

Greedy training algorithms for neural networks and applications to PDEs



Jonathan W. Siegel ^{a,*}, Qingguo Hong ^a, Xianlin Jin ^b, Wenrui Hao ^a, Jinchao Xu ^a

^a Department of Mathematics, Pennsylvania State University, University Park, PA, 16802, USA

^b School of Mathematical Sciences, Peking University, Beijing, China

ARTICLE INFO

Article history:

Received 5 July 2022

Received in revised form 10 March 2023

Accepted 19 March 2023

Available online 23 March 2023

Keywords:

Neural networks

Partial differential equations

Greedy algorithms

Generalization accuracy

ABSTRACT

Recently, neural networks have been widely applied for solving partial differential equations (PDEs). Although such methods have been proven remarkably successful on practical engineering problems, they have not been shown, theoretically or empirically, to converge to the underlying PDE solution with arbitrarily high accuracy. The primary difficulty lies in solving the highly non-convex optimization problems resulting from the neural network discretization, which are difficult to treat both theoretically and practically. It is our goal in this work to take a step toward remedying this. For this purpose, we develop a novel greedy training algorithm for shallow neural networks. Our method is applicable to both the variational formulation of the PDE and also to the residual minimization formulation pioneered by physics informed neural networks (PINNs). We analyze the method and obtain a priori error bounds when solving PDEs from the function class defined by shallow networks, which rigorously establishes the convergence of the method as the network size increases. Finally, we test the algorithm on several benchmark examples, including high dimensional PDEs, to confirm the theoretical convergence rate. Although the method is expensive relative to traditional approaches such as finite element methods, we view this work as a proof of concept for neural network-based methods, which shows that numerical methods based upon neural networks can be shown to rigorously converge.

© 2023 Elsevier Inc. All rights reserved.

1. Introduction

Machine learning based approaches in the computational mathematics community have increased rapidly in recent years. One of the main new applications of machine learning has been to the numerical solution of differential equations. In particular, neural network-based discretization has become a revolutionary tool for solving differential equations [26,84,33] and for learning the underlying physics behind experimental data [70]. This approach has been applied to a wide variety of practical problems with astounding success [10,11,58,72,64]. The benefit of this new approach is that neural networks can lessen or even overcome the curse of dimensionality for high-dimensional problems [40,33,32,44]. This is due to the dimension independent approximation properties of neural networks [6,42], which have been compared with finite element methods (FEMs) and other tradition methods in approximation theory [91,81,25,19,92,51,75].

* Corresponding author.

E-mail address: jus1949@psu.edu (J.W. Siegel).

Broadly speaking, there are three approaches for solving PDEs using neural networks which have been extensively studied recently. The first approach is to use neural networks to parameterize a set of functions in which the PDE is solved. This approach is taken by the deep Ritz method [89] and by physics informed neural networks (PINNs) [70], and has been used to effectively solve the Schrödinger equation [34,15]. Another common approach is to learn the PDE solution operator using neural networks. This allows the efficient approximation of new solutions as parameters of the underlying PDE are varied and is a non-linear analogue of model reduction [73]. The effectiveness of this approach has been shown through the success of DeepONet [52], the Fourier neural operator [48], and the Galerkin Transformer [14], which have recently been theoretically analyzed [43,44]. Finally, there is the approach of learning the underlying PDE itself from data using deep neural networks, which was pioneered by PINNs [70]. In the following, we consider exclusively the first approach, where neural networks are used to parameterize a set of functions in which the equation is solved.

Generally speaking, we classify the numerical error of the neural network discretization into three parts: 1) the modeling error incurred by solving the PDE over a restricted function class; 2) the optimization error incurred by failing to fully optimize over the function class; and 3) discretization error incurred by discretizing the integrals appearing in the weak form of the equation (More details in Section 3). There are some results which bound the modeling error by considering how efficiently the PDE solution can be approximated with neural networks. For instance, the convergence analysis of the finite neuron method is discussed by considering a family of H^m -conforming piecewise polynomials based on artificial neural network [91]. The convergence rate of the deep Ritz method depends on the dimensionality [29]. The error estimate of the deep Ritz Method for elliptic problems with different boundary conditions is established in [62]. The convergence analysis of the least-squares method based on residual minimization in PINNs has been studied in [77] based on strong and variational formulations. The convergence of PINNs to the PDE solution is analyzed in [76] for linear second-order elliptic and parabolic PDEs.

The optimization error arises when the highly non-linear and non-convex optimization problem resulting from discretizing using neural networks is only approximately solved. There have been some results in the literature which work toward bounding this error. For instance, it has been shown that gradient descent applied to a sufficiently wide network will reach a global minimum [54,27,2,94,4]. In addition, the convergence of stochastic gradient descent (SGD) and Adam [41] has been analyzed in Fourier space. This results in the empirical observation that the error converges fastest in the lowest frequency modes which is known as the frequency principle or spectral bias of neural network training [55,69]. In practice, Adam or SGD are typically used to solve the resulting optimization problems, although other methods, such as a randomized Newton's method [16] and novel specialized methods, for instance the Active Neuron Least Squares method [1], have also been explored. Recently, an interesting optimization method which resembles the greedy algorithms we introduce has been developed for shallow ReLU neural networks [1].

However, the important point is that none of these algorithms empirically achieve asymptotic convergence as the network size increases [89,70]. More specifically, the relative L_2 error of the deep Ritz method using the SGD optimizer stabilizes as the number of neurons increases (Table 1 in [89]), and the relative L_2 error of PINNs using the L-BFGS optimizer even increases as the number of neurons increases (Tables A.2 & A.3 in [70]). Moreover, for one-hidden-layer neural network with fixed inner weights with the ReLU activation function, one can prove that the condition number of the mass matrix is $\mathcal{O}(n^4)$, where n is the number of neurons [37]. This implies that gradient descent method converges very slow especially when n is large. We want to stress that this lack of convergence in no way diminishes the practical utility of PINNs and the deep Ritz methods. In many practical problems, these methods achieve more than sufficient accuracy. However, from a mathematical point of view the question of whether neural network methods can be used to provably solve differential equations remains interesting.

Concerning the generalization accuracy, there are also some analytical results along this direction. For instance, a priori generalization analysis of the deep Ritz method is studied using the Barron norm with activation function SP_τ in [53]. The empirical risk of the PDE solution represented by an over-parameterized-two-layer neural network achieves a global minimizer under some assumptions [54]. The generalization error of PINNs can be bounded by the training error [59]. The generalization error of deep learning-based methods is also analyzed for high dimensional Black-Scholes PDEs to overcome the curse of dimensionality in [9].

However, there are significant gaps in the existing convergence and generalization theory. In particular, the wide networks which are required to make gradient descent or SGD converge cannot be guaranteed to generalize well. On the other hand, networks which are small enough or satisfy an appropriate bound on their coefficients to guarantee generalization cannot be provably optimized using gradient descent or its variants. Recently, this gap has been closed for shallow neural networks in [36].

To control these three numerical errors and observe the asymptotic convergence order numerically, we propose provably convergent algorithms in this paper for efficiently solving the neural network optimization problem. The key idea is to use a greedy algorithm to train shallow neural networks instead of gradient descent. Greedy algorithms have previously been proposed for solving PDEs using a basis of separable functions [30,12,3,45], and have been proposed for training shallow neural networks [46]. Our contributions are to develop a convergence analysis when using greedy algorithms for training shallow neural networks to solve PDEs, to show the practical feasibility of this method even in high dimensions, and to demonstrate that the theoretically derived convergence rates are achieved. To the best of our knowledge, this work is the first rigorous analysis without gaps which uses neural networks to solve PDEs and also the first neural network training algorithm which observes asymptotic convergence numerically. Although the method is currently not particularly

efficient, we view it as a proof-of-concept which demonstrates the viability of using neural networks to rigorously solve PDEs. Improving the efficiency of the method and extending it to deeper networks with more complex architectures is a promising future research direction.

The remaining part of the paper is organized as follows: in Section 2, we introduce the problem setup and discuss the class of elliptic PDEs we will solve. We overview the basic machine learning theory for PDEs in Section 3 and introduce the neural network model classes in Section 4, where we also discuss the approximation error associated with using a neural network discretization. In Section 5, we discuss greedy algorithms for non-linear dictionary approximation and their convergence analysis, which bounds the optimization error. In section 7, we show how to bound the discretization error when discretizing the PDE energy. Several numerical examples are used to demonstrate the efficiency of the greedy algorithms in Section 9 and finally a conclusion is given in Section 10.

2. Basic setup and the model problem

2.1. Variational formulation

We follow here largely the setting in [91]. Let $\Omega \subset \mathbb{R}^d$ be a bounded domain with a sufficiently smooth boundary $\partial\Omega$. For any integer $m \geq 1$, we consider the following $2m$ -th order partial differential equation with certain boundary conditions:

$$\begin{cases} Lu = f & \text{in } \Omega, \\ B^k(u) = 0 & \text{on } \partial\Omega \quad (0 \leq k \leq m-1), \end{cases} \quad (2.1)$$

where $B^k(u)$ denotes the Dirichlet, Neumann, or mixed boundary conditions which will be discussed in detail in the following. Here L is the partial differential operator defined as follows

$$Lu = \sum_{|\alpha|=m} (-1)^m \partial^\alpha (a_\alpha(x) \partial^\alpha u) + a_0(x)u, \quad (2.2)$$

where α denotes n -dimensional multi-index $\alpha = (\alpha_1, \dots, \alpha_n)$ with

$$|\alpha| = \sum_{i=1}^n \alpha_i, \quad \partial^\alpha = \frac{\partial^{|\alpha|}}{\partial x_1^{\alpha_1} \cdots \partial x_n^{\alpha_n}}.$$

For simplicity, we assume that a_α are strictly positive and bounded on Ω for $|\alpha| = m$ and $a_0 = 0$, namely, $\exists \alpha_0 > 0, \alpha_1 < \infty$, such that

$$\alpha_0 \leq a_\alpha(x), a_0(x) \leq \alpha_1 \quad \forall x \in \Omega, \quad |\alpha| = m. \quad (2.3)$$

Further, when considering deterministic numerical quadrature in Section 7.3, we will make the additional assumption that a_α are sufficiently smooth.

Given a nonnegative integer k and a bounded domain $\Omega \subset \mathbb{R}^d$, let

$$H^k(\Omega) := \left\{ v \in L^2(\Omega), \partial^\alpha v \in L^2(\Omega), |\alpha| \leq k \right\} \quad (2.4)$$

be standard Sobolev spaces with norm and seminorm given respectively by

$$\|v\|_k := \left(\sum_{|\alpha| \leq k} \|\partial^\alpha v\|_0^2 \right)^{1/2}, \quad |v|_k := \left(\sum_{|\alpha|=k} \|\partial^\alpha v\|_0^2 \right)^{1/2}.$$

For $k=0$, $H^0(\Omega)$ is the standard $L^2(\Omega)$ space with the inner product denoted by $(\cdot, \cdot)_{0,\Omega}$. Similarly, for any subset $K \subset \Omega$, $L^2(K)$ inner product is denoted by $(\cdot, \cdot)_{0,K}$. We note that, by a well-known property of Sobolev spaces, the assumption (2.3) implies that

$$a(v, v) \gtrsim \|v\|_{m,\Omega}^2, \quad \forall v \in H^m(\Omega), \quad (2.5)$$

where $a(u, v) := \sum_{|\alpha|=m} (a_\alpha \partial^\alpha u, \partial^\alpha v)_{0,\Omega} + (a_0 u, v)$.

Next, we discuss the boundary conditions in detail. A popular type of boundary conditions is the Dirichlet boundary condition when $B^k = B_D^k$ are given by the following Dirichlet type trace operators

$$B_D^k(u) := \left. \frac{\partial^k u}{\partial \nu^k} \right|_{\partial\Omega} \quad (0 \leq k \leq m-1), \quad (2.6)$$

with v being the outward unit normal vector of $\partial\Omega$. Using (2.6), we define

$$H_0^m(\Omega) = \{v \in H^m(\Omega) : B_D^k(v) = 0, 0 \leq k \leq m-1\}. \quad (2.7)$$

For the aforementioned Dirichlet boundary condition, the elliptic boundary value problem (2.1) is equivalent to

Minimization Problem D: Find $u \in H_0^m(\Omega)$ such that

$$u = \arg \min_{v \in H_0^m(\Omega)} \mathcal{R}(v), \quad (2.8)$$

with the energy function \mathcal{R} defined by

$$\mathcal{R}(v) = \frac{1}{2}a(v, v) - \int_{\Omega} f v dx. \quad (2.9)$$

Next we consider the following minimization problem over the whole space $H^m(\Omega)$

Minimization Problem N: Find $u \in H^m(\Omega)$ such that

$$u = \arg \min_{v \in H^m(\Omega)} \mathcal{R}(v), \quad (2.10)$$

with energy function \mathcal{R} defined by (2.9).

The optimization problem (2.10) is equivalent to the following pure Neumann boundary value problems for the PDE operator (2.2):

$$\begin{cases} Lu = f & \text{in } \Omega, \\ B_N^k(u) = 0 & \text{on } \partial\Omega \quad (0 \leq k \leq m-1), \end{cases} \quad (2.11)$$

where

$$B_N^k : H^{2m}(\Omega) \mapsto L^2(\partial\Omega) \quad (2.12)$$

such that the following identity holds

$$(Lu, v) = a(u, v) - \sum_{k=0}^{m-1} \langle B_N^k(u), B_D^k(v) \rangle_{0, \partial\Omega}. \quad (2.13)$$

In particular,

- For $m=1$, we have $B_N^k(u) = \frac{\partial u}{\partial n}$.
- For $m=2$ and $d=2$, we have $B_N^0 u = \frac{\partial}{\partial n} \left(\Delta u + \frac{\partial^2 u}{\partial s^2} \right) - \frac{\partial}{\partial s} \left(\kappa_s \frac{\partial u}{\partial s} \right) \Big|_{\partial\Omega}$ and $B_N^1 u = \frac{\partial^2 u}{\partial n^2} \Big|_{\partial\Omega}$.

In order to handle Dirichlet boundary conditions, we consider the mixed boundary value problem:

$$\begin{cases} Lu_{\delta} = f & \text{in } \Omega, \\ B_D^k(u_{\delta}) + \delta B_N^k(u_{\delta}) = 0, & 0 \leq k \leq m-1. \end{cases} \quad (2.14)$$

It is easy to see that (2.14) is equivalent to the following optimization problem:

$$u_{\delta} = \arg \min_{v \in H^m(\Omega)} \mathcal{R}_{\delta}(v) \quad (2.15)$$

where

$$\mathcal{R}_{\delta}(v) = \frac{1}{2}a_{\delta}(v, v) - (f, v) \quad (2.16)$$

and

$$a_{\delta}(u, v) = a(u, v) + \delta^{-1} \sum_{k=0}^{m-1} \langle B_D^k(u), B_D^k(v) \rangle_{0, \partial\Omega}. \quad (2.17)$$

Using the theory developed in [91], we have an estimate between u and u_{δ} as follows

Lemma 1. [91, Lemma 5.4] Define $\|\cdot\|_{a,\delta} = \sqrt{a_\delta(\cdot, \cdot)}$. Let u be the solution of (2.1) with $B^k = B_D^k$, $0 \leq k \leq m-1$, and u_δ be the solution of (2.14). Then

$$\|u - u_\delta\|_{a,\delta} \lesssim \sqrt{\delta} \|u\|_{2m,\Omega}. \quad (2.18)$$

2.2. Residual formulation

The second type of problem we will consider are more general (potentially non-elliptic and non-symmetric linear PDEs given by

$$Lu = f \text{ in } \Omega, \quad (2.19)$$

where the operator L is given by

$$Lu = \sum_{|\alpha| \leq m} a_\alpha(x) \partial^\alpha u. \quad (2.20)$$

We approach such an equation using the residual minimization technique pioneered by physics informed neural networks (PINNs) [70]. This approach minimizes the residual norm

$$\mathcal{R}(v) = \frac{1}{2} \int_{\Omega} (Lv - f)^2 dx + \int_{\partial\Omega} B(v) dx, \quad (2.21)$$

where $B(v)$ are our boundary conditions.

