



HHS Public Access

Author manuscript

Mach Learn Sci Technol. Author manuscript; available in PMC 2024 March 01.

Published in final edited form as:

Mach Learn Sci Technol. 2023 March 01; 4(1): 015013. doi:10.1088/2632-2153/acb416.

Recipes for when Physics Fails: Recovering Robust Learning of Physics Informed Neural Networks

Chandrajit Bajaj¹, Luke McLennan¹, Timothy Andeen², Avik Roy^{3,†}

¹Department of Computer Science & Oden Institute for Computational Engineering and Sciences, The University of Texas at Austin, Austin, TX 78712

²Department of Physics, The University of Texas at Austin, Austin, TX 78712

³Center for AI Innovation, National Center for Supercomputing Applications, University of Illinois at Urbana Champaign, Urbana, IL 61801

Abstract

Physics-informed Neural Networks (PINNs) have been shown to be effective in solving partial differential equations by capturing the physics induced constraints as a part of the training loss function. This paper shows that a PINN can be sensitive to errors in training data and overfit itself in dynamically propagating these errors over the domain of the solution of the PDE. It also shows how physical regularizations based on continuity criteria and conservation laws fail to address this issue and rather introduce problems of their own causing the deep network to converge to a physics-obeying local minimum instead of the global minimum. We introduce Gaussian Process (GP) based smoothing that recovers the performance of a PINN and promises a robust architecture against noise/errors in measurements. Additionally, we illustrate an inexpensive method of quantifying the evolution of uncertainty based on the variance estimation of GPs on boundary data. Robust PINN performance is also shown to be achievable by choice of sparse sets of inducing points based on sparsely induced GPs. We demonstrate the performance of our proposed methods and compare the results from existing benchmark models in literature for time-dependent Schrödinger and Burgers' equations.

Keywords

Physics informed deep learning; Gaussian Processes; nonlinear PDE; Robust deep learning

1. Introduction

Neural Networks (NNs) are finding ubiquitous applications in fundamental sciences. Their abilities to perform classification and regression over large and complicated datasets are making them extremely useful for a variety of purposes, including modeling molecular dynamics modeling, nonlinear dynamical system design and control, and other many body interactions [1, 2, 3]. In many cases, these important problems in physics and engineering are posed in terms of static and time dependent partial differential equations

[†] corresponding author: avroy@illinois.edu.

(PDEs). Analytical solutions to PDEs are scarce. Moreover, PDEs are difficult to solve and require computationally intensive, and highly pre-conditioned numerical linear solvers [4]. Popular discretization methods, Finite Difference Method (FDM), and even Finite Element Method (FEM) are used to obtain point wise or piece-wise linear estimates over a fine grid or meshed domains of interest [5]. Although NN-based approximations to differential equations have been explored for some time [6, 7, 8], interest in such approaches have been reinvigorated recently due to significant improvement in computational platforms that support fast forward and backward gradient propagation utilizing automatic differentiation [9]. Rapid progress has been seen in NN-assisted solutions of Ordinary Differential Equations (ODEs) [10]. Novel architectures like NeuralODEs [11] have been proposed to harness the power of blackbox ODE solvers in conjunction with continuous-depth residual Neural Networks (Resnets) using the method of adjoints [12], while models like ODE2VAEs use variational auto-encoder architectures to learn functions and derivatives via latent space embeddings similar to Cauchy boundary conditions [13].

Solving PDEs using NNs has also seen significant attention by the development of Physics Informed Neural Networks (PINNs) [14]. Trained PINNs have been shown to be effective in solving time dependent partial differential equations for a given set of Cauchy boundary conditions over a finite spatio-temporal domain. They exploit a deep neural network's ability as universal function approximators [15, 16]. Many variants of the trainable PINN architecture have been explored to exploit structure, quality and speed of convergence, and dimensional scalability [17, 18, 19, 20, 21, 22, 23]. However, one of the least explored areas is the robustness of trainable PINNs for various noisy data scenarios and conditions. While most PINN architectures in literature assume perfectly known boundary data, in many practical applications, this data comes from regulated and calibrated measurement processes and is subject to uncertainties or errors pertaining to the limitations of the measurement system or the stochastic nature of the physical processes themselves. Most PINN architectures utilize a finite and small collection of training data on the domain boundary. Given the typical small size of this training data and NNs ability to capture arbitrary non-linearity, vanilla PINNs can often propagate these errors or uncertainties in an unstable fashion. Such unregulated error propagation can significantly limit the applicability of PINNs as industrial strength numerical approximators of PDEs.

In this paper, we extensively investigate the problem of error propagation in PINNs. In Section 2 we introspect the architecture of a PINN, and how it responds when data on the initial timeslice is corrupted with noise. The notion of PINN robustness is thus tied to the PINNs ability to preserve solution features under such noisy perturbations. We analyze examples of learning time-varying non-linear Schrödinger, and the non-convective fluid flow Burgers' equations. We further show the impact of introducing regularization based on continuity criteria inspired from conservative PINN (cPINN) architectures [19, 22] through attempts to satisfy conservation laws. In Section 3 we provide details of a Gaussian Process smoothed PINN (GP-smoothed PINN), and its sparse variant and compare its performance against most typical and popular PINN architectures. We demonstrate how these twin GP-smoothed PINNs recover the intended solution and can outperform methods realizing continuity or conservation regularizers [24] as well as better possess the ability to control uncertainty propagation compared to uncertainty quantification methods proposed

for example in [20]. We briefly examine the efficiency role of sparse inducing points and also demonstrate the importance and choice of various kernels for achieving best model selection in Gaussian process training to prevent data driven under- or overfitting.

2. Robustness of PINNs

In this section, we review the PINN and cPINN architectures and explain with different examples how such models fail to capture the essence of robustness. We also investigate multiple physics-inspired regularization schemes and identify their limitations in addressing the issue of robustness for usual PINN architectures.

2.1. Review of PINN Architectures

A partial differential equation that determines spatio-temporal evolution of a set of scalar (real or complex) fields, collectively represented by $u(\vec{x})$, can be expressed as

$$\mathcal{N}[u(\vec{x}), f(\vec{x})] = 0 \quad (1)$$

where \vec{x} defines a n dimensional spatial or spatio-temporal coordinate system, defined on the domain $\vec{x} \in \mathcal{D} \subset \mathbb{R}^n$. \mathcal{N} represents a set of known, finite-order, differential operators and $f(\vec{x})$ is the source function, usually known as analytical expression in problems intended to solve a forward PDE problem. Eq. 1 is subject to a set of boundary conditions,

$$\mathcal{B}[u(\vec{x} \in \partial D)] = 0 \quad (2)$$

The general idea of a PINN [14] is to obtain an approximation of the field $u(\vec{x}) \approx \tilde{u}(\vec{x})$ by a deep neural network that can solve the system of equations 1, subject to 2.

$$\tilde{u}(\vec{x}) = \text{NN}_\theta(\vec{x}; \mathcal{U}_B, \mathcal{U}_C, \mathcal{U}_D) \quad (3)$$

where,

- $\mathcal{U}_B = \left\{ \left(\vec{x}_i^b, \mathcal{B}[u(\vec{x}_i^b)] \right)_{i=1}^{N_b} \right\}$ represents a set of samples on the domain boundary $\partial \mathcal{D}$,
- $\mathcal{U}_C = \left\{ \left(\vec{x}_i^c, f(\vec{x}_i^c) \right)_{i=1}^{N_c} \right\}$ is a set of measurements that enforces the PDE physics of Eq. 1 on the neural net (Eq. 3), and
- $\mathcal{U}_D = \left\{ \left(\vec{x}_i^d, u(\vec{x}_i^d) \right)_{i=1}^{N_d} \right\}$ is a set of direct measurements. Although, \mathcal{U}_D is not necessary for training a plain PINN as in Ref. [14], they are often used for training networks for targeted simulation.

The parameters of the deep network (θ) in Eq. 27 are obtained by minimization of the loss function,

$$\theta^* = \operatorname{argmin}_{\theta} \mathcal{L}_{PINN} \quad (4)$$

where the loss function can be decomposed as,

$$\mathcal{L}_{PINN} = \alpha_{BC} \mathcal{L}_{BC} + \alpha_{PDE} \mathcal{L}_{PDE} + \alpha_D \mathcal{L}_D \quad (5)$$

Here, \mathcal{L}_{BC} is the MSE loss calculated over \mathcal{U}_B , enforcing the NN to approximate the boundary condition,

$$\mathcal{L}_{BC} = \frac{1}{N_b} \sum_i \left| \mathcal{B} \left[\tilde{u}(\vec{x}_i^b) \right] \right|^2 \quad (6)$$

\mathcal{L}_{PDE} is the MSE loss calculated over \mathcal{U}_C , enforcing the physics on the NN,

$$\mathcal{L}_{PDE} = \frac{1}{N_c} \sum_i \left| \mathcal{N} \left[\tilde{u}(\vec{x}_i^c), f(\vec{x}_i^c) \right] \right|^2 \quad (7)$$

and finally, \mathcal{L}_D , if employed, determines the loss with respect to observation,

$$\mathcal{L}_D = \frac{1}{N_d} \sum_i \left| \tilde{u}(\vec{x}_i^d) - u(\vec{x}_i^d) \right|^2 \quad (8)$$

The α_0 parameters in Eq. 5 are penalty parameters that determine the relative strength of the regularizing terms in the loss function. Authors in [14] assign $\alpha_0 = 1$ identically, but alternate choices have been explored in other works [24, 25].

