
Deep Network Approximation in Terms of Intrinsic Parameters

Zuowei Shen¹ Haizhao Yang² Shijun Zhang^{*1}

Abstract

One of the arguments to explain the success of deep learning is the powerful approximation capacity of deep neural networks. Such capacity is generally accompanied by the explosive growth of the number of parameters, which, in turn, leads to high computational costs. It is of great interest to ask whether we can achieve successful deep learning with a small number of learnable parameters adapting to the target function. From an approximation perspective, this paper shows that the number of parameters that need to be learned can be significantly smaller than people typically expect. First, we theoretically design ReLU networks with a few learnable parameters to achieve an attractive approximation. We prove by construction that, for any Lipschitz continuous function f on $[0, 1]^d$ with a Lipschitz constant $\lambda > 0$, a ReLU network with $n + 2$ intrinsic parameters (those depending on f) can approximate f with an exponentially small error $5\lambda\sqrt{d}2^{-n}$. Such a result is generalized to generic continuous functions. Furthermore, we show that the idea of learning a small number of parameters to achieve a good approximation can be numerically observed. We conduct several experiments to verify that training a small part of parameters can also achieve good results for classification problems if other parameters are pre-specified or pre-trained from a related problem.

1. Introduction

Deep neural networks have recently achieved great success in a large number of real-world applications. However, the success in deep neural networks is generally accompanied by the explosive growth of computation and parameters. This follows a natural problem: how to handle computationally expensive deep learning models with limited com-

puting resources. This problem is challenging and has been widely studied. Numerous model compression and acceleration methods have recently been proposed, e.g., parameter pruning and quantization (Han et al., 2015; Wu et al., 2016; Gupta et al., 2015), low-rank factorization (Denton et al., 2014; Jaderberg et al., 2014), transferred compact convolutional filters (Howard et al., 2017; Sandler et al., 2018; Qin et al., 2019), knowledge distillation (Hinton et al., 2015; Zhang et al., 2019; 2021). See a survey of these methods in (Cheng et al., 2017). This paper explores an approximation perspective for the problem mentioned above. We adopt the approximation perspective since the approximation power is a key ingredient for the performance of deep neural networks. In other words, The goal of this paper is to investigate how to reduce the number of parameters that need to be learned from an approximation perspective. We will provide both theoretical and numerical examples to show that adjusting only a small number of parameters is enough to achieve good results if the network architecture is properly designed.

To design a simple and computable hypothesis function ϕ to approximate a target function $f \in \mathcal{F}$ well via adjusting only a small number of f -dependent parameters, where \mathcal{F} is a given target function space, we have the following two main ideas, motivated by the power of deep neural networks via function composition:

- The main component of ϕ is determined by \mathcal{F} and can be shared as fixed parameters for all functions in \mathcal{F} . These shared parameters can be determined a priori or learned from any function in \mathcal{F} .
- ϕ is constructed via the composition of a few functions, only two of which is determined by f with a small number of f -dependent parameters.

In particular, we have the following construction

$$\phi = \phi_{f,\mathcal{R}} \circ \phi_{\text{out}} \circ \phi_f \circ \phi_{\text{in}}, \quad (1)$$

where ϕ_{in} and ϕ_{out} are f -independent functions designed based on prior knowledge. We call ϕ_{in} and ϕ_{out} **inner-function** and **outer-function**, respectively. ϕ_f and $\phi_{f,\mathcal{R}}$ are f -dependent functions. ϕ_f is the core part of the whole architecture and $\phi_{f,\mathcal{R}}$ is a simple function for the purpose

^{*}Correspondence to: Shijun Zhang (zhangshijun@u.nus.edu).

¹Department of Mathematics, National University of Singapore, Singapore.

²Department of Mathematics, University of Maryland, College Park, United States.

of adjusting the output range. If we use neural networks to implement the architecture in (1), most parameters are f -independent and stored in ϕ_{in} and ϕ_{out} . We will show that good theoretical and numerical approximations can be achieved by adjusting only a small number of parameters in ϕ_f and $\phi_{f,\mathcal{R}}$.

We will focus on the rectified linear unit (ReLU) activation function and use it to demonstrate our ideas. It would be interesting for future work to extend our work to other activation functions. First, we use the architecture in (1) to theoretically design ReLU networks to approximate (Hölder) continuous functions within an exponentially small approximation error in terms of the number of f -dependent parameters. As we shall see later, in an extreme case, adjusting three intrinsic parameters is enough to achieve an arbitrarily small approximation error. Next, we design a ReLU convolutional neural network (CNN), similar to the architecture in (1), to conduct several experiments to numerically verify that training a small number of intrinsic parameters are enough to achieve good results.

In fact, the architecture (1) has a more general form as follows:

$$\phi = \phi_{f,\mathcal{R}} \circ \phi_k \circ \phi_{f,k-1} \circ \phi_{k-1} \circ \cdots \circ \phi_{f,1} \circ \phi_1, \quad (2)$$

where ϕ_1, \dots, ϕ_k are f -independent functions, which are designed based on prior knowledge. $\phi_{f,1}, \dots, \phi_{f,k-1}$ are the core parts of the whole architecture. $\phi_{f,\mathcal{R}}$ is a simple function adjusting the output range. $\phi_{f,1}, \dots, \phi_{f,k-1}$ and $\phi_{f,\mathcal{R}}$ are determined by f . This paper only focuses on the form in (1). The study of the general form in (2) is left as future work.

Let us further discuss why we emphasize the parameters depending on the target function. It was shown in (Yarotsky, 2018; Shen et al., 2020; Zhang, 2020; Shen et al., 2022) that the approximation error $\mathcal{O}(n^{-2/d})$ is (nearly) optimal for ReLU networks with $\mathcal{O}(n)$ parameters to approximate Lipschitz continuous functions on $[0, 1]^d$. To gain better approximation errors, existing results either consider smaller target function spaces (Yarotsky & Zhevn-erchuk, 2020; Yarotsky, 2017; Lu et al., 2021; Barron, 1993; E et al., 2019b; Chen & Wu, 2019; Montanelli et al., 2021) or introduce new activation functions (Shen et al., 2021a;c;b; Yarotsky, 2021). Observe that, in many existing results, most parameters of networks constructed to approximate the target function f are independent of f . We propose a new perspective to study the approximation error in terms of the number of parameters depending on f , which are called **intrinsic parameters**, excluding those independent of f . We prove by construction that the approximation error can be greatly improved from our new perspective.

Our main contributions can be summarized as follows.

- First, we propose a compositional architecture in (1) and use such an architecture to design networks to approximate target functions. In particular, we construct a ReLU network with $n+2$ intrinsic parameters to approximate a Hölder continuous function f on $[0, 1]^d$ with an error $5\lambda d^{\alpha/2} 2^{-\alpha n}$ measured in the L^p -norm for $p \in [1, \infty)$, where $\alpha \in (0, 1]$ and $\lambda > 0$ are the Hölder order and constant, respectively. Such a result is generalized to generic continuous functions. See Theorem 2.1 for more details.
- We generalize the result in the L^p -norm for $p \in [1, \infty)$ to a new one measured in the L^∞ -norm. Such a generalization is at a price of more intrinsic parameters. Refer to Theorem 2.2 for more details.
- We further extend our results and show that the number of intrinsic parameters can be reduced to three. To be precise, ReLU networks with three intrinsic parameters can achieve an arbitrarily small error for approximating Hölder continuous function on $[0, 1]^d$. In this scenario, extremely high precision is required as we shall see later.
- Finally, we conduct several experiments to numerically verify that training a small part of parameters can achieve good results for classification problems if other parameters are pre-trained from a part of samples.

The rest of this paper is organized as follows. We first present our main theorems and discuss related work in Section 2. These theorems are proved in the appendix. Next, we conduct several experiments to numerically verify our theory in Section 3. Finally, Section 4 concludes this paper with a short discussion.

2. Main results and further interpretation

In this section, we first present our main theorems and then discuss related work. The proofs of these theorems can be found in the appendix.

2.1. Main results

Denote $C([0, 1]^d)$ as the space of continuous functions defined on $[0, 1]^d$. Let $\mathcal{H}_W(d_1, d_2)$ denote the function space consisting of all functions realized by ReLU networks with W parameters mapping from \mathbb{R}^{d_1} to \mathbb{R}^{d_2} , i.e.,

$$\mathcal{H}_W(d_1, d_2) := \left\{ g : g : \mathbb{R}^{d_1} \rightarrow \mathbb{R}^{d_2} \text{ is realized by a } \right. \\ \left. \text{ReLU network with } W \text{ parameters} \right\}.$$

Let $\mathcal{H}(d_1, d_2) := \bigcup_{W=1}^{\infty} \mathcal{H}_W(d_1, d_2)$.

For any $f \in C([0, 1]^d)$, our goal is to construct two f -independent functions $\phi_1 \in \mathcal{H}(d, 1)$ and $\phi_2 \in \mathcal{H}(n, 1)$, and use $s \cdot (\phi_2 \circ \phi_f \circ \phi_1) + b$ to approximate f , where $s \in [0, \infty)$, $b \in \mathbb{R}$, and $\phi_f \in \mathcal{H}_n(1, n)$ are learned from f . Under these settings, an approximation error $\omega_f(\sqrt{d}2^{-n}) + 2^{-n+2}\omega_f(\sqrt{d})$ is attained as shown in the theorem below, where the modulus of continuity of a continuous function $f \in C([0, 1]^d)$ is defined as

$$\omega_f(r) := \sup \{ |f(\mathbf{x}) - f(\mathbf{y})| : \|\mathbf{x} - \mathbf{y}\|_2 \leq r, \mathbf{x}, \mathbf{y} \in [0, 1]^d \}$$

for any $r \geq 0$.

Theorem 2.1. *Given any $n \in \mathbb{N}^+$ and $p \in [1, \infty)$, there exist $\phi_1 \in \mathcal{H}_{2dn+4}(d, 1)$ and $\phi_2 \in \mathcal{H}_{2dn+5n}(n, 1)$ such that: For any $f \in C([0, 1]^d)$, there exists a linear map $\mathcal{L} : \mathbb{R} \rightarrow \mathbb{R}^n$ satisfying*

$$\begin{aligned} & \|s \cdot (\phi_2 \circ \mathcal{L} \circ \phi_1) + b - f\|_{L^p([0, 1]^d)} \\ & \leq \omega_f(\sqrt{d}2^{-n}) + 2^{-n+2}\omega_f(\sqrt{d}), \end{aligned}$$

where $s = 2\omega_f(\sqrt{d})$, $b = f(\mathbf{0}) - \omega_f(\sqrt{d})$, and \mathcal{L} is a linear map given by $\mathcal{L}(t) = (a_1t, a_2t, \dots, a_nt)$ with $a_1, a_2, \dots, a_n \in [0, \frac{1}{3})$ determined by f and n .

In Theorem 2.1, s is a scale factor, b is the bias for a vertical shift, and $a_1, a_2, \dots, a_n \in [0, \frac{1}{3})$ are the key intrinsic parameters storing most of information of f . Clearly, $s \cdot (\phi_2 \circ \mathcal{L} \circ \phi_1) + b$ can be implemented by a ReLU network with $n + 2$ intrinsic parameters. ϕ_1 and ϕ_2 are independent of the target function f and can be implemented by ReLU networks. Remark that the architecture $s \cdot (\phi_2 \circ \mathcal{L} \circ \phi_1) + b$ in Theorem 2.1 can be rewritten as $\phi_{f, \mathcal{R}} \circ \phi_{\text{out}} \circ \phi_f \circ \phi_{\text{in}}$, where $\phi_{\text{in}} = \phi_1$, $\phi_f = \mathcal{L}$, $\phi_{\text{out}} = \phi_2$, and $\phi_{f, \mathcal{R}}$ is a linear function given by $\phi_{f, \mathcal{R}}(x) = sx + b$. Clearly, it is a special case of the architecture in (1).

Note that the approximation error in Theorem 2.1 is characterized by the L^p -norm for $p \in [1, \infty)$. In fact, we can extend such a result to a similar one measured in the L^∞ -norm.

Theorem 2.2. *Given any $n \in \mathbb{N}^+$, there exist $\phi_1 \in \mathcal{H}_{3d2^{2n+5}}(d, 3^d)$ and $\phi_2 \in \mathcal{H}_{3d2^{2n+5n}}(3^d n, 1)$ such that: For any $f \in C([0, 1]^d)$, there exists a linear map $\mathcal{L} : \mathbb{R}^{3^d} \rightarrow \mathbb{R}^{3^d n}$ satisfying*

$$\begin{aligned} & \|s \cdot (\phi_2 \circ \mathcal{L} \circ \phi_1) + b - f\|_{L^\infty([0, 1]^d)} \\ & \leq \omega_f(\sqrt{d}2^{-n}) + 2^{-n+2}\omega_f(\sqrt{d}), \end{aligned}$$

where $s = 2\omega_f(\sqrt{d})$, $b = f(\mathbf{0}) - \omega_f(\sqrt{d})$, and \mathcal{L} is given by

$$\mathcal{L}(x_1, \dots, x_{3^d}) = (\mathcal{L}_0(x_1), \dots, \mathcal{L}_0(x_{3^d}))$$

for any $\mathbf{x} = (x_1, \dots, x_{3^d}) \in \mathbb{R}^{3^d}$, where $\mathcal{L}_0 : \mathbb{R} \rightarrow \mathbb{R}^n$ is a linear map given by $\mathcal{L}_0(t) = (a_1t, a_2t, \dots, a_nt)$ with $a_1, a_2, \dots, a_n \in [0, \frac{1}{3})$ determined by f and n .

Simplifying the implicit approximation error in Theorem 2.1 (or 2.2) to make it explicitly depending on n is challenging in general, since the modulus of continuity $\omega_f(\cdot)$ may be complicated. However, the error can be simplified if f is a Hölder continuous function on $[0, 1]^d$ of order $\alpha \in (0, 1]$ with a Hölder constant $\lambda > 0$. That is, f satisfies

$$|f(\mathbf{x}) - f(\mathbf{y})| \leq \lambda \|\mathbf{x} - \mathbf{y}\|_2^\alpha \quad \text{for any } \mathbf{x}, \mathbf{y} \in [0, 1]^d,$$

implying $\omega_f(r) \leq \lambda \cdot r^\alpha$ for any $r \geq 0$. This means we can get an exponentially small approximation error $5\lambda d^{\alpha/2} 2^{-\alpha n}$. In particular, in the special case of $\alpha = 1$, i.e., f is a Lipschitz continuous function with a Lipschitz constant $\lambda > 0$, then the approximation error is further simplified to $5\lambda\sqrt{d}2^{-n}$.

Though the linear mapping \mathcal{L} in Theorem 2.2 is essentially determined by n key parameters a_1, a_2, \dots, a_n , these n key parameters are repeated 3^d times in the final network architecture as shown in Figure 8. Therefore, $s \cdot (\phi_2 \circ \mathcal{L} \circ \phi_1) + b$ can be implemented by a ReLU network with $3^d n + 2$ intrinsic parameters. Remark that we can reduce the number of intrinsic parameters to $n + 2$ via using a fixed ReLU network to copy n key parameters 3^d times.