The advantage of the PINNs approach is exceptional flexibility which allows arbitrary equations, boundary conditions, data assimilation, and unknown terms in the equation itself to be treated in a straightforward manner which can be implemented rapidly. This flexibility has driven multiple recent breakthroughs in scientific computing [10,11,58,72,64].

Our theory will allow us to obtain both a priori and a posteriori bounds on the PDE residual $\mathcal{R}(v)$. Relating this to the solution error is an important problem which has been studied for a variety of PDEs under certain assumptions, including for linear elliptic and parabolic PDEs [76,60], for Kolmogorov PDEs [21], and for the Navier-Stokes equation [20].

3. Basic machine learning theory for PDEs

In this section, we describe the basics of machine learning and statistical learning theory and explain their connections with numerical methods for solving PDEs. Our focus will be on the connections with numerical PDEs, while the statistics and probability theory background can be found in standard references on statistical learning theory [74,61].

3.1. General objective

We consider the following general setup corresponding to classification or regression. Let X, Y and Z denote three sets. Here X represents the input space, Y the label space, and Z is the prediction space. We are trying to ‘learn’ a function $u : X \rightarrow Z$. We suppose that u minimizes the risk, defined by

$$u = \arg \min_{v \in \mathcal{F}} \mathcal{R}(v), \text{ where } \mathcal{R}(v) = \mathbb{E}_{x,y \sim d\mu}[l(x, y, v(x))] = \int_{X \times Y} l(x, y, v(x)) d\mu(x, y), \quad (3.1)$$

over an appropriate function class \mathcal{F} . Here $l : X \times Y \times Z \rightarrow \mathbb{R}$ is an appropriate loss function, and $d\mu$ is a probability measure on $X \times Y$.

For example, in a binary image classification problem we would set $X = [0, 1]^{n \times n}$ and $Y = Z = \{0, 1\}$. Here X represents the set of possible $n \times n$ pixel arrangements, i.e. images, and Y and Z represent the two possible classes. The function $u : X \rightarrow Y$ maps an image x to a label $u(x) \in \{0, 1\}$. A typical loss function would be the indicator function

$$l(x, y, z) = y(1 - z) + z(1 - y) = \begin{cases} 0 & y = z, \\ 1 & y \neq z. \end{cases} \quad (3.2)$$

In this case the risk (3.1) is exactly the classification error, since we calculate

$$\mathcal{R}(v) = \mathbb{E}_{(x,y) \sim d\mu}[l(x, y, v(x))] = \mathbb{P}_{(x,y) \sim d\mu}[y \neq v(x)]. \quad (3.3)$$

The function class \mathcal{F} could be taken as the set of all measurable functions from X to Z , for instance.

To give another example which is more closely related to the situation when solving PDEs, we consider a regression problem, where $X = \mathbb{R}^d$ is the space of regressors, and $Y = Z = \mathbb{R}$ is the space of responses. In this case, we would take

$$l(x, y, z) = \frac{1}{2}(y - z)^2, \quad (3.4)$$

for instance. In this case the risk is exactly the expected ℓ^2 regression error

$$\mathcal{R}(v) = \mathbb{E}_{(x,y) \sim d\mu}[l(x, y, v(x))] = \frac{1}{2} \int_{\mathbb{R}^d \times \mathbb{R}} |y - v(x)|^2 d\mu(x, y). \quad (3.5)$$

To put the solution of PDEs into this framework, we let $X = \Omega \subset \mathbb{R}^d$, $Y = \{0\}$ (i.e. we have no labels) and $Z = \mathbb{R}$, and consider the function class $\mathcal{F} = H^m(\Omega)$. The distribution $d\mu$ on $X \times Y$, which we can simply identify with $X = \Omega$, is the uniform distribution on the domain Ω . We frame the solution of the PDE as the minimization of the risk (3.1) for an appropriate loss function l . For the solution of PDEs, the loss function must depend upon the derivatives of u , so we consider the somewhat more general risk

$$\mathcal{R}(v) = \mathbb{E}_{x \sim d\mu}[l(x, v(x), Dv(x), \dots, D^m v(x))] = \int_X l(x, v(x), Dv(x), \dots, D^m v(x)) d\mu(x). \quad (3.6)$$

There are two prominent approaches for framing a PDE in this manner. One, known as the deep Ritz method [89] is to consider the variational formulation of the PDE. For the elliptic PDE (2.1), this corresponds to setting

$$l(x, v(x), Dv(x), \dots, D^k v(x)) = \frac{1}{2} \left(\sum_{|\alpha|=m} a_\alpha(x) |\partial^\alpha v(x)|^2 + a_0(x) v(x)^2 \right) - f(x) v(x) \quad (3.7)$$

to solve the $2m$ -th order elliptic equation (2.1). In the case of Dirichlet boundary conditions, we must add to this an expectation of an appropriate penalty over the boundary of the domain $X = \Omega$, i.e. our risk becomes

$$\mathcal{R}(v) = \int_X l(x, v(x), Dv(x), \dots, D^m v(x)) d\mu(x) + \delta^{-1} \int_{\partial X} l_{BC}(x, v(x), Dv(x), \dots, D^{m-1} v(x)) d\mu_{BC}(x). \quad (3.8)$$

Here the loss function for the boundary conditions is given by

$$l_{BC}(x, v(x), Dv(x), \dots, D^{m-1} v(x)) = \sum_{k=0}^{m-1} B_D^k(v(x))^2 = \sum_{k=0}^{m-1} \left(\frac{\partial^k v}{\partial \nu^k}(x) \right)^2, \quad (3.9)$$

where ν denotes the outward normal vector. The distribution μ_{BC} is the uniform distribution on the boundary of the domain Ω .

The other main approach we consider, which was pioneered in the breakthrough work on physics informed neural networks (PINNs) [70], sets the loss function to the L^2 residual of the PDE, i.e. in order to solve the m -th order equation $Lu = f$, we set our loss function to

$$l(x, v(x), Dv(x), \dots, D^m v(x)) = \frac{1}{2}(Lu - f)^2 = \frac{1}{2} \left(\sum_{|\alpha| \leq m} a_\alpha(x) \partial^\alpha v(x) - f(x) \right)^2, \quad (3.10)$$

which results in the risk (2.21) when appropriate boundary conditions are added.

3.2. A priori bounds and statistical learning theory

Our goal in the work is to design a method for solving PDEs using shallow neural networks which permits *a priori* estimates. Such a method has the property that it can be guaranteed to work as long as the true solution is well-approximated by a given function class. The field of statistical learning theory is concerned with deriving such *a priori* error estimates for different machine learning methods.

The basic framework of statistical learning theory analyzes the empirical risk minimization procedure. In this method, we draw samples and minimize a potentially modified empirical risk over a restricted function class \mathcal{F}_Θ (depending upon a set of parameters Θ) to obtain the estimate

$$u_{\Theta, N} = \arg \min_{v \in \mathcal{F}_\Theta} \mathcal{R}_N(v), \text{ where } \mathcal{R}_N(v) = \frac{1}{N} \sum_{i=1}^N l'(x_i, v(x_i), z_i). \quad (3.11)$$

Here the modified loss function l' is not necessarily the same as loss function l occurring in (3.6). This is because the loss function l may not be differentiable or even continuous, which makes the numerical optimization of the empirical risk

(3.11) intractable. For example, the classification loss in (3.2) is discontinuous and this presents significant problems when optimizing. As a result, the classification loss may be replaced by a soft margin SVM loss

$$l'(x, y, z) = \max(0, 1 - yz), \quad (3.12)$$

where the model output $z \in Z = \mathbb{R}$ and the label $y \in \{\pm 1\}$. In this case the prediction is not a label, but rather a real number, which can be converted into a label via thresholding. It is easily verified that the SVM loss is a convex upper bound on the classification loss (3.2).

When solving PDEs, the loss function is continuously differentiable so we usually set $l' = l$. In addition, the distribution $d\mu$ is known explicitly and the empirical risk can be approximated using numerical quadrature instead of sampling (recall that we have no labels z in this case as well)

$$u_{\Theta, N} = \arg \min_{v \in \mathcal{F}_{\Theta}} \mathcal{R}_N(v), \text{ where } \mathcal{R}_N(v) = \sum_{i=1}^N w_i l(x_i, v(x_i)), \quad (3.13)$$

where w_i and x_i are quadrature weights and points in domain Ω . When solving equations with Dirichlet boundary conditions, we also need to discretize the integral on the boundary occurring the definition of the risk (3.8). In this case, our empirical risk would become

$$\mathcal{R}_{N, \delta}(v) = \sum_{i=1}^N w_i l(x_i, v(x_i), z_i) + \delta^{-1} \sum_{i=1}^{N_0} \tilde{w}_i l_{BC}(\tilde{x}_i, v(\tilde{x}_i)), \quad (3.14)$$

where the \tilde{w}_i and \tilde{x}_i are quadrature weights and points on the boundary of the domain Ω . Our notation here contains the case where a Monte Carlo discretization is used. In this case the weights $w_i = 1/N$ and the x_i are randomly sampled from a distribution $d\mu$.

In practice, some algorithm is used to approximately solve the optimization problem (3.11) to obtain an estimate $\bar{u}_{\Theta, N}$. The risk can then be bounded as

$$\begin{aligned} \mathcal{R}(\bar{u}_{\Theta, N}) - \mathcal{R}(u) &= [\mathcal{R}(\bar{u}_{\Theta, N}) - \mathcal{R}_N(\bar{u}_{\Theta, N})] + [\mathcal{R}_N(\bar{u}_{\Theta, N}) - \mathcal{R}_N(u_{\Theta, N})] + \\ &\quad [\mathcal{R}_N(u_{\Theta, N}) - \mathcal{R}_N(u_{\Theta})] + [\mathcal{R}_N(u_{\Theta}) - \mathcal{R}(u_{\Theta})] + [\mathcal{R}(u_{\Theta}) - \mathcal{R}(u)], \end{aligned} \quad (3.15)$$

where $u_{\Theta} = \arg \min_{v \in \mathcal{F}_{\Theta}} \mathcal{R}(v)$ is the minimizer of the true risk over the function class \mathcal{F}_{Θ} and u is the global minimizer of the risk (i.e. the function we are trying to learn).

We bound the first and fourth terms in (3.15) by

$$|\mathcal{R}(\bar{u}_{\Theta, N}) - \mathcal{R}_N(\bar{u}_{\Theta, N})| + |\mathcal{R}_N(u_{\Theta}) - \mathcal{R}(u_{\Theta})| \leq 2 \sup_{v \in \mathcal{F}_{\Theta}} |\mathcal{R}(v) - \mathcal{R}_N(v)| \quad (3.16)$$

and note that the term $\mathcal{R}_N(u_{\Theta, N}) - \mathcal{R}_N(u_{\Theta})$ is non-positive by definition to obtain the following fundamental theorem.

Theorem 1. *The true risk (also called generalization error) is bounded by*

$$\mathcal{R}(\bar{u}_{\Theta, N}) - \mathcal{R}(u) \leq \mathcal{R}(u_{\Theta}) - \mathcal{R}(u) + 2 \sup_{u \in \mathcal{F}_{\Theta}} |\mathcal{R}(u) - \mathcal{R}_N(u)| + \mathcal{R}_N(\bar{u}_{\Theta, N}) - \mathcal{R}_N(u_{\Theta, N}). \quad (3.17)$$

When using Monte Carlo sampling to discretize the risk, we take an expectation over the samples x_1, \dots, x_N on both sides of the above equation to get

$$\begin{aligned} &\mathbb{E}_{x_1, \dots, x_N} [\mathcal{R}(\bar{u}_{\Theta, N}) - \mathcal{R}(u)] \\ &\leq \underbrace{\mathcal{R}(u_{\Theta}) - \mathcal{R}(u)}_{\text{modelling error}} + \underbrace{\mathbb{E}_{x_1, \dots, x_N} \left[2 \sup_{u \in \mathcal{F}_{\Theta}} |\mathcal{R}(u) - \mathcal{R}_N(u)| \right]}_{\text{discretization error}} + \underbrace{\mathbb{E}_{x_1, \dots, x_N} [\mathcal{R}_N(\bar{u}_{\Theta, N}) - \mathcal{R}_N(u_{\Theta, N})]}_{\text{optimization error}} \end{aligned} \quad (3.18)$$

The term on the left hand side here is the *generalization error* which we are trying to bound. We will proceed to analyze the three terms on the right hand side.

The term $\mathcal{R}_N(\bar{u}_{\Theta, N}) - \mathcal{R}_N(u_{\Theta, N})$ is the *optimization error* of the method. This measures the failure to completely optimize over the model class \mathcal{F}_{Θ} . In traditional methods for solving PDEs, for example finite element methods, this term corresponds to the error in solving the discrete linear system.

The middle term $2 \sup_{u \in \mathcal{F}_{\Theta}} |\mathcal{R}(u) - \mathcal{R}_N(u)|$ is called the *discretization error* and measures the error incurred by discretizing the integral defining the risk (3.1). In the theory of linear finite elements this term corresponds to numerical quadrature error, which is typically bounded using Strang's lemma [85]. When using a non-linear model class \mathcal{F}_{Θ} , we must develop new methods for bounding this term. The key tool in our analysis is the Rademacher complexity [8].

Finally, the term $\mathcal{R}(u_{\Theta}) - \mathcal{R}(u)$ is called the *modelling error* and measures the failure of the model class \mathcal{F}_{Θ} to capture the true solution, or ground truth u . In statistical learning theory, this term cannot be theoretically controlled since the ground truth is unknown. The validity of this assumption is checked experimentally either by calculating the empirical risk $\mathcal{R}_N(\bar{u}_{\Theta, N})$ of the learned model or by using a new test dataset if the discretization error cannot be bounded. The advantage of being able to bound the other error terms is that one can conclude that if the method does not empirically perform well on the given data, then this must be due to the model class \mathcal{F}_{Θ} not accurately capturing the ground truth.

Bounding this term in the PDE context requires an estimate on how accurately the model class \mathcal{F}_{Θ} can approximate the solution of the PDE. This requires both a regularity result on the solution of the PDE and an approximation theoretic result concerning the model class \mathcal{F}_{Θ} . For the Barron space model class we introduce in Section 4 such bounds have been obtained in [91,53,18,17] for certain equations. In addition, sharp approximation results for neural networks on the Barron space can be found in [83,79].

In typical applications of deep learning, including to PDEs [70,89] the empirical risk (3.11) is minimized using stochastic gradient descent (SGD) or a variant like ADAM [41]. Bounding both the optimization and discretization error for such methods is a significant challenge. There are results which bound the optimization error by showing that sufficiently large neural networks can be trained to match arbitrary training data using SGD [4,28]. However, when using such a large network the function class \mathcal{F}_{Θ} is very large and this precludes the estimation of the discretization error. This makes analyzing the solution error when solving PDEs using neural networks a significant challenge if SGD or ADAM are used for training. Indeed, a convergence of the error as the size of the network increases cannot be found empirically when solving PDEs [89], although these methods have reliably been able to attain an acceptable accuracy for many practical problems [10,11,58,72,64]. Our approach to this problem is to use greedy algorithms for training instead of SGD or ADAM. This allows us to obtain *a priori* estimates on our error, i.e. to bound the optimization and discretization errors.

3.3. Test error bounds

Next, we consider the problem of obtaining bounds on the risk \mathcal{R} when Theorem 1 does not apply. Suppose that a function u^* has been obtained in some manner, potentially via an unknown black-box method. Our goal is to estimate the risk $\mathcal{R}(u^*)$, i.e. to test the single function u^* . Such a situation would occur when we are unable to bound the optimization, discretization, or modelling errors on the right hand side of Theorem 1.

