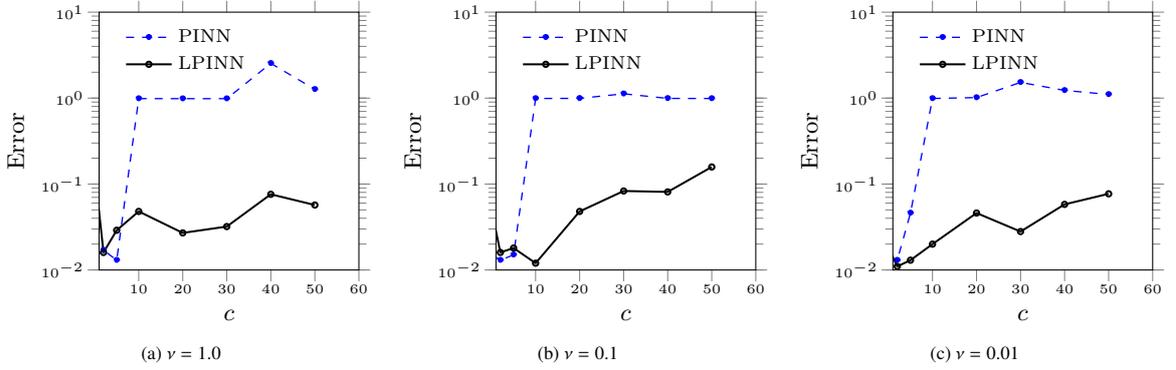
Figure 10: Comparison of the error in PINN (dashed blue) vs. the proposed LPINNs (solid back) for the convection–diffusion equation eq. (23) with different v.

and its reformulation in the Lagrangian frame of reference,

$$\frac{dx}{dt} = c, \tag{24a}$$

$$\frac{\partial w}{\partial t} = v \frac{\partial^2 w}{\partial x^2}. \tag{24b}$$

Figure 10 compares the accuracy of PINN and the proposed LPINNs. Similar to the inviscid case discussed in section 4.1, the error in PINNs increases by increasing c and they fail to train after a critical c. Similarly, the LPINNs increases by increasing c, however, the error remains around 10%, even in the most challenging cases.
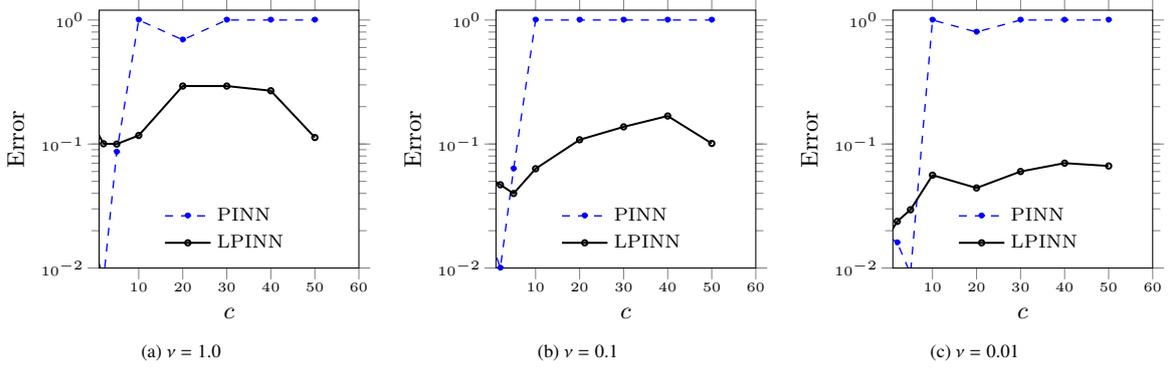
Figure 11: Comparison of the error in PINN (dashed blue) vs. the proposed LPINNs (solid back) for the viscous Burgers' equation eq. (25) with different ν.
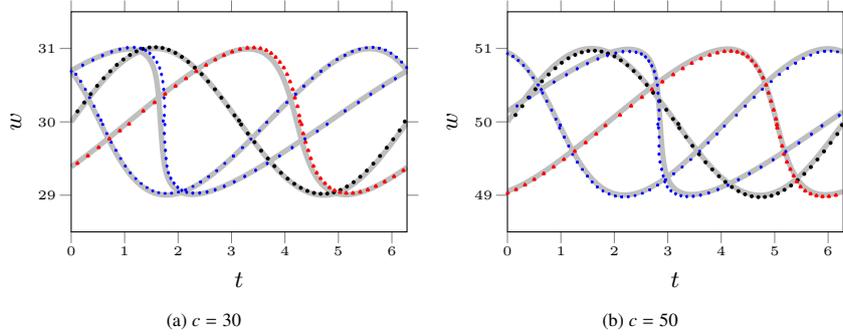
Figure 12: Comparison of the proposed LPINN with pseudo–spectral solver for the viscous Burgers' equation eq. (25) ($\nu = 0.01$) at $t \in \{0, T/3, 2T/3, T\}$ (black circle, blue square, red triangle, blue diamond) for $c = \{30, 50\}$. PINN cannot be trained in both regimes.

