Accelerated Training of Physics Informed Neural Networks (PINNs) using Meshless Discretizations

Anonymous Author(s) Affiliation Address email

Abstract

Physics-informed neural networks (PINNs) are neural networks trained by using 1 physical laws in the form of partial differential equations (PDEs) as soft constraints. 2 We present a new technique for the accelerated training of PINNs that combines 3 modern scientific computing techniques with machine learning: discretely-trained 4 PINNs (DT-PINNs). The repeated computation of the partial derivative terms in 5 the PINN loss functions via automatic differentiation during training is known to 6 be computationally expensive, especially for higher-order derivatives. DT-PINNs 7 are trained by replacing these exact spatial derivatives with high-order accurate 8 numerical discretizations computed using meshless radial basis function-finite 9 differences (RBF-FD) and applied via sparse-matrix vector multiplication. While 10 in principle any high-order discretization may be used, the use of RBF-FD allows 11 for DT-PINNs to be trained even on point cloud samples placed on irregular 12 13 domain geometries. Additionally, though traditional PINNs (vanilla-PINNs) are typically stored and trained in 32-bit floating-point (fp32) on the GPU, we show 14 that for DT-PINNs, using fp64 on the GPU leads to significantly faster training 15 times than fp32 vanilla-PINNs with comparable accuracy. We demonstrate the 16 efficiency and accuracy of DT-PINNs via a series of experiments. First, we explore 17 the effect of network depth on both numerical and automatic differentiation of a 18 neural network with random weights and show that RBF-FD approximations of 19 third-order accuracy and above are more efficient while being sufficiently accurate. 20 We then compare the DT-PINNs to vanilla-PINNs on both linear and nonlinear 21 22 Poisson equations and show that DT-PINNs achieve similar losses with 2-4x faster training times on a consumer GPU. Finally, we also demonstrate that similar 23 results can be obtained for the PINN solution to the heat equation (a space-time 24 problem) by discretizing the spatial derivatives using RBF-FD and using automatic 25 differentiation for the temporal derivative. Our results show that fp64 DT-PINNs 26 27 offer a superior cost-accuracy profile to fp32 vanilla-PINNs, opening the door to a 28 new paradigm of leveraging scientific computing techniques to support machine learning. 29

30 1 Introduction

31 Partial differential equations (PDEs) provide a convenient framework to model a large number of phenomena across science and engineering. In real-world scenarios, PDEs are typically challenging 32 or impossible to solve using analytical techniques, and must instead be approximately solved using 33 a numerical method. A variety of numerical methods to solve these PDEs have been developed 34 including but not limited to finite difference (FD) methods (19) (which work primarily on regular 35 domains partitioned into Cartesian grids) and finite element (FE) methods (36) (which work on 36 37 irregular domains but require partitioning the domain into multidimensional simplices). A modern 38 class of numerical methods called *meshless* or *meshfree* methods generalize finite difference methods in such a way as to remove the dependence on Cartesian grids, thereby allowing for the numerical
solution of PDEs on point clouds. Of these, radial basis function-finite differences (RBF-FD) are
among the most popular and widely-used (3; 5; 37; 8; 9; 1; 11; 12; 13; 10; 25; 26; 14; 33; 18), though
a host of other such methods also exist. Much like FD or FE methods, these meshless methods can
also approximate solutions to a desired *order* of accuracy.

More recently, PDE solvers based on machine learning (ML) have begun to gain in popularity due to 44 the inherent ability of ML techniques such as neural networks (NNs) to recover highly complicated 45 functions from data specified at arbitrary locations (15; 20). We focus on a popular class of ML-based 46 meshless methods called physics-informed neural networks (PINNs) (27). PINNs can be used both 47 to discover/infer PDEs that govern a given data set, and as replacements for meshless numerical 48 methods such as RBF-FD for directly approximating solutions to PDEs. Our focus in this work is 49 on the latter problem, though our techniques extend straightforwardly to inferring PDEs as well. 50 PINNs are typically multilayer feedforward deep NNs (DNNs) that are trained using PDEs and 51 boundary conditions as soft constraints, leveraging automatic differentiation (autograd) for computing 52 spatial and/or temporal derivatives appearing in the PDE terms. The original PINNs, often referred 53 to as vanilla-PINNs, are challenging to train, at least partly because PDE-based constraints lead 54 to complicated loss landscapes (17). These issues are somewhat ameliorated by using domain 55 decomposition (X-PINNs) (16) or gradient-enhanced training (G-PINNs) (38). Other approaches for 56 ameliorating these issues involve curriculum training or sequence-to-sequence learning (17). Many 57 of these extensions can help improve training and test accuracy in addition to reducing training 58 time. Much like other DNNs, PINNs are typically trained in 32-bit floating-point (i.e., fp32 or single 59 precision). 60

In this work, we introduce a new technique for accelerating the training of vanilla-PINNs. Our 61 technique relies on two key features: (a) using RBF-FD to compute highly accurate (nevertheless 62 approximate) **spatial** derivatives in place of autograd, and (b) training the DNN in fp64 rather than 63 fp32. These new discretely-trained PINNs (DT-PINNs) can be trained significantly faster than fp32 64 vanilla-PINNs on consumer desktop GPUs with no loss in accuracy or change in DNN architecture. 65 The use of RBF-FD allows DT-PINNs to retain the meshless nature of vanilla-PINNs, thereby 66 allowing for the solution of PDEs on irregular domains. As RBF-FD uses sparse-matrix vector 67 multiplication (SpMV) to approximate the derivatives, DT-PINNs are also parallelizable on modern 68 GPU architectures. It is important to note that DT-PINNs use autograd for the actual optimization of 69 the PINN weights; only PDE derivatives are discretized using RBF-FD. 70

The NN literature does contain efforts to replace automatic differentiation with numerical differentia-71 tion. For instance, recent work showed that FD approximations can be efficient for learning generative 72 models via score matching (23). Another example is FD-Net, an NN architecture that involves learn-73 ing FD-like filters for faster prediction of PDEs (35). In the PINN literature, fractional-PINNs 74 75 (F-PINNs) use numerical differentiation as autograd cannot compute fractional derivatives (22). Neverthless, to the best of our knowledge, ours is the first work on using meshless high-order accu-76 rate FD-like methods in conjunction with PINNs, allowing them to be trained without any loss in 77 accuracy on irregular domains (just as autograd does). An alternative to replacing autograd would 78 involve eliminating inefficiencies via Taylor-mode differentiation (4). However, we show that at least 79 part of the speedups observed in DT-PINNs is because numerical differentiation results in training 80 completing in fewer epochs than if autograd were to be used. 81

