Neural Network Architecture Beyond Width and Depth

Anonymous Author(s)
Affiliation
Address
email

Abstract

This paper proposes a new neural network architecture by introducing an additional dimension called height beyond width and depth. Neural network architectures with height, width, and depth as hyperparameters are called three-dimensional architectures. It is shown that neural networks with three-dimensional architectures are significantly more expressive than the ones with two-dimensional architectures (those with only width and depth as hyperparameters), e.g., standard fully connected networks. The new network architecture is constructed recursively via a nested structure, and hence we call a network with the new architecture nested network (NestNet). A NestNet of height s is built with each hidden neuron activated by a NestNet of height $\leq s-1$. When $s=1$, a NestNet degenerates to a standard network with a two-dimensional architecture. It is proved by construction that height-s ReLU NestNets with $\mathcal{O}(n)$ parameters can approximate Lipschitz continuous functions on $[0,1]^d$ with an error $\mathcal{O}(n^{-(s+1)/d})$, while the optimal approximation error of standard ReLU networks with $\mathcal{O}(n)$ parameters is $\mathcal{O}(n^{-2/d})$. Furthermore, such a result is extended to generic continuous functions on $[0,1]^d$ with the approximation error characterized by the modulus of continuity. Finally, a numerical example is provided to explore the advantages of the super approximation power of ReLU NestNets.

1 Introduction

In this paper, we design a new neural network architecture by introducing one more dimension, called height, in addition to width and depth in the characterization of dimensions of neural networks. We call neural network architectures with height, width, and depth as hyperparameters three-dimensional architectures. We will show that neural networks with three-dimensional architectures improve the approximation power significantly, compared to standard networks with two-dimensional architectures (those with only width and depth as hyperparameters). The approximation power of standard neural networks has been widely studied in recent years. The optimality of the approximation of standard fully-connected Rectified Linear Unit (ReLU) networks (e.g., see [33, 38, 43, 45]) implies limited room for further improvements. This motivates us to design a new neural network architecture by introducing an additional dimension of height beyond width and depth.

We will focus on ReLU ($\max\{x,0\}$) activation function and use it to demonstrate our ideas. Our new network architecture is constructed recursively via a nested structure, and hence we call a neural network with the new architecture nested network (**NestNet**). A NestNet of height s is built with each hidden neuron activated by a NestNet of height $\leq s-1$. In the case of $s=1$, a NestNet degenerates to a standard network with a two-dimensional architecture. Let us use an simple example to explain the height of a NestNet. We say a network is activated by $(\varrho_1,\cdots,\varrho_r)$ if each hidden neuron of this network is activated by one of $\varrho_1,\cdots,\varrho_r$. Here, $\varrho_1,\cdots,\varrho_r$ are trainable functions mapping $\mathbb{R}$ to $\mathbb{R}$.

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

Then a network of height $s \geq 2$ can be regarded as a $(\varrho_1, \cdots, \varrho_r)$-activated network, where $\varrho_1, \cdots, \varrho_r$ are (realized by) networks of height $\leq s-1$. See an example of a network of height 2 in Figure 1. The network therein can be regarded as a (ϱ_1, ϱ_2)-activated network, where ϱ_1 and ϱ_2 are (realized by) networks of height 1 (i.e., standard networks). The number of parameters in the network of Figure 1 is the sum of the numbers of parameters in $\mathcal{L}_0, \mathcal{L}_1, \mathcal{L}_2$ and ϱ_1, ϱ_2.

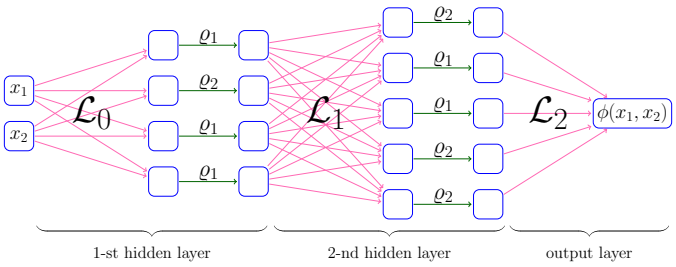

Figure 1: An example of a network of height 2, where ϱ_1 and ϱ_2 are (realized by) networks of height 1 (i.e., standard networks). Here, $\mathcal{L}_0$, $\mathcal{L}_1$ and $\mathcal{L}_2$ are affine linear maps.

Remark that a NestNet can be regarded as a sufficiently large standard network by expanding all of its sub-network activation functions. We propose the nested network architecture since it shares the parameters in the sub-network activation functions. This is the key reason why the NestNet has much better approximation property than the standard network. If we regard the network in Figure 1 as a NestNet of height, then the number of parameters is the sum of the numbers of parameters in $\mathcal{L}_0, \mathcal{L}_1, \mathcal{L}_2$ and ϱ_1, ϱ_2. However, if we expand the network in Figure 1 to a large standard network, then the number of parameters in ϱ_1 and ϱ_2 will be added many times for computing the total number of parameters.

Next, let us discuss our new network architecture from a perspective of hyperparameters. We call the network architecture with only width as a hyperparameter one-dimensional architecture. Its depth and height are both equal to one. Neural networks with this type of architecture are general called shallow networks. See an example in Figure 2(a). We call the network architectures with only width and depth as hyperparameters two-dimensional architecture. Its height is both equal to one. Networks with this type of architecture are general called deep networks. See an example in Figure 2(b). We call the network architectures with height, width, and depth as hyperparameters three-dimensional architecture. This is the architecture introduced in this paper. Networks with this type of architecture are called NestNet. See an example in Figure 2(c). One may refer to Table 1 for the approximation power of networks with these three types of architectures discussed above.