Furthermore, the number of intrinsic parameters can be reduced to three in the case of Hölder continuous functions. In other words, three intrinsic parameters are enough to achieve an arbitrary pre-specified error if sufficiently high precision is provided, as shown in the theorem below.

Theorem 2.3. *Given any $\varepsilon > 0$, $\alpha \in (0, 1]$, and $\lambda > 0$, there exists $\phi \in \mathcal{H}(d + 1, 1)$ such that: For any Hölder continuous function f on $[0, 1]^d$ of order $\alpha \in (0, 1]$ with a Hölder constant $\lambda > 0$, there exist three parameters $s \in [0, \infty)$, $v \in [0, 1)$, and $b \in \mathbb{R}$ satisfying*

$$|s\phi(\mathbf{x}, v) + b - f(\mathbf{x})| \leq \varepsilon \quad \text{for any } \mathbf{x} \in [0, 1]^d.$$

In Theorem 2.3, s is a scale factor, b is the bias for a vertical shift, and v is the key intrinsic parameter storing sufficient information of the target function f . Clearly, ϕ is independent of f , while s , v , and b are determined by f . Let $\phi_2 = \phi$, ϕ_1 be the identity map on \mathbb{R}^d , and $\mathcal{L}_v : \mathbb{R}^d \rightarrow \mathbb{R}^{d+1}$ be an affine linear transform mapping \mathbf{x} to (\mathbf{x}, v) . Then $s\phi(\mathbf{x}, v) + b$ can also be represented as $s\phi_2 \circ \mathcal{L}_v \circ \phi_1(\mathbf{x}) + b$, which is a special case of the architecture in (1).

Remark that Theorem 2.3 is just a theoretical result since the key intrinsic parameter v requires extremely high precision, which is necessary for storing the values of f at sufficiently many points within a sufficiently small error. By using the idea of the binary representation, we can extract the values of f stored in v via an f -independent ReLU network (as a sub-network of the final network realizing ϕ

in Theorem 2.3). In fact, there is a balance between the precision requirement and the number of intrinsic parameters. For example, if we store the values of f in two intrinsic parameters (not one), then the precision requirement is greatly lessened.

2.2. Further interpretation

We will connect our theoretical results to related existing results for a deeper understanding. First, we connect our results to transfer learning. Next, we discuss the error analysis of deep neural networks to reveal the motivation for reducing the number of parameters that need to be trained. Finally, we discuss related work from an approximation perspective.

2.2.1. CONNECTION TO TRANSFER LEARNING

Transfer learning dates back to 1970s (Bozinovski & Fulgosi, 1976; Bozinovski, 2020). It is a research direction in machine learning that applies knowledge gained in one problem to solve a different but related problem. Typically in deep learning, transfer learning uses a pre-trained neural network obtained for one task as an initial guess of the neural network for another task to achieve a short training time. Our theory in this paper could provide insights into the success of transfer learning using neural networks, though the setting of our theory is different from realistic transfer learning. In our theory, the inner-function and outer-function are universally used for all learning tasks for continuous functions, which can be understood as the part of networks that can be transferred to different tasks. Suppose f_1 and f_2 are the target functions for two different but related tasks. If f_1 has been learned via an architecture $\phi_{f_1, \mathcal{R}} \circ \phi_{\text{out}} \circ \phi_{f_1} \circ \phi_{\text{in}}$, then we can “transfer” the prior knowledge (ϕ_{in} and ϕ_{out}) to another task. This means that, by only learning ϕ_{f_2} and $\phi_{f_2, \mathcal{R}}$ from f_2 , we can use $\phi_{f_2, \mathcal{R}} \circ \phi_{\text{out}} \circ \phi_{f_2} \circ \phi_{\text{in}}$ to approximate f_2 well. Therefore, the total number of parameters that need to be learned again is not large if most of the parameters are distributed in the sub-networks corresponding to ϕ_{in} and ϕ_{out} . Our theory may provide a certain theoretical understanding in the spirit of transfer learning from a network approximation perspective. To gain a deeper understanding, one can refer to (Karimpanal & Bouffanais, 2019; Pratt, 1993; Baxter, 1998; Mihalkova et al., 2007; Niculescu-Mizil & Caruana, 2007; Rusu et al., 2017; Mikolov et al., 2018). We will test the proposed architecture in (1) in the context of transfer learning in Section 3.

2.2.2. ERROR ANALYSIS

Let us discuss the motivation for reducing the number of parameters that need to be trained. To this end, let us first talk about the error analysis of deep neural networks. Sup-

pose that a target function space \mathcal{F} is given. To numerically compute (or approximate) the element in \mathcal{F} , we need to design a simple and computable hypothesis space \mathcal{H} and use the elements in \mathcal{H} to approximate those in \mathcal{F} . To evaluate how well a numerical solution in \mathcal{H} approximates an element in \mathcal{F} , we introduce three typical errors: the approximation, generalization, and optimization errors.

Given a target function $f \in \mathcal{F}$ defined on a domain \mathcal{X} , our goal is to learn a hypothesis function $\phi \in \mathcal{H}$ from finitely many samples $\{(\mathbf{x}_i, f(\mathbf{x}_i))\}_{i=1}^n$. To infer $f(\mathbf{x})$ for an unseen sample \mathbf{x} , we need to identify the empirical risk minimizer $\phi_{\mathcal{S}}$, which is given by

$$\phi_{\mathcal{S}} \in \arg \min_{\phi \in \mathcal{H}} R_{\mathcal{S}}(\phi), \quad (3)$$

where $R_{\mathcal{S}}(\phi)$ is the empirical risk defined by

$$R_{\mathcal{S}}(\phi) := \frac{1}{n} \sum_{i=1}^n \ell(\phi(\mathbf{x}_i), f(\mathbf{x}_i))$$

with a loss function typically taken as $\ell(y, y') = \frac{1}{2}|y - y'|^2$.

In fact, the best hypothesis function to infer $f(\mathbf{x})$ is $\phi_{\mathcal{D}}(\mathbf{x})$, but not $\phi_{\mathcal{S}}(\mathbf{x})$, where $\phi_{\mathcal{D}}$ is the expected risk minimizer given by

$$\phi_{\mathcal{D}} \in \arg \min_{\phi \in \mathcal{H}} R_{\mathcal{D}}(\phi),$$

where $R_{\mathcal{D}}(\phi)$ is the expected risk defined by

$$R_{\mathcal{D}}(\phi) := \mathbb{E}_{\mathbf{x} \sim U(\mathcal{X})} [\ell(\phi(\mathbf{x}), f(\mathbf{x}))],$$

where U is a unknown data distribution over \mathcal{X} . The best possible inference error is $R_{\mathcal{D}}(\phi_{\mathcal{D}})$. In real-world applications, $U(\mathcal{X})$ is unknown and only finitely many samples from this distribution are available. Hence, the empirical risk $R_{\mathcal{S}}(\phi)$ is minimized, hoping to obtain $\phi_{\mathcal{S}}(\mathbf{x})$, instead of minimizing the expected risk $R_{\mathcal{D}}(\phi)$ to obtain $\phi_{\mathcal{D}}(\mathbf{x})$. A numerical optimization method to solve (3) may result in a numerical solution (denoted as $\phi_{\mathcal{N}}$) that may not be a global minimizer $\phi_{\mathcal{S}}$. Therefore, the actually learned hypothesis function to infer $f(\mathbf{x})$ is $\phi_{\mathcal{N}}(\mathbf{x})$ and the corresponding inference error is measured by $R_{\mathcal{D}}(\phi_{\mathcal{N}})$, which is bounded by

$$\begin{aligned} R_{\mathcal{D}}(\phi_{\mathcal{N}}) &= \underbrace{[R_{\mathcal{D}}(\phi_{\mathcal{N}}) - R_{\mathcal{S}}(\phi_{\mathcal{N}})]}_{\text{GE}} + \underbrace{[R_{\mathcal{S}}(\phi_{\mathcal{N}}) - R_{\mathcal{S}}(\phi_{\mathcal{S}})]}_{\text{OE}} \\ &\quad + \underbrace{[R_{\mathcal{S}}(\phi_{\mathcal{S}}) - R_{\mathcal{S}}(\phi_{\mathcal{D}})]}_{\leq 0 \text{ by (3)}} + \underbrace{[R_{\mathcal{S}}(\phi_{\mathcal{D}}) - R_{\mathcal{D}}(\phi_{\mathcal{D}})]}_{\text{GE}} + \underbrace{R_{\mathcal{D}}(\phi_{\mathcal{D}})}_{\text{AE}} \\ &\leq \underbrace{R_{\mathcal{D}}(\phi_{\mathcal{D}})}_{\text{Approximation error (AE)}} + \underbrace{[R_{\mathcal{S}}(\phi_{\mathcal{N}}) - R_{\mathcal{S}}(\phi_{\mathcal{S}})]}_{\text{Optimization error (OE)}} \\ &\quad + \underbrace{[R_{\mathcal{D}}(\phi_{\mathcal{N}}) - R_{\mathcal{S}}(\phi_{\mathcal{N}})] + [R_{\mathcal{S}}(\phi_{\mathcal{D}}) - R_{\mathcal{D}}(\phi_{\mathcal{D}})]}_{\text{Generalization error (GE)}}. \end{aligned}$$

As we can see from the above equation, the numerical inference error $R_{\mathcal{D}}(\phi_{\mathcal{N}})$, the distance between the numerical solution $\phi_{\mathcal{N}}$ and the target function f , is bounded by three errors: the approximation, generalization, and optimization errors. See Figure 1 for an illustration.

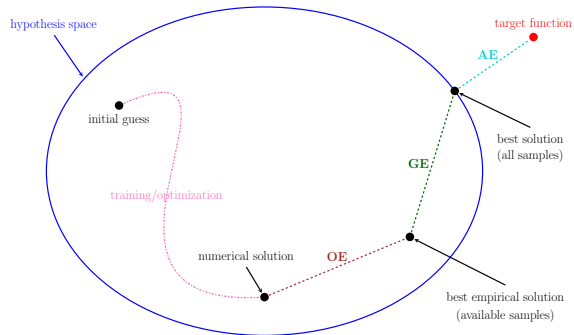


Figure 1. An illustration of the approximation error (AE), the generalization error (GE), and the optimization error (OE).

The constructive approximation established in this paper and the literature provides an upper bound of the approximation error $R_{\mathcal{D}}(\phi_{\mathcal{D}})$. The theoretical guarantee of the convergence of an optimization algorithm to a global minimizer $\phi_{\mathcal{S}}$ and the characterization of the convergence belong to the optimization analysis of neural networks. The generalization error is controlled by two key factors: the complexity of the hypothesis function space and the number of training (available) samples. One could refer to (Beck et al., 2019; E et al., 2019b;a; Kawaguchi, 2016; Nguyen & Hein, 2017; Kawaguchi & Bengio, 2019; He et al., 2020; Li et al., 2019) for further discussions of the generalization and optimization errors.

Theorems 2.1, 2.2, and 2.3 provide upper bounds of $R_{\mathcal{D}}(\phi_{\mathcal{D}})$. These bounds only depends on the number of intrinsic parameters of ReLU networks and the modulus of continuity $\omega_f(\cdot)$. Hence, these bounds are independent of the empirical risk minimization problem in (3) and the optimization algorithm used to compute the numerical solution of (3). In other words, Theorems 2.1, 2.2, and 2.3 quantify the approximation power of ReLU networks in terms of the number of intrinsic parameters. Designing efficient optimization algorithms and analyzing the generalization error for ReLU networks are two other separate future directions.

Generally, making the hypothesis function space \mathcal{H} smaller would result in a larger approximation error, a smaller generalization error, and a smaller optimization error. Thus, there is a balance for the choice of the hypothesis function space. Our theory pre-specifies most parameters based on the prior knowledge of the target function space \mathcal{F} , which leads to a smaller hypothesis function space, denoted by $\tilde{\mathcal{H}}$. Thus, we can expect better approximation and generalization errors. Meanwhile, the approximation error becomes only a little worse since we only pre-specify unimportant parameters (non-intrinsic ones). Therefore, we can expect a good inference (test) error by using our

method to reduce the number of parameters that need to be trained. See Figure 2 for an illustration.




Figure 2. An illustration of our theory. The approximation error becomes only a little bit worse since we just pre-specify unimportant parameters based on the target function space \mathcal{F} . The new hypothesis function space $\tilde{\mathcal{H}}$ is much smaller than the original hypothesis function space \mathcal{H} , leading to better generalization and optimization errors.

The key point of our method is the prior knowledge of the target function space \mathcal{F} . This generally means that \mathcal{F} is small and special. Our method would fail if the target function space \mathcal{F} is pretty large (e.g., $\mathcal{F} = L^1([0, 1]^d)$). In this case, if the new hypothesis function space $\tilde{\mathcal{H}}$ is much smaller than the original one \mathcal{H} , then the approximation error would become much worse. See Figure 3 for an illustration.




Figure 3. An illustration of the case that the target function space \mathcal{F} is large.

2.2.3. RELATED WORK

The expressiveness of deep neural networks has been studied extensively from many perspectives, e.g., in terms of combinatorics (Montufar et al., 2014), topology (Bianchini & Scarselli, 2014), Vapnik-Chervonenkis (VC) dimension (Bartlett et al., 1998; Sakurai, 1999; Harvey et al., 2017),

fat-shattering dimension (Kearns & Schapire, 1994; Anthony & Bartlett, 2009), information theory (Petersen & Voigtlaender, 2018), classical approximation theory (Cybenko, 1989; Hornik et al., 1989; Barron, 1993; Yarotsky, 2018; 2017; Bölcskei et al., 2019; Zhou, 2020; Chui et al., 2018; Gribonval et al., 2019; Gühring et al., 2019; Suzuki, 2019; Nakada & Imaizumi, 2020; Chen et al., 2019; Bao et al., 2019; Li et al., 2019; Montanelli & Yang, 2020; Shen et al., 2019; 2020; Lu et al., 2021; Zhang, 2020), etc. In the early works of approximation theory for neural networks, the universal approximation theorem (Cybenko, 1989; Hornik, 1991; Hornik et al., 1989) without approximation errors showed that, given any $\varepsilon > 0$, there exists a sufficiently large neural network approximating a target function in a certain function space within an error ε . For one-hidden-layer neural networks and sufficiently smooth functions, it is shown in (Barron, 1993; Barron & Klusowski, 2018) that an asymptotic approximation error $\mathcal{O}(N^{-1/2})$ in the L^2 -norm, leveraging an idea that is similar to Monte Carlo sampling for high-dimensional integrals.

