arXiv:2005.10807v2 [math.FA] 2 Oct 2020

KOLMOGOROV WIDTH DECAY AND POOR APPROXIMATORS IN MACHINE LEARNING: SHALLOW NEURAL NETWORKS, RANDOM FEATURE MODELS AND NEURAL TANGENT KERNELS

WEINAN E AND STEPHAN WOJTOWYTSCH

Dedicated to Andrew Majda on the occasion of his 70th birthday

Abstract. We establish a scale separation of Kolmogorov width type between subspaces of a given Banach space under the condition that a sequence of linear maps converges much faster on one of the subspaces. The general technique is then applied to show that reproducing kernel Hilbert spaces are poor L^2-approximators for the class of two-layer neural networks in high dimension, and that multi-layer networks with small path norm are poor approximators for certain Lipschitz functions, also in the L^2-topology.

1. Introduction

It has been known since the early 1990s that two-layer neural networks with sigmoidal or ReLU activation can approximate arbitrary continuous functions on compact sets in the uniform topology [Cyb89, Hor91]. In fact, when approximating a suitable (infinite-dimensional) class of functions in the L^2 topology of *any* compactly supported Radon probability measure, two-layer networks can evade the curse of dimensionality [Bar93]. In this article, we show that

(1) infinitely wide random feature functions with norm bounds are much worse approximators in high dimension compared to two-layer neural networks.
(2) infinitely wide neural networks are subject to the curse of dimensionality when approximating general Lipschitz functions in high dimension.

In both cases, we consider approximation in the $L^2([0,1]^d)$-topology. The second statement applies more generally to any model in which few data samples are needed to estimate integrals uniformly over the hypothesis class. In the first point, we can consider more general kernel methods instead of random features (including certain neural tangent kernels), and the second claim holds true for multi-layer networks as well as deep ResNets of bounded width. We conjecture that Lipschitz functions in the second statement could be replaced with C^k functions for fixed k. Precise statements of the results are given in Corollary 3.4 and Example 4.3.

To prove these results, we show more generally that if X, Y are subspaces of a Banach space Z and a sequence of linear maps A_n converges quickly to a limit A on X, but not on Y, then there must be a Kolmogorov width-type separation between X and Y. The classical notion of Kolmogorov width is considered in Lemma 2.1 and later extended to a stronger notion of separation in Lemma 2.3.

We apply the abstract result to the pairs $X =$ Barron space (for two-layer networks) or $X =$ tree-like function space (for multi-layer networks)/$Y =$ Lipschitz space, and $X =$ RKHS/$Y =$

Date: December 20, 2020.
2020 Mathematics Subject Classification. 68T07, 41A30, 41A65, 46E15, 46E22 .
Key words and phrases. Curse of dimensionality, two-layer network, multi-layer network, population risk, Barron space, reproducing kernel Hilbert space, random feature model, neural tangent kernel, Kolmogorov width, approximation theory.

Barron space. In the first case, the sequence of linear maps is given by a type of Monte-Carlo integration, in the second case by projection onto the eigenspaces of the RKHS kernel.

This article is structured as follows. In Section 2, we prove the abstract result which we apply to Barron/tree-like function spaces and Lipschitz space in Section 3 and to RKHS and Barron space in Section 4. We conclude by discussing our results and some open questions in Section 5. In appendices A and B, we review the natural function spaces for shallow neural networks and kernel methods respectively. In Appendix B, we specifically focus on kernels arising from random feature models and neural tangent kernels for two-layer neural networks.

1.1. **Notation.** We denote the closed ball of radius $r > 0$ around the origin in a Banach space by B_r^X and the unit ball by $B_1^X = B^X$. The space of continuous linear maps between Banach spaces X, Y is denoted by $L(X, Y)$ and the continuous dual space of X by $X^* = L(X, \mathbb{R})$. The support of a Radon measure μ is denoted by $\operatorname{spt} \mu$.

2. An Abstract Lemma

2.1. **Kolmogorov Width Version.** The Kolmogorov width of a function class $\mathcal{F}$ in another function class $\mathcal{G}$ with respect to a metric d on the union of both classes is defined as the biggest distance of an element in $\mathcal{G}$ from the class $\mathcal{F}$:

$$w_d(\mathcal{F}; \mathcal{G}) = \sup_{g \in \mathcal{G}} \operatorname{dist}(g, \mathcal{F}) = \sup_{g \in \mathcal{G}} \inf_{f \in \mathcal{F}} d(f, g).$$

In this article, we consider the case where $\mathcal{G}$ is the unit ball in a Banach space Y, $\mathcal{F}$ is the ball of radius $t > 0$ in a Banach space X and $d = d_Z$ is induced by the norm on a Banach space Z into which both X and Y embed densely. As t increases, points in Y are approximated to higher degrees of accuracy by elements of X. The rate of decay

$$\rho(t) := w_{d_Z}(B_t^X, B_1^Y)$$

provides a quantitative measure of density of X in Y with respect to the topology of Z. For a different point of view on width focusing on approximation by finite-dimensional spaces, see [Lor66, Chapter 9].

In the following Lemma, we show that if there exists a sequence of linear operators on Z which behaves sufficiently differently on X and Y, then ρ must decay slowly as $t \to \infty$.

Lemma 2.1. *Let X, Y, Z, W be Banach spaces such that $X, Y \hookrightarrow Z$. Assume that $A_n, A : Z \to W$ are continuous linear operators such that*

$$\|A_n - A\|_{L(X,W)} \leq C_X\, n^{-\alpha}, \qquad \|A_n - A\|_{L(Y,W)} \geq c_Y\, n^{-\beta}, \qquad \|A_n - A\|_{L(Z,W)} \leq C_Z$$

for $\beta < \alpha$ and constants $C_X, c_Y, C_Z > 0$. Then

$$\rho(t) \geq 2^{-\beta} \frac{(c_Y/2)^{\frac{\alpha}{\alpha-\beta}}}{C_Z\, C_X^{\frac{\beta}{\alpha-\beta}}}\, t^{-\frac{\beta}{\alpha-\beta}} \quad \forall\ t \geq \frac{c_Y}{2C_X} \tag{2.1}$$

and

$$\liminf_{t \to \infty} \left(t^{\frac{\beta}{\alpha-\beta}}\, \rho(t) \right) \geq \frac{(c_Y/2)^{\frac{\alpha}{\alpha-\beta}}}{C_Z\, C_X^{\frac{\beta}{\alpha-\beta}}}. \tag{2.2}$$

Proof. Choose a sequence $y_n \in B^Y$ such $\|(A_n - A)y_n\|_W \geq c_Y\, n^{-\beta}$ and $x_n \in X$ such that

$$x_n \in \operatorname{argmin}_{\{x : \|x\|_X \leq t_n\}} \|x - y_n\|_Z \qquad \text{for} \qquad t_n := \frac{c_Y}{2\, C_X}\, n^{\alpha-\beta}$$

(see Remark 2.2). Then

$$c_Y\, n^{-\beta} \leq \|(A_n - A)y_n\|_W$$

$$\begin{aligned}
&\leq \|(A_n - A)(y_n - x_n)\|_W + \|(A_n - A)x_n\|_W \\
&\leq C_Z \, \|y_n - x_n\|_Z + C_X \, n^{-\alpha} \|x_n\|_X \\
&\leq C_Z \, \|x_n - y_n\|_Z + \frac{c_Y}{2} \, n^{-\beta}.
\end{aligned}$$

We therefore have

$$\|x_n - y_n\|_Z \geq \frac{c_Y}{2\, C_Z} \, n^{-\beta} = \frac{c_Y}{2\, C_Z} \left(\frac{2\, C_X}{c_Y} \, t_n \right)^{\frac{-\beta}{\alpha - \beta}} = \left(\frac{c_Y}{2} \right)^{\frac{\alpha}{\alpha - \beta}} \frac{1}{C_Z \, C_X^{\frac{\beta}{\alpha - \beta}}} \, t_n^{-\frac{\beta}{\alpha - \beta}}.$$

Clearly $t_n \to \infty$ since $\alpha > \beta$. For general $t > 0$, take $t_n = \inf\{t_k : k \in \mathbb{N}, t_k \geq t\}$. Then

$$\begin{aligned}
\rho(t) &\geq \rho(t_n) \\
&\geq \frac{(c_Y/2)^{\frac{\alpha}{\alpha - \beta}}}{C_Z \, C_X^{\frac{\beta}{\alpha - \beta}}} \, t_n^{-\frac{\beta}{\alpha - \beta}} \\
&\geq \frac{(c_Y/2)^{\frac{\alpha}{\alpha - \beta}}}{C_Z \, C_X^{\frac{\beta}{\alpha - \beta}}} \left(\frac{t_{n-1}}{t_n} \right)^{\frac{\beta}{\alpha - \beta}} \, t^{-\frac{\beta}{\alpha - \beta}} \\
&= \frac{(c_Y/2)^{\frac{\alpha}{\alpha - \beta}}}{C_Z \, C_X^{\frac{\beta}{\alpha - \beta}}} \left(\frac{n-1}{n} \right)^{\beta} \, t^{-\frac{\beta}{\alpha - \beta}}.
\end{aligned}$$

As $t \to \infty$, so does n, and the n-dependent term converges to 1. □

Remark 2.2. Generally elements like x_n, y_n may not exist if the extremum is not attained. Otherwise, we can choose x_n such that $\|x_n - y_n\|_Z$ is sufficiently close to its infimum and $\|(A_n - A)y_n\|$ is sufficiently close to its supremum. To simplify our presentation, we assume that the supremum and infimum are attained.

The choice of x_n as a minimizer is valid if

(1) X embeds into Z compactly, so the minimum of the continuous function $\| \cdot - y\|_Z$ is attained on the compact set $\{\| \cdot \|_X \leq t_k\}$, or
(2) the embedding $X \hookrightarrow Z$ maps closed bounded sets to closed sets and Z admits continuous projections onto closed convex sets (for example, Z is uniformly convex).

In the applications below, the first condition will be met.

2.2. **Improved Estimate.** In the previous section, we have shown by elementary means that the estimate

$$\liminf_{t \to \infty} \left(t^{\gamma} \sup_{\|y\| \leq 1} \inf_{\|x\|_X \leq t} \|x - y\|_Z \right) \geq c > 0$$

holds for suitable γ if a sequence of linear maps between Z and another Banach space W behaves very differently on subspaces X and Y of Z. So intuitively, on each scale $t > 0$ there exists an element $y_t \in B^Y$ such that y_t is poorly approximable by elements in X on this scale. In this section, we establish that there exists a single point $y \in Y$ which is poorly approximable across infinitely many scales. This statement has applications in Wasserstein gradient flows for machine learning which we discuss in a companion article [WE20].

Lemma 2.3. *Let X, Y, Z be Banach spaces such that $X, Y \hookrightarrow Z$. Assume that $A_n, A \in L(Z, W)$ are operators such that*

$$\|A_n - A\|_{L(X,W)} \leq C_X \, n^{-\alpha}, \qquad \|A_n - A\|_{L(Y,W)} \geq c_Y \, n^{-\beta}, \qquad \|A_n - A\|_{L(Z,W)} \leq C_Z$$

for $\beta < \frac{\alpha}{2}$ and constants C_X, c_Y, C_Z. Then there exists $y \in B^Y$ such that for every $\gamma > \frac{\beta}{\alpha-\beta}$ we have

$$\limsup_{t\to\infty} \left(t^\gamma \inf_{\|x\|_X \le t} \|x-y\|_Z \right) = \infty.$$

The result is stronger than the previous one in that it fixes a single point y which is poorly approximable in infinitely many scales t_{n_k}. While in each scale t_n there exists a point y_n which is poorly approximable, we only show that y is poorly approximable in infinitely many scales, not in all scales.

Proof of Lemma 2.3. Since $Y \hookrightarrow Z$, there exists a constant $C^Y > 0$ such that $\|A_n - A\|_{Y^*} \le C^Y$.

