

DEEP RELU NETWORKS OVERCOME THE CURSE OF DIMENSIONALITY FOR GENERALIZED BANDLIMITED FUNCTIONS*

Hadrien Montanelli

Centre de Mathématiques Appliquées, École Polytechnique, Palaiseau, France

Email: hadrien.montanelli@polytechnique.edu

Haizhao Yang¹⁾

Department of Mathematics, Purdue University, Indiana, United States

Email: haizhao@purdue.edu

Qiang Du

Department of Applied Physics and Applied Mathematics, Columbia University,

New York, United States

Email: qd2125@columbia.edu

Abstract

We prove a theorem concerning the approximation of generalized bandlimited multivariate functions by deep ReLU networks for which the curse of the dimensionality is overcome. Our theorem is based on a result by Maurey and on the ability of deep ReLU networks to approximate Chebyshev polynomials and analytic functions efficiently.

Mathematics subject classification: 68T01, 33F05, 41A10.

Key words: Machine learning, Deep ReLU networks, Curse of dimensionality, Approximation theory, Bandlimited functions, Chebyshev polynomials.

1. Introduction

The curse of dimensionality is a inevitable issue in high-dimensional scientific computing. Standard numerical algorithms whose cost is exponential in the dimension d are prohibitive when d is large. As a mesh-free function parametrization tool, neural networks are believed to be a suitable approach to conquer the curse of dimensionality. In this paper, we show that deep ReLU networks overcome the curse of dimensionality for *generalized bandlimited functions*, which we shall define at the end of the introduction. Let us first quickly review what networks are.

Shallow networks are approximations \tilde{f}_W of multivariate functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$ of the form

$$\tilde{f}_W(\mathbf{x}) = \sum_{i=1}^W \alpha_i \sigma(\mathbf{w}_i \cdot \mathbf{x} + \theta_i), \quad (1.1)$$

for some *activation function* $\sigma : \mathbb{R} \rightarrow \mathbb{R}$, weights $\alpha_i, \theta_i \in \mathbb{R}$, $\mathbf{w}_i \in \mathbb{R}^d$ and integer $W \geq 1$. Each operation $\sigma(\mathbf{w}_i \cdot \mathbf{x} + \theta_i)$ is called a *unit* and the W units in (1.1) form a *hidden layer*; this is a special form of nonlinear approximation [1, 2]. *Deep networks* are compositions of shallow networks and have several hidden layers, and each unit of each layer performs an operation

* Received October 25, 2019 / Revised version received May 25, 2020 / Accepted July 15, 2020 /
Published online October 12, 2021 /

¹⁾ Corresponding author

of the form $\sigma(\mathbf{w} \cdot \mathbf{x} + \theta)$. Following Yarotsky [3], we allow connections between units in non-neighboring layers. We define the *depth* L of a network as the number of hidden layers and the *size* W as the total number of units. In practice, networks with depth $L = \mathcal{O}(1)$ are considered shallow, while deep networks have typically $L \gg 1$ layers.

Before the revolution of deep learning [4], most research concerned shallow networks with sigmoid activation functions. Nowadays, networks using the *REctifier Linear Unit (ReLU)* activation function $\sigma(x) = \max(0, x)$ have become the most popular tool, partly because sigmoid activation functions lead to severe gradient degeneracy during the optimization process. It was also shown in [5] that deep ReLU networks produce sparsity that helps a wide range of machine learning applications; smooth activation functions, including smoothed ReLU functions, do not. This is why we focus on ReLU networks in this paper.

The theory of approximating functions using shallow networks goes back to 1989 when Cybenko showed that any continuous functions can be approximated by shallow networks [6], while Hornik, Stinchcombe and White proved a similar result for Borel measurable functions [7]. In the 1990s, the attention shifted to the *approximation power*¹⁾ of shallow networks [8–11]. Of particular interest was the absence of the curse of dimensionality in the approximation of functions with fast decaying Fourier coefficients [12].

Fast forward to the 2010s and the success of deep networks, one of the most important theoretical problems is to determine why and when deep networks can lessen or break the curse of dimensionality, especially for ReLU networks. One may focus on a particular set of functions which have a very special structure (such as compositional or polynomial), and show that for this particular set deep networks overcome the curse of dimensionality [13–21]. Alternatively, one may consider a function space that is more generic for multivariate approximation in high dimensions, such as Korobov spaces [22], and prove convergence results for which the curse of dimensionality is lessened [23].

In this paper, we may consider *generalized bandlimited* functions $f : B = [0, 1]^d \rightarrow \mathbb{R}$ of the form

$$f(\mathbf{x}) = \int_{\mathbb{R}^d} F(\mathbf{w})K(\mathbf{w} \cdot \mathbf{x})d\mathbf{w}, \quad \text{supp } F \subset [-M, M]^d, \quad M \geq 1, \quad (1.2)$$

for some square-integrable function $F : [-M, M]^d \rightarrow \mathbb{C}$ and analytic kernel $K : \mathbb{R} \rightarrow \mathbb{C}$. This class of functions contains several examples of Reproducing Kernel Hilbert Spaces (RKHSs), including the space of bandlimited functions. The latter are ubiquitous in science and engineering. In information theory, bandlimited signals are often used for analysis and representation after sampling. In scientific computing, after discretization, functions are bandlimited by the Nyquist–Shannon sampling theorem. Studying the approximation power of ReLU networks for bandlimited functions is particularly important for neural network-based scientific computing in high dimensions. In Section 3, we shall show that for any measure μ such functions can be approximated to accuracy ϵ in the $L^2(B, \mu)$ -norm by deep ReLU networks of depth $L = \mathcal{O}(\log_2^2 \frac{1}{\epsilon})$ and size $W = \mathcal{O}(\frac{1}{\epsilon^2} \log_2^2 \frac{1}{\epsilon})$.

We review some properties of deep ReLU networks in Section 2, providing new proofs of existing results (Propositions 2.2 and 2.3), as well as presenting new results (Propositions 2.4 and 2.5, Theorem 2.1). In Section 3, we recall an existing theorem (Theorem 3.1), before proving our main theorem (Theorem 3.2).

