

# Implicit Stochastic Gradient Descent for Training Physics-Informed Neural Networks

Ye Li, Song-Can Chen, Sheng-Jun Huang

College of Computer Science and Technology/Artificial Intelligence, Nanjing University of Aeronautics and Astronautics  
 MITT Key Laboratory of Pattern Analysis and Machine Intelligence, Nanjing, China  
 yeli20@nuaa.edu.cn, s.chen@nuaa.edu.cn, huangsj@nuaa.edu.cn

## Abstract

Physics-informed neural networks (PINNs) have effectively been demonstrated in solving forward and inverse differential equation problems, but they are still trapped in training failures when the target functions to be approximated exhibit high-frequency or multi-scale features. In this paper, we propose to employ implicit stochastic gradient descent (ISGD) method to train PINNs for improving the stability of training process. We heuristically analyze how ISGD overcome stiffness in the gradient flow dynamics of PINNs, especially for problems with multi-scale solutions. We theoretically prove that for two-layer fully connected neural networks with large hidden nodes, randomly initialized ISGD converges to a globally optimal solution for the quadratic loss function. Empirical results demonstrate that ISGD works well in practice and compares favorably to other gradient-based optimization methods such as SGD and Adam, while can also effectively address the numerical stiffness in training dynamics via gradient descent.

## Introduction

Gradient descent (GD) and practical stochastic gradient descent with mini-batch gradients (SGD) are widely used optimization algorithms, especially in optimizing deep neural networks. Formally, the goal of optimization is to find a weight vector  $\hat{\theta}$  in parameter space  $\mathbb{R}^m$  that minimizes a loss  $L(\theta)$ . The GD algorithm is the updating procedure of model weights in the direction of the steepest loss gradient:

$$\theta_{n+1} = \theta_n - \alpha \cdot \nabla L(\theta_n), \quad (1)$$

where  $\alpha$  is the learning rate. The SGD replaces the gradient  $\nabla L(\theta)$  with a mini-batch gradient  $\nabla \hat{L}_i(\theta)$ , where  $\hat{L}_i$  is the loss computed on mini-batch data instead of the whole dataset. The continuous gradient flow is defined as a curvature  $\theta(t)$  that satisfies the following ordinary differential equation (ODE):

$$\frac{d}{dt}\theta(t) = -\nabla_{\theta}L(\theta(t)). \quad (2)$$

It is easy to show that when the learning rate is sufficiently small, the discrete updates  $\{\theta_n\}_{n=0}^{\infty}$  computed by Eq.(1)

stay close to a function  $\{\theta(t_n)\}_{n=0}^{\infty}$  where  $t_n = n\alpha$ . Variants based on GD/SGD, such as AdaGrad (Duchi, Hazan, and Singer 2011), RMSprop (Tieleman and Hinton 2012), and Adam (Kingma and Ba 2014), have been developed in recent years.

Despite its numerous successes in practical optimization tasks such as optimizing deep neural networks, GD/SGD may suffer from numerical instability in some key hyperparameters, such as the learning rate and batch size. For example, if the learning rate is misspecified, GD/SGD may numerically diverge, and the model training fails. The main reason is the *stiffness* in the gradient flow dynamics. Typically, the gradient flow dynamics is called a *stiff ODE* when the gap between the maximum and minimum eigenvalues of the Hessian matrix is large (Wang, Teng, and Perdikaris 2021). We can simply perform a linearization for the gradient flow (2) and obtain

$$\frac{d}{dt}\tilde{\theta}(t) = -\nabla_{\theta}^2L(\tilde{\theta}(t)) \cdot \tilde{\theta}(t). \quad (3)$$

The largest eigenvalue of the Hessian dictates the fastest time-scale of the ODEs. In the language of numerical analysis, to ensure the numerical stability of GD, we need  $\alpha \leq 2/\lambda_{\max}(\nabla_{\theta}^2L(\theta))$ , where  $\lambda_{\max}(\nabla_{\theta}^2L(\theta))$  is the maximum eigenvalue of the Hessian matrix (Butcher 2016).

From the theory of numerical analysis, GD/SGD is not suitable for stiff ODEs, because a very small learning rate and very large number of iterations are required to maintain numerical stability. One of the outstanding first-order solvers with strong stability for stiff ODEs is the implicit (backward) Euler method:

$$\theta_{n+1} = \theta_n - \alpha \cdot \nabla L(\theta_{n+1}), \quad (4)$$

where a large learning rate can be used. Eq.(4) is also known as the implicit gradient descent (IGD) or implicit stochastic gradient descent (ISGD) method, as the next iteration  $\theta_{n+1}$  appears implicitly on the right side of Eq.(4), and cannot be computed explicitly.

Physics-informed neural networks (PINNs) are neural networks with outputs constrained to approximately satisfy a system of partial differential equations (PDEs) by using a regularization functional  $\mathcal{R}(u_{\theta}(\mathbf{x}))$  that typically represents the residual of PDEs. A general loss function representation




Figure 1: *1D Poisson equation*: Results for the heuristic example in Section obtained by training a conventional PINN (5-layer, 200 hidden units, tanh activations) with gradient descent based Adam optimizer, quasi-Newton based L-BFGS optimizer and our ISGD optimizer. All eigenvalues of  $\nabla_{\theta}^2 L(\theta)$  are computed and arranged in increasing order. **(a)**: smooth solution, the maximum eigenvalue is  $1.1\text{e}+04$ , non-stiff and all three optimizers trained well. **(b)**: solution with multi-scale features, the maximum eigenvalue is  $4.6\text{e}+08$ , stiff and both Adam (green dashed line) and L-BFGS (orange dashed line) failed to train, while our ISGD (red dashed line) trained well.

