

DEEP NEURAL NETWORKS FOR SOLVING LARGE LINEAR SYSTEMS ARISING FROM HIGH-DIMENSIONAL PROBLEMS*

YIQI GU[†] AND MICHAEL K. NG[‡]

Abstract. This paper studies deep neural networks for solving extremely large linear systems arising from high-dimensional problems. Because of the curse of dimensionality, it is expensive to store both the solution and right-hand side vector in such extremely large linear systems. Our idea is to employ a neural network to characterize the solution with many fewer parameters than the size of the solution under a matrix-free setting. We present an error analysis of the proposed method, indicating that the solution error is bounded by the condition number of the matrix and the neural network approximation error. Several numerical examples from partial differential equations, queuing problems, and probabilistic Boolean networks are presented to demonstrate that the solutions of linear systems can be learned quite accurately.

Key words. very large scale linear systems, neural networks, partial differential equations, Riesz fractional diffusion, overflow queuing model, probabilistic Boolean networks

MSC codes. 65F10, 65F50, 65N22, 68T07, 60K25

DOI. 10.1137/22M1488132

1. Introduction. Linear equations appear widely in applied problems such as partial differential equations (PDEs) and numerical optimization. In physical problems, one usually needs to compute some physical quantity, such as temperature distribution and fluid velocity. Let us suppose the problem is addressed in a d -dimensional domain $\Omega \subset \mathbb{R}^d$, in which a grid Γ is set up. A grid function u on Γ is thereafter introduced to approximate the physical quantity. Then, after using a spatial discretization on the problem, u is computed through the following system of linear equations,

$$(1.1) \quad \mathbf{A}u = \mathbf{b},$$

where $\mathbf{A} \in \mathbb{R}^{\widetilde{M} \times \widetilde{N}}$ ($\widetilde{M}, \widetilde{N} \in \mathbb{N}^+$) is a nonsingular matrix; $\mathbf{b} \in \mathbb{R}^{\widetilde{M}}$ is a given vector; $\mathbf{u} = \text{vec}(u) \in \mathbb{R}^{\widetilde{N}}$ is the vector representation of u .

1.1. Existing methods and difficulties. Traditional linear solvers, including direct and iterative methods, have been extensively studied for a long time. Let us consider the high-dimensional problem that Ω is a d -dimensional box and (1.1) is assembled with tensor product structure on an $N \times N \times \cdots \times N$ (d times) grid, where $N \in \mathbb{N}^+$ is the degree of freedom in every dimension. A series of methods have been developed for the linear system with tensor product structure, such as the Krylov subspace method [31], the projection method [4], and Cayley transformation [21]. However, for such linear systems, the numbers of equations and unknowns

*Submitted to the journal's Methods and Algorithms for Scientific Computing section April 1, 2022; accepted for publication (in revised form) April 27, 2023; published electronically September 21, 2023.

<https://doi.org/10.1137/22M1488132>

Funding: This work was supported by Hong Kong Research Grant Council GRF 12300519, 17201020, 17300021, Hong Kong Research Grant Council CRF C1013-21GF, C7004-21GF, and Joint NSFC-RGC N-HKU76921.

[†]School of Mathematical Sciences, University of Electronic Science and Technology of China, Sichuan, 611731, China (yiqigu@uestc.edu.cn).

[‡]Department of Mathematics, Hong Kong Baptist University, Kowloon Tong, Hong Kong (michael-ng@hkbu.edu.hk).

satisfy $\widetilde{M} = \widetilde{N} = N^d$, which are extremely large even if d is moderately large. In spite of some development concerning large-scale problems [3, 55, 41], the practical performance of many existing methods is still limited by the dimension. For high-dimensional problems with larger d , the system size N^d might exceed the machine storage so that even the intermediate solution cannot be stored entirely in memory.

1.2. Motivations and contributions. In recent years, the theory and applications of neural networks (NNs) have been widely studied in a variety of areas, including computer science and applied mathematics. Generally speaking, NNs are a type of function with a nonlinear parametric structure. It has been found in a series of literature that NNs can approximate common functions effectively. In the pioneering work [13, 26, 6], the universal approximation theory of two-layer shallow NNs is discussed. In recent research, quantitative information about the NN approximation error was presented for various types of functions, e.g., smooth functions [37, 34, 52, 19, 39, 50, 17, 18], piecewise smooth functions [42], band-limited functions [40], and continuous functions [53, 44, 45].

One remarkable property of NNs is the capability to approximate high-dimensional functions. Many traditional approximation structures, such as finite elements and polynomials, suffer from the curse of dimensionality. Specifically, when approximating a function of d -dimensional variables, their error bounds will be $O(J^{-\alpha/d})$, where J is the number of free parameters and α characterizes the regularity of the function. However, NNs can avoid such issues for some special function spaces. A typical example is the Barron space, for which the NN approximation error is either independent of d or increasing with d very slowly [5, 6, 30, 20, 16, 48, 49, 9, 18]. So far, NNs have been applied successfully in solving high-dimensional PDEs and inverse problems [23, 43, 51, 54, 28]. For high-dimensional PDEs, one successful application is the physics-informed neural networks (PINNs) [43]. More precisely, the PDE solution is approximated by a general NN, which is then trained through the minimization of the PDE residual (e.g., the PINN) in the least squares sense.

In this work, we follow the idea of NN approximation for high-dimensional functions adopted in previous applications (e.g., PINNs) and propose a novel NN-based method to solve linear systems (1.1). Specifically, we realize that in many real-world problems, the physical quantities are continuously distributed. Our method comes from the fact that if the physical quantity v is smooth enough in Ω , NNs with only a small number of parameters can approximate v with the desired accuracy. Meanwhile, if the true solution \mathbf{u} of (1.1) is a good approximation to v , it can also be approximated accurately by such NNs. This allows us to take an NN ϕ to characterize \mathbf{u} with many fewer parameters than $\widetilde{N} = N^d$. In this method, the unknown numeric vector \mathbf{u} in (1.1) is replaced with a vector function whose variables are NN parameters. Then the task of solving the large linear system is transformed into solving a new small nonlinear system. In the proposal, the new system is solved by a least squares method under a deep learning framework. This approach is able to solve linear systems of very large sizes that may be difficult for existing methods. Error analysis is also conducted for this method, provided that the true physical quantity is in the Barron space.

Several typical problems are solved using the proposed method in numerical experiments. The first problem is the tensor-structured linear system derived from Poisson's equation using the centered finite difference scheme, in which the system with $(N, d) = (10^4, 6)$ is solved effectively. We mention that our method applied to this example is equivalent to PINNs, except for a few differences in the setting. Next, we consider linear systems derived from Riesz fractional diffusion [27], a fractional

differential equation. Beyond PDEs, we apply our method to the discrete problem, the overflow queuing problem [10, 11, 12], where we succeed in solving such linear systems derived from 10 dimensions (i.e., 10 queues), while earlier work can address at most three dimensions numerically. Finally, as the last example, we solve a $2^d \times 2^d$ sparse system from probabilistic Boolean networks. Our method successfully solves a problem with 100 dimensions and $O(10^{30})$ nonzero entries, while previous work [33] addresses merely 30 dimensions and $O(10^4)$ nonzero entries.

1.3. Organization of the paper. This paper is organized as follows. In section 2, we review the fully connected NNs, the conceptual NN-based method, and the practical algorithms with gradient descent. In section 3, we estimate the error of the approximate solution under the Barron space hypothesis. Several examples of physical problems are demonstrated in section 4 to test the performance of the proposed method. Conclusions and discussions about further research work are provided in section 5.

2. NN-based method. In this section, we introduce the concepts of NNs and explain how to use NNs to approximate the solution of linear systems. In this paper, we use bold fonts to denote matrices and vectors.

2.1. Fully connected neural network. Among the many types of NNs, the fully connected neural network (FNN) is the most basic and commonly used in applied mathematics. Mathematically speaking, given $L \in \mathbb{N}^+$ and $M_\ell \in \mathbb{N}^+$ for $\ell = 1, \dots, L-1$, where \mathbb{N}^+ denotes the set of positive integers, we define the simple nonlinear function $h_\ell : \mathbb{R}^{M_{\ell-1}} \rightarrow \mathbb{R}^{M_\ell}$ given by

$$(2.1) \quad h_\ell(\mathbf{x}_\ell) := \sigma(\mathbf{W}_\ell \mathbf{x}_\ell + \mathbf{b}_\ell),$$

where $\mathbf{W}_\ell \in \mathbb{R}^{M_\ell \times M_{\ell-1}}$; $\mathbf{b}_\ell \in \mathbb{R}^{M_\ell}$; $\sigma(\mathbf{y})$ is a given function which is applied entrywise to a vector \mathbf{y} to obtain another vector of the same size, called an activation function. Common activation functions include rectified linear unit (ReLU) $\max\{0, y\}$ and the sigmoidal function $(1 + e^{-y})^{-1}$.

Set $M_0 = d$, the dimension of the input variable; then an FNN $\phi : \mathbb{R}^d \rightarrow \mathbb{R}$ is formulated as the composition of these $L-1$ simple nonlinear functions, namely

$$(2.2) \quad \phi(\mathbf{x}; \theta) = \mathbf{a}^\top h_{L-1} \circ h_{L-2} \circ \cdots \circ h_1(\mathbf{x}) \quad \text{for } \mathbf{x} \in \mathbb{R}^d,$$

where $\mathbf{a} \in \mathbb{R}^{M_{L-1}}$ and $\theta := \{\mathbf{a}, \mathbf{W}_\ell, \mathbf{b}_\ell : 1 \leq \ell \leq L-1\}$ denotes the set of all free parameters. Here M_ℓ is called the width of the ℓ th layer, and L is called the depth. The widths and depth characterize the architecture of an FNN. So, in fixing σ , L , and $\{M_\ell\}_{\ell=1}^{L-1}$, the FNN architecture is completely determined, but the parameters in θ are still free. In the following passage, for simplicity, we only consider the architecture with fixed width $M_\ell = M$ for all $\ell = 1, \dots, L-1$. We use $\mathcal{F}_{L,M,\sigma}$ to denote the set of all FNNs with depth L , width M , and activation function σ .

We can calculate the number of scalar parameters in θ . It is clear that the input layer with $\ell = 1$ has $(d+1)M$ scalars, the output layer \mathbf{a} has M scalars, and other hidden layers with $\ell = 2, \dots, L-1$ have a total of $(L-2)M(M+1)$ scalars. So $|\theta| = (d+1)M + M + (L-2)M(M+1) = (L-2)M^2 + (d+L)M$.

2.2. Problem description. Let us describe the linear system arising from d -dimensional problems with a physical or conceptual domain Ω . As a typical example, we assume that Ω is the d -dimensional unit box $[0, 1]^d$, and a Cartesian grid is set up on Ω . The following discussion can be easily generalized for other domain geometries

and grid settings. Suppose one aims to determine an unknown real function $v(\mathbf{x})$ for $\mathbf{x} \in \Omega$ from a high-dimensional physical problem. One common way is setting a grid on Ω and determining v on every grid point. Specifically, we let $N \in \mathbb{N}^+$ (the set of positive integers) and let $0 \leq x_1 < \dots < x_N \leq 1$ be a one-dimensional grid in $[0, 1]$. We use the vector of the form $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_d)$, where each component α_i is an integer in $[1, N]$, to denote a multi-index. Also, we define $\Lambda = \{\boldsymbol{\alpha} : 1 \leq \alpha_i \leq N \text{ for } i = 1, \dots, d\}$ as the set of all multi-indices. Then, for any $\boldsymbol{\alpha}$, the column vector $\mathbf{x}_{\boldsymbol{\alpha}} := [x_{\alpha_1} \dots x_{\alpha_d}]^\top$ represents a Cartesian grid point in Ω , and

$$(2.3) \quad \Gamma := \{\mathbf{x}_{\boldsymbol{\alpha}} : \boldsymbol{\alpha} \in \Lambda\}$$

is the set of all Cartesian grid points. It is clear that $|\Gamma| = N^d$.