4.3. Burgers' equation

Consider the viscous Burgers' equation,

$$\frac{\partial w(x,t)}{\partial t} - w(x,t)\frac{w(x,t)}{\partial x} = \nu\frac{\partial^2 w(x,t)}{\partial x^2}, \tag{25}$$

and its representation on the Lagrangian frame,

$$\frac{dx}{dt} = w(x,t), \tag{26a}$$

$$\frac{\partial w}{\partial t} = \nu\frac{\partial^2 w}{\partial x^2}. \tag{26b}$$

While the traditional formulation of PINN is successfully demonstrated for Burgers' equation [1], the examined problem lack the main property of challenging cases for training, i.e., the large Kolmogorov n–width associated with the travel of the shock as in fig. 4. In fig. 11, and in a similar to the trend in convection–diffusion equation, the PINN fails to train for $c \geq 10$, while the LPINN is trained for all cases. The higher error in this case compared to the convection–diffusion equation is due to the higher interpolation error close to the shock. Note that in these cases the viscosity, ν, is small enough to form the high gradient shock and is large enough to avoid the intersecting characteristics. fig. 12 shows the accuracy of the proposed LPINN compared to the numerical solver at different simulation time steps.

5. Conclusions

The contributions of this paperis threefold.

1. We described the challenge of training of the traditional architecture of PINNs through the lens of approximation theory using Kolmogorov n–width, and the rate of decay of singular values of the solution. This realization explains many of the successful remedies in training of PINNs. More importantly, it can help with identifying unknown challenging problems, and opens of a wide variety of possibility to address them.

2. We identified Burgers' equation in the presence of traveling shocks as another challenging case. The complexity of the training is explained based on our discussion of the irreducibly of the solution on a linear space.

3. We propose LPINN for 1D linear and non–linear convection–diffusion equations. The reformulation of the equations on the characteristics automatically conforms to the direction of travel of the information in the domain, and satisfies the expected causality. More importantly, the solution on the manifold is of lower dimension, i.e., low Kolmogorov n–width, and is less sensitive to the system parameters. Using the inherent low–dimensionality of the characteristics [48], only a shallow branch is added to the traditional PINN to minimize the composite loss comprised of residual equations of characteristics, and the state variable on the characteristics. While it is suggested that the condition number of the loss is a probable source of complexity in training of PINNs [10], the proposed LPINN architecture is robust with respect to the condition number and show convergence regardless.