of PINNs takes the form

$$L(\theta) = \frac{1}{N_u} \sum_{i=1}^{N_u} |\mathbf{u}_i - u_{\theta}(\mathbf{x}_i)|^2 + \mathcal{R}(u_{\theta}(\mathbf{x})), \quad (5)$$

where the given set input output pairs  $(\mathbf{x}_i, \mathbf{u}_i)$  are corresponding to the initial/boundary conditions of PDEs. The most popular optimizers for training PINNs are gradient descent based Adam optimizer and quasi-Newton based L-BFGS optimizer (Lu et al. 2021). However, the additional regularization term  $\mathcal{R}(u_{\theta}(\mathbf{x}))$  has been shown to *increase the stiffness* of the gradient dynamics (Wang, Teng, and Perdikaris 2021), causing model training failures especially when the target functions to be approximated exhibit high-frequency or multi-scale features. Typically, for stiff solutions, L-BFGS is more likely to be stuck at a bad local minimum, and Adam may need very small learning rate and very large number of iterations. We claim that IGD/ISGD is more stable than GD/SGD and L-BFGS in the PINNs training when fitting multi-scale solutions. As an example, Figure 1 contrasts these approaches. As the solution of Poisson equation changes from smooth to multi-scale, the maximum eigenvalue of Hessian increases significantly and the gradient flow dynamics of PINN becomes stiff, both Adam and L-BFGS become divergent while our IGD/ISGD is still convergent.

## Contributions

Our main contributions can be summarized in the following points:

- We first propose to employ the IGD/ISGD method to train PINNs. We theoretically and numerically show that IGD/ISGD can overcome the stiffness in the gradient flow dynamics of PINNs, especially for PDEs with multi-scale solutions.
- We used practical L-BFGS and Adam optimizer to deal with the implicit updates in IGD/ISGD, which is effective in practice. The computational cost is comparable to Adam. Furthermore, the method is stable for the learning rate and batch size, making it easier for nonexperts to process neural network training tasks.

- The IGD global convergence property is proven. We theoretically prove that for two-layer fully connected neural networks with large hidden nodes, randomly initialized IGD converges to a globally optimal solution at a linear convergence rate for the quadratic loss function.

## Related Work

**Gradient Descent.** Global convergence of gradient descent based methods have been proved when training deep neural network despite the objective function being non-convex (Du et al. 2019; Du 2019; Du et al. 2018; Allen-Zhu, Li, and Song 2019; Zou et al. 2020). The dynamics of neural network weights under GD converge to a point that is close to the minimum norm solution under proper conditions (Sathathi and Srikant 2021).

Toulis and his collaborators (Toulis and Airolidi 2017; Toulis, Tran, and Airolidi 2016; Toulis, Airolidi, and Rennie 2014) first theoretically studied the implicit stochastic gradient descent algorithm, and claimed it to be more stable than standard stochastic gradient descent. However, both the theoretical and practical results of them are only suited to generalized linear models. The implicit scheme was extended to combine with the ResNet architecture with implicit Euler skip connections (called IE-ResNet) by (Li, He, and Lin 2020) to improve the robustness and generalization ability. The IGD/ISGD method was also applied to optimize the k-means clustering problem (Yin et al. 2018) and the objective matrix factorization loss function that appears in recommendation systems (Vo, Hong, and Jung 2020), and the convergence time was effectively improved.

**PINNs.** With the rapid growth of deep learning, using neural networks to represent PDE solutions has attracted the attention of many researchers. Based on the early studies of Psychogiios and Ungar (1992); Lagaris, Likas, and Fotiadis (1998), Raissi, Perdikaris, and Karniadakis (2019) proposed the pioneering work of PINNs to solve both forward and inverse problems involving nonlinear PDEs. PINNs have demonstrated remarkable power in applications including fluid dynamics (Raissi, Yazdani, and Karniadakis 2020; Jin et al. 2021; Mao, Jagtap, and Karniadakis 2020), biomedical engineering (Sahli Costabal et al. 2020), meta-material design (Fang and Zhan 2019; Chen et al. 2020), software

packages (Lu et al. 2021), and numerical simulators (Henning et al. 2021; Cai et al. 2021). Adaptive activation functions can be applied to accelerate PINN training (Jagtap, Kawaguchi, and Em Karniadakis 2020; Jagtap, Kawaguchi, and Karniadakis 2020; Jagtap et al. 2022). However, despite early empirical success, the original formulations of PINNs often struggles to handle problems exhibiting high-frequency and multi-scale behavior.

Recent works by Wang, Teng, and Perdikaris (2021); Wang, Yu, and Perdikaris (2022); Wang, Wang, and Perdikaris (2021) have identified two fundamental weaknesses in conventional PINN formulations. The first is the remarkable discrepancy in the convergence rate between the data-based loss function and the physical-based loss function. The second is related to the spectral bias, which indeed exists in PINN models and is the leading reason that prevents them from accurately approximating high-frequency or multi-scale functions. In fact, they demonstrated that the gradient flow of PINN models becomes increasingly stiff for PDE solutions exhibiting high-frequency or multi-scale behavior. This result motivates us to use robust implicit numerical schemes such as IGD/ISGD for the numerical solution to the gradient flow of PINN models.

## Organization of the Paper

In Section 2, we present the methodology of the proposed IGD/ISGD method. The PINNs framework is also introduced briefly for completeness. Two heuristic examples are presented to show the strong stability of the IGD/ISGD method. In Section 3, we analyze the training dynamics of the IGD/ISGD method when applied to neural network training tasks. In Section 4, we report various computational examples for inferring the solution of ordinary/partial differential equations by PINNs. Finally, we conclude in Section 5 with a summary.