We use a real number $u_{\boldsymbol{\alpha}}$ to approximate $v(\mathbf{x}_{\boldsymbol{\alpha}})$. By using computational methods (e.g., finite difference methods in solving differential equations), we can derive a linear system concerning $\{u_{\boldsymbol{\alpha}} : \boldsymbol{\alpha} \in \Lambda\}$ from the original physical problem, namely

$$(2.4) \quad \sum_{\boldsymbol{\alpha} \in \Lambda} a_{m, \boldsymbol{\alpha}} u_{\boldsymbol{\alpha}} = b_m \quad \text{for } m = 1, \dots, \widetilde{M},$$

where $a_{m, \boldsymbol{\alpha}} \in \mathbb{R}$ is the coefficient of $u_{\boldsymbol{\alpha}}$ in the m th equation (cf. (1.1)) and \widetilde{M} is the number of equations. Many linear systems arising from practical problems are matrix-free such that one can directly get the value of $a_{m, \boldsymbol{\alpha}}$ from $\boldsymbol{\alpha}$ and m instead of loading it from the storage. In the following discussion, we only consider the linear systems that are matrix-free. Moreover, we assume that the right-hand side $\{b_m\}$ can be obtained instantly for specified m when the linear system is being solved, and we do not need to save the entire right-hand side in the storage. For example, in solving differential equations, $\{b_m\}$ are the values of a given function at grid points, which can be computed in real time for m in a small subset of $\{1, \dots, \widetilde{M}\}$. This assumption allows us to implement the memory-saving algorithm proposed in section 2.5.

2.3. A conceptual method. In most cases, the number of grid points N in every dimension is set large for high resolution. Therefore, one difficulty of solving (2.4) is its possibly large size N^d when d is moderately large. At the very worst, N^d exceeds the memory limit, and even a vector in \mathbb{R}^{N^d} cannot be saved entirely in memory. For instance, if one sets $N = 10$ grid points in every dimension on a machine with 32G memory, the bytes of an N^d double-precision vector will exceed the memory limit when $d \geq 10$. This situation forbids many classical linear solvers, including the matrix-free iterative methods. We will propose an NN representation of the unknowns $\{u_{\boldsymbol{\alpha}}\}$, which can be viewed as an approximation of the unknown vector with fewer free elements and hence costs much less storage.

Since the linear system (2.4) is derived from a physical problem having a smooth unknown function v , it is expected that the unknowns $\{u_{\boldsymbol{\alpha}}\}$ are also distributed smoothly on Γ . Namely, the grid mapping $\chi : \Gamma \rightarrow \mathbb{R}^{N^d}$ defined by $\chi(\mathbf{x}_{\boldsymbol{\alpha}}) = u_{\boldsymbol{\alpha}}$ is spatially smooth (it means the data $\{\mathbf{x}_{\boldsymbol{\alpha}}, u_{\boldsymbol{\alpha}}\}_{\boldsymbol{\alpha} \in \Gamma}$ can be fit by a function with few high-frequency components). Thanks to the good approximation property for high-dimensional functions (see, e.g., [30, 48, 49, 9, 18, 45, 36, 46, 47]), NNs can be employed to serve as the functioning of χ . Specifically, we introduce an NN $\phi : \mathbb{R}^d \rightarrow \mathbb{R}$ and let $\phi(\mathbf{x}_{\boldsymbol{\alpha}}; \theta)$ be an approximation of $u_{\boldsymbol{\alpha}}$ for all $\boldsymbol{\alpha} \in \Lambda$. By this setting, the linear system (2.4) can be formulated into

$$(2.5) \quad \sum_{\boldsymbol{\alpha} \in \Lambda} a_{m, \boldsymbol{\alpha}} \phi(\mathbf{x}_{\boldsymbol{\alpha}}; \theta) = b_m \quad \text{for } m = 1, \dots, \widetilde{M},$$

where the NN parameter set θ is the unknown. Note that (2.5) is actually a nonlinear system of equations due to the nonlinear structure of NNs. And the number of unknowns $|\theta| = O(M^2L + Md)$ is essentially different from N^d , the number of unknowns of the original linear system (2.4). In high-dimensional cases, the former number can be much smaller (see Remark 3.1).

For simplicity of notation, we can also formulate (2.4) and (2.5) as a matrix-vector form. Without loss of generality, we assume the unknowns $\{u_\alpha\}$ in (2.4) are ordered lexicographically in a column vector, namely

$$(2.6) \quad \mathbf{u} := [u_{(1,1,\dots,1)} \ u_{(1,1,\dots,2)} \ \cdots \ u_{(N,N,\dots,N)}]^\top \in \mathbb{R}^{N^d},$$

and we let

$$(2.7) \quad \mathbf{A} := \begin{bmatrix} a_{1,(1,1,\dots,1)} & a_{1,(1,1,\dots,2)} & \cdots & a_{1,(N,N,\dots,N)} \\ \vdots & \vdots & \ddots & \vdots \\ a_{\widetilde{M},(1,1,\dots,1)} & a_{\widetilde{M},(1,1,\dots,2)} & \cdots & a_{\widetilde{M},(N,N,\dots,N)} \end{bmatrix} \in \mathbb{R}^{\widetilde{M} \times N^d},$$

$$(2.8) \quad \mathbf{b} := [b_1 \ \cdots \ b_{\widetilde{M}}]^\top \in \mathbb{R}^{\widetilde{M}};$$

then (2.4) can be written as

$$(2.9) \quad \mathbf{A}\mathbf{u} = \mathbf{b}.$$

Similarly, we let

$$(2.10) \quad \boldsymbol{\phi}(\theta) := [\phi(\mathbf{x}_{(1,1,\dots,1)}; \theta) \ \phi(\mathbf{x}_{(1,1,\dots,2)}; \theta) \ \cdots \ \phi(\mathbf{x}_{(N,N,\dots,N)}; \theta)]^\top \in \mathbb{R}^{N^d};$$

then (2.5) can be written as

$$(2.11) \quad \mathbf{A}\boldsymbol{\phi}(\theta) = \mathbf{b}.$$

Now we solve the nonlinear system (2.11) with $O(M^2L + Md)$ unknowns instead of the original linear system (2.9) with N^d unknowns. And the vector $\boldsymbol{\phi}(\theta)$ is an approximation of the original solution vector \mathbf{u} . Usually, the system (2.11) does not have an exact solution θ . So we will find the least squares solution of (2.11) through the following optimization framework:

$$(2.12) \quad \min_{\theta} L(\theta) = \frac{1}{\widetilde{M}} \|\mathbf{A}\boldsymbol{\phi}(\theta) - \mathbf{b}\|_2^2.$$

The loss function L can be decreased via gradient descent methods under the NN learning framework.

To measure the error vector of large linear systems in a fair way, one usually uses the ℓ^2 -norm, which does not increase with the vector size, instead of the Euclidean norm $\|\cdot\|_2$. Specifically, for $\mathbf{u} \in \mathbb{R}^{N^d}$, we define

$$(2.13) \quad \|\mathbf{u}\|_{\ell^2} := \sqrt{\frac{\sum_{\alpha \in \Lambda} |u_\alpha|^2}{N^d}} = N^{-d/2} \|\mathbf{u}\|_2.$$

If a feasible solution θ_0 of (2.12) is found, it satisfies $L(\theta_0) = \frac{1}{\widetilde{M}} \|\mathbf{A}\boldsymbol{\phi}(\theta_0) - \mathbf{b}\|_2^2$, and then the error between $\boldsymbol{\phi}(\theta_0)$ and the true solution \mathbf{u} of (2.9) is estimated by

$$(2.14) \quad \begin{aligned} \|\boldsymbol{\phi}(\theta_0) - \mathbf{u}\|_{\ell^2} &\leq N^{-d/2} \|\boldsymbol{\phi}(\theta_0) - \mathbf{u}\|_2 \leq N^{-d/2} \|\mathbf{A}^{-1}\|_2 \|\mathbf{A}(\boldsymbol{\phi}(\theta_0) - \mathbf{u})\|_2 \\ &= N^{-d/2} \|\mathbf{A}^{-1}\|_2 \|\mathbf{A}\boldsymbol{\phi}(\theta_0) - \mathbf{b}\|_2 \leq \|\mathbf{A}^{-1}\|_2 \sqrt{L(\theta_0)}, \end{aligned}$$

provided that \mathbf{A} is square ($\widetilde{M} = N^d$) and invertible. Therefore, the above method finds an approximate solution with the error bounded by the norm of inverse matrix and the resulting minimized loss.

2.4. Nonunique solutions. If the linear system has more than one solution, the optimization (2.12) may not locate the particular solution we are looking for. One typical problem is the computation of the nontrivial solutions of homogeneous systems. For example, the eigenvectors of \mathbf{A} can be computed through $(\lambda \mathbf{I} - \mathbf{A})\mathbf{u} = 0$ given the eigenvalue λ . Another example is the computation of probability distribution in the overflow queuing problem [10, 11].

Now we assume that the linear system

$$(2.15) \quad \mathbf{A}\mathbf{u} = 0$$

admits nontrivial solutions, and the nullspace is exactly one-dimensional. If we solve (2.15) via the unconstrained optimization (2.12), sometimes the trivial solution $\mathbf{u} = 0$ will be obtained. For instance, due to the implicit regularization of NNs [8], gradient descent in deep learning will probably converge to the smoothest solution, i.e., the zero solution. As we do not wish to admit the solution $\mathbf{u} = 0$, we set constraints on the solution. One simple way is requiring $\|\mathbf{u}\|_p = 1$ with some $p \in [1, \infty]$, and then following (2.12) a penalized optimization for (2.15) is given by

$$(2.16) \quad \min_{\theta} L(\theta) = \frac{1}{M} \|\mathbf{A}\phi(\theta)\|_2^2 + \epsilon^{-1} (\|\phi(\theta)\|_p - 1)^2,$$

where $\epsilon > 0$ is a penalty parameter.

A simpler way is fixing one component of \mathbf{u} , say $u_{\alpha_0} = 1$ for some multi-index α_0 , if u_{α_0} is known to be nonzero in advance. In this case, the penalized optimization for (2.15) is as follows:

$$(2.17) \quad \min_{\theta} L(\theta) = \frac{1}{M} \|\mathbf{A}\phi(\theta)\|_2^2 + \epsilon^{-1} (\phi(\mathbf{x}_{\alpha_0}; \theta) - 1)^2.$$

2.5. Mini-batch gradient descent algorithm. We now describe a class of practical algorithms for the optimization of (2.12). Algorithms for solving (2.16) and (2.17) can be derived in similar ways.

We rewrite \mathbf{A} row-wise as follows:

$$(2.18) \quad \mathbf{A} = [\mathbf{a}_1 \cdots \mathbf{a}_{\widetilde{M}}]^\top,$$

where $\mathbf{a}_k \in \mathbb{R}^{N^d}$ is a row of \mathbf{A} for $k = 1, 2, \dots, \widetilde{M}$. Then the optimization (2.12) can be rewritten as

$$(2.19) \quad \min_{\theta} L(\theta) = \frac{1}{\widetilde{M}} \sum_{k=1}^{\widetilde{M}} |\mathbf{a}_k^\top \phi(\theta) - b_k|^2.$$

To decrease L , a gradient descent method will update θ in every iteration by

$$(2.20) \quad \theta \leftarrow \theta - \tau \nabla_{\theta} L(\theta)$$

with the gradient

$$(2.21) \quad \nabla_{\theta} L(\theta) = \frac{2}{\widetilde{M}} \sum_{k=1}^{\widetilde{M}} (\mathbf{a}_k^\top \phi(\theta) - b_k) \cdot \mathbf{J}[\phi(\theta)]^\top \mathbf{a}_k,$$

where

$$(2.22) \quad \mathbf{J}[\phi(\theta)] = [\nabla_{\theta} \phi(\mathbf{x}_{(1,1,\dots,1)}; \theta) \quad \nabla_{\theta} \phi(\mathbf{x}_{(1,1,\dots,2)}; \theta) \quad \dots \quad \nabla_{\theta} \phi(\mathbf{x}_{(N,N,\dots,N)}; \theta)]^\top$$

is the Jacobian matrix of $\phi(\theta)$ and $\tau > 0$ is some adaptive learning rate.

However, it is sometimes undesirable to use all N^d terms in (2.21) due to computational expense. So, in practice, one can use mini-batch gradient descent by selecting a small batch of all terms for training. More precisely, in every iteration, a small subset \mathcal{S} is selected from $\{1, \dots, \widetilde{M}\}$ according to some principles (i.e., random sampling), and θ is then updated by

$$(2.23) \quad \theta \leftarrow \theta - \tau \nabla_{\theta} L_{\mathcal{S}}(\theta),$$

where $L_{\mathcal{S}}(\theta) = \frac{1}{|\mathcal{S}|} \sum_{k \in \mathcal{S}} |\mathbf{a}_k^{\top} \phi(\theta) - b_k|^2$. This algorithm is known as mini-batch gradient descent and is shown in Algorithm 1.

We remark that for linear systems with moderately small sizes, choosing $\mathcal{S} = \{1, \dots, \widetilde{M}\}$ is tolerable in the sense of computational cost. In this case, Algorithm 1 computes not only the matrix-vector products but also the NN-related quantities $\phi(\theta)$ and $\mathbf{J}[\phi(\theta)]$. Therefore, this method might be computationally more expensive than traditional iterative methods, which only conduct matrix-vector multiplications, for small-scale linear systems.