In typical machine learning problems, the distribution $d\mu$ in (3.1) is unknown and we can only interact with it by drawing i.i.d. samples from $d\mu$. The (true) risk (3.1) is then approximated by the empirical risk

$$\mathcal{R}_{N'}(u^*) = \frac{1}{N'} \sum_{i=1}^{N'} l(x_i, u^*(x_i), z_i), \quad (3.19)$$

where $(x_i, z_i)_{i=1}^{N'}$ are i.i.d. samples from $d\mu$ constituting the test dataset. This is akin to a Monte Carlo discretization of the integral in (3.1). For applications in numerical PDEs, however, the distribution $d\mu$ is typically known explicitly. In these cases, the integral in (3.1) can potentially be more effectively discretized as

$$\mathcal{R}_{N'}(u^*) = \sum_{i=1}^{N'} w_i l(x_i, u^*(x_i), z_i), \quad (3.20)$$

where the w_i are quadrature weights and the (x_i, z_i) are quadrature points. The weights and points can be taken to be accurate to a given high order or may be determined via quasi-Monte Carlo integration methods [49], for instance.

To ensure that we are accurately estimating the true risk of the function u^* we need to obtain a bound on the discretization error

$$|\mathcal{R}_{N'}(u^*) - \mathcal{R}(u^*)|. \quad (3.21)$$

As an example, in the case of the classification loss we can apply Hoeffding's inequality [35] to obtain

$$\mathbb{P}(|\mathcal{R}_{N'}(u^*) - \mathcal{R}(u^*)| \geq \epsilon) \leq 2 \exp(-2\epsilon^2/N'), \quad (3.22)$$

since the loss l is bounded between 0 and 1. This implies that with a large number of samples N' , we can estimate the classification error probability of a given *fixed* model u^* to accuracy $O((N')^{-\frac{1}{2}})$ with high probability.

We remark that in order for this approach to be rigorously correct, the test dataset used to evaluate $\mathcal{R}_N(u^*)$ must be *independent* of the function u^* . This means for instance that the procedure used to determine u^* cannot depend upon the test accuracy (using the same test dataset) $\mathcal{R}_N(u')$ of *any other model* u' , i.e. it cannot depend upon previously published results run on the same test dataset. Of course, in practice this is violated in the deep learning community due to the expense of obtaining datasets and the consequent necessity of reusing test datasets many times for different models. Nonetheless, deviating from the ideal of redrawing a new test dataset for each model has been shown empirically to result in models that exhibit a significant drop in accuracy on new data [71].

The disadvantage of an a posteriori bound is that if the estimated function u^* does not have small risk, there is no way to fix this other than to try again with a different method for estimating u^* (i.e. “to tweak the hyperparameters of the method”) and hope for the best, since we do not know why the method is failing. This is why we are trying to solve PDEs using neural networks in a way which allows *a priori* error estimates to be obtained as described in Section 3.2. This will allow us to obtain error bounds before applying our method, and further, if the method does not work, it allows us to conclude that the model class \mathcal{F}_Θ cannot accurately approximate the PDE solution.

We also note that the simple test error bound derived in (3.22) relied critically upon the fact that the loss function is bounded (in the case of classification). Unfortunately, the loss functions used to solve PDEs are typically not bounded. This means that the test error cannot be used to bound the generalization error in this context. The reason is that one would have to know how smooth the neural network function is in order to use quadrature which guarantees a certain error. Such bounds on the derivative norms of trained neural networks are not available to the best of our knowledge. In our method, the fact that we can control the complexity of the numerical solution, see Section 7, implies that we can obtain bounds on the true (i.e. continuous) energy and on the true residual in the case of the variational formulation and PINNs loss, respectively. This enables us to calculate a posteriori estimates on the energy and on the residuals using a new test dataset (or a new set of quadrature points). This permits the calculation of reference solutions even when the true solution is not known, which to the best of our knowledge cannot be done with other neural networks based methods.

4. Shallow neural network model classes

In this section, we introduce the model class \mathcal{F}_Θ over which we will optimize the empirical loss (3.11). This classical choice is to take \mathcal{F}_Θ to be an n -dimensional subspace of an appropriate Sobolev space. In our approach, we instead take \mathcal{F}_Θ to be non-linear expansions with respect to a suitable collection of functions \mathbb{D} , called a dictionary.

Specifically, for a set $\mathbb{D} \subset C^m(\Omega)$, we consider

$$\Sigma_{n,M}(\mathbb{D}) = \left\{ \sum_{i=1}^n a_i d_i, \quad d_i \in \mathbb{D}, \quad \sum_{i=1}^n |a_i| \leq M \right\}. \quad (4.1)$$

Note that here we restrict the ℓ^1 -norm of the coefficients a_i in the expansion. In addition, we take our dictionary $\mathbb{D} \subset C^m(\Omega)$ (instead of $H^m(\Omega)$ since we will discretize the resulting integrals using quadrature point evaluations). In some cases, we will also need to consider the set

$$\Sigma_{n,\infty}(\mathbb{D}) = \left\{ \sum_{i=1}^n a_i d_i, \quad d_i \in \mathbb{D} \right\} \quad (4.2)$$

with no restriction on the coefficients. We then take the model class \mathcal{F}_Θ to be

$$\mathcal{F}_{n,M} = \Sigma_{n,M}(\mathbb{D}) \quad (4.3)$$

which is parameterized by $\Theta = (n, M)$. Here the dependence on \mathbb{D} is suppressed since the dictionary \mathbb{D} will typically be fixed throughout our analysis.

For shallow neural networks with ReLU ^{k} activation function $\sigma = \max(0, x)^k$ the dictionary \mathbb{D} would be taken as [78]

$$\mathbb{D} = \mathbb{P}_k^d := \{\sigma_k(\omega \cdot x + b) : \omega \in S^{d-1}, b \in [c_1, c_2]\} \subset L^2(B_1^d), \quad (4.4)$$

where $S^{d-1} = \{\omega \in \mathbb{R}^d : |\omega| = 1\}$ is the unit sphere. Here c_1 and c_2 are chosen to satisfy

$$c_1 < \inf\{x \cdot \omega : x \in \Omega, \omega \in S^{d-1}\} < \sup\{x \cdot \omega : x \in \Omega, \omega \in S^{d-1}\} < c_2. \quad (4.5)$$

The default choice $[c_1, c_2] = [-2, 2]$ is used in our experiments in Section 8 for the cases $\Omega \subset B_1^d$, where B_1^d is the closed d -dimensional unit ball. We note that $\mathbb{P}_k^d \subset C^m(\Omega)$ whenever $k > m$ and that in this case $|\mathbb{P}_k^d| = \sup_{g \in \mathbb{P}_k^d} \|g\|_{H^m(\Omega)} < \infty$. In this

case the model class would be given by

$$\mathcal{F}_{n,M} = \Sigma_{n,M}(\mathbb{P}_k^d) = \left\{ \sum_{i=1}^n a_i \sigma_k(\omega_i \cdot x + b_i), \quad \omega_i \in S^{d-1}, \quad b_i \in [c_1, c_2], \quad \sum_{i=1}^n |a_i| \leq M \right\}, \quad (4.6)$$

which is the class of shallow ReLU ^{k} neural networks with width n and coefficients bounded in ℓ^1 by M .

In the case of a general activation function σ , the corresponding dictionary is given by

$$\mathbb{D}_\sigma = \{\sigma(\omega \cdot x + b) : (\omega, b) \in \Theta\}, \quad (4.7)$$

where $\Theta \subset \mathbb{R}^d \times \mathbb{R}$ is compact. In this case, we have $\mathbb{D}_\sigma \subset C^m(\Omega)$ and $|\mathbb{D}_\sigma| < \infty$ whenever $\sigma \in C^m(\Omega)$. In this case, the function class would consist of

$$\mathcal{F}_{n,M} = \Sigma_{n,M}(\mathbb{D}_\sigma) = \left\{ \sum_{i=1}^n a_i \sigma(\omega_i \cdot x + b_i), (\omega_i, b_i) \in \Theta, \sum_{i=1}^n |a_i| \leq M \right\}, \quad (4.8)$$

which is the class of shallow neural networks with activation function σ , bounded inner coefficients and outer coefficients bounded in ℓ^1 by M .

4.1. Barron space regularity

In this section, we introduce the notion of regularity which corresponds to the model class of shallow neural networks introduced in Section 4. As in Section 4, we give this notion in the abstract setting of a general dictionary $\mathbb{D} \subset C^m(\Omega)$.

Consider the closed convex hull of \mathbb{D} , defined by

$$\overline{B}_1(\mathbb{D}) := \overline{\bigcup_{n=1}^{\infty} \Sigma_{n,1}(\mathbb{D})}, \quad (4.9)$$

where $\Sigma_{n,1}$ is defined in (4.1). Note that here the closure is taken in $H^m(\Omega)$. Associated with the convex set $B_1(\mathbb{D})$, we define the gauge norm (also called the Minkowski functional) by

$$\|f\|_{\mathcal{K}_1(\mathbb{D})} = \inf\{c > 0 : f \in cB_1(\mathbb{D})\}. \quad (4.10)$$

The norm $\|\cdot\|_{\mathcal{K}_1(\mathbb{D})}$, which is also called the variation norm corresponding to the dictionary \mathbb{D} , is constructed precisely so that $B_1(\mathbb{D})$ its unit ball. We further define the function space

$$\mathcal{K}_1(\mathbb{D}) := \{f \in H^m(\Omega) : \|f\|_{\mathcal{K}_1(\mathbb{D})} < \infty\}. \quad (4.11)$$

Important fundamental properties of this space, for instance is the fact that if \mathbb{D} is a uniformly bounded dictionary, i.e. if $\sup_{d \in \mathbb{D}} \|d\|_H = K_{\mathbb{D}} < \infty$, then the space $\mathcal{K}_1(\mathbb{D})$ is a Banach space, can be found in [80].

The utility of the space $\mathcal{K}_1(\mathbb{D})$ is due to the fact that its elements can be efficiently approximated by non-linear dictionary expansions. In particular, the following classical bound holds [6,68]

$$\inf_{f_n \in \Sigma_{n,M}(\mathbb{D})} \|f - f_n\|_{H^m(\Omega)} \leq |\mathbb{D}| \|f\|_{\mathcal{K}_1(\mathbb{D})} n^{-\frac{1}{2}}, \quad (4.12)$$

for $M = \|f\|_{\mathcal{K}_1(\mathbb{D})}$. Because of this approximation result, we consider regularity assumptions with respect to the $\mathcal{K}_1(\mathbb{D})$ -norm, i.e. we assume that the variation norm of the PDE solution can be controlled. For the specific variation spaces corresponding to the dictionaries \mathbb{P}_k^d , such regularity results for a variety of PDEs have been obtained [91,53,18,17].

Recently, the spaces $\mathcal{K}_1(\mathbb{P}_k^d)$ for the dictionaries \mathbb{P}_k^d corresponding to shallow ReLU^k neural networks have been characterized in terms of the Radon transform [80,65,63,66] and they are closely related to the Ridgelet spaces [13]. In addition, precise approximation theoretic properties of the space $\mathcal{K}_1(\mathbb{P}_k^d)$, such as the asymptotics of its metric entropy and n -widths can be found in [83]. In [83] it is also shown that the approximation rate (4.12) can be improved to

$$\inf_{f_n \in \Sigma_{n,M}(\mathbb{P}_k^d)} \|f - f_n\|_{H^m(\Omega)} \lesssim \|f\|_{\mathcal{K}_1(\mathbb{P}_k^d)} n^{-\frac{1}{2} - \frac{2k+1}{2d}}, \quad (4.13)$$

with $M \lesssim \|f\|_{\mathcal{K}_1(\mathbb{P}_k^d)}$ for the dictionary $\mathbb{D} = \mathbb{P}_k^d$. Similar results for more general activation functions can be found in [79]. Pointwise properties of functions in $\mathcal{K}_1(\mathbb{P}_1^d)$, which is also called the Barron space [56], have also been obtained in [90].

5. Greedy algorithms

In this section, we address the problem of bounding the optimization error in Theorem 1 when optimizing the empirical loss over the model class $\mathcal{F}_\Theta = \Sigma_{n,M}(\mathbb{D})$ introduced in Section 4. For simplicity, we denote the numerical solution $\bar{u}_{\Theta,N} = \bar{u}_{n,M,N}$ as u_n in this section.

As in Section 4, let $\mathbb{D} \subset H$ be a dictionary in Hilbert space H (in our applications typically $H = H^m(\Omega)$ for some domain Ω). Greedy algorithms for expanding a function $u \in H$ as a linear combination of the dictionary elements \mathbb{D} are fundamental in approximation theory [24,87,86] and signal processing [57,67]. Greedy methods have also been proposed for optimizing shallow neural networks [46,22] and for solving PDEs numerically [30,12,3,45].

The class $\mathcal{K}_1(\mathbb{D})$ which was introduced in Section 4.1 is a natural target space in the analysis of greedy algorithms [87,86]. Given the dictionary \mathbb{D} and a target function u or a convex loss function \mathcal{L} , greedy algorithms either approximate f or approximately minimize \mathcal{L} by a finite linear combination of dictionary elements:

$$u_n = \sum_{i=1}^n a_i g_i, \quad (5.1)$$

with $g_i \in \mathbb{D}$. The two types of greedy algorithm we discuss here are the relaxed greedy algorithm (RGA) and orthogonal greedy algorithm (OGA).

5.1. Relaxed greedy algorithm

We consider the following version of the RGA, which explicitly optimizes \mathcal{L} over the convex hull of the dictionary,

$$u_0 = 0, \quad g_n = \arg \max_{g \in \mathbb{D}} \langle g, \nabla \mathcal{L}(u_{n-1}) \rangle_H, \quad u_n = (1 - \alpha_n)u_{n-1} - M\alpha_n g_n. \quad (5.2)$$

Here the dictionary \mathbb{D} is assumed to be symmetric (i.e. $g \in \mathbb{D}$ implies that $-g \in \mathbb{D}$ as well), the sequence α_n is given by $\alpha_n = \min\left(1, \frac{2}{n}\right)$, and M is a regularization parameter which controls the $\mathcal{K}_1(\mathbb{D})$ -norm of the iterates u_n . This algorithm was first introduced and analyzed by Jones [38] for function approximation (i.e. $\mathcal{L}(u) = \|u - f\|_H^2$), and has been extended to the optimization of general convex objectives as well [93]. The convergence theorem we will use in our analysis, which is closely related to Theorem IV.2 in [93], is the following.

Theorem 2. Suppose that the dictionary \mathbb{D} is symmetric and satisfies $\sup_{d \in \mathbb{D}} \|d\|_H \leq C < \infty$. Let the iterates u_n be given by the RGA (5.2). Assume that the loss function \mathcal{L} is convex and K -smooth (on the Hilbert space H). Recall that K -smoothness means that for any $u, v \in H$ we have

$$\mathcal{L}(u) \leq \mathcal{L}(v) + \langle \nabla \mathcal{L}(v), u - v \rangle_H + \frac{K}{2} \|u - v\|_H^2. \quad (5.3)$$

Then we have $u_n \in \Sigma_{n,M}$ and

$$\mathcal{L}(u_n) - \inf_{\|v\|_{\mathcal{K}_1(\mathbb{D})} \leq M} \mathcal{L}(v) \leq \frac{32(CM)^2 K}{n} \quad (5.4)$$

This theorem will be applied in the case $\mathcal{L} = \mathcal{R}_N$ is the empirical risk to bound the *optimization error*. In particular it yields that

$$\mathcal{R}_N(u_n) - \inf_{v \in \Sigma_{n,M}(\mathbb{D})} \mathcal{R}_N(v) \leq \mathcal{R}_N(u_n) - \inf_{\|v\|_{\mathcal{K}_1(\mathbb{D})} \leq M} \mathcal{R}_N(v) \lesssim n^{-1}. \quad (5.5)$$

We remark that this theorem holds for any convex and K -smooth loss function. This means that the RGA can be applied to non-linear equations in addition to the linear equations introduced in Section 2, provided that the non-linear equations admit a variational formulation with a convex energy function.