¹⁾ For a real-valued function f in \mathbb{R}^d whose smoothness is characterized by some integer $m \geq 1$, and for some prescribed accuracy $\epsilon > 0$, one shows that there exists a shallow network \hat{f}_W of size $W = W(d, m)$ that satisfies $\|f - \hat{f}_W\| \leq \epsilon$ for some norm $\|\cdot\|$.

2. Approximation Properties of Deep ReLU Networks

The ability of deep ReLU networks to implement the multiplication of two real numbers with amplitude M was proved by Yarotsky in [3, Prop. 1]. Liang and Srikant proved a similar result for $M = 1$ using networks with rectifier linear as well as binary step units in [16, Thm. 1]. In the rest of the paper, “with accuracy ϵ ” or “bounded” should be understood in the L^∞ -norm, unless stated otherwise. The extension to the multiplication in d dimensions first appeared in [28, Prop. 2.2].

Proposition 2.1 (Multiplication in two dimensions). *For any scalar $M \geq 1$, $N \geq 1$ and $0 < \epsilon < 1$, there is a deep ReLU network $\tilde{\pi}$ with inputs $(x_1, x_2) \in [-M, M] \times [-N, N]$, that has depth*

$$L = \mathcal{O} \left(\log_2 \frac{MN}{\epsilon} \right),$$

and size

$$W = \mathcal{O} \left(\log_2 \frac{MN}{\epsilon} \right),$$

such that

$$\|\tilde{\pi}(x_1, x_2) - x_1 x_2\|_{L^\infty([-M, M] \times [-N, N])} \leq \epsilon.$$

Equivalently, if the network has depth $L = \mathcal{O}(\log_2 \frac{1}{\epsilon})$ and size $W = \mathcal{O}(\log_2 \frac{1}{\epsilon})$, the approximation error satisfies $\|\tilde{\pi}(x_1, x_2) - x_1 x_2\|_{L^\infty([-M, M] \times [-N, N])} \leq MN\epsilon$.

The proposition of Yarotsky can be generalized to the d -dimensional case as in Prop. 2.2.

Proposition 2.2 (Multiplication in $d \geq 2$ dimensions). *For any scalar $M \geq 1$ and $0 < \epsilon < 1$, and any integer $d \geq 2$, there is a deep ReLU network $\tilde{\Pi}$ with inputs $(x_1, \dots, x_d) \in [-M, M]^d$, that has depth*

$$L = \mathcal{O} \left(d \log_2 \frac{d}{\epsilon} + d^2 \log_2 M \right),$$

and size

$$W = \mathcal{O} \left(d \log_2 \frac{d}{\epsilon} + d^2 \log_2 M \right),$$

such that

$$\left\| \tilde{\Pi}(x_1, \dots, x_d) - x_1 \dots x_d \right\|_{L^\infty([-M, M]^d)} \leq \epsilon.$$

Proof. Let $M \geq 1$ and $0 < \epsilon < 1$ be two scalars, and $d \geq 2$ an integer. For any scalar $A \geq 1$ and $B \geq 1$, let us call $\tilde{\pi}$ the network of Proposition 2.1 that implements the multiplication xy , $x \in [-A, A]$, $y \in [-B, B]$, with accuracy $AB\epsilon_0$, for some scalar $0 < \epsilon_0 < 1$ to be determined later. This network has depth and size $\mathcal{O}(\log_2 \frac{1}{\epsilon_0})$.

We construct the network $\tilde{\Pi}$ that implements the multiplication $x_1 x_2 \dots x_d$ as follows,

$$\begin{aligned} y_1 &= \tilde{\pi}(x_1, x_2), & |y_1| &\leq M^2(1 + \epsilon_0), \\ y_2 &= \tilde{\pi}(y_1, x_3), & |y_2| &\leq M^3(1 + \epsilon_0)^2, \\ y_3 &= \tilde{\pi}(y_2, x_4), & |y_3| &\leq M^4(1 + \epsilon_0)^3, \\ &\vdots & &\vdots \\ y_{d-1} &= \tilde{\pi}(y_{d-2}, x_d), & |y_{d-1}| &\leq M^d(1 + \epsilon_0)^{d-1}, \end{aligned}$$

and by setting $\tilde{\Pi}(x_1, \dots, x_d) = y_{d-1}$.

The network $\tilde{\Pi}$ has accuracy

$$\begin{aligned} &|\tilde{\Pi}(x_1, \dots, x_d) - x_1 \dots x_d| \\ &\leq |y_{d-1} - y_{d-2}x_d| + |x_d||y_{d-2} - y_{d-3}x_{d-1}| + \dots + |x_d x_{d-1} \dots x_5||y_3 - y_2x_4| \\ &\quad + |x_d x_{d-1} \dots x_4||y_2 - y_1x_3| + |x_d x_{d-1} \dots x_3||y_1 - x_1x_2|, \\ &< M^d(1 + \epsilon_0)^{d-2}\epsilon_0 + M^d(1 + \epsilon_0)^{d-3}\epsilon_0 + \dots + M^d(1 + \epsilon_0)^2 \\ &\quad + M^d(1 + \epsilon_0) + M^d\epsilon_0, \\ &< dM^d(1 + \epsilon_0)^d\epsilon_0 \quad (\text{crude estimate}). \end{aligned}$$

We choose $\epsilon_0 = \epsilon/(dM^d e)$ to obtain accuracy ϵ .

The depth and the size of the resulting network are equal to $(d - 1)$ times the depth and size of the network defined at the beginning of the proof. With accuracy ϵ_0 defined above, this gives depth and size

$$\mathcal{O}\left(d \log_2 \frac{dM^d e}{\epsilon}\right) = \mathcal{O}\left(d \log_2 \frac{d}{\epsilon} + d^2 \log_2 M\right).$$

The proof is complete. □

The network of Proposition 2.2 computes $x_1 \dots x_d$ as well as all the intermediate products $x_1 \dots x_k$, $2 \leq k \leq d - 1$, to the same accuracy ϵ . This allows us to prove the following result about polynomials, similar to [16, Thm. 2].

