

Nonparametric Regression Using Over-parameterized Shallow ReLU Neural Networks

Yunfei Yang *

*Department of Mathematics, City University of Hong Kong
Kowloon, Hong Kong, China*

YUNFYANG@CITYU.EDU.HK

Ding-Xuan Zhou

*School of Mathematics and Statistics, University of Sydney
Sydney, NSW 2006, Australia*

DINGXUAN.ZHOU@SYDNEY.EDU.AU

Editor: Joseph Salmon

Abstract

It is shown that over-parameterized neural networks can achieve minimax optimal rates of convergence (up to logarithmic factors) for learning functions from certain smooth function classes, if the weights are suitably constrained or regularized. Specifically, we consider the nonparametric regression of estimating an unknown d -variate function by using shallow ReLU neural networks. It is assumed that the regression function is from the Hölder space with smoothness $\alpha < (d+3)/2$ or a variation space corresponding to shallow neural networks, which can be viewed as an infinitely wide neural network. In this setting, we prove that least squares estimators based on shallow neural networks with certain norm constraints on the weights are minimax optimal, if the network width is sufficiently large. As a byproduct, we derive a new size-independent bound for the local Rademacher complexity of shallow ReLU neural networks, which may be of independent interest.

Keywords: neural networks, nonparametric regression, over-parameterization, regularization, rate of convergence

1. Introduction

In nonparametric regression, we are given a data set of n samples $\mathcal{D}_n = \{(X_i, Y_i)\}_{i=1}^n$, which are independent and identically distributed as a $\mathbb{R}^d \times \mathbb{R}$ -valued random vector (X, Y) . The statistical problem is to estimate the so-called regression function $h(x) = \mathbb{E}[Y|X = x]$ from the observed data \mathcal{D}_n . One of the most popular estimators is the least squares

$$\hat{f}_n \in \operatorname{argmin}_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n (f(X_i) - Y_i)^2,$$

where \mathcal{F} is a suitably chosen function class. In recent years, due to the breakthrough of deep learning, there is increasing interest in the literature in studying the performance of the least squares when \mathcal{F} is parameterized by neural networks (Bauer and Kohler, 2019; Chen et al., 2022; Jiao et al., 2023a; Kohler and Langer, 2021; Mao et al., 2021; Mao and Zhou, 2022; Nakada and Imaizumi, 2020; Schmidt-Hieber, 2020; Suzuki, 2019; Yang and Zhou, 2024).

*. Corresponding author.

To analyze rates of convergence for the least squares estimator \widehat{f}_n , one often assume that the regression function h is in certain function class \mathcal{H} , which represents prior knowledge on the problem. Typically, we can decompose the error of \widehat{f}_n as

$$\mathbb{E}_X \left[\left| \widehat{f}_n(X) - h(X) \right|^2 \right] \lesssim \mathcal{E}_{app}(\mathcal{H}, \mathcal{F}) + \mathcal{C}_n(\mathcal{F}). \quad (1)$$

Here, for two quantities A and B , $A \lesssim B$ (or $B \gtrsim A$) denotes the statement that $A \leq cB$ for some constant $c > 0$ (we will also denote $A \asymp B$ when $A \lesssim B \lesssim A$). In the error decomposition (1), the approximation error $\mathcal{E}_{app}(\mathcal{H}, \mathcal{F})$ quantifies how well we can approximate $h \in \mathcal{H}$ by functions in \mathcal{F} , and the sample complexity $\mathcal{C}_n(\mathcal{F})$ measures how well the learned function $\widehat{f}_n \in \mathcal{F}$ generalizes to unseen data. For example, the complexity $\mathcal{C}_n(\mathcal{F})$ is often measured by the covering number, VC dimension (Vapnik and Chervonenkis, 1971) or Rademacher complexity (Bartlett and Mendelson, 2002) of \mathcal{F} . In most analysis for neural networks, the approximation error and the complexity are bounded by the size of the networks, e.g. the number of parameters or neurons. In particular, when $\mathcal{H} = \mathcal{H}^\alpha$ is the unit ball of Hölder space with smoothness index $\alpha > 0$, the results of Schmidt-Hieber (2020); Kohler and Langer (2021) established the rate $n^{-\frac{2\alpha}{d+2\alpha}}$ (up to logarithmic factors) for deep neural networks, which is known to be minimax optimal (Stone, 1982). However, in these analyses, one has to use under-parameterized neural networks (the number of parameters is less than the sample size n) in order to get optimal trade-offs between $\mathcal{E}_{app}(\mathcal{H}, \mathcal{F})$ and $\mathcal{C}_n(\mathcal{F})$. But, in practical applications of deep neural networks, the number of parameters is often much larger than the number of samples. Hence, these results cannot fully explain the success of neural networks used in practice. To obtain rates of convergence in the over-parameterized regime, the recent works of Jiao et al. (2023b) and Yang and Zhou (2024) proposed to bound $\mathcal{E}_{app}(\mathcal{H}, \mathcal{F})$ and $\mathcal{C}_n(\mathcal{F})$ by the size of the weights, i.e. certain norms of the weight matrices, rather than the size of networks. Using this idea, in Yang and Zhou (2024), we obtained nearly optimal rates for over-parameterized neural networks.