To alleviate concerns about replacing autograd with RBF-FD, we first compare fp64 RBF-FD 82 approximation of different orders of accuracy against fp32 autograd for DNNs with different depths 83 and show the cost benefits of using higher-order accurate RBF-FD. Then, to illustrate the features 84 of DT-PINNs, we focus for brevity on two purely spatial PDEs (the nonlinear and linear Poisson 85 equations) and one space-time PDE (the heat equation). We use these settings to compare DT-86 PINNs and vanilla-PINNs for relative errors, timings, and speedups on a simple desktop GPU. We 87 demonstrate through our experiments that DT-PINNs offer a superior cost-accuracy profile over 88 vanilla-PINNs. 89

The remainder of this paper is organized as follows. In Section 2, we review both vanilla-PINNs and RBF-FD. Next, in Section 3, we discuss how to train DT-PINNs to solve both the Poisson and heat equations. Then, in Section 4, we present experimental results comparing RBF-FD and autograd, and comparing DT-PINNs against vanilla-PINN on the Poisson and heat equations. We summarize our results and discuss possible future work in Section 5. Finally, the appendix contains code snippets
 and key implementation details.

Notation: We use x to refer to spatial coordinates in d dimensions. On the other hand, a bolded quantity such as c or u indicates a vector with more than d elements (an array). Finally, the \sim symbol

on top of a quantity indicates that the quantity is an approximation.

99 2 Review

We now provide a brief mathematical review of both vanilla-PINNs and RBF-FD discretizations.
 Unless we note otherwise, all derivatives in this section are spatial or temporal. We focus on three
 prototypical PDEs: the nonlinear Poisson equation, the linear Poisson equation, and the heat equation.

103 2.1 Physics-informed neural networks

Let $\Omega \subset \mathbb{R}^d$ be a domain with boundary given by $\partial \Omega$; here, d is the spatial dimension. We will focus on the solution of the nonlinear Poisson equation on Ω using PINNs. Let $x \in \mathbb{R}^d$, and let $u : \mathbb{R}^d \to \mathbb{R}$ be the solution to

$$\Delta u(x) = e^{u(x)} + f(x), \ x \in \Omega, \tag{1}$$

$$(\alpha n(x) \cdot \nabla + \beta) u(x) = g(x), \ x \in \partial\Omega,$$
⁽²⁾

where Δ is the Laplacian in \mathbb{R}^d , ∇ is the \mathbb{R}^d gradient, n(x) is the unit outward normal vector on the boundary $\partial\Omega$, f(x) and g(x) are known functions, and $\alpha, \beta \in \mathbb{R}$ are known coefficients. If the $e^{u(x)}$ term is dropped from (1), we obtain the simpler *linear* Poisson equation:

$$\Delta u(x) = f(x), \ x \in \Omega. \tag{3}$$

The vanilla-PINN technique for solving either Poisson problem involves approximating the unknown 110 solution u(x) by a DNN $\tilde{u}(x, \mathbf{w})$ (where \mathbf{w} is a vector of unknown NN weights), so that $\|\tilde{u}(x, \mathbf{w}) - \tilde{u}(x, \mathbf{w})\|$ 111 $|u(x)|| \le \epsilon$ for some norm ||.|| and some tolerance ϵ . In the absence of existing solution data, this 112 is accomplished by enforcing (3) and (2) as soft constraints on $\tilde{u}(x)$ to find the weights w during training. Denote by $X = \{x_k\}_{k=1}^N$ the set of training points at which these constraints are enforced; 113 114 in the context of PDEs, these are also called **collocation points**. For convenience, we divide X into 115 two sets: N_i interior points in the set X_i and N_b boundary points in the set X_b ; then, $X = X_i \cup X_b$, 116 and $N = N_i + N_b$. Further, let $\mathcal{B} = \alpha n(x) \cdot \nabla + \beta$. The vanilla-PINN training loss $e(x, \mathbf{w})$ can then 117 be written as: 118

$$e(x, \mathbf{w}) = \underbrace{\frac{1}{N_i} \sum_{j=1}^{N_i} \left(\Delta \tilde{u}(x, \mathbf{w}) \big|_{x=x_j} - e^{\tilde{u}(x_j)} - f(x_j) \right)^2}_{\text{PDE loss in interior}} + \underbrace{\frac{1}{N_b} \sum_{i=1}^{N_b} \left(\mathcal{B}\tilde{u}(x, \mathbf{w}) \big|_{x=x_i} - g(x_i) \right)^2}_{\text{Boundary condition loss on boundary}}$$
(4)

where Δ and the ∇ term in \mathcal{B} are both applied through autograd. The tanh activation function is typically used, *L-BFGS* is used as the optimizer for finding the weights w, and training is typically

done in fp32 (17). For the linear Poisson equation, one simply omits the $e^{\bar{u}}$ term from the loss above.

For time-dependent PDEs, the PINN becomes a function of space and time $\tilde{u}(x,t)$. We focus on the forced heat equation, given by

$$\frac{\partial u(x,t)}{\partial t} = \Delta u(x,t) + f(x,t), \ x \in \Omega,$$
(5)

$$\mathcal{B}u(x,t) = g(x,t), \ x \in \partial\Omega,$$
(6)

$$u(x,0) = u_0(x),$$
 (7)

where (7) is an initial condition and $u_0(x)$ is some known function. While the Δ term is handled 124 via autograd, there are two options to handle temporal derivatives: in a continuous fashion or 125 a time-discrete fashion. In the former, one samples the full space-time interval $\Omega \times [0,T]$ with 126 collocation/training points, and then uses autograd to compute all spatial and temporal derivatives. 127 The loss terms are also augmented with the initial condition (7), which is enforced on the full space-128 time solution. In the time-discrete approach, one typically discretizes the time derivative using an 129 appropriate scheme (such as a Runge-Kutta method), and then proceeds in a step by step fashion. We 130 focus on the continuous approach in this work. 131