Proposition 2.3 (Polynomials). *For any scalar $M \geq 1$, $C \geq 0$ and $0 < \epsilon < 1$, any integer $n \geq 2$, and any polynomial p_n of degree n with input $x \in [-M, M]$ of the form*

$$p_n(x) = \sum_{k=0}^n c_k x^k, \quad \max_{0 \leq k \leq n} |c_k| \leq C,$$

there is a deep ReLU network \tilde{p}_n with inputs $(x_1, \dots, x_n) \in [-M, M]^n$, that has depth

$$L = \mathcal{O}\left(n \log_2 \frac{Cn}{\epsilon} + n^2 \log_2 M\right),$$

and size

$$W = \mathcal{O}\left(n \log_2 \frac{Cn}{\epsilon} + n^2 \log_2 M\right),$$

such that

$$\|\tilde{p}_n(x, \dots, x) - p_n(x)\|_{L^\infty([-M, M])} \leq \epsilon.$$

Proof. Let $M \geq 1$, $C \geq 0$ and $0 < \epsilon < 1$ be three scalars, $n \geq 2$ an integer, and consider a polynomial

$$p_n(x) = \sum_{k=0}^n c_k x^k, \quad \max_{0 \leq k \leq n} |c_k| \leq C.$$

We construct $\tilde{p}(x_1, \dots, x_n)$ as follows,

$$\tilde{p}_n(x_1, \dots, x_n) = c_0 + c_1 x_1 + \sum_{k=2}^n c_k y_{k-1}(x_1, \dots, x_k),$$

where $y_{k-1}(x_1, \dots, x_k)$ approximates $x_1 \dots x_k$ with the network of Proposition 2.2 to accuracy $0 < \epsilon_0 < 1$ to be determined later. (Note that when the inputs are the same $y_{k-1}(x, \dots, x)$ approximates x^k .)

The network \tilde{p}_n has accuracy

$$|\tilde{p}_n(x, \dots, x) - p_n(x)| \leq C \sum_{k=2}^n |y_{k-1}(x, \dots, x) - x^k| < nC\epsilon_0.$$

We choose $\epsilon_0 = \epsilon/(Cn)$ to obtain accuracy ϵ .

The resulting network has depth and size

$$\mathcal{O} \left(n \log_2 \frac{Cn^2 M^n}{\epsilon} \right) = \mathcal{O} \left(n \log_2 \frac{Cn}{\epsilon} + n^2 \log_2 M \right).$$

The proof is complete. □

The Chebyshev polynomials of the first kind play a central role in approximation theory [24]. They are defined on the interval $[-1, 1]$ via the three-term recurrence relation

$$T_n(x) = 2xT_{n-1}(x) - T_{n-2}(x), \quad n \geq 2, \tag{2.1}$$

with $T_0 = 1$ and $T_1(x) = x$. We show next how deep ReLU networks can efficiently implement Chebyshev polynomials, using the recurrence (2.1).

Proposition 2.4 (Chebyshev polynomials). *For any scalar $0 < \epsilon < 1$, any integer $n \geq 2$ and any Chebyshev polynomial T_n of degree n with input $x \in [-1, 1]$, there is a deep ReLU network \tilde{T}_n with inputs $(x_1, \dots, x_n) \in [-1, 1]^n$, that has depth*

$$L = \mathcal{O} \left(n \log_2 \frac{n}{\epsilon} + n^2 \right),$$

and size

$$W = \mathcal{O} \left(n \log_2 \frac{n}{\epsilon} + n^2 \right),$$

such that

$$\left\| \tilde{T}_n(x, \dots, x) - T_n(x) \right\|_{L^\infty([-1,1])} \leq \epsilon.$$

Proof. Let $0 < \epsilon < 1$ be a scalar and $n \geq 2$ be an integer. For any scalar $A \geq 1$ and $B \geq 1$, let us call $\tilde{\pi}$ the network of Proposition (2.1) that implements the multiplication xy , $x \in [-A, A]$, $y \in [-B, B]$, with accuracy $AB\epsilon_0$ for some scalar $0 < \epsilon_0 < 1$ to be determined later. This network has depth and size $\mathcal{O}\left(\log_2 \frac{1}{\epsilon_0}\right)$.

We construct the network \tilde{T}_n that approximates $T_n(x)$ as follows,

$$\begin{aligned} \tilde{T}_0 &= 1, & |\tilde{T}_0| &\leq 1, \\ \tilde{T}_1(x) &= x, & |\tilde{T}_1| &\leq 1, \\ \tilde{T}_2(x, x) &= 2\tilde{\pi}(x, \tilde{T}_1) - \tilde{T}_0, & |\tilde{T}_2| &< (1 + \epsilon_0)^2, \\ \tilde{T}_3(x, x, x) &= 2\tilde{\pi}(x, \tilde{T}_2) - \tilde{T}_1, & |\tilde{T}_3| &< 3(1 + \epsilon_0)^3, \\ \vdots & & \vdots & \\ \tilde{T}_n(x, \dots, x) &= 2\tilde{\pi}(x, \tilde{T}_{n-1}) - \tilde{T}_{n-2}, & |\tilde{T}_n| &< 3^{n-2}(1 + \epsilon_0)^n. \end{aligned}$$

Let us now estimate the accuracy e_n of the network $\tilde{T}_n(x, \dots, x)$, where $e_n = |\tilde{T}_n(x, \dots, x) - T_n(x)|$. We have

$$\begin{aligned} e_n &= |2\tilde{\pi}(x, \tilde{T}_{n-1}) - \tilde{T}_{n-2} - 2xT_{n-1} + T_{n-2}|, \\ &\leq 2|\tilde{\pi}(x, \tilde{T}_{n-1}) - x\tilde{T}_{n-1}| + 2|x||\tilde{T}_{n-1} - T_{n-1}| + e_{n-2}, \\ &\leq 2\epsilon_0|\tilde{T}_{n-1}| + 2e_{n-1} + e_{n-2}, \\ &< 2\epsilon_0 3^{n-3}(1 + \epsilon_0)^{n-1} + 2e_{n-1} + e_{n-2}, \\ &< n4^n(1 + \epsilon_0)^n \epsilon_0 \quad (\text{crude estimate}). \end{aligned}$$

We choose $\epsilon_0 = \epsilon/(n4^n e)$ to obtain accuracy ϵ .