Recently, it is proved in (Shen et al., 2020; Yarotsky, 2018; Zhang, 2020) that the (nearly) optimal approximation error would be $\mathcal{O}(n^{-2/d})$ when using ReLU networks with n parameters to approximate functions in the unit ball of Lipschitz continuous function space on $[0, 1]^d$. Clearly, such an error suffers from the curse of dimensionality. To bridge this gap, one could either consider smaller function spaces, e.g., smooth functions (Lu et al., 2021; Yarotsky & Zhevnerchuk, 2020) and band-limited functions (Montanelli et al., 2021), or introducing new network architectures, e.g., Floor-ReLU networks (Shen et al., 2021a), Floor-Exponential-Step (FLES) networks (Shen et al., 2021b), and (Sin, ReLU, 2^x)-activated networks (Jiao et al., 2021). This paper proposes a new perspective to characterize the approximation error in terms of the number of intrinsic parameters. Such a method is inspired by an observation that most parameters of most constructed networks in the mentioned papers are independent of the target function. Thus, most parameters can be assigned or computed in advance, i.e., we can approximate the target function by only adjusting a small number of parameters. As shown in Theorem 2.1, we can first design an inner-function ϕ_1 and an outer-function ϕ_2 , both of which can be implemented by ReLU networks. Then, for any continuous function $f \in C([0, 1]^d)$, $s \cdot (\phi_2 \circ \mathcal{L} \circ \phi_1) + b$ can approximate f with an error $\omega_f(\sqrt{d}2^{-n}) + 2^{-n+2}\omega_f(\sqrt{d})$ by the following two steps: 1) determining s and b , 2) designing a linear map \mathcal{L} defined by $\mathcal{L}(t) = (a_1t, \dots, a_nt)$, where a_1, \dots, a_n are determined by the target function f . Therefore, we overcome the curse of dimensionality in the sense of the approximation error characterized by the number of intrinsic parameters when the variation of $\omega_f(r)$ as $r \rightarrow 0$

is moderate (e.g., $\omega_f(r) \lesssim r^\alpha$ for Hölder continuous functions).

3. Experiments

In this section, we conduct several experiments to provide numerical examples that training a small number of parameters is enough to achieve good results. We first discuss the goal of our experiments. Next, we extend the architecture in (1) to a simple convolutional neural network (CNN) architecture for classification problems. Finally, we use the proposed CNN architecture to conduct several experiments and present the numerical results for three comment datasets: MNIST, Kuzushiji-MNIST (KMNIST), and Fashion-MNIST (FMNIST).

3.1. Goal of experiments

First, let us discuss why we adopt classification problems as our experiment examples. The goal of a classification problem with $J \in \mathbb{N}^+$ classes is to identify a classification function f defined by

$$f(\mathbf{x}) = j \quad \text{for any } \mathbf{x} \in E_j \text{ and } j = 0, 1, \dots, J-1,$$

where $E_j \subseteq E$ is the minimum closed set containing all samples with a label j and $E \subseteq \mathbb{R}^d$ is a bounded closed set (e.g., $E = [0, 1]^d$). Clearly, E_0, E_1, \dots, E_{J-1} are pairwise disjoint. Such a classification function can be continuously extended to E , which means a classification problem can also be regarded as a continuous function approximation problem. We take the case $J = 2$ as an example to illustrate the extension. The multiclass case is similar. Define

$$\tilde{f}(\mathbf{x}) := \frac{\text{dist}(\mathbf{x}, E_0)}{\text{dist}(\mathbf{x}, E_0) + \text{dist}(\mathbf{x}, E_1)} \quad \text{for any } \mathbf{x} \in [0, 1]^2,$$

where

$$\text{dist}(\mathbf{x}, E_i) := \inf_{\mathbf{y} \in E_i} \|\mathbf{x} - \mathbf{y}\|_2$$

for any $\mathbf{x} \in [0, 1]^2$ and $i = 0, 1$. It is easy to verify that \tilde{f} is continuous on $[0, 1]^d$ and

$$\tilde{f}(\mathbf{x}) = f(\mathbf{x}) \quad \text{for any } \mathbf{x} \in E_0 \cup E_1.$$

That means \tilde{f} is a continuous extension of f . Remark that, in our experiments, we use networks to approximate an equivalent variant \hat{f} of the original classification function f mentioned above, where \hat{f} is given by

$$\hat{f}(\mathbf{x}) = \mathbf{e}_j \quad \text{for any } \mathbf{x} \in E_j \text{ and } j = 0, 1, \dots, J-1,$$

where $\{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_J\}$ is the standard basis of \mathbb{R}^J , i.e., $\mathbf{e}_j \in \mathbb{R}^J$ denotes the vector with a 1 in the j -th coordinate and 0's elsewhere.

Table 1. Network architecture.

layers	activation	size	dropout	batch normalization	#parameters	remark
input $\in \mathbb{R}^{28 \times 28}$		28 \times 28				
Conv-1: $1 \times (3 \times 3)$, 32	ReLU	$(26 \times 26) \times 32$		yes	320	low-level features, block 1 (ϕ_{in})
Conv-2: $32 \times (3 \times 3)$, 32	ReLU, MaxPool	$(12 \times 12) \times 32$	0.25	yes	9248	
Conv-3: $32 \times (3 \times 3)$, 64	ReLU	$(10 \times 10) \times 64$		yes	18496	high-level features, block 2 (ϕ_f)
Conv-4: $64 \times (3 \times 3)$, 64	ReLU, Flatten	$(8 \times 8) \times 64$	0.25	yes	36992	
FC-1: 4096, 512	ReLU	512	0.5	yes	2097664	initial classifier, block 3 (ϕ_{out})
FC-2: 512, 64	ReLU	64		yes	32832	final classifier, block 4 ($\phi_{f,\mathcal{R}}$)
FC-3: 64, 10	Softmax	10		yes	650	
output $\in \mathbb{R}^{10}$						

In our theoretical results, we need to specify the network architecture and the values of the parameters corresponding to ϕ_{in} and ϕ_{out} . It is conjectured that there are more choices for the network architecture and the values of the parameters corresponding to ϕ_{in} and ϕ_{out} . To verify that numerically, we extend the architecture in (1) to a CNN architecture and pre-specify the values of the parameters corresponding to ϕ_{in} and ϕ_{out} via pre-training all parameters with a part of training samples.

Our theory only focuses on the approximation error, while the test error (accuracy) in the experiment is bounded by three errors, as discussed in Section 2.2. A good approximation error may not guarantee a good test error. However, the approximation error is bounded by the test error, as we can see from the discussion of the error analysis previously. Thus, a good test error implies that the approximation error is well controlled. It remains to show that training a small number of parameters are enough to achieve a good test error. If this can be numerically observed, then the proposed CNN can approximate the classification function well via only adjusting only a small number of parameters.

3.2. Network architecture

We will design a CNN architecture to classify the images in three datasets: MNIST, KMNIST, and FMNIST. Each of these three datasets has ten classes and their elements have a size 28×28 . Thus, we can take the same CNN architecture for these three datasets. For simplicity, we consider a basic CNN architecture: four convolutional layers and three fully connected layers. The whole CNN architecture is summarized in Figure 4. We present more details of the proposed CNN architecture and connect it to the architecture in (1) in Table 1.

To illustrate the connection between the proposed CNN architecture and the architecture in (1), we divide the proposed CNN architecture into four main blocks. Block 1 has two convolutional layers, extracting the low-level features; Block 2 has two convolutional layers, extracting the high-level features; Block 3 has one fully connected layer,




Figure 4. An illustration of the CNN architecture. Conv and FC are short of convolutional and fully connected layers, respectively.

regarded as the initial classifier; Block 4 has two fully connected layers, regarded as the final classifier. See a summary in the following equation.

$$\begin{aligned}
 \text{input} &\rightarrow \underbrace{\text{Conv-1} \rightarrow \text{Conv-2}}_{\text{block 1 (low-level features)}} \rightarrow \underbrace{\text{Conv-3} \rightarrow \text{Conv-4}}_{\text{block 2 (high-level features)}} \\
 &\rightarrow \underbrace{\text{FC-1}}_{\text{block 3 (initial classifier)}} \rightarrow \underbrace{\text{FC-2} \rightarrow \text{FC-3}}_{\text{block 4 (final classifier)}} \rightarrow \text{output}
 \end{aligned}$$

Remark that the above CNN architecture can be considered as a special case of the architecture in (1). Blocks 1, 2, 3, and 4 in the proposed CNN architecture are approximately equivalent to ϕ_{in} , ϕ_f , ϕ_{out} , and $\phi_{f,\mathcal{R}}$ in the architecture in (1), respectively.

3.3. Numerical results

Our goal is to numerically verify our theoretical result that adjusting a small number of parameters is enough to achieve a good approximation. It is natural to ask how to pre-specify the values of the non-training parameters. Since it is difficult to manually specify the values of the parameters in a CNN architecture, we use a part of training samples to pre-training all parameters. Then, we propose




Figure 5. Test accuracy over epochs.

two optimization strategies to train the proposed CNN as follows.

- (S1) The normal strategy: We use all training samples to train all parameters.
- (S2) A strategy based on the architecture in (1): We first use the training samples with labels $0, 1, \dots, 4$ to pre-training all parameters, and then use all training samples to continue training the parameters in Blocks 2 and 4.

Before presenting the numerical results, let us present the hyperparameters for training the proposed CNN architecture. To reduce overfitting and speed up optimization, we take two common regularization methods: dropout (Hinton et al., 2012; Srivastava et al., 2014) and batch normalization (Ioffe & Szegedy, 2015). We use the cross-entropy loss function to evaluate the loss. The number of epochs and the batch size are set to 300 and 128, respectively. We adopt RAdam (Liu et al., 2020) as the optimization method. The weight decay of the optimizer is 0.0001 and the learning rate is $0.002 \times 0.93^{i-1}$ in the i -th epoch. Remark that all training (test) samples are standardized before training, i.e., we rescale the samples to have a mean of 0 and a standard deviation of 1.

For the three mentioned datasets, we use two proposed optimization strategies to train the proposed CNN and use all test samples to obtain the test accuracy. As we can see from Figure 5, the test accuracy becomes steady after 50 epochs. Thus, it is reasonable to take the largest test accuracy over epochs as the target test accuracy. The test accuracy comparison of two optimization strategies is summarized in Table 2.

As we can see from Table 2, for all three datasets, the second optimization strategy (S2) trains much less parameters at the price of a slightly lower test accuracy, compared to the first optimization strategy (S1). As discussed in Section 2.2, the test accuracy (error) is controlled by three errors: the approximation, generalization, and optimization

Table 2. Test accuracy comparison.

strategy	MNIST	KMNIST	FMNIST	#training-parameters
(S1)	0.9964	0.9832	0.9458	2.2×10^6
(S2)	0.9962	0.9744	0.9360	8.9×10^4

errors. A good approximation error cannot guarantee a good test accuracy. However, a good test accuracy numerically implies that the three errors are all well controlled. Our numerical results suggest that training a small number of parameters is enough to achieve a good test accuracy. Therefore, the proposed CNN can approximate the classification function well via only adjusting only a small number of parameters.

4. Conclusion

This paper aims to achieve a good approximation via adjusting only a small number of parameters based on the target function f while using a ReLU network to approximate f . We first propose a composition architecture in (1), and then use such an architecture to construct ReLU networks to approximate the target function f . In Theorem 2.1, we prove that a ReLU network with $n + 2$ intrinsic parameters can approximate a continuous function f on $[0, 1]^d$ with an error $\omega_f(\sqrt{d}2^{-n}) + 2^{-n+2}\omega_f(\sqrt{d})$ measured in the L^p -norm for $p \in [1, \infty)$. Moreover, such a result can be generalized from the L^p -norm to the L^∞ -norm at a price of adding $\mathcal{O}(n)$ intrinsic parameters, as shown in Theorem 2.2. Next, we show in Theorem 2.3 that three intrinsic parameters are enough to achieve an arbitrarily small error in the case of Hölder continuous functions, though this result requires high precision to encode these three parameters on computers. Finally, we conduct several experiments to provide numerical examples of the architecture in (1). Remark that this paper only focuses on the approximation error characterized by the number of intrinsic parameters, the study of the optimization error and generalization error

will be left as future work.

Acknowledgments

Z. Shen is supported by Distinguished Professorship of National University of Singapore. H. Yang was partially supported by the US National Science Foundation under award DMS-1945029.

References

- Anthony, M. and Bartlett, P. L. *Neural Network Learning: Theoretical Foundations*. Cambridge University Press, New York, NY, USA, 1st edition, 2009. ISBN 052111862X, 9780521118620.
- Bao, C., Li, Q., Shen, Z., Tai, C., Wu, L., and Xiang, X. Approximation analysis of convolutional neural networks. *Semantic Scholar e-Preprint*, art. Corpus ID: 204762668, 2019.
- Barron, A. R. Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information Theory*, 39(3):930–945, May 1993. ISSN 0018-9448. doi: 10.1109/18.256500.
- Barron, A. R. and Klusowski, J. M. Approximation and estimation for high-dimensional deep learning networks. *arXiv e-prints*, art. arXiv:1809.03090, September 2018.
- Bartlett, P., Maierov, V., and Meir, R. Almost linear VC-dimension bounds for piecewise polynomial networks. *Neural Computation*, 10(8):2159–2173, 1998.
- Baxter, J. Theoretical models of learning to learn. In *Learning to Learn*. Springer, 1998. doi: 10.1007/978-1-4615-5529-2_4.
- Beck, C., Jentzen, A., and Kuckuck, B. Full error analysis for the training of deep neural networks. *CoRR*, abs/1910.00121, 2019. URL <http://arxiv.org/abs/1910.00121>.
- Bianchini, M. and Scarselli, F. On the complexity of neural network classifiers: A comparison between shallow and deep architectures. *IEEE Transactions on Neural Networks and Learning Systems*, 25(8):1553–1565, Aug 2014. ISSN 2162-237X. doi: 10.1109/TNNLS.2013.2293637.
- Bölcskei, H., Grohs, P., Kutyniok, G., and Petersen, P. Optimal approximation with sparsely connected deep neural networks. *SIAM Journal on Mathematics of Data Science*, 1(1):8–45, 2019. doi: 10.1137/18M118709X.
- Bozinovski, S. Reminder of the first paper on transfer learning in neural networks, 1976. *Informatica*, 44(3): 291–302, 2020. doi: 10.31449/inf.v44i3.2828.
- Bozinovski, S. and Fulgosi, A. The influence of pattern similarity and transfer learning upon the training of a base perceptron B2. (original in croatian). In *Proceedings of the Symposium Informatica 3-121-5*, Bled, Slovenia, 1976.
- Chen, L. and Wu, C. A note on the expressive power of deep rectified linear unit networks in high-dimensional spaces. *Mathematical Methods in the Applied Sciences*, 42(9):3400–3404, 2019. doi: 10.1002/mma.5575.
- Chen, M., Jiang, H., Liao, W., and Zhao, T. Efficient approximation of deep ReLU networks for functions on low dimensional manifolds. In Wallach, H., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/file/fd95ec8df5dbeea25aa8e6c808bad583-Paper.pdf>.
- Cheng, Y., Wang, D., Zhou, P., and Zhang, T. A survey of model compression and acceleration for deep neural networks. *arXiv e-prints*, art. arXiv:1710.09282, October 2017.
- Chui, C. K., Lin, S.-B., and Zhou, D.-X. Construction of neural networks for realization of localized deep learning. *Frontiers in Applied Mathematics and Statistics*, 4: 14, 2018. ISSN 2297-4687. doi: 10.3389/fams.2018.00014.
- Cybenko, G. Approximation by superpositions of a sigmoidal function. *MCSS*, 2:303–314, 1989.
- Denton, E., Zaremba, W., Bruna, J., LeCun, Y., and Fergus, R. Exploiting linear structure within convolutional networks for efficient evaluation. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 1*, NIPS’14, pp. 1269–1277, Cambridge, MA, USA, 2014. MIT Press.
- E, W., Ma, C., and Wang, Q. A priori estimates of the population risk for residual networks. *ArXiv*, abs/1903.02154, 2019a.
- E, W., Ma, C., and Wu, L. A priori estimates of the population risk for two-layer neural networks. *Communications in Mathematical Sciences*, 17(5):1407–1425, 2019b.
- Gribonval, R., Kutyniok, G., Nielsen, M., and Voigtlaender, F. Approximation spaces of deep neural networks. *arXiv e-prints*, art. arXiv:1905.01208, May 2019.
- Gühning, I., Kutyniok, G., and Petersen, P. Error bounds for approximations with deep ReLU neural networks in $W^{s,p}$ norms. *arXiv e-prints*, art. arXiv:1902.07896, Feb 2019.