132 2.2 Radial basis function-finite differences (RBF-FD)

We now briefly review RBF-FD methods. Given some function $f : \mathbb{R}^d \to \mathbb{R}$, the goal of any FD formula is to approximate the action of a linear operator \mathcal{L} on that function (*i.e.*, to approximate $\mathcal{L}f$) at some location x_1 . This is typically accomplished by using a weighted linear combination of f at x_1 and its n - 1 nearest neighbors. Mathematically, this can be written as:

 $\mathcal{L}f(x)|_{x=x_1} \approx \sum_{k=1}^n c_k f(x_k),$

(8)

where the real numbers c_k are called FD weights, and the set of points x_1, \ldots, x_n is called an FD 137 stencil. In general, given a set of samples $X = \{x_j\}_{j=1}^N$, one can repeat the above procedure to find 138 FD weights at every single point. These weights can be assembled into an $N \times N$ differentiation 139 matrix L so that $\mathcal{L}f(x)|_X \approx L f(x)|_X$. If $n \ll N$, L will be a sparse matrix with at most n 140 non-zero elements per row. If X lies on a Cartesian grid, the entries of L (*i.e.*, the FD weights c_k) 141 are known in advance. However, if X is a more general point cloud, the weights c_k are unknown, 142 and standard FD (generated from polynomial interpolation) cannot be used to generate the entries 143 of L; this follows from the Mairhuber-Curtis theorem (7). The RBF-FD method involves using an 144 interpolatory combination of RBFs and polynomials instead. Without loss of generality, we describe 145 the RBF-FD procedure for x_1 and its n-1 nearest neighbors. Let $\phi(r) = r^m$, where m is odd, be 146 a radial kernel (a polyharmonic spline), and $q_i(x), j = 1, \dots, \binom{\ell+d}{d}$ be a basis for polynomials of 147 total degree ℓ in d dimensions; we use tensor-product Legendre polynomials. The RBF-FD weights 148 for the operator \mathcal{L} at the point x_1 are computed by solving the following dense (block) linear system 149 on this stencil: 150

$$\begin{bmatrix} A & P \\ P^T & 0 \end{bmatrix} \begin{bmatrix} \mathbf{c} \\ \lambda \end{bmatrix} = \begin{bmatrix} \mathcal{L}a \\ \mathcal{L}q \end{bmatrix},\tag{9}$$

151 where

$$A_{ij} = \phi(\|x_i - x_j\|), i, j = 1, \dots, n, \quad P_{ij} = q_j(x_i), i = 1, \dots, n, j = 1, \dots, \binom{\ell+d}{d}, \quad (10)$$

$$\mathcal{L}a = \mathcal{L}\phi\left(\|x - x_j\|\right)|_{x = x_1}, \quad \mathcal{L}q = \mathcal{L}q_j(x)|_{x = x_1}, j = 1, \dots, \binom{\ell + d}{d}, \tag{11}$$

152 where c is the (column) vector of n RBF-FD weights. The vector λ is a set of Lagrange multipliers enforcing the condition $P^T \mathbf{c} = \mathcal{L}q$, thereby ensuring that (a) the RBF-FD weights \mathbf{c} can exactly 153 differentiate all polynomials up to total degree ℓ ; and that (b) the error in the RBF-FD approximation to \mathcal{L} when applied to all other functions is $O(h^{\ell+1-\theta})$, where $0 \le h \le 1$ is a measure of sample 154 155 spacing in the stencil, and θ is the number of derivatives in the differential operator \mathcal{L} (6). We set 156 $\ell = p + \theta - 1$ based on the desired order of convergence p so that the error is $O(h^p)$. We then set the stencil size to $n = 2\binom{\ell+d}{d} + 1$ as this ensures that the above linear system has a solution (2), and 157 158 also set $m = \ell$ if ℓ is odd, and $m = \ell - 1$ if ℓ is even (29). L becomes more dense for higher values 159 of p and dimension d, as $n = O(p^d)$. When this procedure is repeated for each point in the set X, 160 the cost scales as O(N) for fixed n, with further speedups possible by computing multiple sets of 161 weights using each stencil (28; 29; 31; 34; 32). 162

Ghost points When tackling boundary conditions involving derivatives (such as in (2)) using RBF-FD, it is common to include a set of N_b ghost points outside the domain boundary $\partial\Omega$ into the set of samples to ensure that RBF-FD stencils at the boundary are less one-sided; this aids in numerical stability and accuracy. Ghost points allow us to also enforce the PDE at both the interior and boundary points. We therefore define and use the extended set $\tilde{X} = X_i \cup X_b \cup X_g$, where X_g is the set of ghost points. For the remainder of this article, let the RBF-FD differentiation matrix for Δ be L(dimensions $(N_i + N_b) \times (N_i + 2N_b)$), and for \mathcal{B} be B (dimensions $N_b \times (N_i + 2N_b)$).

170 **3** Discretely-Trained PINNs (DT-PINNs)

Having described both vanilla-PINNs and RBF-FD, we are now ready to describe DT-PINNs. In short, DT-PINNs are PINNs that are trained using the sparse differentiation matrices L and B in place of the autograd operations used to compute the Laplacian and boundary operators in the loss function (4) (and its heat equation equivalent). All operations are carried out in fp64. **Poisson Equation** Focusing first on the nonlinear Poisson equation (1), recall that $\tilde{u}(x, \mathbf{w})$ is the PINN approximation to the true solution u(x). Let the evaluation of $\tilde{u}(x, \mathbf{w})$ on the set \tilde{X} be \tilde{u} , *i.e.*, \tilde{u} is obtained by evaluating $\tilde{u}(x, \mathbf{w})$ at interior, boundary, *and* ghost points. Further define the vector e, which is the loss function evaluated at only the interior and boundary points, *i.e.*, $\mathbf{e} = e(x, \mathbf{w})|_X$.