Proof. Since $u_0 = 0$ and u_k is a convex combination of u_{k-1} and $-Mg_k$, we see by induction that $u_k \in \Sigma_{k,M}$. The K -smoothness of the objective \mathcal{L} implies that

$$L(u_k) \leq L(u_{k-1}) + \langle \nabla \mathcal{L}(u_{k-1}), u_k - u_{k-1} \rangle + \frac{K}{2} \|u_k - u_{k-1}\|_H^2. \quad (5.6)$$

Using the iteration (5.2), we see that $u_k - u_{k-1} = -s_k u_{k-1} - M s_k g_k$. Plugging this into the above equation, we get

$$\mathcal{L}(u_k) \leq \mathcal{L}(u_{k-1}) - s_k \langle \nabla \mathcal{L}(u_{k-1}), u_{k-1} + Mg_k \rangle + \frac{K s_k^2}{2} \|u_{k-1} + Mg_k\|_H^2. \quad (5.7)$$

Since the dictionary elements g_k satisfy $\|g_k\|_H \leq C$ and $\|u_{k-1}\|_{\mathcal{K}_1(\mathbb{D})} \leq M$, we see that $\|u_{k-1}\|_H \leq CM$ as well. Plugging this into the previous equation implies the bound

$$\mathcal{L}(u_k) \leq \mathcal{L}(u_{k-1}) - s_k \langle \nabla \mathcal{L}(u_{k-1}), u_{k-1} + Mg_k \rangle + 2(CM)^2 K s_k^2. \quad (5.8)$$

Now let z with $\|z\|_{\mathcal{K}_1(\mathbb{D})} \leq M$ be arbitrary. Then also $\|z\|_{\mathcal{K}_1(\mathbb{D})} \leq M$ and the $\arg \max$ characterization of g_k (6.2) implies that

$$\langle \nabla \mathcal{L}(u_{k-1}), -z \rangle \leq \langle \nabla \mathcal{L}(u_{k-1}), Mg_k \rangle. \quad (5.9)$$

Using this in equation (5.8) gives

$$\mathcal{L}(u_k) \leq \mathcal{L}(u_{k-1}) - s_k \langle \nabla \mathcal{L}(u_{k-1}), u_{k-1} - z \rangle + 2(CM)^2 K s_k^2. \quad (5.10)$$

The convexity of \mathcal{L} means that $\mathcal{L}(u_{k-1}) - \mathcal{L}(z) \leq \langle \nabla \mathcal{L}(u_{k-1}), u_{k-1} - z \rangle$. Using this and subtracting $\mathcal{L}(z)$ from both sides of the above equation gives

$$\mathcal{L}(u_k) - \mathcal{L}(z) \leq (1 - s_k)(\mathcal{L}(u_{k-1}) - \mathcal{L}(z)) + 2(CM)^2 K s_k^2. \quad (5.11)$$

Expanding the above recursion (using that $s_k \leq 1$), we get that

$$\mathcal{L}(u_n) - \mathcal{L}(z) \leq \left(\prod_{k=1}^n (1 - s_k) \right) (\mathcal{L}(u_0) - \mathcal{L}(z)) + 2(CM)^2 K \sum_{i=1}^n \left(\prod_{k=i+1}^n (1 - s_k) \right) s_i^2. \quad (5.12)$$

Using the choice $s_k = \max\left(1, \frac{2}{k}\right)$, for which $s_1 = 1$, we get

$$\mathcal{L}(u_n) - \mathcal{L}(z) \leq 2(CM)^2 K \sum_{i=1}^n \left(\prod_{k=i+1}^n (1 - s_k) \right) s_i^2. \quad (5.13)$$

Finally, we bound the product $\prod_{k=i+1}^n (1 - s_k)$ using that $\log(1 + x) \leq x$ as

$$\log \left(\prod_{k=i+1}^n (1 - s_k) \right) \leq - \sum_{k=i+1}^n s_k = - \sum_{k=i+1}^n \frac{2}{k} \leq - \int_{i+1}^{n+1} \frac{2}{x} dx \leq 2(\log(i+1) - \log(n+1)), \quad (5.14)$$

for $i \geq 1$. Thus, $\prod_{k=i+1}^n (1 - s_k) \leq \frac{(i+1)^2}{(n+1)^2}$. Using this in equation (5.13), we get

$$\mathcal{L}(u_n) - \mathcal{L}(z) \leq 2(CM)^2 K \sum_{i=1}^n \frac{(i+1)^2}{(n+1)^2} s_i^2 \leq 8(CM)^2 K \frac{1}{(n+1)^2} \sum_{i=1}^n \frac{(i+1)^2}{i^2}. \quad (5.15)$$

Crudely bounding $\frac{(i+1)^2}{i^2} \leq 4$ for $i \geq 1$, we get

$$\mathcal{L}(u_n) - \mathcal{L}(z) \leq 32(CM)^2 K \frac{n}{(n+1)^2} \leq \frac{32(CM)^2 K}{n}. \quad (5.16)$$

Taking the infimum over z with $\|z\|_{\mathcal{K}_1(\mathbb{D})} \leq M$ gives the result. \square

5.2. The orthogonal greedy algorithm

The OGA only applies to function approximation, not to general convex optimization, and is given by

$$u_0 = 0, \quad g_n = \arg \max_{g \in \mathbb{D}} |\langle g, u_{n-1} - u \rangle_H|, \quad u_n = P_n(u), \quad (5.17)$$

where P_n is the orthogonal projection onto the span of g_1, \dots, g_n . Note here that the residual $u_{n-1} - u$ is the gradient $\nabla \mathcal{L}(u_{n-1})$ for the quadratic function $\mathcal{L}(u_{n-1}) = \frac{1}{2} \|u_{n-1} - u\|_H^2$. We remark that since this algorithm only applies to function approximation in a Hilbert space, our methods based upon the OGA can only be used to solve linear PDEs.

This algorithm was first analyzed in [24], where an $O(n^{-\frac{1}{2}})$ convergence rate is derived. Recently, it has been shown that this convergence rate can be significantly improved for the dictionaries whose convex hull $B_1(\mathbb{D})$ has small entropy [82]. In this section, we explain how to use the orthogonal greedy algorithm to solve linear PDEs and analyze the optimization error this induces.

When solving linear PDEs, the discretized energy function $\mathcal{R}_N(v)$ (or $\mathcal{R}_{n,\delta}(v)$) defined in (3.11) or (3.14) is a quadratic function of v . In particular, we have

$$\mathcal{R}_N(v) = \frac{1}{2} \left(\sum_{i=1}^N w_i a_0(x_i) v(x_i)^2 + \sum_{i=1}^N \sum_{|\alpha|=m} w_i a_\alpha(x_i) (\partial^\alpha v(x_i))^2 \right) - \sum_{i=1}^N w_i f(x_i) v(x_i), \quad (5.18)$$

with Neumann boundary conditions and an analogous expression with Dirichlet boundary conditions. In the following we assume that the quadrature weights $w_i > 0$. Then the loss $\mathcal{L} = \mathcal{R}_N(v)$ is equivalent to

$$\mathcal{R}_N(v) = \|I_{m,N}(v) - u_N\|_{a,N}^2, \quad (5.19)$$

where the evaluation map $I_{m,N} : \mathcal{C}^m(\Omega) \rightarrow \mathbb{R}^P$ is given by evaluating the function v and all derivatives of order m at the quadrature points x_i . Specifically, this map is given by

$$(I_{m,N}(v))_{\alpha,i} = \partial^\alpha v(x_i), \quad (5.20)$$

where the index set (α, i) runs over all multi-indices α such that either $|\alpha| = 0$ (the terms with no derivatives) or $|\alpha| = m$ and indices $i = 1, \dots, N$. Consequently $P = N \times N \binom{m+d-1}{d-1}$. The norm $\|\cdot\|_{a,N}$ on \mathbb{R}^P is given by the weighted norm

$$\|x\|_{a,N}^2 = \sum_{(\alpha,i)} w_i a_\alpha(x_i) x_{(\alpha,i)}^2. \quad (5.21)$$

Finally, u_N is the minimizer of the quadratic (5.18) in \mathbb{R}^P . Specifically, the components of u_N are given by

$$(u_N)_{(\alpha,i)} = \begin{cases} 0 & |\alpha| = m \\ a_0(x_i)^{-1} f(x_i) & |\alpha| = 0. \end{cases} \quad (5.22)$$

Crucially, u_N can be determined solely from knowledge of the right hand side f and the coefficients a_0 and does not require knowledge of the true solution u .

Using the orthogonal greedy algorithm to minimize the quadratic objective (5.19) results in the iteration

$$u_{0,N} = 0, \quad g_n = \arg \max_{g \in \mathbb{D}} |\langle I_{m,N}(g), u_{n-1,N} - u_N \rangle_{a,N}|, \quad u_{n,N} = P_n(u_N), \quad (5.23)$$

where the projection P_n is onto the span of the elements $I_{m,N}(g_1), \dots, I_{m,N}(g_n)$ with respect to the norm $\|\cdot\|_{a,N}$ on \mathbb{R}^P .

In a similar manner the PINNs risk (2.21) can be handled using the orthogonal greedy algorithm as long as the equation is linear. In this case the discretized risk is given by

$$\mathcal{R}_N(v) = \frac{1}{2} \left(\sum_{i=1}^N w_i \left(a_0(x_i) v(x_i) + \sum_{|\alpha|=m} a_\alpha(x_i) \partial^\alpha v(x_i) - f(x_i) \right)^2 \right). \quad (5.24)$$

This defines a quadratic function, and thus an inner product (possibly with kernel) on \mathbb{R}^P . We then maximize and project with respect to this inner product and the dictionary is embedded into \mathbb{R}^P via the map $I_{m,N}$, resulting in an analogous method to (5.23). Using the PINN risk allows us to tackle non-symmetric linear problems which may not have a variational formulation.

The estimation of the optimization error follows from the estimates derived in [82]. In particular, we quote the following theorem. Note that this theorem gives an upper bound and for certain dictionaries it is possible that the convergence rate of the OGA may be even faster.

Theorem 3. Let H be a Hilbert space and $\mathbb{D} \subset H$ a dictionary such that the metric entropy of the convex hull of \mathbb{D} satisfies

$$\epsilon_n(B_1(\mathbb{D}))_H \leq C n^{-\frac{1}{2}-\gamma} \quad (5.25)$$

for some $\gamma > 0$. Then for any $v \in \mathcal{K}_1(\mathbb{D})$, we have

$$\|u_n - u\|_H^2 \leq \|v - u\|_H^2 + K \|v\|_{\mathcal{K}_1(\mathbb{D})}^2 n^{-1-2\gamma}, \quad (5.26)$$

where K is a constant only depending upon C and γ .

Here the metric entropy $\epsilon_n(B_1(\mathbb{D}))_H$ is a measure of compactness of the set $B_1(\mathbb{D})$ with respect to the norm of H . For a precise definition and development of its properties, see for instance [50], Chapter 15. The important point is that the dictionary \mathbb{P}_k^d satisfies [83]

$$\epsilon_n(B_1(\mathbb{P}_k^d))_{H^m(\Omega)} \lesssim n^{-\frac{1}{2}-\frac{2(k-m)+1}{2d}}. \quad (5.27)$$

Thus, for this dictionary the value of γ in Theorem 3 is $\gamma = \frac{2(k-m)+1}{2d}$. When solving the discrete equation (5.19) using the orthogonal greedy algorithm it is important to note that up to logarithmic factors this entropy bound also holds in the $C^m(\Omega)$ -norm when $k = m + 1$ [5,83]. It is conjectured but not yet proven that this also holds for larger values of k . This means that since the evaluation map $I_{m,N} : C^m(\Omega) \rightarrow \mathbb{R}^P$ is bounded uniformly in N we have

$$\epsilon_n(I_{m,N}(B_1(\mathbb{P}_k^d)))_{a,N} \leq C n^{-\frac{1}{2}-\gamma} \quad (5.28)$$

holds uniformly in N for $\gamma = \frac{2(k-m)+1}{2d}$. Hence, denoting by $\bar{u}_{n,N} \in \Sigma_{n,\infty}(\mathbb{P}_k^d)$ the solution produced by the OGA at step n , we have for any M that

$$\mathcal{R}_N(\bar{u}_{n,N}) - \inf_{v \in \Sigma_{n,M}(\mathbb{P}_k^d)} \mathcal{R}_N(v) \leq \mathcal{R}_N(\bar{u}_{n,N}) - \inf_{\|v\|_{\mathcal{K}_1(\mathbb{P}_k^d)} \leq M} \mathcal{R}_N(v) \lesssim n^{-1-\frac{2(k-m)+1}{d}}. \quad (5.29)$$

This follows by taking the infimum over $\|v\|_{\mathcal{K}_1(\mathbb{P}_k^d)} \leq M$ in the conclusion of Theorem 3. This is precisely the optimization error bound we desire. Of course, this analysis applies to more general dictionaries \mathbb{D} as well, provided that the metric entropy $\epsilon_n(\mathbb{D})$ can be estimated.

Although the OGA attains the best convergence rate of the greedy algorithms, it is also the most computationally expensive since it requires an orthogonal projection at every step. In addition, it can only be applied to function approximation, which corresponds in our case to linear PDEs. A final drawback of the OGA is that the $\mathcal{K}_1(\mathbb{D})$ -norm of the iterates cannot be a priori bounded for general dictionaries as shown in [82]. As a result, we can only have a priori guarantee that the numerical solution satisfies $\bar{u}_{n,N} \in \Sigma_{n,\infty}(\mathbb{D})$. This means that our a priori generalization analysis only holds when using the RGA to optimize the empirical loss. Despite this, we have empirically observed the improved convergence rate of the OGA a posteriori and it significantly outperforms the RGA in our experiments.

6. Solving the argmax sub-problem

In order to implement the relaxed and orthogonal greedy algorithms, we need to be able to numerically solve the substep

$$g_n = \arg \max_{g \in \mathbb{D}} |\langle g, \nabla \mathcal{L}(u_{n-1}) \rangle|. \quad (6.1)$$

In fact, for the convergence analysis it is sufficient that the argmax in (6.1) is not solved exactly, but rather is approximated in the following sense

$$|\langle g_n, \nabla \mathcal{L}(u_{n-1}) \rangle| \geq \frac{1}{R} \max_{g \in \mathbb{D}} |\langle g, \nabla \mathcal{L}(u_{n-1}) \rangle| \quad (6.2)$$

for some fixed $R > 1$. This is a more tractable problem for most dictionaries.

6.1. Exactly solving the argmax sub-problem

We remark that in low dimensions and for certain dictionaries the argmax subproblem can be efficiently solved exactly. This is due to the fact that the objective

$$|\langle g, \nabla \mathcal{L}(u_{n-1}) \rangle| = \left| \sum_{i=1}^N \sum_{|\alpha|=m} (a_\alpha \partial^\alpha u_{n-1}(x_i), \partial^\alpha g(x_i)) + (a_0 u_{n-1}(x_i) - f(x_i), g(x_i)) \right| \quad (6.3)$$

is really a sum over a finite number N of quadrature points. In this case the set of possible hyperplane partitions of the quadrature points x_i can be enumerated and this can be used to exactly determine the arg max in (6.1). This algorithm is unfortunately intractable in high dimensions since its complexity scales as $O(N^d \log(N))$ [8, 82]. As a result, for higher dimensional problems we must resort to heuristics to approximate the subproblem (6.1) or consider different dictionaries for which this problem can be solved more efficiently. In our high dimensional numerical experiments, we use a special dictionary for which this argmax can be efficiently solved.