- Gupta, S., Agrawal, A., Gopalakrishnan, K., and Narayanan, P. Deep learning with limited numerical precision. In Bach, F. and Blei, D. (eds.), *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pp. 1737–1746, Lille, France, 07–09 Jul 2015. PMLR. URL <https://proceedings.mlr.press/v37/gupta15.html>.
- Han, S., Pool, J., Tran, J., and Dally, W. J. Learning both weights and connections for efficient neural networks. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1*, NIPS’15, pp. 1135–1143, Cambridge, MA, USA, 2015. MIT Press.
- Harvey, N., Liaw, C., and Mehrabian, A. Nearly-tight VC-dimension bounds for piecewise linear neural networks. In Kale, S. and Shamir, O. (eds.), *Proceedings of the 2017 Conference on Learning Theory*, volume 65 of *Proceedings of Machine Learning Research*, pp. 1064–1068, Amsterdam, Netherlands, 07–10 Jul 2017. PMLR. URL <http://proceedings.mlr.press/v65/harvey17a.html>.
- He, J., Jia, X., Xu, J., Zhang, L., and Zhao, L. Make ℓ_1 regularization effective in training sparse CNN. *Computational Optimization and Applications*, 77(1):163–182, 2020. doi: 10.1007/s10589-020-00202-1.
- Hinton, G., Vinyals, O., and Dean, J. Distilling the knowledge in a neural network. *arXiv e-prints*, art. arXiv:1503.02531, March 2015.
- Hinton, G. E., Srivastava, N., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. Improving neural networks by preventing co-adaptation of feature detectors. *CoRR*, abs/1207.0580, 2012. URL <http://arxiv.org/abs/1207.0580>.
- Hornik, K. Approximation capabilities of multilayer feedforward networks. *Neural Networks*, 4(2):251–257, 1991. ISSN 0893-6080. doi: 10.1016/0893-6080(91)90009-T.
- Hornik, K., Stinchcombe, M., and White, H. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366, 1989. ISSN 0893-6080. doi: 10.1016/0893-6080(89)90020-8.
- Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., and Adam, H. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *CoRR*, abs/1704.04861, 2017. URL <http://arxiv.org/abs/1704.04861>.
- Ioffe, S. and Szegedy, C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37*, ICML’15, pp. 448–456. JMLR.org, 2015.
- Jaderberg, M., Vedaldi, A., and Zisserman, A. Speeding up convolutional neural networks with low rank expansions. In *Proceedings of the British Machine Vision Conference*. BMVA Press, 2014. doi: 10.5244/C.28.88.
- Jiao, Y., Lai, Y., Lu, X., Wang, F., Zhijian Yang, J., and Yang, Y. Deep neural networks with ReLU-Sine-Exponential activations break curse of dimensionality on hölder class. *arXiv e-prints*, art. arXiv:2103.00542, February 2021.
- Karimpanal, T. G. and Bouffanais, R. Self-organizing maps for storage and transfer of knowledge in reinforcement learning. *Adaptive Behavior*, 27(2):111–126, 2019. doi: 10.1177/1059712318818568.
- Kawaguchi, K. Deep learning without poor local minima. In Lee, D. D., Sugiyama, M., Luxburg, U. V., Guyon, I., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 29*, pp. 586–594. Curran Associates, Inc., 2016.
- Kawaguchi, K. and Bengio, Y. Depth with nonlinearity creates no bad local minima in resnets. *Neural Networks*, 118:167–174, 2019. ISSN 0893-6080. doi: 10.1016/j.neunet.2019.06.009.
- Kearns, M. J. and Schapire, R. E. Efficient distribution-free learning of probabilistic concepts. *J. Comput. Syst. Sci.*, 48(3):464–497, June 1994. ISSN 0022-0000. doi: 10.1016/S0022-0000(05)80062-5.
- Li, Q., Lin, T., and Shen, Z. Deep learning via dynamical systems: An approximation perspective. *arXiv e-prints*, art. arXiv:1912.10382, December 2019.
- Li, Q., Tai, C., and E, W. Stochastic modified equations and dynamics of stochastic gradient algorithms I: Mathematical foundations. *Journal of Machine Learning Research*, 20(40):1–47, 2019. URL <http://jmlr.org/papers/v20/17-526.html>.
- Liu, L., Jiang, H., He, P., Chen, W., Liu, X., Gao, J., and Han, J. On the variance of the adaptive learning rate and beyond. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=rkgz2aEKDr>.
- Lu, J., Shen, Z., Yang, H., and Zhang, S. Deep network approximation for smooth functions. *SIAM Journal on Mathematical Analysis*, 53(5):5465–5506, 2021. doi: 10.1137/20M134695X.

- Mihalkova, L., Huynh, T., and Mooney, R. J. Mapping and revising markov logic networks for transfer learning. In *Proceedings of the 22nd Conference on Artificial Intelligence (AAAI-07)*, volume 1, pp. 608–614, Vancouver, Canada, 2007. doi: 10.5555/1619645.1619743.
- Mikolov, T., Joulin, A., and Baroni, M. A roadmap towards machine intelligence. In Gelbukh, A. (ed.), *Computational Linguistics and Intelligent Text Processing*, pp. 29–61, Cham, 2018. Springer International Publishing. ISBN 978-3-319-75477-2.
- Montanelli, H. and Yang, H. Error bounds for deep ReLU networks using the Kolmogorov-Arnold superposition theorem. *Neural Networks*, 129:1–6, 2020. ISSN 0893-6080. doi: 10.1016/j.neunet.2019.12.013.
- Montanelli, H., Yang, H., and Du, Q. Deep ReLU networks overcome the curse of dimensionality for generalized bandlimited functions. *Journal of Computational Mathematics*, 39(6):801–815, 2021. ISSN 1991-7139. doi: 10.4208/jcm.2007-m2019-0239.
- Montufar, G. F., Pascanu, R., Cho, K., and Bengio, Y. On the number of linear regions of deep neural networks. In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N. D., and Weinberger, K. Q. (eds.), *Advances in Neural Information Processing Systems 27*, pp. 2924–2932. Curran Associates, Inc., 2014.
- Nakada, R. and Imaizumi, M. Adaptive approximation and generalization of deep neural network with intrinsic dimensionality. *Journal of Machine Learning Research*, 21(174):1–38, 2020. URL <http://jmlr.org/papers/v21/20-002.html>.
- Nguyen, Q. N. and Hein, M. The loss surface of deep and wide neural networks. *CoRR*, abs/1704.08045, 2017. URL <http://arxiv.org/abs/1704.08045>.
- Niculescu-Mizil, A. and Caruana, R. Inductive transfer for bayesian network structure learning. In Meila, M. and Shen, X. (eds.), *Proceedings of the Eleventh International Conference on Artificial Intelligence and Statistics*, volume 2 of *Proceedings of Machine Learning Research*, pp. 339–346, San Juan, Puerto Rico, 21–24 Mar 2007. PMLR. URL <https://proceedings.mlr.press/v2/niculescu-mizil07a.html>.
- Petersen, P. and Voigtlaender, F. Optimal approximation of piecewise smooth functions using deep ReLU neural networks. *Neural Networks*, 108:296–330, 2018. ISSN 0893-6080. doi: 10.1016/j.neunet.2018.08.019.
- Pratt, L. Y. Discriminability-based transfer between neural networks. In Hanson, S., Cowan,
- J., and Giles, C. (eds.), *Advances in Neural Information Processing Systems*, volume 5. Morgan-Kaufmann, 1993. URL <https://proceedings.neurips.cc/paper/1992/file/67e103b0761e60683e83c559be18d40c-Paper.pdf>.
- Qin, Z., Li, Z., Zhang, Z., Bao, Y., Yu, G., Peng, Y., and Sun, J. Thundernet: Towards real-time generic object detection on mobile devices. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 6717–6726, 2019. doi: 10.1109/ICCV.2019.00682.
- Rusu, A. A., Vecerík, M., Rothörl, T., Heess, N., Pascanu, R., and Hadsell, R. Sim-to-real robot learning from pixels with progressive nets. In *1st Annual Conference on Robot Learning, CoRL 2017, Mountain View, California, USA, November 13-15, 2017, Proceedings*, volume 78 of *Proceedings of Machine Learning Research*, pp. 262–270. PMLR, 2017. URL <http://proceedings.mlr.press/v78/rusu17a.html>.
- Sakurai, A. Tight bounds for the VC-dimension of piecewise polynomial networks. In *Advances in Neural Information Processing Systems*, pp. 323–329. Neural information processing systems foundation, 1999. ISBN 0262112450.
- Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., and Chen, L.-C. Mobilenetv2: Inverted residuals and linear bottlenecks. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4510–4520, 2018. doi: 10.1109/CVPR.2018.00474.
- Shen, Z., Yang, H., and Zhang, S. Nonlinear approximation via compositions. *Neural Networks*, 119:74–84, 2019. ISSN 0893-6080. doi: 10.1016/j.neunet.2019.07.011.
- Shen, Z., Yang, H., and Zhang, S. Deep network approximation characterized by number of neurons. *Communications in Computational Physics*, 28(5):1768–1811, 2020. ISSN 1991-7120. doi: 10.4208/cicp.OA-2020-0149.
- Shen, Z., Yang, H., and Zhang, S. Deep network with approximation error being reciprocal of width to power of square root of depth. *Neural Computation*, 33(4): 1005–1036, 03 2021a. ISSN 0899-7667. doi: 10.1162/neco.a_01364.
- Shen, Z., Yang, H., and Zhang, S. Neural network approximation: Three hidden layers are enough. *Neural Networks*, 141:160–173, 2021b. ISSN 0893-6080. doi: 10.1016/j.neunet.2021.04.011.
- Shen, Z., Yang, H., and Zhang, S. Deep network approximation: Achieving arbitrary accuracy with fixed number

- of neurons. *arXiv e-prints*, art. arXiv:2107.02397, July 2021c.
- Shen, Z., Yang, H., and Zhang, S. Optimal approximation rate of ReLU networks in terms of width and depth. *Journal de Mathématiques Pures et Appliquées*, 157:101–135, 2022. ISSN 0021-7824. doi: 10.1016/j.matpur.2021.07.009.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(56):1929–1958, 2014. URL <http://jmlr.org/papers/v15/srivastava14a.html>.
- Suzuki, T. Adaptivity of deep ReLU network for learning in Besov and mixed smooth Besov spaces: optimal rate and curse of dimensionality. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=H1ebTsActm>.
- Wu, J., Leng, C., Wang, Y., Hu, Q., and Cheng, J. Quantized convolutional neural networks for mobile devices. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4820–4828, Los Alamitos, CA, USA, jun 2016. IEEE Computer Society. doi: 10.1109/CVPR.2016.521.
- Yarotsky, D. Error bounds for approximations with deep ReLU networks. *Neural Networks*, 94:103–114, 2017. ISSN 0893-6080. doi: 10.1016/j.neunet.2017.07.002.
- Yarotsky, D. Optimal approximation of continuous functions by very deep ReLU networks. In Bubeck, S., Perchet, V., and Rigollet, P. (eds.), *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pp. 639–649. PMLR, 06–09 Jul 2018. URL <http://proceedings.mlr.press/v75/yarotsky18a.html>.
- Yarotsky, D. Elementary superexpressive activations. In Meila, M. and Zhang, T. (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 11932–11940. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/yarotsky21a.html>.
- Yarotsky, D. and Zhevnerchuk, A. The phase diagram of approximation rates for deep neural networks. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. F., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 13005–13015. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/979a3f14bae523dc5101c52120c535e9-Paper.pdf>.
- Zhang, L., Song, J., Gao, A., Chen, J., Bao, C., and Ma, K. Be your own teacher: Improve the performance of convolutional neural networks via self distillation. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 3712–3721, Los Alamitos, CA, USA, nov 2019. IEEE Computer Society. doi: 10.1109/ICCV.2019.00381.
- Zhang, L., Bao, C., and Ma, K. Self-distillation: Towards efficient and compact neural networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–1, 2021. doi: 10.1109/TPAMI.2021.3067100.
- Zhang, S. Deep neural network approximation via function compositions. *PhD Thesis, National University of Singapore*, 2020. URL: <https://scholarbank.nus.edu.sg/handle/10635/186064>.
- Zhou, D.-X. Universality of deep convolutional neural networks. *Applied and Computational Harmonic Analysis*, 48(2):787–794, 2020. ISSN 1063-5203. doi: 10.1016/j.acha.2019.06.004.

A. Proofs of main theorems

In this section, we first list all notations used throughout this paper. Then, we prove Theorems 2.1, 2.2, and 2.3 based on an auxiliary theorem, Theorem A.1, which will be proved in Section B.

A.1. Notations

Let us summarize the main notations as follows.