179 Then, the DT-PINN loss function can be written as:

$$\mathbf{e} = \underbrace{\frac{1}{N_i + N_b} \| L \tilde{\boldsymbol{u}} - \exp(\tilde{\boldsymbol{u}}) - \mathbf{f} \|_2^2}_{\text{PDE loss in interior and on boundary}} + \underbrace{\frac{1}{N_b} \| B \tilde{\boldsymbol{u}} - \mathbf{g} \|_2^2}_{\text{Boundary condition loss on boundary}}, \quad (12)$$

where L and B were defined previously, $\exp(\tilde{u})$ is the element-wise exponential of the vector \tilde{u} , and $\mathbf{f} = f(x)|_X$, and $\mathbf{g} = g(x)|_{X_b}$; here, **f** has dimension $(N_i + N_b) \times 1$, and **g** has dimension $N_b \times 1$. For efficiency, L and B can be precomputed using RBF-FD before the training process begins, and then simply multiplied with the vector \tilde{u} to obtain its numerical derivatives. The loss function (12) is then minimized over **w** as usual using autograd in conjunction with a suitable optimizer. For the linear Poisson equation (3), we simply drop the $\exp(\tilde{u})$ term.

Heat Equation When using DT-PINNs for the heat equation, we use a mixed training technique so 186 that all temporal derivatives are discretized with autograd, and the spatial derivatives are discretized 187 with RBF-FD. We carefully order the evaluations of the network so that L and B multiply the right 188 quantities. Let $\tilde{u}(x, t, \mathbf{w})$ be the PINN, and recall that we have N_t time steps over the interval [0, T]; 189 in addition, we also have the initial condition at time t = 0, making for a total of $N_t + 1$ steps. 190 Define $\tilde{u}_k = \tilde{u}|_{x=\tilde{X},t=k \Delta t}$, where Δt is the timestep. This vector is the evaluation of \tilde{u} on all spatial 191 locations (including ghost nodes) for the k-th time slice. This definition in turn allows us to define 192 193 two vectors, \tilde{u}_{Δ} and \tilde{u}_{B} as follows:

$$\tilde{\boldsymbol{u}}_{\Delta} = \begin{bmatrix} L\tilde{\boldsymbol{u}}_{0} \\ L\tilde{\boldsymbol{u}}_{1} \\ \vdots \\ L\tilde{\boldsymbol{u}}_{N_{t}} \end{bmatrix}, \quad \tilde{\boldsymbol{u}}_{\mathcal{B}} = \begin{bmatrix} B\tilde{\boldsymbol{u}}_{0} \\ B\tilde{\boldsymbol{u}}_{1} \\ \vdots \\ B\tilde{\boldsymbol{u}}_{N_{t}} \end{bmatrix}.$$
(13)

The vector $\tilde{\boldsymbol{u}}_{\Delta}$ has dimensions $(N_t + 1)(N_i + N_b) \times 1$, and $\tilde{\boldsymbol{u}}_{\mathcal{B}}$ has dimensions $N_t N_b \times 1$. Next, we define the data vectors \mathbf{f} and \mathbf{g} as follows:

$$\mathbf{f} = \begin{bmatrix} \mathbf{f}_0 \\ \mathbf{f}_1 \\ \vdots \\ \mathbf{f}_{N_t} \end{bmatrix}, \quad \mathbf{g} = \begin{bmatrix} \mathbf{g}_0 \\ \mathbf{g}_1 \\ \vdots \\ \mathbf{g}_{N_t} \end{bmatrix}, \quad (14)$$

where $\mathbf{f}_k = f(x,t)|_{x=X,t=k \triangle t}$, and $\mathbf{g}_k = g(x,t)|_{x=X_b,t=k \triangle t}$. Finally, we define two more vectors: $u_0 = u_0(x)|_X$, the vector evaluating the initial condition on the set X (interior and boundary points); and \tilde{u}_t , the vector of evaluations of $\frac{\partial \tilde{u}}{\partial t}$ at spatial locations (interior and boundary) for each time slice:

$$\tilde{\boldsymbol{u}}_{t} = \begin{bmatrix} \left(\frac{\partial \tilde{\boldsymbol{u}}}{\partial t}\right)_{0} \\ \left(\frac{\partial \tilde{\boldsymbol{u}}}{\partial t}\right)_{1} \\ \vdots \\ \left(\frac{\partial \tilde{\boldsymbol{u}}}{\partial t}\right)_{N_{t}} \end{bmatrix},$$
(15)

where $\left(\frac{\partial \tilde{u}}{\partial t}\right)_k = \frac{\partial \tilde{u}}{\partial t}\Big|_{x=X,t=k \Delta t}$. This vector is computed using autograd. With these different vectors defined, we can finally write the DT-PINN loss vector **e** for the heat equation as

$$\mathbf{e} = \underbrace{\frac{1}{N_i + N_b} \| \boldsymbol{u}_0 - \tilde{\boldsymbol{u}} \|_{\boldsymbol{x} = \boldsymbol{X}, t = 0} \|_2^2}_{\text{Initial condition}} + \underbrace{\frac{1}{(N_t + 1)(N_i + N_b)} \| \tilde{\boldsymbol{u}}_t - \tilde{\boldsymbol{u}}_\Delta - \mathbf{f} \|_2^2}_{\text{PDE loss in interior and on boundary}} + \underbrace{\frac{1}{(N_t + 1)N_b} \| \tilde{\boldsymbol{u}}_B - \mathbf{g} \|_2^2}_{\text{Boundary condition loss on boundary}}$$
(16)

201 4 Results

We now present experimental results comparing DT-PINN and vanilla-PINN performance on the linear Poisson equation (3), the nonlinear Poisson equation (1), and the forced heat equation (5).