The depth and the size of the resulting network are equal to $(n + 1)$ times the depth and size of the network defined at the beginning of the proof. With accuracy ϵ_0 defined above, this gives depth and size

$$\mathcal{O}\left(n \log_2 \frac{n4^n e}{\epsilon}\right) = \mathcal{O}\left(n \log_2 \frac{n}{\epsilon} + n^2\right).$$

The proof is complete. □

Note that we could have proven Proposition 2.4 using Proposition 2.3 and an estimate for the size C of the coefficients of the expansion of T_n in the monomial basis. (The leading term of T_n grows like 2^{n-1} , while the other terms grow at most like c^n , for some $c < 4$.)

Since Proposition 2.4 implements T_n , as well as the intermediate T_k 's, $0 \leq k \leq n - 1$, to the same accuracy ϵ , we have the following result about truncated Chebyshev series.

Proposition 2.5 (Truncated Chebyshev series). *For any scalar $C > 0$ and $0 < \epsilon < 1$, any integer $n \geq 2$, and any truncated Chebyshev series f_n of degree n with input $x \in [-1, 1]$, real coefficients c_k 's, $1 \leq k \leq n$, of the form*

$$f_n(x) = \sum_{k=0}^n c_k T_k(x), \quad \max_{0 \leq k \leq n} |c_k| \leq C,$$

there is a deep ReLU network \tilde{f}_n with inputs $(x_1, \dots, x_n) \in [-1, 1]^n$, that has depth

$$L = \mathcal{O} \left(n \log_2 \frac{Cn}{\epsilon} + n^2 \right),$$

and size

$$W = \mathcal{O} \left(n \log_2 \frac{Cn}{\epsilon} + n^2 \right),$$

such that

$$\left\| \tilde{f}_n(x, \dots, x) - f_n(x) \right\|_{L^\infty([-1,1])} \leq \epsilon.$$

Proof. Let $C \geq 0$ and $0 < \epsilon < 1$ be two scalars, $n \geq 2$ an integer. Consider a truncated Chebyshev series of the form

$$f_n(x) = \sum_{k=0}^n c_k T_k(x), \quad \max_{0 \leq k \leq n} |c_k| \leq C,$$

for some real coefficients c_k 's, $1 \leq k \leq n$.

We construct \tilde{f}_n as follows,

$$\tilde{f}_n(x_1, \dots, x_n) = c_0 + c_1 x_1 + \sum_{k=2}^n c_k \tilde{T}_k(x_1, \dots, x_k), \tag{2.2}$$

where \tilde{T}_k approximates T_k with the network of Proposition 2.4 to accuracy $0 < \epsilon_0 < 1$ to be determined later.

The network \tilde{f}_n has accuracy

$$|\tilde{f}_n(x, \dots, x) - f_n(x)| \leq C \sum_{k=2}^n |\tilde{T}_k - T_k| < nC\epsilon_0.$$

We choose $\epsilon_0 = \epsilon/(Cn)$ to obtain accuracy ϵ . The resulting network has depth

$$\mathcal{O} \left(n \log_2 \frac{Cn^2}{\epsilon} + n^2 \right) = \mathcal{O} \left(n \log_2 \frac{Cn}{\epsilon} + n^2 \right),$$

and size

$$\mathcal{O} \left(n \log_2 \frac{Cn^2}{\epsilon} + n^2 \right) = \mathcal{O} \left(n \log_2 \frac{Cn}{\epsilon} + n^2 \right).$$

The proof is complete. □

Chebyshev series lie at the heart of approximation theory. In particular, it is possible to show that Lipschitz continuous functions f with input $x \in [-M, M]$ have a unique absolutely and uniformly convergent Chebyshev series, and we write $f(x) = \sum_{k=0}^\infty c_k T_k(x/M)$ [24, Thm. 3.1]. For analytic functions, the truncated Chebyshev series defined as $f_n(x) = \sum_{k=0}^n c_k T_k(x/M)$ are *exponentially accurate* approximations [24, Thm. 8.2].

More precisely, for some scalars $M \geq 1$ and $s > 1$, let us define

$$a_s^M = M \frac{s + s^{-1}}{2}, \quad b_s^M = M \frac{s - s^{-1}}{2},$$

and the Bernstein s -ellipse scaled to $[-M, M]$,

$$E_s^M = \left\{ x + iy \in \mathbb{C} : \frac{x^2}{(a_s^M)^2} + \frac{y^2}{(b_s^M)^2} = 1 \right\}.$$

(It has foci $\sqrt{(a_s^M)^2 - (b_s^M)^2} = \pm M$, semi-major axis a_s^M and semi-minor axis b_s^M .) If a function f is analytic on the interval $[-M, M]$, and analytically continuable to the ellipse E_s^M , where it satisfies $|f(x)| < C_f$, for some $C_f > 0$, then, for each $n \geq 0$, the truncated Chebyshev series f_n satisfies

$$\|f_n - f\|_{L^\infty([-M, M])} \leq \frac{2C_f s^{-n}}{s - 1}. \tag{2.3}$$

Using Proposition 2.5 and Eq. (2.3), we prove a result about the approximation of analytic functions by deep ReLU networks.

Theorem 2.1 (Deep networks for analytic functions). *For any scalar $M \geq 1$, $s > 1$, $C_f > 0$ and $0 < \epsilon < 1$, and any real-valued analytic function f with input $x \in [-M, M]$ that is analytically continuable to the open ellipse E_s^M , where it satisfies $|f(x)| \leq C_f$, there is a deep ReLU network \tilde{f}_n with inputs $(x_1, \dots, x_n) \in [-M, M]^n$, that has depth*

$$L = \mathcal{O} \left(\frac{1}{\log_2^2 s} \log_2^2 \frac{C_f}{\epsilon} \right),$$

and size

$$W = \mathcal{O} \left(\frac{1}{\log_2^2 s} \log_2^2 \frac{C_f}{\epsilon} \right),$$

such that

$$\left\| \tilde{f}_n(x, \dots, x) - f(x) \right\|_{L^\infty([-M, M])} \leq \epsilon.$$

Proof. Let $M \geq 1$, $s > 1$, $C_f > 0$ and $0 < \epsilon < 1$ be four scalars, and f be an analytic function defined on $[-M, M]$ that is analytically continuable to the open Bernstein s -ellipse E_s^M , where it satisfies $|f(x)| \leq C_f$. We first approximate f by a truncated Chebyshev series f_n , and then approximate f_n by a deep ReLU network \tilde{f}_n using Proposition 2.5.