Table 1: Comparison for the approximation error of 1-Lipshitz continuous functions on $[0,1]^d$ approximated by ReLU NestNets and standard ReLU networks.

	dimension(s)	#parameters	approximation error	remark	reference
one-hidden-layer network	width varies (depth=height=1)	$\mathcal{O}(n)$	n^{-1} for $d=1$	linear combination	
deep network	width and depth vary (height=1)	$\mathcal{O}(n)$	$n^{-2/d}$	composition	[33, 38, 43, 45]
NestNet of height s	width, depth, and height vary	$\mathcal{O}(n)$	$n^{-(s+1)/d}$	nested composition	this paper

Our main contributions can be summarized as follows. We first propose a three-dimensional neural network architecture by introducing one more dimension called height beyond width and depth. We show that neural networks with three-dimensional architectures are significantly more expressive than the ones with two-dimensional architectures. In particular, we prove by construction that height-s ReLU NestNets with $\mathcal{O}(n)$ parameters can approximate Lipschitz continuous functions on $[0,1]^d$ with an error $\mathcal{O}(n^{-(s+1)/d})$, which is much better than the optimal error $\mathcal{O}(n^{-2/d})$ achieved by standard ReLU networks with $\mathcal{O}(n)$ parameters. In the case of $s+1 \geq d$, the approximation error is bounded by $\mathcal{O}(n^{-(s+1)/d}) \leq \mathcal{O}(n^{-1})$, which means we overcome the curse of dimensionality. Furthermore, we extend our result to generic continuous functions with the approximation error characterized by the modulus of continuity. See Theorem 2.1 and Corollary 2.2 for more details. Finally, we conduct a simple experiment to show the numerical advantages of the super approximation power of ReLU NestNets.

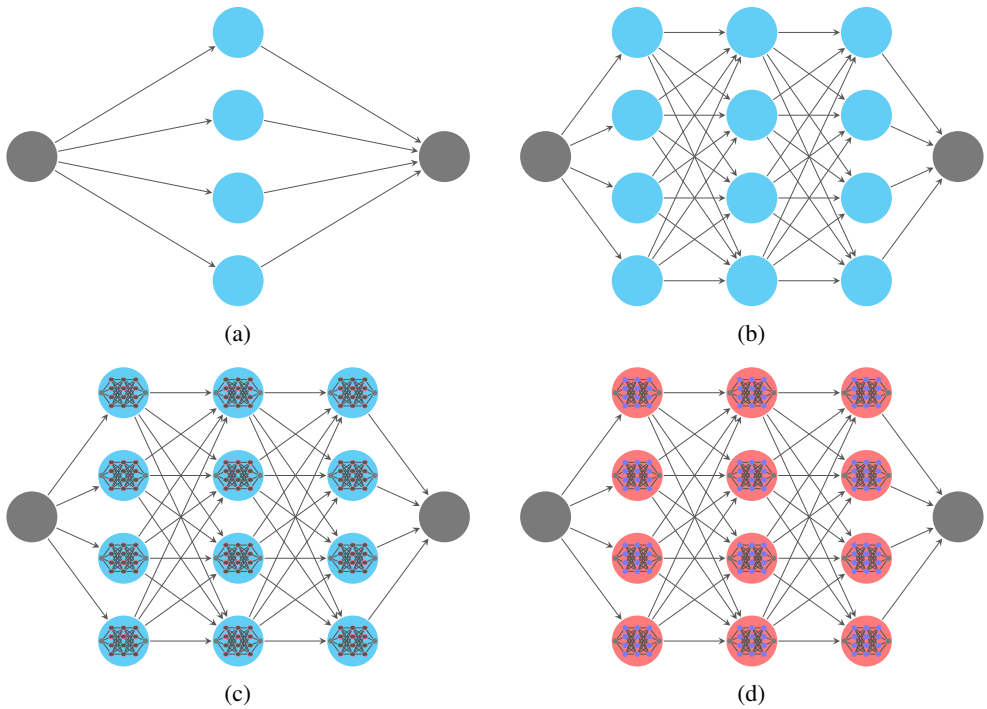

Figure 2: Illustrations of neural networks with one-, two-, and three-dimensional architecture. (a) One-dimensional case (width=4, depth=height=1). (b) Two-dimensional case (width=4, depth=3, height=1). (c) Three-dimensional case (width=4, depth=3, height=3). (d) Zoom-in of an activation function of the network in (c). The network in (d) can also be regarded as a network of height 2.

The rest of this paper is organized as follows. In Section 2, we first present our main results and then discuss related work. The proofs of the main results can be found 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 related work

In this section, we first present our main results and then discuss related work. The proofs of the main results are placed in the appendix.

2.1 Main results

In the rest of this paper, we will use $\mathcal{NN}_s\{n\}$ for $s \in \mathbb{N}^+$ and $n \in \mathbb{N}$ to denote the set of functions realized by height-s ReLU NestNets with as most n parameters. We will present the mathematical definition of $\mathcal{NN}_s\{n\}$. We first discuss some notation regarding affine linear maps. We use $\mathscr{L}$ to denote the set of all affine linear maps, i.e.,

$$\mathscr{L} \coloneqq \left\{\mathcal{L} : \mathcal{L}(\boldsymbol{x}) = \boldsymbol{W}\boldsymbol{x} + \boldsymbol{b},\ \boldsymbol{W} \in \mathbb{R}^{d_2 \times d_1},\ \boldsymbol{b} \in \mathbb{R}^{d_2},\ d_1, d_2 \in \mathbb{N}^+\right\}.$$

Let $\#\mathcal{L}$ denote the number of parameters in $\mathcal{L} \in \mathscr{L}$, i.e.,

$$\#\mathcal{L} = (d_1 + 1)d_2 \quad \text{if } \mathcal{L}(\boldsymbol{x}) = \boldsymbol{W}\boldsymbol{x} + \boldsymbol{b} \text{ for } \boldsymbol{W} \in \mathbb{R}^{d_2 \times d_1} \text{ and } \boldsymbol{b} \in \mathbb{R}^{d_2}.$$

We use $\vec{g} = (\varrho_1, \cdots, \varrho_k)$ to denote an activation function vector, where $\varrho_i : \mathbb{R} \to \mathbb{R}$ is an activation function for each $i \in \{1, \cdots, k\}$. When $\vec{g} = (\varrho_1, \cdots, \varrho_k)$ is applied to a vector input $\boldsymbol{x} = (x_1, \cdots, x_k)$, we have

$$\vec{g}(\boldsymbol{x}) = \Big(\varrho_1(x_1),\ \cdots,\ \varrho_k(x_k)\Big) \quad \text{for any } \boldsymbol{x} = (x_1, \cdots, x_k) \in \mathbb{R}^k.$$

Let $\text{set}(\vec{g})$ denote the function set containing all entries (functions) in $\vec{g}$. For example, if $\vec{g} = (g_1, g_2, g_3, g_2, g_1)$, then $\text{set}(\vec{g}) = \{g_1, g_2, g_3\}$.