Setup All experiments were run for 5000 epochs on an NVIDIA GeForce RTX 2070. All results are 204 reproducible with the consistent random seeds we used in the experiments. We used the L-BFGS 205 optimizer with manually fine-tuned learning rates for both vanilla-PINNs and DT-PINNs. Both 206 DT-PINNs and vanilla-PINNs used a constant NN depth of s = 4 layers with 50 nodes each across 207 all runs. We use quasi-uniform collocation points generated using a node generator (30). For the 208 Poisson experiments, we report errors on a test set of $N_{test} = 21748$ points. For the heat equation, 209 we report results directly at the collocation points for convenience. For all experiments, the spatial 210 domain Ω is set to 211

$$\Omega = \{ x \in \mathbb{R}^d \mid \|x\|_2^2 \le 1 \}.$$
(17)

This corresponds to the unit disk in 2D, whose boundary is the unit circle $||x||_2^2 = 1$. In the timedependent 2D heat equation experiment, the space-time domain is chosen to be $\Omega \times [0, 1]$, *i.e.*, the time domain runs from time t = 0 to time t = 1. The interval [0, 1] is evenly divided into 24 time steps so that $N_t = 24$ (excluding t = 0), and the time-step was set to $\Delta t = \frac{1}{24}$. We measure all errors against a *manufactured* (specified) solution u, and specify f so that the solution holds true. The boundary condition term g is computed by applying the operator \mathcal{B} to u; we use $\alpha = \beta = 1$ for all tests. To compare DT-PINNs and vanilla-PINNs to the manufactured solutions u, we report the relative ℓ_2 error

$$e_{\ell_2} = \frac{\|\tilde{\boldsymbol{u}} - \boldsymbol{u}\|_2}{\|\boldsymbol{u}\|_2},\tag{18}$$

where u is the true solution vector, and \tilde{u} is either the DT-PINN or vanilla-PINN solution vector.

4.1 Effect of neural network depth



Figure 1: Autograd properties as a function of network depth s. The figure shows (a) effect of neural network depth s on the relative error (with respect to fp64 autograd) and (b) time taken for one application of autograd on fp32 and fp64, compared to the time taken for SpMV using RBF-FD.

We first study the effect of PINN depth (fixing the number of nodes per layer) s on computing the 222 Laplacian Δ of the output with respect to the spatial variable x using either autograd or RBF-FD. We 223 compute errors against fp64 autograd for fp32 autograd and for RBF-FD with p = 2, 3, 4, and 5. All 224 errors were computed on N = 19638 quasi-uniform collocation points. The results are in shown in 225 Figure 1a. We see fp32 autograd is the most accurate, and that increasing p increases the accuracy of 226 RBF-FD by about two orders of magnitude. Only the p = 5 case appears to match fp32 autograd in 227 accuracy, but errors are reasonably low for p = 3 and p = 4 also. In Figure 1b, we report the time 228 taken for the same test. It is immediately clear that fp64 autograd is significantly more expensive 229 than the fp32 variant, though both costs scale slowly with the network depth s. More importantly, the 230 time taken for fp64 RBF-FD (for all orders) is both lower than both fp32 and fp64 autograd and is 231 independent of the network depth s. 232



Figure 2: fp64 DT-PINNs and fp32 vanilla-PINN results on the linear Poisson equation (3) for different numbers of collocation points (N) and orders of accuracy (p). We show (a) the relative error in the PINN solution; (b) the time taken to converge to lowest relative error; and (c) the speedup attained by fp64 DT-PINNs relative to fp32 vanilla-PINN for those times.

233 4.2 Linear Poisson equation

Next, we study the performance of fp64 DT-PINNs and fp32 vanilla-PINNs on the linear Poisson equation (3) on the domain (17). Letting $x = [x_1, x_2]$, we specify the true solution u to be

$$u(x) = u(x_1, x_2) = 1 + \sin(\pi x_1) \cos(\pi x_2), \tag{19}$$

and enforce this by setting $f = \Delta u$. We then solve for \tilde{u} as described in Section 3. The results of this experiment are shown in Figure 2. We present relative errors (Figure 2a), wall clock time (Figure 2b), and speedup (Figure 2c). We do not present results for fp64 vanilla PINNs as they were almost identical to the fp32 vanilla-PINNs in terms of relative errors while being far slower. It is important to note that fp64 DT-PINNs were completely stored and trained in fp64, a format widely known to be significantly slower on the GPU than fp32.

Figure 2a shows the relative errors for DT-PINNs as a function of the number of collocation points 242 N. DT-PINNs for p = 3, 4, 5 produce similar relative errors to vanilla-PINN for the same value of N. 243 In contrast, the DT-PINN using p = 2 is generally less accurate, showing that higher-order accuracy 244 is needed to reach the same relative errors as fp32 vanilla-PINNs. Examining Figures 2b and 2c, 245 we also see that all fp64 DT-PINNs (using RBF-FD) can be trained much more rapidly than fp32 246 vanilla-PINNs. In fact, Figure 2c shows a maximum training speedup of 4x for DT-PINNs even if 247 p = 2 is ignored. In general, fp64 DT-PINNs for p > 2 are trained much more quickly than 248 vanilla-PINNs without a significant loss in accuracy. We also note that using fp32 DT-PINNs did 249 not lead to greater speedups over the fp64 DT-PINNs (results not shown); in addition, fp32 DT-PINNs 250 were not as accurate as their fp64 counterparts. The reason for the superior performance of fp64 251 DT-PINNs becomes clearer when we examine the number of training epochs as a function of the 252 number of collocation points N (Figure 3). Figures 3a and 3b both illustrate that both fp32 and fp64 253 DT-PINNs reach their lowest relative errors in fewer epochs than both fp32 and fp64 vanilla-PINNs. 254 This implies that the discrete losses are more easily minimized with respect to the NN weight vectors 255 w. However, Figure 3b shows that fp64 DT-PINNs take fewer epochs to train as N is increased (while 256 also achieving lower errors); in contrast, fp32 DT-PINNs exhibit no such behavior. We also attempted 257 to train fp32 vanilla-PINNs using ghost points, but using ghost points offered no improvement (results 258 not shown). 259

260 4.3 Nonlinear Poisson equation

Next, to understand the influence of nonlinearities in terms not including the differential operator, we test the performance of DT-PINNs on the nonlinear Poisson equation (1). To measure errors, we use the manufactured solution given by (19), and set $f = \Delta u - e^u$. The results are shown in Figure 4; we omit p = 2 to make the trends in the higher-order methods clearer. First, Figure 4a shows



Figure 3: Number of training epochs to achieve the lowest relative error as a function of number of collocation points N and order p for (a) fp32 DT-PINNs and (b) fp64 DT-PINNs.