The purpose of this paper is to show that over-parameterized neural networks can achieve the minimax optimal rates for estimating certain smooth functions, if the weights are suitably constrained or regularized. Specifically, we consider shallow neural networks with ReLU activation

$$f_\theta(x) := \sum_{i=1}^N a_i \sigma((x^\top, 1)w_i), \quad x \in \mathbb{R}^d,$$

where $\theta := (a_1, w_1^\top, \dots, a_N, w_N^\top)^\top$ denotes the vector of trainable parameters and $\sigma(t) := \max\{t, 0\}$ is the rectified linear unit (ReLU). Following the ideas in Jiao et al. (2023b); Yang and Zhou (2024), we constrain the weights by $\kappa(\theta) := \sum_{i=1}^N |a_i| \|w_i\|_2 \leq M$ and study the performance of the constrained least squares

$$\operatorname{argmin}_{\kappa(\theta) \leq M} \frac{1}{n} \sum_{i=1}^n (f_\theta(X_i) - Y_i)^2.$$

As in Yang and Zhou (2024), we assume that the regression function h belongs to the unit ball of Hölder class \mathcal{H}^α with $\alpha < (d+3)/2$ or the unit ball of the variation space $\mathcal{F}_\sigma(1)$, which can be regarded as an infinitely wide neural network with weight constraint

$\kappa(\theta) \leq 1$ (see Section 2 for details). The recent paper (Yang and Zhou, 2024) proved the rates $n^{-\frac{2\alpha}{d+3+2\alpha}}$ for \mathcal{H}^α and $n^{-1/2}$ for $\mathcal{F}_\sigma(1)$ up to logarithmic factors, under the condition that M is chosen properly and the network width N is sufficiently large (i.e. weights constrained over-parameterized neural networks). We improve these rates to $n^{-\frac{2\alpha}{d+2\alpha}}$ and $n^{-\frac{d+3}{2d+3}}$, which are minimax optimal for \mathcal{H}^α and $\mathcal{F}_\sigma(1)$ respectively (Stone, 1982; Parhi and Nowak, 2023; Yang and Zhou, 2024). The improvement is from the use of localization technique (Bartlett et al., 2005; Koltchinskii, 2006) in estimating the complexity $\mathcal{C}_n(\mathcal{F})$ of neural networks. One of our main technical contributions is a new size-independent bound for the local Rademacher complexity of shallow neural networks (see Definition 4 and Theorem 5):

$$\mathcal{R}_n(\mathcal{NN}(N, M); \delta) \lesssim \frac{\delta^{\frac{3}{d+3}} M^{\frac{d}{d+3}}}{\sqrt{n}} \sqrt{\log(nM/\delta)},$$

where $\mathcal{NN}(N, M)$ is the set of neural network functions f_θ with width at most N and constraint $\kappa(\theta) \leq M$. Our bound generalizes the bound $\mathcal{O}(M/\sqrt{n})$ for Rademacher complexity in Golowich et al. (2020) which was used in Yang and Zhou (2024) to control the sample complexity.

Besides the constrained least squares, we also study the corresponding regularized least squares

$$\operatorname{argmin}_\theta \frac{1}{n} \sum_{i=1}^n (f_\theta(X_i) - Y_i)^2 + \lambda \kappa(\theta),$$

where $\lambda > 0$ is a tunable parameter. It is proven that this optimization problem has essentially the same solutions as the problem

$$\operatorname{argmin}_\theta \frac{1}{n} \sum_{i=1}^n (f_\theta(X_i) - Y_i)^2 + \frac{\lambda}{2} \|\theta\|_2^2.$$

Our second main result shows that the solutions of these two optimization problems also achieve the minimax optimal rates of learning functions in \mathcal{H}^α with $\alpha < (d+3)/2$ and $\mathcal{F}_\sigma(1)$, if λ is chosen properly. This result gives a theoretical guarantee for the use of some regularization methods, such as Neyshabur et al. (2015a).

The rest of the paper is organized as follows. In Section 2, we introduce some function classes of shallow neural networks and their approximation results. Section 3 presents our main results on the rates of convergence for over-parameterized neural networks. In Section 4, we give some discussions on our results and related works. The omitted proofs of some theorems and lemmas are given in Section 5.

2. The Variation Space of Shallow Neural Networks

Let $\mathbb{B}^d = \{x \in \mathbb{R}^d : \|x\|_2 \leq 1\}$ be the unit ball of \mathbb{R}^d , and $\mathbb{S}^d = \{x \in \mathbb{R}^{d+1} : \|x\|_2 = 1\}$ be the unit sphere of \mathbb{R}^{d+1} . We denote by $\mathcal{M}(\mathbb{S}^d)$ the space of signed Radon measures on \mathbb{S}^d with the total variation norm $\|\mu\| = |\mu|(\mathbb{S}^d)$ for $\mu \in \mathcal{M}(\mathbb{S}^d)$. It is well known that $\mathcal{M}(\mathbb{S}^d)$ is the dual space of the space of continuous functions $C(\mathbb{S}^d)$ (Folland, 1999, Section 7.3). We are interested in the function class \mathcal{F}_σ that contains continuous functions $f_\mu : \mathbb{B}^d \rightarrow \mathbb{R}$ of

the integral form

$$f_\mu(x) := \int_{\mathbb{S}^d} \sigma((x^\top, 1)v) d\mu(v), \quad x \in \mathbb{B}^d, \mu \in \mathcal{M}(\mathbb{S}^d), \quad (2)$$

where σ is the ReLU function. Note that the integral representation is not unique for $f \in \mathcal{F}_\sigma$. For example, if μ is supported on the set

$$S_- := \left\{ (x_1, \dots, x_{d+1})^\top \in \mathbb{S}^d : x_{d+1} \leq -\sqrt{2}/2 \right\}, \quad (3)$$

then $f_\mu = 0$ is the zero function on \mathbb{B}^d . Following Bach (2017) and Yang and Zhou (2024), we can define a norm on \mathcal{F}_σ by

$$\gamma(f) := \inf \left\{ \|\mu\| : f = f_\mu, \mu \in \mathcal{M}(\mathbb{S}^d) \right\}.$$

By the compactness of closed balls in $\mathcal{M}(\mathbb{S}^d)$ (due to Prokhorov's theorem), the infimum defining $\gamma(f)$ is attained by some measure μ . One can show that \mathcal{F}_σ equipped with the norm γ is a Banach space. Since the norm is defined through the variation norm, we call \mathcal{F}_σ the variation space (of shallow ReLU neural networks). For any $M > 0$, we denote

$$\mathcal{F}_\sigma(M) := \{f \in \mathcal{F}_\sigma : \gamma(f) \leq M\}.$$

Each function in this class can be thought of as an infinitely wide neural network with a constraint on its weights. There are several other definitions and characterizations of function spaces corresponding to infinitely wide neural networks in recent studies (Bartolucci et al., 2023; Ongie et al., 2020; Parhi and Nowak, 2022, 2023; Savarese et al., 2019; Siegel and Xu, 2020, 2022, 2023; Siegel, 2023). We will give more discussions in Section 4.