Definition of y. Choose sequences $y_n \in B^Y$ and $w_n^* \in B^{W^*}$ such

$$w_n^* \circ (A_n - A)(y_n) \ge c_Y\, n^{-\beta}.$$

Consider two sequences n_k, m_k of strictly increasing integers such that

$$\sum_{k=1}^\infty \frac{1}{n_k} \le 1.$$

We will impose further conditions below. Set

$$y := \sum_{k=0}^\infty \frac{\varepsilon_k}{n_k}\, y_{m_k}$$

where the signs $\varepsilon_k \in \{-1, 1\}$ are chosen inductively such that

$$\varepsilon_K \cdot w_{m_k}^* \circ (A_{m_K} - A) \left(\sum_{k=1}^{K-1} \frac{\varepsilon_k}{n_k} y_{m_k} \right) \ge 0.$$

Clearly

$$\|y\|_Y \le \sum_{k=1}^\infty \frac{|\varepsilon_k|}{n_k}\, \|y_{m_k}\|_Y = \sum_{k=1}^\infty \frac{1}{n_k} = 1.$$

To shorten notation, define $L_k = w_{m_k}^* \circ (A_{m_k} - A) \in Z^*$ and note that the estimates for $A_{m_k} - A$ transfer to L_k. If $\varepsilon_K = 1$ we have

$$\begin{aligned} L_k y &= L_k \left(\sum_{k=1}^{K-1} \frac{\varepsilon_k}{n_k} y_{m_k} \right) + \frac{1}{n_K}\, L_k\, y_{m_K} + L_k \left(\sum_{k=K+1}^\infty \frac{\varepsilon_k}{n_k} y_{m_k} \right) \\ &\ge 0 + \frac{1}{n_K}\, L_k y_{m_K} - C^Y \sum_{l=K+1}^\infty \frac{1}{n_l} \\ &\ge \frac{1}{n_K} \left(c_Y\, m_K^{-\beta} - C^Y\, n_K \sum_{l=k+1}^\infty \frac{1}{n_l} \right) \end{aligned}$$

where the infinite tail of the series is estimated by $\|L_k\|_{L(Y,W)} \le C^Y$ and $\|y_{m_k}\|_Y \le 1$. Similarly if $\varepsilon_K = -1$ we obtain

$$L_k y \le -\frac{1}{n_K} \left(c_Y\, m_K^{-\beta} - C^Y\, n_K \sum_{l=K+1}^\infty \frac{1}{n_l} \right).$$

Slow approximation rate. Choose

$$t_k := \frac{c_Y\, m_k^{\alpha-\beta}}{2\, C_X\, n_k}, \qquad x_k \in \operatorname{argmin}_{\|x\|_X \le t_k} \|x-y\|_Z.$$

Then

$$\begin{aligned}\frac{1}{n_k}\left(c_Y\, m_k^{-\beta} - C^Y\, n_k \sum_{l=k+1}^{\infty} \frac{1}{n_l}\right) &\le \big|L_k y\big| \\ &\le \big|L_k(y-x_k)\big| + \big|L_k x_k\big| \\ &\le C_Z\, \|y-x_k\|_Z + \|A_{m_k} - A\|_{X^*}\, \|x_k\|_X \\ &\le C_Z\, \|y-x_k\|_Z + C_X\, m_k^{-\alpha}\, t_k.\end{aligned}$$

Since t_k was chosen precisely such that

$$C_X m_k^{-\alpha} t_k = \frac{c_Y}{2\, n_k}\, m_k^{-\beta},$$

we obtain that

$$\frac{1}{2C_Z\, n_k}\left(c_Y\, m_k^{-\beta} - 2C^Y\, n_k \sum_{l=k+1}^{\infty} \frac{1}{n_l}\right) \le \|y-x_k\|_Z = \min_{\|x\|_X \le t_k} \|x-y\|_Z. \tag{2.3}$$

For this lower bound to be meaningful, the first term in the bracket has to dominate the second term. We specify the scaling relationship between n_k and m_k as

$$m_k = n_k^{\frac{k}{\alpha-\beta}}.$$

In this definition, m_k is not typically an integer unless $\frac{1}{\alpha-\beta}$ is an integer (or, to hold for a subsequence, rational). In the general case, we choose the integer $\tilde{m}_k$ closest to m_k. To simplify the presentation, we proceed with the non-integer m_k and note that the results are insensitive to perturbations of order 1.

We obtain

$$t_k = \frac{c_Y}{2\, C_X}\frac{m_k^{\alpha-\beta}}{n_k} = \frac{c_Y}{2C_X}\, n_k^{k-1}, \qquad \frac{m_k^{-\beta}}{n_k} = n_k^{-\frac{\beta k}{\alpha-\beta}-1} = n_k^{-\frac{\beta(k-1)+\alpha}{\alpha-\beta}} = \left(\frac{2C_X}{c_Y}\, t_k\right)^{-\frac{\beta}{\alpha-\beta}-\frac{\alpha}{(k-1)(\alpha-\beta)}}.$$

In particular, note that $t_k \to \infty$ as $k \to \infty$. In order for

$$n_k \sum_{l=k+1}^{\infty} \frac{1}{n_l}$$

to be small, we need n_k to grow super-exponentially. Note that $\frac{\beta}{\alpha-\beta} \le 1$ since $\beta \le \frac{\alpha}{2}$. We specify $n_k = 2^{(k^k)}$ and compute

$$\begin{aligned}\sum_{l=k+1}^{\infty} \frac{1}{n_l} = \sum_{l=1}^{\infty} 2^{-\left((k+l)^k\,(k+l)^l\right)} &\le \sum_{l=1}^{\infty} 2^{-(k^k\,(k+l)^l)} = \sum_{l=1}^{\infty}\left(\frac{1}{n_k}\right)^{\left((k+l)^l\right)} \\ \le \frac{2}{n_k^{k+1}} &\ll n_k^{-\frac{\beta k}{\alpha-\beta}-1} = \frac{m_k^{-\beta}}{n_k}\end{aligned}$$

for large enough k. Thus we can neglect the negative term on the left hand side of (2.3) at the price of a slightly smaller constant. Thus

$$\frac{c_Y}{4C_Z}\left(\frac{2C_X}{c_Y}\, t_k\right)^{-\frac{\beta}{\alpha-\beta}-\frac{\alpha}{(k-1)(\alpha-\beta)}} = \frac{c_Y}{4C_Z\, n_k}\, m_k^{-\beta} \le \min_{\|x\|_X \le t_k} \|x-y\|_Z.$$

Finally, we conclude that for all $\gamma > \frac{\beta}{\alpha-\beta}$ we have

$$\limsup_{t\to\infty}\left(t^{\gamma} \inf_{\|x\|_X \le t\}} \|x-y\|_Z\right) \ge \limsup_{k\to\infty}\left(t_k^{\gamma} \inf_{\|x\|_X \le t_k\}} \|x-y\|_Z\right)$$

$$\geq C_{X,Y,Z} \limsup_{k\to\infty} t_k^{\gamma - \frac{\beta}{\alpha-\beta} - \frac{\alpha}{(k-1)(\alpha-\beta)}}$$
$$= \infty.$$

□

3. Approximating Lipschitz Functions by Functions of Low Complexity

In this section, we apply Lemma 2.3 to the situation where general Lipschitz functions are approximated by functions in a space with much lower complexity. Examples include function spaces for infinitely wide neural networks with a single hidden layer and spaces for deep ResNets of bounded width. For simplicity, we first consider uniform approximation and then modify the ideas to also cover L^2-approximation.

3.1. Approximation in L^∞.

Consider the case where

(1) Z is the space of continuous functions on the unit cube $Q = [0,1]^d \subset \mathbb{R}^d$ with the norm

$$\|\phi\|_Z = \sup_{x\in Q} \phi(x),$$

(2) Y is the space of Lipschitz-continuous functions with the norm

$$\|\phi\|_Y = \sup_{x\in Q} \phi(x) + \sup_{x\neq y} \frac{|\phi(x)-\phi(y)|}{|x-y|}, \text{ and}$$

(3) X is a Banach space of functions such that
- X embeds continuously into Z,
- the Monte-Carlo estimate

$$\mathbb{E}_{X_i\sim \mathcal{L}^d|_Q \text{ iid}} \left\{ \sup_{\phi\in B^X} \left[\frac{1}{n}\sum_{i=1}^n \phi(X_i) - \int_Q \phi(x)\ \mathrm{d}x \right] \right\} \leq \frac{C_X}{\sqrt{n}}$$

holds. Here and in the following when no other measure is specified, we assume that integrals are taken with respect to Lebesgue measure.

Examples of admissible spaces for X are Barron space for two-layer ReLU networks and the compositional function space for deep ReLU ResNets of finite width, see [EMW19a, EMW18, EMW19b]. A brief review of Barron space is provided in Appendix A. The Monte-Carlo estimate is proved by estimating the Rademacher complexity of the unit ball in the respective function space. For Barron space, $C_X = 2\sqrt{2\ \log(2d)}$ and for compositional function space $C_X = e^2\ \sqrt{\log(2d)}$, see [EMW19b, Theorems 6 and 12].

We observe the following: If X is a vector of iid random variables sampled from the uniform distribution on Q, then

$$\sup_{\phi \text{ is 1-Lipschitz}} \left(\frac{1}{n}\sum_{i=1}^n \phi(X_i) - \int_Q \phi(x)\ \mathrm{d}x \right) = W_1\left(\mathcal{L}^d|_Q,\ \frac{1}{n}\sum_{i=1}^n \delta_{X_i} \right)$$

is the 1-Wasserstein distance between d-dimensional Lebesgue measure on the cube and the empirical measure generated by the random points – see [Vil08, Chapter 5] for further details on Wasserstein distances and the link between Lipschitz functions and optimal transport theory. The distance on $\mathbb{R}^d$ for which the Wasserstein transportation cost is computed is the same for which ϕ is 1-Lipschitz.

Empirical measures converge to the underlying distribution slowly in high dimension [FG15], by which we mean that

$$\mathbb{E}_{X\sim\left(\mathcal{L}^d|_Q\right)^n} W_1\left(\mathcal{L}^d|_Q,\, \frac{1}{n}\sum_{i=1}^n \delta_{X_i}\right) \geq c_d\, n^{-1/d}$$

for some dimension-dependent constant d. Observe that also

$$\begin{aligned}\sup_{\phi \text{ is 1-Lipschitz}} \left(\frac{1}{n}\sum_{i=1}^n \phi(X_i) - \int_Q \phi(x)\ \mathrm{d}x\right) &= \sup_{\phi \text{ is 1-Lipschitz},\, \phi(0)=0} \left(\frac{1}{n}\sum_{i=1}^n \phi(X_i) - \int_Q \phi(x)\ \mathrm{d}x\right)\\ &\leq \big[1+\operatorname{diam}(Q)\big] \sup_{\|\phi\|_Y\leq 1,\, \phi(0)=0} \left(\frac{1}{n}\sum_{i=1}^n \phi(X_i) - \int_Q \phi(x)\ \mathrm{d}x\right)\\ &\leq \big[1+\operatorname{diam}(Q)\big] \sup_{\|\phi\|_Y\leq 1} \left(\frac{1}{n}\sum_{i=1}^n \phi(X_i) - \int_Q \phi(x)\ \mathrm{d}x\right)\end{aligned}$$

where $\operatorname{diam}(Q)$ is a diameter of the d-dimensional unit cube with respect to the norm for which ϕ is 1-Lipschitz. Here we used that replacing ϕ by $\phi + c$ for $c \in \mathbb{R}$ does not change the difference of the two expectations, and that on the space of functions with $\phi(0) = 0$ the equivalence

$$[\phi]_Y := \sup_{x\neq y} \frac{|\phi(x)-\phi(y)|}{|x-y|} \leq \|\phi\|_Y \leq \big(1+\operatorname{diam}(Q)\big)\,[\phi]_Y$$

holds. By ω_d we denote the Lebesgue measure of the unit ball in $\mathbb{R}^d$ with respect to the correct norm.

Lemma 3.1. *For every $n \in \mathbb{N}$ we can choose n points $x_1, \dots, x_n$ in Q such that*

$$\sup_{\phi\in B^Y} \left[\frac{1}{n}\sum_{i=1}^n \phi(x_i) - \int_Q \phi(x)\ \mathrm{d}x\right] \geq \frac{d}{d+1}\,\frac{1}{\big[(d+1)\,\omega_d\big]^{\frac{1}{d}}}\,\frac{1}{1+\operatorname{diam}(Q)}\ n^{-1/d}$$

and

$$\sup_{\phi\in B^X} \left[\frac{1}{n}\sum_{i=1}^n \phi(x_i) - \int_Q \phi(x)\ \mathrm{d}x\right] \leq \frac{C_X}{n^{1/2}}.$$

Proof. First, we prove the following. *Claim:* Let $x_1, \dots, x_n$ be any collection of n points in Q. Then

$$W_1\left(\mathcal{L}^d|_Q,\, \frac{1}{n}\sum_{i=1}^n \delta_{x_i}\right) \geq \frac{d}{d+1}\,\frac{1}{\big[(d+1)\,\omega_d\big]^{\frac{1}{d}}}\, n^{-1/d}.$$

Proof of claim: Choose $\varepsilon > 0$ and consider the set

$$U = \bigcup_{i=1}^n B_{\varepsilon\, n^{-1/d}}(x_i).$$

We observe that

$$\mathcal{L}^d(U\cap Q) \leq \mathcal{L}^d(U) \leq \sum_{i=1}^n \mathcal{L}^d\big(B_{\varepsilon\, n^{-1/d}}(x_i)\big) = n\,\omega_d\,\big(\varepsilon\, n^{-1/d}\big)^d = \omega_d\,\varepsilon^d.$$

So any transport plan between $\mathcal{L}^d|_Q$ and the empirical measure needs to transport mass $\geq 1 - \omega_d\,\varepsilon^d$ by a distance of at least $\varepsilon\, n^{-1/d}$. We conclude that

$$W_1\left(\mathcal{L}^d|_Q,\, \frac{1}{n}\sum_{i=1}^n \delta_{x_i}\right) \geq \sup_{\varepsilon\in(0,1)} \big(1-\omega_d\,\varepsilon^d\big)\,\varepsilon\, n^{-1/d}.$$

The infimum is attained when

$$0 = 1 - (d+1)\omega_d \varepsilon^d \quad \Leftrightarrow \quad \varepsilon = \left[(d+1)\,\omega_d\right]^{-\frac{1}{d}} \quad \Rightarrow \quad 1 - \omega_d \varepsilon^d = 1 - \frac{1}{d+1} = \frac{d}{d+1}.$$

This concludes the proof of the claim.