However, Algorithm 1 is able to address very large scale linear systems that traditional methods may not handle. On the one hand, in every iteration, we choose a small subset (indexed by \mathcal{S}) of all equations for computation, so only those matrix rows and vector entries that are necessary for the current batch are assessed. We do not need to save the entire matrix or vector in memory (especially in the case that \mathbf{A} is matrix-free and b_k can be computed in real time for any k). On the other hand, to save the intermediate solution, it suffices to save θ in memory, whose size can be much smaller than the number of entries of \mathbf{u} . And we only need to compute a few entries of $\phi(\theta)$ at the position where \mathbf{a}_k is nonzero, rather than the entire $\phi(\theta)$.

The complexity of every iteration in the `while`-loop can be estimated. For a special problem, the dimension d is always fixed, so we do not involve it in the estimation. In usual situations, each b_k can be evaluated with complexity that is independent of L , M , and N (e.g., evaluating a given function at some grid point), so the evaluation of b_k for all $k \in \mathcal{S}$ costs $O(|\mathcal{S}|)$ FLOPS. Next, we use N_{nz} to denote the maximal number of nonzero entries of \mathbf{a}_k for all $k \in \{1, \dots, \widetilde{M}\}$. For example, the matrix derived from the centered finite difference scheme on Poisson's equation satisfies $N_{\text{nz}} = 2d + 1$ (see section 4.1). By a simple calculation, we can derive that it costs $O(LM^2)$ FLOPS to compute $\phi(\mathbf{x}; \theta)$ or $\nabla_{\theta} \phi(\mathbf{x}; \theta)$ for each point $\mathbf{x} \in \mathbb{R}^d$, namely doing forward and back propagations of NNs. Then the complexity of computing the vector $\phi(\theta)$ or

Algorithm 1 NN-based mini-batch gradient descent to solve the linear system $\mathbf{A}\mathbf{u} = \mathbf{b}$

Require: hyperparameters $L, M, \sigma, \{\tau_i\}$; initial guess θ_0 .

Output: an approximate solution $\phi(\theta) \approx \mathbf{u}$.

initialize $\phi(\mathbf{x}; \theta) \in \mathcal{F}_{L, M, \sigma}$ with $\theta \leftarrow \theta_0$

$i \leftarrow 1$

while stopping criterion is not satisfied **do**

 generate $\mathcal{S} \subset \{1, \dots, \widetilde{M}\}$

 evaluate b_k for $k \in \mathcal{S}$

$\theta \leftarrow \theta - \frac{2\tau_i}{|\mathcal{S}|} \sum_{k \in \mathcal{S}} (\mathbf{a}_k^{\top} \phi(\theta) - b_k) \cdot \mathbf{J}[\phi(\theta)]^{\top} \mathbf{a}_k$

$i \leftarrow i + 1$

end while

return $\phi(\theta)$

the matrix $\nabla_{\theta} \phi(\theta)$ is $O(|\theta|LM^2) = O(L^2M^4)$. Hence, the complexity of computing $\mathbf{a}_k^{\top} \phi(\theta)$ and $\nabla_{\theta} \phi(\theta)^{\top} \mathbf{a}_k$ for all $k \in \mathcal{S}$ is $O(|\mathcal{S}|N_{\text{nz}}L^2M^4)$. Therefore, every iteration in the `while`-loop costs $O(|\mathcal{S}|N_{\text{nz}}L^2M^4)$ FLOPS.

It is worth mentioning that the optimization (2.12) can also be solved by other mini-batch gradient-based optimizers (e.g., Adam [29]). The algorithms using these optimizers can be developed like Algorithm 1. Besides, we remark that the proposed method is not limited to linear systems with tensor product structures. In Algorithm 1, \mathbf{A} is a general sparse matrix. We do not require \mathbf{A} to have any more special structures (e.g., banded or Toeplitz matrices). This is essentially different from some existing methods for large-scale linear systems that rely on special properties of the matrix.

3. Error analysis. We will give an error analysis of the proposed method in this section. The method is developed based on the approximation property of NNs for smooth functions. Therefore, the analysis should depend on the regularity hypothesis of the target function v . A series of recent literatures [37, 34, 52, 19, 39, 50, 17, 42, 40, 53, 44, 45] have developed many results of the NN approximation theory. Here we consider the NN approximation for Barron-type functions, which is studied extensively in [20, 16, 48, 49, 9, 18]. The authors show that the approximation error of two-layer FNNs for Barron-type functions is independent of the dimension or increases with it very slowly, hence overcoming the curse of dimensionality. Among various types of Barron spaces, we specifically use the definition described in [18], which corresponds to infinitely wide two-layer ReLU FNNs. The definition and properties of the Barron space/functions will be introduced in this section. Without loss of generality, it is still assumed that $\Omega = [0, 1]^d$ in the following discussion.

As discussed in the section 2.1, the class of two-layer (i.e., $L = 2$) ReLU FNNs can be reformulated as follows:

$$(3.1) \quad \mathcal{F}_{2,M,\text{ReLU}} = \left\{ \phi : \phi(\mathbf{x}) = \frac{1}{M} \sum_{i=1}^M a_i \sigma(\mathbf{b}_i^{\top} \mathbf{x} + c_i) \quad \forall (a_i, \mathbf{b}_i, c_i) \in \mathbb{R} \times \mathbb{R}^d \times \mathbb{R} \right\}.$$

Without ambiguity, we specify $\sigma(y) = \max\{0, y\}$ being the ReLU activation function throughout this section.

We consider functions $f_{\pi} : \mathbb{R}^d \rightarrow \mathbb{R}$ that admit the following representation:

$$(3.2) \quad f(\mathbf{x}) = \int_{\Omega'} a \sigma(\mathbf{b}^{\top} \mathbf{x} + c) \pi(da, d\mathbf{b}, dc) = \mathbb{E}_{\pi}[a \sigma(\mathbf{b}^{\top} \mathbf{x} + c)] \quad \forall \mathbf{x} \in \mathbb{R}^d,$$

where $\Omega' = \mathbb{R} \times \mathbb{R}^d \times \mathbb{R}$ and π is a probability distribution on $(\Omega, \Sigma_{\Omega'})$, with $\Sigma_{\Omega'}$ being a Borel σ -algebra on Ω' . This representation can be seen as a continuum analogue of the two-layer ReLU FNNs in $\mathcal{F}_{2,M,\text{ReLU}}$ as $M \rightarrow \infty$. We remark that in general, there are more than one π 's such that (3.2) is satisfied.

Now let us introduce the Barron space and its norm with respect to $\mathcal{F}_{2,M,\text{ReLU}}$. For functions that admit the representation (3.2), its Barron norm is defined as

$$(3.3) \quad \|f\|_{\mathcal{B}} := \inf_{\pi} \left(\int_{\Omega'} a^2 (\|\mathbf{b}\|_1 + |c|)^2 \pi(da, d\mathbf{b}, dc) \right)^{1/2} = \inf_{\pi} \left(\mathbb{E}_{\pi}[a^2 (\|\mathbf{b}\|_1 + |c|)^2] \right)^{1/2},$$

where the infimum is taken over all π such that (3.2) holds for all $\mathbf{x} \in \mathbb{R}^d$. The infimum of the empty set is considered as $+\infty$. The set of all functions with finite Barron norm is denoted by \mathcal{B} . Note that \mathcal{B} equipped with the Barron norm is shown to be a Banach

space that is called a Barron space [18]. Some examples of Barron functions are given in [6], including Gaussian density, positive definite functions, smooth functions with high-order derivatives, etc. (see [18] for a mathematical connection between the Barron definitions in [6] and in this paper). The following result characterizes the approximation error of NNs in $\mathcal{F}_{2,M,\text{ReLU}}$ for functions in \mathcal{B} .

LEMMA 3.1 (see [16, Theorem 12]). *For any $f \in \mathcal{B}$ and any $M \in \mathbb{N}^+$, there exists a two-layer ReLU FNN ϕ in $\mathcal{F}_{2,M,\text{ReLU}}$ such that*

$$(3.4) \quad \|f - \phi\|_{L^\infty([0,1]^d)} \leq 4\|f\|_{\mathcal{B}} \sqrt{\frac{d+1}{M}}.$$

However, the solution of the linear system is a grid function defined merely at the set of grid points Γ rather than a continuous domain. Therefore, we define the following norm for grid functions based on the above Barron norm; namely, for any $\mathbf{u} \in \mathbb{R}^{N^d}$,

$$(3.5) \quad \|\mathbf{u}\|_{\mathcal{B},\Gamma} := \inf \|f\|_{\mathcal{B}},$$

where the infimum is taken over all $f \in \mathcal{B}$ such that $f(\mathbf{x}_\alpha) = u_\alpha$ for all $\alpha \in \Lambda$. Briefly, $\|\mathbf{u}\|_{\mathcal{B},\Gamma}$ is the minimal Barron norm among all Barron functions that interpolate \mathbf{u} at Γ . Since Γ is finite, there always exists some $C^\infty(\mathbb{R}^{N^d})$ function with a compact support that interpolates \mathbf{u} at Γ , which is a Barron function (see [6]). So the infimum in (3.5) will never be taken on the empty set. And it is trivial to show that (3.5) is a well-defined norm.

Now let us consider our method, i.e., the NN-based minimization (2.12). The following result shows that the error of our proposed method is bounded by the product of the condition number of the matrix and the NN approximation error. Recall that $\|\cdot\|_{\ell^2}$ is defined by (2.13).

THEOREM 3.2. *Suppose \mathbf{A} in (2.12) is square and invertible. Let θ^* be a minimizer of (2.12) with ϕ being an FNN of depth L and width M . Let \mathbf{u} be the solution of the linear system (2.9). Then it satisfies*

$$(3.6) \quad \|\phi(\theta^*) - \mathbf{u}\|_{\ell^2} \leq 4\kappa(\mathbf{A})\|\mathbf{u}\|_{\mathcal{B},\Gamma} \sqrt{\frac{d+1}{M}},$$

where $\kappa(\mathbf{A}) := \|\mathbf{A}\|_2 \|\mathbf{A}^{-1}\|_2$ is the condition number of \mathbf{A} .

Proof. Let $\tilde{v} \in \mathcal{B}$ be the function taking the infimum in (3.5). Then, by Lemma 3.1, there exists some θ' such that $\phi(x; \theta')$ satisfies

$$(3.7) \quad \|\tilde{v}(x) - \phi(x; \theta')\|_{L^\infty(\Omega)} \leq 4\|\tilde{v}\|_{\mathcal{B}} \sqrt{\frac{d+1}{M}}.$$

Similar to (2.6), we denote

$$(3.8) \quad \tilde{\mathbf{v}} := [\tilde{v}(\mathbf{x}_{(1,1,\dots,1)}) \tilde{v}(\mathbf{x}_{(1,1,\dots,2)}) \cdots \tilde{v}(\mathbf{x}_{(N,N,\dots,N)})]^\top \in \mathbb{R}^{N^d}.$$

Since θ^* is the minimizer of (2.12), we have

$$(3.9) \quad \begin{aligned} \|\phi(\theta^*) - \mathbf{u}\|_2 &\leq \|\mathbf{A}^{-1}\|_2 \|\mathbf{A}\phi(\theta^*) - \mathbf{b}\|_2 \leq \|\mathbf{A}^{-1}\|_2 \|\mathbf{A}\phi(\theta') - \mathbf{b}\|_2 \\ &\leq \|\mathbf{A}^{-1}\|_2 \|\mathbf{A}\|_2 \|\phi(\theta') - \mathbf{u}\|_2 = \kappa(\mathbf{A}) \|\phi(\theta') - \mathbf{u}\|_2. \end{aligned}$$

Noting that $\tilde{v}(\mathbf{x}_\alpha) = u_\alpha$ for all $\alpha \in \Lambda$, we have

$$(3.10) \quad \|\phi(\theta') - \mathbf{u}\|_2 = \|\phi(\theta') - \tilde{\mathbf{v}}\|_2 \leq N^{d/2} \|\phi(\theta') - \tilde{\mathbf{v}}\|_\infty \leq 4N^{d/2} \|\tilde{\mathbf{v}}\|_{\mathcal{B}} \sqrt{\frac{d+1}{M}},$$

where the NN approximation (3.7) is used. Then (3.6) directly follows (3.9), (3.10), and the fact that $\|\phi(\theta^*) - \mathbf{u}\|_{\ell^2} = N^{-d/2} \|\phi(\theta^*) - \mathbf{u}\|_2$. \square

Remark 3.1. We can estimate how wide a two-layer NN should be to obtain acceptable accuracy ϵ under Theorem 3.2. As a typical example, we consider the linear system derived from the centered finite difference method in solving a d -dimensional Poisson's equation (see section 4.1). Supposing the finite difference scheme has a p th order truncation error, then it can be proved that $\kappa(\mathbf{A}) \leq CN^p$, where C is independent of d and N . From (3.6), it suffices to let $M \geq 16C^2 N^{2p} \|u\|_{\mathcal{B},\Gamma}^2 (d+1)/\epsilon^2$. In this case, the number of unknowns $|\theta| = (d+2)M \sim O(N^{2p} \|u\|_{\mathcal{B},\Gamma}^2 d^2/\epsilon^2)$. This number can be compared with N^d , the number of unknowns in the original linear system (2.4). Supposing $\|u\|_{\mathcal{B},\Gamma}$ only increases with d mildly, then $|\theta|$ is less than N^d if $d > 2p$, and their difference is much more significant as d increases. This implies our method contains fewer unknowns than traditional linear solvers in high-dimensional problems.