The neural networks with finite neurons are corresponding to discrete measures $\mu \in \mathcal{M}(\mathbb{S}^d)$ that are supported on finitely many points. Let us denote the collection of these measures by $\mathcal{M}_{disc}(\mathbb{S}^d)$, which is a linear subspace of $\mathcal{M}(\mathbb{S}^d)$, and define the corresponding function class $\mathcal{F}_{\sigma, disc} := \{f_\mu : \mu \in \mathcal{M}_{disc}(\mathbb{S}^d)\}$. Similar to the norm γ defined on \mathcal{F}_σ , we can also define a norm on $\mathcal{F}_{\sigma, disc}$ by

$$\kappa(f) := \inf \left\{ \|\mu\| : f = f_\mu, \mu \in \mathcal{M}_{disc}(\mathbb{S}^d) \right\}.$$

By definition, $\gamma(f) \leq \kappa(f)$ for all $f \in \mathcal{F}_{\sigma, disc}$. But we do not know whether these two norms are equal for $f \in \mathcal{F}_{\sigma, disc}$. Nevertheless, it is good enough to bound $\gamma(f)$ by $\kappa(f)$ in our theory. We will see soon that the norm $\kappa(f)$ can be estimated in practice, and hence it can be used in statistical estimations.

For any $N \in \mathbb{N}$, the function class $\mathcal{NN}(N) \subseteq \mathcal{F}_{\sigma, disc}$ of shallow neural networks with (at most) N neurons can be parameterized in the form

$$f_\theta(x) := \sum_{i=1}^N a_i \sigma((x^\top, 1)w_i), \quad a_i \in \mathbb{R}, w_i \in \mathbb{R}^{d+1}, x \in \mathbb{B}^d, \quad (4)$$

where θ denotes the vector of parameters

$$\theta := (a_1, w_1^\top, \dots, a_N, w_N^\top)^\top \in \mathbb{R}^{(d+2)N}.$$

Note that, in this parameterization, we do not restrict the internal weights w_i to \mathbb{S}^d . Thanks to the homogeneity of the ReLU function, it is easy to see that $f_\theta = f_\mu$ with discrete measure $\mu = \sum_{i=1}^N c_i \delta_{v_i} \in \mathcal{M}_{disc}(\mathbb{S}^d)$, where $v_i = w_i / \|w_i\|_2$ and $c_i = a_i \|w_i\|_2$ if $\|w_i\|_2 \neq 0$, and $c_i = 0$ if $\|w_i\|_2 = 0$. This implies that we can estimate $\kappa(f_\theta)$ by the following quantity

$$\kappa(\theta) := \sum_{i=1}^N |a_i| \|w_i\|_2.$$

With a slight abuse of notation, we have denoted it by the same notation as the norm κ . The reason is given by the following theorem, which shows that $\kappa(f_\theta)$ is equivalent to the minimum of $\kappa(\theta)$ over all parameterization (4).

Theorem 1 *For any $f \in \mathcal{NN}(N)$,*

$$\kappa(f) \leq \inf \left\{ \kappa(\theta) : f = f_\theta, \theta \in \mathbb{R}^{(d+2)N} \right\} \leq 3\kappa(f).$$

The proof is given in Subsection 5.1. In the proof, we actually show that any $f \in \mathcal{NN}(N)$ can be reduced to the following form

$$f(x) = \sum_{k=1}^K c_k \sigma((x^\top, 1)v_k) + (x^\top, 1)w, \quad (5)$$

where $0 \leq K \leq N$, $w \in \mathbb{R}^{d+1}$ and $v_1, \dots, v_K \in S_0$ satisfy $v_i \neq \pm v_k$ for $1 \leq i \neq k \leq K$. The set S_0 is defined by $S_0 := \mathbb{S}^d \setminus (S_- \cup S_+)$, where S_- is defined by (3) and $S_+ := \{-v : v \in S_-\}$. Note that the set S_0 contains vectors $v \in \mathbb{S}^d$ such that the function $x \mapsto \sigma((x^\top, 1)v)$ is nonlinear on \mathbb{B}^d . We prove that, for the reduced parameterization (5),

$$\|w\|_2 \vee \sum_{k=1}^K |c_k| \leq \kappa(f) \leq 2\|w\|_2 + \sum_{k=1}^K |c_k|, \quad (6)$$

where we use the notation $a \vee b := \max\{a, b\}$. The inequality (6) gives a way to estimate $\kappa(f)$ in practice. Roughly speaking, the norm $\kappa(f)$ is a combination of the variation of the nonlinear part of f and the norm of the linear part.