Proof of the Lemma: Using the claim, any n points $x_1, \dots, x_n$ such that

$$\sup_{\phi \in B^X} \left[\frac{1}{N} \sum_{i=1}^N \phi(x_i) - \int_Q \phi(x)\ \mathrm{d}x \right] \leq \mathbb{E} \left\{ \sup_{\phi \in B^X} \left[\frac{1}{N} \sum_{i=1}^N \phi(X_i) - \int_Q \phi(x)\ \mathrm{d}x \right] \right\} \leq \frac{C_X}{n^{1/2}}$$

satisfy the conditions. □

For any n, we fix such a collection of points $x_1^n, \dots, x_n^n$ and define

$$A_n : Z \to \mathbb{R}, \quad A_n(\phi) = \frac{1}{n} \sum_{i=1}^n \phi(x_i^n), \qquad A : Z \to \mathbb{R}, \quad A(\phi) = \int_Q \phi(x)\,\mathrm{d}x.$$

Clearly

$$|A\phi|,\ |A_n\phi| \leq \|\phi\|_{C^0} = \|\phi\|_Z.$$

Thus we can apply Lemma 2.3 with

$$\frac{\beta}{\alpha - \beta} = \frac{\frac{1}{d}}{\frac{1}{2} - \frac{1}{d}} = \frac{1}{d\,\frac{d-2}{2d}} = \frac{2}{d-2}.$$

Corollary 3.2. *There exists a 1-Lipschitz function ϕ on Q such that*

$$\limsup_{t \to \infty} \left(t^\gamma \inf_{\|f\|_X \leq t} \|\phi - f\|_{L^\infty(Q)} \right) = \infty.$$

for all $\gamma > \frac{2}{d-2}$.

3.2. **Approximation in L^2.** Point evaluation functionals are no longer well defined if we choose $Z = L^2(Q)$. We therefore need to replace A_n by functionals of the type

$$A_n(\phi) = \frac{1}{n} \sum_{i=1}^n \fint_{B_{\varepsilon_n}(X_i^n)} \phi\,\mathrm{d}x$$

for sample points X_i^n and find a balance between the radii ε_n shrinking too fast (causing the norms $\|A_n\|_{Z^*}$ to blow up) and ε_n shrinking too slowly (leading to better approximation properties on Lipschitz functions).

We interpret Q as the unit cube for function spaces, but as a d-dimensional flat torus when considering balls. Namely the ball $B_\varepsilon(x)$ in Q is to be understood as projection of the ball of radius $\varepsilon > 0$ around $[x]$ on $\mathbb{R}^d/\mathbb{Z}^d$ onto Q. This allows us to avoid boundary effects.

Lemma 3.3. *For every $n \in \mathbb{N}$ we can choose n points $x_1, \dots, x_n$ in Q such that the estimates*

$$\sup_{\phi \in B^X} \left[\frac{1}{n} \sum_{i=1}^n \fint_{B_{\varepsilon_n}(x_i)} \phi\,\mathrm{d}x - \int_Q \phi(x)\ \mathrm{d}x \right] \leq \frac{3\,C_X}{n^{1/2}}$$

$$\sup_{\phi \in B^Y} \left[\frac{1}{n} \sum_{i=1}^n \fint_{B_{\varepsilon_n}(x_i)} \phi\,\mathrm{d}x - \int_Q \phi(x)\ \mathrm{d}x \right] \geq c_d\, n^{-1/d}$$

$$\sup_{\phi \in B^Z} \left[\frac{1}{n} \sum_{i=1}^n \fint_{B_{\varepsilon_n}(x_i)} \phi\,\mathrm{d}x - \int_Q \phi(x)\ \mathrm{d}x \right] \leq C_d$$

hold. c_d, C_d are dimension dependent constants and

$$\varepsilon_n = \gamma_d\, n^{-1/d}$$

for a dimension-dependent $\gamma_d > 0$.

Proof of Lemma 3.3. **L^2-estimate.** In all of the following, we rely on the interpretation of balls as periodic to avoid boundary effects. For a sample $S = (X_1, \dots, X_n)$ denote

$$A_S(\phi) = \frac{1}{n} \sum_{i=1}^{n} \fint_{B_{\varepsilon_n}(X_i)} \phi \, \mathrm{d}x$$

Observe that

$$\begin{aligned}
\sup_{\|\phi\|_{L^2} \le 1} (A_S(\phi) - A(\phi)) &= \sup_{\|\phi\|_{L^2} \le 1} \left[\frac{1}{n} \sum_{i=1}^{n} \fint_{B_{\varepsilon_n}(X_i)} \phi \, \mathrm{d}x - \int \phi \, \mathrm{d}x \right] \\
&= \sup_{\|\phi\|_{L^2} \le 1, \int \phi = 0} \frac{1}{n} \sum_{i=1}^{n} \fint_{B_{\varepsilon_n}(X_i)} \phi \, \mathrm{d}x \\
&= \sup_{\|\phi\|_{L^2} \le 1, \int \phi = 0} \frac{1}{n |B_{\varepsilon_n}|} \int \left(\sum_{i=1}^{n} 1_{B_{\varepsilon_n}(X_i)} \right) \phi \, \mathrm{d}x \\
&\le \frac{1}{n \, \omega_d \, \varepsilon_n^d} \left\| \sum_{i=1}^{n} 1_{B_{\varepsilon_n}(X_i)} \right\|_{L^2}.
\end{aligned}$$

We compute

$$\begin{aligned}
\left\| \sum_{i=1}^{n} 1_{B_{\varepsilon_n}(x_i)} \right\|_{L^2}^2 &= \sum_{i=1}^{n} \left\| 1_{B_{\varepsilon_n}(x_i)} \right\|_{L^2}^2 + \sum_{i \ne j} \int 1_{B_{\varepsilon_n}(x_i)} 1_{B_{\varepsilon_n}(x_j)} \, \mathrm{d}x \\
&= n \, \omega_d |\varepsilon_n|^d + \sum_{i \ne j} \left| B_{\varepsilon_n}(x_i) \cap B_{\varepsilon_n}(x_j) \right| \qquad (3.1)
\end{aligned}$$

It is easy to see that

$$\begin{aligned}
\mathbb{E}_{(X_i, X_j) \sim U_Q \times Q} \left| B_{\varepsilon_n}(X_i) \cap B_{\varepsilon_n}(X_j) \right| &= \mathbb{E}_{X \sim U_Q} \left| B_{\varepsilon_n}(X) \cap B_{\varepsilon_n}(0) \right| \\
&= \int_{B_{2\varepsilon_n}} \left| B_{\varepsilon_n}(x) \cap B_{\varepsilon_n}(0) \right| \mathrm{d}x \\
&= \varepsilon_n^d \int_{B_2} \left| B_{\varepsilon_n}(\varepsilon_n x) \cap B_{\varepsilon_n}(0) \right| \mathrm{d}x \\
&= \varepsilon_n^d \int_{B_2} \varepsilon_n^d \left| B_1(x) \cap B_1(0) \right| \mathrm{d}x \\
&= \varepsilon_n^{2d} \, \omega_d 2^d \, \frac{1}{\omega_d 2^d} \int_{B_2} \left| B_1(x) \cap B_1(0) \right| \mathrm{d}x \\
&= \varepsilon_n^{2d} \, \bar{c}_d \, 2^d \omega_d. \qquad (3.2)
\end{aligned}$$

where

$$\bar{c}_d := \frac{1}{\omega_d 2^d} \int_{B_2} \left| B_1(x) \cap B_1(0) \right| \mathrm{d}x$$

is a dimension-dependent constant. Thus combining (3.1) and (3.2) we find that

$$\begin{aligned}
\mathbb{E}_{S \sim (\mathcal{L}^d|_Q)^n} \left\| \sum_{i=1}^{n} 1_{B_{\varepsilon_n}(X_i)} \right\|_{L^2}^2 &= \omega_d \, \varepsilon_n^d \left[n + n(n-1) \, \bar{c}_d \, (2\varepsilon_n)^d \right] \\
&\le \omega_d \, n \varepsilon_n^d \left[1 + \bar{c}_d 2^d \, n \varepsilon_n^d \right].
\end{aligned}$$

This allows us to estimate

$$\begin{aligned}
\mathbb{E}_S\left[\sup_{\|\phi\|_{L^2}\leq 1}\big(A_S\phi - A\phi\big)\right] &\leq \frac{1}{n\,\omega_d\,\varepsilon_n^d}\mathbb{E}_S\left\|\sum_{i=1}^n 1_{B_{\varepsilon_n}(X_i)}\right\|_{L^2} \\
&\leq \frac{1}{n\,\omega_d\,\varepsilon_n^d}\left(\mathbb{E}_S\left\|\sum_{i=1}^n 1_{B_{\varepsilon_n}(X_i)}\right\|_{L^2}^2\right)^{\frac{1}{2}} \\
&\leq \frac{1}{n\,\omega_d\,\varepsilon_n^d}\sqrt{\omega_d\,n\varepsilon_n^d\big[1+c_d 2^d\,n\varepsilon_n^d\big]} \\
&= \sqrt{\frac{1+\bar{c}_d 2^d\,n\varepsilon_n^d}{\omega_d\,n\varepsilon_n^d}} \\
&= \sqrt{\frac{1+\bar{c}_d 2^d\,\gamma_d^d}{\omega_d\,\gamma_d^d}}
\end{aligned}$$

when we choose

$$\varepsilon_n = \gamma_d\,n^{-1/d}$$

for a dimension-dependent constant γ_d.

Lipschitz estimate. If $E\subset\mathbb{R}^d$ is open and bounded, denote by U_E the uniform distribution on E. Note that

$$W_1\big(\delta_0, U_{B_\varepsilon}\big) = \fint_{B_\varepsilon}|x|\,\mathrm{d}x = \fint_{B_1}\varepsilon|x|\,\mathrm{d}x = \varepsilon\fint_{B_1}|x|\,\mathrm{d}x.$$

Since the set of all transport plans between two measures given as convex combinations of measures is larger than the set of all plans which transport one term of the combination to another, we find that

$$\begin{aligned}
W_1\left(\frac{1}{n}\sum_{i=1}^n\delta_{x_i}, \frac{1}{n}\sum_{i=1}^n U_{B_{\varepsilon_n}(x_i)}\right) &= \inf\left\{\int|x-y|\,\mathrm{d}\pi_{x,y}\;\middle|\;\pi\in\mathcal{P}\left(\frac{1}{n}\sum_{i=1}^n\delta_{x_i}, \frac{1}{n}\sum_{i=1}^n U_{B_{\varepsilon_n}(x_i)}\right)\right\} \\
&\leq \inf\left\{\int|x-y|\,\mathrm{d}\pi_{x,y}\;\middle|\;\pi = \frac{1}{n}\sum_{i=1}^n\pi_i,\;\;\pi_i\in\mathcal{P}\left(\delta_{x_i}, U_{B_{\varepsilon_n}(x_i)}\right)\right\} \\
&= \frac{1}{n}\sum_{i=1}^n W_1(\delta_{x_i}, U_{B_{\varepsilon_n}(x_i)}) \\
&= \varepsilon_n\fint_{B_1}|x|\,\mathrm{d}x \\
&= \gamma_d\,n^{-1/d}\fint_{B_1}|x|\,\mathrm{d}x.
\end{aligned}$$

We find by the triangle inequality that

$$\begin{aligned}
W_1\left(\mathcal{L}^d|_Q, \frac{1}{n}\sum_{i=1}^n U_{B_{\varepsilon_n}(x_i)}\right) &\geq W_1\left(\mathcal{L}^d|_Q, \frac{1}{n}\sum_{i=1}^n\delta_{x_i}\right) - W_1\left(\frac{1}{n}\sum_{i=1}^n\delta_{x_i}, \frac{1}{n}\sum_{i=1}^n U_{B_{\varepsilon_n}(x_i)}\right) \\
&\geq \left[\frac{d}{d+1}\,\frac{1}{\big[(d+1)\,\omega_d\big]^{\frac{1}{d}}} - \gamma_d\fint_{B_1}|x|\,\mathrm{d}x\right]\,n^{-1/d}
\end{aligned}$$

When we choose γ_d small enough, we conclude as before that

$$\sup_{\phi \text{ is 1-Lipschitz}} \left(\frac{1}{n} \sum_{i=1}^{n} \fint_{B_{\varepsilon_n}(x_i)} \phi \, \mathrm{d}x - \int_Q \phi \, \mathrm{d}x \right) \geq c_d \, n^{-1/d}$$

for some positive $c_d > 0$. Recall that this holds for *all* empirical measures, and that the Lipschitz constant is an equivalent norm for our purposes.