Note that the error bound (3.6) involves the norm $\|\mathbf{u}\|_{\mathcal{B},\Gamma}$ and the condition number $\kappa(\mathbf{A})$, which both depend on d and N . But the relations are usually implicit. For $\|\mathbf{u}\|_{\mathcal{B},\Gamma}$, it is small if \mathbf{u} has a small interpolant in the sense of the Barron norm. And one can infer from [6] that a function has a small Barron norm if its Fourier transform decays to zero quickly as the frequency increases. Roughly speaking, such functions have "smooth" images. Consequently, we can simply conclude that $\|\mathbf{u}\|_{\mathcal{B},\Gamma}$ is small if the d -dimensional mesh of \mathbf{u} looks smooth. This is usually true in physical problems because \mathbf{u} is an approximation of a smooth physical function. However, we do not have an explicit formula to estimate $\|\mathbf{u}\|_{\mathcal{B},\Gamma}$ for specific problems.

In Theorem 3.2, we specify θ^* as the global minimizer of (2.12). However, in NN optimization, it is usually difficult to find global minimizers numerically due to nonconvexity of loss functions. To the best of our knowledge, there is no optimizers that guarantee to identify a global minimizer. Consequently, the overall error is also affected by the optimization error, which is the difference between the theoretical global minimizer θ^* and the actually found solution θ .

In this paper, we only discuss the approximation property of two-layer shallow networks and take it to figure out the error estimate. We need to mention that the approximation properties of deep networks with $L > 2$ have also been studied (see [36, 45, 46, 47]). For example, it is proved in [36] that if f is a C^s smooth function in $[0, 1]^d$ with $s \in \mathbb{N}^+$, then there exists a ReLU FNN ϕ with width $O(J \log(J))$ and depth $O(K \log(K))$ such that $\|f - \phi\|_{L^\infty([0,1]^d)} \leq O(J^{-2s/d} K^{-2s/d})$ for all $J, K \in \mathbb{N}^+$. By these approximation properties, error estimates for deep networks can also be derived similarly as in Theorem 3.2, and the error bound can be much sharper if v is more special.

4. Numerical experiments. In this section, several linear systems from physical problems are solved by the proposed NN-based method. Due to the high nonlinearity of NNs concerning their parameters, the least squares optimization in this method is very nonlinear and, hence, difficult to solve. In practice, we can only find local minimizers rather than global minimizers, which causes a certain amount of optimization error. The optimization error then limits the accuracy of the numerical

solution. Consequently, for small linear systems, our NN-based method performs less accurately than well-developed traditional methods (e.g., conjugate gradient method and GMRES), which can achieve errors around machine precision for well-conditioned systems. Nevertheless, the proposed method is capable of extremely large linear systems that traditional methods cannot deal with. In this paper, we choose extremely large systems as numerical examples that are exclusively solved by our method, and we do not have any comparison tests.

Our algorithm is implemented on PyTorch with the CUDA platform. The implementation is not picky for hyperparameters, and the numerical results are usually stable, so we do not spend much time on tuning parameters. Specifically, for NN initialization, the results do not show any significant differences between the default initialization used by us and other common initialization methods (e.g., Xavier's initialization [22] and He's initialization [24]). Moreover, we observe in experiments that as long as the batch size is moderately large, enlarging batch sizes will slightly speed up the error decay but can hardly improve the errors of the final numerical solution. In our experiments, suitable batch sizes are chosen according to the time-memory trade-off. For the number of iterations, we empirically set it to make sure the mean of the losses of the last 100 iterations does not exceed 0.001% of the initial loss. More details about the experiment settings can be listed as follows.

- *Environment.* The method is tested in a Python environment. PyTorch library (version 1.10.1) with CUDA toolkit (version 11.3) is utilized for NN implementation and GPU-based parallel computing. We upload the programs¹ for readers and researchers to generate experimental results.
- *Optimizer and hyperparameters.* Algorithm 1 is implemented. In each iteration of Algorithm 1, we randomly select a batch of grid points from Γ with uniform distribution and take their indices to form \mathcal{S} . The learning rates are set to decay from 10^{-3} to 10^{-5} with linearly decreasing powers; namely, we let τ_i be the learning rate of the i th iteration and I be the maximum number of iterations, and then

$$(4.1) \quad \tau_i = 10^{-3 - \frac{2i}{I}} \quad \text{for } i = 1, \dots, I.$$

We remark that we have also tried an Adam optimizer, which can obtain smaller errors but cost slightly more computational time compared with the standard mini-batch gradient descent. Due to the high accuracy of Adam, the results obtained by Adam in comparative tests are not as illustrative as those obtained by mini-batch gradient descent, so we only present the results of the latter.

- *Stopping criterion.* We set sufficiently many iterations for each example, which guarantee that the mean of the losses of the last 100 iterations is less than 0.001% of the initial loss. This stopping criterion is rigorous enough by our empirical experiences, from which promising conclusions can be derived.
- *Network setting.* FNNs with ReLU activation functions are used in the experiments. We implement the method with various depth L and width M to investigate their effects. The network parameters are initialized using the default setting of PyTorch library; namely,

$$(4.2) \quad a, W_\ell, b_\ell \sim U(-M^{-1/2}, M^{-1/2}), \quad \ell = 1, \dots, L-1,$$

which are generated with uniform distribution.

¹The programs are available online from <https://www.math.hkbu.edu.hk/~mng/files/dlearncode.zip>.

- *Testing set and error evaluation.* We prescribe a set of N_{test} grid points from Γ with uniform distribution and take their indices as the testing set, denoted as $\mathcal{T} := \{(\alpha_1^n, \dots, \alpha_d^n)\}_{n=1}^{N_{\text{test}}}$. For the examples given true solutions, we define the following ∞ -error over \mathcal{T} between the numerical solution $\phi(\theta)$ and true solution \mathbf{u} ,

$$(4.3) \quad \|\phi(\theta) - \mathbf{u}\|_{\infty, \mathcal{T}} := \max_{\alpha \in \mathcal{T}} |u_{\alpha} - \phi(\mathbf{x}_{\alpha}; \theta)|,$$

and the ℓ^2 -error over \mathcal{T} ,

$$(4.4) \quad \|\phi(\theta) - \mathbf{u}\|_{\ell^2, \mathcal{T}} := \left(\frac{1}{N_{\text{test}}} \sum_{\alpha \in \mathcal{T}} |u_{\alpha} - \phi(\mathbf{x}_{\alpha}; \theta)|^2 \right)^{\frac{1}{2}}.$$

And for the examples whose true solutions are unknown, we define the following ℓ^2 -residual over \mathcal{T} :

$$(4.5) \quad \begin{aligned} \text{Res}_{\ell^2}(\mathcal{T}) &:= \left(\frac{1}{N_{\text{test}}} \sum_{\alpha \in \mathcal{T}} [\mathbf{b} - \mathbf{A}\phi(\theta)]_{\zeta(n)}^2 \right)^{\frac{1}{2}} \\ &= \left(\frac{1}{N_{\text{test}}} \sum_{\alpha \in \mathcal{T}} |b_{\zeta(n)} - \mathbf{a}_{\zeta(n)}^{\top} \phi(\theta)|^2 \right)^{\frac{1}{2}}, \end{aligned}$$

where $\zeta(n) = \sum_{k=1}^d (\alpha_k^n - 1)N^{d-k} + 1$ is the position of the multi-index $\alpha = (\alpha_1^n, \dots, \alpha_d^n)$ in the lexicographical sequence. Note that the solution error is bounded by the product of the matrix inverse norm and the residual, i.e., $\|\phi(\theta) - \mathbf{u}\| \leq \|\mathbf{A}^{-1}\| \|\mathbf{A}\phi(\theta) - \mathbf{b}\|$, so small residuals imply small errors for well-conditioned problems. For all examples, we set $N_{\text{test}} = \min\{10^4, N^d\}$, noting that N^d is the number of all grid points so N_{test} is at most N^d .

- *Randomness.* To check the effect of the uncertainty of results caused by the randomness of the algorithm in the NN initialization and training data, we compute the mean of the errors/residuals of the last 100 iterations as the “final” error/residual. Moreover, we repeat each experiment using 10 different random seeds (commands `numpy.random.seed(n)` for stochastic NumPy subroutines and `torch.manual_seed(n)` for PyTorch subroutines) and list the mean and standard deviation of the results (shown as “mean \pm standard deviation”). It shows in the following resulting tables that the standard deviations are always dominated by the means, so our implementation is numerically stable and convincing.

In this section, the matrices arising from numerical examples are quite well conditioned. The condition numbers are increasing with N and d mildly, and they do not grow with d exponentially (i.e., the curse of dimensionality). Hence, we do not encounter an ill-conditioning issue. However, for problems with huge condition numbers, the convergence speed of the gradient descent optimizer will be vastly slowed down. If the condition number is large, one can only reduce the loss function to a small extent in every line search toward the gradient, even with the best step size. Hence, it will cost a huge number of iterations to obtain the desired accuracy.

4.1. Poisson’s equation. We consider the Poisson’s equation

$$(4.6) \quad \begin{cases} -\Delta v(\mathbf{x}) = f(\mathbf{x}) & \text{in } \Omega := [-1, 1]^d, \\ v(\mathbf{x}) = 0 & \text{on } \partial\Omega, \end{cases}$$

which is an elliptic PDE describing a variety of steady-state physical phenomena. The physical solution of (4.6) is set as

$$(4.7) \quad v(\mathbf{x}) = \prod_{i=1}^d \sin(\pi x_i),$$

where x_i is the i th component of \mathbf{x} .

A widely used approach for (4.6) is the second-order centered finite difference scheme with uniform grid spacing, which leads to the following matrix with tensor product structure (see [21] for more details on the structure):

$$(4.8) \quad \mathbf{A} = \sum_{n=1}^d \underbrace{\mathbf{I}_N \otimes \cdots \otimes \mathbf{I}_N}_{n-1 \text{ terms}} \otimes \mathbf{T} \otimes \underbrace{\mathbf{I}_N \otimes \cdots \otimes \mathbf{I}_N}_{d-n \text{ terms}},$$

where \otimes denotes the Kronecker product; $\mathbf{I}_N \in \mathbb{R}^{N \times N}$ is the identity matrix; $\mathbf{T} = [T_{i,j}] \in \mathbb{R}^{N \times N}$ is given by

$$(4.9) \quad T_{i,j} = \begin{cases} -2/h^2, & j = i, \\ 1/h^2, & j = i \pm 1, \\ 0 & \text{else;} \end{cases}$$

$h = 2/(N + 1)$ is the grid size.

In the implementation of our method, the batch size $|\mathcal{S}|$ of the mini-batch gradient descent is set to be $\min\{10^4, N^d\}$, and the maximum number of iterations is set to be 5×10^4 . Note that if the total number of grid points N^d is much larger than the sizes of training and testing sets, then with high probability, the randomly selected training and testing sets will be almost disjoint. But for small systems, the training and testing sets may overlap. And they even coincide if N^d is smaller than 10^4 , in which case the training/testing set consists of all the grid points.

The condition number $\kappa(\mathbf{A})$ of this problem is of $O(N^2)$ and is independent of d [7]. So the theoretical error bound given by (3.6) is $O(\|\mathbf{u}\|_{\mathcal{B},\Gamma} N^2 d^{1/2} M^{-1/2})$, which is increasing with N at least in the second order and decreasing with M in the half order. We will show the error bound is coarse for the actual numerical results in the following tests.