- Let \mathbb{R} , \mathbb{Q} , and \mathbb{Z} denote the set of real numbers, rational numbers, and integers, respectively.
- Let \mathbb{N} and \mathbb{N}^+ denote the set of natural numbers and positive natural numbers, respectively. That is, $\mathbb{N}^+ = \{1, 2, 3, \dots\}$ and $\mathbb{N} = \mathbb{N}^+ \cup \{0\}$.
- Vectors and matrices are denoted in a bold font. Standard vectorization is adopted in the matrix and vector computation. For example, adding a scalar and a vector means adding the scalar to each entry of the vector.
- For $\theta \in [0, 1)$, suppose its binary representation is $\theta = \sum_{\ell=1}^{\infty} \theta_{\ell} 2^{-\ell}$ with $\theta_{\ell} \in \{0, 1\}$, we introduce a special notation $\text{bin}0.\theta_1\theta_2 \cdots \theta_L$ to denote the L -term binary representation of θ , i.e., $\text{bin}0.\theta_1\theta_2 \cdots \theta_L := \sum_{\ell=1}^L \theta_{\ell} 2^{-\ell}$.
- For any $p \in [1, \infty)$, the p -norm of a vector $\mathbf{x} = (x_1, x_2, \dots, x_d) \in \mathbb{R}^d$ is defined by

$$\|\mathbf{x}\|_p := (|x_1|^p + |x_2|^p + \cdots + |x_d|^p)^{1/p}.$$

- The expression “a network with (of) width N and depth L ” means
 - The maximum width of this network for all **hidden** layers is no more than N .
 - The number of **hidden** layers of this network is no more than L .
- Similar to “min” and “max”, let $\text{mid}(x_1, x_2, x_3)$ be the middle value (median) of three inputs x_1 , x_2 , and x_3 . For example, $\text{mid}(2, 1, 3) = 2$ and $\text{mid}(3, 2, 3) = 3$.
- Given any $K \in \mathbb{N}^+$ and $\delta \in (0, \frac{1}{K}]$, define a trifling region $\Omega([0, 1]^d, K, \delta)$ of $[0, 1]^d$ as

$$\Omega([0, 1]^d, K, \delta) := \bigcup_{i=1}^d \left\{ \mathbf{x} = (x_1, x_2, \dots, x_d) : x_i \in \bigcup_{k=1}^{K-1} \left(\frac{k}{K} - \delta, \frac{k}{K} \right) \right\}. \quad (4)$$

In particular, $\Omega([0, 1]^d, K, \delta) = \emptyset$ if $K = 1$. See Figure 6 for two examples of trifling regions.




Figure 6. Two examples of trifling regions. (a) $K = 5, d = 1$. (b) $K = 4, d = 2$.

- Given a univariate activation function σ , let us introduce the architecture of a σ -activated network, i.e., a network with each hidden neuron activated by σ . To be precise, a σ -activated network with a vector input $\mathbf{x} \in \mathbb{R}^d$, an output $\phi(\mathbf{x}) \in \mathbb{R}$, and $L \in \mathbb{N}^+$ hidden layers can be briefly described as follows:

$$\mathbf{x} = \tilde{\mathbf{h}}_0 \xrightarrow[\mathcal{L}_0]{\mathbf{A}_0, \mathbf{b}_0} \mathbf{h}_1 \xrightarrow{\sigma} \tilde{\mathbf{h}}_1 \cdots \xrightarrow[\mathcal{L}_{L-1}]{\mathbf{A}_{L-1}, \mathbf{b}_{L-1}} \mathbf{h}_L \xrightarrow{\sigma} \tilde{\mathbf{h}}_L \xrightarrow[\mathcal{L}_L]{\mathbf{A}_L, \mathbf{b}_L} \mathbf{h}_{L+1} = \phi(\mathbf{x}), \quad (5)$$

where $N_0 = d \in \mathbb{N}^+$, $N_1, N_2, \dots, N_L \in \mathbb{N}^+$, $N_{L+1} = 1$, $\mathbf{A}_i \in \mathbb{R}^{N_{i+1} \times N_i}$ and $\mathbf{b}_i \in \mathbb{R}^{N_{i+1}}$ are the weight matrix and the bias vector in the i -th affine linear transform \mathcal{L}_i , respectively, i.e.,

$$\mathbf{h}_{i+1} = \mathbf{A}_i \cdot \tilde{\mathbf{h}}_i + \mathbf{b}_i =: \mathcal{L}_i(\tilde{\mathbf{h}}_i) \quad \text{for } i = 0, 1, \dots, L$$

and

$$\tilde{h}_{i,j} = \sigma(h_{i,j}) \quad \text{for } j = 1, 2, \dots, N_i \text{ and } i = 1, 2, \dots, L.$$

Here, $\tilde{h}_{i,j}$ and $h_{i,j}$ are the j -th entry of $\tilde{\mathbf{h}}_i$ and \mathbf{h}_i , respectively, for $j = 1, 2, \dots, N_i$ and $i = 1, 2, \dots, L$. If σ is applied to a vector entrywisely, i.e.,

$$\sigma(\mathbf{y}) = (\sigma(y_1), \dots, \sigma(y_d)) \quad \text{for any } \mathbf{y} = (y_1, \dots, y_d) \in \mathbb{R}^d,$$

then ϕ can be represented in a form of function compositions as follows:

$$\phi(\mathbf{x}) = \mathcal{L}_L \circ \sigma \circ \dots \circ \sigma \circ \mathcal{L}_1 \circ \sigma \circ \mathcal{L}_0(\mathbf{x}) \quad \text{for any } \mathbf{x} \in \mathbb{R}^d.$$

See Figure 7 for an example.




Figure 7. An example of a σ -activated network with width 5 and depth 2.

A.2. Proof of Theorem 2.1

To prove Theorems 2.1 and 2.2, we introduce an auxiliary theorem below with a similar result ignoring the approximation inside the trifling region.

Theorem A.1. *Given any $n \in \mathbb{N}^+$, there exist $\phi_1 \in \mathcal{H}_{2^{dn+4}}(d, 1)$ and $\phi_2 \in \mathcal{H}_{2^{dn+5}n}(n, 1)$ such that: For any continuous function $f : [0, 1]^d \rightarrow [0, 1]$, there exists a linear map $\mathcal{L} : \mathbb{R} \rightarrow \mathbb{R}^n$ satisfying $\|\phi_2 \circ \mathcal{L} \circ \phi_1\|_{L^\infty(\mathbb{R}^d)} \leq 1$ and*

$$|\phi_2 \circ \mathcal{L} \circ \phi_1(\mathbf{x}) - f(\mathbf{x})| \leq \omega_f(\sqrt{d}2^{-n}) + 2^{-n} \quad \text{for any } \mathbf{x} \in [0, 1]^d \setminus \Omega([0, 1]^d, K, \delta),$$

where $K = 2^n$, δ is an arbitrary number in $(0, \frac{1}{3K}]$, and \mathcal{L} is given by $\mathcal{L}(t) = (a_1 t, a_2 t, \dots, a_n t)$. Here, $a_i \in [0, \frac{1}{3}]$ is determined by f and n , and $2^m a_i \in \mathbb{N}$ for $i = 1, 2, \dots, n$, where $m = 2^{dn+1}$.

The proof of Theorem A.1 can be found later in Section B. Let us first prove Theorem 2.1 based on Theorem A.1.

Proof of Theorem 2.1. We may assume f is not a constant function since it is a trivial case. Then $\omega_f(r) > 0$ for any $r > 0$. Set $s = 2\omega_f(\sqrt{d}) > 0$ and $b = f(\mathbf{0}) - \omega_f(\sqrt{d})$. Then, by defining

$$\tilde{f} := \frac{f - b}{s} = \frac{f - f(\mathbf{0}) + \omega_f(\sqrt{d})}{2\omega_f(\sqrt{d})},$$

we have $\tilde{f}(\mathbf{x}) \in [0, 1]$ for any $\mathbf{x} \in [0, 1]^d$. By applying Theorem A.1 to \tilde{f} , there exist two functions, $\phi_1 : \mathbb{R}^d \rightarrow \mathbb{R}$ and $\phi_2 : \mathbb{R}^n \rightarrow \mathbb{R}$, both of which are independent of f and can be implemented by ReLU networks with $\leq 2^{dn+4}$ and $\leq 2^{dn+5}n$ parameters, respectively, and a linear map $\mathcal{L} : \mathbb{R} \rightarrow \mathbb{R}^n$ satisfying $\|\phi_2 \circ \mathcal{L} \circ \phi_1\|_{L^\infty(\mathbb{R}^d)} \leq 1$ and

$$|\phi_2 \circ \mathcal{L} \circ \phi_1(\mathbf{x}) - \tilde{f}(\mathbf{x})| \leq \omega_{\tilde{f}}(\sqrt{d}2^{-n}) + 2^{-n} \quad \text{for any } \mathbf{x} \in [0, 1]^d \setminus \Omega([0, 1]^d, K, \delta),$$

where $K = 2^n$, δ is an arbitrary number in $(0, \frac{1}{3K}]$, and \mathcal{L} is given by $\mathcal{L}(t) = (a_1 t, a_2 t, \dots, a_n t)$ with $a_1, a_2, \dots, a_n \in [0, \frac{1}{3})$ determined by \tilde{f} and n . Since \tilde{f} is derived from f , a_1, a_2, \dots, a_n are essentially determined by f and n .

Then choose a small δ satisfying

$$dK\delta 2^p \leq 2^{-pn} = (2^{-n})^p.$$

Note that the Lebesgue measure of $\Omega([0, 1]^d, K, \delta)$ is bounded by $dK\delta$ and

$$|\phi_2 \circ \mathcal{L} \circ \phi_1(\mathbf{x}) - \tilde{f}(\mathbf{x})| \leq |\phi_2 \circ \mathcal{L} \circ \phi_1(\mathbf{x})| + |\tilde{f}(\mathbf{x})| \leq 1 + 1 = 2 \quad \text{for any } \mathbf{x} \in [0, 1]^d.$$

Then, $\|\phi_2 \circ \mathcal{L} \circ \phi_1 - \tilde{f}\|_{L^p([0, 1]^d)}^p$ is bounded by

$$\begin{aligned} & \int_{[0, 1]^d \setminus \Omega([0, 1]^d, K, \delta)} |\phi_2 \circ \mathcal{L} \circ \phi_1(\mathbf{x}) - \tilde{f}(\mathbf{x})|^p d\mathbf{x} + \int_{\Omega([0, 1]^d, K, \delta)} |\phi_2 \circ \mathcal{L} \circ \phi_1(\mathbf{x}) - \tilde{f}(\mathbf{x})|^p d\mathbf{x} \\ & \leq (\omega_{\tilde{f}}(\sqrt{d}2^{-n}) + 2^{-n})^p + dK\delta 2^p \leq (\omega_{\tilde{f}}(\sqrt{d}2^{-n}) + 2^{-n})^p + (2^{-n})^p \leq (\omega_{\tilde{f}}(\sqrt{d}2^{-n}) + 2^{-n+1})^p, \end{aligned}$$

implying $\|\phi_2 \circ \mathcal{L} \circ \phi_1 - \tilde{f}\|_{L^p([0, 1]^d)} \leq \omega_{\tilde{f}}(\sqrt{d}2^{-n}) + 2^{-n+1}$. Note that $\omega_f(r) = s \cdot \omega_{\tilde{f}}(r)$ for any $r \geq 0$. Therefore, we have

$$\begin{aligned} & \|s \cdot (\phi_2 \circ \mathcal{L} \circ \phi_1) + b - f\|_{L^p([0, 1]^d)} = \left\| s \cdot (\phi_2 \circ \mathcal{L} \circ \phi_1) + b - (s \cdot \tilde{f} + b) \right\|_{L^p([0, 1]^d)} \\ & = s \|\phi_2 \circ \mathcal{L} \circ \phi_1 - \tilde{f}\|_{L^p([0, 1]^d)} \leq s \cdot \omega_{\tilde{f}}(\sqrt{d}2^{-n}) + 2^{-n+1}s = \omega_f(\sqrt{d}2^{-n}) + 2^{-n+2}\omega_f(\sqrt{d}). \end{aligned}$$

So we finish the proof. \square

A.3. Proof of Theorem 2.2

Next, let us prove Theorem 2.2. To this end, we need to introduce the following lemma, which is actually Lemma 3.4 of (Lu et al., 2021) (or Lemma 3.11 of (Zhang, 2020)).

Lemma A.2 (Lemma 3.4 of (Lu et al., 2021)). *Given any $\varepsilon > 0$, $K \in \mathbb{N}^+$, and $\delta \in (0, \frac{1}{3K}]$, assume $f \in C([0, 1]^d)$ and $g : \mathbb{R}^d \rightarrow \mathbb{R}$ is a general function with*

$$|g(\mathbf{x}) - f(\mathbf{x})| \leq \varepsilon \quad \text{for any } \mathbf{x} \in [0, 1]^d \setminus \Omega([0, 1]^d, K, \delta).$$

Then

$$|\phi(\mathbf{x}) - f(\mathbf{x})| \leq \varepsilon + d \cdot \omega_f(\delta) \quad \text{for any } \mathbf{x} \in [0, 1]^d,$$

where $\phi := \phi_d$ is defined by induction through

$$\phi_{i+1}(\mathbf{x}) := \text{mid}(\phi_i(\mathbf{x} - \delta \mathbf{e}_{i+1}), \phi_i(\mathbf{x}), \phi_i(\mathbf{x} + \delta \mathbf{e}_{i+1})) \quad \text{for } i = 0, 1, \dots, d-1,$$

where $\phi_0 = g$ and $\{\mathbf{e}_i\}_{i=1}^d$ is the standard basis in \mathbb{R}^d .

With Lemma A.2 in hand, we are ready to present the proof of Theorem 2.2.

Proof of Theorem 2.2. We may assume f is not a constant function since it is a trivial case. Then $\omega_f(r) > 0$ for any $r > 0$. Set $s = 2\omega_f(\sqrt{d}) > 0$ and $b = f(\mathbf{0}) - \omega_f(\sqrt{d})$. Then, by defining

$$\tilde{f} := \frac{f - b}{s} = \frac{f - f(\mathbf{0}) + \omega_f(\sqrt{d})}{2\omega_f(\sqrt{d})},$$

we have $\tilde{f}(\mathbf{x}) \in [0, 1]$ for any $\mathbf{x} \in [0, 1]^d$. By applying Theorem A.1 to \tilde{f} , there exist two functions, $\psi_{1,0} : \mathbb{R}^d \rightarrow \mathbb{R}$ and $\psi_{2,0} : \mathbb{R}^n \rightarrow \mathbb{R}$, both of which are independent of \tilde{f} (or f) and can be implemented by ReLU networks with $\leq 2^{dn+4}$ and $\leq 2^{dn+5}n$ parameters, respectively, and a linear map $\mathcal{L}_{\mathbf{a},0} : \mathbb{R} \rightarrow \mathbb{R}^n$ satisfying $\|\psi_0\|_{L^\infty(\mathbb{R}^d)} \leq 1$ and

$$|\psi_0(\mathbf{x}) - \tilde{f}(\mathbf{x})| \leq \omega_{\tilde{f}}(\sqrt{d}2^{-n}) + 2^{-n} =: \varepsilon \quad \text{for any } \mathbf{x} \in [0, 1]^d \setminus \Omega([0, 1]^d, K, \delta),$$

where $\psi_0 := \psi_{2,0} \circ \mathcal{L}_{a,0} \circ \psi_{1,0}$, $K = 2^n$, δ is an arbitrary number in $(0, \frac{1}{3K}]$, and $\mathcal{L}_{a,0}$ is given by $\mathcal{L}_{a,0}(t) = (a_1 t, a_2 t, \dots, a_n t)$ with $\mathbf{a} = (a_1, a_2, \dots, a_n)$ determined by f and n . Moreover, $a_i \in [0, \frac{1}{3}]$ and $2^m a_i \in \mathbb{N}$ for $i = 1, 2, \dots, n$, where $m = 2^{dn+1}$.¹

Choose a small δ satisfying $d \cdot \omega_{\tilde{f}}(\delta) \leq 2^{-n}$. With $\psi_0 = \psi_{2,0} \circ \mathcal{L}_{a,0} \circ \psi_{1,0}$ in hand, we can define ψ_1, \dots, ψ_d by induction via

$$\psi_{i+1}(\mathbf{x}) := \text{mid}\left(\psi_i(\mathbf{x} - \delta \mathbf{e}_{i+1}), \psi_i(\mathbf{x}), \psi_i(\mathbf{x} + \delta \mathbf{e}_{i+1})\right) \quad \text{for } i = 0, 1, \dots, d-1.$$

The detailed iterative equations for $\psi_i : \mathbb{R}^d \rightarrow \mathbb{R}$, $\boldsymbol{\psi}_{1,i} : \mathbb{R}^d \rightarrow \mathbb{R}^{3^i}$, $\mathcal{L}_{a,i} : \mathbb{R}^{3^i} \rightarrow \mathbb{R}^{3^{i+1}}$, and $\psi_{2,i} : \mathbb{R}^{3^i} \rightarrow \mathbb{R}$, for $i = 1, 2, \dots, d$, are listed as follows.