## Methodology

### Physics-Informed Neural Networks

PINNs are neural networks that imbeds differential equations into neural network training. The initial/boundary condition data of the differential equations are treated as the supervised learning component in the objective loss function, while the residual of the differential equations is applied as an unsupervised regularization factor in the objective loss function. We consider a parametrized PDE system given by:

$$\begin{aligned} \mathcal{F}(\mathbf{x}, u, u_x, \dots, \lambda) &= 0, & \mathbf{x} \in \Omega, \\ u(\mathbf{x}) &= g_0(\mathbf{x}), & \mathbf{x} \in \partial\Omega, \end{aligned}$$

where  $\mathbf{x}$  are the spatial and time coordinates,  $u = u(\mathbf{x})$  is the solution to the PDE with boundary/initial data  $g_0(\mathbf{x})$ ,  $\mathcal{F}$  denotes the PDE residual, and  $\lambda$  is the PDE parameter. For example,  $\mathcal{F} = -u_{xx} - f(x) = 0$  is the simplest 1D Poisson equation for a given function  $f(x)$ . The vanilla PINN uses a fully connected feed-forward neural network  $u_\theta(x)$  to approximate the solution  $u(x)$  by minimizing the following loss function:

$$L(\theta) = \omega_d L_{data} + \omega_f L_{PDE}, \quad (6)$$

where

$$\begin{aligned} L_{data} &= \frac{1}{N_d} \sum_{j=1}^{N_d} |u_\theta(\mathbf{x}_d^j) - g_0(\mathbf{x}_d^j)|^2, \\ L_{PDE} &= \frac{1}{N_f} \sum_{i=1}^{N_f} |\mathcal{F}(\mathbf{x}_f^i)|^2. \end{aligned}$$

Here,  $\{\mathbf{x}_d^j\}_{j=1}^{N_d}$  represents the training data points on  $\partial\Omega$  while  $\{\mathbf{x}_f^i\}_{i=1}^{N_f}$  represents the set of residual points in  $\Omega$ .  $\omega_f$  and  $\omega_d$  are the user-specified weighting coefficients for different loss terms. The first term  $L_{data}$  includes the known boundary/initial conditions and experimental data, which is the usual supervised data-driven part of the neural network. To compute the residuals in the loss function, automatic differentiation is applied to compute the derivatives of the solution with respect to the independent variables. This constitutes the physics-informed part of the neural network as given by the second term  $L_{PDE}$ .

The resulting optimization problem is to find the minimum of the loss function by optimizing the trainable parameters  $\theta$ . Gradient descent based first-order optimizers such as SGD and Adam (Kingma and Ba 2014), or quasi-Newton based optimizers like L-BFGS (Liu and Nocedal 1989), are widely used in PINNs training. However, as Wang, Yu, and Perdikaris (2022) claimed, "...PINNs using fully connected architectures often fail to achieve stable training and produce accurate predictions, especially when the underlying PDE solutions contain high-frequencies or multi-scale features". The gradient flow dynamics of PINNs will become stiff as multi-scale phenomena appear, so explicit GD based optimizers may be unstable, and L-BFGS is more likely to be stuck at a bad local minimum. As we mentioned in the previous section, implicit schemes like IGD/ISGD are more stable to overcome the stiffness problems. Two illustrative examples are presented to show the robustness of IGD/ISGD in the next section.

### Heuristic Examples with Stability

In this section, we present two heuristic examples to show the stability of IGD/ISGD and the instability of GD/IGD.

**Analytical Stiff Problem.** The first example is to theoretically analysis the learning rate constraint in the gradient flow dynamics of stiffness problems. We denote a fabricated loss function by

$$L(\theta_1, \theta_2) = \frac{K_1}{2} (\theta_1 - \theta_1^*)^2 + \frac{K_2}{2} (\theta_2 - \theta_2^*)^2,$$

where  $\theta_i \in \mathbb{R}$ ,  $i = 1, 2$  are two parameters to be optimized,  $K_i > 0$ ,  $i = 1, 2$  are two constants. The eigenvalues of the Hessian matrix of  $L(\theta_1, \theta_2)$  are characterized by  $K_1$  and  $K_2$ . When  $K_1$  and  $K_2$  differ in scales, for example,  $K_1 = 10^{-4}$  and  $K_2 = 10^4$ , the gradient flow of the loss function suffers from the stiffness phenomenon.

A direct computation shows that the loss function update procedure of GD has the following relation:

$$\frac{L(\theta_1^{n+1}, \theta_2^{n+1})}{L(\theta_1^n, \theta_2^n)} \leq \max\{(1 - \alpha K_1)^2, (1 - \alpha K_2)^2\}. \quad (7)$$

Typically, we need  $D = \max\{(1-\alpha K_1)^2, (1-\alpha K_2)^2\} \leq 1$  to guarantee loss decay, which implies  $\alpha \leq \frac{2}{\max\{K_1, K_2\}}$ . When  $K_1 = 10^{-4}$  and  $K_2 = 10^4$ , we have  $\alpha \leq 10^{-4}$  and  $D \leq 1-10^{-8}$ , meaning that the loss decays very slowly, and very large number of iterations (at least  $\mathcal{O}(10^8)$ ) are needed to converge. For a large learning rate  $\alpha$ , the loss decay rate  $D$  may be greater than 1, and the loss may increase as the iterations increase, causing numerical instability in the gradient flow dynamics computation.

For IGD method, the loss function update procedure has the following relation:

$$\frac{L(\theta_1^{n+1}, \theta_2^{n+1})}{L(\theta_1^n, \theta_2^n)} \leq \max\left\{\frac{1}{(1+\alpha K_1)^2}, \frac{1}{(1+\alpha K_2)^2}\right\}. \quad (8)$$

The loss decay rate  $D = \max\left\{\frac{1}{(1+\alpha K_1)^2}, \frac{1}{(1+\alpha K_2)^2}\right\}$  satisfies  $D < 1$  automatically for all learning rates  $\alpha > 0$  and regardless of the scales of  $K_1, K_2$ , and  $D$  is even smaller for larger  $\alpha$ . This shows the strong stability of IGD to deal with stiffness phenomena.

**1D Poisson Equation with Multi-Scale Solution.** This heuristic example is to show the advantage of IGD/ISGD when the gradient flow dynamics of PINN is stiff. We consider a simple 1D Poisson equation

$$-\Delta u(x) = f(x), \quad x \in (0, 1) \quad (9)$$

subject to the boundary condition

$$u(0) = u(1) = 0.$$

We consider two fabricated solutions: one is  $u_L(x) = \sin(2\pi x)$  exhibiting low frequency on the whole domain, and another is  $u_H(x) = \sin(2\pi x) + 0.1 \sin(50\pi x)$  exhibiting low frequency in the macro-scale and high frequency in the micro-scale. Though this example is simple and pedagogical, it resembles many practical scenarios with multi-scale phenomenons.