4.1.1. Test for small sizes. In the first test, we set the right-hand side \mathbf{b} as the grid representation of the true physical function $f = -\Delta v = d\pi^2 \prod_{i=1}^d \sin(\pi x_i)$. Also, we let \mathbf{v} be the grid representation of v ; namely,

$$(4.10) \quad \mathbf{v} := [v(\mathbf{x}_{(1,1,\dots,1)}) \ v(\mathbf{x}_{(1,1,\dots,2)}) \ \cdots \ v(\mathbf{x}_{(N,N,\dots,N)})]^\top \in \mathbb{R}^{N^d}.$$

Therefore, the true solution \mathbf{u} of the linear system $\mathbf{A}\mathbf{u} = \mathbf{b}$ is an approximation of the physical solution \mathbf{v} , up to a discretization error $O(N^{-2})$ (Theorem 4.2 in [32]). On the other hand, our method will numerically solve the linear system, obtaining $\phi(\theta)$, which is an approximation of \mathbf{u} .

The linear system is solved for $d = 3$ and various N ($N = 6, 12, 24, 48, 96$). We implement Algorithm 1 with the network size $(L, M) = (3, 200)$, in which \mathbf{A} is applied as a matrix-free operator. Also, we use MATLAB backslash to obtain a high-accuracy solution seen as the “true” solution \mathbf{u} for error evaluation. In the test, we note that the size of the linear system is at most 96^3 by 96^3 , which is still tractable by MATLAB

sparse solver with high accuracy (the two-norm residual of the MATLAB solution is 1.1×10^{-9}). The ∞ -error and ℓ^2 -error between any two quantities ($\phi(\theta)$, \mathbf{u} and \mathbf{v}) over the testing set are presented in Table 4.1. The error curves with respect to N are shown in Figure 4.1. We remark that $\|\cdot\|_{\infty, \mathcal{T}}$ and $\|\cdot\|_{\ell^2, \mathcal{T}}$ follow the definitions in (4.3) and (4.4), respectively.

It is observed in Table 4.1 and in Figure 4.1 that on the one hand, the solution error $\|\phi(\theta) - \mathbf{u}\|$ is always below $O(10^{-4})$ for various N , though it increases mildly with N since the linear system is larger and hence more difficult to solve as N becomes larger. On the other hand, as expected by the theory, the discretization error $\|\mathbf{u} - \mathbf{v}\|$ decreases in the rate $O(N^{-2})$ but is still larger than the corresponding solution error $\|\phi(\theta) - \mathbf{u}\|$. And hence the error $\|\phi(\theta) - \mathbf{v}\|$, which characterizes the accuracy of our method in solving the original continuous problem, is dominated by the discretization error. These results imply that for moderately small N , the accuracy of solving the linear system is high enough compared with the discretization itself. But if N is further larger, the discretization error is likely to continue decreasing and be dominated by the solution error of the linear system. However, we cannot test larger N due to the memory limitation in the use of MATLAB direct solvers.

TABLE 4.1
Errors for various N in the Poisson's equation with physical right-hand sides.

	$\ \phi(\theta) - \mathbf{u}\ _{\infty, \mathcal{T}}$	$\ \mathbf{u} - \mathbf{v}\ _{\infty, \mathcal{T}}$	$\ \phi(\theta) - \mathbf{v}\ _{\infty, \mathcal{T}}$
$N = 6$	$2.397\text{e-}05 \pm 3.299\text{e-}05$	$6.480\text{e-}02$	$6.483\text{e-}02 \pm 2.914\text{e-}05$
$N = 12$	$6.114\text{e-}05 \pm 2.479\text{e-}05$	$1.927\text{e-}02$	$1.929\text{e-}02 \pm 2.327\text{e-}05$
$N = 24$	$1.701\text{e-}04 \pm 1.935\text{e-}05$	$5.249\text{e-}03$	$5.346\text{e-}03 \pm 1.741\text{e-}05$
$N = 48$	$1.731\text{e-}04 \pm 1.745\text{e-}05$	$1.336\text{e-}03$	$1.468\text{e-}03 \pm 1.262\text{e-}05$
$N = 96$	$1.774\text{e-}04 \pm 1.372\text{e-}05$	$3.383\text{e-}04$	$4.638\text{e-}04 \pm 1.281\text{e-}05$

(a) ∞ -errors

	$\ \phi(\theta) - \mathbf{u}\ _{\ell^2, \mathcal{T}}$	$\ \mathbf{u} - \mathbf{v}\ _{\ell^2, \mathcal{T}}$	$\ \phi(\theta) - \mathbf{v}\ _{\ell^2, \mathcal{T}}$
$N = 6$	$1.518\text{e-}05 \pm 2.129\text{e-}05$	$3.116\text{e-}02$	$3.116\text{e-}02 \pm 4.927\text{e-}08$
$N = 12$	$1.704\text{e-}05 \pm 1.408\text{e-}05$	$7.780\text{e-}03$	$7.852\text{e-}03 \pm 1.724\text{e-}07$
$N = 24$	$5.048\text{e-}05 \pm 7.770\text{e-}06$	$1.971\text{e-}03$	$1.985\text{e-}03 \pm 2.177\text{e-}07$
$N = 48$	$5.017\text{e-}05 \pm 6.368\text{e-}06$	$4.983\text{e-}04$	$5.024\text{e-}04 \pm 7.598\text{e-}07$
$N = 96$	$4.978\text{e-}05 \pm 4.852\text{e-}06$	$1.249\text{e-}04$	$1.351\text{e-}04 \pm 2.032\text{e-}06$

(b) ℓ^2 -errors

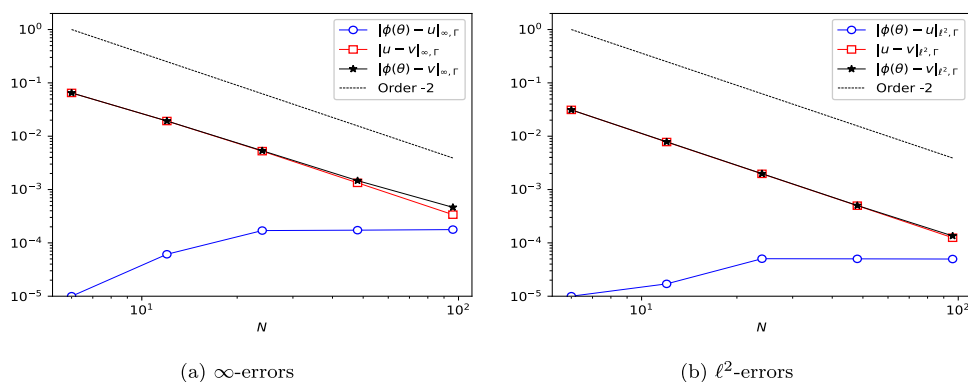


FIG. 4.1. Errors versus N in the Poisson's equation with physical right-hand sides (\mathbf{u} and $\phi(\theta)$ are the true and numerical solutions of the linear system, respectively; \mathbf{v} is the grid representation of the true solution of the original Poisson's equation).

4.1.2. Test for large sizes. In this experiment, we turn to solve linear systems of larger sizes in d . Unlike the preceding experiment, we cannot use high-accuracy traditional solvers to find a “true” solution of the linear system. So we artificially set the true solution as $\mathbf{u} = \mathbf{v}$ given by (4.10). Namely, we ignore the discretization error and directly take the physical solution as the solution of the linear system. The right-hand side is thereafter computed as $\mathbf{b} = \mathbf{A}\mathbf{u}$.

We solve $\mathbf{A}\mathbf{u} = \mathbf{b}$ for $d = 3$ and 6 using Algorithm 1. The experiment is repeated for various N ($N = 10^2, 10^3, 10^4$) and FNN sizes (L, M) ($L = 2, 3$ and $M = 100, 200$). The solution errors over the testing set are shown in Table 4.2. It is observed that the solution error decreases as L or M increases. Also, it is surprising that the solution error does not differ too much for different values of N , which implies that the solution error bound given in (3.6) is coarse in terms of N . Moreover, the running time (i.e., training time for FNN) of the experiment is also reported in Table 4.3. We can find that the running time increases with the network size, but it is almost unchanged for

TABLE 4.2
Errors for various d , N , L , and M in the Poisson's equation.

(L, M)	$N = 10^2$	$N = 10^3$	$N = 10^4$
(2, 100)	4.578e-03 \pm 1.042e-03	4.612e-03 \pm 1.094e-03	4.443e-03 \pm 8.440e-04
(2, 200)	1.986e-03 \pm 2.323e-04	2.034e-03 \pm 1.683e-04	2.036e-03 \pm 1.606e-04
(3, 100)	2.577e-04 \pm 2.449e-05	2.609e-04 \pm 2.893e-05	2.646e-04 \pm 3.021e-05
(3, 200)	1.818e-04 \pm 1.105e-05	1.839e-04 \pm 1.194e-05	1.815e-04 \pm 1.149e-05
(a) $\ \phi(\theta) - \mathbf{u}\ _{\infty, \mathcal{T}} (d = 3)$			
(L, M)	$N = 10^2$	$N = 10^3$	$N = 10^4$
(2, 100)	1.703e-01 \pm 3.143e-02	1.709e-01 \pm 3.449e-02	1.718e-01 \pm 2.643e-02
(2, 200)	4.471e-02 \pm 5.377e-03	4.505e-02 \pm 6.190e-03	4.407e-02 \pm 5.752e-03
(3, 100)	1.301e-02 \pm 1.366e-03	1.393e-02 \pm 1.377e-03	1.339e-02 \pm 1.074e-03
(3, 200)	4.850e-03 \pm 3.314e-04	4.827e-03 \pm 2.576e-04	4.960e-03 \pm 4.270e-04
(b) $\ \phi(\theta) - \mathbf{u}\ _{\infty, \mathcal{T}} (d = 6)$			
(L, M)	$N = 10^2$	$N = 10^3$	$N = 10^4$
(2, 100)	9.506e-04 \pm 2.932e-04	9.337e-04 \pm 3.043e-04	8.818e-04 \pm 2.080e-04
(2, 200)	3.400e-04 \pm 3.561e-05	3.400e-04 \pm 2.936e-05	3.429e-04 \pm 2.426e-05
(3, 100)	6.511e-05 \pm 7.664e-06	6.434e-05 \pm 6.953e-06	6.448e-05 \pm 7.084e-06
(3, 200)	5.225e-05 \pm 4.369e-06	5.174e-05 \pm 3.905e-06	5.092e-05 \pm 3.951e-06
(c) $\ \phi(\theta) - \mathbf{u}\ _{\ell^2, \mathcal{T}} (d = 3)$			
(L, M)	$N = 10^2$	$N = 10^3$	$N = 10^4$
(2, 100)	1.678e-02 \pm 1.690e-03	1.643e-02 \pm 1.939e-03	1.638e-02 \pm 1.810e-03
(2, 200)	6.739e-03 \pm 3.228e-04	6.660e-03 \pm 4.591e-04	6.581e-03 \pm 3.476e-04
(3, 100)	2.128e-03 \pm 1.454e-04	2.084e-03 \pm 1.522e-04	2.087e-03 \pm 1.504e-04
(3, 200)	8.007e-04 \pm 4.432e-05	7.819e-04 \pm 4.124e-05	7.807e-04 \pm 4.382e-05
(d) $\ \phi(\theta) - \mathbf{u}\ _{\ell^2, \mathcal{T}} (d = 6)$			

TABLE 4.3
Running time (seconds) for various d , N , L , and M in the Poisson's equation.

(L, M)	$d = 3$			$d = 6$		
	$N = 10^2$	$N = 10^3$	$N = 10^4$	$N = 10^2$	$N = 10^3$	$N = 10^4$
(2, 100)	5.5e + 02	5.7e + 02	5.5e + 02	1.1e + 03	1.1e + 03	1.1e + 03
(2, 200)	6.1e + 02	6.1e + 02	6.2e + 02	1.2e + 03	1.2e + 03	1.2e + 03
(3, 100)	7.1e + 02	6.9e + 02	6.9e + 02	1.3e + 03	1.4e + 03	1.3e + 03
(3, 200)	1.0e + 03	1.0e + 03	1.0e + 03	2.0e + 03	2.0e + 03	2.0e + 03

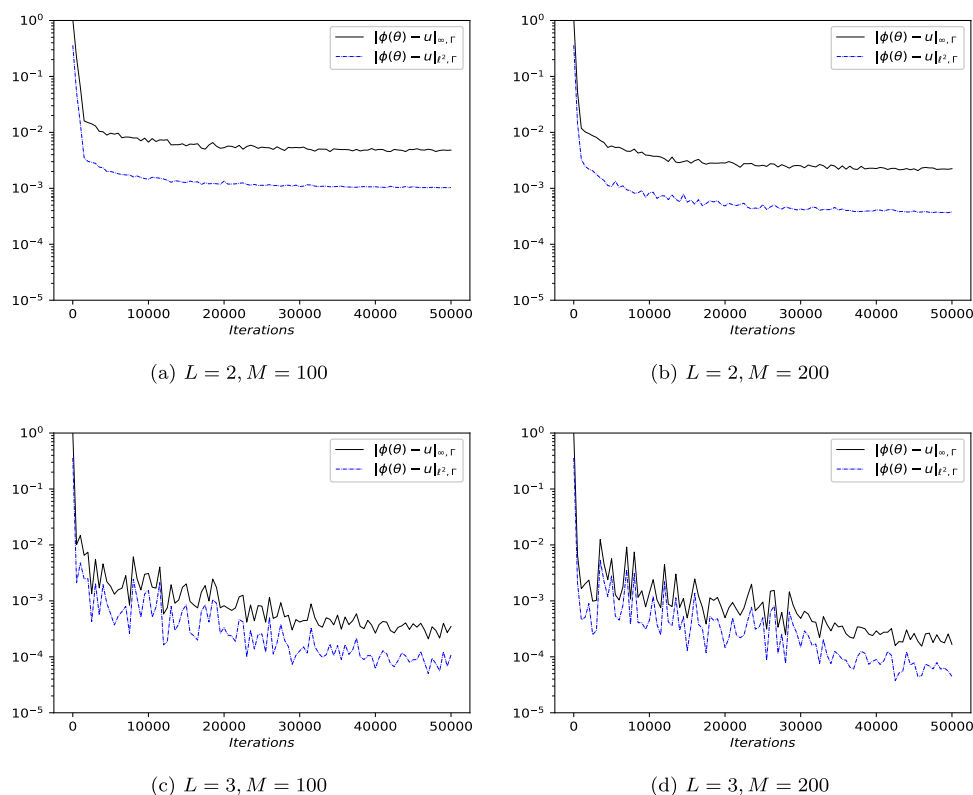


FIG. 4.2. Errors versus iterations for various FNN sizes in the Poisson's equation ($d = 3$).