Now, we are ready to introduce the central object of this paper

$$\mathcal{NN}(N, M) := \{f_\theta \in \mathcal{NN}(N) : \kappa(\theta) \leq M\}.$$

Since $\gamma(f_\theta) \leq \kappa(f_\theta) \leq \kappa(\theta)$, we have $\mathcal{NN}(N, M) \subseteq \mathcal{F}_\sigma(M)$. It was shown in Yang and Zhou (2024, Proposition 2.2) that $\mathcal{F}_\sigma(M)$ is the closure of $\cup_{N \in \mathbb{N}} \mathcal{NN}(N, M)$ in $C(\mathbb{B}^d)$. The rates of approximation by $\mathcal{NN}(N, M)$ for smooth function classes are also studied in many recent works (Bach, 2017; Klusowski and Barron, 2018; Siegel and Xu, 2022; Mao and Zhou, 2023; Siegel, 2023; Yang and Zhou, 2024). To describe these approximation results, let us first recall the classical notion of smoothness of functions. Given a smoothness index $\alpha > 0$, we write $\alpha = r + \beta$ where $r \in \mathbb{N}_0 := \mathbb{N} \cup \{0\}$ and $\beta \in (0, 1]$. Let $C^{r, \beta}(\mathbb{R}^d)$ be the Hölder space with the norm

$$\|f\|_{C^{r, \beta}(\mathbb{R}^d)} := \max \left\{ \|f\|_{C^r(\mathbb{R}^d)}, \max_{\|s\|_1=r} |\partial^s f|_{C^{0, \beta}(\mathbb{R}^d)} \right\},$$

where $s = (s_1, \dots, s_d) \in \mathbb{N}_0^d$ is a multi-index and

$$\begin{aligned} \|f\|_{C^r(\mathbb{R}^d)} &:= \max_{\|s\|_1 \leq r} \|\partial^s f\|_{L^\infty(\mathbb{R}^d)}, \\ |f|_{C^{0,\beta}(\mathbb{R}^d)} &:= \sup_{x \neq y \in \mathbb{R}^d} \frac{|f(x) - f(y)|}{\|x - y\|_2^\beta}. \end{aligned}$$

Here, we use $\|\cdot\|_{L^\infty}$ to denote the supremum norm, since we only consider continuous functions in this paper. We write $C^{r,\beta}(\mathbb{B}^d)$ for the Banach space of all restrictions to \mathbb{B}^d of functions in $C^{r,\beta}(\mathbb{R}^d)$. The norm of this space is given by $\|f\|_{C^{r,\beta}(\mathbb{B}^d)} = \inf\{\|g\|_{C^{r,\beta}(\mathbb{R}^d)} : g \in C^{r,\beta}(\mathbb{R}^d) \text{ and } g = f \text{ on } \mathbb{B}^d\}$. For convenience, we will denote the unit ball of $C^{r,\beta}(\mathbb{B}^d)$ by

$$\mathcal{H}^\alpha := \left\{ f \in C^{r,\beta}(\mathbb{B}^d) : \|f\|_{C^{r,\beta}(\mathbb{B}^d)} \leq 1 \right\}.$$

Note that, for $\alpha = 1$, \mathcal{H}^α is a class of Lipschitz continuous functions.

In the recent work (Yang and Zhou, 2024), we characterized how well one can use neural networks in $\mathcal{F}_\sigma(M)$ to approximate functions in \mathcal{H}^α . It was shown by Yang and Zhou (2024, Theorem 2.1) that if $\alpha > (d+3)/2$ then $\mathcal{H}^\alpha \subseteq \mathcal{F}_\sigma(M)$ for some constant M , and if $\alpha < (d+3)/2$ then

$$\sup_{h \in \mathcal{H}^\alpha} \inf_{f \in \mathcal{F}_\sigma(M)} \|h - f\|_{L^\infty(\mathbb{B}^d)} \lesssim M^{-\frac{2\alpha}{d+3-2\alpha}}.$$

For the critical value $\alpha = (d+3)/2$, it seems that the inclusion $\mathcal{H}^\alpha \subseteq \mathcal{F}_\sigma(M)$ also holds for some constant M , but Yang and Zhou (2024) only obtained the exponential approximation rate $\mathcal{O}(\exp(-\alpha M^2))$. Combining these results with the rate of approximation for $\mathcal{F}_\sigma(M)$ by its subclass $\mathcal{F}_\sigma(N, M)$ derived in Bach (2017); Siegel (2023), one can obtain the following approximation bounds for shallow neural networks. Recall that we use the notation $a \vee b := \max\{a, b\}$.

Theorem 2 (Yang and Zhou 2024) *We have the following approximation bounds.*

- (1) *For \mathcal{H}^α with $0 < \alpha < (d+3)/2$, it holds that*

$$\sup_{h \in \mathcal{H}^\alpha} \inf_{f \in \mathcal{NN}(N, M)} \|h - f\|_{L^\infty(\mathbb{B}^d)} \lesssim N^{-\frac{\alpha}{d}} \vee M^{-\frac{2\alpha}{d+3-2\alpha}}.$$

- (2) *For $\mathcal{F}_\sigma(1)$, it holds that*

$$\sup_{h \in \mathcal{F}_\sigma(1)} \inf_{f \in \mathcal{NN}(N, 1)} \|h - f\|_{L^\infty(\mathbb{B}^d)} \lesssim N^{-\frac{d+3}{2d}}.$$

As discussed by Yang and Zhou (2024), for \mathcal{H}^α with $0 < \alpha < (d+3)/2$, the approximation rate is optimal in terms of the number of neurons N and the norm constraint M . We will use this result to establish the optimal rates of convergence for shallow neural networks in the nonparametric regression problem.