We represent the unknown solution  $u(x)$  by a 5-layer fully-connected neural network  $u_\theta(x)$  with 200 units per hidden layer.  $N_r = 1000$  training points  $\{x_i, f(x_i)\}$  are uniformly sampled in the interval  $(0, 1)$ . Figure 1 shows the results obtained by training PINN with gradient descent based Adam optimizer (Kingma and Ba 2014) with default settings for a maximum  $10^7$  epochs, quasi-Newton based L-BFGS optimizer (Liu and Nocedal 1989) with default settings, and our ISGD method with learning rate 0.1 for a maximum  $10^4$  epochs. We observe that all three optimizers can train PINN well for smooth solution  $u_L(x)$  when there is non-stiff. As multi-scale solution  $u_H(x)$  appears, the maximum eigenvalue of Hessian has a significant rise from  $1.1\text{e}+04$  to  $4.6\text{e}+08$ . The gradient flow dynamics of PINN becomes stiff, and the popular Adam optimizer is incapable of training PINN to the correct solution even after a million training epochs. The L-BFGS optimizer is also failed to train. As a comparison, our ISGD method can train PINN well both for smooth  $u_L(x)$  as well as multi-scale  $u_H(x)$  with larger learning rate and smaller iterations.

## Loss Decay of GD/IGD

Wang, Teng, and Perdikaris (2021) shows that the loss decay

of GD is

$$\begin{aligned} & L(\theta_{n+1}) - L(\theta_n) \\ &= \alpha \|\nabla_\theta L(\theta_n)\|_2^2 \left( -1 + \frac{1}{2} \alpha \sum_{i=1}^N \lambda_i y_i^2 \right), \end{aligned} \quad (10)$$

where  $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_N$  are eigenvalues of the Hessian matrix  $\nabla_\theta^2 L(\xi)$ , and  $\mathbf{y} = (y_1, \dots, y_N)$  is a normalized vector. When  $\{\theta_n\}_{n=0}^\infty$  reaches a local or global minimum, the Hessian matrix  $\nabla_\theta^2 L(\xi)$  is semi-positive definite and all  $\lambda_i \geq 0$  for all  $i = 1, \dots, N$ . Moreover, for the multi-scale solution  $u_H(x)$ , computational results show that many eigenvalues of  $\nabla_\theta^2 L(\xi)$  are very large (see Figure 1), i.e., stiff during gradient flow dynamics. As a result, it is very possible that  $L(\theta_{n+1}) - L(\theta_n) > 0$ , which implies that the GD method fails to decrease the loss. A similar computation approach (see the Appendix of Li, Chen, and Huang (2023)) shows that the loss decay of IGD is

$$\begin{aligned} & L(\theta_{n+1}) - L(\theta_n) \\ &= \alpha \|\nabla_\theta L(\theta_{n+1})\|_2^2 \left( -1 - \frac{1}{2} \alpha \sum_{i=1}^N \lambda_i y_i^2 \right), \end{aligned} \quad (11)$$

means that the loss will always decay regardless of the stiffness of the gradient flow dynamics of PINNs. In addition, the linear convergence rate of IGD is strictly proven in Section 3.

## Implementation of the IGD/ISGD Method

Although the IGD/ISGD method Eq.(4) looks simple and theoretically stable, one difficulty that can not be ignored is the implicitness of the nonlinear Eq.(4). It can also be expressed as the celebrated proximal point algorithm (Yin et al. 2018; Rockafellar 1976):

$$\theta_{n+1} = \arg \min_{\theta} \left\{ \frac{1}{2} \|\theta - \theta_n\|^2 + \alpha \cdot L(\theta) \right\}. \quad (12)$$

Hence, when  $\alpha$  is sufficiently small,  $\theta_{n+1}$  is approximately close to its previous updates  $\theta_n$  with the original loss as a regularizer. This sub-optimization task requires additional computation and brings difficulties for the whole optimization process.

To reduce the computational burden, we take a practical ‘‘ISGD,L-BFGS’’ (or ‘‘ISGD,Adam’’) optimizer for PINNs training with multi-scale solutions. Here ‘‘ISGD,L-BFGS’’ means that we first use ISGD with large learning rate for a certain number of iterations, and then switch to L-BFGS with default settings. In the sub-optimization problem (12), we also apply L-BFGS to compute  $\theta_{n+1}$ . The optimizer L-BFGS does not require learning rate, and the neural network is trained until convergence, so the number of iterations is also ignored for L-BFGS (Liu and Nocedal 1989). Here, the successful application of L-BFGS in ‘‘ISGD,L-BFGS’’ optimizer is that both the sub-optimization problem and the subsequent optimization problem have good initial point  $\theta_n$ , thus are easier for L-BFGS to achieve good convergence properties. The ‘‘ISGD,Adam’’ optimizer is to replace L-BFGS by Adam optimizer with default settings in the

---

**Algorithm 1** Practical “ISGD,Adam” optimization for the loss  $L(\theta)$  with stiff solutions

---

**Input:** initial  $\theta_0$ ; ISGD learning rate  $\alpha$  and maximum iterations  $K_0$ ; the inner Adam learning rate  $\gamma$  and maximum iterations  $K_1$ ; the outer Adam learning rate  $\eta$  and maximum iterations  $K_2$

**Output:** the optimized  $\theta^*$

```

1: Let  $n = 0$ .
2: while  $n < K_0$  do
3:   Let  $\tilde{\theta}_0 = \theta_n$  and  $k = 0$ .
4:   while  $k \leq K_1$  do
5:     Update  $\tilde{\theta}_{k+1} = \text{Adam} \left( \left\{ \frac{1}{2} \left\| \tilde{\theta} - \theta_n \right\|^2 + \alpha \cdot L(\tilde{\theta}) \right\} \Big|_{\tilde{\theta}=\tilde{\theta}_k}, \gamma \right)$ 
       and  $k \leftarrow k + 1$ 
6:   end while
7:   Update  $\theta_{n+1} = \tilde{\theta}_{K_1}$  and  $n \leftarrow n + 1$ 
8: end while
9: while  $K_0 \leq n < K_0 + K_2$  do
10:  Update  $\theta_{n+1} = \text{Adam}(L(\theta)|_{\theta=\theta_n}, \eta)$  and  $n \leftarrow n + 1$ 
11: end while
12: Denote  $\theta^* = \theta_{K_0+K_2}$ 
13: return the optimized  $\theta^*$ 

```

---

“ISGD,L-BFGS” optimizer when the parameters of PINNs are too large for the quasi-Hessian matrix computation. The details are illustrated in Algorithm 1.