Future work should investigate the problem with shocks where the characteristics intersect. In such cases, a vanishing viscosity approach or removing the triple point value using Rankine–Hugoniot condition seem to be straightforward solutions [54]. Similar strategies are even applied to PINNs in presence [18, 19]. An extension of the proposed LPINN to higher spatial dimensions is possible using Radon transform [55]. In cases where characteristics curves are not real, e.g., wave equation, a manifold can be identified by registration–based or feature tracking approaches, e.g., [30, 32, 39, 56, 57], and to replace the mass–spring–damper grid models in arbitrary Lagrangian–Eulerian (ALE) formulation of PINNs [58]. Specifically, an optimal and low–rank manifold can be constructed offline by identifying an optimally morphing grid [32]. Moreover, LPINN architecture provides an opportunity for one–shot discovery of an optimal manifold.

Acknowledgements

References

[1] M. Raissi, P. Perdikaris, G. 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 (2019) 686–707.

[2] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, L. Yang, Physics–informed machine learning, Nature Reviews Physics 3 (2021) 422–440.

[3] L. Lu, R. Pestourie, W. Yao, Z. Wang, F. Verdugo, S. G. Johnson, Physics–informed neural networks with hard constraints for inverse design, SIAM Journal on Scientific Computing 43 (2021) B1105–B1132.

[4] A. Arzani, J.-X. Wang, R. M. D'Souza, Uncovering near-wall blood flow from sparse data with physics-informed neural networks, Physics of Fluids 33 (2021) 071905.

[5] S. Mowlavi, S. Nabi, Optimal control of PDEs using physics-informed neural networks, preprint arXiv:2111.09880 (2021).

[6] O. Fuks, 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 1 (2020) 19–37.

[7] Z. Mao, A. D. Jagtap, G. E. Karniadakis, Physics–informed neural networks for high-speed flows, Computer Methods in Applied Mechanics and Engineering 360 (2020) 112789.

[8] S. Wang, X. Yu, P. Perdikaris, When and why PINNs fail to train: A neural tangent kernel perspective, Journal of Computational Physics 449 (2022) 110768.

[9] S. Wang, Y. Teng, P. Perdikaris, Understanding and mitigating gradient flow pathologies in physics-informed neural networks, SIAM Journal on Scientific Computing 43 (2021) A3055–A3081.

[10] A. S. Krishnapriyan, A. Gholami, S. Zhe, R. M. Kirby, M. W. Mahoney, Characterizing possible failure modes in physics-informed neural networks, Advances in Neural Information Processing Systems (2021) 1–22.

[11] A. Bihlo, R. O. Popovych, Physics–informed neural networks for the shallow–water equations on the sphere, Journal of Computational Physics 456 (2021) 111024.

[12] S. Wang, S. Sankaran, P. Perdikaris, Respecting causality is all you need for training physics–informed neural networks, preprint arXiv:2203.07404 (2022).

[13] S. Basir, I. Senocak, Critical investigation of failure modes in physics-informed neural networks, in: AIAA SCITECH 2022 Forum, American Institute of Aeronautics and Astronautics, Reston, Virginia, 2022.

[14] Colby, L. Wight, , 10128, , C. L. Wight, Jia, Zhao, , 10129, , J. Zhao, Solving Allen-Cahn and Cahn-Hilliard equations using the adaptive physics informed neural networks, Communications in Computational Physics 29 (2021) 930–954.

[15] M. R. Hestenes, Multiplier and gradient methods, Journal of Optimization Theory and Applications 4 (1969) 303–320.

[16] R. G. Patel, I. Manickam, N. A. Trask, M. A. Wood, M. Lee, I. Tomas, E. C. Cyr, Thermodynamically consistent physics-informed neural networks for hyperbolic systems, Journal of Computational Physics 449 (2022) 110754.

[17] Y. Shin, J. Darbon, G. E. Karniadakis, On the convergence of physics informed neural networks for linear second-order elliptic and parabolic type PDEs, Communications in Computational Physics 28 (2020) 2042–2074.

[18] C. G. Fraces, H. Tchelepi, Physics informed deep learning for flow and transport in porous media, preprint arXiv:2104.02629 (2021).