3. Nonparametric Regression

We consider the problem of nonparametric regression in a classical setting. Suppose we have a data set of $n \geq 2$ samples $\mathcal{D}_n = \{(X_i, Y_i)\}_{i=1}^n$, which are independently and identically generated from a probability distribution supported on $\mathbb{B}^d \times \mathbb{R}$. Let μ be the marginal distribution of the covariate X and $h(x) := \mathbb{E}[Y|X = x]$ denote the regression function. In this paper, we assume that $|Y| \leq B$ with some fixed constant $B > 0$ and $h \in \mathcal{H}$, where we will study the cases that \mathcal{H} is the unit ball of the variation space $\mathcal{F}_\sigma(1)$ or the Hölder class \mathcal{H}^α with $\alpha < (d+3)/2$. Note that \mathcal{H}^α with $\alpha > (d+3)/2$ is included in the first case by scaling, because $\mathcal{H}^\alpha \subseteq \mathcal{F}_\sigma(M)$ for some M (Yang and Zhou, 2024). Let us denote the noises by

$$\eta_i := Y_i - h(X_i), \quad i = 1, \dots, n,$$

then $\mathbb{E}[\eta_i] = 0$ and $|\eta_i| \leq 2B$. We denote the sample data points by the sequence $X_{1:n} := (X_1, \dots, X_n)$ (similarly, $\eta_{1:n} := (\eta_1, \dots, \eta_n)$ denotes the sequence of noises). The empirical distribution is denoted by $\mu_n := \frac{1}{n} \sum_{i=1}^n \delta_{X_i}$, and the associated $L^2(\mu_n)$ norm is given by $\|f\|_{L^2(\mu_n)}^2 := \frac{1}{n} \sum_{i=1}^n f(X_i)^2$.

3.1 Least Squares

In practice, one popular way to estimate the regression function $h \in \mathcal{H}$ is by the constrained least squares

$$\hat{f}_n \in \operatorname{argmin}_{f \in \mathcal{F}_n} \frac{1}{n} \sum_{i=1}^n (f(X_i) - Y_i)^2. \quad (7)$$

We are interested in the case that $\mathcal{F}_n = \mathcal{NN}(N_n, M_n)$ is parameterized by a shallow ReLU neural network. For simplicity, we assume here and in the sequel that the minimum above indeed exists. The performance of the estimation is measured by the expected risk

$$\mathcal{L}(\hat{f}_n) := \mathbb{E}_{(X,Y)}[(\hat{f}_n(X) - Y)^2].$$

It is equivalent to evaluating the estimator by the excess risk

$$\|\hat{f}_n - h\|_{L^2(\mu)}^2 = \mathcal{L}(\hat{f}_n) - \mathcal{L}(h).$$

In statistical analysis of learning algorithms, we often require that the hypothesis class is uniformly bounded. We define the truncation operator \mathcal{T}_b with level $b > 0$ for real-valued functions f as

$$\mathcal{T}_b f(x) := \begin{cases} f(x), & \text{if } |f(x)| \leq b, \\ \operatorname{sgn}(f(x))b, & \text{if } |f(x)| > b. \end{cases}$$

For a class of real-valued functions \mathcal{F} , we use the notation $\mathcal{T}_b \mathcal{F} := \{\mathcal{T}_b f : f \in \mathcal{F}\}$. It is easy to see that, if $b \geq \sup_{h \in \mathcal{H}} \|h\|_{L^\infty(\mathbb{B}^d)}$, then $\|\mathcal{T}_b \hat{f}_n - h\|_{L^2(\mu)} \leq \|\hat{f}_n - h\|_{L^2(\mu)}$ for any $h \in \mathcal{H}$. Hence, truncating the output of the estimator \hat{f}_n appropriately does not increase the excess risk. In the following, we will simply take $b = B$. Our first result gives convergence rates for the least squares estimator \hat{f}_n with \mathcal{F}_n being a shallow neural network.

Theorem 3 *Let \hat{f}_n be the estimator (7) with $\mathcal{F}_n = \mathcal{NN}(N_n, M_n)$.*

(1) If $\mathcal{H} = \mathcal{H}^\alpha$ with $\alpha < (d + 3)/2$, we choose

$$M_n \asymp n^{\frac{d+3-2\alpha}{2d+4\alpha}}, \quad N_n \gtrsim n^{\frac{d}{d+2\alpha}},$$

then the bound

$$\|\mathcal{T}_B \widehat{f}_n - h\|_{L^2(\mu)}^2 \lesssim n^{-\frac{2\alpha}{d+2\alpha}} \log n$$

holds with probability at least $1 - c_1 \exp(-c_2 n^{\frac{d}{d+2\alpha}} \log n)$ for some constants $c_1, c_2 > 0$.

(2) If $\mathcal{H} = \mathcal{F}_\sigma(1)$, we let $M \geq 1$ be a constant and choose

$$M_n = M, \quad N_n \gtrsim n^{\frac{d}{2d+3}},$$

then the bound

$$\|\mathcal{T}_B \widehat{f}_n - h\|_{L^2(\mu)}^2 \lesssim n^{-\frac{d+3}{2d+3}} \log n$$

holds with probability at least $1 - c_1 \exp(-c_2 n^{\frac{d}{2d+3}} \log n)$ for some constants $c_1, c_2 > 0$.

Note that the (implied) constants in the theorem depend on d, α and B . Notice that we can always take $N_n \gtrsim n$, which gives justification for the use of over-parameterized neural networks in practice. The results also hold for $\mathcal{F}_n = \mathcal{F}_\sigma(M_n)$, which corresponds to $N_n = \infty$. The required minimal widths N_n of neural networks match the results of Yang and Zhou (2024, Theorem 4.2) in the under-parameterized regime. For over-parameterized shallow neural networks, Theorem 3 improves the rates of convergence in Yang and Zhou (2024, Theorem 4.7). Ignoring logarithmic factors, the rates $n^{-\frac{2\alpha}{d+2\alpha}}$ for \mathcal{H}^α and $n^{-\frac{d+3}{2d+3}}$ for $\mathcal{F}_\sigma(1)$ are minimax optimal, as shown by Stone (1982); Parhi and Nowak (2023); Yang and Zhou (2024).