6.2. Numerical approximation of the argmax sub-problem

Next we describe the numerical heuristics we use in our experiments to approximately solve the argmax sub-problem in (6.1). Note that here the inner product in (6.1) is the energy inner product associated with the elliptic PDE we are solving. Our first step is to make the target function $|\langle g, \nabla \mathcal{L}(u_{n-1}) \rangle|$ differentiable, so we instead consider the following equivalent optimization problem:

$$g_n = \arg \min_{g \in \mathbb{D}} -\frac{1}{2} \langle g, \nabla \mathcal{L}(u_{n-1}) \rangle^2, \quad (6.4)$$

where

$$\langle \sigma(\omega \cdot x + b), \nabla \mathcal{L}(u_{n-1}) \rangle = \sum_{|\alpha|=m} (a_\alpha \partial^\alpha u_{n-1}, \partial^\alpha \sigma(\omega \cdot x + b)) + (a_0 u_{n-1} - f, \sigma(\omega \cdot x + b)). \quad (6.5)$$

Here we choose the dictionary $\mathbb{D} \subset \mathbb{R}^d$ as $\mathbb{D} = \mathbb{P}_k^d$, which is naturally parameterized by $\omega \in S^{d-1}$ and $b \in [-c, c]$ [80]. We also enforce the constraint $\|\omega\| = 1$ by taking $\omega = \pm 1$ for 1D case and $\omega = (\cos \theta, \sin \theta)$ based on the polar coordinates for 2D case. The low-dimensional optimization problem in (6.4) is typically non-convex so it may be very difficult to obtain the global minimum. Our approach is to obtain a good initial guess by choosing many samples initially on $\omega - b$ parameter space and evaluating the objective function at each of them. More specifically, we sample $b_i = -c + \frac{2ci}{N_b}$, $(i = 0, \dots, N_b)$, $w_0 = -1$, $w_1 = 1$ (1D case), and $\theta_j = \frac{2\pi j}{N_\theta}$, $(j = 0, \dots, N_\theta)$ (2D case) to find the best initial samples by evaluating (6.4) at each (b_i, w_j) . We then further optimize the best initial sample points using gradient descent or Newton's method. For the RGA, we optimize $g_n = \arg \min_{g \in \mathbb{D}} -\langle g, \nabla \mathcal{L}(u_{n-1}) \rangle$ instead of (6.1).

7. Uniform error bounds

In this section, we explain how to bound the discretization error

$$\sup_{f \in \mathcal{F}_\Theta} |\mathcal{R}_N(f) - \mathcal{R}(f)| \quad (7.1)$$

in Theorem 1 when solving elliptic PDEs. Recall that the loss function we consider in this work corresponds to the variational formulation of an elliptic PDE and is given in equation (3.7).

7.1. Uniform Monte Carlo error

The tool which we use to analyze the discretization error when the Monte Carlo discretization in equation (3.11) is used is the Rademacher complexity [8]. Given a class of functions $\mathcal{F} : \Omega \rightarrow \mathbb{R}$, and a collection of sample points $x_1, \dots, x_N \in \Omega$, the empirical Rademacher complexity of \mathcal{F} is defined by

$$\tilde{R}_N(\mathcal{F}) = \mathbb{E}_{\xi_1, \dots, \xi_N} \left[\sup_{h \in \mathcal{F}} \frac{1}{N} \sum_{i=1}^N \xi_i h(x_i) \right], \quad (7.2)$$

where ξ_1, \dots, ξ_N are Rademacher random variables, i.e. uniformly distributed signs. The Rademacher complexity is obtained by averaging over the samples x_i , which we take to be uniformly distributed over Ω , i.e. we have

$$R_N(\mathcal{F}) = \mathbb{E}_{x_1, \dots, x_N \sim \mu} \mathbb{E}_{\xi_1, \dots, \xi_N} \left[\sup_{h \in \mathcal{F}} \frac{1}{N} \sum_{i=1}^N \xi_i h(x_i) \right], \quad (7.3)$$

where μ is the uniform distribution on Ω . For the mixed boundary value problem, we will also need the Rademacher complexity with respect to the uniform distribution on the boundary $\partial\Omega$, which we denote by $R_{\partial, N}(\mathcal{F})$.

The utility of the Rademacher complexity is its role in giving a law of large numbers which is uniform over the class \mathcal{F} , detailed by the following theorem.

Theorem 4. [88, Proposition 4.11] Let \mathcal{F} be a set of functions. Then

$$\mathbb{E}_{x_1, \dots, x_N \sim \mu} \sup_{h \in \mathcal{F}} \left| \frac{1}{N} \sum_{i=1}^N h(x_i) - \int h(x) d\mu \right| \leq 2R_N(\mathcal{F}). \quad (7.4)$$

In order to apply Theorem 4 to the solution of PDEs via the class of Barron functions, we need to estimate the Rademacher complexity $R_N(\mathcal{L}_M)$ of the model class

$$\mathcal{L}_M = \{l(u, Du, \dots, D^m u) : u \in \mathcal{F}_{n, M}\}, \quad (7.5)$$

where the loss function l is given in equation (3.7) and the model class $\mathcal{F}_{n, M}$ is described in section 4. We remark that the Rademacher complexity of the Barron class $\mathcal{F}_{n, M}$ corresponding to shallow ReLU networks has been estimated in [56], so the novelty of our contribution is to generalize these bounds to the class \mathcal{L}_M obtained by composing with the loss function (3.7).

For this we will utilize the following fundamental lemma.

Lemma 2. Let \mathcal{F}, \mathcal{S} be classes of functions on Ω . Then the following bounds hold.

- $R_N(\text{conv}(\mathcal{F})) = R_N(\mathcal{F})$.
- Define the set $\mathcal{F} + \mathcal{S} = \{h(x) + g(x) : h \in \mathcal{F}, g \in \mathcal{S}\}$. We have

$$R_N(\mathcal{F} + \mathcal{S}) = R_N(\mathcal{F}) + R_N(\mathcal{S}). \quad (7.6)$$

- Suppose that $\phi : \mathbb{R} \rightarrow \mathbb{R}$ is L -Lipschitz. Let $\phi \circ \mathcal{F} = \{\phi(h(x)) : h \in \mathcal{F}\}$. Then

$$R_N(\phi \circ \mathcal{F}) \leq L R_N(\mathcal{F}). \quad (7.7)$$

- Suppose that $f : \Omega \rightarrow \mathbb{R}$ is a fixed function. Let $f \cdot \mathcal{F} = \{f(x)h(x) : h \in \mathcal{F}\}$. Then

$$R_N(f \cdot \mathcal{F}) \leq \|f(x)\|_{L^\infty(\Omega)} R_N(\mathcal{F}). \quad (7.8)$$

Proof. The first, second, and third of these statements are well-known facts, see [74, Lemma 26.7] for the first, [61, Page 56] for the second and [74, Lemma 26.9] for the third, so we only prove the fourth.

Suppose that $\|f(x)\|_{L^\infty(\Omega)} \leq 1$, the general result follows by a scaling argument. Let $x_1, \dots, x_N \in \Omega$ and consider the empirical Rademacher complexity

$$\tilde{R}_N(f \cdot \mathcal{F}) = \mathbb{E}_{\xi_1, \dots, \xi_N} \left[\sup_{h \in \mathcal{F}} \frac{1}{N} \sum_{i=1}^N \xi_i f(x_i) h(x_i) \right]. \quad (7.9)$$

We observe that the right-hand side of the above equation, being an average of a supremum of linear functions, is a convex function of $\vec{f} = (f(x_1), \dots, f(x_N))$. Consequently, its maximum must be achieved at the extreme points of the set $\{\vec{y} : \|\vec{y}\|_\infty \leq 1\}$, which correspond to the points where each component is ± 1 . Thus we only need to consider the case where $f(x_i) = \epsilon_i \in \{\pm 1\}$. But then

$$\mathbb{E}_{\xi_1, \dots, \xi_N} \left[\sup_{h \in \mathcal{F}} \frac{1}{N} \sum_{i=1}^N \xi_i \epsilon_i h(x_i) \right] = \mathbb{E}_{\xi_1, \dots, \xi_N} \left[\sup_{h \in \mathcal{F}} \frac{1}{N} \sum_{i=1}^N \xi_i h(x_i) \right] = \tilde{R}_N(\mathcal{F}), \quad (7.10)$$

since the ϵ_i simply permute the choices of sign ξ_i in the expectation. Taking an average over the sample points x_1, \dots, x_N completes the proof. \square

Utilizing this lemma, we prove the following bound on the Rademacher complexity of the set \mathcal{L}_M .

Theorem 5. Let $\mathbb{D} \subset C^k(\Omega)$ for $k \geq m$ be a dictionary. Suppose that $\|a_\alpha\|_{L^\infty(\Omega)}, \|a_0\|_{L^\infty(\Omega)} \leq K$ and $\sup_{d \in \mathbb{D}} \|d\|_{W^{m,\infty}} \leq C$. Then the Rademacher complexity of the set \mathcal{L}_M is bounded by

$$R_N(\mathcal{L}_M) \leq CKM \sum_{|\alpha|=m} R_N(\partial^\alpha \mathbb{D}) + CKMR_N(\mathbb{D}) + \|f\|_{L^\infty(\Omega)} MR_N(\mathbb{D}), \quad (7.11)$$

where $\partial^\alpha \mathbb{D} = \{\partial^\alpha d : d \in \mathbb{D}\}$.

Theorem 5 implies that to bound the Rademacher complexity of the set of interest, we only need to bound the Rademacher complexity of the derivatives of the dictionary \mathbb{D} , which is a much simpler task. In the Section 7.2 we will detail how to do this for the specific dictionaries corresponding to shallow neural networks.

Proof. The proof is a straightforward application of Lemma 2. We begin by noting that

$$\mathcal{L}_M \subset \sum_{|\alpha|=m} a_\alpha \cdot [\phi \circ B_M(\partial^\alpha \mathbb{D})] + a_0 \cdot [\phi \circ B_M(\mathbb{D})] + f \cdot B_M(\mathbb{D}), \quad (7.12)$$

where $\phi(x) = \frac{1}{2}x^2$ and $B_M(\mathbb{D}) = MB_1(\mathbb{D}) = \{f : \|f\|_{\mathcal{K}_1(\mathbb{D})} \leq M\}$.

Utilizing the first part of Lemma 2, we see that for all α

$$R_N(B_M(\partial^\alpha \mathbb{D})) \leq MR_N(\partial^\alpha \mathbb{D}). \quad (7.13)$$

The third part of the Lemma, combined with the bound $\|d\|_{W^{m,\infty}} \leq C$ and the fact that ϕ is locally Lipschitz, imply that

$$R_N(\phi \circ B_M(\partial^\alpha \mathbb{D})) \leq CMR_N(\partial^\alpha \mathbb{D}). \quad (7.14)$$

Finally, the second and fourth parts of the Lemma, combined with the bounds on a_α and a_0 complete the proof. \square

We are primarily interested in the following corollary of this result, which uniformly bound the Monte Carlo discretization error when discretizing elliptic PDEs. The next corollary provides a bound on the discretization error incurred in such a discretization.

Corollary 1. Suppose the empirical and true risk are defined as in (3.6) and (3.19) for the loss function (3.7) corresponding to the variational form of an elliptic PDE. Then, under the assumptions of Theorem 5, we have that

$$\mathbb{E}_{x_1, \dots, x_N} \sup_{v \in B_M(\mathbb{D})} |\mathcal{R}_N(v) - \mathcal{R}(v)| \leq 2CKM \sum_{|\alpha|=m} R_N(\partial^\alpha \mathbb{D}) + 2CKMR_N(\mathbb{D}) + 2\|f\|_{L^\infty(\Omega)} MR_N(\mathbb{D}). \quad (7.15)$$

Proof. This follows immediately by combining Theorem 5 with Theorem 4. \square

7.2. Rademacher bounds for neural networks

In this section, we show how the Rademacher complexity can be bounded for dictionaries corresponding to shallow neural networks. Specifically, we consider dictionaries of the form

$$\mathbb{D}_\sigma = \{\sigma(\omega \cdot x + b) : (\omega, b) \in \Theta\} \subset H^m(\Omega), \quad (7.16)$$

where the parameter set $\Theta \subset \mathbb{R}^{d+1}$ is compact. Of particular importance are the dictionaries corresponding to ReLU^k activation functions, \mathbb{P}_k^d , which were introduced in [78] and described in more detail in Section 4.1. Our main result is the following bound on the Rademacher complexity. This generalizes the results of [56], which calculate the Rademacher complexity of the unit ball in the Barron space for ReLU neural networks (see also [31], Theorem 2 and [39], Theorem 3).

Theorem 6. Suppose that $\sigma \in W^{m+1,\infty}$. Then for any α with $|\alpha| \leq m$, we have

$$R_N(\partial^\alpha \mathbb{D}) \lesssim N^{-\frac{1}{2}}, \quad R_{\partial,N}(\partial^\alpha \mathbb{D}) \lesssim N^{-\frac{1}{2}} \quad (7.17)$$

where the implied constant is independent of N .

Proof. This result follows immediately upon noting that

$$\partial^\alpha \mathbb{D} = \{\omega^\alpha \sigma^{(\alpha)}(\omega \cdot x + b) : (\omega, b) \in \Theta\}. \quad (7.18)$$

Since Θ is a compact set, $|\omega^\alpha|$ is bounded. Further, since $\sigma \in W^{m+1,\infty}$, we have that $\sigma^{(\alpha)}$ is Lipschitz. Using the third point in Lemma 2, we obtained

$$R_N(\partial^\alpha \mathbb{D}) \lesssim R_N(\{\omega \cdot x + b : (\omega, b) \in \Theta\}), \quad (7.19)$$

and likewise for $R_{\partial,N}(\partial^\alpha \mathbb{D})$.

It is well-known that the Rademacher complexity of the set of linear functions is bounded by [74, Section 26.2]

$$R_N(\{\omega \cdot x + b : (\omega, b) \in \Theta\}) \lesssim N^{-\frac{1}{2}}, \quad (7.20)$$

for any distribution on x which is bounded almost surely. This applies both to the uniform distribution on Ω as well as to the uniform distribution on $\partial\Omega$, which completes the proof. \square

7.3. Numerical quadrature

In this section we bound the discretization error when the Gauss-Legendre quadrature rule is used to compute the energy inner-product (6.4) and the error $\|u - u_n\|_a$. Let $\mathcal{T}_h \subset \Omega$ be a partition on Ω with mesh size h , where $h = \mathcal{O}(N^{-\frac{1}{d}})$ and N is the number of quadrature points. For each $T_l \in \mathcal{T}_h$, $l = 1, \dots, L$, the quadrature rule satisfies

$$\int_{T_l} p(x) dx = \sum_{i=0}^t p(x_{l,i}) \omega_i, \quad \forall p \in \mathcal{P}_{2t+1}(T_l), \quad (7.21)$$

where $\mathcal{P}_{2t+1}(T)$ is the space of polynomials with degree less equal than $2t+1$. Therefore, we have

$$\int_{\Omega} f(x) dx = \sum_{l=1}^L \int_{T_l} f(x) dx = \sum_{l=1}^L \sum_{i=0}^t f(x_{l,i}) \omega_i = \sum_{j=1}^N f(x_j) \omega_j, \quad \forall f \in \mathcal{P}_{2t+1}(\mathcal{T}_h), \quad (7.22)$$

where $\mathcal{P}_{2t+1}(\mathcal{T}_h) = \{g \in L^2(\Omega) : g|_T \in \mathcal{P}_{2t+1}(T), \forall T \in \mathcal{T}_h\}$ is the space of piece-wise polynomial functions on the partition \mathcal{T}_h . We define the error operator

$$E_t(f) = \int_{\Omega} f(x) dx - \sum_{j=1}^N f(x_j) \omega_j = \sum_{l=1}^L E_{t,l}(f) = \sum_{l=1}^L \left(\int_{T_l} f(x) dx - \sum_{i=0}^t f(x_{l,i}) \omega_i \right) \quad (7.23)$$

for $f \in W^{k+1,\infty}(\Omega)$. It is clear that $E_{t,l} \in (W^{k+1,\infty}(T_l))^*$ if $k \leq 2t+1$.