### Training Dynamics Analysis of IGD/ISGD

In this section, we analyze the neural network training dynamics of our IGD/ISGD method. The technical proofs are given in the Appendix of Li, Chen, and Huang (2023).

**Quadratic Loss.** We show that randomly initialized IGD method with a constant positive step size converges to the global minimum at a linear rate. For simplicity of proof, we demonstrate a two-layer neural network with the quadratic loss functions. The global convergence property can be extended to an arbitrary  $N$ -layer neural network with quadratic loss with the technique introduced in Du et al. (2019). Formally, we consider a neural network of the following form:

$$u(\mathbf{W}, \mathbf{a}, \mathbf{x}) = \frac{1}{\sqrt{m}} \sum_{r=1}^m a_r \sigma(\mathbf{w}_r^T \mathbf{x}), \quad (13)$$

where  $\mathbf{x} \in \mathbb{R}^d$  is the input data,  $\mathbf{w}_r \in \mathbb{R}^d$  is the weight vector of the first layer,  $a_r \in \mathbb{R}$  is the weight vector of the output layer, and  $\sigma(z)$  is the activation function. We focus on the empirical risk minimization problem with a quadratic loss. Given a training data set  $\{(\mathbf{x}_i, y_i)\}_{i=1}^N$ , we minimize

$$L(\mathbf{W}, \mathbf{a}) = \sum_{i=1}^N \frac{1}{2} |y_i - u(\mathbf{W}, \mathbf{a}, \mathbf{x}_i)|^2. \quad (14)$$

For simplicity, we fix the second layer and apply the IGD method to optimize the first layer

$$\mathbf{W}(n+1) = \mathbf{W}(n) - \alpha \frac{\partial L(\mathbf{W}(n+1), \mathbf{a})}{\partial \mathbf{W}(n+1)}, \quad (15)$$

where  $\alpha > 0$  is the learning rate.

The training dynamics of  $u(\mathbf{W}(n), \mathbf{a}, \mathbf{x}_i)$  strongly relies on the Gram matrix  $\mathbf{H}(n+1)$  defined by

$$\mathbf{H}_{ij}(n+1) = \sum_{r=1}^m \left\langle \frac{\partial u_i(n+1)}{\partial \mathbf{w}_r}, \frac{\partial u_j(n+1)}{\partial \mathbf{w}_r} \right\rangle, \quad (16)$$

and its limit Gram matrix  $\mathbf{H}^\infty$  defined by

$$\mathbf{H}_{ij}^\infty = \mathbf{x}_i^T \mathbf{x}_j \mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \sigma'(\mathbf{w}^T \mathbf{x}_i) \sigma'(\mathbf{w}^T \mathbf{x}_j). \quad (17)$$

The positivity of  $\mathbf{H}^\infty$  is the key to prove convergence. We first state some technical assumptions.

**Assumption 1.** *The activation function  $\sigma(\cdot)$  is smooth, analytic, and is not a polynomial function. Moreover, both  $\sigma(\cdot)$  and its derivatives are Lipschitz continuous, i.e., there exists a constant  $C > 0$  such that  $|\sigma(0)| \leq C$  and for any  $z_1, z_2 \in \mathbb{R}$ ,*

$$\begin{aligned} |\sigma(z_1) - \sigma(z_2)| &\leq c|z_1 - z_2|, \\ |\sigma'(z_1) - \sigma'(z_2)| &\leq c|z_1 - z_2|. \end{aligned}$$

Here and below, we use the same constant  $C$  without confusion for simplicity to represent different constants independent of  $m, N, \lambda_{\min}(\mathbf{H}^\infty)$ .

**Assumption 2.** *No two input data are parallel, i.e., for any  $i \neq j$ , we need  $\mathbf{x}_i \neq c\mathbf{x}_j$  for any constant  $c$ .*

Now we present our main theorem.

**Theorem 1.** *Assume Assumption 1 and Assumption 2 hold and for all  $i \in [N]$ ,  $\|\mathbf{x}_i\|_2 \leq C$ ,  $|y_i| \leq C$ , and the hidden numbers  $m \geq \max \left\{ \frac{16CN^2}{\lambda_{\min}(\mathbf{H}^\infty)^2} \log\left(\frac{2N}{\delta}\right), \frac{16C^2N^4}{\lambda_{\min}(\mathbf{H}^\infty)^4} \right\}$ , and the learning rate  $\alpha \leq \frac{C\lambda_{\min}(\mathbf{H}^\infty)}{N^2}$  for some constant  $C$ , and we i.i.d. initialize  $\mathbf{w}_r \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ ,  $a_r \sim \text{unif}[-1, 1]$  for  $r \in [m]$ , then with probability  $1 - \delta$  we have for  $n = 0, 1, 2, \dots$*

$$L(n) \leq \left( \frac{1}{1 + \frac{\alpha\lambda_{\min}(\mathbf{H}^\infty)}{2}} \right)^n L(0). \quad (18)$$

where the quadratic loss  $L(n) \doteq L(\mathbf{W}(n), \mathbf{a})$  is defined by Eq.(14).