Similar to the classical analysis of nonparametric regression (Györfi et al., 2002; Wainwright, 2019), our proof of Theorem 3 bounds the error $\|\mathcal{T}_B \widehat{f}_n - h\|_{L^2(\mu)}^2$ as the inequality (1) by two terms: the approximation error $\mathcal{E}_{app}(\mathcal{H}, \mathcal{F}_n)$ and the complexity $\mathcal{C}_n(\mathcal{F}_n)$. The approximation error can be handled by Theorem 2. We control the complexity by using the localization technique (Bartlett et al., 2005; Koltchinskii, 2006).

Definition 4 (Local complexity) Let $\xi_{1:n}$ be a sequence of independent zero-mean random variables. For a given radius $\delta > 0$ and a sequence of sample points $X_{1:n}$, we define the local complexity of a function class \mathcal{F} at scale δ with respect to $\xi_{1:n}$ by

$$\mathcal{G}_n(\mathcal{F}; \delta, \xi_{1:n}) := \mathbb{E}_{\xi_{1:n}} \left[\sup_{f \in \mathcal{F}, \|f\|_{L^2(\mu_n)} \leq \delta} \left| \frac{1}{n} \sum_{i=1}^n \xi_i f(X_i) \right| \right].$$

If each ξ_i is the Rademacher random variable (taking values ± 1 with equal probability $1/2$), $\mathcal{G}_n(\mathcal{F}; \delta, \xi_{1:n})$ is the local Rademacher complexity and is denoted by $\mathcal{R}_n(\mathcal{F}; \delta)$.

This definition is a direct generalization of the local Rademacher and Gaussian complexities (Bartlett et al., 2005; Koltchinskii, 2006; Wainwright, 2019). It is obvious from the definition that $\mathcal{G}_n(\mathcal{F}; \delta, \xi_{1:n})$ also depends on the sample points $X_{1:n}$, which is omitted in the notation because we will provide bounds for local complexities that hold uniformly on $X_{1:n}$.

When $\xi_{1:n}$ is the noise sequence $\eta_{1:n}$, the local complexity $\mathcal{G}_n(\mathcal{F}; \delta, \eta_{1:n})$ measures how well the function class \mathcal{F} correlates with the noise on the samples $X_{1:n}$. Note that, $\mathcal{G}_n(\mathcal{F}; \delta, \eta_{1:n})$ is a data-dependent quantity. If the sample points $X_{1:n}$ are random, then $\mathcal{G}_n(\mathcal{F}; \delta, \eta_{1:n})$ is also a random variable.

To use the local complexities, we often require that the function class \mathcal{F} is star-shaped (around the origin), meaning that, for any $f \in \mathcal{F}$ and $a \in [0, 1]$, the function $af \in \mathcal{F}$. If the star-shaped condition fails to hold, one can consider the star hull

$$\text{star}(\mathcal{F}) := \{af : f \in \mathcal{F}, a \in [0, 1]\}.$$

One of the key properties of local complexities of a star-shaped class \mathcal{F} is that the function $\delta \mapsto \mathcal{G}_n(\mathcal{F}; \delta, \xi_{1:n})/\delta$ is non-increasing on $(0, \infty)$. This can be easily proven by observing that the rescaled function $\tilde{f} = \frac{\delta_1}{\delta_2} f \in \mathcal{F}$ if $f \in \mathcal{F}$ and $0 < \delta_1 \leq \delta_2$ (see Wainwright 2019, Lemma 13.6 for example). Consequently, for any constant $c > 0$, the inequality $\mathcal{G}_n(\mathcal{F}; \delta, \xi_{1:n}) \leq c\delta^2$ always has positive solutions $\delta = \delta_n(\mathcal{F})$. Our proof of Theorem 3 uses these solutions to control the complexities of neural networks $\mathcal{C}_n(\mathcal{F}_n) \lesssim \delta_n^2(\mathcal{F}_n)$ (see Lemmas 6 and 7 below). Note that the function classes $\mathcal{F}_\sigma(M)$ and $\mathcal{NN}(N, M)$ are star-shaped. The following theorem gives an estimate for the local complexities of shallow neural networks, which holds uniformly for all sample points $X_{1:n}$ and network width.

Theorem 5 *Let $\xi_{1:n}$ be a sequence of independent sub-Gaussian random variables with parameter $\varsigma > 0$ in the sense that*

$$\mathbb{E}[\exp(\lambda\xi_i)] \leq \exp(\varsigma^2\lambda^2/2), \quad \forall \lambda \in \mathbb{R}.$$

Then, for any $0 < \delta \leq M$,

$$\mathcal{G}_n(\mathcal{F}_\sigma(M); \delta, \xi_{1:n}) \lesssim \frac{\varsigma \delta^{\frac{3}{d+3}} M^{\frac{d}{d+3}}}{\sqrt{n}} \sqrt{\log(nM/\delta)},$$

where the implied constant is independent of $\xi_{1:n}$ and the sample points $X_{1:n}$ in \mathbb{B}^d . The bound also holds for the function class $\text{star}(\mathcal{T}_B\mathcal{F}_\sigma(M))$ for any constant $B > 0$.

The proof of Theorem 5 is given in Subsection 5.2. Note that, when $\delta = M$, we obtain the bound $\mathcal{O}(M\sqrt{\log n}/\sqrt{n})$ which is the same as that for (global) Rademacher complexity derived in Golowich et al. (2020, Theorem 3.2), if we ignore the logarithmic factor. The key point of Theorem 5 is that the bound is independent of the network width so that we can apply it to over-parameterized neural networks.