Theorem 7. Let $B_M(\mathbb{D}) = \{f \in \mathcal{K}_1(\mathbb{D}) : \|f\|_{\mathcal{K}_1(\mathbb{D})} \leq M\}$. Suppose the integrand $f \in B_M(\mathbb{D})$ where $\sup_d \|d\|_{W^{k+1,\infty}(\Omega)} \leq C$, and the Gauss-Legendre quadrature rule is accurate for \mathcal{P}_k . Then it holds that

$$|E_t(f)| \leq C_k C M N^{-\frac{k+1}{d}}, \quad (7.24)$$

where $t \geq \lceil \frac{k-1}{2} \rceil + 1$ and N is the number of quadrature points.

Proof. Using the Bramble-Hilbert Lemma, we get

$$|E_{n,\hat{T}}(\hat{f})| \leq C \|E_{n,\hat{T}}\|_{W^{k+1,\infty}(\hat{T})}^* |\hat{f}|_{W^{k+1,\infty}(\hat{T})} \quad (7.25)$$

on the reference domain \hat{T} . By the standard scaling argument, it gives on \mathcal{T}_h that

$$|E_n(f)| \leq C_k h^{k+1} \|f\|_{W^{k+1,\infty}(\Omega)} \leq C_k CMh^{k+1}. \quad (7.26)$$

The relation $h = \mathcal{O}(N^{-\frac{1}{d}})$ gives the result. \square

The accuracy with respect to N in (7.24) allows us to use the numerical quadrature (7.22) to compute the generalization errors such as $\|u - u_n\|_{L^2}$ and $\|u - u_n\|_a$. Since $\text{ReLU}^k(\omega \cdot x + b)$ is in the $W^{k,\infty}(\Omega)$ Sobolev space where Ω is bounded, then we have $\mathcal{K}_1(\mathbb{P}_k^d) \subset W^{k,\infty}(\Omega)$ which satisfies the condition of Theorem 7.

8. Balancing the error terms

In this section, we combine the estimates of the optimization, discretization, and modelling errors obtained in the previous sections to obtain a complete convergence theory and explain how to choose the hyperparameters N and n in each of the different situations discussed. Specifically, when using the relaxed greedy algorithm we have the following convergence theorem.

Theorem 8. Suppose that the Relaxed Greedy Algorithm (RGA) (5.2) is applied to the discretized loss function \mathcal{R}_N corresponding to the risk formulation (3.6) of the PDE (2.1). Suppose further that the solution u satisfies $\|u\|_{\mathcal{K}_1(\mathbb{D})} \leq M$ and that

- Monte Carlo quadrature is used and the assumptions of Theorem 5 are satisfied. If the dictionary \mathbb{D} satisfies $R_N(\partial^\alpha \mathbb{D}) \lesssim N^{-\frac{1}{2}}$ and we set $N = n^2$, we have the convergence rate

$$\mathcal{R}(u_{\Theta,N}) - \mathcal{R}(u) \lesssim n^{-1}. \quad (8.1)$$

In particular, since the objective error is comparable to the squared H^m -error, we also have

$$\|u_{\Theta,N} - u\|_{H^m(\Omega)} \lesssim n^{-\frac{1}{2}}. \quad (8.2)$$

- Numerical quadrature of order k is used and \mathbb{D} is uniformly bounded in $W^{k+1,\infty}(\Omega)$. If we set $N = n^{\frac{d}{2(k+1)}}$, then we have

$$\mathcal{R}(u_{\Theta,N}) - \mathcal{R}(u) \lesssim n^{-1}. \quad (8.3)$$

We also have

$$\|u_{\Theta,N} - u\|_{H^m(\Omega)} \lesssim n^{-\frac{1}{2}}. \quad (8.4)$$

Note in particular that the assumptions of this theorem hold when using shallow neural network dictionaries. We remark that when using the orthogonal greedy algorithm, the $\mathcal{K}_1(\mathbb{D})$ -norm cannot be a priori bounded and this is a missing ingredient in obtaining an a priori bound on the discretization error. Nonetheless, we obtain good performance in practice for the orthogonal greedy algorithm. In addition, if the solution u does not satisfy the bound $\|u\|_{\mathcal{K}_1(\mathbb{D})} \leq M$, then the method will nonetheless still optimize the risk over this set. In this case, a bound on the error can be obtained by determining how efficiently the solution can be approximated by a function u which satisfies $\|u\|_{\mathcal{K}_1(\mathbb{D})} \leq M$. The proper theory here is the theory of interpolation spaces (see for instance [23], Chapter 6, or [7]), but we do not go into detail here.

9. Numerical experiments

In this section, we provide numerical experiments demonstrating the effectiveness of the proposed algorithms on a variety of problems. For all the experiments below, the energy functions are discretized using Gaussian quadrature with the default setting $t = 2$ and $L = 4000$ in (7.22) for 1D and $t = 2 \times 2$, $L = 400 \times 400$ for 2D. For simplicity, we use u_n to denote the numerical solution and use u to represent the analytical solution in this section. We also define $\|u - u_n\|_a$ to be the discretization error in the energy norm which is defined in Section 5.2. For a specific type of second order elliptic equation discussed in (2.9), the energy norm is identical to the H^1 norm. In addition, the discretization error in the L^2 norm is also reported. For the detailed computation of these norms, we refer to the technique presented in Section 7.3 and the definition of norms (5.21). We remark that when calculating these norms we used a new and large set of quadrature points different from the ones used for training the network.

Section 9 is organized as follows. In Example 1 we test our method on a simple one-dimensional problem with both Dirichlet and Neumann boundary conditions. We do this with both the energy and PINN loss formulation of the problem and

Table 1
 L^2 and H^1 numerical error of OGA v.s. the number of neurons n for
Example 1.

n	$\ u - u_n\ _{L^2}$	order(n^{-3})	$\ u - u_n\ _{H^1}$	order(n^{-2})
16	7.86e-04	-	2.79e-02	-
32	7.70e-05	3.35	5.89e-03	2.24
64	8.45e-06	3.19	1.36e-03	2.11
128	9.68e-07	3.13	3.22e-04	2.08
256	1.18e-07	3.04	7.81e-05	2.04
512	1.44e-08	3.03	1.94e-05	2.01
1024	1.83e-09	2.97	4.86e-06	1.99
2048	2.50e-10	2.88	1.28e-06	1.93

Table 2
Numerical results of OGA for **Example 1** with Dirichlet boundary condition. Here we take $\delta = 0.1 \times n^{-2}$.

n	$\ u - u_n\ _{L^2}$	order	$\ u - u_n\ _{a,\delta}$	order(n^{-1})
16	6.72e-04	-	4.40e-02	-
32	1.70e-04	2.01	2.20e-02	1.00
64	4.17e-05	2.00	1.10e-02	1.00
128	1.04e-05	2.00	5.49e-03	1.00
256	2.63e-06	1.99	2.75e-03	1.00
512	8.10e-07	1.70	1.37e-03	1.00

compare our method with the common SGD, ADAM and L-BFGS optimizers to demonstrate its effectiveness. In Example 2 we present a 1D benchmark to verify the empirical adaptive property of greedy algorithms. Next we consider solving high order and high dimensional PDEs using the OGA. Examples 3 and 4 confirm our theoretical convergence rates for two dimensional elliptic problems with second and fourth order. In Example 5 we develop a method using a restricted dictionary designed for high dimensional problems. We show that our method can tackle high-dimensional problems as long as the solution is well-approximated by the convex hull of the dictionary. Finally, we give an example of non-linear PDEs in Section 9.2 using the relaxed greedy algorithm (RGA). We note that Theorem 2 holds for any convex and smooth energy function. Therefore, we can get convergence for non-linear equations provided the equation has a variation formulation with a convex energy.

9.1. Linear PDEs

Example 1 (1D elliptic equation). We consider the 1D elliptic equation

$$\begin{aligned} -u'' + u &= f, \quad x \in (-1, 1), \\ u'(-1) &= u'(1) = 0, \end{aligned} \tag{9.1}$$

with the source term $f = (1 + \pi^2) \cos(\pi x)$ which has the analytical solution $u(x) = \cos(\pi x)$. The energy function is discretized using Gaussian quadrature with $t = 2$ and $L = 4000$ in (7.22) and the discrete energy is minimized using the orthogonal greedy algorithm OGA with dictionary \mathbb{P}_2^1 (i.e. corresponding to ReLU²). The convergence rate is shown in Table 1. We obtain second order convergence in $H^1((-1, 1))$ which matches the theoretical convergence rate of the orthogonal greedy algorithm. In addition, we obtain third order convergence in $L^2((-1, 1))$ which matches the theoretically predicted approximation rates of shallow neural networks [83].

Next we consider the same equation with Dirichlet boundary conditions and consider the forcing term $f(x) = (1 + \frac{\pi^2}{4}) \cos(\frac{\pi}{2}x)$ so that the analytical solution is given by $u(x) = \cos(\frac{\pi}{2}x)$. We use the orthogonal greedy algorithm with $\mathbb{D} = \mathbb{P}_2^1$ to minimize a discretized version of the penalized energy $\mathcal{R}_{N,\delta}$, which is discretized using the same Gaussian quadrature. To balance the errors, we let δ scale as n^{-2} . The convergence order is given in Table 2 and matches the expected rate obtained by combining the convergence order of the orthogonal greedy algorithm with the error incurred by the penalization.

We also use the first example with Neumann's boundary conditions to compare with the deep Ritz method [89] using SGD and ADAM [41] as the optimizers. The numerical solution of the deep Ritz method is represented by a single hidden layer neural network with ReLU² activation function. We run both SGD and ADAM optimizers for 10000 epochs using Gauss quadrature points with random initialization. The initial learning rate for each experiment is 1×10^{-3} and is decreased by 5 every 3000 epochs. The numerical errors shown in Table 3 are the average results of 30 independent experiments, where we can see that both SGD and ADAM do not achieve any convergence order numerically as n gets larger (i.e. the size of the network gets larger).

Next, we compare with the widely used PINN method [70] which has been proven exceptionally successful in practical engineering applications. Specifically, we optimize a discretization of the PINN risk (2.21), given by

Table 3

The numerical convergence test of the deep Ritz method with both Adam and SGD optimizers on one-hidden-layer ReLU² neural network v.s. the number of neurons n for **Example 1**.

n	Adam				SGD			
	$\ u - u_n\ _{L^2}$	order	$\ u - u_n\ _{H^1}$	order	$\ u - u_n\ _{L^2}$	order	$\ u - u_n\ _{H^1}$	order
16	1.61e-02	-	1.45e-01	-	1.30e-02	-	1.52e-01	-
32	3.71e-03	2.12	5.84e-02	1.32	9.35e-03	0.47	1.13e-01	0.43
64	1.80e-03	1.04	3.46e-02	0.76	7.11e-03	0.39	8.64e-02	0.38
128	5.52e-04	1.70	1.43e-02	1.27	5.91e-03	0.27	7.22e-02	0.26
256	2.26e-04	1.29	6.99e-03	1.04	5.75e-03	0.04	7.03e-02	0.04
512	1.88e-04	0.27	3.90e-03	0.84	4.41e-03	0.38	5.40e-02	0.38
1024	2.09e-04	-0.16	2.56e-03	0.61	1.52e-03	1.54	1.99e-02	1.43
2048	4.11e-04	-0.97	2.51e-03	0.03	3.22e-03	-1.09	3.56e-02	-0.84

Table 4

The numerical convergence test of the PINN model with L-BFGS optimizer on one-hidden-layer ReLU³ neural network v.s. the number of neurons n for **Example 1**.

n	PINN-loss	order	$\ u - u_n\ _{L^2}$	order	$\ u - u_n\ _{H^1}$	order
16	9.19e-02	-	5.08e-03	-	3.17e-02	-
32	7.65e-02	0.27	4.11e-03	0.31	2.54e-02	0.32
64	1.86e-03	5.37	1.44e-04	4.84	1.62e-03	3.97
128	1.84e-04	3.33	5.20e-05	1.47	3.02e-04	2.43
256	1.13e-05	4.03	5.22e-06	3.32	4.50e-05	2.74
512	5.58e-06	1.02	3.58e-06	0.54	3.10e-05	0.54
1024	3.28e-05	-2.55	1.73e-05	-2.27	1.28e-04	-2.04
2048	1.52e-05	1.11	1.30e-05	0.42	8.15e-05	0.65

Table 5

The loss function and numerical error of OGA in L^2 and H^1 norms v.s. the number of neurons n for **Example 1**.

n	PINN-loss	order	$\ u - u_n\ _{L^2}$	order	$\ u - u_n\ _{H^1}$	order
16	2.64e-03	-	5.05e-04	-	2.02e-03	5.28
32	1.50e-04	4.14	1.04e-04	2.29	2.54e-04	2.99
64	8.10e-06	4.21	2.28e-05	2.18	4.43e-05	2.52
128	5.10e-07	3.99	5.03e-06	2.18	1.26e-05	1.81
256	3.16e-08	4.01	3.49e-06	0.53	5.40e-06	1.22
512	1.98e-09	4.00	5.45e-07	2.68	1.76e-06	1.62
1024	1.11e-10	4.16	1.81e-07	1.59	5.74e-07	1.62
2048	5.54e-12	4.32	6.67e-08	1.44	2.14e-07	1.43

$$MSE = MSE_f + MSE_{bc}, \quad (9.2)$$

where

$$MSE_f = \frac{1}{N_f} \sum_{i=1}^{N_f} |\Delta u_n(x_i) + u_n(x_i) - f(x_i)|^2 \text{ and } MSE_{bc} = |u'(-1)|^2 + |u'(1)|^2. \quad (9.3)$$

Here the collocation points $\{x_i\}_{i=1}^{N_f}$ are randomly chosen from the uniform distribution on $[-1, 1]$ where $N_f = 10000$. The numerical solution u_n is computed by optimizing the MSE loss with two stages. First the network is trained using ADAM's optimizer up to 10000 steps. In the next stage we change the optimizer into L-BFGS, where the learning rate is determined by the line search with the strong Wolfe's condition, and stop when the update of MSE loss is less than 10^{-16} . Since the gradient based training method requires the computation of first order derivatives, we change the activation function into ReLU³ to satisfy the regularity demand. The result shown in Table 4 is the mean error of 30 independent experiments. Similar to the results of the Deep Ritz method, we do not observe any stable numerical convergence with respect to n .

For comparison, next we use the orthogonal greedy algorithm to train the neural network using the same loss function (9.3) with $N_f = 10000$ and the dictionary \mathbb{P}_3^1 . We compute the numerical errors using the quadrature with a number of points that is large enough to get a good accuracy. We observe convergence for both the loss function and the numerical error in Table 5:

Example 2 (Adaptivity in 1D). Next, we test the OGA using the dictionary \mathbb{P}_2^1 on the 1D elliptic equation (9.1) where f is chosen so that the exact solution is given by:

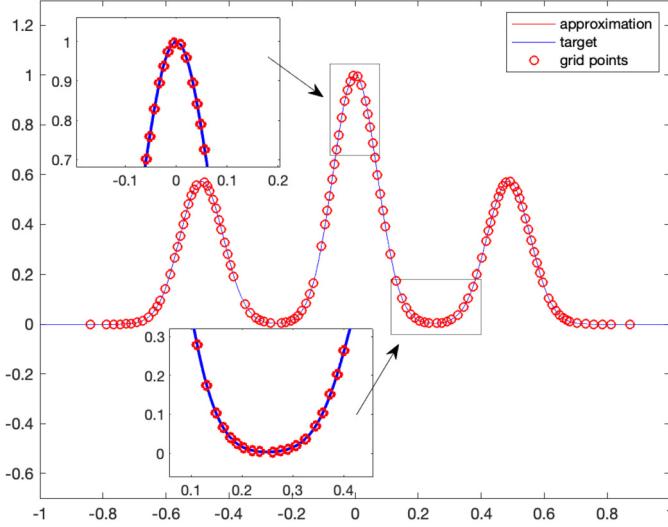


Fig. 1. Grid points of a 1-hidden layer neural network solution with $N = 128$ for **Example 2**.