2.1.1. Conservative PINNs (cPINNs)—Variations of PINN architectures have been explored in a number of recent works. For example, authors in [21] explore a multi-staged PDE solver for a long range solution and Ref. [19] introduces adaptive, hyperparameterized activation functions that would accelerate the convergence of such networks.

Although the training of PINNs doesn't strictly require a discretized grid of evaluation points as often required by traditional numerical techniques like finite difference method and finite element methods, they can certainly benefit from such grid structures by parallelizing the training of PINNs over a collection of subdomains. Hence, the domain of integration and its boundary are divided into subdomains, i.e.

$$\mathcal{D} = \bigcup_{i=1}^K d_i, \quad \partial \mathcal{D} = \bigcup_{i=1}^K \partial d_i, \quad \tilde{u} = \bigcup_{i=1}^K \tilde{u}_i \quad (9)$$

An additional benefit to such parallelized structure is to include flux continuity at subdomain boundaries in the loss function, providing additional safeguard against error propagation. In other words, the loss function in Eqn. 5 is modified as

$$\mathcal{L}_{cPINN} = \sum_{j=1}^d \mathcal{L}_{PINN}^j + \alpha_I \mathcal{L}_I \quad (10)$$

where \mathcal{L}_{PINN}^j is the PINN loss for the j -th subdomain as defined in Eq. 5. The other term, \mathcal{L}_I^j acts as a regularizer that enforces functional and flux continuity at the interface of j -th subdomain interface.

$$\mathcal{L}_I^j = \frac{1}{N_{Ij}} \sum_{i=1}^{N_{Ij}} \left(\left| \tilde{u}_j(\vec{x}_i) - \tilde{u}_{j+1}(\vec{x}_i) \right|^2 + \left| \nabla \tilde{u}_j(\vec{x}_i) \cdot \mathbf{n}_i^j - \nabla \tilde{u}_{j+1}(\vec{x}_i) \cdot \mathbf{n}_i^{j+1} \right|^2 \right) \quad (11)$$

here, $\{\vec{x}_i\}$ is a collection of points on the interface of j -th and $j+1$ -th subdomains and \mathbf{n}_i^j is the unit vector normal to the interface of j -th subdomain at the location of \vec{x}_i .

Figure 1 shows a schematic representation of the different spatio-temporal regions that contribute to evaluating the loss function in Eqn. 5 along with a generic architecture for such models.

2.2. Error Propagation Through PINNs

An NN can approximate non-linear functions with increasing degrees of accuracy. It is typically expected in the case of a PINN that $N_b \ll N_c$ i.e. the size of the training data on the domain boundary is usually much smaller than the size of the physicsenforcing collocation points. This unavoidable feature of PINNs make them susceptible to overfitting on the boundary and eventually propagate those errors across the domain. In the following subsections, we investigate the physical nature of error propagation through PINNs and the impact of different regularizations on training the PINN architecture. We will use two popular examples that have been widely used in the literature to illustrate this issue of error propagation.

2.2.1. Nonlinear Schrödinger equation—We consider the example considered in Ref. [14] of a nonlinear Schrödinger partial differential equation which describes the spatio-temporal evolution of a 1D complex field $h(x, t) = u(x, t) + iv(x, t)$ as

$$i \frac{\partial h}{\partial t} + \frac{1}{2} \frac{\partial^2 h}{\partial x^2} + |h|^2 h = 0 \quad (12)$$

which can also be interpreted as a set of coupled partial differential equations given as

$$\begin{aligned} -\frac{\partial v}{\partial t} + \frac{1}{2} \frac{\partial^2 u}{\partial x^2} + (u^2 + v^2)u &= 0 \\ \frac{\partial u}{\partial t} + \frac{1}{2} \frac{\partial^2 v}{\partial x^2} + (u^2 + v^2)v &= 0 \end{aligned} \quad (13)$$

The domain boundary is defined as $(x, t) \in [-5, 5] \times [0, \frac{\pi}{2}]$. The boundary conditions can be classified as

- i. Initial condition; the known value of the field on the initial time slice given as a collection of measurements $\{(x_j, h(x_j, 0))_{j=0}^{N_{b,t}}\}$. Unbeknownst to the NN, these measurements are taken from the analytical solutions with possible sources of additive corruption-

$$\begin{aligned} u(x_j, 0) &= 2\text{sech}(x_j) + \Theta_u \epsilon_i^u \\ v(x_j, 0) &= \Theta_v \epsilon_i^v \end{aligned} \quad (14)$$

where ϵ^u, ϵ^v are randomly chosen from a normal distribution with $\mathcal{N}(0, \sigma^2)$. The parameters Θ_u and Θ_v represent acceptance of errors which are set to 0 (1) for error-free (error-inclusive) initial conditions.

- ii. Periodic boundary condition on spatial slices enforced on a discretized spatial boundary. A total of $N_{b,s}$ points are chosen to enforce the following spatial boundary conditions

$$\begin{aligned} h(+5, t) &= h(-5, t) \\ \frac{\partial h}{\partial x}(+5, t) &= \frac{\partial h}{\partial x}(-5, t) \end{aligned} \quad (15)$$

The loss function is constructed according to Eqn. 5 with $\alpha_0 = 1.0$. To set the benchmark for the performance of a PINN for this problem, we solved for the PDE with error-free boundary data ($\Theta_u = \Theta_v = 0$). We train a fully connected MLP with 6 hidden layers and 70 nodes per layer, with two inputs corresponding to the space and time coordinate and two outputs corresponding to the real and imaginary parts of the complex field. The MLP was trained for $N_b = 100$ points on the domain boundary, 50 of which were taken from uniformly sampling the space coordinate (x) on the initial timeslice to impose the initial condition in Eq. 14, and the 50 points were taken on from a uniform distribution on the time coordinate to impose the periodic boundary condition in Eq. 15. A fine grid of $N_c = 20000$ collocation points was chosen to impose the physics of the PDE.

Figure 2 shows the evolution of the complex field in the Schrödinger equation as evaluated by a vanilla PINN at four different timeslices, taken at $t = 0, 0.39, 0.78, 1.37$. The performance of the PINN for error-free data on the initial timeslice is shown in Figures 2(a)-(d). We use the same architecture to repeat the exercise while training on corrupted data on initial timeslice by letting $\Theta_u = \Theta_v = 1$ and additive errors generated by drawing samples from zero-mean Gaussian distribution with $\sigma = 0.1$. The performance of the PINN in evaluating the complex field magnitude at the same time instances is shown in Figures 2(e)-(h). The effect of introducing corrupted data on initial timeslice becomes evident when comparing Figures 2(a)-(d) with Figures 2(e)-(h). The PINN tends to overfit on the initial timeslice and eventually propagates these errors on the following timeslices.

2.2.2. Burgers' Equation—We consider the Burgers' equation in one spatial dimension with Dirichlet boundary conditions as a second example. Widely used in fluid dynamics and nonlinear acoustics, this nonlinear PDE has been widely studied as a benchmark example in the PINN literature [14, 20]. The PDE and the boundary conditions for 1D Burgers' equation are given as :

$$\frac{\partial u}{\partial t} + u \frac{\partial u}{\partial x} = \nu \frac{\partial^2 u}{\partial x^2} \quad (16)$$

$$u(-1, t) = u(1, t) = 0 \quad (17)$$

$$u(x, 0) = -\sin(\pi x) + \Theta_u e^u \quad (18)$$

where the domain boundary is given as $(x, t) \in [-1, 1] \times [0, 1]$. To set the benchmark for this problem, we solve Eqn. 16 subject to boundary conditions in Eqns. 17 and 18 with $\nu = \frac{0.01}{\pi}$ using a PINN without noise ($\Theta_u = 0$) in the initial data. We used an MLP with 4 hidden layers, each with 40 nodes. We trained with $N_c = 10000$ collocation points to enforce the physics (Eqn. 16) and $N_b = 150$ points on the boundary, 50 points for enforcing the initial condition in Eqn 18 and 50 points on each of the spatial boundaries at $x = -1, 1$ to enforce each of the Dirichlet conditions in Eqn. 17. The loss function is constructed according to Eq. 5 with $\alpha_0 = 1.0$. Figure 3(a)-(d) shows the evolution of the field $u(x, t)$ in the Burgers' equation as evaluated by a vanilla PINN at four different timeslices, taken at $t = 0, 0.25, 0.50, 1.0$.