Now we are ready to define $\mathcal{NN}_s\{n\}$ for $n, s \in \mathbb{N}$ recursively in $s \in \mathbb{N}$. In the degenerate case $n = 0$ or $s = 0$, we define

$$\mathcal{NN}_0\{n\} \coloneqq \Big\{\mathrm{id}_{\mathbb{R}},\ \mathrm{ReLU}\Big\} \eqqcolon \mathcal{NN}_s\{0\} \quad \text{for } n, s \in \mathbb{N},$$

where $\mathrm{id}_{\mathbb{R}} : \mathbb{R} \to \mathbb{R}$ is the identity map. That is, we regard the identity map and ReLU as height-s ReLU NestNets with 0 parameters or as height-0 ReLU NestNets with n parameters.

Next, let us present the recursive step. For $n, s \in \mathbb{N}^+$, a (vector-valued) function $\boldsymbol{\phi} \in \mathcal{NN}_s\{n\}$ has the following form:

$$\boldsymbol{\phi} = \mathcal{L}_m \circ \vec{g}_m \circ \cdots \circ \mathcal{L}_1 \circ \vec{g}_1 \circ \mathcal{L}_0, \tag{1}$$

where $\mathcal{L}_0, \cdots, \mathcal{L}_m \in \mathscr{L}$ are affine linear maps. Moreover, Equation (1) shall satisfy the following two conditions:

- Condition on activation functions:

$$\bigcup_{i=1}^{m} \mathrm{set}(\vec{g}_i) = \{\varrho_1, \cdots, \varrho_r\} \quad \text{and} \quad \varrho_j \in \bigcup_{i=0}^{s-1} \mathcal{NN}_i\{n_j\} \quad \text{for } j = 1, \cdots, r. \tag{2}$$

  Here, $\vec{g}_i$ is an activation function vector for each $i \in \{1, \cdots, m\}$. All entries in $\vec{g}_1, \cdots, \vec{g}_m$ form an activation function set $\{\varrho_1, \cdots, \varrho_r\}$, and ϱ_j for each $j \in \{1, \cdots, r\}$ can be realized by a height-i NestNet with $\le n_j$ parameters for some $i = i_j \le s - 1$. This condition means each hidden neuron is activated by a NestNet of lower height.

- Condition on the number of parameters:

$$\sum_{i=0}^{m} \#\mathcal{L}_i + \sum_{j=1}^{r} n_j \le n. \tag{3}$$

  This condition means the total number of parameters is no more than n. The total number of parameters is calculated by adding two parts. The first one is the number of parameters in affine linear maps $\mathcal{L}_0, \cdots, \mathcal{L}_m$. The other part is the number of parameters in the activation set $\{\varrho_1, \cdots, \varrho_r\}$ formed by the entries in activation function vectors $\vec{g}_1, \cdots, \vec{g}_m$.

Then, with two conditions in Equations (2) and (3), we can define $\mathcal{NN}_s\{n\}$ for $n, s \in \mathbb{N}^+$ as follows:

$$\mathcal{NN}_s\{n\} \coloneqq \Bigg\{\boldsymbol{\phi} : \boldsymbol{\phi} = \mathcal{L}_m \circ \vec{g}_m \circ \cdots \circ \mathcal{L}_1 \circ \vec{g}_1 \circ \mathcal{L}_0, \quad \mathcal{L}_0, \cdots, \mathcal{L}_r \in \mathscr{L}, \quad \bigcup_{i=1}^{m} \mathrm{set}(\vec{g}_i) = \{\varrho_1, \cdots, \varrho_r\}$$
$$\varrho_j \in \bigcup_{i=0}^{s-1} \mathcal{NN}_i\{n_j\} \text{ for } j = 1, \cdots, r, \quad \sum_{i=0}^{m} \#\mathcal{L}_i + \sum_{j=1}^{r} n_j \le n\Bigg\}.$$

Remark that, in the definition above, m can be equal to 0. In this case, the function $\boldsymbol{\phi}$ degenerates to an affine linear map.

In the NestNet example in Figure 1, the function ϕ therein is in $\bigcup_{n\in\mathbb{N}} \mathcal{NN}_2\{n\}$ and the activation function vectors $\vec{g}_1$ and $\vec{g}_2$ can be represented as

$$\vec{g}_1 = \big(\varrho_1,\ \varrho_2,\ \varrho_1,\ \varrho_1\big) \quad \text{and} \quad \vec{g}_2 = \big(\varrho_2,\ \varrho_1,\ \varrho_1,\ \varrho_2,\ \varrho_2\big).$$

Moreover, the activation function set containing all entries in $\vec{g}_1$ and $\vec{g}_2$ is a subset of $\bigcup_{n\in\mathbb{N}} \mathcal{NN}_1\{n\}$, i.e., $\{\varrho_1, \varrho_2\} \subseteq \bigcup_{n\in\mathbb{N}} \mathcal{NN}_1\{n\}$.

Let $C([0,1]^d)$ denote the set of continuous function on $[0,1]^d$. By convention, the modulus of continuity of a continuous function $f \in C([0,1]^d)$ is defined as

$$\omega_f(r) \coloneqq \sup\Big\{|f(\boldsymbol{x}) - f(\boldsymbol{y})| : \|\boldsymbol{x} - \boldsymbol{y}\|_2 \le r,\ \boldsymbol{x}, \boldsymbol{y} \in [0,1]^d\Big\} \quad \text{for any } r \ge 0.$$

Under these settings, we can find a function in $\mathcal{NN}_s\big\{\mathcal{O}(n)\big\}$ to approximate $f \in C([0,1]^d)$ with an approximation error $\mathcal{O}\big(\omega_f(n^{-(s+1)/d})\big)$, as shown in the main theorem below.