Table 6
 L^2 and H^1 numerical errors and convergence orders of OGA for
Example 2.

n	$\ u - u_n\ _{L^2}$	order(n^{-3})	$\ u - u_n\ _{H^1}$	order(n^{-2})
16	5.05e-02	-	1.43e+00	-
32	1.96e-03	4.69	1.62e-01	3.14
64	2.08e-04	3.23	3.93e-02	2.05
128	1.99e-05	3.39	8.42e-03	2.22
256	2.34e-06	3.09	2.04e-03	2.04
512	2.85e-07	3.04	4.83e-04	2.08

$$u(x) = (1+x)^2(1-x^2) \left(0.5 \exp\left(-\frac{(x+0.5)^2}{K}\right) + \exp\left(-\frac{x^2}{K}\right) + 0.5 \exp\left(-\frac{(x-0.5)^2}{K}\right) \right), \quad (9.4)$$

for $x \in \Omega = (-1, 1)$ and $K = 0.01$. The exact solution has three peaks as shown in Fig. 1. In this example, we illustrate the adaptivity of the neural network discretization by identifying the grid points $x = (x_1, \dots, x_N)^T$ such that $w_1 x + b_1 = 0$, i.e. where the second derivative of the numerical solution changes. Since $u(x)$ has three peaks, we see that the grid points are gathered mainly at places with a larger curvature and are adaptive to fit the three peaks shown in Fig. 1. Furthermore, both the numerical error and the convergence order are shown in Table 6, where we see the theoretical convergence order achieved numerically. Note that the adaptivity is mainly a result of the neural network function class we are using and is likely to be present for other training algorithms as well. This merely demonstrates that a greedy algorithm is able to adapt to sharp changes in the solution.

Example 3 (2D elliptic equation). We consider the 2D elliptic equation in $\Omega = (0, 1)^2$ given by

$$\begin{aligned} -\Delta u + u &= f, \quad x \in (0, 1)^2 \\ \frac{\partial u}{\partial n} &= 0, \quad x \in \partial(0, 1)^2, \end{aligned} \quad (9.5)$$

where the right hand side f is chosen so that the exact solution is given by $u(x, y) = \cos(2\pi x) \cos(2\pi y)$. We discretize the energy using Gaussian quadrature of order 2 with 400 points in each direction. We optimize the discrete energy using the orthogonal greedy algorithm with the dictionary \mathbb{P}_2^2 . The convergence orders with both L^2 and H^1 errors are shown in Table 7 and confirm the theoretical orders of 1.75 and 1.25 for L^2 and H^1 errors, respectively. Note that the convergence order appears even to be slightly better than predicted by our theory. This demonstrates that we have only proved an upper bound, and for certain dictionaries the convergence rate of the orthogonal greedy algorithm may even be faster than predicted by Theorem 3. Due to the computational difficulty of solving the argmax subproblem (6.1) to the required high degree of accuracy, we were not able to run this example beyond 356 neurons with the variational loss.

Next we consider the dictionary \mathbb{P}_3^2 and optimize the PINN formulation instead of the energy formulation of the problem. Specifically, we optimize a discretization of the PINN risk (2.21), given by

$$MSE = MSE_f + MSE_{bc},$$

Table 7

Convergence order test of OGA with both L^2 and H^1 errors for **Example 3**.

n	$\ u - u_n\ _{L^2}$	order($n^{-1.75}$)	$\ u - u_n\ _{H^1}$	order($n^{-1.25}$)
16	5.13e-02	-	9.74e-01	-
32	9.72e-03	2.40	3.07e-01	1.66
64	2.26e-03	2.10	1.07e-01	1.53
128	5.86e-04	1.95	4.04e-02	1.40
256	1.42e-04	2.04	1.51e-02	1.42
356	7.68e-05	1.87	9.82e-03	1.30

Table 8

The loss function and numerical error of OGA in L^2 and H^1 norms v.s. the number of neurons n for **Example 3** using the PINN loss.

n	PINN-loss	order	$\ u - u_n\ _{L^2}$	order	$\ u - u_n\ _{H^1}$	order
16	9.51e+01	-	2.93e-01	-	1.41e+00	-
32	1.34e+01	2.83	7.04e-02	2.06	4.09e-01	1.79
64	1.79e+00	2.90	1.49e-02	2.24	1.09e-01	1.91
128	1.91e-01	3.23	4.67e-03	1.68	3.78e-02	1.52
256	2.67e-02	2.84	5.88e-04	2.99	8.13e-03	2.22
512	3.27e-03	3.03	5.60e-04	0.07	2.17e-03	1.90
1024	5.33e-04	2.62	5.22e-04	0.10	7.24e-04	1.58
2048	7.96e-05	2.74	9.43e-05	2.47	2.08e-04	1.80

where MSE_f is the discrete L^2 -residual in the domain $(0, 1)^2$:

$$MSE_f = \frac{1}{N_f} \sum_{i=1}^{N_f} | -\Delta u_n(x_i^f) + u_n(x_i^f) - f(x_i^f) |^2, \quad (9.6)$$

and MSE_{bc} is the residual on the boundary $\partial(0, 1)^2$:

$$MSE_{bc} = \frac{1}{N_{bc}} \sum_{j=1}^{N_{bc}} \left| \frac{\partial u_n}{\partial n}(x_j^{bc}) \right|^2. \quad (9.7)$$

Here we take $N_f = 20000$ and 2000 samples on each edge of $\partial(0, 1)^2$ so that $N_{bc} = 8000$. The samples $\{x_i^f\}_{i=1}^{N_f}$ and $\{x_j^{bc}\}_{j=1}^{N_{bc}}$ are randomly chosen in the corresponding domains from the uniform distribution. The following table shows the numerical result where the neural network is trained by OGA. We compute the numerical errors using the quadrature with a number of points large enough to get a good accuracy. We see that although the PINN loss converges with a good order as expected, the solution errors converge somewhat more slowly and less reliably than the loss. However, a good accuracy is nonetheless finally obtained even in terms of the solution error. (See Table 8.)

Example 4 (2D fourth-order differential equation). We consider the fourth-order equation

$$\begin{aligned} \Delta^2 u + u &= f, \quad x \in (-1, 1)^2, \\ B_N^0(u) &= 0, \quad x \in \partial(-1, 1)^2, \\ B_N^1(u) &= 0, \quad x \in \partial(-1, 1)^2. \end{aligned} \quad (9.8)$$

We choose the right hand side so that the exact solution is $u(x, y) = (x^2 - 1)^4(y^2 - 1)^4$. We discretize the energy using Gaussian quadrature of order 2 with 400 points in each direction and using the orthogonal greedy algorithm with the dictionary \mathbb{P}_3^2 to optimize the discrete energy. We plot the convergence orders for the L^2 energy norms in Table 9. Each of these errors is calculated by using finer Gaussian quadrature. For this example, we were again only able to run the algorithm with 256 neurons due to the computational difficulty of the argmax subproblem (6.1).

Example 5 (High-dimensional example). We consider the following 10d elliptic equation:

$$\begin{aligned} -\nabla \cdot (\alpha \nabla u) + u &= f, \quad x \in (0, 1)^{10} \\ \frac{\partial u}{\partial n} &= 0, \quad x \in \partial(0, 1)^{10}, \end{aligned} \quad (9.9)$$

with

Table 9

The convergence order of OGA with $\|\cdot\|_{L^2}$ and $\|\cdot\|_a$ errors for **Example 4**.

n	$\ u - u_n\ _{L^2}$	order($n^{-2.25}$)	$\ u - u_n\ _a$	order($n^{-1.25}$)
16	1.72e-01	-	4.84e+00	-
32	1.89e-02	3.18	1.60e+00	1.59
64	3.36e-03	2.50	5.85e-01	1.46
128	4.24e-04	2.99	2.04e-01	1.52
256	8.25e-05	2.36	8.19e-02	1.32

Table 10

The convergence order of OGA on a high-dimensional problem with $\|\cdot\|_{L^2}$ and $\|\cdot\|_{H^1}$ errors for **Example 5**.

n	$\ u - u_n\ _{L^2}$	order(n^{-3})	$\ u - u_n\ _{H^1}$	order(n^{-2})
16	5.02e-01	-	3.18e+00	-
32	4.70e-02	3.42	5.99e-01	2.41
64	4.63e-03	3.34	1.10e-01	2.44
128	4.44e-04	3.38	2.27e-02	2.28
256	5.41e-05	3.04	5.19e-03	2.13

$$\alpha = \sqrt{1 + \sum_{i=1}^{10} (x_i - \frac{1}{2})^2}. \quad (9.10)$$

We choose the right hand side f so that the exact solution is given by

$$u = \sum_{i=1}^{10} \cos(\pi x_i). \quad (9.11)$$

In order to be able to solve the argmax problem arising in (6.1) we use the restricted dictionary

$$\mathbb{P}_2^{10,r} = \{\sigma(\omega \cdot x + b), \omega = \pm e_i, i = 1, \dots, 10, b \in [-2, 2]\}, \quad (9.12)$$

where $\sigma = \text{ReLU}^2$. We note that the solution u was specifically chosen to lie in the convex hull of $\mathbb{P}_2^{10,r}$, which is given by

$$B_1(\mathbb{P}_2^{10,r}) = \left\{ f(x) = \sum_{i=1}^{10} f_i(x_i), \sum_{i=1}^{10} \|f_i\|_{BV}^2 \leq 1 \right\}. \quad (9.13)$$

Note that the equation itself is not separable due to the complicated coefficients α , and all that is required for the method to work is that the solution be well approximated by the dictionary. We discretize the energy using 100 million quasi-Monte-Carlo samples in $[0, 1]^{10}$ generated by the Halton sequence, and optimize the energy using the orthogonal greedy algorithm with dictionary $\mathbb{P}_2^{10,r}$. The results are shown in Table 10. The point of this example is to demonstrate that the proposed method converges as expected even in high-dimensions as long as the solution is well-approximated by the dictionary \mathbb{D} . For this example we were only able to run the algorithm for 256 iterations due to the very large number of quasi-Monte Carlo samples required for the high dimensional problem.

9.2. Nonlinear PDEs

Next, we test the convergence order of the RGA on a nonlinear Poisson-Boltzmann PDE to confirm the theoretically derived first order convergence in Theorem 2. We also test both sigmoid and ReLU² activation functions and compare RGA with OGA in the 1D example. For all the RGA's results, we report the generalization error defined in the left hand side of (3.17) in a relative sense, i.e., $\frac{\mathcal{R}(u_0) - \mathcal{R}(u)}{\mathcal{R}(u_0) - \mathcal{R}(\bar{u})}$, which is computed by using the numerical quadrature scheme.

Example 6 (2D Poisson-Boltzmann equation, [47]). We consider the 2D Poisson-Boltzmann equation on the sphere $\{(x, y) | x^2 + y^2 \leq 4\}$, with Neumann boundary conditions, namely,

$$\begin{cases} -\Delta u + \kappa \sinh(u) = f, & (x, y) \in \Omega = B_2(0), \\ \frac{\partial u}{\partial n} = 0, & (x, y) \in \partial B_2(0). \end{cases} \quad (9.14)$$

The energy functional for this problem is

Table 11
Convergence order of RGA for the 2D Poisson-Boltzmann equation in **Example 6**.

n	$\frac{\mathcal{R}(u_n) - \mathcal{R}(u)}{\mathcal{R}(u_0) - \mathcal{R}(u)}$	order(n^{-1})
16	8.18e+00	-
32	4.19e+00	0.96
64	2.96e+00	0.50
128	6.95e-01	2.09
256	2.54e-01	1.45
512	7.70e-02	1.72
1024	2.90e-02	1.41
2048	1.39e-02	1.06

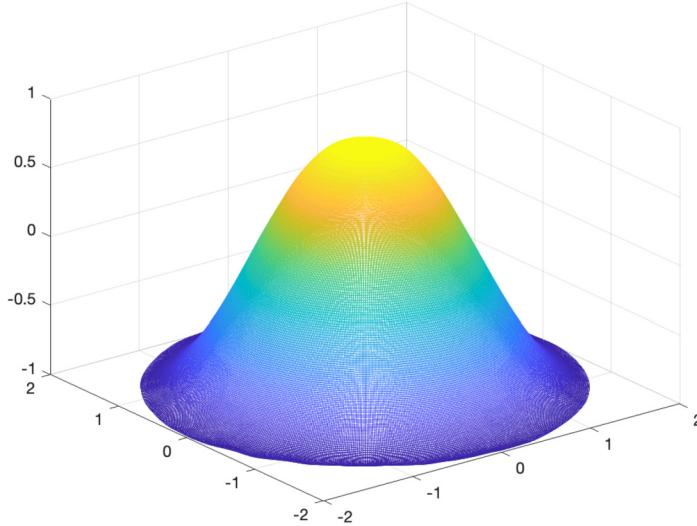


Fig. 2. Numerical solution of the 2D Poisson-Boltzmann equation obtained by RGA with $n = 4096$ for **Example 6**.

$$\mathcal{R}(u) = \int_{\Omega} \left(\frac{1}{2} |\nabla u|^2 + \kappa \cosh(u) - fu \right) dx,$$

which is a strictly convex and coercive energy with respect to u as long as $\kappa > 0$ and this implies the existence and uniqueness of the solution. We set $\kappa = 1$ and consider the radially symmetric solution $u(x, y) = \cos(\frac{\pi}{2} \sqrt{x^2 + y^2})$, which gives the source terms f . We use Monte-Carlo quadrature with the number of samples $N = O(n^2) = \frac{n^2}{10}$ to approximate the integration. The dictionary for the RGA algorithm is taken as

$$\mathbb{D} = \{\sigma(w_1 x + w_2 y + b) | (w_1, w_2, b) \in [-20, 20]^3\},$$

where σ is the sigmoidal activation function and we set $M = 20$ in (5.2). The convergence order test is shown in Table 11 and the numerical solution is plotted in Fig. 2.

10. Conclusion

The process of training neural networks is the main bottleneck in applying neural networks to solve PDEs, both in terms of the effort required to tune hyperparameters and in the computational complexity required for the training process. In order to solve the resulting highly non-convex optimization problems, typically SGD or ADAM is used to train the neural networks. Despite their breakthrough empirical performance, these algorithms are often difficult to properly tune and require multiple tries and additional tricks to obtain good performance. Moreover, it is often difficult or impossible to exhibit convergence as the network size increases. In order to overcome this problem, we develop a greedy training algorithm for neural network discretizations of PDEs, which we theoretically prove will exhibit convergence to the solution as network size increases. Empirical experiments validate the theory and exhibit the predicted convergence. This demonstrates that neural network can, at least in principle, be used to rigorously solve PDEs. One drawback of the approach is the large amount of computational effort required in each greedy step. Improving the efficiency of this step is an ongoing research problem.