**PINN Loss.** For the PINN loss Eq.(6), it has been observed that the Gram matrix  $\mathbf{H}^\infty$  may not guarantee strict positivity (see Wang, Yu, and Perdikaris 2022, Figure 1), and the proof technique may fail. However, as demonstrated in the next section, the convergence and strong stability of IGD/ISGD for training PINNs are numerically verified.

### Computational Results

In this section, we compare the performance of SGD optimizer, Adam optimizer and our ISGD optimizer in training PINNs to solve different differential equations. The hyper-parameters used in the three optimizers are listed in Table 1. We note #Iterations =  $(K_0 \cdot K_1 + K_2) \cdot \text{batches}$ , where  $K_0, K_1, K_2$  are hyper-parameters in Algorithm 1. The wall-clock computational time is proportional to #Iterations, so the computational time is comparable for three optimizers in all numerical examples. More computational results are given in the Appendix of Li, Chen, and Huang (2023).

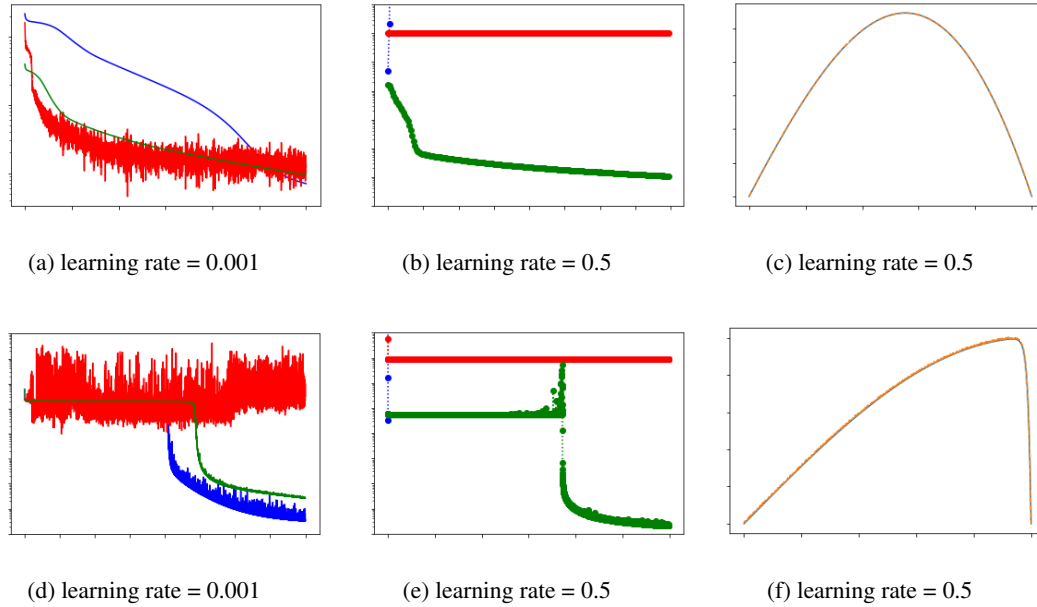


Figure 2: The training results of three optimizers SGD (blue line), Adam (red line) and ISGD (green line) for ODE Eq.(19). Top row:  $\epsilon = 2.0$ . Down row:  $\epsilon = 0.01$ .

Example	Optimizer	Learning rate	#Iterations
4.1 ( $\epsilon = 2$ )	SGD(Adam)	0.001	120,000
	ISGD, Adam	0.5, 0.001	102,000
4.1 ( $\epsilon = 0.01$ )	SGD(Adam)	0.001	400,000
	ISGD, Adam	0.5, 0.001	360,000
4.2	SGD(Adam)	0.0005	2,000,000
	ISGD, Adam	0.5, 0.0005	1,100,000
4.3	SGD(Adam)	0.0005	1,000,000
	ISGD, Adam	0.5, 0.0005	550,000

Table 1: Hyper-parameters used in the three optimizers for the following 3 examples. “SGD(Adam)” represents SGD shares the same hyper-parameters with Adam. “ISGD, Adam” is referred in Algorithm 1.

### PINN for Ordinary Differential Equations

Singularly perturbed ordinary differential equations have been successfully applied to many fields including gas dynamics, chemical reaction, fluid mechanics, elasticity, etc. To find the solution is a hot and difficult problem because it contains a very small parameter  $\epsilon$ . We consider the second-order linear singularly perturbed boundary value differential equation

$$\begin{cases} -\epsilon y''(x) + y'(x) = f(x), & x \in (0, 1), \\ y(0) = 0, & y(1) = 0. \end{cases} \quad (19)$$

The true solution is chosen as  $y(x) = \frac{1-\epsilon^{\frac{x}{\epsilon}}}{\epsilon^{\frac{1}{\epsilon}}-1} + \sin(\frac{\pi}{2}x)$ , and  $f(x)$  is given according to Eq.(19).  $\epsilon > 0$  is a constant; when  $\epsilon$  is very small, a boundary layer exists near the boundary  $x = 1$ . Let  $y_{\theta}(x)$  be the neural network approximation of

$y(x)$ , then the PINN loss function can be defined as


$$\begin{aligned} L(\theta) = & \frac{1}{2} [|y_{\theta}(0) - y(0)|^2 + |y_{\theta}(1) - y(1)|^2] \\ & + \frac{1}{N} \sum_{i=1}^N \left| -\epsilon y_{\theta}''(x_i) + y_{\theta}'(x_i) - f(x_i) \right|^2. \end{aligned}$$

We choose  $N = 400$  randomly sampled points to compute the loss function, a batch size of 40 for a small learning rate  $\alpha = 0.001$ , and a full batch size for a large learning rate  $\alpha = 0.5$ . A neural network with 4 hidden layers, every 50 units with  $\tanh$  activations, is applied in all the computations. The results are shown in Figure 2. For the case  $\epsilon = 2$ , the true solution is smooth. As shown in Fig. 2(a)(b), we find that the ISGD optimizer can significantly improve training convergence and remain stable for different learning rates. For the case  $\epsilon = 0.01$ , as shown in Fig. 2(f), the true solution has a boundary layer near  $x = 1$ , and the large gradient creates difficulties for the optimizers. As shown in Fig. 2(d)(e), more epochs and a smaller learning rate are required to be convergent for this singularity phenomenon. While the SGD and Adam optimizers are not convergent for large learning rates, the ISGD can still have stable convergent results, demonstrating the robustness of the proposed method.