Since f is analytic in the open Bernstein s -ellipse E_s^M then, for any integer $n \geq 2$,

$$\|f_n(x) - f(x)\|_{L^\infty([-M, M])} \leq \frac{2C_f s^{-n}}{s - 1} = \mathcal{O}(C_f s^{-n}).$$

Therefore, if we take $n = \mathcal{O} \left(\frac{1}{\log_2 s} \log_2 \frac{2C_f}{\epsilon} \right)$, then the above term is bounded by $\epsilon/2$.

Let us now approximate f_n by a deep ReLU network \tilde{f}_n . We first write

$$f_n(x) = \sum_{k=0}^n c_k T_k \left(\frac{x}{M} \right),$$

with

$$\max_{0 \leq k \leq n} |c_k| = \mathcal{O}(C_f s), \text{ via [24, Thm. 8.1]}. \tag{2.4}$$

We then define our network \tilde{f}_n as in Proposition 2.5, with extra scaling x/M , and such that

$$|\tilde{f}_n(x, \dots, x) - f_n(x)| \leq \frac{\epsilon}{2}.$$

This yields

$$|\tilde{f}_n(x, \dots, x) - f(x)| \leq |\tilde{f}_n(x, \dots, x) - f_n(x)| + |f_n(x) - f(x)| \leq \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon.$$

To compute the depth and the size of the resulting network, we note that (i) the extra scaling x/M adds a layer and increases the size by $\mathcal{O}(n)$, (ii) the coefficients satisfy Equation (2.4), and (iii) the truncated series was computed to accuracy $\epsilon/2$. Therefore, the network $\tilde{f}_n(x, \dots, x)$ has depth

$$\mathcal{O}\left(n \log_2 \frac{2C_f s n}{\epsilon} + n^2 + 1\right) = \mathcal{O}\left(n \log_2 \frac{2C_f s n}{\epsilon} + n^2\right),$$

and size

$$\mathcal{O}\left(n \log_2 \frac{2C_f s n}{\epsilon} + n^2 + n\right) = \mathcal{O}\left(n \log_2 \frac{2C_f s n}{\epsilon} + n^2\right).$$

Using $n = \mathcal{O}\left(\frac{1}{\log_2 s} \log_2 \frac{2C_f}{\epsilon}\right)$, this gives depth and size

$$\begin{aligned} &\mathcal{O}\left(\left(\frac{1}{\log_2 s} \log_2 \frac{2C_f}{\epsilon}\right) \log_2 \left(\frac{2C_f s}{\epsilon} \frac{1}{\log_2 s} \log_2 \frac{2C_f}{\epsilon}\right) + \frac{1}{\log_2^2 s} \log_2^2 \frac{2C_f}{\epsilon}\right) \\ &= \mathcal{O}\left(\frac{1}{\log_2^2 s} \log_2^2 \frac{C_f}{\epsilon}\right). \end{aligned}$$

The proof is complete. □

Theorem 2.1 above could be generalized to multiple dimensions, which would be interesting future work. In [25], it was shown that deep ReLU networks can approximate multivariate analytic functions with exponential convergence, a result similar to our theorem above. However, we would like to emphasize that it is not possible to directly apply the result in [25] to prove our main theorem in Section 3, because it is only valid on an open interval contained in $[-1, 1]$, instead of an arbitrary closed interval $[-M, M]$.

Let us now highlight that, in general, the constants s and C_f depend on M . Let us look at two examples, a function with a singularity on the imaginary axis and an *entire* function (*i.e.*, a function that is analytic over the whole complex plane). A typical example of an analytic function with singularities on the imaginary axis is the Runge-like function $f(x) = 1/(1 + \frac{x^2}{\beta^2})$, $\beta > 1$, whose singularities are located at $x = \pm i\beta$. The function f is analytic on the interval $[-M, M]$ and analytically continuable to the open Bernstein s -ellipse E_s^M with

$$s(M) = \frac{\sqrt{(4M^2 - 2)r^2 + r^4 + 1} + r^2 - 1}{2Mr}$$

and $r = \beta + \sqrt{\beta^2 + 1}$. Since f increases along the imaginary axis we may take

$$C_f(M) = f\left(M \frac{s(M) - s(M)^{-1}}{2}\right).$$

The complex exponential $f(x) = e^{ix}$ is an entire function. Hence, any $s > 1$ works but $C_f(s, M)$ must grow with s and M . As f increases along the imaginary axis we may choose

$$C_f(s, M) = f\left(M \frac{s - s^{-1}}{2}\right) = e^{M \frac{s - s^{-1}}{2}}. \tag{2.5}$$

In this case the network of Thm. 2.1 has depth and size

$$\mathcal{O}\left(\frac{1}{\log_2^2 s} \left(M \frac{s - s^{-1}}{2} + \log_2 \frac{1}{\epsilon}\right)^2\right).$$

We would also like to mention that the ReLU activation function is not an optimal choice for constructing neural networks to approximate smooth functions. For example, Thm. 2.3 of [9] shows that one-hidden-layer shallow networks with $\mathcal{O}(\log(\frac{1}{\epsilon}))$ parameters can approximate analytic functions with ϵ accuracy when a smooth activation function is used. The disadvantage of the ReLU activation function in this scenario is not unexpected since it is not a natural choice to use a function that is not differentiable to approximate a smooth function. However, from the point of view of deep learning and optimization, ReLU is a much better choice [25]. The study in this paper should be regarded as a complement to existing approximation theory, using a more modern approach.

3. Approximation of Generalized Bandlimited Functions by Deep ReLU Networks

A famous theorem of Carathéodory states that if a point $x \in \mathbb{R}^d$ lies in the the convex hull of a set P , then x can be written as the convex combination of at most $d + 1$ points in P . Maurey’s theorem [26] is an extension of Carathéodory’s result to the infinite-dimensional case. It was used in the context of shallow network approximations by Barron in 1993 [12]. We recall Maurey’s theorem below.

Theorem 3.1 (Maurey’s theorem). *Let H be a Hilbert space with norm $\|\cdot\|$. Suppose there exists $G \subset H$ such that for every $g \in G$, $\|g\| \leq b$ for some $b > 0$. Then, for every f in the convex hull of G and every integer $n \geq 1$, there is a f_n in the convex hull of n points in G and a constant $c > b^2 - \|f\|^2$ such that $\|f - f_n\|^2 \leq \frac{c}{n}$.*

We are now ready to prove our main theorem about the approximation of generalized bandlimited functions of the form (1.2) by deep ReLU networks. Let us first define a Hilbert space of such functions.