Next, we repeat the exercise for Burgers' equation by introducing additive corruption with $\Theta_u = 1$ and $\sigma = 0.5$. Figures 3(e)-(h) show the corresponding line shapes of the Burgers' field for the aforementioned timestamps. Similar to what was found for the Schrödinger Equation, the vanilla PINN architecture fails to auto-correct for the corruption in initial data and ends up overfitting on the initial timeslice and eventually propagates these errors to later timeslices.

2.3. Regularization of PINNs

Both Figures 2 and 3 show a PINN's inherent inability to self-correct when trained with error-corrupted data on the initial timeslice. The PINN rather learns to overfit on the initial timeslice and eventually propagates the initial errors to later timeslices. This propagation of error to later timeslices is a direct consequence of having no regularization in the loss term in Eqn. 5 to constrain overfitting at the domain boundary.

In principle, this is not very different from overfitting in classic regression problems with a high degree polynomial or introduction of bias in an un-regularized regression by outliers. This naively indicates that additional regularization might be useful to limit the propagation of errors. However, our investigations indicates that some of the most physically intuitive choices for regularizing constraints have little impact on error propagation and boundary overfitting.

We consider two unique choices of regularizers. First, inspired by the cPINN architecture, we impose the constraint of functional and flux continuity at arbitrary spatio-temporal boundaries to constrain the PDE solution. Second, we explore the possibility of using physical conservation laws as additional sources of regularization. In the following subsections, we explore the impact of using such regularization schemes in training PINNs with corrupted boundary data.

2.3.1. Functional and flux continuity at subdomain interfaces—The cPINN architecture inspires a useful regularization that imposes continuity of the field and its flux across domain boundaries. From a physics standpoint, these regularizations can be thought of as additional conservation laws that ensures continuity and differentiability of a field across subdomains. In this subsection, we explore the impact of including this term in the training loss function in Eqn. 10 in controlling propagation of uncorrelated errors at sampling points on the initial timeslice.

One immediate concern is that the convergence of the cPINN is susceptible to exact choices of how many subdomains are chosen and where those boundaries are located. Based on the location of the subdomain boundaries, the cPINN's capacity to converge to the global minimum of the training loss function, can be significantly impacted. To illustrate this, we compare the performance of cPINNs in solving the Schrödinger equation with two and three equal spatial subdomains trained on error-free boundary data. The results are shown in Figure 4. Evidently, a three subdomain cPINN better recovers the analytical solution. However, the failure of a two subdomain cPINN, to capture the solution of a PDE even for error-free boundary data is intriguing. This observation yields a deeper insight to the impact of adding additional regularizers with the PINN loss function in Eqn. 5. As can be seen in Figure 4(a), the two subdomain cPINN moderately deviates from the analytical solution on the initial timeslice at the expense of converging to the local minima introduced by the inclusion of the interface loss. Figure 5 shows how the PINNs trained on different subdomains converge to identical functional and flux values at the subdomain boundary but eventually experiences large deviations from the analytical solution which requires $\frac{\partial u}{\partial x}(x, 0) = \frac{\partial v}{\partial x}(x, 0) = 0$. It is apparent that choice of subdomain boundary at $x = 0$ plays an important role in causing such deviating solutions. As the real and imaginary fields reach local extrema at $x = 0$ for all values of t , even small deviations in estimating the local gradients at the interface can create a cascading effect in the evaluation of the complex field across subdomain boundaries as we can see in Figures 5.

The result of such instabilities in convergence, is evidential consequence of such cPINN-inspired regularizations and the architecture's failure to be robust to such noisy perturbations. As also shown in Figure 6, regularization of functional and flux continuity at subdomain boundaries does not provide the necessary safeguard to ensure robustness.

When we repeat the same set of exercises for the Burgers' equation, we see very similar results, as shown in Figures 7. Similar to what we have seen for the Schrödinger equation, placing the subdomain boundaries at functionally critical points can destabilize and deteriorate the quality of the solution learned by the deep network. It can be seen

in Figures 7(a)-(d) from the deviation of the Burgers' field's predicted behavior at later timeslices even without any error introduced on the initial timeslice when two subdomains are considered with an interface at $x = 0$. However, the function is almost identically recovered when the data is perfect on the initial timeslice with three subdomains as shown in Figures 7(e)-(h). This tendency of a PINN-like architecture to converge to a local minima instead of the global minima almost infallibly deteriorates the quality of convergence when error is introduced on the initial timeslice, and thus the solution departs significantly from its ideal behavior. This is apparent in both the two subdomain (Figures 7(i)-(l)) and the three subdomain cPINNs (Figures 7(m)-(p)).

2.3.2. Conservation law constraints—Physical laws are often subject to a number of conservation laws. While in many cases these conservation laws emerge as direct consequences of the mathematical structure of the PDE, explicitly enforcing such conservation laws will back-propagate additional constraining gradients for the NN hyperparameters. We can explicitly include these conservation laws into the loss function.

For example, one of the major consequences of non-linear Schrödinger equation is global conservation of the squared absolute value of the complex Schrödinger field $|h(x, t)|^2$

$$\int |h(x, t)|^2 dx = \int (u(x, t)^2 + v(x, t)^2) dx = C \quad (19)$$

where C is a constant. A number of other conserved quantities follow for the 1D nonlinear Schrödinger equation we are considering [26, 27, 28], which include:

$$\int \left(u \frac{\partial v}{\partial x} + v \frac{\partial u}{\partial x} \right) dx \quad (20)$$

$$\int \left(\left| \frac{\partial h}{\partial x} \right|^2 - |h|^4 \right) dx \quad (21)$$

We can constrain the solution by explicitly including these conservation laws as regularizers in the loss function. For instance, the probability conservation law in Eqn. 19, along with the requirement of probability confinement within the spatially bounded region for the domain of Eqn. 12 requires that

$$0 = \frac{d}{dt} \int_a^b |h|^2 dx = \frac{d}{dt} \int_a^b (u^2 + v^2) dx = 2 \int_a^b (uu_t + vv_t) dx \quad (22)$$

where (a, b) is the spatial domain. We can approximate this integral using time-sliced collocation points:

$$\int_a^b (uu_t + vv_t) dx \approx \frac{b-a}{n_{c,t}} \sum_{i=1}^{n_{c,t}} (u(x_i, t)u_t(x_i, t) + v(x_i, t)v_t(x_i, t)) \quad (23)$$

where $n_{c,t}$ represents the number of collocation points chosen over the spatial subdomain at timeslice t . We can define the conservation loss to be

$$\mathcal{L}_c = \frac{1}{N_t} \sum_{t=0}^{N_t} \frac{1}{n_{c,t}} \sum_{j=0}^{n_{c,t}} (u(x_j, t)u(x_j, t) + v(x_j, t)v(x_j, t))^2 \quad (24)$$

We train a PINN with this conservative constraint on data appended with the loss function in Eqn. 5. As seen in Figure 8, the constrained PINN better represents the conservation laws, with the overall range for dispersion of cumulative probability effectively reduced with the inclusion of conservative constraints in the PINN loss function. However, the observed evolution of the field in Figure 9, obtained by applying this constraint does not result in superior accuracy. Errors are still propagated from the initial timeslice throughout the spatio-temporal domain.

As a second example of the implication of conservation laws as regularizers, we take the example of Cole-Hopf transformation [29, 30] for the Burgers' equation. Based on proper mathematical wisdom developing analytical solutions to PDEs, methods like the Cole-Hopf transform have been found useful to convert one family of PDEs into another whose analytical solution is known and rather simple to compute. The Cole-Hopf transformation converts the nonlinear Burgers' equation to a linear Heat equation. This transformation is defined by making the following change of variables:

$$u(x, t) = -2\nu \frac{\partial v}{\partial x}(x, t) \quad (25)$$

The transformed field $v(x, t)$ satisfies the heat equation $\partial_t v = \nu \partial_{xx} v$. The viscous case of the Burgers' equation is for $\nu > 0$ causing a non-linear dissipative shock for small values of ν . The inviscid case yields the equation having a non-linear hyperbolic conservation law. For our purpose, the conservation law for the modified field can be viewed as a regularizer to the conventional PINN loss function in Eqn. 5.

$$\mathcal{L}_{CH} = \sum_{i,j} (v_i(x_i, t_j) - \nu v_{xx}(x_i, t_j))^2 \quad (26)$$

In addition to using the Cole-Hopf loss term as a regularizer for a vanilla PINN architecture, we additionally consider including the continuity criteria for equally split two and three subdomains. Figure 10(a)-(d) shows the time evolution of the Burgers' field $u(x, t)$ when the PINN is trained with the loss function including the Cole-Hopf term in Eqn. 26. The functional approximation obtained from the PINN is much smoother compared to what we have observed in our previous exploration of regularized evaluation of the solution to the Burgers' equation (Figures 3 and 7). It can be directly traced back to the fact that the Cole-Hopf transformed field is indeed an anti-derivative of the Burgers' field and in the neural architecture when implemented as a discrete integration acts as a smoothing operation in somewhat canceling out the effect of the error. The smoothing operation however eventually

leads to underfitting and the evolution of the field at later timeslices is affected in a similar fashion.