### PINN for Poisson Equation

Poisson equation is an elliptic partial differential equation of broad utility in theoretical physics. We consider the Poisson equation on the domain  $\Omega = [0, 1] \times [0, 1]$

$$\begin{cases} -\frac{\partial^2 u}{\partial x^2} - \frac{\partial^2 u}{\partial y^2} = f(x, y), & (x, y) \in \Omega, \\ u(x, y) = 0, & (x, y) \in \partial\Omega. \end{cases} \quad (20)$$



(a) learning rate = 0.0005 (b) learning rate = 0.5



(c) prediction  $u_\theta(x, y)$  (d) error  $|u(x, y) - u_\theta(x, y)|$

Figure 3: PINN training for the Poisson equation (20).

The true solution is chosen as  $u(x, y) = \sin(\pi x) \sin(\pi y) + 0.1 \sin(10\pi x) \sin(10\pi y)$  with multi-scale features. The PINN loss function is defined as

$$L(\theta) = \frac{1}{N_b} \sum_{i=1}^{N_b} |u_\theta(x_i, y_i) - u(x_i, y_i)|^2 + \frac{1}{N_f} \sum_{j=1}^{N_f} \left| \frac{\partial^2 u_\theta(x_j, y_j)}{\partial x^2} + \frac{\partial^2 u_\theta(x_j, y_j)}{\partial y^2} - f(x_j, y_j) \right|^2.$$


We choose  $N_b = 400$  randomly sampled points on  $\partial\Omega$ , and  $N_f = 4,000$  randomly sampled points in  $\Omega$  to compute the loss function. A neural network with 6 hidden layers, every 100 units with  $\tanh$  activations, is applied in all the computations. The three optimizer training results for  $\alpha = 0.0005$  and 0.5 are shown in Fig. 3(a) and Fig. 3(b), respectively. We see that neither SGD nor Adam can train well as learning rate increases, but our ISGD trains well for different values of  $\alpha$ . The PINN prediction is plotted in Fig. 3(c), and the absolute error is shown in Fig. 3(d), with an absolute error less than 0.2%. We see that the PINN trained by the ISGD optimizer can obtain stable and accurate results for the Poisson equation (20).

### PINN for Helmholtz Equation


The Helmholtz equation is one of the fundamental equations of mathematical physics arising in many physical problems, such as vibrating membranes, acoustics, and electromagnetism equations. We solve the two-dimensional Helmholtz equation given by

$$\begin{cases} \frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2} + k^2 u(x, y) = f(x, y), & (x, y) \in \Omega, \\ u(x, y) = 0, & (x, y) \in \partial\Omega. \end{cases} \quad (21)$$

The exact solution for  $k = 4$  is  $u(x, y) = \sin(\pi x) \sin(4\pi y)$ , and the force term  $f(x, y)$  is given by the Eq.(21). We choose  $N_b = 400$  randomly sampled points on  $\partial\Omega$ , and



(a) learning rate = 0.0005 (b) learning rate = 0.5



(c) prediction  $u_\theta(x, y)$  (d) error  $|u(x, y) - u_\theta(x, y)|$

Figure 4: PINN training for Helmholtz equation (21).

$N_f = 4,000$  randomly sampled points in  $\Omega$  to compute the loss function. A neural network with 6 hidden layers, every 100 units with  $\tanh$  activations, is applied in all the computations. The three optimizer training results for  $\alpha = 0.0005$  and 0.5 are shown in Fig. 4(a) and Fig. 4(b), respectively. The PINN solution is plotted in Fig. 4(c), and the absolute error is shown in Fig. 4(d), with an absolute error less than 0.7%. We see that the PINN trained by the ISGD optimizer can obtain stable and accurate results for the Helmholtz equation (21).

### Conclusion

To overcome the numerical instability of traditional gradient descent methods to some key hyper-parameters, a stable IGD/ISGD method was proposed, analyzed and tested in this paper. The IGD/ISGD method includes implicit updates, and the L-BFGS or Adam optimizer can be combined to forward the updates. The global convergence of IGD/ISGD are theoretically analyzed and proven. We apply the IGD/ISGD method to train deep as well as physics-informed neural networks, showing that the IGD/ISGD method can effectively deal with stiffness phenomenon in the training dynamics via gradient descent. The techniques proposed in this paper stabilize the training of neural network models. This may result in making it easier for non-experts to train such models for beneficial applications, such as solving PDEs.

### Acknowledgments

The first author is supported by the National Natural Science Foundation of China (No.62106103), Fundamental Research Funds for the Central Universities (No.ILA22023) and 173 Program Technical Field Fund (No.2021-JCJQ-JJ-0018).