Definition 3.1 (Generalized bandlimited functions). *Let $d \geq 2$ be an integer, $M \geq 1$ be a scalar, and $B = [0, 1]^d$. Suppose $K : \mathbb{R} \rightarrow \mathbb{C}$ is analytic and bounded by a constant $D_K \in (0, 1]$ on $[-dM, dM]$, and that K satisfies the assumption of Thm. 2.1 for some $s > 1$ and $C_K > 0$. We define the Hilbert space $\mathcal{H}_{K,M}(B)$ of generalized bandlimited functions via*

$$\mathcal{H}_{K,M}(B) = \left\{ f(\mathbf{x}) = \int_{[-M,M]^d} F(\mathbf{w})K(\mathbf{w} \cdot \mathbf{x})d\mathbf{w} \mid F : [-M, M]^d \rightarrow \mathbb{C} \text{ is in } L^2([-M, M]^d) \right\},$$

with an inner product

$$\langle f, g \rangle_{\mathcal{H}_{K,M}(B)} := \int_{[-M,M]^d} F_f(\mathbf{w})\overline{F_g(\mathbf{w})}d\mathbf{w}$$

and norm $\|f\|_{\mathcal{H}_{K,M}(B)} := \|F_f\|_{L^2([-M,M]^d)}$, where

$$F_f = \arg \min_{F \in S_f} \|F\|_{L^2([-M,M]^d)}, \quad S_f = \left\{ F \mid f(\mathbf{x}) = \int_{[-M,M]^d} F(\mathbf{w})K(\mathbf{w} \cdot \mathbf{x})d\mathbf{w} \right\}.$$

Note that

$$|f(\mathbf{x})| \leq D_K \int_{[-M,M]^d} |F_f(\mathbf{w})|d\mathbf{w} \leq (2M)^{d/2} \|F_f\|_{L^2([-M,M]^d)} = (2M)^{d/2} \|f\|_{\mathcal{H}_{K,M}(B)}.$$

The above inequality shows that if we consider an evaluation functional $L_{\mathbf{x}}$ defined on $\mathcal{H}_{K,M}(B)$ as follows,

$$f(\mathbf{x}) = L_{\mathbf{x}}(f) := \int_{[-M,M]^d} F_f(\mathbf{w})K(\mathbf{w} \cdot \mathbf{x})d\mathbf{w},$$

then $L_{\mathbf{x}}$ is bounded on $\mathcal{H}_{K,M}(B)$. Hence, $\mathcal{H}_{K,M}(B)$ is a RKHS; it contains the space of bandlimited functions, which corresponds to $K(t) = e^{it}$. For simplicity, we will use F instead of F_f for $f \in \mathcal{H}_{K,M}(B)$, when the dependency on f is clear.

Theorem 3.2 (Deep networks for $\mathcal{H}_{K,M}$). *Suppose f is an arbitrary real-valued function in $\mathcal{H}_{K,M}(B)$, for some function K , scalars $M \geq 1$, $s > 1$ and $C_K > 0$, and integer $d \geq 2$. Let us assume that $\int_{\mathbb{R}^d} |F(\mathbf{w})|d\mathbf{w} = \int_{[-M,M]^d} |F(\mathbf{w})|d\mathbf{w} = C_F$. Then, for any measure μ and any scalar $0 < \epsilon < 1$, there exists a deep ReLU network \tilde{f} with inputs $\mathbf{x} \in B = [0, 1]^d$, that has depth*

$$L = \mathcal{O} \left(\frac{1}{\log_2^2 s} \log_2^2 \frac{C_F C_K \sqrt{\mu(B)}}{\epsilon} \right),$$

and size

$$W = \mathcal{O} \left(\frac{C_F^2 \mu(B)}{\epsilon^2 \log_2^2 s} \log_2^2 \frac{C_F C_K \sqrt{\mu(B)}}{\epsilon} \right),$$

such that

$$\|\tilde{f} - f\|_{L^2(\mu, B)} = \sqrt{\int_B |\tilde{f}(\mathbf{x}) - f(\mathbf{x})|^2 d\mu(\mathbf{x})} \leq \epsilon.$$

Proof. Let f be an arbitrary function in $\mathcal{H}_{K,M}$, and μ be an arbitrary measure. Let $F(\mathbf{w}) = |F(\mathbf{w})|e^{i\theta(\mathbf{w})}$. Since f is real-valued, we may write

$$\begin{aligned} f(\mathbf{x}) &= \operatorname{Re} \left(\int_{\mathbb{R}^d} F(\mathbf{w})K(\mathbf{w} \cdot \mathbf{x})d\mathbf{w} \right), \\ &= \operatorname{Re} \left(\int_{\mathbb{R}^d} C_F e^{i\theta(\mathbf{w})} K(\mathbf{w} \cdot \mathbf{x}) \frac{|F(\mathbf{w})|}{C_F} d\mathbf{w} \right), \\ &= \int_{[-M,M]^d} C_F \left[\cos(\theta(\mathbf{w}))K_R(\mathbf{w} \cdot \mathbf{x}) - \sin(\theta(\mathbf{w}))K_I(\mathbf{w} \cdot \mathbf{x}) \right] \frac{|F(\mathbf{w})|}{C_F} d\mathbf{w}, \end{aligned}$$

where $K_R(\mathbf{w} \cdot \mathbf{x}) = \text{Re}(K(\mathbf{w} \cdot \mathbf{x}))$ and $K_I(\mathbf{w} \cdot \mathbf{x}) = \text{Im}(K(\mathbf{w} \cdot \mathbf{x}))$. The integral above represents f as an infinite convex combination of functions in the set

$$G_{K,M} = \left\{ \gamma [\cos(\beta)\text{Re}(K(\mathbf{w} \cdot \mathbf{x})) - \sin(\beta)\text{Im}(K(\mathbf{w} \cdot \mathbf{x}))], |\gamma| \leq C_F, \beta \in \mathbb{R}, \mathbf{w} \in [-M, M]^d \right\}.$$