Theorem 2.1. *Given a continuous function $f \in C([0,1]^d)$, for any $n, s \in \mathbb{N}^+$ and $p \in [0, \infty]$, there exists $\phi \in \mathcal{NN}_s\big\{C_{s,d}(n+1)\big\}$ such that*

$$\big\|\phi(\boldsymbol{x}) - f(\boldsymbol{x})\big\|_{L^p([0,1]^d)} \le 7\sqrt{d}\,\omega_f\big(n^{-(s+1)/d}\big),$$

where $C_{s,d} = 10^3 d^2 (s+7)^2$ if $p \in [1, \infty)$ and $C_{s,d} = 10^{d+3} d^2 (s+7)^2$ if $p = \infty$.

Remark that the constant $C_{s,d}$ in Theorem 2.1 is valid for all $n \in \mathbb{N}^+$. As we shall see later, $C_{s,d}$ can be greatly reduced if one only cares about large $n \in \mathbb{N}^+$. Generally, it is challenging to simplify the approximation error in Theorem 2.1 to make it explicitly depend on n due to the complexity of $\omega_f(\cdot)$. However, the approximation error can be simplified to an explicit one depending on n in the case of special target function spaces like Hölder continuous function space. To be exact, if f is a Hölder continuous function on $[0,1]^d$ of order $\alpha \in (0,1]$ with a Hölder constant $\lambda > 0$, then

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

implying $\omega_f(r) \le \lambda r^\alpha$ for any $r \ge 0$. This means we can get an exponentially small approximation error $7\lambda\sqrt{d}\, n^{-(s+1)\alpha/d}$ as shown in Corollary 2.2 below.

Corollary 2.2. *Suppose f is a Hölder continuous function on $[0,1]^d$ of order $\alpha \in (0,1]$ with a Hölder constant $\lambda > 0$. For any $n, s \in \mathbb{N}^+$ and $p \in [0,\infty]$, there exists $\phi \in \mathcal{NN}_s\big\{C_{s,d}(n+1)\big\}$ such that*

$$\big\|\phi(\boldsymbol{x}) - f(\boldsymbol{x})\big\|_{L^p([0,1]^d)} \le 7\lambda\sqrt{d}\, n^{-(s+1)\alpha/d},$$

where $C_{s,d} = 10^3 d^2 (s+7)^2$ if $p \in [1,\infty)$ and $C_{s,d} = 10^{d+3} d^2 (s+7)^2$ if $p = \infty$.

In Corollary 2.2, if $\alpha = 1$, i.e., f is a Lipschitz continuous function with a Lipschitz constant $\lambda > 0$, then the approximation error can be further simplified to $7\lambda\sqrt{d}\, n^{-(s+1)/d}$. See Table 1 for the comparison of the approximation error of 1-Lipshitz continuous functions on $[0,1]^d$ approximated by ReLU NestNets and standard ReLU networks.

2.2 Related work

We will connect our results to related existing ones for a deeper understanding. We first compare our results to existing ones from an approximation perspective. Next, we discuss the connection between our NestNet architecture and existing trainable activation functions.

Discussion from an approximation perspective

The study of the approximation power of deep neural networks has become an active topic in recent years. This topic has been extensively studied from many perspectives, e.g., in terms of combinatorics [27], topology [7], information theory [29], fat-shattering dimension [1, 21], Vapnik-Chervonenkis (VC) dimension [6, 14, 31], classical approximation theory [3, 4, 8, 9, 10, 11, 12, 13, 18, 23, 24, 25, 28, 32, 33, 34, 37, 40, 42, 43, 45, 46], etc. To the best of our knowledge, the study of neural network approximation has two main stages: shallow (one-hidden-layer) networks and deep networks.

In the early works of neural network approximation, the approximation power of shallow networks is investigated. In particular, the universal approximation theorem [11, 17, 18], without approximation error estimate, showed that a sufficiently large neural network can approximate a target function in a certain function space arbitrarily well. For one-hidden-layer neural networks of width n and sufficiently smooth functions, an asymptotic approximation error $\mathcal{O}(n^{-1/2})$ in the L^2-norm is proved in [4, 5], leveraging an idea that is similar to Monte Carlo sampling for high-dimensional integrals.

Recently, a large number of works focus on the study of deep neural networks. It is shown in [33, 43, 45] that the optimal approximation error is $\mathcal{O}(n^{-2/d})$ by using ReLU networks with n parameters to approximate Lipschitz functions on $[0,1]^d$. This optimal approximation error follows a natural question: How can we a better approximation error? Generally, there are two ideas to get better errors. The first one is to consider smaller function spaces, e.g., smooth functions [24, 44] and band-limited functions [26]. The other one is to introducing new networks, e.g., Floor-ReLU networks [35], Floor-Exponential-Step (FLES) networks [36], and (Sin, ReLU, 2^x)-activated networks [20].

This paper proposes a three-dimensional neural network architecture by introducing one more dimension called height beyond width and depth. We show that neural networks with three-dimensional architectures are significantly more expressive than the ones with two-dimensional architectures. To be exact, we prove by construction that height-s ReLU NestNets with $\mathcal{O}(n)$ parameters can approximate Lipschitz continuous functions on $[0,1]^d$ with an error $\mathcal{O}(n^{-(s+1)/d})$, which is better than the optimal error $\mathcal{O}(n^{-2/d})$ achieved by standard ReLU networks with $\mathcal{O}(n)$ parameters. Such a result can be generalized to generic continuous functions. As shown in Theorem 2.1, the approximation

error is $\mathcal{O}\big(\omega_f(n^{-(s+1)/d})\big)$ if the target function f is a continuous function on $[0,1]^d$. Therefore, we overcome the curse of dimensionality if $s+1 \geq d$ and the variation of $\omega_f(r)$ as $r \to 0$ is moderate (e.g., $\omega_f(r) \lesssim r^\alpha$ for Hölder continuous functions).