[19] E. Abreu, J. B. Florindo, A study on a feedforward neural network to solve partial differential equations in hyperbolic-transport problems, in: M. Paszynski, D. Kranzlmüller, V. V. Krzhizhanovskaya, J. J. Dongarra, P. M. A. Sloot (Eds.), Computational Science – ICCS 2021, Springer International Publishing, Cham, 2021, pp. 398–411.

[20] B. Hillebrecht, B. Unger, Certified machine learning: A posteriori error estimation for physics–informed neural networks, preprint arXiv:2203.17055 (2022) 1–8.

[21] W. E, S. Wojtowytsch, Kolmogorov width decay and poor approximators in machine learning: shallow neural networks, random feature models and neural tangent kernels, Research in the Mathematical Sciences 8 (2021) 5.

[22] S. Wojtowytsch, W. E, Can shallow neural networks beat the curse of dimensionality? a mean field training perspective, IEEE Transactions on Artificial Intelligence 1 (2021) 121–129.

[23] A. Quarteroni, A. Manzoni, F. Negri, Reduced Basis Methods for Partial Differential Equations, volume 92 of *UNITEXT*, Springer International Publishing, 2016.

[24] A. Pinkus, n-Widths in Approximation Theory, Springer, Berlin, Heidelberg, 1985.

[25] J. M. Melenk, On n-widths for elliptic problems, Journal of Mathematical Analysis and Applications 247 (2000) 272–289.

[26] S. M. Djouadi, On the optimality of the proper orthogonal decomposition and balanced truncation, Proceedings of the IEEE Conference on Decision and Control (2008) 4221–4226.

[27] S. M. Djouadi, On the connection between balanced proper orthogonal decomposition, balanced truncation, and metric complexity theory for infinite dimensional systems, Proceedings of the 2010 American Control Conference, ACC 2010 (2010) 4911–4916.

[28] B. Unger, S. Gugercin, Kolmogorov n-widths for linear dynamical systems, Advances in Computational Mathematics (2019).

[29] J. A. Evans, Y. Bazilevs, I. Babuška, T. J. Hughes, n-Widths, sup–infs, and optimality ratios for the k-version of the isogeometric finite element method, Computer Methods in Applied Mechanics and Engineering 198 (2009) 1726–1741.

[30] M. A. Mirhoseini, M. J. Zahr, Model reduction of convection-dominated partial differential equations via optimization-based implicit feature tracking, preprint arXiv:2109.14694 (2021) 1–38.

[31] T. Taddei, A. T. Patera, A localization strategy for data assimilation; Application to state estimation and parameter estimation, SIAM Journal on Scientific Computing 40 (2018) B611–B636.

[32] R. Mojgani, M. Balajewicz, Low-rank registration based manifolds for convection-dominated PDEs, in: Proceedings of the AAAI Conference on Artificial Intelligence, volume 35, 2021, pp. 399–407.

[33] S. E. Ahmed, S. M. Rahman, O. San, A. Rasheed, I. M. Navon, Memory embedded non-intrusive reduced order modeling of non-ergodic flows, Physics of Fluids 31 (2019) 126602.

[34] S. Dutta, P. Rivera-Casillas, B. Styles, M. W. Farthing, Reduced order modeling using advection-aware autoencoders, Mathematical and Computational Applications 27 (2022).

[35] R. Mojgani, Reduced order modeling of convection-dominated flows, dimensionality reduction and stabilization, Ph.D. thesis, University of Illinois at Urbana-Champaign, Urbana, IL, USA, 2020.

[36] M. Nonino, F. Ballarin, G. Rozza, Y. Maday, Overcoming slowly decaying Kolmogorov n-width by transport maps: application to model order reduction of fluid dynamics and fluid–structure interaction problems, preprint arXiv:1911.06598 (2019).

[37] B. Peherstorfer, Model reduction for transport-dominated problems via online adaptive bases and adaptive sampling, SIAM Journal on Scientific Computing 42 (2020) A2803–A2836.