X-estimate. We compute that

$$\begin{aligned}
&\mathbb{E}_{X_i \sim \mathcal{L}^d|_Q \text{ iid}} \left\{ \sup_{\phi \in B^X} \left[\frac{1}{n} \sum_{i=1}^{n} \int_{B_{\varepsilon_n}(X_i)} \phi \, \mathrm{d}x - \int_Q \phi(x) \, \mathrm{d}x \right] \right\} \\
&\quad = \mathbb{E}_{X_i \sim \mathcal{L}^d|_Q \text{ iid}} \left\{ \sup_{\phi \in B^X} \left[\frac{1}{n} \sum_{i=1}^{n} \int_{B_{\varepsilon_n}(X_i)} \phi \, \mathrm{d}x - \int_Q \phi(x) \, \mathrm{d}x \right] \right\} \\
&\quad = \mathbb{E}_{X_i \sim \mathcal{L}^d|_Q \text{ iid}} \left\{ \sup_{\phi \in B^X} \left[\frac{1}{n} \sum_{i=1}^{n} \int_{B_{\varepsilon_n}(X_i)} \left(\phi(y) - \int_Q \phi(x) \, \mathrm{d}x \right) \mathrm{d}y \right] \right\} \\
&\quad = \mathbb{E}_{X_i \sim \mathcal{L}^d|_Q \text{ iid}} \left\{ \sup_{\phi \in B^X} \left[\int_{B_{\varepsilon_n}} \frac{1}{n} \sum_{i=1}^{n} \left(\phi(X_i + y) - \int_Q \phi(x) \, \mathrm{d}x \right) \mathrm{d}y \right] \right\} \\
&\quad \leq \mathbb{E}_{X_i \sim \mathcal{L}^d|_Q \text{ iid}} \left\{ \int_{B_{\varepsilon_n}} \sup_{\phi \in B^X} \left[\frac{1}{n} \sum_{i=1}^{n} \left(\phi(X_i + y) - \int_Q \phi(x) \, \mathrm{d}x \right) \right] \mathrm{d}y \right\} \\
&\quad = \mathbb{E}_{X_i \sim \mathcal{L}^d|_Q \text{ iid}} \left\{ \sup_{\phi \in B^X} \left[\frac{1}{n} \sum_{i=1}^{n} \left(\phi(X_i) - \int_Q \phi(x) \, \mathrm{d}x \right) \right] \right\} \\
&\quad \leq \frac{C_X}{\sqrt{n}}
\end{aligned}$$

since X_i and $X_i + y$ have the same law (where $X_i + y$ is interpreted as a shift on the flat torus).

Conclusion. Since the random variables

$$\sup_{\phi \in B^X} \left[\frac{1}{n} \sum_{i=1}^{n} \left(\phi(X_i) - \int_Q \phi(x) \, \mathrm{d}x \right) \right]$$

are non-negative, we find that by Chebyshev's inequality that

$$(\mathcal{L}^d|_Q)^n \left(\left\{ (X_1, \ldots, X_n) \;\middle|\; \sup_{\phi \in B^X} [A_X(\phi) - A(\phi)] > \frac{3C_X}{\sqrt{n}} \right\} \right) \leq \frac{1}{3}$$

and similarly

$$(\mathcal{L}^d|_Q)^n \left(\left\{ (X_1, \ldots, X_n) \;\middle|\; \sup_{\phi \in B^Z} [A_X(\phi) - A(\phi)] > 3\gamma_d \, n^{-1/d} \fint_{B_1} |x| \, \mathrm{d}x \right\} \right) \leq \frac{1}{3}.$$

Since the Y-estimate is satisfied for any empirical measure, we conclude that there exists a set of points $x_1, \ldots, x_n$ in Q such that $A_n := A_{(x_1, \ldots, x_n)}$ satisfies the conditions of the theorem. □

Corollary 3.4. *There exists a 1-Lipschitz function ϕ on Q such that*

$$\limsup_{t \to \infty} \left(t^{\gamma} \inf_{\|f\|_X \leq t} \|\phi - f\|^2_{L^2(Q)} \right) = \infty.$$

for all $\gamma > \frac{4}{d-2}$.

Example 3.5. In [EW20], 'tree-like' function spaces $\mathcal{W}^L$ for fully connected ReLU networks with L hidden layers and bounded path-norm are introduced. There it is shown that the unit ball B^L in $\mathcal{W}^L$ satisfies the Rademacher complexity estimate

$$\mathrm{Rad}(B^L, S) \leq 2^{L+1} \sqrt{\frac{2\,\log(d+2)}{n}}$$

on any sample set of n elements inside $[-1,1]^d$. According to [SSBD14, Lemma 26.2], we have

$$\mathbb{E}_{(x_1,\dots,x_n)\sim\mathbb{P}^n}\left[\sup_{\|h\|_{\mathcal{W}^L}\leq 1}\left|\int_Q h(x)\,\mathbb{P}(\mathrm{d}x) - \frac{1}{n}\sum_{i=1}^n h(x_i)\right|\right] \leq 2\,\mathbb{E}_{(x_1,\dots,x_n)\sim\mathbb{P}^n}\,\mathrm{Rad}\big(B, \{x_1,\dots,x_n\}\big).$$

for any data distribution $\mathbb{P}$ on Q, so $\mathcal{W}^L$ can be chosen as X in this section. A more detailed consideration of Rademacher complexities in the case of Barron functions (for different activation functions) can be found in Appendix A in the proof of Lemma A.10. For ReLU activation, Barron space coincides with the tree-like function space $\mathcal{W}^1$ with one hidden layer.

Corollary 3.6. *Let $X = \mathcal{W}^L(Q)$ the tree-like function space, $Y = C^{0,1}(Q)$ the space of Lipschitz-continuous functions (with respect to the ℓ^∞-norm on $\mathbb{R}^d$) and $Z = L^2(Q)$. Then*

$$\rho(t) = w_{d_Z}(B_t^X, B_1^Y) \geq 2^{-\frac{1}{d}} \frac{\big(c_d/2\big)^{\frac{d}{d-2}}}{C_d\left(3\cdot 2^{L+2}\sqrt{\frac{2\,\log(2d+2)}{n}}\right)^{\frac{2}{d-2}}}\, t^{-\frac{2}{d-2}} = \bar{c}_d\, 2^{-\frac{L}{d-2}}\, t^{-\frac{2}{d-2}}$$

for a dimension-dependent constant $\bar{c}_d > 0$.

In particular, slightly deeper neural networks do not possess drastically larger the approximation power compared in the class of Lipschitz functions.

4. Approximating Two-Layer Neural Networks by Kernel Methods

A brief review of reproducing kernel Hilbert spaces and our notation is given in Appendix B. For bounded kernels, the RKHS $\mathcal{H}_k$ embeds continuously into $L^2(\mathbb{P})$. We assume additionally that $\mathcal{H}_k$ embeds into $L^2(\mathbb{P})$ compactly. In Appendix B, we show that this assumption is met for common random feature kernels and two-layer neural tangent kernels. Compactness allows us to apply the Courant-Hilbert Lemma, which is often used in the eigenvalue theory of elliptic operators.

Lemma 4.1. [Dob10, Satz 8.39] *Let H be a real, separable infinite-dimensional Hilbert space with two symmetric and continuous bilinear forms $B, K : H \times H \to \mathbb{R}$. Assume that*

(1) K is continuous in the weak topology on H,
(2) $K(u,u) > 0$ for all $u \neq 0$, and
(3) B is coercive relative K, i.e.

$$B(u,u) \geq c\,\|u\|_H^2 - \tilde{c}\,K(u,u)$$

for constants $c, \tilde{c} > 0$.

Then the eigenvalue problem

$$B(u, u_i) = \lambda_i\, K(u, u_i) \quad \forall\ u \in H$$

has countably many solutions (λ_i, u_i). Every eigenvalue has finite multiplicity, and if sorted in ascending order then

$$\lim_{i\to\infty} \lambda_i = +\infty.$$

The space spanned by the eigenvectors u_i is dense in H and the eigenvectors satisfy the orthogonality relation

$$K(u_i, u_j) = \delta_{ij}, \qquad B(u_i, u_j) = \lambda_i\, K(u_i, u_j) = \lambda_i \delta_{ij}.$$

We can expand the bilinear forms as

$$B(u, v) = \sum_{i=1}^{\infty} \lambda_i\, K(u_i, u)\, K(u_i, v), \qquad K(u, v) = \sum_{i=1}^{\infty} K(u_i, u)\, K(u_i, v).$$

The pairs (λ_i, u_i) are defined as solutions to the sequence of variational problems

$$\lambda_i = B(u_i, u_i) = \inf \left\{ \frac{B(u, u)}{K(u, u)} \;\middle|\; K(u, u_j) = 0 \quad \forall\, 1 \le j \le i-1 \right\}.$$

Consider $H = L^2(\mathbb{P})$ and

$$B(u, v) = \langle u, v \rangle_{L^2(\mathbb{P})}, \qquad K(u, v) = \int_{\mathbb{R}^d \times \mathbb{R}^d} u(x)\, v(x')\, k(x, x')\, \mathbb{P}(\mathrm{d}x)\, \mathbb{P}(\mathrm{d}x') = \langle \overline{K} u, v \rangle_{L^2(\mathbb{P})}$$

where

$$\overline{K} u(x) = \int_{\mathbb{R}^d} k(x, x')\, u(x')\, \mathbb{P}(\mathrm{d}x').$$

It is easy to see that the assumptions of the Courant-Hilbert Lemma are indeed satisfied. In particular, note that

$$K(u, u_i) = \mu_i\, B(u, u_i) \quad \forall\, u \in L^2(\mathbb{P}) \qquad \Leftrightarrow \qquad \overline{K} u_i = \mu_i\, u_i.$$

By definition, all eigenfunctions lie in the reproducing kernel Hilbert space.

Let $X = \mathcal{H}_k$ be a suitable RKHS, $Y = \mathcal{B}(\mathbb{P})$ Barron space (see Appendix A for a brief review), $Z = W = L^2(\mathbb{P})$. Then $X, Y \hookrightarrow Z$. Furthermore, consider the sequence of n-dimensional spaces

$$X_n = \operatorname{span}\{u_1, \dots, u_n\} \subseteq X$$

spanned by the first n eigenfunctions of $\overline{K}$ and the maps

$$A_n : Z \to W, \qquad A_n = P_{X_n}, \qquad A = \mathrm{id}_{L^2(\mathbb{P})}$$

where P_V denotes the Z-orthogonal projection onto the subspace V. Due to the orthogonality statement in the Courant-Hilbert Lemma, P_V is also the X-orthogonal projection for this specific sequence of spaces.

Lemma 4.2. *If $\mathbb{P} = \mathcal{L}^d|_{[0,1]^d}$, we have the estimates*

$$\|A_n - A\|_{L(X,W)} \le \frac{1}{\lambda_{n+1}^{1/2}}, \qquad \|A_n - A\|_{L(Y,W)} \ge \frac{c}{d}\, n^{-1/d}, \qquad \|A_n - A\|_{L(Z,W)} \le 1$$

where $c > 0$ is a universal constant.

Proof of Lemma 4.2. **X-estimate.** Normalize the eigenfunctions f_i of k to be $L^2(\mathbb{P})$-orthonormal. For $f \in \mathcal{H}_{k_\pi}$ we have the expansion $f = \sum_{i=1}^{\infty} a_i\, f_i$ and thus $(A - A_n) f = \sum_{i=n+1}^{\infty} a_i\, f_i$ such that

$$\|(A_n - A) f\|_W^2 = \sum_{i=n+1}^{\infty} |a_i|^2 \le \sum_{i=n+1}^{\infty} \frac{\lambda_i}{\lambda_{n+1}} |a_i|^2 = \frac{\|A_n f\|_X^2}{\lambda_{n+1}} \le \frac{\|f\|_X}{\lambda_{n+1}} \le \frac{1}{\lambda_{n+1}}$$

Y-estimate. See [Bar93, Theorem 6]. There it is shown that *any* sequence of n-dimensional spaces suffers from the curse of dimensionality when approximating a subset of Barron space.

Z-estimate. The orthogonal projection $A - A_n = P_{X_n^\perp}$ has norm one. □

Thus, if an RKHS has rapidly decreasing eigenvalues independently of the dimension of the ambient space (which is favourable from the perspective of statistical learning theory), then it suffers from a slow approximation property.

Example 4.3. In Appendix B, we give examples of random feature kernels and neural tangent kernels for which $\lambda_k \leq c_d\, k^{-\frac{1}{2}+\frac{3}{2d}}$. In Appendix A, we briefly discuss Barron space. Applying Lemma 2.1 with

$$\alpha = \frac{1}{4} - \frac{3}{4d}, \qquad \beta = \frac{1}{d} \qquad \Rightarrow \qquad \frac{\beta}{\alpha-\beta} = \frac{\frac{1}{d}}{\frac{d}{4d} - \frac{3}{4d} - \frac{4}{4d}} = \frac{4}{d-7},$$

we see that the L^2-width of the RKHS in Barron space is bounded from below by

$$\rho(t) := \sup_{\|\phi\|_{\mathcal{B}(\mathbb{P})}\leq 1} \inf_{\|\psi\|_{\mathcal{H}_{k,\mathbb{P}}}\leq t} \|\phi-\psi\|_{L^2(\mathbb{P})} \geq c_d t^{-\frac{\beta}{\alpha-\beta}} = c_d\, t^{-\frac{4}{d-10}}.$$

Due to Lemma 2.3, there exists a function ϕ in Barron space such that

$$\limsup_{t\to\infty} \left(t^\gamma \cdot \inf_{\|\psi\|_{\mathcal{H}_{k,\mathbb{P}}}\leq t} \|\phi-\psi\|_{L^2(\mathbb{P})} \right) = \infty$$

for all $\gamma > \frac{4}{d-7}$.