Now, let us come back to the proof of Theorem 3. We begin with a decomposition of the excess risk $\|\mathcal{T}_B\hat{f}_n - h\|_{L^2(\mu)}^2$, using similar ideas as Nakada and Imaizumi (2020) and Schmidt-Hieber (2020) (see also Wainwright (2019, Chapters 13 and 14)). For any $f \in \mathcal{F}_n$ (we will take f to minimize the approximation error $\|f - h\|_{L^\infty(\mathbb{B}^d)}$), observe that

$$\begin{aligned} \|\mathcal{T}_B\hat{f}_n - h\|_{L^2(\mu)}^2 &\leq 2\|\mathcal{T}_B\hat{f}_n - \mathcal{T}_Bf\|_{L^2(\mu)}^2 + 2\|\mathcal{T}_Bf - h\|_{L^2(\mu)}^2 \\ &\leq 2\|\mathcal{T}_{2B}(\hat{f}_n - f)\|_{L^2(\mu)}^2 + 2\|f - h\|_{L^2(\mu)}^2, \end{aligned} \tag{8}$$

where we use $|\mathcal{T}_B \widehat{f}_n(x) - \mathcal{T}_B f(x)| \leq |\mathcal{T}_{2B}(\widehat{f}_n - f)(x)|$ and $\|h\|_{L^\infty(\mathbb{B}^d)} \leq B$ in the second inequality. Note that, for the empirical error,

$$\|\mathcal{T}_{2B}(\widehat{f}_n - f)\|_{L^2(\mu_n)}^2 \leq \|\widehat{f}_n - f\|_{L^2(\mu_n)}^2 \leq 2\|\widehat{f}_n - h\|_{L^2(\mu_n)}^2 + 2\|f - h\|_{L^2(\mu_n)}^2. \quad (9)$$

By the definition of \widehat{f}_n ,

$$\frac{1}{n} \sum_{i=1}^n (\widehat{f}_n(X_i) - Y_i)^2 \leq \frac{1}{n} \sum_{i=1}^n (f(X_i) - Y_i)^2.$$

Using $Y_i = h(X_i) + \eta_i$, we have the base inequality

$$\|\widehat{f}_n - h\|_{L^2(\mu_n)}^2 \leq \|f - h\|_{L^2(\mu_n)}^2 + \frac{2}{n} \sum_{i=1}^n \eta_i (\widehat{f}_n(X_i) - f(X_i)). \quad (10)$$

Our proof can be divided into the following three steps.

Step 1. Estimating $\|\widehat{f}_n - h\|_{L^2(\mu_n)}^2$ by using the base inequality (10).

Step 2. Bounding $\|\mathcal{T}_{2B}(\widehat{f}_n - f)\|_{L^2(\mu)}^2$ by its empirical counterpart $\|\mathcal{T}_{2B}(\widehat{f}_n - f)\|_{L^2(\mu_n)}^2$.

Step 3. Combining inequalities (8), (9), Step 1, Step 2 with the approximation result (Theorem 2).

In Step 1, we can treat the sample points $X_{1:n}$ as fixed so that the randomness is only from the noises $\eta_{1:n}$. The effect of the noise can be measured by the local complexity with respect to the noise. The detail is given in the following lemma, which provides an oracle inequality for the least squares estimator. It is convenient to use the notation $\partial\mathcal{F} := \{f_1 - f_2 : f_1, f_2 \in \mathcal{F}\}$.

Lemma 6 *For any fixed sample points $X_{1:n}$, let \widehat{f}_n be the estimator (7) and δ_n be any positive solution to the inequality*

$$\mathcal{G}_n(\text{star}(\partial\mathcal{F}_n); \delta_n, \eta_{1:n}) \leq \delta_n^2.$$

Then, there are constants $c_1, c_2 > 0$ such that the bound

$$\|\widehat{f}_n - h\|_{L^2(\mu_n)}^2 \leq 3\|f - h\|_{L^2(\mu_n)}^2 + 32\delta_n^2, \quad \forall f \in \mathcal{F}_n,$$

holds with probability at least $1 - c_1 \exp(-c_2 n \delta_n^2 / B^2)$.

This lemma is similar to Wainwright (2019, Theorem 13.13) which assumes that the noise is standard Gaussian. Our proof is essentially the same as that of Wainwright (2019, Theorem 13.13). Since we will use a similar argument in the proof of Lemma 10, we give a proof in Subsection 5.3 for completeness. Note that δ_n^2 in Lemma 6 corresponds to the complexity $\mathcal{C}_n(\mathcal{F})$ in the decomposition (1). But Lemma 6 only bounds the error in the $L^2(\mu_n)$ norm, rather than the $L^2(\mu)$ norm. So, we need the second step.

In Step 2, we need to quantify the effect of the random sample points $X_{1:n}$. This can be done by using the following uniform law of large number with localization from Wainwright (2019, Theorem 14.1 and Proposition 14.25). In the lemma, it is required that the function class \mathcal{F} is star-shaped and B -uniformly bounded for some constant $B > 0$, meaning that $\|f\|_{L^\infty(\mathbb{B}^d)} \leq B$ for all $f \in \mathcal{F}$.

Lemma 7 *Given a star-shaped and B -uniformly bounded function class \mathcal{F} , let ϵ_n be any positive solution to the inequality*

$$\mathcal{R}_n(\mathcal{F}; \epsilon_n) \leq \frac{\epsilon_n^2}{B}.$$

Then, there are constants $c_1, c_2 > 0$ such that the bound

$$\|f\|_{L^2(\mu)}^2 \leq 2\|f\|_{L^2(\mu_n)}^2 + \epsilon_n^2, \quad \forall f \in \mathcal{F},$$

holds with probability at least $1 - c_1 \exp(-c_2 n \epsilon_n^2 / B^2)$.

We can now prove Theorem 3. We will use c_1, c_2 to denote constants, which may be different in different bounds.