- $\psi_i = \psi_{2,i} \circ \mathcal{L}_{a,i} \circ \boldsymbol{\psi}_{1,i}$.
- $\boldsymbol{\psi}_{1,i}(\mathbf{y}) = \left(\psi_{1,i-1}(\mathbf{y} - \delta \mathbf{e}_i), \psi_{1,i-1}(\mathbf{y}), \psi_{1,i-1}(\mathbf{y} + \delta \mathbf{e}_i)\right)$ for any $\mathbf{y} \in \mathbb{R}^d$.
- $\mathcal{L}_{a,i}(\mathbf{y}_1, \mathbf{y}_2, \mathbf{y}_3) = \left(\mathcal{L}_{a,i-1}(\mathbf{y}_1), \mathcal{L}_{a,i-1}(\mathbf{y}_2), \mathcal{L}_{a,i-1}(\mathbf{y}_3)\right)$ for any $\mathbf{y}_1, \mathbf{y}_2, \mathbf{y}_3 \in \mathbb{R}^{3^{i-1}}$.
- $\psi_{2,i}(\mathbf{y}_1, \mathbf{y}_2, \mathbf{y}_3) = \text{mid}\left(\psi_{2,i-1}(\mathbf{y}_1), \psi_{2,i-1}(\mathbf{y}_2), \psi_{2,i-1}(\mathbf{y}_3)\right)$ for any $\mathbf{y}_1, \mathbf{y}_2, \mathbf{y}_3 \in \mathbb{R}^{3^{i-1}d}$.

See the illustrations in Figure 8.

By Lemma A.2, we have

$$|\phi(\mathbf{x}) - \tilde{f}(\mathbf{x})| \leq \varepsilon + d \cdot \omega_{\tilde{f}}(\delta) \leq \omega_{\tilde{f}}(\sqrt{d}2^{-n}) + 2^{-n+1} \quad \text{for any } \mathbf{x} \in [0, 1]^d,$$

where $\phi := \psi_d = \psi_{2,d} \circ \mathcal{L}_{a,d} \circ \boldsymbol{\psi}_{1,d}$. By defining $\mathcal{L} := \mathcal{L}_{a,d}$, $\boldsymbol{\phi}_1 := \boldsymbol{\psi}_{1,d}$, and $\phi_2 := \psi_{2,d}$, we have

$$|\phi_2 \circ \mathcal{L} \circ \boldsymbol{\phi}_1(\mathbf{x}) - \tilde{f}(\mathbf{x})| \leq \omega_{\tilde{f}}(\sqrt{d}2^{-n}) + 2^{-n+1} \quad \text{for any } \mathbf{x} \in [0, 1]^d,$$

As shown in Figure 8, $\mathcal{L} := \mathcal{L}_{a,d}$ is a linear map from \mathbb{R}^{3^d} to $\mathbb{R}^{3^{d+1}}$ determined by $\mathbf{a} = (a_1, a_2, \dots, a_n) \in [0, \frac{1}{3}]^n$, which depends on f and n . Moreover, as shown in Figure 8, $\boldsymbol{\phi}_1 := \boldsymbol{\psi}_{1,d}$ and $\phi_2 := \psi_{2,d}$ are independent of f and can be implemented by ReLU networks with

$$\leq 3^d(2^{dn+4}) + 3d(d+1)(3^{d-1} + 3^{d-2} + \dots + 3^0) \leq 3^d 2^{dn+5}$$

and

$$\leq 3^d(2^{dn+5}n) + 280(3^{d-1} + 3^{d-2} + \dots + 3^0) \leq 3^d 2^{dn+8}n$$

parameters, respectively.²

Note that $\omega_f(r) = s \cdot \omega_{\tilde{f}}(r)$ for any $r \geq 0$. Therefore, we have

$$\begin{aligned} & \|s \cdot (\phi_2 \circ \mathcal{L} \circ \boldsymbol{\phi}_1) + b - f\|_{L^\infty([0,1]^d)} = \left\| s \cdot (\phi_2 \circ \mathcal{L} \circ \boldsymbol{\phi}_1) + b - (s \cdot \tilde{f} + b) \right\|_{L^\infty([0,1]^d)} \\ & = s \|\phi_2 \circ \mathcal{L} \circ \boldsymbol{\phi}_1 - \tilde{f}\|_{L^\infty([0,1]^d)} \leq s \cdot \omega_{\tilde{f}}(\sqrt{d}2^{-n}) + 2^{-n+1}s = \omega_f(\sqrt{d}2^{-n}) + 2^{-n+2}\omega_f(\sqrt{d}). \end{aligned}$$


So we finish the proof. \square

A.4. Proof of Theorem 2.3


To simplify the proof of Theorem 2.3, we introduce two lemmas below. First, we need to establish a lemma showing how to store many parameters in one intrinsic parameter via a fixed network.

¹This property will be used in the proof of Theorem 2.3.


²As shown Lemma 3.1 of (Lu et al., 2021), “mid(\cdot, \cdot, \cdot)” can be implemented by a ReLU network with width 14 and depth 2, which has $\leq (3+1) \times 14 + (14+1) \times 14 + (14+1) = 280$ parameters.



(a) An illustration of the network architecture implementing $\psi_1 = \psi_{2,1} \circ \mathcal{L}_{a,1} \circ \psi_{1,1}$ based on $\psi_0 = \psi_{2,0} \circ \mathcal{L}_{a,0} \circ \psi_{1,0}$. The top architecture is in detail, while the bottom one is just a sketch of the top one. The orange numbers indicate the number of neurons in each layer.



(b) An illustration of the network architecture implementing $\psi_2 = \psi_{2,2} \circ \mathcal{L}_{a,2} \circ \psi_{1,2}$ based on $\psi_1 = \psi_{2,1} \circ \mathcal{L}_{a,1} \circ \psi_{1,1}$. The top architecture is in detail, while the bottom one is just a sketch of the top one. The orange numbers indicate the number of neurons in each layer.



(c) An illustration of the network architecture implementing $\psi_3 = \psi_{2,3} \circ \mathcal{L}_{a,3} \circ \psi_{1,3}$ based on $\psi_2 = \psi_{2,2} \circ \mathcal{L}_{a,2} \circ \psi_{1,2}$. The top architecture is in detail, while the bottom one is just a sketch of the top one. The orange numbers indicate the number of neurons in each layer.

Figure 8. Illustrations of the implementations of ϕ_1 , ϕ_2 , and ϕ_3 . The inductive implementations of ϕ_4, \dots, ϕ_d are similar.

Lemma A.3. *Given any $m, n \in \mathbb{N}$, there exists a vector-valued function $\phi : \mathbb{R} \rightarrow \mathbb{R}^n$ realized by a ReLU network such that: For any $a_i \in [0, 1)$ with $2^m a_i \in \mathbb{N}$ for $i = 0, 1, \dots, n$, there exists a real number $v \in [0, 1)$ such that*

$$\phi(v) = (a_1, a_2, \dots, a_n).$$

Next, we establish another lemma using a ReLU network to uniformly approximate multiplication operation $\psi(x, y) = xy$ well.

Lemma A.4. *For any $M > 0$ and $\eta > 0$, there exists a function $\psi_\eta : \mathbb{R}^2 \rightarrow \mathbb{R}$ realized by a ReLU network such that*

$$\psi_\eta(x, y) \rightrightarrows \psi(x, y) = xy \quad \text{on } [-M, M]^2 \quad \text{as } \eta \rightarrow 0^+,$$

where \rightrightarrows denotes the uniform convergence.

The proof of Lemma A.3 is placed later in this section. Lemma A.4 is just a direct result of Lemma 4.2 of (Lu et al., 2021). With Lemmas A.3 and A.4 in hand, we are ready to prove Theorem 2.3.

Proof of Theorem 2.3. For any $\varepsilon > 0$, choose a large $n = n(\varepsilon, \alpha, \lambda) \in \mathbb{N}^+$ such that

$$5\lambda d^{\alpha/2} 2^{-\alpha n} \leq \varepsilon/2.$$

Since f is a Hölder continuous function on $[0, 1]^d$ of order $\alpha \in (0, 1]$ with a Hölder constant $\lambda > 0$, we have $\omega_f(r) \leq \lambda r^\alpha$ for any $r \geq 0$. By Theorem 2.2, there exist two functions $\phi_1 : \mathbb{R}^d \rightarrow \mathbb{R}^{3^d}$ and $\phi_2 : \mathbb{R}^{3^d n} \rightarrow \mathbb{R}$, implemented by f -independent ReLU networks, such that

$$\|s(\phi_2 \circ \mathcal{L} \circ \phi_1) + b - f\|_{L^\infty([0,1]^d)} \leq \omega_f(\sqrt{d}2^{-n}) + 2^{-n+2}\omega_f(\sqrt{d}) \leq 5\lambda d^{\alpha/2} 2^{-\alpha n} \leq \varepsilon/2, \quad (6)$$

where $s = 2\omega_f(\sqrt{d}) \leq 2\lambda d^{\alpha/2}$, $b = f(\mathbf{0}) - \omega_f(\sqrt{d})$, and $\mathcal{L} : \mathbb{R}^{3^d} \rightarrow \mathbb{R}^{3^d n}$ is a linear map given by

$$\mathcal{L}(y_1, \dots, y_{3^d}) = \left(\mathcal{L}_0(y_1), \dots, \mathcal{L}_0(y_{3^d}) \right) \quad \text{for any } \mathbf{y} = (y_1, \dots, y_{3^d}) \in \mathbb{R}^{3^d},$$

where $\mathcal{L}_0 : \mathbb{R} \rightarrow \mathbb{R}^n$ is a linear map given by $\mathcal{L}_0(t) = (a_1 t, a_2 t, \dots, a_n t)$ with $a_1, a_2, \dots, a_n \in [0, \frac{1}{3})$ determined by f and n . Since a_1, a_2, \dots, a_n are repeated 3^d times in the definition of \mathcal{L} , there are $3^d n$ parameters in total. We will show how to store these $3^d n$ parameters in one intrinsic parameter v via an f -independent ReLU network as shown in the following two steps.

- Regard a_1, a_2, \dots, a_n as inputs, but not parameters. See the difference in Figure 9. Then, we only need to store a_1, a_2, \dots, a_n one time by copying them 3^d times.
- As stated in the proof of Theorem 2.2, a_1, a_2, \dots, a_n have finite binary representations. Then, we can store them in a key parameter v and use an f -independent ReLU network to extract them from v .

The details of these two steps can be found below.

Step 1: Regard a_1, a_2, \dots, a_n as inputs and copy them 3^d times.

Since a_1, a_2, \dots, a_n are regarded as inputs, the implementation of $\mathcal{L}_0(t) = (a_1 t, a_2 t, \dots, a_n t)$ requires multiplication operations. This means that we need to approximate $\psi(x, y) = xy$ well via an f -independent ReLU network. See Figure 9 for illustrations.

Denote $\mathbf{a} = (a_1, a_2, \dots, a_n)$ and

$$\phi_1(\mathbf{x}) = \left(\phi_{1,1}(\mathbf{x}), \phi_{1,2}(\mathbf{x}), \dots, \phi_{1,3^d}(\mathbf{x}) \right) \quad \text{for any } \mathbf{x} \in [0, 1]^d.$$

Then define

$$M := 1 + \sup \left\{ |\phi_{1,j}(\mathbf{x})| : \mathbf{x} \in [0, 1]^d, j = 1, 2, \dots, 3^d \right\}.$$




Figure 9. Illustrations of two methods getting/approximating $a_i t$ for $i = 1, 2, \dots, n$ and $t \in \{\phi_{1,j}(\mathbf{x}) : \mathbf{x} \in [0, 1]^d, j = 1, 2, \dots, 3^d\}$. (a) By regarding a_i as a parameter, one can easily get the product of an input t and an parameter a_i . (b) By regarding a_i as an input, one needs to use a ReLU network to approximate the multiplication operation for approximating $a_i t$ well.