5. Discussion

From the viewpoint of functional analysis and more precisely function spaces, a fundamental task in machine learning is balancing the approximation and estimation errors of a hypothesis class. In overly expressive function classes, it may be difficult to assess the performance of a function from a small data sample, whereas too restrictive function classes lack the expressivity to perform well in many problems. In this article, we made a first step in trying to quantify the competition between estimation and approximation. Our results show that linear function classes in which the estimation error is strongly controlled (including, but not limited to those developed for infinitely wide neural networks), the approximation error must suffer from the curse of dimensionality within the class of Lipschitz-functions. Additionally, we emphasize that kernel methods (including some neural tangent kernels) are subject to the curse of dimensionality where adaptive methods like shallow neural networks are not.

In a companion article [WE20], we show that the Barron norm and the RKHS norm increase at most linearly in time during gradient flow training in the mean field scaling regime. This means that the L^2-population risk of a shallow neural network or kernel function can only decay like $t^{-\alpha_d}$ for general Lipschitz or Barron target functions respectively, where α_d is close to zero in high dimension. It is therefore of crucial importance to understand the function spaces associated with neural network architectures under the natural path norms.

While the theory of function spaces for low-dimensional analysis (Sobolev, Besov, BV, BD, etc.) is well studied, the spaces for high-dimensional (but not infinite-dimensional) analysis is at its very beginning. To the best of our knowledge, the currently available models for neural networks consider infinitely wide two-layer or multi-layer networks or infinitely deep networks with bounded width [EMW19c, EW20, EMW19b]. A different perspective on the approximation spaces of deep networks focussing on the number of parameters (but not their size) is developed in [GKNV19].

Even for existing function spaces, it is hard to check whether a given function belongs to the space. Barron's original work [Bar93] shows (in modern terms) that every sufficiently smooth function on an extension domain belongs to Barron space, where the required degree of smoothness depends on the dimension. More precisely, if f is a function on $\mathbb{R}^d$ such that its Fourier

transform $\hat{f}$ satisfies

$$C_f := \int_{\mathbb{R}^d} \left|\hat{f}(\xi)\right| |\xi| \, \mathrm{d}\xi < \infty,$$

then for every compact set $K \subset \mathbb{R}^d$ there exists a Barron function g such that $f \equiv g$ on $\mathbb{R}^d$. In particular, if $f \in H^s(\mathbb{R}^d)$ where $s > \frac{d}{2} + 1$, then $C_f < \infty$. In particular, if $f \in C^k(\Omega)$ for $k > \frac{d}{2} + 1$ and Ω has smooth boundary, then by standard extension results we see that $f \in \mathcal{B}$. This holds for networks with *any* sigmoidal activation function or ReLU activation. The criterion is unsatisfying in two ways:

(1) The function f has to be defined on the whole space for Fourier-analytic considerations to apply. Given a function defined on a compact set, one has to find a good extension to the whole space.
(2) The constant C_f merely gives an upper bound on the Barron-norm of a two layer network. If $f \notin C^1$, then $C_f = +\infty$. If $f(x) = \sum_{i=1}^m a_i \operatorname{ReLU}(w_i^T x + b_i)$ is a two-layer neural network with finitely many nodes, then f is only Lipschitz continuous and not C^1-smooth (unless it is linear). Thus the criterion misses many functions of practical importance.

5.1. **Open Problems.** Many questions in this field remain open.

(1) The slow approximation property in L^∞ is based purely on the slow convergence of empirical measures, while the L^2-construction also uses the translation-invariance of Lebesgue measure for convenience. Does a 'curse of dimensionality' type phenomenon affect L^2-approximation when $\mathbb{P}$ has a density with respect to Lebesgue measure, or more generally is a regular measure concentrated on or close to a high-dimensional manifold in an even higher-dimensional ambient space?
(2) We used Lipschitz functions for convenience, but we believe that a similar phenomenon holds for C^k functions for any fixed k which does not scale with dimension. To apply the same approach, we need to answer how quickly the 1-Wasserstein-type distances

$$\widetilde{W}_{C^k,\lambda}(\mu,\nu) = \sup\left\{ \int f\,\mu(\mathrm{d}x) - \int f\,\nu(\mathrm{d}x) \;\middle|\; f \text{ is 1-Lipschitz, } |D^k f|_\infty \le \lambda \right\}$$

decay in expectation when $\nu = \mu_n$ is an empirical measure sampled iid from μ. Other concepts of Wasserstein-type distance would lead to similar results.
(3) To prove the curse of dimensionality phenomenon, we used a multi-scale construction with modifications on quickly diverging scales. The statement about the upper limit is only a 'worst case curse of dimensionality' and describes slow convergence on an infinite set of vastly different scales. Replacing the upper limit in Theorem 2.3 by a lower limit would lead to a much stronger statement with more severe implications for applications.
(4) Our proof of slow approximation was purely functional analytic and abstract. Is it possible to give a concrete example of a Lipschitz function which is poorly approximated by Barron functions of low norm in $L^2(\mathbb{P})$ for a suitable data measure $\mathbb{P}$?

More generally, is it possible to find general criteria to establish how well a given Lipschitz function (or even a given collection of data) can be approximated by a certain network architecture?
(5) We proved that a Lipschitz function ϕ exists for which

$$\inf_{\|f\|_{\mathrm{Barron}} \le t} \|f - \phi\|_{L^2} \ge c_d\, t^{-1/d}$$

on a suitable collection of scales t_k. Can we more generally characterize the class

$$Y_\alpha := \left\{ y \in Y \;\middle|\; \limsup_{t\to\infty} \left[t^\alpha \mathrm{dist}_Z(y, t\cdot B^X)\right] < \infty \right\}$$

for $\alpha > 0$, X = Barron space, Y = Lipschitz space and $Z = L^2$? This question arises naturally when considering training algorithms which increase the complexity-controlling norm only slowly. It is related, but not identical to considerations in the theory of real interpolation spaces.

(6) The draw-back of the functional analytic approach of this article is that it does not encompass many common loss functionals such as logistic loss. Is it possible to show by different means that they are subject to similar problems in high dimension?

(7) We give relevant examples of reproducing kernel Hilbert spaces in which the eigenvalues of the kernel decay at a dimension-independent rate (at leading order). To the best of our knowledge, a general perspective on the decay of eigenvalues of kernels in practical applications without strong symmetry assumptions is still missing.

Appendix A. A Brief Review of Barron Space

For the convenience of the reader, we recall Barron space for two-layer neural networks as introduced by E, Ma and Wu [EMW19b, EMW18]. We focus on the functional analytic properties of Barron space, for results with a focus on machine learning we refer the reader to the original sources. The same space is denoted as $\mathcal{F}_1$ in [Bac17], but described from a different perspective. For functional analytic notions, we refer the reader to [Bre11].

Let $\mathbb{P}$ be a probability measure on $\mathbb{R}^d$ and σ a Lipschitz-continuous function such that either

(1) $\sigma = \mathrm{ReLU}$ or
(2) σ is sigmoidal, i.e. $\lim_{z\to\pm\infty} \sigma(z) = \pm 1$ (or 0 and 1).

Consider the class $\mathcal{F}_m$ of two-layer networks with m neurons

$$f_\Theta(x) = \frac{1}{m}\sum_{i=1}^m a_i\,\sigma\big(w_i^T x + b_i\big), \qquad \Theta = \{(a_i, w_i, b_i) \in \mathbb{R}^{d+2}\}_{i=1}^m.$$

It is well-known that the closure of $\mathcal{F} = \bigcup_{m=1}^\infty \mathcal{F}_m$ in the uniform topology is the space of continuous functions, see e.g. [Cyb89]. Barron space is a different closure of the same function class where the path-norm

$$\|f_\Theta\|_{\mathrm{path}} = \frac{1}{m}\sum_{i=1}^m |a_i|\,\big[|w_i|_{\ell^q} + |b_i|\big]$$

remains bounded. Here we assume that data space $\mathbb{R}^d$ is equipped with the ℓ^p-norm and take the dual ℓ^q-norm on w. The concept of path norm corresponds to ReLU activation, a slightly different path norm for bounded Lipschitz activation is discussed below.

The same class is often discussed without the normalizing factor of $\frac{1}{m}$. With the factor, the following concept of infinitely wide two-layer networks emerges more naturally.

Definition A.1. Let π be a Radon probability measure on $\mathbb{R}^{d+2}$ with finite second moments, which we denote by $\pi \in \mathcal{P}_2(\mathbb{R}^{d+2})$. We denote by

$$f_\pi(x) = \int_{\mathbb{R}^{d+2}} a\,\sigma(w^T x + b)\,\pi(\mathrm{d}a \otimes \mathrm{d}w \otimes \mathrm{d}b)$$

the two-layer network associated to π.

If $\sigma = \mathrm{ReLU}$, it is clear that f_π is Lipschitz-continuous on $\mathbb{R}^d$ with Lipschitz-constant $\leq \|f_\pi\|_{\mathcal{B}(\mathbb{P})}$, so f_π lies in the space of (possibly unbounded) Lipschitz functions $\mathrm{Lip}(\mathbb{R}^d)$. If σ is a bounded Lipschitz function, the integral converges in $C^0(\mathbb{R}^d)$ without assumptions on the moments of π. If the second moments of π are bounded, f_π is a Lipschitz function also in the case of bounded Lipschitz activation σ.

For technical reasons, we will extend the definition to distributions π for which only the *mixed second moments*

$$\int_{\mathbb{R}^{d+2}} |a| \left[|w| + \psi(b)\right] \pi(\mathrm{d}a \otimes \mathrm{d}w \otimes \mathrm{d}b)$$

are finite, where $\psi = |b|$ in the ReLU case and $\psi = 1$ otherwise. Now we introduce the associated function space.

Definition A.2. We denote

$$\|\cdot\|_{\mathcal{B}(\mathbb{P})} : \mathrm{Lip}(\mathbb{R}^d) \to [0,\infty), \qquad \|f\|_{\mathcal{B}(\mathbb{P})} = \inf_{\{\pi \,|\, f_\pi = f \ \mathbb{P}-\mathrm{a.e.}\}} \int_{\mathbb{R}^{d+2}} |a| \left[|w|_{\ell^q} + |b|\right] \pi(\mathrm{d}a \otimes \mathrm{d}w \otimes \mathrm{d}b)$$

if $\sigma = \mathrm{ReLU}$ and

$$\|\cdot\|_{\mathcal{B}(\mathbb{P})} : \mathrm{Lip}(\mathbb{R}^d) \to [0,\infty), \qquad \|f\|_{\mathcal{B}(\mathbb{P})} = \inf_{\{\pi \,|\, f_\pi = f \ \mathbb{P}-\mathrm{a.e.}\}} \int_{\mathbb{R}^{d+2}} |a| \left[|w|_{\ell^q} + 1\right] \pi(\mathrm{d}a \otimes \mathrm{d}w \otimes \mathrm{d}b)$$

otherwise. In either case, we denote

$$\mathcal{B}(\mathbb{P}) = \left\{ f \in \mathrm{Lip}(\mathbb{R}^d) : \|f\|_{\mathcal{B}(\mathbb{P})} < \infty \right\}.$$

Here $\inf \emptyset = +\infty$.

Remark A.3. It depends on the activation function σ whether or not the infimum in the definition of the norm is attained. If $\sigma = \mathrm{ReLU}$, this can be shown to be true by using homogeneity. Instead of a probability measure π on the whole space, one can use a signed measure μ on the unit sphere to express a two layer network. The compactness theorem for Radon measures provides the existence of a measure minimizer, which can then be lifted to a probability measure (see below for similar arguments). On the other hand, if σ is a classical sigmoidal function such that

$$\lim_{z \to -\infty} \sigma(z) = 0, \qquad \lim_{x \to \infty} \sigma(z) = 1, \qquad 0 < \sigma(z) < 1 \quad \forall\, z \in \mathbb{R},$$

then the function $f(z) \equiv 1$ has Barron norm 1, but the infimum is not attained. This holds true for *any* data distribution $\mathbb{P}$.

Proof.

(1) For any $x \in \mathrm{spt}(\mathbb{P})$ we have

$$1 = f(x) = \int_{\mathbb{R}^{d+2}} a\, \sigma(w^T x + b)\, \mathrm{d}\pi \le \int_{\mathbb{R}^{d+2}} |a|\, \mathrm{d}\pi \le \int_{\mathbb{R}^{d+2}} |a| \left[|w| + 1\right] \mathrm{d}\pi.$$

Taking the infimum over all π, we find that $1 \le \|f\|_{\mathcal{B}(\mathbb{P})}$. For any measure π, the inequality above is strict since $|\sigma| < 1$, so there is no π which attains equality.