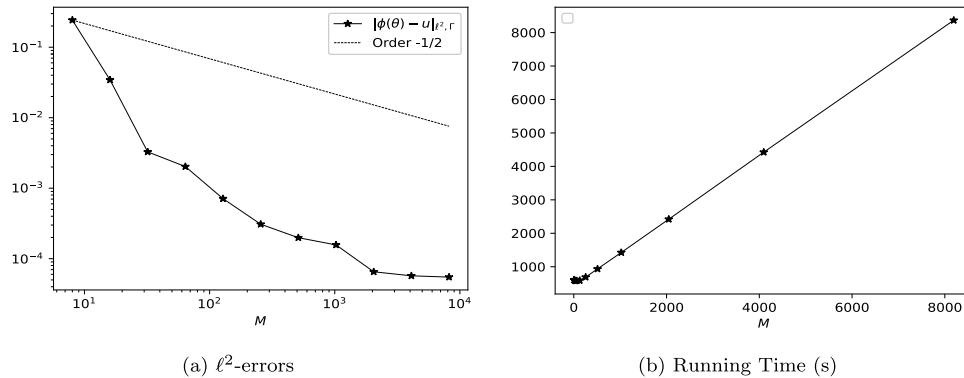
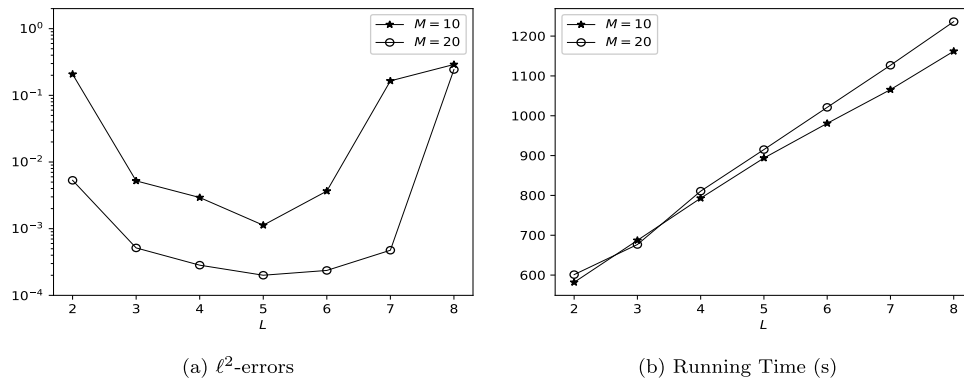
different values of N . So the degree of discretization N does not have a strong effect on the efficiency of the method.

It is noted that the number of unknown parameters $|\theta|$ is at most 41200 when $(L, M) = (3, 200)$, which is much less than N^d , the size of the linear system. Specifically, in the case that $d = 6$ and $N = 10^4$, the size $N^d = 10^{24}$ is extremely large,² which prevents one from using traditional linear solvers, yet the proposed method is still effective, and obtaining ∞ -errors is at best of $O(10^{-3})$.

We also present the error curve (over the testing set) versus iterations in Figure 4.2 to visualize the dynamics of the optimization. It is observed the error decreases rapidly in the first few iterations and decreases slowly afterward. This means a rough solution can be obtained within many fewer iterations. Also, It can be seen that the error decreases of $L = 3$ are more oscillatory than that of $L = 2$. The reason is that the loss function (2.12) with deeper NNs is highly nonlinear in terms of θ , so the optimization is more difficult to be solved using the mini-batch gradient method.

4.1.3. Error versus M . We conduct a third experiment for Poisson's equation to investigate the relation between the solution error and the network width M and compare the result with the solution error bound given in (3.6). With the same setting as in the second experiment, artificial true solution $\mathbf{u} = \mathbf{v}$ is still used. We implement our method to solve the problem with $d = 3$ and $N = 1000$ using $L = 2$

²The ratio of $41200/10^{24}$ is about $4e-20$.

FIG. 4.3. ℓ^2 -errors and running time (seconds) versus M in the Poisson's equation.FIG. 4.4. ℓ^2 -errors and running time (seconds) versus L in the Poisson's equation.

and various M ($M = 2^3, 2^4, \dots, 2^{13}$). The ℓ^2 -errors versus M are shown in Figure 4.3. It is observed that when M is relatively small ($M \leq 2^{11}$ in the figure), the numerical error order is clearly faster than $-\frac{1}{2}$. This implies that the error bound in (3.6) is slightly coarse in terms of M . Indeed, the solution error estimate in (3.6) is derived from the NN approximation error $O(M^{-1/2})$ for the class of Barron functions, but some special functions (e.g., the analytic function v in Poisson's equation) might be approximated by NNs more tightly. But the decrease of the error curve becomes flat as M continues increasing and exceeds some threshold ($M \geq 2^{11}$ in the figure), in which case the solution error is dominated by the optimization error of the gradient descent optimizer. In addition, the running time for various M is shown in Figure 4.3. It is clear that when M is moderately large, the running time increases almost linearly with M .

4.1.4. Error versus L . Similar to the third experiment, we further investigate how the solution error behaves with deeper NNs by a fourth test. We still use the artificial true solution $\mathbf{u} = \mathbf{v}$. The problem with $d = 3$ and $N = 1000$ is solved again with various L ($L = 2, 3, \dots, 8$). The network width M is set as 10 or 20. The ℓ^2 -errors versus L are shown in Figure 4.4. A U-shaped curve is observed such that the smallest error is attained when $L = 5$. Consequently, for this experiment, our method achieves the best accuracy when $L = 5$. It is reasonable to infer that although

deeper NNs usually have smaller approximation errors in theory, they may not bring better numerical results because they are more difficult to train in the practical deep learning. In other words, smaller L leads to larger approximation errors, and larger L tends to bring larger optimization errors. We remark that the optimization difficulty is probably a consequence of the vanishing gradient problem, and it could be overcome by residual NNs [25] to some extent. The running time for various L is also reported in Figure 4.4. Similar to the relation with M , the running time increases almost linearly with the depth L .

4.1.5. Some remarks. In this example, we implement our method to solve the linear system derived from Poisson's equation with the finite difference method. We remark that this process is almost equivalent to using PINNs [43] to solve the same problem. One slight difference is that we address equidistant grid points, but in PINNs, the distribution of the points can be more general (e.g., uniformly distributed random points). Another is that we consider using finite difference schemes to compute the derivatives numerically, but in PINNs, the differentiation can be either analytically or numerically.

We also note that when the linear system is large, the mini-batch gradient descent cannot take all equations as the training set. For instance, in the test with $N = 10^4$ and $d = 6$, there are a total of 10^{24} equations, but we only use 5×10^4 iterations with batch size 10^4 . So most of the equations are actually not involved in the computation. In spite of a very tiny ratio of the training set to all data points, this method is still effective in finding a solution that is quite accurate globally. This is because the approximate NN can generalize well if the solution is smooth and less oscillatory. The residual of the nontraining equations is also reduced along with the minimization of the training set. It also partially explains why the results do not alter too much even if N increases quickly as long as the nature of the physical solution is unchanged.

4.2. Riesz fractional diffusion. In the second numerical example, we consider the following Riesz fractional diffusion equation:

$$(4.11) \quad -\sum_{n=1}^d c_n \frac{\partial^{\alpha_n} v}{\partial |x_n|^{\alpha_n}} = y(\mathbf{x}) \quad \text{in } \Omega := [-1, 1]^d,$$

where $c_n > 0, 1 < \alpha_n < 2$ for all n and $\frac{\partial^{\alpha_n} v}{\partial |x_n|^{\alpha_n}}$ is the Riesz fractional derivative (see [27]). The physical solution of (4.11) is set as

$$(4.12) \quad v(\mathbf{x}) = \sin\left(\sum_{n=1}^d x_n\right),$$

and the corresponding y is therefore given by (4.11).

Employing the finite difference method on (4.11) leads to a linear system whose matrix \mathbf{A} is given by

$$(4.13) \quad \mathbf{A} = \sum_{n=1}^d \underbrace{\mathbf{I}_N \otimes \cdots \otimes \mathbf{I}_N}_{n-1 \text{ terms}} \otimes \mathbf{T}^{(n)} \otimes \underbrace{\mathbf{I}_N \otimes \cdots \otimes \mathbf{I}_N}_{d-n \text{ terms}},$$

with $\mathbf{T}^{(n)}$ being the Toeplitz matrix

$$(4.14) \quad \mathbf{T}^{(n)} := \begin{bmatrix} 2t_1^{(n)} & t_0^{(n)} + t_2^{(n)} & t_3^{(n)} & \cdots & t_{N-1}^{(n)} & t_N^{(n)} \\ t_0^{(n)} + t_2^{(n)} & 2t_1^{(n)} & t_0^{(n)} + t_2^{(n)} & t_3^{(n)} & \cdots & t_{N-1}^{(n)} \\ \vdots & t_0^{(n)} + t_2^{(n)} & 2t_1^{(n)} & \cdots & \cdots & \vdots \\ \vdots & \cdots & \cdots & \cdots & \cdots & t_3^{(n)} \\ t_{N-1}^{(n)} & \cdots & \cdots & \cdots & 2t_1^{(n)} & t_0^{(n)} + t_2^{(n)} \\ t_N^{(n)} & t_{N-1}^{(n)} & \cdots & \cdots & t_0^{(n)} + t_2^{(n)} & 2t_1^{(n)} \end{bmatrix} \in \mathbb{R}^{N \times N},$$

and $t_0^{(n)} := \frac{c_n}{2 \cos(\frac{\alpha_n \pi}{2}) h^{\alpha_n}}$, $t_i^{(n)} = (1 - \frac{\alpha_n + 1}{i}) t_{i-1}^{(n)}$ for $1 \leq i \leq N$. The condition number of \mathbf{A} is of $O(N^\alpha)$ with some $1 < \alpha < 2$, and some preconditioning techniques have been developed (see [35, 15, 38, 14, 27]). In prior work, the linear system with at most $d = 3$ is solved.

In this experiment, we set $c_n = 1$ and $\alpha_n = 1.5$ for all n . We solve $\mathbf{A}\mathbf{u} = \mathbf{b}$ for $d = 5, 10$ and $N = 10$ using Algorithm 1, where the true solution is set as $\mathbf{u} = \mathbf{v}$, with \mathbf{v} being the grid representation of v . We set $|\mathcal{S}| = 2 \times 10^4$ and the number of iterations to be 2×10^4 . Results are shown in Table 4.4. It is clear that the deeper networks with $L = 3$ outperform the shallow ones with $L = 2$ in general. It is also noted in the case $d = 10$ that the wider networks with $M = 200$ performs worse than the narrower one with $M = 100$. It implies that larger networks are sometimes more difficult to train than smaller ones, resulting in greater optimization errors. Despite being more accurate in approximation, larger networks may not provide better numerical results in practice.

4.3. Overflow queuing model. The next example is the overflow queuing model proposed in [10, 11]. Suppose there are d queues with individual queue size N ; then one aims to find nontrivial solutions of the following $N^d \times N^d$ linear system:

$$(4.15) \quad (\mathbf{A} + \mathbf{R})\mathbf{u} = 0.$$

TABLE 4.4
Errors for various d, L , and M in the linear systems arising from the Riesz fractional diffusion.