Figure 10(e)-(h) [(i)-(l)] shows the results when we trained the Cole-Hopf regularized PINN with functional and flux continuity imposed at the interface of two [three] sub-domains. These results establish the inadequacy of the Cole-Hopf regularizer in establishing the robustness of the neural architecture and reinforces their tendency to converge to local minimum instead of reaching the intended global minimum.

3. Gaussian Process (GP) based Error Correction for PINNs

As the previous section illustrates, physics-inspired regularization alone does not eliminate propagation of errors in PINNs. In fact, application of such constraints forces the PINN to converge to a local minimum that satisfies the physics of conservation laws for overfitted boundary conditions and eventually propagates the overfitting across the spatio-temporal domain. In this section, we seek to explore alternate solution to this problem by using smoothing techniques that safeguard the quality of fit by cross-validated regulation of smoothed boundary data. The model of corruption considered in the PDE model for 1D Schrödinger and Burgers' equations is a ubiquitous approximation for many physics processes. In such processes, the spatial evolution of a local field is often expected to be smooth. When the physical data on domain boundary is subject to such errors, it is often convenient to model these measurements as a realization of a stochastic process. For instance, the initial condition in Eqn. 14 can be modeled by a pair of continuous stochastic processes, U_x, V_x where the index representing the spatial coordinate of the PDE domain. The mean and covariance for such processes are given as

$$\begin{aligned} \mathbb{E}(U_i) &= 2\text{sech}(x = x_i) \\ \mathbb{E}(V_i) &= 0 \\ \text{Cov}(U_i, U_j) &= \text{Cov}(V_i, V_j) = \sigma^2 \delta_{ij} \end{aligned}$$

To obtain a functional estimate of these stochastic processes, Gaussian Process Regression [31] is a powerful, nonparametric method. Given the set of samples on the initial timeslice, \mathcal{U}_B , the DNN structure in Eqn. 27 is replaced by,

$$\tilde{u}(\vec{x}) = \text{NN}_\theta(\vec{x}; \hat{\mathcal{U}}_B, \mathcal{U}_C, \mathcal{U}_D) \quad (27)$$

where

$$\hat{\mathcal{U}}_B = \left\{ \left(\vec{x}_i^b, \mathcal{B} \left[\hat{u}(\vec{x}_i^b) \right] \right)_{i=1}^{N_b} \right\} \quad (28)$$

and

$$\hat{u}(\vec{x}_i^b) = \text{GP} \left(\vec{x}_i^b \mid \left\{ \left(\vec{x}_i^b, u(\vec{x}_i^b) \right)_{i=1}^{N_b} \right\} \right) \quad (29)$$

represents the GP-predicted estimate of the boundary data. The choice of the kernel function, representing the pairwise covariance of observations is given as a sum of RBF and white noise kernels,

$$k(x_i, x_j) = A \exp\left(-\frac{|\vec{x}_i - \vec{x}_j|^2}{2l^2}\right) + \sigma^2 \delta_{ij} \quad (30)$$

where A, l, σ are hyperparameters obtained by maximizing the log-marginal likelihood.

Smoothing techniques are commonly applied in problems where robustness is a desired quality. Compared to other parametric smoothing techniques like fixed order polynomials or smoothing splines, GP regression has often been proved to be more robust against underfitting and overfitting [32, 33]. Robustness guarantees for GPs have been extensively explored in literature [34, 35]. GPs have also been explored in connection with physics-inspired kernel building [36, 37] and found to be effective in predicting physical phenomena like phase transitions in quantum systems [38].

Using a GP-smoothing on the boundary data allows for the PDE solver to regain its performance by training itself over the smoothed data on initial timeslice. Compared to other approaches [39, 40, 41] that employ Gaussian Processes to solve differential equations, our method uniquely harnesses the smoothing interpolating functionality of a GP while exploiting the universal approximator feature of a neural network. While Gaussian Processes with a proper choice of a kernel can be very useful in approximating smooth analytical solutions, their $O(n^3)$ complexity makes them infeasible for optimizing such solutions over a large set of collocation points and such complexity grows significantly with high dimensional problems. However, restricting their use on the domain boundary reduces the complexity by an order of magnitude while almost identically recovering the analytical solution. Figure 11 shows the performance of a GP-smoothed PINN in solving the Schrödinger equation, where the DNN can recover the analytical form despite corruption in initial data.

Since the loss function in Eq. 5 is not a direct metric of validating the performance of the the PINN, the validation loss is measured in terms of *mean squared error (MSE)* loss compared with respect to the analytical solution. In Fig. 12, we compare the evolution of the loss function during training and the validation MSE loss. It can be seen that GP-smoothed PINN performs almost as well as error-free PINN, and significantly better than a PINN trained with corrupted boundary data but no smoothing.

3.1. Kernel Selection for Gaussian Processes

The choice of kernel for fitting a GP to the boundary data is very important- improper choices can lead to underfitting or overfitting and eventually propagate large errors through the PINN architecture. In order to make the best choice for a kernel, we explored a k -fold cross-validation technique on the initial time-slice data for the 1D Schrödinger Equation. The dataset on initial boundary is split into k equal subsets where $k - 1$ of them are used for training and and one subset is kept aside for validation. Training and validation data are

used to optimize the GP hyperparameters. We examine the performance of the optimized Gaussian Processes for using the RBF kernel, the Matèrn kernel [31] with $\nu = 0.1, 1.5$, and 4.0, and the Rational Quadratic (RQ) kernel [31]. Each kernel is appended with a localized white noise kernel. The average training and validation MSE losses, measured with respect to a fixed set of noise-corrupted sampling points on the initial timeslice (Eqn. 14), with $k = 10$ for different choices of kernels are summarized in Table 1. It can be seen that RBF and RQ kernels have similar performance while the Matèrn kernels tend to overfit.

3.2. Evolution of Measurement Uncertainty

While a generic PINN fails to recover the physics-motivated evolution of noisy boundary data following a PDE, a GP-smoothed PINN not only can recover the physical evolution but also provide a controlled estimate of uncertainty at every point in the spatio-temporal domain. The uncertainty evaluated by a GP-smoothed PINN is obtained by evaluating the deviation in the NN parameters for $\pm 1\sigma$ variation of the training data on domain boundary

$$\tilde{u}(\vec{x}) \pm \delta\tilde{u}(\vec{x}) = \text{NN}_{\theta \pm \delta\theta}(\vec{x}; \hat{\mathcal{U}}_B^\pm, \mathcal{U}_C, \mathcal{U}_D) \quad (31)$$

where

$$\hat{\mathcal{U}}_B^\pm = \left\{ \left(\vec{x}_i^b, \mathcal{B} \left[\hat{u}(\vec{x}_i^b) \pm \delta\hat{u}(\vec{x}_i^b) \right] \right)_{i=1}^{N_b} \right\} \quad (32)$$

The uncertainty associated with the boundary data, $\delta\hat{u}$ is obtained from the covariance estimate of the optimized GP. The deviation of the NN parameters, $\delta\theta$ can be obtained from minimizing the loss function evaluated with $\hat{\mathcal{U}}_B^\pm$.

$$\left(\theta \pm \delta\theta \right)^* = \underset{\theta}{\text{argmin}} \mathcal{L}_{\text{PINN}}(\theta; \hat{\mathcal{U}}_B^\pm) \quad (33)$$

Analytical estimate of $\delta\theta^*$ is a computationally intractable task since it requires inversion of the very large Hessian matrix $\frac{\partial^2 \mathcal{L}}{\partial \theta^2}$. However, a rather inexpensive technique is to start with a PINN architecture with parameters θ already optimized for the mean value of the boundary data $\hat{\mathcal{U}}_B$ and re-train the network with the modified boundary data. This reoptimization converges more quickly and provides an estimate of evolution of uncertainty at all points of the space-time domain. The evolution of uncertainty for a GP-smoothed PINN evaluated solution of Eqn. 12 is shown in Figure 11. The network was re-trained for an additional 1000 iterations to optimize for the uncertainty bands. In general, the number of additional required to converge for estimating the uncertainty bands depends on the size of the corrupting error, which can be quantitatively estimated from the optimized value of the σ parameter in Eqn. 29.