CRediT authorship contribution statement

Jonathan W. Siegel: Conceptualization, Formal analysis, Investigation, Methodology, Supervision, Writing – original draft. **Qingguo Hong:** Formal analysis, Writing – original draft. **Xianlin Jin:** Investigation, Software, Writing – original draft. **Wenrui Hao:** Supervision, Writing – review & editing. **Jinchao Xu:** Conceptualization, Formal analysis, Funding acquisition, Project administration, Resources, Supervision, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

References

- [1] Mark Ainsworth, Yeonjong Shin, Active neuron least squares: a training method for multivariate rectified neural networks, *SIAM J. Sci. Comput.* 44 (4) (2022) A2253–A2275.
- [2] Zeyuan Allen-Zhu, Yuanzhi Li, Zhao Song, A convergence theory for deep learning via over-parameterization, in: *International Conference on Machine Learning*, PMLR, 2019, pp. 242–252.
- [3] Amine Ammar, et al., A new family of solvers for some classes of multidimensional partial differential equations encountered in kinetic theory modeling of complex fluids, *J. Non-Newton. Fluid Mech.* 139 (3) (2006) 153–176.
- [4] Sanjeev Arora, et al., Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks, in: *International Conference on Machine Learning*, PMLR, 2019, pp. 322–332.
- [5] Francis Bach, Breaking the curse of dimensionality with convex neural networks, *J. Mach. Learn. Res.* 18 (1) (2017) 629–681.
- [6] Andrew R. Barron, Universal approximation bounds for superpositions of a sigmoidal function, *IEEE Trans. Inf. Theory* 39 (3) (1993) 930–945.
- [7] Andrew R. Barron, et al., Approximation and learning by greedy algorithms, *Ann. Stat.* 36 (1) (2008) 64–94.
- [8] Peter L. Bartlett, Shahar Mendelson, Rademacher and Gaussian complexities: risk bounds and structural results, *J. Mach. Learn. Res.* 3 (2002) 463–482.
- [9] Julius Berner, Philipp Grohs, Arnulf Jentzen, Analysis of the generalization error: empirical risk minimization over deep artificial neural networks overcomes the curse of dimensionality in the numerical approximation of Black–Scholes partial differential equations, *SIAM J. Math. Data Sci.* 2 (3) (2020) 631–657.
- [10] Shengze Cai, et al., Physics-informed neural networks (PINNs) for fluid mechanics: a review, *Acta Mech. Sin.* (2022) 1–12.
- [11] Shengze Cai, et al., Physics-informed neural networks for heat transfer problems, *J. Heat Transf.* 143 (2021) 6.
- [12] Eric Cancès, Virginie Ehrlacher, Tony Lelièvre, Greedy algorithms for high-dimensional non-symmetric linear problems, in: *ESAIM: Proceedings*, vol. 41, EDP Sciences, 2013, pp. 95–131.
- [13] Emmanuel Jean Candes, *Ridgelets: Theory and Applications*, Stanford University, 1998.
- [14] Shuhao Cao, Choose a transformer: Fourier or Galerkin, *Adv. Neural Inf. Process. Syst.* 34 (2021) 24924–24940.
- [15] Giuseppe Carleo, Matthias Troyer, Solving the quantum many-body problem with artificial neural networks, *Science* 355 (6325) (2017) 602–606.
- [16] Qipin Chen, Wenrui Hao, A randomized Newton's method for solving differential equations based on the neural network discretization, preprint, arXiv:1912.03196, 2019.
- [17] Ziang Chen, Jianfeng Lu, Yulong Lu, On the representation of solutions to elliptic pdes in Barron spaces, *Adv. Neural Inf. Process. Syst.* 34 (2021).
- [18] Ziang Chen, et al., A regularity theory for static Schrödinger equations on \mathbb{R}^d in spectral Barron spaces, preprint, arXiv:2201.10072, 2022.
- [19] Ingrid Daubechies, et al., Nonlinear approximation and (deep) ReLU networks, in: *Constructive Approximation*, 2021, pp. 1–46.
- [20] Tim De Ryck, Ameya D. Jagtap, Siddhartha Mishra, Error estimates for physics informed neural networks approximating the Navier–Stokes equations, preprint, arXiv:2203.09346, 2022.
- [21] Tim De Ryck, Siddhartha Mishra, Error analysis for physics informed neural networks (PINNs) approximating Kolmogorov PDEs, preprint, arXiv:2106.14473, 2021.
- [22] Anton Dereventsov, Armenak Petrosyan, Clayton Webster, Greedy shallow networks: a new approach for constructing and training neural networks, preprint, arXiv:1905.10409, 2019.
- [23] Ronald A. DeVore, George G. Lorentz, *Constructive Approximation*, vol. 303, Springer Science & Business Media, 1993.
- [24] Ronald A. DeVore, Vladimir N. Temlyakov, Some remarks on greedy algorithms, *Adv. Comput. Math.* 5 (1) (1996) 173–187.
- [25] Ronald DeVore, Boris Hanin, Guergana Petrova, Neural network approximation, preprint, arXiv:2012.14501, 2020.
- [26] M.W.M.G. Dissanayake, Nhan Phan-Thien, Neural-network-based approximations for solving partial differential equations, *Commun. Numer. Methods Eng.* 10 (3) (1994) 195–201.
- [27] Simon S. Du, et al., Gradient descent provably optimizes over-parameterized neural networks, in: *International Conference on Learning Representations*, 2018.
- [28] Simon Du, et al., Gradient descent finds global minima of deep neural networks, in: *International Conference on Machine Learning*, PMLR, 2019, pp. 1675–1685.
- [29] Chengguang Duan, et al., Convergence rate analysis for deep Ritz method, preprint, arXiv:2103.13330, 2021.
- [30] Leonardo E. Figueroa, Endre Süli, Greedy approximation of high-dimensional Ornstein–Uhlenbeck operators, *Found. Comput. Math.* 12 (5) (2012) 573–623.
- [31] Wei Gao, Zhi-Hua Zhou, Dropout Rademacher complexity of deep neural networks, *Sci. China Inf. Sci.* 59 (7) (2016) 1–12.
- [32] Philipp Grohs, et al., A proof that artificial neural networks overcome the curse of dimensionality in the numerical approximation of Black–Scholes partial differential equations, preprint, arXiv:1809.02362, 2018.
- [33] Jiequn Han, Arnulf Jentzen, E. Weinan, Solving high-dimensional partial differential equations using deep learning, *Proc. Natl. Acad. Sci.* 115 (34) (2018) 8505–8510.
- [34] Jan Hermann, Zeno Schätzle, Frank Noé, Deep-neural-network solution of the electronic Schrödinger equation, *Nat. Chem.* 12 (10) (2020) 891–897.
- [35] Wassily Hoeffding, Probability inequalities for sums of bounded random variables, in: *The Collected Works of Wassily Hoeffding*, Springer, 1994, pp. 409–426.

[36] Qingguo Hong, Jonathan W. Siegel, Jinchao Xu, A priori analysis of stable neural network solutions to numerical PDEs, preprint, arXiv:2104.02903, 2021.

[37] Qingguo Hong, et al., On the activation function dependence of the spectral bias of neural networks, preprint, arXiv:2208.04924, 2022.

[38] Lee K. Jones, A simple lemma on greedy approximation in Hilbert space and convergence rates for projection pursuit regression and neural network training, *Ann. Stat.* 20 (1) (1992) 608–613.

[39] Sham M. Kakade, Karthik Sridharan, Ambuj Tewari, On the complexity of linear prediction: Risk bounds, margin bounds, and regularization, 2008.

[40] Yuehaw Khoo, Jianfeng Lu, Lexing Ying, Solving parametric PDE problems with artificial neural networks, *Eur. J. Appl. Math.* 32 (3) (2021) 421–435.

[41] Diederik P. Kingma, Jimmy Ba, Adam: a method for stochastic optimization, preprint, arXiv:1412.6980, 2014.

[42] Jason M. Klusowski, Andrew R. Barron, Approximation by combinations of ReLU and squared ReLU ridge functions with ℓ^1 and ℓ^0 controls, *IEEE Trans. Inf. Theory* 64 (12) (2018) 7649–7656.

[43] Nikola Kovachki, Samuel Lanthaler, Siddhartha Mishra, On universal approximation and error bounds for Fourier neural operators, *J. Mach. Learn. Res.* 22 (1) (2021) 13237–13312.

[44] Samuel Lanthaler, Siddhartha Mishra, George Em Karniadakis, Error estimates for deepnets: a deep learning framework in infinite dimensions, preprint, arXiv:2102.09618, 2021.

[45] Claude Le Bris, Tony Lelievre, Yvon Maday, Results and questions on a nonlinear approximation approach for solving high-dimensional partial differential equations, *Constr. Approx.* 30 (3) (2009) 621–651.

[46] Wee Sun Lee, Peter L. Bartlett, Robert C. Williamson, Efficient agnostic learning of neural networks with bounded fan-in, *IEEE Trans. Inf. Theory* 42 (6) (1996) 2118–2132.

[47] Zhilin Li, C.V. Pao, Zhonghua Qiao, A finite difference method and analysis for 2D nonlinear Poisson–Boltzmann equations, *J. Sci. Comput.* 30 (1) (2007) 61–81 (English).

[48] Zongyi Li, et al., Fourier neural operator for parametric partial differential equations, preprint, arXiv:2010.08895, 2020.

[49] Marcello Longo, et al., Higher-order quasi-Monte Carlo training of deep neural networks, *SIAM J. Sci. Comput.* 43 (6) (2021) A3938–A3966.

[50] George G. Lorentz, Manfred v Golitschek, Yuly Makovoz, *Constructive Approximation: Advanced Problems*, vol. 304, Springer, 1996.

[51] Jianfeng Lu, et al., Deep network approximation for smooth functions, *SIAM J. Math. Anal.* 53 (5) (2021) 5465–5506.

[52] Lu Lu, et al., Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators, *Nat. Mach. Intell.* 3 (3) (2021) 218–229.

[53] Yulong Lu, Jianfeng Lu, Min Wang, A priori generalization analysis of the deep Ritz method for solving high dimensional elliptic partial differential equations, in: *Conference on Learning Theory*, PMLR, 2021, pp. 3196–3241.

[54] Tao Luo, Haizhao Yang, Two-layer neural networks for partial differential equations: optimization and generalization theory, preprint, arXiv:2006.15733, 2020.

[55] Tao Luo, et al., Theory of the frequency principle for general deep neural networks, preprint, arXiv:1906.09235, 2019.

[56] Chao Ma, Lei Wu, et al., The Barron space and the flow-induced function spaces for neural network models, *Constr. Approx.* 55 (1) (2022) 369–406.

[57] Stéphane G. Mallat, Zhifeng Zhang, Matching pursuits with time-frequency dictionaries, *IEEE Trans. Signal Process.* 41 (12) (1993) 3397–3415.

[58] Zhiping Mao, Ameya D. Jagtap, George Em Karniadakis, Physics-informed neural networks for high-speed flows, *Comput. Methods Appl. Mech. Eng.* 360 (2020) 112789.

[59] Siddhartha Mishra, Roberto Molinaro, Estimates on the generalization error of physics informed neural networks (PINNs) for approximating a class of inverse problems for PDEs, preprint, arXiv:2007.01138, 2020.

[60] Siddhartha Mishra, Roberto Molinaro, Estimates on the generalization error of physics-informed neural networks for approximating a class of inverse problems for PDEs, *IMA J. Numer. Anal.* 42 (2) (2022) 981–1022.

[61] Mehryar Mohri, Afshin Rostamizadeh, Ameet Talwalkar, *Foundations of Machine Learning*, MIT Press, 2018.

[62] Johannes Müller, Marius Zeinhofer, Error estimates for the deep Ritz method with boundary penalty, preprint, arXiv:2103.01007, 2021.

[63] Greg Ongie, et al., A function space view of bounded norm infinite width ReLU nets: the multivariate case, in: *International Conference on Learning Representations (ICLR 2020)*, 2019.

[64] Guofei Pang, Lu Lu, George Em Karniadakis, fPINNs: fractional physics-informed neural networks, *SIAM J. Sci. Comput.* 41 (4) (2019) A2603–A2626.

[65] Rahul Parhi, Robert D. Nowak, Banach space representer theorems for neural networks and ridge splines, preprint, arXiv:2006.05626, 2020.

[66] Rahul Parhi, Robert D. Nowak, What kinds of functions do deep neural networks learn? Insights from variational spline theory, preprint, arXiv:2105.03361, 2021.

[67] Yagya Chandra Pati, Ramin Rezaifar, Perinkulam Sambamurthy Krishnaprasad, Orthogonal matching pursuit: recursive function approximation with applications to wavelet decomposition, in: *Proceedings of 27th Asilomar Conference on Signals, Systems and Computers*, IEEE, 1993, pp. 40–44.

[68] Gilles Pisier, Remarques sur un résultat non publié de B. Maurey, in: *Séminaire Analyse fonctionnelle (dit “Maurey-Schwartz”)*, 1981, pp. 1–12.

[69] Nasim Rahaman, et al., On the spectral bias of neural networks, in: *International Conference on Machine Learning*, PMLR, 2019, pp. 5301–5310.

[70] Maziar Raissi, Paris Perdikaris, George 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) 686–707.

[71] Benjamin Recht, et al., Do CIFAR-10 classifiers generalize to CIFAR-10?, preprint, arXiv:1806.00451, 2018.

[72] Francisco Sahli Costabal, et al., Physics-informed neural networks for cardiac activation mapping, *Front. Phys.* 8 (2020) 42.

[73] Nihar Sawant, Boris Kramer, Benjamin Peherstorfer, Physics-informed regularization and structure preservation for learning stable reduced models from data with operator inference, preprint, arXiv:2107.02597, 2021.

[74] Shalev-Shwartz Shai, Ben-David Shai, *Understanding Machine Learning: From Theory to Algorithms*, Cambridge University Press, 2014.

[75] Zuowei Shen, Haizhao Yang, Shijun Zhang, Optimal approximation rate of ReLU networks in terms of width and depth, *J. Math. Pures Appl.* 157 (2022) 101–135.

[76] Yeonjong Shin, Jerome Darbon, George Em Karniadakis, On the convergence of physics informed neural networks for linear second-order elliptic and parabolic type PDEs, preprint, arXiv:2004.01806, 2020.

[77] Yeonjong Shin, Zhongqiang Zhang, George Em Karniadakis, Error estimates of residual minimization using neural networks for linear PDEs, preprint, arXiv:2010.08019, 2020.

[78] Jonathan W. Siegel, Jinchao Xu, Improved approximation properties of dictionaries and applications to neural networks, preprint, arXiv:2101.12365, 2021.

[79] Jonathan W. Siegel, Jinchao Xu, Approximation rates for neural networks with general activation functions, *Neural Netw.* 128 (2020) 313–321.

[80] Jonathan W. Siegel, Jinchao Xu, Characterization of the variation spaces corresponding to shallow neural networks, preprint, arXiv:2106.15002, 2021.

[81] Jonathan W. Siegel, Jinchao Xu, High-order approximation rates for neural networks with ReLU^k activation functions, preprint, arXiv:2012.07205, 2020.

[82] Jonathan W. Siegel, Jinchao Xu, Optimal convergence rates for the orthogonal greedy algorithm, *IEEE Trans. Inf. Theory* 68 (5) (2022) 3354–3361.

[83] Jonathan W. Siegel, Jinchao Xu, Sharp bounds on the approximation rates, metric entropy, and n -widths of shallow neural networks, preprint, arXiv: 2101.12365, 2021.

[84] Justin Sirignano, Konstantinos Spiliopoulos, DGM: a deep learning algorithm for solving partial differential equations, *J. Comput. Phys.* 375 (2018) 1339–1364.

[85] Gilbert Strang, *Variational crimes in the finite element method*, in: *The Mathematical Foundations of the Finite Element Method with Applications to Partial Differential Equations*, Elsevier, 1972, pp. 689–710.

- [86] Vladimir Temlyakov, Greedy Approximation, vol. 20, Cambridge University Press, 2011.
- [87] Vladimir N. Temlyakov, Greedy approximation, *Acta Numer.* 17 (235) (2008) 409.
- [88] Martin J. Wainwright, High-Dimensional Statistics: A Non-asymptotic Viewpoint, vol. 48, Cambridge University Press, 2019.
- [89] E. Weinan, Bing Yu, The deep Ritz method: a deep learning-based numerical algorithm for solving variational problems, *Commun. Math. Stat.* 6 (1) (2018) 1–12.
- [90] Stephan Wojtowytch, et al., Representation formulas and pointwise properties for Barron functions, *Calc. Var. Partial Differ. Equ.* 61 (2) (2022) 1–37.
- [91] Jinchao Xu, Finite neuron method and convergence analysis, *Commun. Comput. Phys.* (ISSN 1991-7120) 28 (5) (2020) 1707–1745, <https://doi.org/10.4208/cicp.OA-2020-0191>, http://global-sci.org/intro/article_detail/cicp/18394.html.
- [92] Dmitry Yarotsky, Error bounds for approximations with deep ReLU networks, *Neural Netw.* 94 (2017) 103–114.
- [93] Tong Zhang, Sequential greedy approximation for certain convex optimization problems, *IEEE Trans. Inf. Theory* 49 (3) (2003) 682–691.
- [94] Difan Zou, et al., Gradient descent optimizes over-parameterized deep ReLU networks, *Mach. Learn.* 109 (3) (2020) 467–492.