Connection to trainable activation functions

The key idea of trainable activation functions is to add a small number of trainable parameters to existing activation functions. Let us present several existing trainable activation functions as follows. An ReLU-like function is introduced in [15] by modifying the negative part of ReLU using a trainable parameter α, i.e., the Parametric ReLU (PReLU) is defined as

$$\mathrm{PReLU}(x) \coloneqq \begin{cases} x & \text{if } x \geq 0 \\ \alpha x & \text{if } x < 0. \end{cases}$$

A variant of ELU unit is introduce in [41] by adding two trainable parameters $\beta, \gamma > 0$, i.e., the Parametric ELU (PELU) is given by

$$\mathrm{PELU}(x) \coloneqq \begin{cases} \frac{\beta}{\gamma} & \text{if } x \geq 0 \\ \beta\big(\exp(\frac{x}{\gamma}) - 1\big)x & \text{if } x < 0. \end{cases}$$

Authors in [30] propose a type of Flexible ReLU (FReLU), which is defined via

$$\mathrm{FReLU}(x) \coloneqq \mathrm{ReLU}(x+\alpha) + \beta,$$

where α and β are two trainable parameters. One may refer to [2] for a survey of modern trainable activation functions. To the best of our knowledge, most existing trainable activation functions can be regarded as **parametric** variants of the original activation functions. That is, they are attained via parameterizing the original activation functions with a small number of (typically 1 or 2) trainable parameters.

By contrast, activation functions in our NestNets are much more flexible. They can be (realized by) either complicated or simple sub-NestNets. That is, we can freely determine the number of parameters in the activation functions of NestNets. In other words, in NestNets, we can randomly distribute the parameters in the affine linear maps and activation functions. In short, compared to the networks with existing trainable activation functions, our NestNets are much flexible and have much more freedom on the choice of activation functions.

3 Experiments

In this section we will conduct a simple experiment to explore the numerical advantages of the super approximation power of ReLU NestNets. To this end, we first discuss the experiment setup in Section 3.1 and then present the experiment results in Section 3.2.

3.1 Experiment setup

We will design a binary classification experiment with sufficiently many data samples. The idea comes from the Archimedean spiral, which can be described by the equation $r = a + b\theta$ in polar coordinates (r, θ), where a and b are given real numbers. In other words, we will design two sets $\mathcal{S}_0$ and $\mathcal{S}_1$ in $[0,1]^2$ based on the Archimedean spiral, where the label for $\mathcal{S}_i$ is i for $i = 0, 1$.

Let us first define two curves (Archimedean spirals) as follows:

$$\widetilde{\mathcal{C}}_i \coloneqq \Big\{(x,y) : x = r_i \cos\theta,\ y = r_i \sin\theta,\ r_i = a_i + b_i\theta,\ \text{for } \theta \in [0, s\pi]\Big\},$$

for $i = 0, 1$, where $a_0 = 0$, $a_1 = 1$, $b_0 = b_1 = \frac{1}{\pi}$, and $s = 20$. To simplify the discussion below, we normalize $\widetilde{\mathcal{C}}_i$ as $\mathcal{C}_i \subseteq [0,1]^2$, where $\mathcal{C}_i$ is defined by

$$\mathcal{C}_i \coloneqq \Big\{(x,y) : x = \tfrac{\widetilde{x}}{2(s+2)} + \tfrac{1}{2},\ y = \tfrac{\widetilde{y}}{2(s+2)} + \tfrac{1}{2},\ (\widetilde{x}, \widetilde{y}) \in \widetilde{\mathcal{C}}_i\Big\},$$

for $i = 0, 1$. Then, we can define the two desired sets as follows:

$$\mathcal{S}_i \coloneqq \Big\{(u,v) : \sqrt{(u-x)^2 + (v-y)^2} \leq \varepsilon,\ (x,y) \in \mathcal{C}_i\Big\},$$

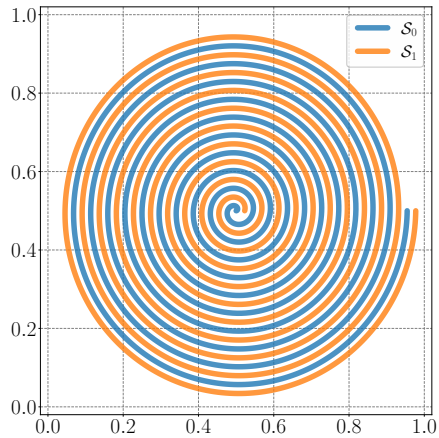

Figure 3: An illustration for $\mathcal{S}_0$ and $\mathcal{S}_1$.

for $i = 0, 1$, where $\varepsilon = 0.008$ in our experiments. See an illustration for $\mathcal{S}_0$ and $\mathcal{S}_1$ in Figure 3.

The goal of our experiments is to train a network to approximate the binary classification function f well, where f is defined by

$$f(\boldsymbol{x}) \coloneqq \begin{cases} 0 & \text{if } \boldsymbol{x} \in \mathcal{S}_0 \\ 1 & \text{if } \boldsymbol{x} \in \mathcal{S}_1. \end{cases}$$

It is easy to verify that $\mathcal{S}_0$ and $\mathcal{S}_1$ are closed subsets of $[0,1]^2$. Then f can be continuously extended to $[0,1]^2$. Define

$$\widetilde{f}(\boldsymbol{x}) \coloneqq \frac{\operatorname{dist}(\boldsymbol{x}, \mathcal{S}_0)}{\operatorname{dist}(\boldsymbol{x}, \mathcal{S}_0) + \operatorname{dist}(\boldsymbol{x}, \mathcal{S}_1)} \quad \text{for any } \boldsymbol{x} \in [0,1]^2,$$

where

$$\operatorname{dist}(\boldsymbol{x}, \mathcal{S}_i) \coloneqq \inf_{\boldsymbol{y} \in \mathcal{S}_i} \|\boldsymbol{x} - \boldsymbol{y}\|_2 \quad \text{for any } \boldsymbol{x} \in [0,1]^2 \text{ and } i = 0, 1.$$

It is easy to verify that $\widetilde{f}$ is continuous on $[0,1]^d$ and

$$\widetilde{f}(\boldsymbol{x}) = f(\boldsymbol{x}) \quad \text{for any } \boldsymbol{x} \in \mathcal{S}_0 \cup \mathcal{S}_1.$$