3.3. Sparse GP (SGP) based Error Correction

GP-based smoothing can provide robustness for PINNs as shown in the previous section. However, optimizing a GP is an expensive process with a complexity of $o(n^3)$, with n being the number of points considered to optimize the GP. Even though we are restricting the GPs to be optimized only over the domain boundary, this can still be a major bottleneck for our method for high dimensional problems. As the dimension of domain boundary $\partial D \subset \mathbb{R}^{d-1}$ increases, it will require more and more points on the boundary to satisfy the boundary condition. Sparse Gaussian Processes have been extensively studied in literature to significantly reduce the complexity for high dimensional problems. A multitude of variants of sparsity inducing GPs have found their applications in the context of sample efficient reinforcement learning [42], deep kriging with big data [43], and variational learning of GPs [44]. We consider a hybrid approach for sparsity inducing smoothing GP on the domain boundary following the algorithm suggested in Ref. [45] to obtain inexpensive selection of inducing points (IPs). Originally designed for Sparse Variational GPs (SVGPs), this algorithm is effective in the context of our problem of selecting a smaller subset of IPs on the domain boundary.

The SGP algorithm we use is explained in algorithm. The sparsity optimizations for GP is done in two steps. In the first step, a small number of data points (n_0) are randomly taken to optimize the GP hyperparameters. In the second step, additional IPs are chosen from the data based on the kernel distance between the new IP candidate (z) and the already selected set of IPs (X_0). The new IP candidate is included in X_0 if the Algorithm 1 Selection of IPs for SGP kernel distance between z and all existing IPs is smaller than some predefined threshold (ρ). To reduce the complexity of this approach, iterative re-optimization of the kernel hyperparameters is avoided and only after the desired set of IPs have been chosen, the the GP hyperparameters are finally reoptimized to smooth the corrupted dataset on the domain boundary. The total number of IPs chosen is bounded by $M \leq N_{b,t}$. Figure 13 shows how sparse GPs can be almost equally useful in recovering the Schrödinger field dynamics. When the number of IPs is set too low, e.g. only 10 IPs for both $u(x, 0)$ and $v(x, 0)$, the recovery of physical dynamics is not as satisfactory. However, with a somewhat larger set of IPs including 29 and 20 IPs for $u(x, 0)$ and $v(x, 0)$ respectively, the PINN's performance improves significantly and becomes comparable to that of the full GP-smoothed PINN shown in Figure 11.

Algorithm 1 Selection of IPs for SGP

```

procedure IPSelect ( $n_0, M, X, y, \rho$ )
  Randomly Select  $X_0, y_0$  From  $X, y$  with  $|X_0| = |y_0| = n_0$ 
   $k^* = \operatorname{argmax}_k \log p(y_0 | \operatorname{GP}(k, X_0, y_0))$ 
  for  $z \in X - X_0$  do
    if  $\max\{k^*(z, x_0) | x_0 \in X_0\} < \rho$  then
       $X_0 \leftarrow X_0 \cup \{z\}$ 

```

```

if  $|X_0| = M$  then
    break
return  $X_0$ 

```

Table 2 summarizes the validation MSE obtained with different models and compare them with the benchmark model of a vanilla PINN with no errors. While the performance of a PINN significantly deteriorates with the introduction of even modest errors with $\sigma = 0.1$, both GP-smoothed PINN and SVGP-smoothed PINN perform similar to the benchmark model.

We demonstrate the effectiveness of GP and SGP smoothing in recovering the physical field dynamics for 1D Burgers' equation in Figures 14(a)–(d) and Figures 14(e)–(h). The SGP employed 41 IPs on the initial timeslice and shows remarkable performance recovery. We also compare the results from the UQ-PINN architecture proposed in Ref. [20] and both GP and SGP smoothed PINNs perform noticeably better than the solution obtained from the UQ-PINN architecture. Table 3 summarizes the validation MSE loss obtained from different PINN architectures and it can be seen that both GP and SGP smoothed recover a similar level of accuracy as observed by the error-free PINN.

4. Additional Examples

To demonstrate the effectiveness of GP and SGP smoothing for higher dimensional PDEs, we consider a couple of 2D PDEs in this section.

4.1. 2D Heat Equation

The 2D heat equation and the corresponding spatio-temporal boundary conditions are given as:

$$\frac{\partial u}{\partial t} = \frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2} \quad (34)$$

$$u(x, y, 0) = 3\sin(\pi x)\sin(\pi y) + \sin(3\pi x)\sin(\pi y) + \Theta_\epsilon \epsilon^u \quad (35)$$

$$u(0, y, t) = u(1, y, t) = u(x, 0, t) = u(x, 1, t) = 0 \quad (36)$$

where the domain boundary is given as $(x, y, t) \in [0,1] \times [0,1] \times [0,0.1]$ and Θ_ϵ is the acceptance function for the noise term in the initial condition. The analytical solution to this equation is given as $u(x, y, t) = 3\sin(\pi x)\sin(\pi y)e^{-2\pi^2 t^2} + \sin(3\pi x)\sin(\pi y)e^{-10\pi^2 t^2}$. An MLP with four hidden layers, each with 256 nodes, has been used. The physics is enforced with $N_c = 50000$ collocation points. 64 points are chosen on each of the four spatial boundaries and 1024 points on the initial timeslice for the initial condition, giving a total of $N_b = 1280$ points on the spatio-temporal boundary. Like the previous examples, the loss function is constructed

according to Eq. 5 with $\alpha_0 = 1.0$. For the SGP process, the IPs are chosen from the pool of 1024 points on the initial time slice according to Algorithm 1 with the number of IPs bounded by $M = 768$. The models are trained for 20000 epochs with ADAM optimizer with a learning rate of 10^{-3} .

As shown in Figure 15. GP-smoothing recovers the performance of the error-free PINN and SGP also considerably brings down the MSE when compared to that of the PINN trained with noisy data without any smoothing applied. We can see the initial condition each PINN architecture is trained with along with the point-wise error estimate in the PINN's solution for different models. The smoothing effect on the initial timeslice can be seen in Figure 16, where we can see that while the noisy initial condition almost completely obliterates the distributive feature of $u(x, t = 0)$, smoothing with GP or SGP allows its significant recovery. This translates into better convergence to actual solution for the latter couple of models on both initial and latter timeslices (Figure 17).

4.2. D Burgers' Equation

The two dimensional Burgers' equation is given by the following pair of PDEs-

$$\begin{aligned} \frac{\partial u}{\partial t} + u \frac{\partial u}{\partial x} + v \frac{\partial u}{\partial y} &= \nu \left(\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2} \right) \\ \frac{\partial v}{\partial t} + u \frac{\partial v}{\partial x} + v \frac{\partial v}{\partial y} &= \nu \left(\frac{\partial^2 v}{\partial x^2} + \frac{\partial^2 v}{\partial y^2} \right) \end{aligned} \quad (37)$$

where we consider $\nu = \frac{0.01}{\pi}$ and train the network to learn the following analytical solution [46, 47]-

$$u(x, y, t) = \frac{3}{4} - \frac{1}{4 \left(1 + \exp\left(\frac{-t - 4x + 4y}{32\nu}\right) \right)} \quad (38)$$

$$v(x, y, t) = \frac{3}{4} + \frac{1}{4 \left(1 + \exp\left(\frac{-t - 4x + 4y}{32\nu}\right) \right)} \quad (39)$$

The domain boundary is chosen as $(x, y, t) \in [0,1] \times [0,1] \times [0,1]$. The network is trained with the initial condition sampled from the functions $u(x, y, 0) + \Theta_u e^u$ and $v(x, y, 0) + \Theta_v e^v$ respectively for u and v where Θ_u and Θ_v are the acceptance functions for the noise terms in the initial condition. The spatial boundary conditions are obtained from plugging in the boundary coordinates in the analytical solution given in Equations 38 and 39. An MLP with four hidden layers, each with 256 nodes, has been used to simultaneously predict the two fields. The physics is enforced with $N_c = 50000$ collocation points. We choose 64 points on each of the four spatial boundaries and 1024 points on the initial timeslice to enforce the spatio-temporal boundary condition with a total $N_b = 1280$ measurements. The choice of loss function and optimizer follows the example of the previous examples. A

pool of 1024 uniformly sampled points on the initial timeslice are used for SGP and IPs are chosen according to Algorithm 1 with the number of IPs bounded by $M = 768$. The resulting performances of the four models after training for 20000 epochs are shown in Figure 18. As we can see from the accompanying table in Figure 18, the noise-free PINN replicates the analytical solution almost perfectly and when error is introduced, similar level of performance cannot be retrieved. However, the MSE is noticeably reduced by the smoothing performed by GP and SGP.