By Lemma A.4, there exists a function $\psi_\eta : \mathbb{R}^2 \rightarrow \mathbb{R}$ realized by a ReLU network such that

$$\psi_\eta(x, y) \rightrightarrows \psi(x, y) = xy \quad \text{on } [-M, M]^2 \quad \text{as } \eta \rightarrow 0^+.$$

Note that \mathcal{L} and \mathcal{L}_0 depend on $\mathbf{a} = (a_1, a_2, \dots, a_n)$. For clarity, we denote $\mathcal{L}_{\mathbf{a}} = \mathcal{L}$ and $\mathcal{L}_{\mathbf{a},0} = \mathcal{L}_0$. For any $a_1, a_2, \dots, a_n \in [0, \frac{1}{3}] \subseteq [-M, M]$ and

$$t \in \left\{ \phi_{1,j}(\mathbf{x}) : \mathbf{x} \in [0, 1]^d, j = 1, 2, \dots, 3^d \right\} \subseteq [-M, M],$$

we can use

$$\mathcal{L}_{\mathbf{a},0,\eta}(t) := \left(\psi_\eta(a_1, t), \psi_\eta(a_2, t), \dots, \psi_\eta(a_n, t) \right)$$

to approximate

$$\mathcal{L}_{\mathbf{a},0}(t) = (a_1 t, a_2 t, \dots, a_n t) = \left(\psi(a_1, t), \psi(a_2, t), \dots, \psi(a_n, t) \right).$$

Then

$$\mathcal{L}_{\mathbf{a},\eta}(y_1, \dots, y_{3^d}) := \left(\mathcal{L}_{\mathbf{a},0,\eta}(y_1), \dots, \mathcal{L}_{\mathbf{a},0,\eta}(y_{3^d}) \right)$$

can also approximate

$$\mathcal{L}_{\mathbf{a}}(y_1, \dots, y_{3^d}) = \left(\mathcal{L}_{\mathbf{a},0}(y_1), \dots, \mathcal{L}_{\mathbf{a},0}(y_{3^d}) \right)$$

well. Define $\tilde{\phi}_\eta : \mathbb{R}^{n+3^d} \rightarrow \mathbb{R}^{3^d n}$ and $\tilde{\phi} : \mathbb{R}^{n+3^d} \rightarrow \mathbb{R}^{3^d n}$ via

$$\tilde{\phi}_\eta(\mathbf{y}, \mathbf{a}) := \mathcal{L}_{\mathbf{a},\eta}(y_1, \dots, y_{3^d}) \quad \text{for any } \mathbf{a} \in [0, \frac{1}{3}]^n \text{ and } \mathbf{y} = (y_1, \dots, y_{3^d}) \in [-M, M]^{3^d}$$

and

$$\tilde{\phi}(\mathbf{y}, \mathbf{a}) := \mathcal{L}_{\mathbf{a}}(y_1, \dots, y_{3^d}) \quad \text{for any } \mathbf{a} \in [0, \frac{1}{3}]^n \text{ and } \mathbf{y} = (y_1, \dots, y_{3^d}) \in [-M, M]^{3^d}.$$

Since

$$\psi_\eta(x, y) \rightrightarrows \psi(x, y) = xy \quad \text{on } [-M, M]^2 \quad \text{as } \eta \rightarrow 0^+,$$

it is easy to verify that

$$\tilde{\phi}_\eta(\mathbf{y}, \mathbf{a}) \rightrightarrows \tilde{\phi}(\mathbf{y}, \mathbf{a}) \quad \text{for } \mathbf{y} \in [-M, M]^{3^d} \text{ and } \mathbf{a} \in [0, \frac{1}{3}]^n \quad \text{as } \eta \rightarrow 0^+.$$

Note that $\phi_1(\mathbf{x}) \in [-M, M]^{3^d}$ for any $\mathbf{x} \in [0, 1]^d$. Then,

$$\phi_2 \circ \tilde{\phi}_\eta(\phi_1(\mathbf{x}), \mathbf{a}) \rightrightarrows \phi_2 \circ \tilde{\phi}(\phi_1(\mathbf{x}), \mathbf{a}) \quad \text{for } \mathbf{x} \in [0, 1]^d \text{ and } \mathbf{a} \in [0, \frac{1}{3}]^n \quad \text{as } \eta \rightarrow 0^+.$$

The fact $\tilde{\phi}(\mathbf{y}, \mathbf{a}) = \mathcal{L}_{\mathbf{a}}(\mathbf{y}) = \mathcal{L}(\mathbf{y})$ implies $\phi_2 \circ \tilde{\phi}(\phi_1(\mathbf{x}), \mathbf{a}) = \phi_2 \circ \mathcal{L} \circ \phi_1(\mathbf{x})$. Therefore,

$$\phi_2 \circ \tilde{\phi}_\eta(\phi_1(\mathbf{x}), \mathbf{a}) \rightrightarrows \phi_2 \circ \mathcal{L} \circ \phi_1(\mathbf{x}) \quad \text{for } \mathbf{x} \in [0, 1]^d \text{ and } \mathbf{a} \in [0, \frac{1}{3}]^n \quad \text{as } \eta \rightarrow 0^+.$$

Choose a small $\eta = \eta(n) > 0$ such that

$$\left| \phi_2 \circ \tilde{\phi}_\eta(\phi_1(\mathbf{x}), \mathbf{a}) - \phi_2 \circ \mathcal{L} \circ \phi_1(\mathbf{x}) \right| \leq 2^{-n} \quad \text{for any } \mathbf{x} \in [0, 1]^d \text{ and } \mathbf{a} \in [0, \frac{1}{3}]^n. \quad (7)$$

Recall that ψ_η can be realized by an f -independent ReLU network. It is easy to verify that $\tilde{\phi}_\eta$ can also be realized by an f -independent ReLU network.

Step 2: Store a_1, a_2, \dots, a_n in a key parameter v .

As we can see from the proof of Theorem 2.2, $a_i \in [0, \frac{1}{3}]$ with $2^m a_i \in \mathbb{N}$ for $i = 1, 2, \dots, n$, where $m = 2^{dn+1}$. That is,

$$a_i \in \left\{ \text{bin} 0.\theta_1 \cdots \theta_m : \theta_\ell \in \{0, 1\}, \ell = 1, 2, \dots, m \right\}.$$

Then, by Lemma A.3, there exists a real number $v \in [0, 1)$ and a vector function $\phi_0 : \mathbb{R} \rightarrow \mathbb{R}^n$ implemented by a ReLU network independent of a_1, a_2, \dots, a_n such that

$$\phi_0(v) = (a_1, a_2, \dots, a_n) = \mathbf{a}.$$

Next, we can define the final network-generated function $\phi : \mathbb{R}^{d+1} \rightarrow \mathbb{R}$ by

$$\phi(\mathbf{x}, v) := \phi_2 \circ \tilde{\phi}_\eta(\phi_1(\mathbf{x}), \phi_0(v)) = \phi_2 \circ \tilde{\phi}_\eta(\phi_1(\mathbf{x}), \mathbf{a})$$

for any $\mathbf{x} \in [0, 1]^d$ and $v \in [0, 1)$.

Since the ReLU network realizing ϕ_0 is independent of a_1, a_2, \dots, a_n , and hence independent of f . Recall that ϕ_1, ϕ_2 , and $\tilde{\phi}_\eta$ can be implemented by f -independent ReLU networks. Hence,

$$\phi(\mathbf{x}, v) = \phi_2 \circ \tilde{\phi}_\eta(\phi_1(\mathbf{x}), \phi_0(v))$$

can also be implemented by an f -independent ReLU network. It remains to estimate the error. By Equations (6) and (7), we have

$$\begin{aligned} |s\phi(\mathbf{x}, v) + b - f(\mathbf{x})| &\leq \left| s\phi(\mathbf{x}, v) + b - (s\phi_2 \circ \mathcal{L} \circ \phi_1(\mathbf{x}) + b) \right| + \left| s\phi_2 \circ \mathcal{L} \circ \phi_1(\mathbf{x}) + b - f(\mathbf{x}) \right| \\ &\leq s \left| \phi(\mathbf{x}, v) - \phi_2 \circ \mathcal{L} \circ \phi_1(\mathbf{x}) \right| + \varepsilon/2 \\ &\leq 2\lambda d^{\alpha/2} \left| \phi_2 \circ \tilde{\phi}_\eta(\phi_1(\mathbf{x}), \mathbf{a}) - \phi_2 \circ \mathcal{L} \circ \phi_1(\mathbf{x}) \right| + \varepsilon/2 \\ &\leq 2\lambda d^{\alpha/2} 2^{-n} + \varepsilon/2 \leq 5\lambda d^{\alpha/2} 2^{-\alpha n} + \varepsilon/2 \leq \varepsilon/2 + \varepsilon/2 = \varepsilon. \end{aligned}$$

So we finish the proof. □

Finally, let us prove Lemma A.3 to end this section.

Proof of Lemma A.3. Since $a_i \in [0, 1)$ with $2^m a_i \in \mathbb{N}$ for $i = 1, 2, \dots, n$, a_i can be represented as a binary form

$$a_i = \text{bin} 0.a_{i,1} a_{i,2} \cdots a_{i,m}.$$

Denote

$$v = \sum_{i=1}^n 2^{-m(i-1)} a_i = \text{bin} 0. \underbrace{a_{1,1} \cdots a_{1,m}}_{\text{store } a_1} \underbrace{a_{2,1} \cdots a_{2,m}}_{\text{store } a_2} \cdots \underbrace{a_{n,1} \cdots a_{n,m}}_{\text{store } a_n},$$

which requires pretty high precision. It is easy to extract a_i from v via the floor function ($\lfloor \cdot \rfloor$), i.e.,

$$a_i = \lfloor 2^{mi} v \rfloor / 2^m - \lfloor 2^{m(i-1)} v \rfloor \quad \text{for } i = 1, 2, \dots, n.$$

Next, we need to use a ReLU network to replace the floor function. Let $g : \mathbb{R} \rightarrow \mathbb{R}$ be the continuous piecewise linear function with the following breakpoints:

$$(\ell, \ell) \quad \text{and} \quad (\ell + 1 - \delta, \ell) \quad \text{for } \ell = 0, 1, \dots, 2^{mn} - 1, \quad \text{where } \delta = 2^{-mn}.$$

Clearly, g can be realized by a ReLU network independent of a_1, a_2, \dots, a_n and

$$g(x) = \lfloor x \rfloor \quad \text{for any } x \in \bigcup_{\ell=0}^{2^{mn}-1} [\ell, \ell + 1 - \delta].$$

Note that $v \in \bigcup_{\ell=0}^{2^{mn}-1} [\ell, \ell + 1 - \delta]$. By defining

$$g_i(t) := g(2^{mi}t)/2^m - g(2^{m(i-1)}t) \quad \text{for } i = 1, 2, \dots, n \text{ and any } t \in \mathbb{R},$$

we have

$$a_i = g(2^{mi}v)/2^m - g(2^{m(i-1)}v) = g_i(v) \quad \text{for } i = 1, 2, \dots, n.$$

Next, The target function ϕ can be defined via

$$\phi(t) = \left(g_1(t), g_2(t), \dots, g_n(t) \right) \quad \text{for any } t \in \mathbb{R}.$$

Thus, we have

$$\phi(v) = \left(g_1(v), g_2(v), \dots, g_n(v) \right) = (a_1, a_2, \dots, a_n).$$

and ϕ can be realized by a ReLU network independent of a_1, a_2, \dots, a_n . So we finish the proof. \square

B. Proof of Theorem A.1

In this section, we first present the proof sketch of Theorem A.1 in Section B.1, and then give the detailed proof in Section B.2 based on Proposition B.1, which will be proved later in Section B.3.

B.1. Sketch of proof

Before proving Theorem A.1, let us present the key steps as follows.

1. Set $K = 2^n$, divide $[0, 1]^d$ into K^d cubes Q_β for $\beta \in \{0, 1, \dots, K-1\}^d$ and the trifling region $\Omega([0, 1]^d, K, \delta)$, and denote \mathbf{x}_β as the vertex of Q_β with minimum $\|\cdot\|_1$ norm for each β . See Figure 10 for illustrations.
2. Design a ReLU sub-network to implement a function $\phi_1 : \mathbb{R}^d \rightarrow \mathbb{R}$, independent of f , projecting the whole Q_β to a number determined by β in $\{4^j : j = 1, 2, \dots, K^d\}$ for each β .
3. Design a linear map $\mathcal{L} : \mathbb{R} \rightarrow \mathbb{R}^n$, given by $\mathcal{L}(t) = (a_1t, a_2t, \dots, a_nt)$, for later use, where a_1, a_2, \dots, a_n are determined by f and n .
4. Design a ReLU sub-network to implement a function $\phi_2 : \mathbb{R}^n \rightarrow \mathbb{R}$, independent of f , such that $\phi_2 \circ \mathcal{L} \circ \phi_1(\mathbf{x}) \approx f(\mathbf{x}_\beta)$ for any $\mathbf{x} \in Q_\beta$ and each $\beta \in \{0, 1, \dots, K-1\}^d$. Then $\phi_2 \circ \mathcal{L} \circ \phi_1$ approximates f well outside of $\Omega([0, 1]^d, K, \delta)$.
5. Estimate the approximation error of $\phi_2 \circ \mathcal{L} \circ \phi_1 \approx f$.

As we shall see later, the constructions of ϕ_1 and \mathcal{L} are not difficult. The most technical part is to design ϕ_2 implemented by a ReLU network, which relies on the following proposition.

Proposition B.1. *Given any $J \in \mathbb{N}^+$, there exists a function ϕ implemented by a ReLU network with width 2 and depth $2J + 2$ such that: For any $\theta_1, \theta_2, \dots, \theta_J \in \{0, 1\}$, we have $\phi(x) \in [0, 1]$ for any $x \in \mathbb{R}$ and*

$$\phi(4^j a) = \theta_j \quad \text{for } j = 1, 2, \dots, J, \quad \text{where } a = \sum_{j=1}^J \theta_j 4^{-j}. \quad (8)$$

The proof of this proposition can be found in Section B.3. We shall point out that the function ϕ in this proposition is independent of $\theta_1, \theta_2, \dots, \theta_J \in \{0, 1\}$.




Figure 10. Illustrations of $\Omega([0, 1]^d, K, \delta)$, Q_β and \mathbf{x}_β for $\beta \in \{0, 1, \dots, K-1\}^d$. (a) $K = 4$, $d = 1$. (b) $K = 4$, $d = 2$.

B.2. Constructive proof

Now we are ready to give the detailed proof of Theorem A.1.

Proof of Theorem A.1. The proof consists of five steps.

Step 1: Set up.

Set $K = 2^n$ and let $\delta > 0$ be a small number determined later. Then define $\mathbf{x}_\beta := \beta/K$ and divide $[0, 1]^d$ into K^d cubes Q_β for $\beta \in \{0, 1, \dots, K-1\}^d$ and a small region $\Omega([0, 1]^d, K, \delta)$. Namely,

$$Q_\beta := \left\{ \mathbf{x} = (x_1, x_2, \dots, x_d) : x_i \in \left[\frac{\beta_i}{K}, \frac{\beta_i+1}{K} - \delta \right] \text{ for } i = 1, \dots, d \right\},$$

for $\beta = (\beta_1, \beta_2, \dots, \beta_d) \in \{0, 1, \dots, K-1\}^d$. Clearly, $\Omega([0, 1]^d, K, \delta) = [0, 1]^d \setminus (\bigcup_{\beta \in \{0, 1, \dots, K-1\}^d} Q_\beta)$. See Figure 10 for illustrations.

Step 2: Construct ϕ_1 .

Let g be a ‘‘step function’’ such that

- $g(\frac{k}{K}) = g(\frac{k+1}{K} - \delta) = k$ for $k = 0, 1, \dots, K-1$ and $g(1) = K-1$.
- g is linear between any two adjacent points of

$$\left\{ \frac{k}{K} : k = 0, 1, \dots, K \right\} \cup \left\{ \frac{k+1}{K} - \delta : k = 0, 1, \dots, K-1 \right\}.$$

Then, for any $\mathbf{x} = (x_1, \dots, x_d) \in Q_\beta$ and $\beta = (\beta_1, \dots, \beta_d) \in \{0, 1, \dots, K-1\}^d$, we have

$$g(x_i) = \beta_i \quad \text{for } i = 1, 2, \dots, d.$$

Also, such a function g can be easily realized by a one-hidden-layer ReLU network with width $2K$.