(2) We consider a family of measures

$$\pi_\lambda = \delta_{a=\sigma(\lambda)^{-1}}\, \delta_{w=0}\, \delta_{b=\lambda} \qquad \Rightarrow \qquad f_{\pi_\lambda} \equiv 1 \quad \text{and} \quad \int_{\mathbb{R}^{d+2}} |a| \left[|w|_{\ell^q} + 1\right] \mathrm{d}\pi_\lambda = \frac{1}{\sigma(\lambda)} \to 1$$

as $\lambda \to \infty$.

Thus $\|f\|_{\mathcal{B}(\mathbb{P})} = 1$, but there is no minimizing parameter distribution π.

Remark A.4. The space $\mathcal{B}(\mathbb{P})$ does not depend on the measure $\mathbb{P}$, but only on the system of null sets for $\mathbb{P}$.

We note that the space $\mathcal{B}(\mathbb{P})$ is reasonably well-behaved from the point of view of functional analysis.

Lemma A.5. *$\mathcal{B}(\mathbb{P})$ is a Banach space with norm $\|\cdot\|_{\mathcal{B}(\mathbb{P})}$. If $\mathrm{spt}(\mathbb{P})$ is compact, $\mathcal{B}(\mathbb{P})$ embeds continuously into the space of Lipschitz functions $C^{0,1}(\mathrm{spt}\,\mathbb{P})$.*

Proof. **Scalar multiplication.** Let $f \in \mathcal{B}(\mathbb{P})$. For $\lambda \in \mathbb{R}$ and $\pi \in \mathcal{P}_2(\mathbb{R}^{d+2})$, define the push-forward

$$T_{\lambda\sharp}\pi \in \mathcal{P}_2(\mathbb{R}^{d+2}) \qquad \text{along } T_\lambda : \mathbb{R}^{d+2} \to \mathbb{R}^{d+2}, \quad T_\lambda(a,w,b) = (\lambda a, w, b).$$

Then

$$f_{T_{\lambda\sharp}\pi} = \lambda\, f_\pi, \qquad \int_{\mathbb{R}^{d+2}} |a| \left[|w|+1\right] T_{\lambda\sharp}\pi(\mathrm{d}a \otimes \mathrm{d}w \otimes \mathrm{d}b) = |\lambda| \int_{\mathbb{R}^{d+2}} |a| \left[|w|+1\right] \pi(\mathrm{d}a \otimes \mathrm{d}w \otimes \mathrm{d}b)$$

and similarly in the ReLU case. Thus scalar multiplication is well-defined in $\mathcal{B}(\mathbb{P})$. Taking the infimum over π, we find that $\|\lambda f\|_{\mathcal{B}(\mathbb{P})} = |\lambda|\, \|f\|_{\mathcal{B}(\mathbb{P})}$.

Vector addition. Let $g, h \in \mathcal{B}(\mathbb{P})$. Choose π_g, π_h such that $g = f_{\pi_g}$ and $h = f_{\pi_h}$. Consider

$$\pi = \frac{1}{2}\left[T_{2\sharp}\pi_g + T_{2\sharp}\pi_h\right]$$

like above. Then $f_\pi = g + h$ and

$$\int_{\mathbb{R}^{d+2}} |a| \left[|w|+1\right] \pi(\mathrm{d}a \otimes \mathrm{d}w \otimes \mathrm{d}b) = \int_{\mathbb{R}^{d+2}} |a| \left[|w|+1\right] \pi_g(\mathrm{d}a \otimes \mathrm{d}w \otimes \mathrm{d}b) + \int_{\mathbb{R}^{d+2}} |a| \left[|w|+1\right] \pi_h(\mathrm{d}a \otimes \mathrm{d}w \otimes \mathrm{d}b).$$

Taking infima, we see that $\|g+h\|_{\mathcal{B}(\mathbb{P})} \leq \|g\|_{\mathcal{B}(\mathbb{P})} + \|h\|_{\mathcal{B}(\mathbb{P})}$. The same holds in the ReLU case.

Positivity and embedding. Recall that the norm on the space of Lipschitz functions on a compact set K is

$$\|f\|_{C^{0,1}(K)} = \sup_{x \in K} |f(x)| + \sup_{x,y \in K,\, x \neq y} \frac{|f(x) - f(y)|}{|x-y|}.$$

It is clear that $\|\cdot\|_{\mathcal{B}(\mathbb{P})} \geq 0$. If σ is a bounded Lipschitz function, then

$$\|f_\pi\|_{L^\infty(\operatorname{spt}\mathbb{P})} \leq \|f_\pi\|_{\mathcal{B}(\mathbb{P})}\, \|\sigma\|_{L^\infty(\mathbb{R})}$$

like in Remark A.3. If $\sigma = \mathrm{ReLU}$, then

$$\sup_{x \in \mathbb{R}^d} \frac{|f_\pi(x)|}{1+|x|} \leq \|f_\pi\|_{\mathcal{B}(\mathbb{P})}.$$

In either case

$$|f_\pi(x) - f_\pi(y)| \leq [\sigma]_{\mathrm{Lip}}\, \|f_\pi\|_{\mathcal{B}(\mathbb{P})}$$

for all x, y in $\operatorname{spt}(\mathbb{P})$. In particular, $\|f\|_{\mathcal{B}(\mathbb{P})} > 0$ whenever $f \neq 0$ in $\mathcal{B}(\mathbb{P})$ and if $\operatorname{spt}(\mathbb{P})$ is compact, $\mathcal{B}(\mathbb{P})$ embeds into the space of Lipschitz functions on $\mathcal{B}(\mathbb{P})$. If σ is a bounded Lipschitz function, $\mathcal{B}(\mathbb{P})$ embeds into the space of bounded Lipschitz functions also on unbounded sets.

Completeness. Completeness is proved most easily by introducing a different representation for Barron functions. Consider the space

$$V = \left\{ \mu \,\middle|\, \mu \text{ (signed) Radon measure on } \mathbb{R}^{d+2} \text{ s.t. } \int_{\mathbb{R}^{d+2}} |a| \left[|w| + |b|\right] |\mu|(\mathrm{d}a \otimes \mathrm{d}w \otimes \mathrm{d}b) < \infty \right\}$$

where $|\mu|$ is the total variation measure of μ. Equipped with the norm

$$\|\mu\|_V = \int_{\mathbb{R}^{d+2}} |a| \left[|w| + |b|\right] |\mu|(\mathrm{d}a \otimes \mathrm{d}w \otimes \mathrm{d}b),$$

V is a Banach space when we quotient out measures supported on $\{|a| = 0\} \cup \{|b| = |w| = 0\}$ or restrict ourselves to the subspace of measures such that $|\mu|(\{|a| = 0\}) = |\mu|(\{|b| = |w| = 0\}) = 0$. The only non-trivial question is whether V is complete. By definition μ_n is a Cauchy sequence in V if and only if $\nu_n := |a| \left[|w| + |b|\right] \cdot \mu_n$ is a Cauchy sequence in the space of finite Radon measures. Since the space of finite Radon measures is complete, ν_n converges (strongly) to a measure ν

which satisfies $\nu(\{a=0\}\cup\{(w,b)=0\})=0$. We then obtain $\mu := |a|^{-1}\left[|w|+|b|\right]^{-1}\cdot\nu$. For $\mu\in V$ we write

$$f_\mu(x)=\int_{\mathbb{R}^{d+2}} a\,\sigma(w^T x+b)\,\mu(\mathrm{d}a\otimes\mathrm{d}w\otimes\mathrm{d}b)$$

and consider the subspace

$$V^0_{\mathbb{P}}=\{\mu\in V\mid f_\mu=0\ \ \mathbb{P}-\text{almost everywhere}\}.$$

Since the map

$$V\mapsto C^{0,1}(\operatorname{spt}\mathbb{P}),\qquad \mu\mapsto f_\mu$$

is continuous by the same argument as before, we find that $V^0_{\mathbb{P}}$ is a closed subspace of V. In particular, $V/V^0_{\mathbb{P}}$ is a Banach space. We claim that $\mathcal{B}(\mathbb{P})$ is isometric to the quotient space $V/V^0_{\mathbb{P}}$ by the map $[\mu]\mapsto f_\mu$ where μ is any representative in the equivalence class $[\mu]$.

It is clear that any representative in the equivalence class induces the same function f_μ such that the map is well-defined. Consider the Hahn decomposition $\mu=\mu^+-\mu^-$ of μ as the difference of non-negative Radon measures. Set

$$m^\pm := \|\mu^\pm\|,\qquad \pi=\frac{1}{2}\left[\frac{1}{m^+}T_{2m^+\sharp}\mu^+ + \frac{1}{m^-}T_{-2m^-\sharp}\mu^-\right].$$

Then π is a probability Radon measure such that $f_\pi=f_\mu$ and

$$\int_{\mathbb{R}^{d+2}}|a|\left[|w|+|b|\right]\pi(\mathrm{d}a\otimes\mathrm{d}w\otimes\mathrm{d}b)=\int_{\mathbb{R}^{d+2}}|a|\left[|w|+|b|\right]|\mu|(\mathrm{d}a\otimes\mathrm{d}w\otimes\mathrm{d}b).$$

In particular, $\|f_\mu\|_{\mathcal{B}(\mathbb{P})}\le\|\mu\|_V$. Taking the infimum of the right hand side, we conclude that $\|f_\mu\|_{\mathcal{B}(\mathbb{P})}\le\|[\mu]\|_{V/V^0_{\mathbb{P}}}$. The opposite inequality is trivial since every probability measure is in particular a signed Radon measure.

Thus $\mathcal{B}(\mathbb{P})$ is isometric to a Banach space, hence a Banach space itself. We presented the argument in the context of ReLU activation, but the same proof holds for bounded Lipschitz activation. □

A few remarks are in order.

Remark A.6. The requirement that $\mathbb{P}$ have compact support can be relaxed when we consider the norm

$$\|f\|_{C^{0,1}(\mathbb{P})}=\sup_{x,y\in\operatorname{spt}\mathbb{P},\,x\neq y}\frac{|f(x)-f(y)|}{|x-y|}+\|f\|_{L^1(\mathbb{P})}$$

on the space of Lipschitz functions. Since Lipschitz functions grow at most linearly, this is well-defined for all data distributions $\mathbb{P}$ with finite first moments.

Remark A.7. For general Lipschitz-activation σ which is neither bounded nor ReLU, the Barron norm is defined as

$$\|f\|_{\mathcal{B}(\mathbb{P})}=\inf_{\{\pi\mid f_\pi=f\ \mathbb{P}-\text{a.e.}\}}\int_{\mathbb{R}^{d+2}}|a|\left[|w|_{\ell^q}+|b|+1\right]\pi(\mathrm{d}a\otimes\mathrm{d}w\otimes\mathrm{d}b).$$

Similar results hold in this case.

Remark A.8. In general, $\mathcal{B}(\mathbb{P})$ for ReLU activation is not separable. Consider $\mathbb{P}=\mathcal{L}^1|_{[0,1]}$ to be Lebesgue measure on the unit interval in one dimension. For $\alpha\in(0,1)$ set $f_\alpha(x)=\sigma(x-\alpha)$. Then for $\beta>\alpha$ we have

$$1=\frac{\beta-\alpha}{\beta-\alpha}=\frac{(f_\alpha-f_\beta)(\beta)-(f_\alpha-f_\beta)(\alpha)}{\beta-\alpha}\le[f_\beta-f_\alpha]_{\mathrm{Lip}}\le\|f_\beta-f_\alpha\|_{\mathcal{B}(\mathbb{P})}.$$

Thus there exists an uncountable family of functions with distance ≥ 1, meaning that $\mathcal{B}(\mathbb{P})$ cannot be separable.

Remark A.9. In general, $\mathcal{B}(\mathbb{P})$ for ReLU activation is not reflexive. We consider $\mathbb{P}$ to be the uniform measure on $[0,1]$ and demonstrate that $\mathcal{B}(\mathbb{P})$ is the space of functions whose first derivative is in BV (i.e. whose second derivative is a Radon measure) on $[0,1]$. The space of Radon measures on $[0,1]$ is denoted by $\mathcal{M}[0,1]$ and equipped with the total variation norm.