One of the noticeable aspects of the 2D Burgers' equations is the presence of a shockwave front at $4y - 4x = t$ around where both fields experience rather sharp, yet continuous gradients. We show the exact initial condition used to train the noise-free PINN in Figures 19a and 19i. When corrupted with noise and left unsmoothed, the shockwave feature is almost completely lost and as can be seen from Figure 20 (third column from the left), the network struggles to retrieve the shockwave front for latter timeslices as well. On the other hand, while GP and SGP to some extent recovers the initial field distributions, the gradients near the shockwave front are oversmoothed. It should be noted that this oversmoothing is not due to some limitation of GP or SGP itself, but rather the choice of samples used to train these processes. Being agnostic to the physical distribution, the enforcing points on the initial timeslice are uniformly sampled, which led to an under-representation of sharp gradients near the shockwave front. As a consequence to the missing perception of sharpness around the shockwave front, the GP-PINN and SGP-PINN accumulate their largest deviations in the latter timeslices around the shockwave front, as seen in Figure 20 (fourth and fifth columns from the left). The opposite signs of the errors on the two sides of the shockwave front represent that the PINN-reconstructed solution after GP or SGP smoothing has a more smoothly shifting wavefront.

5. Conclusion

As it often happens, measurements associated physical processes are subject to errors. When these measurements are used to learn the evolution of a system respecting some underlying physics dictated by a PDE using NNs, these errors can significantly distort the predicted behavior via nonlinear propagation of errors. In this paper, we explored the behavior of a PINN when it is trained with noise-corrupted datasets. Our work shows that deep PDE-solvers can be subject to overfitting and dynamically propagating errors observed on the domain boundaries even when physics-inspired regularizers are introduced to constrain the solution. To circumvent this issue, we proposed GP-smoothed deep network that can help recover the system's behavior over a finite space-time domain while providing a controlled prediction and bounded uncertainty. We further showed that the computational complexity of fitting a Gaussian Process can be significantly reduced by incorporating sparsely choosing inducing points for sparse GPs. This opens up opportunities to explore uncertainty propagation in predictive estimation using cPINNs or cPINN-like architectures as well as learning an optimal policy of selecting sparsely chosen inducing points.

Acknowledgement

The research for C.B. and L.M was supported in part by NIH-R01GM117594, by the Peter O'Donnell Foundation, and in part from a grant from the Army Research Office accomplished under Cooperative Agreement Number W911NF-19-2-0333. The work of L.M. was additionally supported by the the Moncreif Summer Undergraduate Internship Program. The work of T.A. and A.R. is supported by U.S. Department of Energy, Office of High Energy Physics under Grant No. DE-SC0007890. Studies performed by A.R. utilize resources supported by the National Science Foundation's Major Research Instrumentation program, grant 1725729, as well as the University of Illinois at Urbana-Champaign

References

- [1]. Frank Noé, Alexandre Tkatchenko, Klaus-Robert Müller, and Cecilia Clementi. Machine learning for molecular simulation. *Annual Review of Physical Chemistry*, 71(1):361–390, 2020.
- [2]. Puscasu Gheorghe, Codres Bogdan, Stancu Alexandru, and Murariu Gabriel. Nonlinear system identification based on internal recurrent neural networks. *International Journal of Neural Systems*, 19(02):115–125, 2009. [PubMed: 19496207]
- [3]. Carleo Giuseppe, Nomura Yusuke, and Imada Masatoshi. Constructing exact representations of quantum many-body systems with deep neural networks. *Nature Communications*, 9(5322), 2018.
- [4]. Benzi Michele. Preconditioning techniques for large linear systems: a survey. *Journal of computational Physics*, 182(2):418–477, 2002.
- [5]. Grossmann Christian, Roos Hans-Görg, and Stynes Martin. *Numerical treatment of partial differential equations*, volume 154. Springer, 2007.
- [6]. Dissanayake MWMG and Nhan Phan-Thien. Neural-network-based approximations for solving partial differential equations. *communications in Numerical Methods in Engineering*, 10(3):195–201, 1994.
- [7]. Aarts Lucie P and Van Der Veer Peter. Neural network method for solving partial differential equations. *Neural Processing Letters*, 14(3):261–271, 2001.
- [8]. Hayati Mohsen and Karami Behnam. Feedforward neural network for solving partial differential equations. *Journal of Applied Sciences*, 7(19):2812–2817, 2007.
- [9]. Baydin Atilim Gunes, Pearlmutter Barak A, Radul Alexey Andreyevich, and Siskind Jeffrey Mark. Automatic differentiation in machine learning: a survey. *Journal of machine learning research*, 18,2018.
- [10]. Lagaris IE, Likas A, and Fotiadis DI. Artificial neural networks for solving ordinary and partial differential equations. *IEEE Transactions on Neural Networks*, 9(5):987–1000, 1998. [PubMed: 18255782]
- [11]. Chen Ricky T. Q., Rubanova Yulia, Bettencourt Jesse, and Duvenaud David. Neural ordinary differential equations. *Advances in Neural Information Processing Systems*, 2018.
- [12]. Semenovitch Pontryagin Lev, Mishchenko EF, Boltyanskii VG, and Gamkrelidze RV. *Mathematical theory of optimal processes*. Gordon and Breach Science Publishers, 1987.
- [13]. Yıldız Ça atay, Heinonen Markus, and Lähdesmäki Harri. Ode2vae: Deep generative second order odes with bayesian neural networks. *arXiv preprint - arXiv:1905.10994*, 2019.
- [14]. Raissi Maziar, Perdikaris Paris, and Karniadakis George E. 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.
- [15]. Hornik Kurt, Stinchcombe Maxwell, White Halbert, et al. Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366, 1989.
- [16]. Pinkus Allan. Approximation theory of the mlp model in neural networks. *Acta numerica*, 8:143–195, 1999.
- [17]. Pang Guofei, Lu Lu, and Karniadakis George Em. fpinns: Fractional physics-informed neural networks. *SIAM Journal on Scientific Computing*, 41(4):A2603–A2626, 2019.
- [18]. Yang Liu, Meng Xuhui, and Karniadakis George Em. B-pinns: Bayesian physics-informed neural networks for forward and inverse pde problems with noisy data. *arXiv preprint arXiv:2003.06097*, 2020.