Let h be a function satisfying $h(j) = 4^j$ for $j = 1, 2, \dots, K^d$. Such a function h can be easily realized by a one-hidden-layer ReLU network with width K^d . Then the desired function $\phi_1 : \mathbb{R}^d \rightarrow \mathbb{R}$ can be defined via

$$\phi_1(\mathbf{x}) := h\left(1 + \sum_{i=1}^d g(x_i)K^{i-1}\right) = h \circ \varrho\left(g(x_1), \dots, g(x_d)\right) \quad \text{for any } \mathbf{x} = (x_1, \dots, x_d) \in \mathbb{R}^d,$$

where $\varrho : \mathbb{R}^d \rightarrow \mathbb{R}$ is a linear function defined by

$$\varrho(\mathbf{y}) = 1 + \sum_{i=1}^d y_i K^{i-1} \quad \text{for any } \mathbf{y} = (y_1, \dots, y_d) \in \mathbb{R}^d.$$

Clearly, ϱ is a bijection (one-to-one map) from $\beta \in \{0, 1, \dots, K-1\}^d$ to $\varrho(\beta) = 1 + \sum_{i=1}^d \beta_i K^{i-1} \in \{1, 2, \dots, K^d\}$.

Then, for any $\mathbf{x} \in Q_\beta$ and $\beta \in \{0, 1, \dots, K-1\}^d$, we have

$$\phi_1(\mathbf{x}) = h \circ \varrho(g(x_1), \dots, g(x_d)) = h \circ \varrho(\beta_1, \dots, \beta_d) = 4^{\varrho(\beta)} = 4^{1+\sum_{i=1}^d \beta_i K^{i-1}}. \quad (9)$$

Apparently, ϕ_1 is independent of f and it can be realized by a ReLU network with

$$\leq d(2K \times 3 + 1) + (K + 1) + 3K^d + 1 \leq 11K^d + 2 = 11 \times 2^{dn} + 2 \leq 2^{dn+4}$$

parameters.

Step 3: Construct \mathcal{L} .

For each $\beta \in \{0, 1, \dots, K-1\}^d$, it follows from $f(\mathbf{x}_\beta) \in [0, 1]$ that there exist $\xi_{\beta,1}, \dots, \xi_{\beta,n}$ such that

$$|f(\mathbf{x}_\beta) - \text{bin}_{0,\xi_{\beta,1} \cdots \xi_{\beta,n}}| \leq 2^{-n}. \quad (10)$$

Given any $j \in \{1, 2, \dots, K^d\}$, there exists a unique $\beta \in \{0, 1, \dots, K-1\}^d$ such that $j = 1 + \sum_{i=1}^d \beta_i K^{i-1} = \varrho(\beta)$. Thus, for any $\ell \in \{1, 2, \dots, n\}$, we can define

$$\theta_{j,\ell} := \xi_{\beta,\ell}, \quad \text{for } j = \varrho(\beta) \text{ and } \beta \in \{0, 1, \dots, K-1\}^d. \quad (11)$$

Then the desired linear map \mathcal{L} can be defined via

$$\mathcal{L}(t) := (a_1 t, a_2 t, \dots, a_n t) \quad \text{for any } t \in \mathbb{R},$$

where $a_\ell = \sum_{j=1}^{K^d} \theta_{j,\ell} 4^{-j}$ for $\ell = 1, 2, \dots, n$. Clearly, for $\ell = 1, 2, \dots, n$, we have

$$a_\ell = \sum_{j=1}^{K^d} \theta_{j,\ell} 4^{-j} \in [0, \frac{1}{3})$$

and

$$2^m a_\ell = 2^{2^{dn+1}} a_\ell = 4^{2^{dn}} a_\ell = 4^{K^d} a_\ell \in \mathbb{N}, \quad \text{where } m = 2^{dn+1}.$$

Step 4: Construct ϕ_2 .

Fix $\ell \in \{1, 2, \dots, n\}$, by Proposition B.1 (set $J = K^d$ and $\theta_j = \theta_{j,\ell}$ therein), there exists a function $\phi_{2,\ell}$ implemented by a ReLU network with width 2 and depth $2K^d + 2$ such that $\phi_{2,\ell}(t) \in [0, 1]$ for any $t \in \mathbb{R}$ and

$$\phi_{2,\ell}(4^j a_\ell) = \theta_{j,\ell} \quad \text{for } j = 1, 2, \dots, K^d, \quad \text{where } a_\ell = \sum_{j=1}^{K^d} \theta_{j,\ell} 4^{-j}. \quad (12)$$

Note that $\phi_{2,\ell}$ is independent of $\theta_{j,\ell}$ for $j = 1, 2, \dots, K^d$, so it is also independent of f . Then the desired function $\phi_2 : \mathbb{R}^n \rightarrow \mathbb{R}$ can be defined via

$$\phi_2(\mathbf{y}) := \sum_{\ell=1}^n 2^{-\ell} \phi_{2,\ell}(y_\ell) \quad \text{for any } \mathbf{y} = (y_1, \dots, y_n) \in \mathbb{R}^n.$$

Then $\phi_2(\mathbf{y}) \in [0, 1]$ for any $\mathbf{y} \in \mathbb{R}^n$ and ϕ_2 can be implemented by a ReLU network, independent of f , with

$$\leq \left(6(2K^n + 2) + (2 + 1)\right)n + n + 1 = 6n(2^{nd+1} + 2) + 4n + 1 \leq 2^{nd+5}n$$

parameters.

Step 5: Estimate the approximation error.

It remains to estimate the approximation error. By Equations (9), (11), and (12), for $\ell \in \{1, 2, \dots, n\}$, $\mathbf{x} \in Q_\beta$, $j = \varrho(\beta) = 1 + \sum_{i=1}^d \beta_i K^{i-1} \in \{1, 2, \dots, K^d\}$, and $\beta \in \{0, 1, \dots, K-1\}^d$, we have

$$\begin{aligned} \phi_2 \circ \mathcal{L} \circ \phi_1(\mathbf{x}) &= \phi_2 \circ \mathcal{L} \circ h \circ \varrho(g(x_1), \dots, g(x_d)) = \phi_2 \circ \mathcal{L} \circ h \circ \varrho(\beta) \\ &= \phi_2 \circ \mathcal{L}(4^{\varrho(\beta)}) = \phi_2 \circ \mathcal{L}(4^j) = \phi_2(4^j a_1, 4^j a_2, \dots, 4^j a_n) \\ &= \sum_{\ell=1}^n 2^{-\ell} \phi_{2,\ell}(4^j a_\ell) = \sum_{\ell=1}^n 2^{-\ell} \theta_{j,\ell} = \sum_{\ell=1}^n 2^{-\ell} \xi_{\beta,\ell}. \end{aligned}$$

Then by Equation (10), for any $\mathbf{x} \in Q_\beta$ and $\beta \in \{0, 1, \dots, K-1\}^d$, we get

$$\begin{aligned} |f(\mathbf{x}) - \phi_2 \circ \mathcal{L} \circ \phi_1(\mathbf{x})| &= |f(\mathbf{x}) - f(\mathbf{x}_\beta)| + |f(\mathbf{x}_\beta) - \phi_2 \circ \mathcal{L} \circ \phi_1(\mathbf{x})| \\ &\leq \omega_f\left(\frac{\sqrt{d}}{K}\right) + \left|f(\mathbf{x}_\beta) - \sum_{\ell=1}^n 2^{-\ell} \xi_{\beta,\ell}\right| \\ &\leq \omega_f\left(\frac{\sqrt{d}}{K}\right) + |f(\mathbf{x}_\beta) - \text{bin} 0.\xi_{\beta,1} \dots \xi_{\beta,n}| \leq \omega_f(\sqrt{d} 2^{-n}) + 2^{-n}. \end{aligned}$$

That is,

$$|f(\mathbf{x}) - \phi_2 \circ \mathcal{L} \circ \phi_1(\mathbf{x})| \leq \omega_f(\sqrt{d} 2^{-n}) + 2^{-n} \quad \text{for any } \mathbf{x} \in [0, 1]^d \setminus \Omega([0, 1]^d, K, \delta).$$

Moreover, the fact $\phi_2(\mathbf{y}) \in [0, 1]$ for any $\mathbf{y} \in \mathbb{R}^n$ implies $\|\phi_2 \circ \mathcal{L} \circ \phi_1\|_{L^\infty(\mathbb{R}^d)} \leq 1$. So we finish the proof. \square

B.3. Proof of Proposition B.1

Before proving Proposition B.1, let us introduce a notation to simplify the proof. We use \mathcal{T}_m for $m \in \mathbb{N}^+$ to denote a “sawtooth” function satisfying the following conditions.

- $\mathcal{T}_m : [0, 2m] \rightarrow [0, 1]$ is linear between any two adjacent integers of $\{0, 1, \dots, 2m\}$.
- $\mathcal{T}_m(2j) = 0$ for $j = 0, 1, \dots, m$ and $\mathcal{T}_m(2j+1) = 1$ for $j = 0, 1, \dots, m-1$.




Figure 11. Illustrations of teeth functions \mathcal{T}_1 , \mathcal{T}_2 , \mathcal{T}_3 and \mathcal{T}_4 .

To simplify the proof of Proposition B.1, we first introduce a lemma based on the “sawtooth” function.

Lemma B.2. Given any $J \in \mathbb{N}^+$ and $\theta_j \in \{0, 1\}$ for $j = 1, 2, \dots, J$, set $a = \sum_{j=1}^J \theta_j 4^{-j}$. Then

$$\mathcal{T}_{4^J}(4^J a) \in [0, 1/3] \quad \text{if } \theta_j = 0 \quad \text{and} \quad \mathcal{T}_{4^J}(4^J a) \in [2/3, 1] \quad \text{if } \theta_j = 1, \quad (13)$$

where $\mathcal{T}_{4^J} : [0, 2 \times 4^J] \rightarrow [0, 1]$ is a “sawtooth” function with 4^J “teeth” defined just above.

Proof. Fix $j \in \{1, 2, \dots, J\}$, we have

$$\begin{aligned}
 4^j a &= 4^j \sum_{i=1}^J \theta_i 4^{-i} = \underbrace{\sum_{i=1}^{j-1} \theta_i 4^{j-i}}_{= 4k \text{ for some } k \in \mathbb{N} \text{ with } k \leq \frac{4^{j-1}}{3}} + \overbrace{\theta_j}^{0 \text{ or } 1} + \underbrace{\sum_{i=j+1}^J \theta_i 4^{j-i}}_{\in [0, \frac{1}{3})}. \tag{14}
 \end{aligned}$$

Clearly,

$$\sum_{i=1}^{j-1} \theta_i 4^{j-i} \in \{4k : k \in \mathbb{N}, k \leq 4^{j-1}/3\} \quad \text{and} \quad 0 \leq \sum_{i=j+1}^J \theta_i 4^{j-i} \leq \sum_{i=j+1}^J 4^{j-i} \leq 1/3.$$

If $\theta_j = 0$, then Equation (14) implies

$$4^j a \in [4k, 4k + 1/3] \quad \text{for some } k \in \mathbb{N} \text{ with } k \leq 4^{j-1}/3 \leq 4^j - 1,$$

which implies $\mathcal{T}_{4^j}(4^j a) \in [0, 1/3]$.

Similarly, if $\theta_j = 1$, then Equation (14) implies

$$4^j a \in [4k + 1, 4k + 1 + 1/3] \quad \text{for some } k \in \mathbb{N} \text{ with } k \leq 4^{j-1}/3 \leq 4^j - 1,$$

which implies $\mathcal{T}_{4^j}(4^j a) \in [2/3, 1]$. So we finish the proof. \square

It is worth mentioning that the ‘‘sawtooth’’ function \mathcal{T}_m can be replaced by other functions that also have a key property ‘‘the function values near even integers are much larger than the ones near odd integers’’.

With Lemma B.2 in hand, we are ready to prove Proposition B.1.

Proof of Proposition B.1. By Lemma B.2, we have

$$\mathcal{T}_{4^j}(4^j a) \in [0, 1/3] \quad \text{if } \theta_j = 0 \quad \text{and} \quad \mathcal{T}_{4^j}(4^j a) \in [2/3, 1] \quad \text{if } \theta_j = 1.$$

Define $g(x) := 3\sigma(x - 1/3) - 3\sigma(x - 2/3)$ for any $x \in \mathbb{R}$, where σ is ReLU, i.e., $\sigma(x) = \max\{0, x\}$. See an illustration of g in Figure 12. Clearly,

$$g(x) = 0 \text{ if } x \leq 1/3 \quad \text{and} \quad g(x) = 1 \text{ if } x \geq 2/3.$$

Hence, $\phi := g \circ \mathcal{T}_{4^j}$ is the desired function. Obviously, $\phi(x) \in [0, 1]$ for any $x \in \mathbb{R}$. To verify Equation (8), we fix $j \in \{1, 2, \dots, J\}$. If $\theta_j = 0$, then $\mathcal{T}_{4^j}(4^j a) \in [0, 1/3]$, implying $\phi(4^j a) = g \circ \mathcal{T}_{4^j}(4^j a) = 0 = \theta_j$. If $\theta_j = 1$, then $\mathcal{T}_{4^j}(4^j a) \in [2/3, 1]$, implying $\phi(4^j a) = g \circ \mathcal{T}_{4^j}(4^j a) = 1 = \theta_j$.

It remains to show $\phi = g \circ \mathcal{T}_{4^j}$ can be realized by a ReLU network with the expected width and depth. Clearly, \mathcal{T}_{4^j} is a continuous piecewise linear function, which means it can be implemented by a one-hidden-layer ReLU network. To make our construction more efficient, we introduce another method to implement ϕ .




Figure 12. Illustrations of g and h on $[0, 1]$.

Define $h(x) := 1 - 2\sigma(x - 1/2) - 2\sigma(1/2 - x)$ for any $x \in [0, 1]$. See an illustration of h in Figure 12. Then, it is easy to verify that $\mathcal{T}_1(2x) = h(x)$ for any $x \in [0, 1]$ and h can be implemented by a one-hidden-layer ReLU network with width 2. For any $J \in \mathbb{N}^+$, it is easy to verify that

$$\mathcal{T}_{4^J}(2^{2J+1}x) = \underbrace{h \circ h \circ \dots \circ h}_{2J+1 \text{ times}}(x), \quad \text{for any } x \in [0, 1].$$

So, \mathcal{T}_{4^J} limited on $[0, 2 \times 4^J] = [0, 2^{2J+1}]$ can be realized by a ReLU network with width 2 and depth $2J + 1$. It follows that $\phi = g \circ \mathcal{T}_{4^J}$ can be implemented by a ReLU network with width 2 and depth $2J + 2$, which means we finish the proof. \square