That means $\widetilde{f}$ is a continuous extension of f. Moreover, the modulus of continuity of the extension $\widetilde{f}$ can be easily estimated based on

$$\operatorname{dist}(\mathcal{S}_0, \mathcal{S}_1) \coloneqq \inf_{\boldsymbol{x} \in \mathcal{S}_0,\, \boldsymbol{y} \in \mathcal{S}_1} \|\boldsymbol{x} - \boldsymbol{y}\|_2 > 0.$$

Next, let us discuss the network architecture in out experiments. We consider two type of networks: NestNets and standard networks. We adopt four-hidden-layer fully connected network architecture of width 20, 35, or 50. As for the activation function, we adopt σ for the standard network and ϱ for the NesNet, where σ is ReLU ($\max\{x, 0\}$) and ϱ is realized by a trainable one-hidden-layer ReLU network of width 3. To be exact, ϱ is given by

$$\varrho(x) = \boldsymbol{w}_1^T \cdot (x\boldsymbol{w}_0 + \boldsymbol{b}_0) + b_1 \quad \text{for any } x \in \mathbb{R},$$

where $\boldsymbol{w}_0, \boldsymbol{w}_1, \boldsymbol{b}_0 \in \mathbb{R}^3$ and $b_1 \in \mathbb{R}$ are trainable parameters. There are 10 parameters in ϱ. The initial settings for ϱ in our experiments are $\boldsymbol{w}_0 = (1,1,1)$, $\boldsymbol{w}_1 = (1,1,-1)$, $\boldsymbol{b}_0 = (-0.2, -0.1, 0.0)$, and $b_1 = 0$. See Figure 6 for illustrations. To reduce overfitting and speed up optimization, we take two common regularization methods: dropout [16, 39] and batch normalization [19]. The full architecture is shown in Figure 4.

Finally, we are ready to present our experiment strategy. For each $i \in \{0, 1\}$, we randomly choose 3×10^5 training samples and 3×10^4 test samples in $\mathcal{S}_i$ with label i. Then, we use these 6×10^5 training samples to train our network and use these 6×10^4 test samples to compute the test accuracy. We use the cross-entropy loss function to evaluate the loss between the network and the binary classification function f. The number of epochs and the batch size are set to 80 and 64, respectively. We adopt Adam [22] as the optimization method. For the i-th epoch, the learning rate for the parameters in ϱ is $0.1 \times 0.003 \times 0.85^{i-1}$ and the learning rate for other parameters is $0.003 \times 0.85^{i-1}$. See Tables 2 for more details.

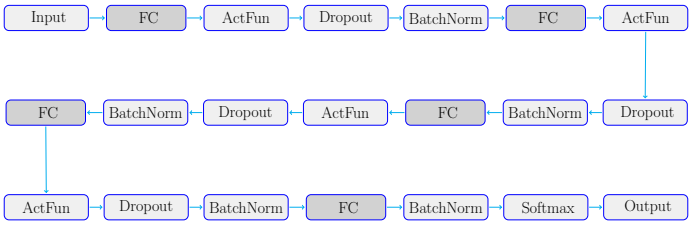

Figure 4: An illustration of the network architecture in the experiments. FC and ActFun are short of fully connected layer and activation function, respectively. ActFun=ReLU (σ) for standard networks; ActFun=sub-network (ϱ) for NestNets.

Table 2: Hyperparameters in the experiments.

	standard network	NestNet	
loss function	cross-entropy loss		
number of training samples	6×10^5 in total, with half in $\mathcal{S}_i$ for $i = 0, 1$		
number of test samples	6×10^4 in total, with half in $\mathcal{S}_i$ for $i = 0, 1$		
number of training epochs	80		
training batch size	256		
optimizer	Adam [22]		
learning rate in i-th epoch	$0.003 \times 0.9^{i-1}$	parameters in ϱ	others
		$0.1 \times 0.003 \times 0.9^{i-1}$	$0.003 \times 0.9^{i-1}$

3.2 Experiment results

We will compare the test accuracies for NestNets and standard networks. We adopt the average of largest 20 test accuracies over all 80 epochs as the final test accuracy. As we can see from Table 3 and Figure 5, via adding 10 more parameters (stored in ϱ), NestNets achieve much better test accuracies than standard networks. In an extreme comparison, the test accuracy attained by the NestNet with 1.4×10^3 parameters is still better than that of the standard network with 7.9×10^3 parameters. This numerically verifies that the NestNet has much better approximation power than the standard network.

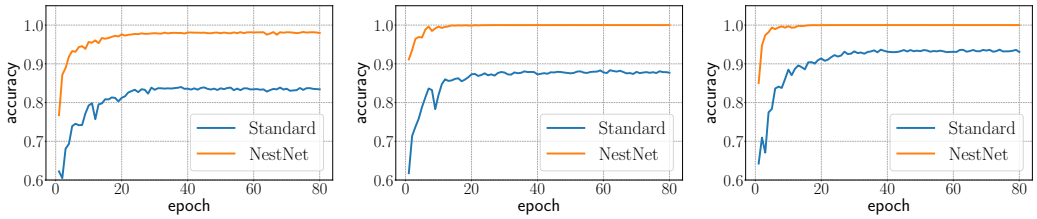

(a) (b) (c)

Figure 5: Test accuracy over epochs. (a) Width=20. (b) Width=35. (c) Width=50.

Table 3: Error comparison of standard networks and NestNets.

	width	depth	activation	#parameters	test accuracy
standard network	20	4	ReLU (σ)	1.4×10^3	0.837581
NestNet			sub-network (ϱ)		0.981322
standard network	35	4	ReLU (σ)	4.0×10^3	0.880078
NestNet			sub-network (ϱ)		0.999985
standard network	50	4	ReLU (σ)	7.9×10^3	0.935018
NestNet			sub-network (ϱ)		1.000000

Next, let us discuss the trainable sub-network activation function ϱ. Recall that ϱ is given by

$$\varrho(x) = \boldsymbol{w}_1^T \cdot (x\boldsymbol{w}_0 + \boldsymbol{b}_0) + b_1 \quad \text{for any } x \in \mathbb{R},$$

where $\boldsymbol{w}_0, \boldsymbol{w}_1, \boldsymbol{b}_0 \in \mathbb{R}^3$ and $b_1 \in \mathbb{R}$ are trainable parameters. There are 10 parameters in ϱ. We will compare the initial guess and the trained solution for the parameters in ϱ. Our initial settings for ϱ are to make it be a variant of ReLU except for a small neighborhood near 0. As we can see from Table 4, the parameters in ϱ only change a little during the training process. However, the approximation error comparison in Table 3 implies that these little changes of the parameters in ϱ make a big difference. We visualize the difference between the initial ϱ and the trained ϱ in Figure 6.