- [19]. Ameya D Jagtap Kenji Kawaguchi, and Karniadakis George Em. Adaptive activation functions accelerate convergence in deep and physics-informed neural networks. *Journal of Computational Physics*, 404:109136, 2020.
- [20]. Yang Yibo and Perdikaris Paris. Adversarial uncertainty quantification in physics-informed neural networks. *Journal of Computational Physics*, 394:136–152, 2019.
- [21]. Meng Xuhui, Li Zhen, Zhang Dongkun, and Karniadakis George Em. Ppinn: Parareal physics-informed neural network for time-dependent pdes. *Computer Methods in Applied Mechanics and Engineering*, 370:113250, 2020.
- [22]. Ameya D Jagtap Ehsan Kharazmi, and Karniadakis George Em. Conservative physics-informed neural networks on discrete domains for conservation laws: Applications to forward and inverse problems. *Computer Methods in Applied Mechanics and Engineering*, 365:113028, 2020.
- [23]. Shin Yeonjong. On the convergence of physics informed neural networks for linear second-order elliptic and parabolic type pdes. *Communications in Computational Physics*, 28(5):2042–2074, Jun 2020.
- [24]. Ameya D Jagtap Ehsan Kharazmi, and Karniadakis George Em. Conservative physics-informed neural networks on discrete domains for conservation laws: Applications to forward and inverse problems. *Computer Methods in Applied Mechanics and Engineering*, 365:113028, 2020.
- [25]. Kharazmi Ehsan, Zhang Zhongqiang, and Karniadakis George Em. Variational physics-informed neural networks for solving partial differential equations. *arXiv preprint arXiv:1912.00873*, 2019.
- [26]. Fujiwara Kazumasa and Miyazaki Hayato. The derivation of conservation laws for nonlinear schrodinger equations with a power type nonlinearity. *arXiv preprint - arXiv:1407.5282*, 2014.
- [27]. Watanabe Shinsuke, Miyakawa Makoto, and Yajima Nobuo. Method of conservation laws for solving nonlinear schrodinger equation. *Journal of the Physical Society of Japan*, 46:1653–1659, 1978.
- [28]. Barrett Johnner. “the local conservation laws of the nonlinear schrodinger equation”. Oregon State University, 2013.
- [29]. Cole Julian D. On a quasi-linear parabolic equation occurring in aerodynamics. *Quarterly of applied mathematics*, 9(3):225–236, 1951.
- [30]. Hopf Eberhard. The partial differential equation $u_t + uu_x = \mu u_{xx}$. *Communications on Pure and Applied mathematics*, 3(3):201–230, 1950.
- [31]. Williams Christopher K and Rasmussen Carl Edward. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA, 2006.
- [32]. Schulz Eric, Speekenbrink Maarten, and Krause Andreas. A tutorial on gaussian process regression: Modelling, exploring, and exploiting functions. *Journal of Mathematical Psychology*, 85:1–16, 2018.
- [33]. Gramacy Robert B. *Surrogates: Gaussian process modeling, design, and optimization for the applied sciences*. Chapman and Hall/CRC, 2020.
- [34]. Cardelli Luca, Kwiatkowska Marta, Laurenti Luca, and Patane Andrea. Robustness guarantees for bayesian inference with gaussian processes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 7759–7768, 2019.
- [35]. Blaas Arno, Patane Andrea, Laurenti Luca, Cardelli Luca, Kwiatkowska Marta, and Roberts Stephen. Adversarial robustness guarantees for classification with gaussian processes. In *International Conference on Artificial Intelligence and Statistics*, pages 3372–3382. PMLR, 2020.
- [36]. David K Duvenaud Hannes Nickisch, and Rasmussen Carl. Additive gaussian processes. *Advances in neural information processing systems*, 24, 2011.
- [37]. Duvenaud David, Lloyd James, Grosse Roger, Tenenbaum Joshua, and Zoubin Ghahramani. Structure discovery in nonparametric regression through compositional kernel search. In *International Conference on Machine Learning*, pages 1166–1174. PMLR, 2013.
- [38]. Vargas-Hernández Rodrigo A, Sous John, Berciu Mona, and Krems Roman V. Extrapolating quantum observables with machine learning: Inferring multiple phase transitions from properties of a single phase. *Physical review letters*, 121(25):255702, 2018. [PubMed: 30608785]
- [39]. Graepel Thore. Solving noisy linear operator equations by gaussian processes: Application to ordinary and partial differential equations. In *ICML*, volume 3, pages 234–241, 2003.

- [40]. Raissi Maziar, Perdikaris Paris, and Karniadakis George Em. Numerical gaussian processes for time-dependent and nonlinear partial differential equations. *SIAM Journal on Scientific Computing*, 40(1):A172–A198, 2018.
- [41]. Chen Yifan, Hosseini Bamdad, Owhadi Houman, and Stuart Andrew M. Solving and learning nonlinear pdes with gaussian processes. arXiv preprint arXiv:2103.12959, 2021.
- [42]. Grande Robert, Walsh Thomas, and How Jonathan. Sample efficient reinforcement learning with gaussian processes. *Proceedings of the 31st International Conference on Machine Learning*, 2014.
- [43]. Gadd Charles, Heinonen Markus, Lähdesmäki Harri, and Kaski Samuel. Sample-efficient reinforcement learning using deep gaussian processes. arXiv preprint arXiv:2011.01226, 2020.
- [44]. Tran Dustin, Ranganath Rajesh, and Blei David M. The variational gaussian process. In *4th International Conference on Learning Representations, ICLR 2016*, 2016.
- [45]. Galy-Fajou Theo and Opper Manfred. Adaptive inducing points selection for gaussian processes. *Continual Learning Workshop of ICML*, 2020.
- [46]. Fletcher Clive AJ. Generating exact solutions of the two-dimensional burgers' equations. *International Journal for Numerical Methods in Fluids*, 3:213–216, 1983.
- [47]. Zhu Hongqing, Shu Huazhong, and Ding Meiyu. Numerical solutions of two-dimensional burgers' equations by discrete adomian decomposition method. *Computers & Mathematics with Applications*, 60(3):840–848, 2010.




Figure 1.

A visualization of the domains of (a) a PINN and (b) a cPINN with boundary points (\cdot), collocation points (\times), and interface points (\blacksquare) for a spatio-temporal domain with one spatial dimension. Figure (c) shows the model architecture diagram of a PINN for a generic $n + 1$ dimensional spatio-temporal domain solving for a set of coupled fields u_1, \dots, u_m .




Figure 2.

The PINN-evaluated solution for $|h(x, t)|$ from the Schrödinger equation at different timeslices, $t = 0, 0.39, 0.78, 1.37$, from left to right for error-free boundary data (top) and corrupted boundary data with (bottom). The additive errors on the boundary data are taken independently from samples of zero mean Gaussian distribution with $\sigma = 0.1$. The points marked with the blue cross (x) pointer in the leftmost set of plots indicate the samples on the initial timeslice used to train the PINN.




Figure 3.

The PINN-evaluated solution for $u(x, t)$ from the Burgers' equation at different timeslices, $t = 0, 0.25, 0.50, 1.00$, from left to right for error-free boundary data (top) and corrupted boundary data with (bottom). The additive errors on the boundary data are taken from independently samples of zero mean Gaussian distribution with $\sigma = 0.5$. The points marked with the blue cross (x) pointer in the leftmost set of plots indicate the samples on the initial timeslice used to train the PINN.




Figure 4.

The cPINN-evaluated $|h(x, t)|$ for two (top row) and three (bottom row) equal subdomains at different timeslices, $t = 0, 0.39, 0.78, 1.37$, from left to right when no error is introduced on initial time-slice. The points marked with the blue cross (x) pointer in the leftmost set of plots indicate the samples on the initial timeslice used to train the cPINN.




Figure 5.

The cPINN evaluated lineshape for $|h(x, t)|$ (left), $\frac{\partial u}{\partial x}$ (middle), and $\frac{\partial v}{\partial x}$ (right) at the subdomain interface ($x = 0$) for a two subdomain cPINN as a function of t . PINNO (PINN1) refers to the PINN trained to solve the PDE on a spatial boundary of $-5 \leq x \leq 0$ ($0 \leq x \leq 5$).




Figure 6.

The cPINN-evaluated $|h(x, t)|$ for two (top row) and three (bottom row) subdomains at different timeslices, $t = 0, 0.39, 0.78, 1.37$, from left to right when additive Gaussian errors with $\sigma = 0.1$ is introduced on initial time-slice. The points marked with the blue cross (x) pointer in the leftmost set of plots indicate the samples on the initial timeslice used to train the cPINN.






Figure 7. cPINN-evaluated solution to Burgers' equation with (a-d) two subdomains with no error, (e-h) three subdomains with no error, (i-l) two subdomains with error, and (m-p) three subdomains with error. The error on individual datapoints, whenever applied, have been taken from additive zero-mean Gaussian distribution with $\sigma = 0.5$.

**Figure 8.**

Plots of three of the conservation laws (Eqns. 19– 21) of the nonlinear Schrödinger equation for (a-c) a vanilla PINN and (d-f) a PINN constrained by the first conservation law. The PINNs were trained with initial data corrupted with $\sigma = 0.1$ Gaussian noise. The constrained PINN has better performance of conservation laws 1 and 3, while both PINNs satisfy conservation law 2 nearly exactly.

**Figure 9.**

The PINN-evaluated $|h(x, t)|$ at different timeslices, $t = 0, 0.39, 0.78, 1.37$, from left to right.

In this case the PINN is constrained by the Schrodinger equations first conservation law:

$$\frac{d}{dt} \int |h|^2 dx = 0.$$

The training data on the initial timeslice is subject to measurement errors, modeled by a Gaussian random variable with zero mean and a standard deviation of $\sigma = 0.1$.

The points marked with the blue cross (x) pointer in the leftmost set of plots indicate the samples on the initial timeslice used to train the PINN.




Figure 10. The PINN-evaluated $u(x, t)$ at different timeslices, $t = 0, 0.25, 0.5, 1.0$, from left to right when additive Gaussian errors with $\sigma = 0.5$ is introduced on initial time-slice for when additional regularization in terms of Cole-Hopf transformation is added to the loss function. A single domain PINN used for the top row while the middle and bottom rows show result from 2 and 3 domain cPINNs respectively.




Figure 11.

The GP-smoothed PINN-evaluated $|h(x, t)|$ at different timeslices, $t = 0, 0.39, 0.78, 1.37$, from left to right. In this case the training data on the initial timeslice is subject to measurement errors, modeled by a Gaussian random variable with zero mean and a standard deviation of 0.1. GP-based smoothing was used on initial timeslice before training the PINN. The points marked with the blue cross (x) pointer in the leftmost set of plots indicate the samples on the initial timeslice used to fit the GP. The grey band in the subsequent plots represents the uncertainty associated with the PINN-evaluated approximation of $|h(x, t)|$.