Therefore, f is in the closure of the convex hull of $G_{K,M}$. Since functions in $G_{K,M}$ are bounded in the $L^2(\mu, B)$ -norm by $2C_F D_K \sqrt{\mu(B)} \leq 2C_F \sqrt{\mu(B)}$, Theorem 3.1 tells us that there exist real coefficients b_j 's and β_j 's such that¹⁾

$$f_{\epsilon_0}(\mathbf{x}) = \sum_{j=1}^{\lceil 1/\epsilon_0^2 \rceil} b_j [\cos(\beta_j)K_R(\mathbf{w} \cdot \mathbf{x}) - \sin(\beta_j)K_I(\mathbf{w} \cdot \mathbf{x})], \quad \sum_{j=1}^{\lceil 1/\epsilon_0^2 \rceil} |b_j| \leq C_F,$$

for some $0 < \epsilon_0 < 1$ to be determined later, with

$$\|f_{\epsilon_0}(\mathbf{x}) - f(\mathbf{x})\|_{L^2(\mu, B)} \leq 2C_F \sqrt{\mu(B)} \epsilon_0.$$

We now approximate $f_{\epsilon_0}(\mathbf{x})$ by a deep ReLU network $\tilde{f}(\mathbf{x})$. Note that K_R and K_I are both analytic and satisfy the same assumptions as K . Using Theorem 2.1, they can be approximated to accuracy ϵ_0 using networks \tilde{K}_R and \tilde{K}_I of depth and size

$$\mathcal{O} \left(\frac{1}{\log_2^2 s} \log_2^2 \frac{C_K}{\epsilon_0} \right).$$

We define the deep ReLU network $\tilde{f}(\mathbf{x})$ by

$$\tilde{f}(\mathbf{x}) = \sum_{j=1}^{\lceil 1/\epsilon_0^2 \rceil} b_j [\cos(\beta_j)\tilde{K}_R(\mathbf{w} \cdot \mathbf{x}) - \sin(\beta_j)\tilde{K}_I(\mathbf{w} \cdot \mathbf{x})].$$

This network has depth $L = \mathcal{O} \left(\frac{1}{\log_2^2 s} \log_2^2 \frac{C_K}{\epsilon_0} \right)$ and size $W = \mathcal{O} \left(\frac{1}{\epsilon_0^2 \log_2^2 s} \log_2^2 \frac{C_K}{\epsilon_0} \right)$, and

$$\begin{aligned} & |\tilde{f}(\mathbf{x}) - f_{\epsilon_0}(\mathbf{x})| \\ & \leq \sum_{j=1}^{\lceil 1/\epsilon_0^2 \rceil} |b_j| |\tilde{K}_R(\mathbf{w}_j \cdot \mathbf{x}) - K_R(\mathbf{w}_j \cdot \mathbf{x})| + \sum_{j=1}^{\lceil 1/\epsilon_0^2 \rceil} |b_j| |\tilde{K}_I(\mathbf{w}_j \cdot \mathbf{x}) - K_I(\mathbf{w}_j \cdot \mathbf{x})| \\ & \leq 2C_F \epsilon_0, \end{aligned}$$

which yields

$$\left\| \tilde{f}(\mathbf{x}) - f_{\epsilon_0}(\mathbf{x}) \right\|_{L^2(\mu, B)} \leq 2C_F \sqrt{\mu(B)} \epsilon_0.$$

The total approximation error satisfies

$$\left\| \tilde{f}(\mathbf{x}) - f(\mathbf{x}) \right\|_{L^2(\mu, B)} \leq 4C_F \sqrt{\mu(B)} \epsilon_0.$$

We take

$$\epsilon_0 = \frac{\epsilon}{4C_F \sqrt{\mu(B)}}$$

¹⁾ We use Theorem 3.1 with $b = 2C_F \sqrt{\mu(B)}$, $c = b^2 > b^2 - \|f\|^2$, and $\|\cdot\| = \|\cdot\|_{L^2(\mu, B)}$.

to complete the proof. \square

Let us end this section with comments on the constants C_F , C_K and $\mu(B)$; we start with C_F . If F is a mollifier then $C_F = 1$, whereas if F is a normal distribution truncated to $[-M, M]^d$ then $C_F < 1$. In general, however, C_F might grow algebraically or exponentially with the dimension d .

We continue with C_K . Consider for example the complex exponential kernel $K(t) = e^{it}$, $t \in [-dM, dM]$. Equation 2.5 yields

$$C_K(s, dM) = e^{dM \frac{s-s^{-1}}{2}}, \quad \text{for any } s > 1.$$

The resulting network to approximate a function to accuracy ϵ in the $L^2(\mu, B)$ -norm with such a kernel has depth

$$L = \mathcal{O} \left(\frac{1}{\log_2^2 s} \left(dM \frac{s-s^{-1}}{2} + \log_2 \frac{C_F \sqrt{\mu(B)}}{\epsilon} \right)^2 \right),$$

and size

$$W = \mathcal{O} \left(\frac{C_F^2 \mu(B)}{\epsilon^2 \log_2^2 s} \left(dM \frac{s-s^{-1}}{2} + \log_2 \frac{C_F \sqrt{\mu(B)}}{\epsilon} \right)^2 \right).$$

We conclude with $\mu(B)$. If μ is a probability measure, then $\mu(B) \leq 1$ for any compact domain B . If μ is Lebesgue measure, then $\mu(B) = 1$ for $B = [0, 1]^d$, but grows exponentially with the dimension d if $B = [0, \ell]^d$, $\ell > 1$. This is a common drawback in the approximation theory of neural networks for conquering the curse of dimensionality [12].

4. Discussion

We have proven new upper bounds for the approximation of bandlimited functions of the form (1.2), for which the curse of dimensionality is overcome. Our proof is based on Maurey's theorem and on the ability of deep ReLU networks to approximate Chebyshev polynomials and analytic functions efficiently.

There are many ways in which this work could be profitably continued. The space of bandlimited functions is a type of RKHS and therefore a possible extension would be to look at different types of RKHS. One could also relax the bandlimited assumption (1.2), e.g., to functions F whose derivatives are rapidly decreasing. In this case, the kernel K could be approximated on the real line by Chebyshev polynomials on truncated intervals or Hermite polynomials. The latter is another example of classical orthogonal polynomials, which can be represented by a three-term recurrence relation similar to (2.1) and efficiently implemented by deep ReLU networks.