Figure 4: fp64 DT-PINNs and fp32 vanilla-PINN results on the nonlinear Poisson equation (1) for different numbers of collocation points (N) and orders of accuracy (p). We show (a) the relative error in the PINN solution; (b) the time taken to converge to lowest relative error; and (c) the speedup attained by fp64 DT-PINNs relative to fp32 vanilla-PINN for those times.

that despite some outliers, fp64 DT-PINNs achieve comparable relative errors to fp32 vanilla-PINNs.
Further, Figure 4b shows that DT-PINNs are still trained faster than vanilla-PINNs. However, when
comparing Figure 4c to Figure 2c (linear Poisson equation), we see that the average speedup is higher
for the linear Poisson equation than the nonlinear one, though we still obtain reasonable speedups in
the latter case. This highlights one of the potential limitations of DT-PINNs: they may not offer
speedups over vanilla-PINNs if nonlinear terms not involving differential operators dominate
training times.

272 4.4 Heat equation

Next, we compare fp64 DT-PINNs and fp32 vanilla-PINNs on the 2D heat equation. We specify the true solution u to be

$$u(x,t) = u(x_1, x_2, t) = 1 + \sin(\pi x_1)\cos(\pi x_2)\sin(\pi t),$$
(20)

and specify $f = \frac{\partial u}{\partial t} - \Delta u$ so that the solution u satisfies the heat equation for all space-time. We compute the initial condition as $u_0(x,0) = u(x_1, x_2, 0) = 1$. We trained on N = 828 spatial collocation points over 25 time slices (including time t = 0) for a total of 20,700 spacetime



Figure 5: fp64 DT-PINN and fp32 vanilla-PINN results on heat equation for all p on N = 828 spatial points and $N_t = 24$ time-steps. The figure shows (a) the relative error for fp64 DT-PINNs as a function of approximation order p; and (b) the speedup attained by fp64 DT-PINNs over fp32 vanilla-PINNs as a function of p.

collocation points; we express all results as a function of p. These results are shown in Figure 5. First, Figure 5a shows similar results to the 2D Poisson equation, with p > 2 achieving relative errors similar to fp32 DT-PINNs. Figure 5b shows that we achieve 2-2.5x speedups over vanilla-PINNs. We observed in our experiments that the speedup appears to increase as a function of the number of timesteps N_t (results not shown). We conclude that DT-PINNs for p > 2 offer a superior cost-accuracy profile over vanilla-PINNs even in a time-dependent problem with mixed training.

284 **5** Summary and future work

We presented a novel technique, DT-PINNs, that involves training PINNs by using RBF-FD for 285 spatial derivatives, and using fp64 weights and training instead of fp32. This involved replacing 286 all autograd operations (dense matrix-matrix multiplies) related to PDE loss terms with an SpMV 287 operation. We showed that using an RBF-FD approximation order of p > 2 resulted in DT-PINNs 288 that were comparable in accuracy to vanilla-PINNs while offering 2-4x speedups in training times 289 for both the linear and nonlinear Poisson equations. For p = 2 (a second-order discretization), we 290 achieved much higher speedups (20x or above for small training sets), but with a loss of accuracy for 291 small N. We also showed that DT-PINNs trained in a mixed fashion (autograd for time, RBF-FD for 292 space) also achieved comparable accuracy with a 2-2.5x speedup on the heat equation. DT-PINNs 293 therefore constitute a new paradigm for scientific machine learning that allow practitioners to leverage 294 existing sophisticated scientific computing techniques to accelerate ML training times. 295

There are several possible extensions to our current work. It is likely that using DT-PINNs in 296 conjunction with X-PINNs and G-PINNs will yield even greater speedups in training times. Further, 297 DT-PINNs open the door to leveraging compute more efficiently. For instance, the SpMV operations 298 could be parallelized using distributed memory systems in conjunction with GPUs, thereby allowing 299 scaling to very large training sets; alternatively, the SpMV operation could be parallelized on 300 many-core CPUs while other operations are conducted on the GPU. It may also be profitable to 301 explore mixed-precision training of DT-PINNs using combinations of fp64, fp32, and even fp16 302 in conjunction with RBF-FD. Finally, DT-PINNs can be viewed as vanilla-PINNs with partially 303 linearized constraints; it may be profitable to explore other types of constraint linearization to 304 accelerate training and simplify loss function landscapes. 305

306 **References**

- [1] Barnett, G. A. (2015). A Robust RBF-FD Formulation based on Polyharmonic Splines and
 Polynomials. PhD thesis, University of Colorado Boulder.
- Bayona, V., Flyer, N., Fornberg, B., and Barnett, G. A. (2017). On the role of polynomials in RBF-FD approximations: II. Numerical solution of elliptic PDEs. *J. Comput. Phys.*, 332:257–273.