Table 4: Parameters in ϱ.

	width	$\boldsymbol{w}_0$	$\boldsymbol{b}_0$	$\boldsymbol{w}_1$	b_1
initial		$(1, 1, 1)$	$(-0.2, -0.1, 0.0)$	$(1, 1, -1)$	0
trained	20	$(0.9847, 0.9675, 1.0104)$	$(-0.2223, -0.1620, 0.0607)$	$(0.9902, 0.9834, -1.0138)$	0.0033
	35	$(0.9809, 0.9634, 1.0140)$	$(-0.2265, -0.1675, 0.0667)$	$(0.9838, 0.9813, -1.0173)$	0.0020
	50	$(0.9743, 0.9594, 1.0222)$	$(-0.2303, -0.1634, 0.0535)$	$(0.9758, 0.9726, -1.0239)$	−0.0154

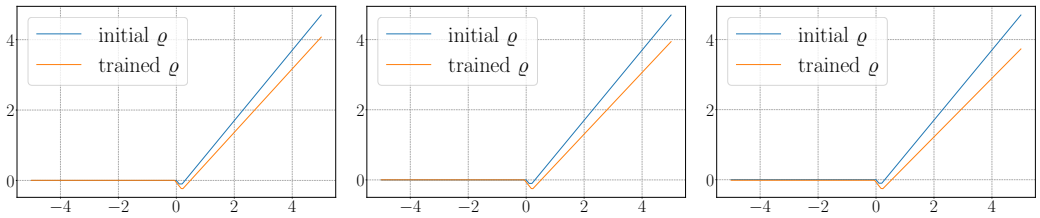

(a) (b) (c)

Figure 6: Illustrations of initial ϱ and trained ϱ. (a) Width=20. (b) Width=35. (c) Width=50.

We use a simple example above to show the numerical advantages of the super approximation power of NestNets, which is regarded as a proof of possibility. It would be of great interest to further explore the numerical performance of NestNets to bridge our theoretical result to applications. We believe that NestNets can be further developed and applied to real-world applications.

4 Conclusion

This paper proposes a three-dimensional neural network architecture by introducing one more dimension called height beyond width and depth. We show that neural networks with three-dimensional architectures are significantly more expressive than the ones with two-dimensional architectures. In particular, we prove by construction that height-s ReLU NestNets with $\mathcal{O}(n)$ parameters can approximate Lipschitz continuous functions on $[0, 1]^d$ with an error $\mathcal{O}(n^{-(s+1)/d})$, which is much better than the optimal error $\mathcal{O}(n^{-2/d})$ achieved by standard ReLU networks with $\mathcal{O}(n)$ parameters. Furthermore, we extend our result to generic continuous functions. Finally, we conduct a simple experiment to show the numerical advantages of the super approximation power of ReLU NestNets.

We would like to remark that our analysis was limited on the ReLU activation function and the (Hölder) continuous function space. It would be interesting to extend our conclusions to other activation functions (e.g., tanh and sigmoid functions) and other function spaces (e.g, Lebesgue and Sobolev spaces). Besides, it would be interesting to explore the numerical performance of NestNets and apply it to real-world applications.

References

[1] Martin Anthony and Peter L. Bartlett. *Neural Network Learning: Theoretical Foundations*. Cambridge University Press, New York, NY, USA, 1st edition, 2009.

[2] Andrea Apicella, Francesco Donnarumma, Francesco Isgrò, and Roberto Prevete. A survey on modern trainable activation functions. *Neural Networks*, 138:14–32, 2021.

[3] Chenglong Bao, Qianxiao Li, Zuowei Shen, Cheng Tai, Lei Wu, and Xueshuang Xiang. Approximation analysis of convolutional neural networks. *Semantic Scholar e-Preprint*, page Corpus ID: 204762668, 2019.

[4] Andrew R. Barron. Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information Theory*, 39(3):930–945, May 1993.

[5] Andrew R. Barron and Jason M. Klusowski. Approximation and Estimation for High-Dimensional Deep Learning Networks, September 2018.

[6] Peter Bartlett, Vitaly Maiorov, and Ron Meir. Almost linear VC-dimension bounds for piecewise polynomial networks. *Neural Computation*, 10(8):2159–2173, 1998.

[7] Monica Bianchini and Franco Scarselli. 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.

[8] Helmut. Bölcskei, Philipp. Grohs, Gitta. Kutyniok, and Philipp. Petersen. Optimal approximation with sparsely connected deep neural networks. *SIAM Journal on Mathematics of Data Science*, 1(1):8–45, 2019.

[9] Minshuo Chen, Haoming Jiang, Wenjing Liao, and Tuo Zhao. Efficient approximation of deep relu networks for functions on low dimensional manifolds. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.

[10] Charles K. Chui, Shao-Bo Lin, and Ding-Xuan Zhou. Construction of neural networks for realization of localized deep learning. *Frontiers in Applied Mathematics and Statistics*, 4:14, 2018.

[11] George Cybenko. Approximation by superpositions of a sigmoidal function. *MCSS*, 2:303–314, 1989.

[12] Rémi Gribonval, Gitta Kutyniok, Morten Nielsen, and Felix Voigtlaender. Approximation spaces of deep neural networks. *arXiv e-prints*, page arXiv:1905.01208, May 2019.

[13] Ingo Gühring, Gitta Kutyniok, and Philipp Petersen. Error bounds for approximations with deep ReLU neural networks in $W^{s,p}$ norms. *arXiv e-prints*, page arXiv:1902.07896, Feb 2019.