Let us conclude this paper with a comment on deep versus shallow networks in the context of parallel computing. Since the depth L grows like $\mathcal{O}(\log_2^2 \frac{1}{\epsilon})$ in Theorem 3.2, the approximation accuracy for deep networks can be root-exponentially improved if L increases. Hence, very deep networks are more efficient than shallow networks when both parallel computing efficiency and approximation efficiency are considered. This is in contrast with the more general case of continuous functions, the approximation of which via very deep networks might be less attractive in terms of parallel computing [27].

Acknowledgements. We thank the members of the CM3 group at Columbia University and our colleague Mikael Slevinsky for fruitful discussions. The first author is much indebted to former PhD supervisor Nick Trefethen for his inspirational contributions to numerical analysis and in particular to the field of approximation theory.

The research of the second author was partially supported by US NSF under the grant award DMS-1945029. The research of the third author is supported in part by US NSF DMS-1719699 and the NSF TRIPODS program CCF-1704833.

References

- [1] R.A. DeVore, R. Howard and C. Micchelli, Optimal nonlinear approximation, *Manuscripta Math.*, **63** (1989), 469–478.
- [2] R.A. DeVore, Nonlinear approximation, *Acta Numer.*, **7** (1998), 51–150.
- [3] D. Yarotsky, Error bounds for approximations with deep ReLU networks, *Neural Netw.*, **94** (2017), 103–114.
- [4] Y. LeCun, Y. Bengio and G. Hinton, Deep learning, *Nature*, **521** (2015), 436–444.
- [5] X. Glorot, A. Bordes and Y. Bengio, Deep sparse rectifier neural networks, G. Gordon, D. Dunson and M. Dudík, editors, Proceedings of the 14th International Conference on Artificial Intelligence and Statistics, Proc. Mach. Learn. Res. 15, pages 315–323, Fort Lauderdale, FL, 2011.
- [6] G. Cybenko, Approximation by superpositions of a sigmoidal function, *Math. Control Signals Syst.*, **2** (1989), 303–314.
- [7] K. Hornik, M. Stinchcombe and H. White, Multilayer feedforward networks are universal approximators, *Neural Netw.*, **2** (1989), 359–366.
- [8] H.N. Mhaskar, Approximation properties of a multilayered feedforward artificial neural network, *Adv. Comput. Math.*, **1** (1993), 61–80.
- [9] H.N. Mhaskar, Neural networks for optimal approximation of smooth and analytic functions, *Neural Comput.*, **8** (1996), 164–177.
- [10] H.N. Mhaskar and C.A. Micchelli, Approximation by superposition of sigmoidal and radial basis functions, *Adv. Appl. Math.*, **13** (1992), 350–373.
- [11] A. Pinkus, Approximation theory of the MLP model in neural networks, *Acta Numer.*, **8** (1999), 143–195.
- [12] A. Barron, Universal approximation bounds for superpositions of a sigmoidal function, *IEEE Trans. Inf. Theory*, **39** (1993), 930–945.
- [13] F. Bach, Breaking the curse of dimensionality with convex neural networks, *J. Mach. Learn. Res.*, **18** (2017), 1–53.
- [14] N. Cohen, O. Sharir and A. Shashua, On the expressive power of deep learning: A tensor analysis, V. Feldman, A. Rakhlin and O. Shamir, editors, 29th Annual Conference on Learning Theory, Proc. Mach. Learn. Res. 49, pages 698–728, Columbia University, New York, 2016.
- [15] R. Eldan and O. Shamir, The power of depth for feedforward neural networks, V. Feldman, A. Rakhlin and O. Shamir, editors, 29th Annual Conference on Learning Theory, Proc. Mach. Learn. Res. 49, pages 907–940, Columbia University, New York, 2016.
- [16] S. Liang and R. Srikant, Why deep neural networks for function approximation?, *arXiv:1610.0416*, (2016).
- [17] H. Montanelli and Yang, Error bounds for deep ReLU networks using the Kolmogorov–Arnold superposition theorem, *Neural Netw.*, **129** (2020), 1–6.
- [18] P. Petersen and F. Voigtlaender, Optimal approximation of piecewise smooth functions using deep ReLU neural networks, *Neural Netw.*, **108** (2018), 296–330.
- [19] T. Poggio, H.N. Mhaskar, L. Rosasco, B. Miranda and Q. Liao, Why and when can deep—but not shallow—networks avoid the curse of dimensionality: A review, *Int. J. Control Autom. and*

- Comput.*, **14** (2017), 503–519.
- [20] U. Shaham, A. Cloninger and R.R. Coifman, Provable approximation properties for deep neural networks, *Appl. Comput. Harm. Anal.*, **44** (2018), 537–557.
- [21] M. Telgarsky, Benefits of depth in neural networks, V. Feldman, A. Rakhlin and O. Shamir, editors, 29th Annual Conference on Learning Theory, Proc. Mach. Learn. Res. 49, pages 1517–1539, Columbia University, New York, 2016.
- [22] N.M. Korobov, On the approximate solution of integral equations, *Dokl. Akad. Nauk SSSR*, **128** (1959), 233–238.
- [23] H. Montanelli and Q. Du, New error bounds for deep ReLU networks using sparse grids, *SIAM J. Math. Data Sci.*, **1** (2019), 78–92.
- [24] L.N. Trefethen, *Approximation Theory and Approximation Practice*, SIAM, Philadelphia, PA, 2013.
- [25] W. E and Q. Wang, Exponential convergence of the deep neural network approximation for analytic functions, *Sci. China Math.*, **61** (2018), 1733–1740.
- [26] G. Pisier, Remarques sur un résultat non publié de B. Maurey, Technical Report 5, Séminaire d’analyse fonctionnelle, École Polytechnique, 1981.
- [27] Z. Shen, H. Yang and S. Zhang, Nonlinear approximation via compositions, *Neural Netw.*, **119** (2019), 74–84.
- [28] L. Chen and C. Wu, A note on the expressive power of deep rectified linear unit networks in high-dimensional spaces, *Math. Methods Appl. Sci.*, **42** (2019).