- [3] Bayona, V., Moscoso, M., Carretero, M., and Kindelan, M. (2010). RBF-FD formulas and
 convergence properties. J. Comput. Phys., 229(22):8281–8295.
- [4] Bettencourt, J., Johnson, M. J., and Duvenaud, D. (2019). Taylor-mode automatic differentiation
 for higher-order derivatives in jax.
- [5] Davydov, O. and Oanh, D. T. (2011). Adaptive meshless centres and RBF stencils for Poisson
 equation. J. Comput. Phys., 230(2):287–304.
- [6] Davydov, O. and Schaback, R. (2018). Minimal numerical differentiation formulas. *Numerische Mathematik*, 140(3):555–592.
- [7] Fasshauer, G. E. (2007). *Meshfree Approximation Methods with MATLAB*. Interdisciplinary
 Mathematical Sciences Vol. 6. World Scientific Publishers, Singapore.
- [8] Flyer, N., Barnett, G. A., and Wicker, L. J. (2016a). Enhancing finite differences with radial basis
 functions: Experiments on the Navier-Stokes equations. *J. Comput. Phys.*, 316:39–62.
- [9] Flyer, N., Fornberg, B., Bayona, V., and Barnett, G. A. (2016b). On the role of polynomials in
 RBF-FD approximations: I. Interpolation and accuracy. *J. Comput. Phys.*, 321:21–38.
- [10] Flyer, N., Lehto, E., Blaise, S., Wright, G. B., and St-Cyr, A. (2012). A guide to RBF-generated finite differences for nonlinear transport: shallow water simulations on a sphere. *J. Comput. Phys.*, 231:4078–4095.
- [11] Flyer, N. and Wright, G. B. (2007). Transport schemes on a sphere using radial basis functions.
 J. Comput. Phys., 226:1059–1084.
- [12] Flyer, N. and Wright, G. B. (2009). A radial basis function method for the shallow water
 equations on a sphere. *Proc. Roy. Soc. A*, 465:1949–1976.
- [13] Fornberg, B. and Lehto, E. (2011). Stabilization of RBF-generated finite difference methods for
 convective PDEs. *J. Comput. Phys.*, 230:2270–2285.
- [14] Fuselier, E. J. and Wright, G. B. (2013). A high-order kernel method for diffusion and reactiondiffusion equations on surfaces. *J. Sci. Comput.*, 56(3):535–565.
- [15] Han, J., Jentzen, A., and E, W. (2018). Solving high-dimensional partial differential equations
 using deep learning. *Proceedings of the National Academy of Sciences*, 115(34):8505–8510.
- [16] Jagtap, A. D. and Karniadakis, G. E. (2020). 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(5):2002–2041.
- [17] Krishnapriyan, A., Gholami, A., Zhe, S., Kirby, R., and Mahoney, M. W. (2021). Characterizing
 possible failure modes in physics-informed neural networks. In Beygelzimer, A., Dauphin, Y.,
 Liang, P., and Vaughan, J. W., editors, *Advances in Neural Information Processing Systems*.
- [18] Lehto, E., Shankar, V., and Wright, G. B. (2017). A radial basis function (RBF) compact
 finite difference (FD) scheme for reaction-diffusion equations on surfaces. *SIAM J. Sci. Comput.*,
 39:A2129–A2151.
- [19] LeVeque, R. J. (2007). Finite difference methods for ordinary and partial differential equations:
 steady-state and time-dependent problems. SIAM.
- [20] Long, Z., Lu, Y., Ma, X., and Dong, B. (2017). Pde-net: Learning pdes from data.
- [21] Okuta, R., Unno, Y., Nishino, D., Hido, S., and Loomis, C. (2017). Cupy: A numpy-compatible
 library for nvidia gpu calculations. In *Proceedings of Workshop on Machine Learning Systems* (*LearningSys*) in *The Thirty-first Annual Conference on Neural Information Processing Systems* (*NIPS*).
- [22] Pang, G., Lu, L., and Karniadakis, G. E. (2019). fpinns: Fractional physics-informed neural networks. *SIAM Journal on Scientific Computing*, 41(4):A2603–A2626.

- Pang, T., Xu, K., LI, C., Song, Y., Ermon, S., and Zhu, J. (2020). Efficient learning of generative models via finite-difference score matching. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H., editors, *Advances in Neural Information Processing Systems*, volume 33, pages 19175–19188. Curran Associates, Inc.
- [24] Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z.,
 Gimelshein, N., Antiga, L., et al. (2019). Pytorch: An imperative style, high-performance deep
 learning library. *Advances in neural information processing systems*, 32.
- [25] Piret, C. (2012). The orthogonal gradients method: A radial basis functions method for solving
 partial differential equations on arbitrary surfaces. J. Comput. Phys., 231(20):4662–4675.
- [26] Piret, C. and Dunn, J. (2016). Fast RBF OGr for solving pdes on arbitrary surfaces. AIP
 Conference Proceedings, 1776(1).
- [27] Raissi, M., Perdikaris, P., and Karniadakis, G. (2019). 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.
- [28] Shankar, V. (2017). The overlapped radial basis function-finite difference (RBF-FD) method: A
 generalization of RBF-FD. *J. Comput. Phys.*, 342:211–228.
- [29] Shankar, V. and Fogelson, A. L. (2018). Hyperviscosity-based stabilization for radial basis
 function-finite difference (rbf-fd) discretizations of advection– diffusion equations. *J. Comput. Phys.*, 372:616 639.
- [30] Shankar, V., Kirby, R., and Fogelson, A. (2018a). Robust node generation for mesh-free discretizations on irregular domains and surfaces. *SIAM Journal on Scientific Computing*, 40(4):A2584–A2608.
- [31] Shankar, V., Narayan, A., and Kirby, R. M. (2018b). Rbf-loi: Augmenting radial basis
 functions (rbfs) with least orthogonal interpolation (loi) for solving pdes on surfaces. *Journal of Computational Physics*, 373:722–735.
- [32] Shankar, V., Wright, G. B., and Fogelson, A. L. (2021). An efficient high-order meshless method
 for advection-diffusion equations on time-varying irregular domains. *Journal of Computational Physics*, 445:110633.
- [33] Shankar, V., Wright, G. B., Kirby, R. M., and Fogelson, A. L. (2014). A radial basis function
 (RBF)-finite difference (FD) method for diffusion and reaction–diffusion equations on surfaces. J.
 Sci. Comput., 63(3):745–768.
- [34] Shankar, V., Wright, G. B., and Narayan, A. (2020). A robust hyperviscosity formulation for
 stable RBF-FD discretizations of Advection-Diffusion-Reaction equations on manifolds. *SIAM Journal on Scientific Computing*, 42(4):A2371–A2401.
- [35] Shi, Z., Gulgec, N. S., Berahas, A. S., Pakzad, S. N., and Takáč, M. (2020). Finite difference
 neural networks: Fast prediction of partial differential equations. In 2020 19th IEEE International
 Conference on Machine Learning and Applications (ICMLA), pages 130–135. IEEE.
- [36] Strang, G., Fix, G. J., and Griffin, D. (1974). An analysis of the finite-element method.
- [37] Wright, G. B. and Fornberg, B. (2006). Scattered node compact finite difference-type formulas
 generated from radial basis functions. *J. Comput. Phys.*, 212(1):99–123.
- [38] Yu, J., Lu, L., Meng, X., and Karniadakis, G. E. (2022). Gradient-enhanced physics-informed
 neural networks for forward and inverse pde problems. *Computer Methods in Applied Mechanics and Engineering*, 393:114823.