Assume that f is a Barron function on $[0,1]$. Then

$$\begin{aligned} f(x) &= \int_{\mathbb{R}^3} a\,\sigma(wx+b)\,\pi(\mathrm{d}a\otimes\mathrm{d}w\otimes\mathrm{d}b) \\ &= \int_{\{w\neq 0\}} a\,|w|\,\sigma\left(\frac{w}{|w|}x+\frac{b}{|w|}\right)\,\pi(\mathrm{d}a\otimes\mathrm{d}w\otimes\mathrm{d}b) + \int_{\{w=0\}} a\,\sigma(b)\pi(\mathrm{d}a\otimes\mathrm{d}w\otimes\mathrm{d}b) \\ &= \int_{\mathbb{R}} \sigma\left(-x+\tilde{b}\right)\,\mu_1(\mathrm{d}\tilde{b}) + \int_{\mathbb{R}} \sigma\left(x+b\right)\,\mu_2(\mathrm{d}\tilde{b}) + \int_{\{w=0\}} a\,\sigma(b)\pi(\mathrm{d}a\otimes\mathrm{d}w\otimes\mathrm{d}b) \end{aligned}$$

where

$$\mu_1 = T_\sharp\big(a\,|w|\cdot\pi\big), \qquad T:\{w<0\}\to\mathbb{R}, \quad T(a,w,b)=\frac{b}{|w|}.$$

and similarly for μ_2. The Barron norm is expressed as

$$\|f\|_{\mathcal{B}(\mathbb{P})} = \inf_{\mu_1,\mu_2}\left[\int_{\mathbb{R}} 1+|\tilde{b}|\,|\mu_1|(\mathrm{d}\tilde{b}) + \int_{\mathbb{R}} 1+|\tilde{b}|\,|\mu_2|(\mathrm{d}\tilde{b})\right] + \int_{\{w=0\}} |ab|\,\pi(\mathrm{d}a\otimes\mathrm{d}w\otimes\mathrm{d}b).$$

Since $\sigma''=\delta$, we can formally calculate that

$$f''=\mu_1+\mu_2$$

This is easily made rigorous in the distributional sense. Since $[0,1]$ is bounded by 1, we obtain in addition to the bounds on $f(0)$ and the Lipschitz constant of f that

$$\|f''\|_{\mathcal{M}[0,1]} \leq 2\,\|f\|_{\mathcal{B}(\mathbb{P})}.$$

On the other hand, if f has a second derivative, then

$$\begin{aligned} f(x) &= f(0)+f'(0)\,x + \int_0^x\int_0^t f''(\xi)\mathrm{d}\xi\,\mathrm{d}t \\ &= f(0)+f'(0)\,x + \int_0^x (x-t)\,f''(t)\mathrm{d}t \\ &= f(0)\,\sigma(1)+f'(0)\,\sigma(x) + \int_0^x f''(t)\,\sigma(x-t)\,\mathrm{d}t \end{aligned}$$

for all $x\in[0,1]$. This easily extends to measure valued derivatives and we conclude that

$$\|f\|_{\mathcal{B}(\mathbb{P})} \leq |f(0)|+|f'(0)|+\|f''\|_{\mathcal{M}[0,1]}.$$

We can thus express $\mathcal{B}(\mathbb{P}) = BV([0,1])\times\mathbb{R}\times\mathbb{R}$ with an equivalent norm

$$\|f\|'_{\mathcal{B}(\mathbb{P})} = |f(0)|+|f'(0)|+\|f''\|_{\mathcal{M}[0,1]}.$$

Thus $\mathcal{B}(\mathbb{P})$ is not reflexive since BV is not reflexive.

Finally, we demonstrate that integration by empirical measures converges quickly on $\mathcal{B}(\mathbb{P})$. Assume that $p=\infty$, $q=1$.

Lemma A.10. *The uniform Monte-Carlo estimate*

$$\mathbb{E}_{X_i\sim\mathbb{P}}\left\{\sup_{\|\phi\|_{\mathcal{B}(\mathbb{P})}\leq 1}\left[\frac{1}{n}\sum_{i=1}^n \phi(X_i) - \int_Q \phi(x)\,\mathbb{P}(\mathrm{d}x)\right]\right\} \leq 2L\sqrt{\frac{2\,\log(2d)}{n}}$$

holds for any probability distribution $\mathbb{P}$ *such that* $\mathrm{spt}(\mathbb{P})\subseteq[-1,1]^d$. *Here* L *is the Lipschitz-constant of* σ.

Proof. A single data point may underestimate or overestimate the average integral, and the proof of convergence relies on these cancellations. A convenient tool to formalize cancellation and decouple this randomness from other effects is through Rademacher complexity [SSBD14, Chapter 26].

We denote $S = \{X_1, \dots, X_n\}$ and assume that S is drawn iid from the distribution $\mathbb{P}$. We consider an auxiliary random vector ξ such that the entries ξ_i are iid (and independent of S) variables which take the values ± 1 with probability $1/2$. Furthermore, abbreviate by B the unit ball in $\mathcal{B}(\mathbb{P})$. Furthermore

$$\begin{aligned}\operatorname{Rep}(B,S) &= \sup_{\phi\in B}\left[\frac{1}{n}\sum_{i=1}^n \phi(X_i) - \int_Q \phi(x)\,\mathbb{P}(\mathrm{d}x)\right]\\ \operatorname{Rad}(B,S) &= \mathbb{E}_\xi \sup_{\phi\in B}\left[\frac{1}{n}\sum_{i=1}^n \xi_i\,\phi(X_i)\right].\end{aligned}$$

According to [SSBD14, Lemma 26.2], the Rademacher complexity Rad bounds the representativeness of the set S by

$$\mathbb{E}_S \operatorname{Rep}(B,S) \le 2\,\mathbb{E}_S \operatorname{Rad}(B,S).$$

The unit ball in Barron space is given by convex combinations of functions

$$\phi_{w,b}(x) = \pm\frac{\sigma(w^T x + b)}{|w|+1} \quad \text{or} \quad \pm\frac{\sigma(w^T x + b)}{|w|+|b|}$$

respectively, so for fixed ξ, the linear map $\phi \mapsto \frac{1}{n}\sum_{i=1}^n \xi_i\,\phi(X_i)$ at one of the functions in the convex hull, i.e.

$$\operatorname{Rad}(B,S) = \mathbb{E}_\xi \sup_{(w,b)}\left[\frac{1}{n}\sum_{i=1}^n \xi_i\,\frac{\sigma(w^T X_i + b)}{|w|+1}\right].$$

According to the Contraction Lemma [SSBD14, Lemma 26.9], the Lipschitz-nonlinearity σ can be neglected in the computation of the complexity. If L is the Lipschitz-constant of σ, then

$$\begin{aligned}\mathbb{E}_\xi \sup_{(w,b)}\left[\frac{1}{n}\sum_{i=1}^n \xi_i\,\frac{\sigma(w^T X_i + b)}{|w|+1}\right] &\le L\,\mathbb{E}_\xi \sup_{(w,b)}\left[\frac{1}{n}\sum_{i=1}^n \xi_i\,\frac{w^T X_i + b}{|w|+1}\right]\\ &= L\,\mathbb{E}_\xi \sup_{w\in\mathbb{R}^d}\frac{w^T}{|w|+1}\frac{1}{n}\sum_{i=1}^n \xi_i X_i\\ &= L\,\mathbb{E}_\xi \sup_{|w|\le 1} w^T\frac{1}{n}\sum_{i=1}^n \xi_i X_i\\ &\le L\sup_i |X_i|_\infty\sqrt{\frac{2\,\log(2d)}{n}}\end{aligned}$$

where we used [SSBD14, Lemma 26.11] for the complexity bound of the linear function class and [SSBD14, Lemma 26.6] to eliminate the scalar translation. A similar computation can be done in the ReLU case. □

Appendix B. A Brief Review of Reproducing Kernel Hilbert Spaces, Random Feature Models, and the Neural Tangent Kernel

For an introduction to kernel methods in machine learning, see e.g. [SSBD14, Chapter 16] or [CS09] in the context of deep learning. Let $k : \mathbb{R}^d \times \mathbb{R} \to \mathbb{R}$ be a symmetric positive definite

kernel. The *reproducing kernel Hilbert space* (RKHS) $\mathcal{H}_k$ associated with k is the completion of the collection of functions of the form

$$h(x) = \sum_{i=1}^{n} a_i\, k(x, x_i)$$

under the scalar product $\langle k(x,x_i), k(x,x_j)\rangle_{\mathcal{H}_k} = k(x_i, x_j)$. Note that, due to the positive definiteness of the kernel, the representation of a function is unique.

B.1. **Random Feature Models.** The random feature models considered in this article are functions of the same form

$$f(x) = \frac{1}{m}\sum_{i=1}^{m} a_i\, \sigma\big(w_i^T x + b_i\big)$$

as a two-layer neural network. Unlike neural networks, (w_i, b_i) are not trainable variables and remain fixed after initialization. Random feature models are linear (so easier to optimize) but less expressive than shallow neural networks. While two-layer neural networks of infinite width are modelled as

$$f(x) = \int_{\mathbb{R}^{d+2}} a\, \sigma(w^T x + b)\, \pi(\mathrm{d}a \otimes \mathrm{d}w \otimes \mathrm{d}b)$$

for variable probability measures π, an infinitely wide random feature model is given as

$$f(x) = \int_{\mathbb{R}^{d+1}} a(w,b)\, \sigma(w^T x + b)\, \pi^0(\mathrm{d}w \otimes \mathrm{d}b)$$

for a fixed distribution π^0 on $\mathbb{R}^{d+1}$. One can think of Barron space as the union over all random feature spaces. It is well known that neural networks can represent the same function in different ways. For ReLU activation, there is a degree of degeneracy due to the identity

$$0 = x + \alpha - (x + \alpha) = \sigma(x) - \sigma(-x) + \sigma(\alpha) - \sigma(-\alpha) - \sigma(x+\alpha) + \sigma(-x-\alpha).$$

This can be generalized to higher dimension and integrated in α to show that the random feature representation of functions where π^0 is the uniform distribution on the sphere has a similar degeneracy. For given π^0, denote

$$N = \left\{ a \in L^2(\pi^0) \;\middle|\; \int_{\mathbb{R}^{d+1}} a(w,b)\, \sigma(w^T x + b)\, \pi^0(\mathrm{d}w \otimes \mathrm{d}b) = 0 \quad \mathbb{P} - \text{a.e.} \right\}$$

Lemma B.1. [RR08, Proposition 4.1] *The space of random feature models is dense in the RKHS for the kernel*

$$k(x,x') = \int_{\mathbb{R}^{d+1}} \sigma(w^T x + b)\, \sigma(w^T x' + b)\, \pi^0(\mathrm{d}w \otimes \mathrm{d}b)$$

and if

$$f(x) = \int_{\mathbb{R}^{d+1}} a(w,b) \sigma(w^T x + b)\, \pi^0(\mathrm{d}w \otimes \mathrm{d}b)$$

for $a \in L^2(\pi^0)/N$, then

$$\|f\|_{\mathcal{H}_k} = \|a\|_{L^2(\pi^0)/N}.$$

Note that the proof in the source uses a different normalization. The result in this form is achieved by setting $a = \alpha/p, b = \beta/p$ in the notation of [RR08].

Corollary B.2. *Any function in the random feature RKHS is Lipschitz-continuous and $[f]_{\mathrm{Lip}} \leq \|f\|_{\mathcal{H}_k}$. Thus $\mathcal{H}_k$ embeds compactly into $C^0(\operatorname{spt}\mathbb{P})$ by the Arzelà-Ascoli theorem and a fortiori into $L^2(\mathbb{P})$ for any compactly supported measure $\mathbb{P}$.*

Let $\mathbb{P}$ be a compactly supported data distribution on $\mathbb{R}^d$. Then the kernel k acts on $L^2(\mathbb{P})$ by the map

$$\overline{K} : L^2(\mathbb{P}) \to L^2(\mathbb{P}), \qquad \overline{K}u(x) = \int_{\mathbb{R}^d} u(x')\, k(x,x')\, \mathbb{P}(\mathrm{d}x').$$

Computing the eigenvalues of the kernel k for a given parameter distribution π^0 and a given data distribution $\mathbb{P}$ is a non-trivial endeavor. The task simplifies considerably under the assumption of symmetry, but remains complicated. The following results are taken from [Bac17, Appendix D], where more general results are proved for α-homogeneous activation for $\alpha \geq 0$. We specify $\alpha = 1$.

Lemma B.3. *Assume that $\pi^0 = \mathbb{P} = \alpha_d^{-1} \cdot \mathcal{H}^d|_{S^d}$ where S^d is the Euclidean unit sphere, α_d is its volume and $\sigma = ReLU$. Then the eigenfunctions of the kernel k are the spherical harmonics. The k-th eigenvalue λ_k (counted without repetition) occurs with the same multiplicity $N(d,k)$ as eigenfunctions to the k-th eigenvalue of the Laplace-Beltrami operator on the sphere (the spherical harmonics). Precisely*

$$N(d,k) = \frac{2k+d-1}{k}\binom{k+d-2}{d-1} \qquad \text{and} \qquad \lambda_k = \frac{d-1}{2\pi}\, 2^{-k}\, \frac{\Gamma(d/2)\,\Gamma(k-1)}{\Gamma(k/2)\,\Gamma(\frac{k+d+2}{2})}$$

for $k \geq 2$.