	$d = 5$	$d = 10$
(2, 100)	1.059e-01 ± 2.054e-02	4.063e-02 ± 1.206e-02
(2, 200)	8.205e-02 ± 3.100e-02	6.516e-02 ± 2.599e-02
(3, 100)	2.947e-03 ± 4.333e-04	8.435e-03 ± 6.248e-03
(3, 200)	2.349e-03 ± 2.032e-04	1.675e-02 ± 6.920e-03
(a) $\ \phi(\theta) - \mathbf{u}\ _{\infty, \mathcal{T}}$		
	$d = 5$	$d = 10$
(2, 100)	4.584e-02 ± 1.382e-02	9.701e-03 ± 2.949e-03
(2, 200)	2.368e-02 ± 1.393e-02	1.443e-02 ± 5.987e-03
(3, 100)	8.140e-04 ± 8.405e-05	1.565e-03 ± 5.339e-04
(3, 200)	6.305e-04 ± 5.490e-05	2.637e-03 ± 4.855e-04
(b) $\ \phi(\theta) - \mathbf{u}\ _{\ell^2, \mathcal{T}}$		

Here the first part **A** has the tensor product structure (4.13), in which $\mathbf{T}^{(n)}$ is given by

$$\mathbf{T}^{(n)} := \begin{bmatrix} \lambda_n & -\mu_n & & & & & & & & & \\ -\lambda_n & \lambda_n + \mu_n & -2\mu_n & & & & & & & & \\ & -\lambda_n & \lambda_n + 2\mu_n & -3\mu_n & & & & & & & \\ & & & \ddots & \ddots & \ddots & & & & & \\ & & & & -\lambda_n & \lambda_n + s_n \mu_n & -s_n \mu_n & & & & \\ & & & & & \ddots & \ddots & \ddots & & & \\ & & & & & & -\lambda_n & \lambda_n + s_n \mu_n & -s_n \mu_n & & \\ & & & & & & & -\lambda_n & s_n \mu_n & & \\ & & & & & & & & & -\lambda_n & s_n \mu_n \end{bmatrix} \in \mathbb{R}^{N \times N},$$

where $s_n \in \mathbb{N}^+$, $\lambda_n, \alpha \in \mathbb{R}^+$ are physical parameters, and $\mu_n := s_n^{-1}(\lambda_n + (N - 1)^{-\alpha})$. The second part $\mathbf{R} = \sum_{m \neq n} \mathbf{R}_{mn}$, with

$$(4.16) \quad \mathbf{R}_{mn} = \bigotimes_{k=1}^d (\mathbf{e}_m \mathbf{e}_m^\top)^{\delta_{mk}} \mathbf{R}_m^{\delta_{nk}},$$

where \mathbf{e}_m is the m th unit vector in \mathbb{R}^N , δ_{mn} is the Kronecker delta, and

$$(4.17) \quad \mathbf{R}_m := \lambda_m \begin{bmatrix} 1 & & & & & & \\ -1 & 1 & & & & & \\ & & \ddots & \ddots & & & \\ & & & -1 & 1 & & \\ & & & & -1 & 0 & \end{bmatrix}.$$

The (normalized) nontrivial solution \mathbf{u} represents the steady-state probability distribution of the queue system. More precisely, $u_{(i_1, i_2, \dots, i_d)}$ is the probability that i_k customers are in the k th queue for $k = 1, \dots, d$. It has been shown that $\mathbf{A} + \mathbf{R}$ has a one-dimensional nullspace.

This problem has an analogue in the continuous case. It is equivalent to the finite difference approximation to an elliptic PDE with a transport term in a rectangular domain, accompanied with the Neumann boundary condition (except for an oblique derivative condition on one particular side). Despite the true solution \mathbf{u} not being given, we can expect that \mathbf{u} is an analogue to the PDE solution up to the local truncation error of the finite difference scheme, and hence \mathbf{u} can be interpolated by smooth Barron functions so that $\|\mathbf{u}\|_{\mathcal{B}, \Gamma}$ in the error bound (3.6) can be small.

In prior work, one could at most solve the linear system with $d = 2$. Due to the one-dimensional nullspace of the matrix, here we use the least squares model with penalty (2.17) to solve (4.15) for $d = 5, 10$, in which the first component of \mathbf{u} is fixed as 1 and the penalty parameter ε is set as 1.0. Since the original problem does not involve any physical domains, a fictitious domain $[0, 1]^d$ is introduced for the implementation. We set $N = 100$, $\alpha = 1$, $\lambda_n = 0.01$, and $s_n = 8n$ for $n = 1, \dots, d$. The algorithm is implemented with $|\mathcal{S}| = 2 \times 10^4$, and the number of iterations is 2×10^4 . The residuals of the obtained solutions are listed in Table 4.5.

We are also interested in the actual distribution of the numerical solution. Specifically, in the case $d = 10$, $L = 2$, and $M = 200$, we find that $x_{\max} = x_{(2,6,10,15,18,24,27,32,36,40)}$ is the location where the approximate solution $\phi(\theta)$ takes its maximum 2.155. It is intuitive to find that the n th index of x_{\max} is equal or

TABLE 4.5
Residuals for various d , L , and M in the queuing problem.

(L, M)	$d = 5$	$d = 10$
(2, 100)	$7.625\text{e-}04 \pm 3.078\text{e-}05$	$1.268\text{e-}03 \pm 2.582\text{e-}04$
(2, 200)	$7.135\text{e-}04 \pm 1.729\text{e-}05$	$1.279\text{e-}03 \pm 1.555\text{e-}04$
(3, 100)	$6.519\text{e-}04 \pm 8.867\text{e-}06$	$1.073\text{e-}03 \pm 3.346\text{e-}05$
(3, 200)	$6.492\text{e-}04 \pm 1.422\text{e-}05$	$1.100\text{e-}03 \pm 2.585\text{e-}05$




FIG. 4.5. Two-dimensional slices of the numerical solution $\phi(x, \theta)$ passing through the maximal point x_{\max} in the queuing problem ($d = 10, L = 3, M = 200$).

close to $s_n/2$. To show a more clear distribution of the solution, we present the two-dimensional slices passing through x_{\max} in Figure 4.5.

Note that the true solutions are unknown. To verify that the images in Figure 4.5 are believable, we take a low-dimensional test for comparison. We solve the queuing linear systems with $d = 2$ or 3 by MATLAB high-accuracy solvers, and the numerical solutions are accurate since the residual is as small as the machine precision. In the low-dimensional cases, we have observed the same property as in the previous high-dimensional cases; namely, the numerical solution attains its maximum at the position $s_n/2$ for the n th index. So we believe that the probability distribution shown in Figure 4.5 is a good simulation, at least in the sense of locating the maxima.

4.4. Probabilistic Boolean networks. Let us consider the steady-state probability distribution of probabilistic Boolean networks, which are widely applied in real-world problems such as genomic signal processing [33]. In this problem, one aims to find the eigenvector associated with the principle eigenvalue 1 of the transition probability matrix, and the normalized eigenvector exactly represents the steady-state probability distribution. The transition probability matrix is of size 2^d by 2^d , where d is the number of genes.

In our experiment, we generate a sparse Toeplitz matrix $\tilde{\mathbf{T}} = [t_{ij}] \in \mathbb{R}^{2^d \times 2^d}$ by

$$(4.18) \quad t_{ij} = \begin{cases} v_k & \text{if } j = i + k \text{ for some } k \in \mathcal{I}, \\ 0 & \text{otherwise,} \end{cases}$$

where \mathcal{I} is a prescribed sparse subset of $\{-2^d, -2^d + 1, \dots, 2^d\}$ and $\{v_k\}_{k \in \mathcal{I}}$ is a prescribed set of positive constants. Followed by column normalization on $\tilde{\mathbf{T}}$, we obtain the transition probability matrix \mathbf{T} whose each column adds up to 1. In practice, we casually choose $\mathcal{I} = \{-13, -5, 2, 6\}$ and $\{v_{13}, v_{-5}, v_2, v_6\} = \{1, 4, 3, 2\}$.

The proposed method is implemented to find the principle eigenvector of \mathbf{T} . First, we introduce the fictitious domain $[0, 1]^d$ and take two grid points $1/3$ and $2/3$ in each dimension. Then we take the following penalized model to compute $(\mathbf{I} - \mathbf{T})\mathbf{u} = 0$:

$$(4.19) \quad \min_{\theta} \frac{1}{N^d} \|(\mathbf{I} - \mathbf{T})\phi(\theta)\|_2^2 + \varepsilon^{-1} (N^{-d} \cdot \mathbf{1}^\top \phi(\theta) - 1)^2,$$

where $\mathbf{1}$ is the all-ones column vector in \mathbb{R}^{2^d} ; namely, we require the mean of the approximate solution to be 1.

This example is slightly different from previous ones. In each dimension, the solution of the linear system characterizes the Boolean state of one object, so the number of grid points in each dimension is always two. Therefore, the solution \mathbf{u} is a grid function defined at a $2 \times 2 \times \dots \times 2$ (d times) grid, which can be interpolated by tensor product linear polynomials. Indeed, recall that the tensor product polynomial space of degree 1 is given by $\mathcal{P} := \text{span}\{x_1^{p_1} x_2^{p_2} \dots x_d^{p_d} : p_i = 0 \text{ or } 1 \text{ for } i = 1, \dots, d\}$. Since the degree of freedom of \mathcal{P} is exactly 2^d , there exists a unique polynomial f in \mathcal{P} such that $u_{\alpha} = f(\mathbf{x}_{\alpha})$ for all $\alpha \in \Lambda$. Moreover, f is analytic and smooth in Ω and hence can be extended to a Barron function on \mathbb{R}^d with a relatively small Barron norm. So we can expect that $\|\mathbf{u}\|_{\mathcal{B}, \Gamma}$ in the error bound (3.6) is small.

The model (4.19) is implemented using a mini-batch gradient descent algorithm with $\varepsilon = 1.0$, $|\mathcal{S}| = 2 \times 10^4$, and number of iteration 2×10^4 . In each iteration, the term $N^{-d} \cdot \mathbf{1}^\top \phi(\theta)$ is only calculated at the training points; specifically, we calculate $|\mathcal{S}|^{-1} \sum_{x \in \mathcal{S}} \phi(x; \theta)$ instead. The cases $d = 50$ and 100 are tested, where the matrix $\mathbf{I} - \mathbf{T}$ has $O(10^{15})$ and $O(10^{30})$ nonzero entries, respectively. The resulting residuals are listed in Table 4.6. We remark that our experiment computes much larger systems than the previous work [33], which solves the same problem with at most 30 dimensions and 5×10^4 nonzero entries.

TABLE 4.6
Residuals for various d , L , and M in the probabilistic Boolean network problem.

(L, M)	$d = 50$	$d = 100$
(2, 100)	6.146e-04 \pm 1.286e-04	1.783e-02 \pm 1.279e-03
(2, 200)	5.809e-04 \pm 8.542e-05	1.761e-02 \pm 8.862e-04
(3, 100)	5.239e-04 \pm 8.815e-05	3.504e-03 \pm 8.978e-04
(3, 200)	5.031e-04 \pm 8.915e-05	3.840e-03 \pm 2.047e-04

5. Conclusion. This work develops a novel NN-based method for extremely large linear systems. The main advantage lies in the saving of storage. Specifically, we create an NN representation for the unknown vector, containing many fewer free elements than the original linear system. The system is then modified to a nonlinear least squares optimization, and it can be solved by gradient descent under a deep learning framework. The proposed method allows us to deal with problems out of storage if using traditional linear solvers. An error estimate is also provided using the approximation property of NNs.

Several physical problems are considered in the numerical experiments. This method is successfully implemented to solve the corresponding linear systems. Compared with prior work on these problems, we solve systems of much larger sizes, usually intractable for other existing methods. However, the accuracy of this NN-based method is generally not as high as traditional ones due to the optimization error. Hence, it is not recommended for small linear systems.

Moreover, as mentioned in section 4.1.5, the effectiveness of the method relies on the smoothness of the physical problem. It is required that the solution and the right-hand side of the linear system do not oscillate globally or locally. Otherwise, the proposed method will fail. For example, suppose the physical solution is zero except for a few localized spikes in small regions. In that case, the corresponding right-hand side of the linear system will also be zero but with a small number of components. So, with high probability, the equations selected as the training set will have exactly zero right-hand sides, and the numerical solution is identically zero. In other words, this method cannot capture the very local property of the solution that is far away from the global tendency.