[14] Nick Harvey, Christopher Liaw, and Abbas Mehrabian. Nearly-tight VC-dimension bounds for piecewise linear neural networks. In Satyen Kale and Ohad Shamir, editors, *Proceedings of the 2017 Conference on Learning Theory*, volume 65 of *Proceedings of Machine Learning Research*, pages 1064–1068, Amsterdam, Netherlands, 07–10 Jul 2017. PMLR.

[15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 1026–1034, 2015.

[16] Geoffrey E. Hinton, Nitish Srivastava, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Improving neural networks by preventing co-adaptation of feature detectors. *CoRR*, abs/1207.0580, 2012.

[17] Kurt Hornik. Approximation capabilities of multilayer feedforward networks. *Neural Networks*, 4(2):251–257, 1991.

[18] Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366, 1989.

[19] Sergey Ioffe and Christian Szegedy. 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, page 448–456. JMLR.org, 2015.

[20] Yuling Jiao, Yanming Lai, Xiliang Lu, Fengru Wang, Jerry Zhijian Yang, and Yuanyuan Yang. Deep neural networks with ReLU-Sine-Exponential activations break curse of dimensionality on hölder class. *arXiv e-prints*, page arXiv:2103.00542, February 2021.

[21] Michael J. Kearns and Robert E. Schapire. Efficient distribution-free learning of probabilistic concepts. *J. Comput. Syst. Sci.*, 48(3):464–497, June 1994.

[22] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.

[23] Qianxiao Li, Ting Lin, and Zuowei Shen. Deep learning via dynamical systems: An approximation perspective. *arXiv e-prints*, page arXiv:1912.10382, December 2019.

[24] Jianfeng Lu, Zuowei Shen, Haizhao Yang, and Shijun Zhang. Deep network approximation for smooth functions. *SIAM Journal on Mathematical Analysis*, 53(5):5465–5506, 2021.

[25] Hadrien Montanelli and Haizhao Yang. Error bounds for deep ReLU networks using the Kolmogorov-Arnold superposition theorem. *Neural Networks*, 129:1–6, 2020.

[26] Hadrien Montanelli, Haizhao Yang, and Qiang Du. Deep ReLU networks overcome the curse of dimensionality for bandlimited functions. *Journal of Computational Mathematics*, 2020.

[27] Guido F Montufar, Razvan Pascanu, Kyunghyun Cho, and Yoshua Bengio. On the number of linear regions of deep neural networks. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2924–2932. Curran Associates, Inc., 2014.

[28] Ryumei Nakada and Masaaki Imaizumi. Adaptive approximation and generalization of deep neural network with intrinsic dimensionality. *Journal of Machine Learning Research*, 21(174):1–38, 2020.

[29] Philipp Petersen and Felix Voigtlaender. Optimal approximation of piecewise smooth functions using deep ReLU neural networks. *Neural Networks*, 108:296–330, 2018.

[30] Suo Qiu, Xiangmin Xu, and Bolun Cai. Frelu: Flexible rectified linear units for improving convolutional neural networks. In *2018 24th International Conference on Pattern Recognition (ICPR)*, pages 1223–1228, Los Alamitos, CA, USA, aug 2018. IEEE Computer Society.

[31] Akito Sakurai. Tight bounds for the VC-dimension of piecewise polynomial networks. In *Advances in Neural Information Processing Systems*, pages 323–329. Neural information processing systems foundation, 1999.

[32] Zuowei Shen, Haizhao Yang, and Shijun Zhang. Nonlinear approximation via compositions. *Neural Networks*, 119:74–84, 2019.

[33] Zuowei Shen, Haizhao Yang, and Shijun Zhang. Deep network approximation characterized by number of neurons. *Communications in Computational Physics*, 28(5):1768–1811, 2020.

[34] Zuowei Shen, Haizhao Yang, and Shijun Zhang. Deep network approximation: Achieving arbitrary accuracy with fixed number of neurons. *arXiv e-prints*, page arXiv:2107.02397, July 2021.

[35] Zuowei Shen, Haizhao Yang, and Shijun Zhang. Deep network with approximation error being reciprocal of width to power of square root of depth. *Neural Computation*, 33(4):1005–1036, 03 2021.

[36] Zuowei Shen, Haizhao Yang, and Shijun Zhang. Neural network approximation: Three hidden layers are enough. *Neural Networks*, 141:160–173, 2021.

[37] Zuowei Shen, Haizhao Yang, and Shijun Zhang. ReLU network approximation in terms of intrinsic parameters. *arXiv e-prints*, page arXiv:2111.07964, November 2021.

[38] Zuowei Shen, Haizhao Yang, and Shijun Zhang. Optimal approximation rate of ReLU networks in terms of width and depth. *Journal de Mathématiques Pures et Appliquées*, 157:101–135, 2022.

[39] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(56):1929–1958, 2014.

[40] Taiji Suzuki. 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.

[41] Ludovic Trottier, Philippe Giguère, and Brahim Chaib-draa. Parametric exponential linear unit for deep convolutional neural networks. *2017 16th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 207–214, 2017.

[42] Dmitry Yarotsky. Error bounds for approximations with deep ReLU networks. *Neural Networks*, 94:103–114, 2017.

[43] Dmitry Yarotsky. Optimal approximation of continuous functions by very deep ReLU networks. In Sébastien Bubeck, Vianney Perchet, and Philippe Rigollet, editors, *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pages 639–649. PMLR, 06–09 Jul 2018.

[44] Dmitry Yarotsky and Anton Zhevnerchuk. The phase diagram of approximation rates for deep neural networks. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 13005–13015. Curran Associates, Inc., 2020.

[45] Shijun Zhang. Deep neural network approximation via function compositions. *PhD Thesis, National University of Singapore*, 2020. URL: https://scholarbank.nus.edu.sg/handle/10635/186064.

[46] Ding-Xuan Zhou. Universality of deep convolutional neural networks. *Applied and Computational Harmonic Analysis*, 48(2):787–794, 2020.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See Sections 1 and 4.
   (c) Did you discuss any potential negative societal impacts of your work? [N/A] Our work is theoretical and we do not see any potential negative societal impacts.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes] See the appendix.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [N/A]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [N/A]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]