We can extract a decay rate for the eigenvalues counted with repetition by estimating the height and width of the individual plateaus of eigenvalues. Denote by μ_i the eigenvalues of k counted as often as they occur, i.e.

$$\mu_i = \lambda_k \quad \Leftrightarrow \quad \sum_{j=1}^{k-1} N(d,j) < i \leq \sum_{j=1}^{k} N(d,j).$$

By Stirling's formula one can estimate that for fixed d

$$\lambda_k \sim k^{-\frac{d-3}{2}}$$

as $k \to \infty$ where we write $a_k \sim b_k$ if and only if

$$0 < \liminf_{k\to\infty} \frac{a_k}{b_k} \leq \limsup_{k\to\infty} \frac{a_k}{b_k} < \infty.$$

On the other hand

$$N(d,k) = \frac{d}{k}\binom{k+d-1}{d} = \frac{d}{k}\frac{(k+d-1)\dots k}{d!} \sim \frac{(k+d)^{d-1}}{(d-1)!}.$$

In particular, if $\mu_i = \lambda_k$, then

$$i \sim C(d) \sum_{j=1}^{k} j^{d-1} \sim C(d) \int_1^{k+1} t^{d-1}\,\mathrm{d}t \sim C(d)\, k^d.$$

Thus

$$\mu_i = \lambda_k \sim k^{-\frac{d-3}{2}} \sim i^{-\frac{d-3}{2d}} = i^{-\frac{1}{2}+\frac{3}{2d}}.$$

Corollary B.4. *Consider the random feature model for ReLU activation when both parameter and data measure are given by the uniform distribution on the unit sphere. Then the eigenvalues of the kernel decay like $i^{-\frac{1}{2}+\frac{3}{2d}}$.*

Remark B.5. It is easy to see that up to a multiplicative constant, all radially symmetric parameter distributions π^0 lead to the same random feature kernel. This can be used to obtain an explicit formula in the case when π^0 is a standard Gaussian in d dimensions. Since k only depends on the angle between x and x', we may assume that $x = e_1$ and $x' = \cos\phi\, e_1 + \sin\phi\, e_2$ with

$\phi \in [0, \pi]$. Now, one can use that the projection of the standard Gaussian onto the $w_1 w_2$-plane is a lower-dimensional standard Gaussian. Thus the kernel does not depend on the dimension. An explicit computation in two dimensions shows that

$$k(x, x') = \frac{\pi - \phi}{\pi} \cos\phi + \frac{\sin\phi}{\pi},$$

see [CS09, Section 2.1].

B.2. **The Neural Tangent Kernel.** The neural tangent kernel [JGH18] is a different model for infinitely wide neural networks. For two layer networks, it is obtained as the limiting object in a scaling regime for parameters which makes the Barron norm infinite. When training networks on empirical risk, in a certain scaling regime between the number of data points, the number of neurons, and the initialization of parameters, it can be shown that parameters do not move far from their initial position according to a parameter distribution $\bar\pi$, and that the gradient flow optimization of neural networks is close to the optimization of a kernel method for all times [DZPS18, EMW19d]. This kernel is called the *neural tangent kernel* (NTK) and is obtained as the sum of derivatives of the feature function with respect to all trainable parameters. It linearizes the dynamics at the initial parameter distribution. For networks with one hidden layer this is

$$\begin{aligned}
k(x, x') &= \int_{\mathbb{R}\times\mathbb{R}^d\times\mathbb{R}} \nabla_{(a,w,b)}\big(a\sigma(w^T x + b)\big) \cdot \nabla_{(a,w,b)}\big(a\sigma(w^T x' + b)\big)\, \bar\pi(\mathrm{d}a \otimes \mathrm{d}w \otimes \mathrm{d}b) \\
&= \int_{\mathbb{R}\times\mathbb{R}^d\times\mathbb{R}} \sigma(w^T x + b)\, \sigma(w^T x' + b) \\
&\qquad + a^2 \sigma'(w^T x + b)\, \sigma'(w^T x' + b) \big[\langle x, x'\rangle + 1\big]\, \bar\pi(\mathrm{d}a \otimes \mathrm{d}w \otimes \mathrm{d}b) \\
&= k_{RF}(x, x') + \int_{\mathbb{R}\times\mathbb{R}^d\times\mathbb{R}} a^2\, \sigma'(w^T x + b)\, \sigma'(w^T x' + b)(x, 1) \cdot (x, 1)\, \bar\pi(\mathrm{d}a \otimes \mathrm{d}w \otimes \mathrm{d}b)
\end{aligned}$$

where k_{RF} denotes the random feature kernel with distribution $P_{(w,b),\sharp}\bar\pi$. The second term is obtained on the right hand side is a positive definite kernel in itself. This can be seen most easily by recalling that

$$\begin{aligned}
&\sum_{i,j} c_i c_j \int_{\mathbb{R}^{d+2}} \nabla_{(w,b)}\big(a\sigma(w^T x_i + b)\big) \cdot \nabla_{(w,b)}\big(a\sigma(w^T x_j + b)\big)\, \bar\pi(\mathrm{d}a \otimes \mathrm{d}w \otimes \mathrm{d}b) \\
&\qquad = \int_{\mathbb{R}^{d+2}} \nabla_{(w,b)} a^2 \left(\sum_i c_i\, \sigma(w^T x_i + b)\right) \cdot \nabla_{(w,b)} \left(\sum_j c_j\, \sigma(w^T x_j + b)\right) \bar\pi(\mathrm{d}a \otimes \mathrm{d}w \otimes \mathrm{d}b) \\
&\qquad = \int_{\mathbb{R}^{d+2}} \left|\nabla_{(w,b)} a^2 \left(\sum_i c_i\, \sigma(w^T x_i + b)\right)\right|^2 \bar\pi(\mathrm{d}a \otimes \mathrm{d}w \otimes \mathrm{d}b) \\
&\qquad \geq 0.
\end{aligned}$$

On the other hand, if we assume that $|a| = a_0$ and $|(w, b)| = 1$ almost surely, we find that

$$\begin{aligned}
&\sum_{i,j} c_i c_j \int_{\mathbb{R}\times\mathbb{R}^d\times\mathbb{R}} a^2\, \sigma'(w^T x_i + b)\, \sigma'(w^T x_j + b)\, (x_i, 1)\, I\, (x_j, 1)^T\, \bar\pi(\mathrm{d}a \otimes \mathrm{d}w \otimes \mathrm{d}b) \\
&\qquad = \int_{\mathbb{R}\times\mathbb{R}^d\times\mathbb{R}} a^2 \left|\sum_i c_i\, \sigma'(w^T x_i + b) \begin{pmatrix} x_i \\ 1 \end{pmatrix}\right|^2 \bar\pi(\mathrm{d}a \otimes \mathrm{d}w \otimes \mathrm{d}b) \\
&\qquad \leq a_0^2 \int_{\mathbb{R}\times\mathbb{R}^d\times\mathbb{R}} \left|\sum_i c_i\, \sigma'(w^T x_i + b) \begin{pmatrix} x_i \\ 1 \end{pmatrix} \cdot \begin{pmatrix} w \\ b \end{pmatrix}\right|^2 \bar\pi(\mathrm{d}a \otimes \mathrm{d}w \otimes \mathrm{d}b)
\end{aligned}$$

$$= a_0^2 \int_{\mathbb{R}\times\mathbb{R}^d\times\mathbb{R}} \left| \sum_i c_i\, \sigma'(w^T x_i + b)(w^T x_i + b) \right|^2 \bar{\pi}(\mathrm{d}a \otimes \mathrm{d}w \otimes \mathrm{d}b)$$

$$= a_0^2 \int_{\mathbb{R}\times\mathbb{R}^d\times\mathbb{R}} \left| \sum_i c_i\, \sigma(w^T x_i + b) \right|^2 \bar{\pi}(\mathrm{d}a \otimes \mathrm{d}w \otimes \mathrm{d}b)$$

since σ is positively one-homogeneous. Thus the NTK satisfies

$$k_{RF} \leq k \leq (1 + |a_0|^2)\, k_{RF}$$

in the sense of quadratic forms. In particular, the eigenvalues of the NTK and the random feature kernel decay at the same rate. Clearly, in exchange for larger constants it suffices to assume that (a, w, b) are bounded. In practice, the intialization of (w, b) is Gaussian, which concentrates close to the Euclidean sphere of radius $\sqrt{d}$ in d dimensions.

The neural tangent kernel is also defined for deep networks, see for example [ADH+19, LXS+19, Yan19, DLL+18, EMWW19].

References

[ADH+19] S. Arora, S. S. Du, W. Hu, Z. Li, R. R. Salakhutdinov, and R. Wang. On exact computation with an infinitely wide neural net. In *Advances in Neural Information Processing Systems*, pages 8139–8148, 2019.

[Bac17] F. Bach. Breaking the curse of dimensionality with convex neural networks. *The Journal of Machine Learning Research*, 18(1):629–681, 2017.

[Bar93] A. R. Barron. Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information theory*, 39(3):930–945, 1993.

[Bre11] H. Brezis. *Functional analysis, Sobolev spaces and partial differential equations*. Universitext. Springer, New York, 2011.

[CS09] Y. Cho and L. K. Saul. Kernel methods for deep learning. In *Advances in neural information processing systems*, pages 342–350, 2009.

[Cyb89] G. Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989.

[DLL+18] S. S. Du, J. D. Lee, H. Li, L. Wang, and X. Zhai. Gradient descent finds global minima of deep neural networks. *arXiv:1811.03804 [cs.LG]*, 2018.

[Dob10] M. Dobrowolski. *Angewandte Funktionalanalysis: Funktionalanalysis, Sobolev-Räume und elliptische Differentialgleichungen*. Springer-Verlag, 2010.

[DZPS18] S. S. Du, X. Zhai, B. Poczos, and A. Singh. Gradient descent provably optimizes over-parameterized neural networks. *arXiv:1810.02054 [cs.LG]*, 2018.

[EMW18] W. E, C. Ma, and L. Wu. A priori estimates of the population risk for two-layer neural networks. *Comm. Math. Sci.*, 17(5):1407 – 1425 (2019), arxiv:1810.06397 [cs.LG] (2018).

[EMW19a] W. E, C. Ma, and Q. Wang. A priori estimates of the population risk for residual networks. *arXiv:1903.02154 [cs.LG]*, 2019.

[EMW19b] W. E, C. Ma, and L. Wu. Barron spaces and the compositional function spaces for neural network models. *arXiv:1906.08039 [cs.LG]*, 2019.

[EMW19c] W. E, C. Ma, and L. Wu. Machine learning from a continuous viewpoint. *arxiv:1912.12777 [math.NA]*, 2019.

[EMW19d] W. E, C. Ma, and L. Wu. A comparative analysis of optimization and generalization properties of two-layer neural network and random feature models under gradient descent dynamics. *Sci. China Math.*, https://doi.org/10.1007/s11425-019-1628-5, arXiv:1904.04326 [cs.LG] (2019).

[EMWW19] W. E, C. Ma, Q. Wang, and L. Wu. Analysis of the gradient descent algorithm for a deep neural network model with skip-connections. *arXiv:1904.05263 [cs.LG]*, 2019.

[EW20] W. E and S. Wojtowytsch. On the Banach spaces associated with multi-layer ReLU networks of infinite width. *arXiv:2007.15623 [stat.ML]*, 2020.

[FG15] N. Fournier and A. Guillin. On the rate of convergence in Wasserstein distance of the empirical measure. *Probability Theory and Related Fields*, 162(3-4):707–738, 2015.

[GKNV19] R. Gribonval, G. Kutyniok, M. Nielsen, and F. Voigtlaender. Approximation spaces of deep neural networks. *arXiv:1905.01208 [math.FA]*, 2019.

[Hor91] K. Hornik. Approximation capabilities of multilayer feedforward networks. *Neural networks*, 4(2):251–257, 1991.

[JGH18] A. Jacot, F. Gabriel, and C. Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in neural information processing systems*, pages 8571–8580, 2018.

[Lor66] G. Lorentz. *Approximation of Functions*. Holt, Rinehart and Winston, New York, 1966.

[LXS+19] J. Lee, L. Xiao, S. Schoenholz, Y. Bahri, R. Novak, J. Sohl-Dickstein, and J. Pennington. Wide neural networks of any depth evolve as linear models under gradient descent. In *Advances in neural information processing systems*, pages 8570–8581, 2019.

[RR08] A. Rahimi and B. Recht. Uniform approximation of functions with random bases. In *2008 46th Annual Allerton Conference on Communication, Control, and Computing*, pages 555–561. IEEE, 2008.

[SSBD14] S. Shalev-Shwartz and S. Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.

[Vil08] C. Villani. *Optimal transport: old and new*, volume 338. Springer Science & Business Media, 2008.

[WE20] S. Wojtowytsch and W. E. Can shallow neural networks beat the curse of dimensionality? A mean field training perspective. *arXiv:2005.10815 [cs.LG]*, 2020.

[Yan19] G. Yang. Scaling limits of wide neural networks with weight sharing: Gaussian process behavior, gradient independence, and neural tangent kernel derivation. *arXiv:1902.04760 [cs.NE]*, 2019.

Weinan E, Department of Mathematics and Program in Applied and Computational Mathematics, Princeton University, Princeton, NJ 08544, USA

Email address: weinan@math.princeton.edu

Stephan Wojtowytsch, Program in Applied and Computational Mathematics, Princeton University, Princeton, NJ 08544

Email address: stephanw@princeton.edu