400 Checklist

401	1. For all authors
402	(a) Do the main claims made in the abstract and introduction accurately reflect the paper's
403	contributions and scope? [Yes]
404	(b) Did you describe the limitations of your work? [Yes]
405	(c) Did you discuss any potential negative societal impacts of your work? [No] To the
406	best of the authors' knowledge, there are no negative societal impacts of our work
407	including potential malicious or unintended uses, environmental impact, security, or
408	privacy concerns.
409	(d) Have you read the ethics review guidelines and ensured that your paper conforms to
410	them? [Yes]
411	2. If you are including theoretical results
412	(a) Did you state the full set of assumptions of all theoretical results? [N/A]
413	(b) Did you include complete proofs of all theoretical results? [N/A]
414	3. If you ran experiments
415	(a) Did you include the code, data, and instructions needed to reproduce the main experi-
416	mental results (either in the supplemental material or as a URL)? [Yes]
417	(b) Did you specify all the training details (e.g., data splits, hyperparameters, how they
418	were chosen)? [Yes]
419	(c) Did you report error bars (e.g., with respect to the random seed after running experi-
420	ments multiple times)? [N/A] Reporting error bars over multiple random runs is not
421	the norm in physics-informed neural network literature. Moreover, the authors did
422	not have the range of compute required for adequate number of random runs at the
423	time of writing. Nevertheless, since our examples are deterministic, all results can be
424	reproduced if the same random seed is used.
425	(d) Did you include the total amount of compute and the type of resources used (e.g., type
426	of GPUs, internal cluster, or cloud provider)? [Yes]
427	4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets
428	(a) If your work uses existing assets, did you cite the creators? [N/A]
429	(b) Did you mention the license of the assets? [N/A]
430	(c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
431	We release the datasets and codebase we used as part of the supplementary material.
432	(d) Did you discuss whether and how consent was obtained from people whose data you're
433	using/curating? [N/A]
434	(e) Did you discuss whether the data you are using/curating contains personally identifiable
435	information or offensive content? [N/A]
436	5. If you used crowdsourcing or conducted research with human subjects
437	(a) Did you include the full text of instructions given to participants and screenshots, if
438	applicable? [N/A]
439	(b) Did you describe any potential participant risks, with links to Institutional Review
440	Board (IRB) approvals, if applicable? [N/A]
441	(c) Did you include the estimated hourly wage paid to participants and the total amount
442	spent on participant compensation? [N/A]

443 A Appendix

From a practical implementation standpoint, DT-PINNs differ from vanilla-PINNs in several ways
 beyond the replacement of autograd with differentiation matrices. We discuss some of the details of
 our implementation within the PyTorch framework (24).

Efficient sparse matrix storage and operations on the GPU Since DT-PINNs replace all autograd computations in the loss function with an SpMV, we use the compressed sparse row (CSR) format for storing *L* and *B*; this ensures both ease of parallelization and compact storage. As of this writing, PyTorch support for SpMV operations in the CSR format on the GPU is limited and somewhat inefficient. To overcome this limitation, we used the CuPy library (21) for all SpMV operations. We also use the DLPack library to enable memory and context sharing between CuPy and PyTorch.

Custom autograd implementation As of this writing, PyTorch is in general unable to apply autograd 453 with respect to the weight vector w to loss terms involving CuPy CSR matrices. To overcome this 454 issue, we wrote custom autograd classes in PyTorch. Consider the second term in (12), $\|L\tilde{u} - f\|_2^2$. 455 According to the PyTorch API, a custom autograd class for handling this term must supply two 456 methods: a forward method that tells PyTorch how to evaluate it, and a backward method that tells 457 PyTorch how to compute the product $(\nabla_{\tilde{u}} (L\tilde{u} - \mathbf{f})) (\nabla_{\mathbf{w}} \tilde{u})$. As $\nabla_{\tilde{u}} (L\tilde{u} - \mathbf{f}) = L^T$, our custom 458 PyTorch backward method can be expressed as a CuPy-based SpMV of L^T and $\nabla_{\mathbf{w}} \tilde{\boldsymbol{u}}$, the latter of 459 which PyTorch automatically supplies; loss terms involving B are handled similarly. The forward 460 method to evaluate the loss term is also an SpMV between L and \tilde{u} . The overall procedure is similar 461 for the heat equation as well. 462

463 Code listings We first show the CuPy custom autograd methods for the SpMV operations in the 464 autograd term. The required imports are:

```
466 | import torch
467 2 from torch.utils.dlpack import to_dlpack, from_dlpack
468 3 import cupy
468 4 from cupy.sparse import csr_matrix
```

The following Python class defines the **forward** and **backward** methods to connect the CuPy SpMV with the native PyTorch tensors.

```
473
    class Cupy_L(torch.autograd.Function):
474
        @staticmethod
475
        def forward(ctx, pinn_pred, sparse_mat):
476
477
             return from_dlpack(sparse_mat.dot(
                    cupy.from_dlpack(to_dlpack(pinn_pred))
478
                                    ).toDlpack())
479
480
481
        @staticmethod
        def backward(ctx, grad_output):
482
            return from_dlpack(L_t.dot(
483
                    cupy.from_dlpack(to_dlpack(grad_output))
4841
                                    ).toDlpack()), None
4851
4861
    # class Cupy_B can be similarly defined
4881
```

489 where L_t is the transpose of the *L* matrix:

465

```
490
492
L_t = csr_matrix(self.L.transpose(), dtype=np.float64)
```

⁴⁹³ Each DT-PINN training step in linear Poisson can then be written as follows:

```
494
   # we use the L-BFGS optimizer provided by PyTorch
495
   def closure():
496
        self.optimizer.zero_grad()
497
        u_tilde = self.pinn_weights.forward(self.X_tilde)
498
499
        pde_residual = L_mul(u_tilde, self.L) - self.f
500
        bdry_residual = B_mul(u_tilde, self.B) - self.g
501
502
        pde_loss = torch.mean(torch.square(torch.flatten(pde_residual)))
503
        bdry_loss = torch.mean(torch.square(torch.flatten(bdry_residual)))
5041
5051
        train_loss = interior_loss + boundary_loss
5061
        train_loss.backward(retain_graph=True)
507
        return train_loss.item()
5081
5091
   loss_value = self.optimizer.step(closure)
5101
```

si2 where self.X_tilde is \tilde{X} and L_mul and B_mul are:

513 514 | L_mul = Cupy_L.apply 515 2 B_mul = Cupy_B.apply