Figure 12. The training loss function from Eqn. 5(left) and the MSE validation loss from as a function of number of iterations for a PINN solving.




Figure 13.

The SGP-smoothed PINN-evaluated $|h(x, t)|$ at different timeslices, $t = 0, 0.39, 0.79, 1.37$, from left to right. In this case the training data on the initial timeslice is subject to measurement errors, modeled by a Gaussian random variable with zero mean and a standard deviation of 0.1. SGP based smoothing was used on initial timeslice before training the PINN using 10 inducing points for each $u(x, 0)$ and $v(x, 0)$ for the top row and 29 inducing points for $u(x, 0)$ and 20 for $v(x, 0)$ for the bottom row. The points marked with the blue cross (x) pointer in the leftmost set of plots indicate the samples on the initial timeslice used to train the PINN. The grey band in the subsequent plots represents the uncertainty associated with the PINN-evaluated approximation of $|h(x, t)|$.




Figure 14.

The PINN-evaluated solution to Burgers’ equation when measurements on initial timeslice is corrupted with zero-mean, $\sigma = 0.5$ Gaussian noise for GP-smoothed PINN (top), SGP-smoothed PINN (middle), and UQ-PINN 20] (bottom), at timeslices $t = 0, 0.25, 0.51, 1$, from left to right. For the GP- and SGP-smoothed PINNs uncertainty bounds are calculated by retraining the PINN using the initial condition of the GP/SGP mean function plus or minus one standard deviation.




Figure 15.

The time evolution of the MSE loss for different models used in solving the 2D heat equation. Except the vanilla PINN (no error), all models were trained with data sampled from the initial timeslice corrupted with additive Gaussian noise with zero mean and $\sigma = 1.0$. MSE error evaluated over 50k points chosen over the entire spatio-temporal domain for the different models is given in the accompanying table.




Figure 16.

The initial condition at $t = 0$ for 2D heat equation for (a) error-free PINN, (b) noisy PINN without smoothing, (c) noisy PINN with GP smoothing, (d) noisy PINN with SGP smoothing. Except the vanilla PINN (no error), all models were trained with data sampled from the initial timeslice corrupted with additive Gaussian noise with zero mean and $\sigma = 1.0$. The figures in the bottom row show the error in PINN-evaluated solution at the initial timeslice for the corresponding architecture in the top row.





Figure 17.

The leftmost column shows the exact solution for 2D heat equation at $t = 0.0368$ (top row) and $t = 0.0632$ (bottom row). The remaining columns show the point-wise error in PINN-evaluated solution for noise-free PINN (second column), noisy PINN without smoothing (third column), GP-PINN (fourth column), and SGP-PINN (final i.e. fifth column). Except the vanilla PINN (no error), all models were trained with data sampled from the initial timeslice corrupted with additive Gaussian noise with zero mean and $\sigma = 1.0$.



| Model | MSE |
|--------------------------------------|---------|
| PINN (no error) | 5.16e-7 |
| PINN ($\sigma = 0.5$) | 4.4e-2 |
| GP-PINN ($\sigma = 0.5$, 1024 IPs) | 5.3e-4 |
| SGP-PINN ($\sigma = 0.5$, 512 IPs) | 7.0e-4 |

Figure 18.

The figure demonstrates the time evolution of the MSE loss for different models used in solving the 2D Burgers' Equation. Except the vanilla PINN (no error), all models were trained with data sampled from the initial timeslice corrupted with additive Gaussian noise with zero mean and $\sigma = 0.5$. MSE error evaluated over 50k points chosen over the entire spatio-temporal domain for the different models is given in the accompanying table




Figure 19.

The initial condition $u(x, y, t = 0)$ for 2D Burgers' equation for (a) errorfree PINN, (b) noisy PINN without smoothing, (c) noisy PINN with GP smoothing, (d) noisy PINN with SGP smoothing. The figures in the second row show the error in PINN-evaluated solution at the initial timeslice for the corresponding architecture in the top row. The final two rows show the equivalent distributions for the $v(x, y, t)$ field. Except the vanilla PINN (no error), all models were trained with data sampled from the initial timeslice corrupted with additive Gaussian noise with zero mean and $\sigma = 0.5$.




Figure 20. The leftmost column shows the exact solution for 2D Burgers' equation at $t = 0.1579$ (first and third row from top for (x, y, t)) and $t = 0.6316$ (second and fourth row from top for $v(x, y, t)$). The remaining columns show the point-wise error in PINN-evaluated solution for noise-free PINN (second column), noisy PINN without smoothing (third column), GP-PINN (fourth column), and SGP-PINN (final i.e. fifth column). Except the vanilla PINN (no error), all models were trained with data sampled from the initial timeslice corrupted with additive Gaussian noise with zero mean and $\sigma = 0.5$.

Table 1.

Comparison of average training and validation MSE losses for training and validation data on the domain boundary as given in Eqn. 14 with $\sigma = 0.1$.

| kernel | MSE loss on training data | MSE loss on validation data |
|------------------------|---------------------------|-----------------------------|
| RBF | 0.00762 | 0.0110 |
| Matèrn ($\nu = 0.1$) | 3.2×10^{-8} | 0.0465 |
| Matèrn ($\nu = 1.5$) | 0.00598 | 0.0116 |
| Matèrn ($\nu = 4.0$) | 0.00588 | 0.0126 |
| Rational Quadratic | 0.00721 | 0.0112 |

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 2.

Comparison of PINNs using different strategies for robustness to solve the 1D nonlinear Schrödinger equation. The introduction of error in the initial condition causes a significant increase in MSE for the standard PINN. GP-smoothing reduces the MSE to nearly as low as the PINN with no error. SGP-smoothing is also effective in reducing error and uses fewer inducing points (IPs). However, if the SGP does not have a sufficient number of IPs the error increases as seen when 10 IPs are used. Multiple domain cPINNs have worse performance. Results quoted for L_1 and L_2 regularizations are taken from the best performance observed over choices of $\lambda \in \{10^{-n}\}_{n=1}^5$.

| Model | MSE |
|---|--------|
| PINN (no error) | 0.0105 |
| PINN ($\sigma = 0.1$) | 0.0289 |
| PINN ($\sigma = 0.1$, L_1 regularization with $\lambda = 10^{-4}$) | 0.1613 |
| PINN ($\sigma = 0.1$, L_2 regularization with $\lambda = 10^{-4}$) | 0.2681 |
| cPINN-2 (no error) | 0.2745 |
| cPINN-2 ($\sigma = 0.1$, no smoothing) | 0.4782 |
| cPINN-3 (no error) | 0.0258 |
| cPINN-3 ($\sigma = 0.1$, no smoothing) | 0.4178 |
| GP-smoothed PINN ($\sigma = 0.1$, 50 IPs for u and v) | 0.0125 |
| SGP-smoothed PINN ($\sigma = 0.1$, 10 IPs for u and v) | 0.0231 |
| SGP-smoothed PINN ($\sigma = 0.1$, 29 and 20 IPs for u and v) | 0.0123 |

Table 3.

Comparison of PINNs using different strategies for robustness to solve the 1D Burgers' equation. The introduction of error in the initial condition causes a significant increase in MSE for the standard PINN. GP-smoothing reduces the MSE to nearly as low as the PINN with no error. SGP-smoothing is also effective in reducing error and uses fewer inducing points (IPs). Results quoted for L_1 and L_2 regularizations are taken from the best performance observed over choices of $\lambda \in \{10^{-n}\}_{n=1}^5$.

| Model | MSE |
|---|----------|
| PINN (no error) | 0.0116 |
| PINN ($\sigma = 0.5$) | 0.1982 |
| PINN ($\sigma = 0.1$, L_1 regularization with $\lambda = 10^{-4}$) | 0.0392 |
| PINN ($\sigma = 0.1$, L_2 regularization with $\lambda = 10^{-4}$) | 0.0293 |
| PINN ($\sigma = 0.5$, Cole-Hopf regularizer) | 0.1125 |
| cPINN-2 (no error) | 0.0161 |
| cPINN-2 ($\sigma = 0.5$, no smoothing) | 0.0834 |
| cPINN-2 ($\sigma = 0.5$, Cole-Hopf regularizer) | 0.0891 |
| cPINN-3 (no error) | 2.782e-5 |
| cPINN-3 ($\sigma = 0.5$, no smoothing) | 0.0854 |
| cPINN-3 ($\sigma = 0.5$, Cole-Hopf regularizer) | 0.0329 |
| UQ-PINN [20] ($\sigma = 0.5$) | 0.1248 |
| GP-smoothed PINN ($\sigma = 0.5$, 50 IPs) | 0.0384 |
| SGP-smoothed PINN ($\sigma = 0.5$, 41 IPs) | 0.0080 |