[38] D. Rim, B. Peherstorfer, K. T. Mandli, Manifold approximations via transported subspaces: Model reduction for transport-dominated problems, preprint arXiv:1912.13024 (2020) 1–30.

[39] T. Taddei, A registration method for model order reduction: Data compression and geometry reduction, SIAM Journal on Scientific Computing 42 (2020) A997–A1027.

[40] S. E. Ahmed, O. San, Breaking the Kolmogorov barrier in model reduction of fluid flows, Fluids 5 (2020).

[41] J. Barnett, C. Farhat, Quadratic approximation manifold for mitigating the Kolmogorov barrier in nonlinear projection-based model order reduction, arXiv preprint arXiv:2204.02462 (2022) 1–38.

[42] P. Krah, S. Büchholz, M. Häringer, J. Reiss, Front transport reduction for complex moving fronts, preprint arXiv:2202.08208 (2022) 1–26.

[43] J. Ren, W. R. Wolf, X. Mao, Model reduction of traveling-wave problems via radon cumulative distribution transform, Phys. Rev. Fluids 6 (2021) L082501.

[44] S. Venturi, T. Casey, SVD perspectives for augmenting DeepONet flexibility and interpretability, preprint arXiv:10.48550/arxiv.2204.12670 (2022) 1–56.

[45] K. Lee, K. T. Carlberg, Model reduction of dynamical systems on nonlinear manifolds using deep convolutional autoencoders, Journal of Computational Physics 404 (2020) 108973.

[46] Y. Kim, Y. Choi, D. Widemann, T. Zohdi, A fast and accurate physics–informed neural network reduced order model with shallow masked autoencoder, Journal of Computational Physics 451 (2022) 110841.

[47] A. D. Jagtap, G. E. Karniadakis, Extended physics-informed neural networks (XPINNs): A generalized space-time domain decomposition

based deep learning framework for nonlinear partial differential equations, Communications in Computational Physics 28 (2020) 2002–2041.

[48] R. Mojgani, M. Balajewicz, Lagrangian basis method for dimensionality reduction of convection dominated nonlinear flows, preprint arXiv:1701.04343 (2017).

[49] H. Lu, D. M. Tartakovsky, Lagrangian dynamic mode decomposition for construction of reduced-order models of advection-dominated phenomena, Journal of Computational Physics 407 (2020) 109229.

[50] S. Dong, N. Ni, A method for representing periodic functions and enforcing exactly periodic boundary conditions with deep neural networks, Journal of Computational Physics 435 (2021) 110242.

[51] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, preprint arXiv:1412.6980 (2014).

[52] F. M. Rohrhofer, S. Posch, C. Gößnitzer, B. C. Geiger, Understanding the difficulty of training physics-informed neural networks on dynamical systems, preprint arXiv:2203.13648 (2022) 1–12.

[53] Z. Yao, A. Gholami, K. Keutzer, M. W. Mahoney, PyHessian: Neural networks through the lens of the hessian, in: 2020 IEEE International Conference on Big Data (Big Data), 2020, pp. 581–590.

[54] R. J. LeVeque, Numerical methods for conservation laws, volume 214, Springer, 1992.

[55] D. Rim, Dimensional splitting of hyperbolic partial differential equations using the Radon transform, SIAM Journal on Scientific Computing 40 (2018) A4184–A4207.

[56] T. Taddei, L. Zhang, Space-time registration-based model reduction of parameterized one-dimensional hyperbolic PDEs, ESAIM: Mathematical Modelling and Numerical Analysis 55 (2021) 99–130.

[57] N. Sarna, P. Benner, Data-driven model order reduction for problems with parameter-dependent jump-discontinuities, Computer Methods in Applied Mechanics and Engineering 387 (2021) 114168.

[58] H. Wessels, C. Weißenfels, P. Wriggers, The neural particle method – An updated Lagrangian physics informed neural network for computational fluid dynamics, Computer Methods in Applied Mechanics and Engineering 368 (2020) 113127.