## References

- Allen-Zhu, Z.; Li, Y.; and Song, Z. 2019. A convergence theory for deep learning via over-parameterization. In *International Conference on Machine Learning*, 242–252. PMLR.
- Butcher, J. C. 2016. *Numerical methods for ordinary differential equations*. John Wiley & Sons.
- Cai, S.; Wang, Z.; Wang, S.; Perdikaris, P.; and Karniadakis, G. E. 2021. Physics-informed neural networks for heat transfer problems. *Journal of Heat Transfer*, 143(6): 060801.
- Chen, Y.; Lu, L.; Karniadakis, G. E.; and Dal Negro, L. 2020. Physics-informed neural networks for inverse problems in nano-optics and metamaterials. *Optics Express*, 28(8): 11618–11633.
- Du, S. S. 2019. *Gradient descent for non-convex problems in modern machine learning*. Ph.D. thesis, Carnegie Mellon University.
- Du, S. S.; Lee, J.; Li, H.; Wang, L.; and Zhai, X. 2019. Gradient descent finds global minima of deep neural networks. In *International Conference on Machine Learning*, 1675–1685. PMLR.
- Du, S. S.; Zhai, X.; Póczos, B.; and Singh, A. 2018. Gradient descent provably optimizes over-parameterized neural networks. In *International Conference on Learning Representations*.
- Duchi, J.; Hazan, E.; and Singer, Y. 2011. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(7): 2121–2159.
- Fang, Z.; and Zhan, J. 2019. Deep physical informed neural networks for metamaterial design. *IEEE Access*, 8: 24506–24513.
- Hennigh, O.; Narasimhan, S.; Nabian, M. A.; Subramaniam, A.; Tangsali, K.; Fang, Z.; Rietmann, M.; Byeon, W.; and Choudhry, S. 2021. NVIDIA SimNet™: An AI-accelerated multi-physics simulation framework. In *International Conference on Computational Science*, 447–461. Springer.
- Jagtap, A. D.; Kawaguchi, K.; and Em Karniadakis, G. 2020. Locally adaptive activation functions with slope recovery for deep and physics-informed neural networks. *Proceedings of the Royal Society A*, 476(2239): 20200334.
- Jagtap, A. D.; Kawaguchi, K.; and Karniadakis, G. E. 2020. Adaptive activation functions accelerate convergence in deep and physics-informed neural networks. *Journal of Computational Physics*, 404: 109136.
- Jagtap, A. D.; Shin, Y.; Kawaguchi, K.; and Karniadakis, G. E. 2022. Deep Kronecker neural networks: A general framework for neural networks with adaptive activation functions. *Neurocomputing*, 468: 165–180.
- Jin, X.; Cai, S.; Li, H.; and Karniadakis, G. E. 2021. NSFnets (Navier-Stokes flow nets): Physics-informed neural networks for the incompressible Navier-Stokes equations. *Journal of Computational Physics*, 426: 109951.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Lagaris, I. E.; Likas, A.; and Fotiadis, D. I. 1998. Artificial neural networks for solving ordinary and partial differential equations. *IEEE Transactions on Neural Networks*, 9(5): 987–1000.
- Li, M.; He, L.; and Lin, Z. 2020. Implicit Euler skip connections: Enhancing adversarial robustness via numerical stability. In *International Conference on Machine Learning*, 5874–5883. PMLR.
- Li, Y.; Chen, S.-C.; and Huang, S.-J. 2023. Implicit stochastic gradient descent for training physics-informed neural networks. *arXiv preprint arXiv:2303.01767*.
- Liu, D. C.; and Nocedal, J. 1989. On the limited memory BFGS method for large scale optimization. *Mathematical programming*, 45(1): 503–528.
- Lu, L.; Meng, X.; Mao, Z.; and Karniadakis, G. E. 2021. DeepXDE: A deep learning library for solving differential equations. *SIAM Review*, 63(1): 208–228.
- Mao, Z.; Jagtap, A. D.; and Karniadakis, G. E. 2020. Physics-informed neural networks for high-speed flows. *Computer Methods in Applied Mechanics and Engineering*, 360: 112789.
- Psichogios, D. C.; and Ungar, L. H. 1992. A hybrid neural network-first principles approach to process modeling. *AIChE Journal*, 38(10): 1499–1511.
- Raissi, M.; Perdikaris, P.; and Karniadakis, G. E. 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.
- Raissi, M.; Yazdani, A.; and Karniadakis, G. E. 2020. Hidden fluid mechanics: Learning velocity and pressure fields from flow visualizations. *Science*, 367(6481): 1026–1030.
- Rockafellar, R. T. 1976. Monotone operators and the proximal point algorithm. *SIAM Journal on Control and Optimization*, 14(5): 877–898.
- Sahli Costabal, F.; Yang, Y.; Perdikaris, P.; Hurtado, D. E.; and Kuhl, E. 2020. Physics-informed neural networks for cardiac activation mapping. *Frontiers in Physics*, 8: 42.
- Satpathi, S.; and Srikant, R. 2021. The dynamics of gradient descent for overparametrized neural networks. In *Learning for Dynamics and Control*, 373–384. PMLR.
- Tieleman, T.; and Hinton, G. 2012. Lecture 6.5-rmsprop, coursera: Neural networks for machine learning. *University of Toronto, Technical Report*.
- Toulis, P.; Airolidi, E.; and Rennie, J. 2014. Statistical analysis of stochastic gradient methods for generalized linear models. In *International Conference on Machine Learning*, 667–675. PMLR.
- Toulis, P.; and Airolidi, E. M. 2017. Asymptotic and finite-sample properties of estimators based on stochastic gradients. *Annals of Statistics*, 45(4): 1694–1727.
- Toulis, P.; Tran, D.; and Airolidi, E. 2016. Towards stability and optimality in stochastic gradient descent. In *Artificial Intelligence and Statistics*, 1290–1298. PMLR.

- Vo, N. D.; Hong, M.; and Jung, J. J. 2020. Implicit stochastic gradient descent method for cross-domain recommendation system. *Sensors*, 20(9): 2510.
- Wang, S.; Teng, Y.; and Perdikaris, P. 2021. Understanding and mitigating gradient flow pathologies in physics-informed neural networks. *SIAM Journal on Scientific Computing*, 43(5): A3055–A3081.
- Wang, S.; Wang, H.; and Perdikaris, P. 2021. On the eigenvector bias of Fourier feature networks: From regression to solving multi-scale PDEs with physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering*, 384: 113938.
- Wang, S.; Yu, X.; and Perdikaris, P. 2022. When and why PINNs fail to train: A neural tangent kernel perspective. *Journal of Computational Physics*, 449: 110768.
- Yin, P.; Pham, M.; Oberman, A.; and Osher, S. 2018. Stochastic backward Euler: An implicit gradient descent algorithm for k-means clustering. *Journal of Scientific Computing*, 77(2): 1133–1146.
- Zou, D.; Cao, Y.; Zhou, D.; and Gu, Q. 2020. Gradient descent optimizes over-parameterized deep ReLU networks. *Machine Learning*, 109(3): 467–492.