One direction of future work could be the convergence analysis of the gradient descent in solving the least squares optimization. Namely, we could investigate whether the gradient descent necessarily finds good minimizers. In recent years, some research work has been conducted on the convergence of gradient descent in NN regression, yet it is significantly different from this situation. On the one hand, the loss function in this method is the residual of the linear system rather than the simple ℓ^2 or entropy loss discussed in prior work. On the other hand, most of the previous analysis is based on the overparametrization hypothesis, in which the NN has many more parameters than terms in the loss function. But, in this method, we expect to use NNs with many fewer parameters than equations or unknowns to save the storage.

Numerical results in section 4.1.2 demonstrate that the accuracy and efficiency of the method are hardly affected by the degree of discretization N . In fact, as N increases, although the linear system becomes larger, the mini-batch gradient descent will not carry more burden because the computational amount only depends on the batch size and the number of iterations. Also, we note that our method aims to learn a “smooth” solution, which is a discretization of the physical solution of the original continuous problem, instead of unstructured discrete data. Hence, we do not need a very wide or deep (i.e., overparametrized) NN as the learner of the target solution. Recent literature [2, 1] implies that in learning problems, if the target function is smooth enough (e.g., a polynomial or an NN-like function), it suffices to use a small learner network whose size does not increase with the number of training samples, and it can be successfully trained by gradient descent. Therefore, the analysis of the gradient descent optimization in our method could be made in similar ways, avoiding the overparametrization framework.

We are also inspired by the last numerical example, in which the solution represents binary probability distribution and does not characterize any smooth physical quantities. One open question is whether the smoothness hypothesis of the solution is necessary for the success of this method if we regard the linear system as a $2^d \times 2^d$ structure. In this case, the approximate network only has to fit two points in every dimension. It is simply required that the network acts as a straight line in any dimension. Therefore, it is interesting to investigate whether general $2^d \times 2^d$ linear systems can be handled by this method without many hypotheses on the solution.

REFERENCES

- [1] Z. ALLEN-ZHU, Y. LI, AND Y. LIANG, *Learning and generalization in overparameterized neural networks, going beyond two layers*, in Proceedings of the 33rd International Conference on Neural Information Processing Systems 2019, pp. 6158–6169.
- [2] A. ANDONI, R. PANIGRAHY, G. VALIANT, AND L. ZHANG, *Learning polynomials with neural networks*, in Proceedings of the 31st International Conference on Machine Learning, 2014.
- [3] Z.-Z. BAI, *Motivations and realizations of Krylov subspace methods for large sparse linear systems*, J. Comput. Appl. Math., 283 (2015), pp. 71–78.
- [4] J. BALLANI AND L. GRASEDYCK, *A projection method to solve linear systems in tensor format*, Numer. Linear Algebra Appl., 20 (2013), pp. 27–43.
- [5] A. R. BARRON, *Neural net approximation*, in Proceedings of the 7th Yale Workshop on Adaptive and Learning Systems, 1992.
- [6] A. R. BARRON, *Universal approximation bounds for superpositions of a sigmoidal function*, IEEE Trans. Inform. Theory, 39 (1993), pp. 930–945.
- [7] A. BÖTTCHER AND S. M. GRUDSKY, *On the condition numbers of large semidefinite Toeplitz matrices*, Linear Algebra Appl., 279 (1998), pp. 285–301.
- [8] Y. CAO, Z. FANG, Y. WU, D. X. ZHOU, AND Q. GU, *Towards understanding the spectral bias of deep learning*, in Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, 2021.
- [9] A. CARAGEA, P. PETERSEN, AND F. VOIGTLAENDER, *Neural Network Approximation and Estimation of Classifiers with Classification Boundary in a Barron Class*, preprint, <https://arxiv.org/abs/2011.09363>, 2020.
- [10] R. H. CHAN, *Iterative methods for overflow queueing models I*, Numer. Math., 51 (1987), pp. 143–180.
- [11] R. H. CHAN, *Iterative methods for overflow queueing models II*, Numer. Math., 54 (1988), pp. 57–78.
- [12] R. H. CHAN AND M. K. NG, *Conjugate gradient methods for Toeplitz systems*, SIAM Rev., 38 (1996), pp. 427–482, <https://doi.org/10.1137/S0036144594276474>.
- [13] G. CYBENKO, *Approximation by superpositions of a sigmoidal function*, Math. Control Signals Systems, 2 (1989), pp. 303–314.
- [14] M. DONATELLI, R. KRAUSE, M. MAZZA, AND K. TROTTI, *Multigrid preconditioners for anisotropic space-fractional diffusion equations*, Adv. Comput. Math., 46 (2020), 49.
- [15] M. DONATELLI, M. MAZZA, AND S. SERRA-CAPIZZANO, *Spectral analysis and structure preserving preconditioners for fractional diffusion equations*, J. Comput. Phys., 307 (2016), pp. 262–279.
- [16] W. E, C. MA, S. WOJTOWYTSCH, AND L. WU, *Towards a Mathematical Understanding of Neural Network-Based Machine Learning: What We Know and What We Don't*, preprint, <https://arxiv.org/abs/2009.10713>, 2020.
- [17] W. E, C. MA, AND L. WU, *A priori estimates of the population risk for two-layer neural networks*, Commun. Math. Sci., 17 (2019), pp. 1407–1425.
- [18] W. E, C. MA, AND L. WU, *The Barron space and the flow-induced function spaces for neural network models*, Constr. Approx., 55 (2022), pp. 369–406.
- [19] W. E AND Q. WANG, *Exponential convergence of the deep neural network approximation for analytic functions*, Sci. China Math., 61 (2018), pp. 1733–1740.
- [20] W. E AND S. WOJTOWYTSCH, *Representation Formulas and Pointwise Properties for Barron Functions*, preprint, <https://arxiv.org/abs/2006.05982>, 2020.

- [21] H.-Y. FAN, L. ZHANG, E.-K. CHU, AND Y. WEI, *Numerical solution to a linear equation with tensor product structure*, Numer. Linear Algebra Appl., 24 (2017), e2106.
- [22] X. GLOROT AND Y. BENGIO, *Understanding the difficulty of training deep feedforward neural networks*, in Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, 2010.
- [23] J. HAN, A. JENTZEN, AND W. E, *Solving high-dimensional partial differential equations using deep learning*, Proc. Natl. Acad. Sci. USA, 115 (2018), pp. 8505–8510.
- [24] K. HE, X. ZHANG, S. REN, AND J. SUN, *Delving deep into rectifiers: Surpassing human-level performance on imagenet classification*, in Proceedings of the IEEE International Conference on Computer Vision, 2015.
- [25] K. HE, X. ZHANG, S. REN, AND J. SUN, *Deep residual learning for image recognition*, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016.
- [26] K. HORNIK, M. STINCHCOMBE, AND H. WHITE, *Multilayer feedforward networks are universal approximators*, Neural Netw., 2 (1989), pp. 359–366.
- [27] X. HUANG, X.-L. LIN, M. K. NG, AND H.-W. SUN, *Spectral Analysis for Preconditioning of Multi-dimensional Riesz Fractional Diffusion Equations*, preprint, <https://arxiv.org/abs/2102.01371>, 2021.
- [28] R. T. KELLER AND Q. DU, *Discovery of dynamics using linear multistep methods*, SIAM J. Numer. Anal., 59 (2021), pp. 429–455, <https://doi.org/10.1137/19M130981X>.
- [29] D. P. KINGMA AND J. BA, *Adam: A Method for Stochastic Optimization*, preprint, <https://arxiv.org/abs/1412.6980>, 2014.
- [30] J. M. KLUSOWSKI AND A. R. BARRON, *Approximation by combinations of ReLU and squared ReLU ridge functions with ℓ^1 and ℓ^0 controls*, IEEE Trans. Inform. Theory, 64 (2018), pp. 7649–7656.
- [31] D. KRESSNER AND C. TOBLER, *Krylov subspace methods for linear systems with tensor product structure*, SIAM J. Matrix Anal. Appl., 31 (2010), pp. 1688–1714, <https://doi.org/10.1137/090756843>.
- [32] S. LARSSON AND V. THOMÉE, *Partial Differential Equations with Numerical Methods*, Springer-Verlag, Berlin, 2003.
- [33] W. LI, L.-B. CUI, AND M. K. NG, *On computation of the steady-state probability distribution of probabilistic Boolean networks with gene perturbation*, J. Comput. Appl. Math., 236 (2012), pp. 4067–4081.
- [34] S. LIANG AND R. SRIKANT, *Why deep neural networks for function approximation?*, in Proceedings of the International Conference on Learning Representations, 2017.
- [35] F.-R. LIN, S.-W. YANG, AND X.-Q. JIN, *Preconditioned iterative methods for fractional diffusion equation*, J. Comput. Phys., 256 (2014), pp. 109–117.
- [36] J. LU, Z. SHEN, H. YANG, AND S. ZHANG, *Deep network approximation for smooth functions*, SIAM J. Math. Anal., 53 (2021), pp. 5465–5506, <https://doi.org/10.1137/20M134695X>.
- [37] Z. LU, H. PU, F. WANG, Z. HU, AND L. WANG, *The expressive power of neural networks: A view from the width*, in Proceedings of the 31st International Conference on Neural Information Processing Systems, Curran Associates, Red Hook, NY, 2017, pp. 6232–6240.
- [38] H. MOGHADERI, M. DEGHAN, AND M. DONATELLI, *Spectral analysis and multigrid preconditioners for two-dimensional space-fractional diffusion equations*, J. Comput. Phys., 350 (2017), pp. 992–1011.
- [39] H. MONTANELLI AND Q. DU, *New error bounds for deep ReLU networks using sparse grids*, SIAM J. Math. Data Sci., 1 (2019), pp. 78–92, <https://doi.org/10.1137/18M1189336>.
- [40] H. MONTANELLI, H. YANG, AND Q. DU, *Deep ReLU networks overcome the curse of dimensionality for bandlimited functions*, J. Comput. Math., 39 (2021), pp. 801–815.
- [41] Y.-Q. NIU AND B. ZHENG, *A greedy block Kaczmarz algorithm for solving large-scale linear systems*, Appl. Math. Lett., 104 (2020), 106294.
- [42] P. PETERSEN AND F. VOIGTLÄENDER, *Optimal approximation of piecewise smooth functions using deep ReLU neural networks*, Neural Netw., 108 (2018), pp. 296–330.
- [43] M. RAISSI, P. PERDIKARIS, AND G. E. KARNIADAKIS, *Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations*, J. Comput. Phys., 378 (2019), pp. 686–707.
- [44] Z. SHEN, H. YANG, AND S. ZHANG, *Nonlinear approximation via compositions*, Neural Netw., 119 (2019), pp. 74–84.
- [45] Z. SHEN, H. YANG, AND S. ZHANG, *Deep network approximation characterized by number of neurons*, Commun. Comput. Phys., 28 (2020), pp. 1768–1811.
- [46] Z. SHEN, H. YANG, AND S. ZHANG, *Deep network with approximation error being reciprocal of width to power of square root of depth*, Neural Comput., 33 (2021), pp. 1005–1036.

- [47] Z. SHEN, H. YANG, AND S. ZHANG, *Neural network approximation: Three hidden layers are enough*, *Neural Netw.*, 141 (2021), pp. 160–173.
- [48] J. W. SIEGEL AND J. XU, *Approximation rates for neural networks with general activation functions*, *Neural Netw.*, 128 (2020), pp. 313–321.
- [49] J. W. SIEGEL AND J. XU, *High-Order Approximation Rates for Neural Networks with $ReLU^k$ Activation Functions*, preprint, <https://arxiv.org/abs/2012.07205>, 2020.
- [50] T. SUZUKI, *Adaptivity of deep ReLU network for learning in Besov and mixed smooth Besov spaces: Optimal rate and curse of dimensionality*, in *Proceedings of the International Conference on Learning Representations*, 2018.
- [51] R. TIPIREDDY, P. PERDIKARIS, P. STINIS, AND A. TARTAKOVSKY, *A Comparative Study of Physics-Informed Neural Network Models for Learning Unknown Dynamics and Constitutive Relations*, preprint, <https://arxiv.org/abs/1904.04058>, 2019.
- [52] D. YAROTSKY, *Error bounds for approximations with deep ReLU networks*, *Neural Netw.*, 94 (2017), pp. 103–114.
- [53] D. YAROTSKY, *Optimal approximation of continuous functions by very deep ReLU networks*, in *Proceedings of the Conference on Learning Theory*, 2018.
- [54] Y. ZANG, G. BAO, X. YE, AND H. ZHOU, *Weak adversarial networks for high-dimensional partial differential equations*, *J. Comput. Phys.*, 411 (2020), 109409.
- [55] J.-J. ZHANG, *A new greedy Kaczmarz algorithm for the solution of very large linear systems*, *Appl. Math. Lett.*, 91 (2019), pp. 207–212.