Proof of Theorem 3

Step 1. We apply Lemma 6 to $\mathcal{F}_n = \mathcal{NN}(N_n, M_n)$. Since $\text{star}(\partial\mathcal{F}_n) \subseteq \mathcal{F}_\sigma(2M_n)$ and $\eta_{1:n}$ is a sequence of independent sub-Gaussian random variables with parameter $\varsigma = 2B$, by Theorem 5, we have

$$\mathcal{G}_n(\text{star}(\partial\mathcal{F}_n); \delta_n, \eta_{1:n}) \lesssim \frac{B\delta_n^{\frac{3}{d+3}} M_n^{\frac{d}{d+3}}}{\sqrt{n}} \sqrt{\log(nM_n/\delta_n)},$$

for all sample points $X_{1:n}$ in \mathbb{B}^d . Hence, we can choose

$$\delta_n \asymp n^{-\frac{d+3}{4d+6}} M_n^{\frac{d}{2d+3}} \sqrt{\log(nM_n)}.$$

Then, with probability at least $1 - c_1 \exp(-c_2 n^{\frac{d}{2d+3}} M_n^{\frac{2d}{2d+3}} \log(nM_n))$, it holds

$$\|\widehat{f}_n - h\|_{L^2(\mu_n)}^2 \lesssim \|f - h\|_{L^2(\mu_n)}^2 + n^{-\frac{d+3}{2d+3}} M_n^{\frac{2d}{2d+3}} \log(nM_n), \quad \forall f \in \mathcal{F}_n. \quad (11)$$

Step 2. Observe that

$$\mathcal{T}_{2B}(\widehat{f}_n - f) \in \mathcal{T}_{2B}\partial\mathcal{F}_n \subseteq \mathcal{T}_{2B}\mathcal{F}_\sigma(2M_n) \subseteq \text{star}(\mathcal{T}_{2B}\mathcal{F}_\sigma(2M_n)).$$

We can apply Lemma 7 to the star-shaped and $2B$ -uniformly bounded function class $\text{star}(\mathcal{T}_{2B}\mathcal{F}_\sigma(2M_n))$. By Theorem 5,

$$\mathcal{R}_n(\text{star}(\mathcal{T}_{2B}\mathcal{F}_\sigma(2M_n)); \epsilon_n) \lesssim \frac{\epsilon_n^{\frac{3}{d+3}} M_n^{\frac{d}{d+3}}}{\sqrt{n}} \sqrt{\log(nM_n/\epsilon_n)},$$

and hence we can choose

$$\epsilon_n \asymp n^{-\frac{d+3}{4d+6}} M_n^{\frac{d}{2d+3}} \sqrt{\log(nM_n)}.$$

Then, with probability at least $1 - c_1 \exp(-c_2 n^{\frac{d}{2d+3}} M_n^{\frac{2d}{2d+3}} \log(nM_n))$, it holds

$$\|\mathcal{T}_{2B}(\widehat{f}_n - f)\|_{L^2(\mu)}^2 \lesssim \|\mathcal{T}_{2B}(\widehat{f}_n - f)\|_{L^2(\mu_n)}^2 + n^{-\frac{d+3}{2d+3}} M_n^{\frac{2d}{2d+3}} \log(nM_n), \quad \forall f \in \mathcal{F}_n. \quad (12)$$

Step 3. Combining inequalities (8), (9), (11) and (12), we get

$$\|\mathcal{T}_B \widehat{f}_n - h\|_{L^2(\mu)}^2 \lesssim \|f - h\|_{L^2(\mu)}^2 + \|f - h\|_{L^2(\mu_n)}^2 + n^{-\frac{d+3}{2d+3}} M_n^{\frac{2d}{2d+3}} \log(nM_n),$$

which holds with probability at least $1 - c_1 \exp(-c_2 n^{\frac{d}{2d+3}} M_n^{\frac{2d}{2d+3}} \log(nM_n))$ simultaneously for all $f \in \mathcal{F}_n$. Taking infimum over $f \in \mathcal{F}_n$ shows

$$\|\mathcal{T}_B \widehat{f}_n - h\|_{L^2(\mu)}^2 \lesssim \inf_{f \in \mathcal{F}_n} \|f - h\|_{L^\infty(\mathbb{B}^d)}^2 + n^{-\frac{d+3}{2d+3}} M_n^{\frac{2d}{2d+3}} \log(nM_n). \quad (13)$$

We continue the proof in two cases.

(1) If $\mathcal{H} = \mathcal{H}^\alpha$ with $\alpha < (d+3)/2$, then by Theorem 2,

$$\sup_{h \in \mathcal{H}} \inf_{f \in \mathcal{F}_n} \|f - h\|_{L^\infty(\mathbb{B}^d)}^2 \lesssim N_n^{-\frac{2\alpha}{d}} \vee M_n^{-\frac{4\alpha}{d+3-2\alpha}}.$$

If $M_n \asymp n^{\frac{d+3-2\alpha}{2d+4\alpha}}$ and $N_n \gtrsim n^{\frac{d}{d+2\alpha}}$, then inequality (13) implies

$$\|\mathcal{T}_B \widehat{f}_n - h\|_{L^2(\mu)}^2 \lesssim n^{-\frac{2\alpha}{d+2\alpha}} \log n,$$

holds with probability at least $1 - c_1 \exp(-c_2 n^{\frac{d}{d+2\alpha}} \log n)$.

(2) If $\mathcal{H} = \mathcal{F}_\sigma(1)$, then by Theorem 2, for any $M_n = M \geq 1$,

$$\sup_{h \in \mathcal{H}} \inf_{f \in \mathcal{F}_n} \|f - h\|_{L^\infty(\mathbb{B}^d)}^2 \lesssim N_n^{-\frac{d+3}{d}}.$$

If $N_n \gtrsim n^{\frac{d}{2d+3}}$, then inequality (13) implies

$$\|\mathcal{T}_B \widehat{f}_n - h\|_{L^2(\mu)}^2 \lesssim n^{-\frac{d+3}{2d+3}} \log n,$$

holds with probability at least $1 - c_1 \exp(-c_2 n^{\frac{d}{2d+3}} \log n)$. ■

Finally, we make a remark on the boundedness assumption of the noise, which can certainly be weakened.

Remark 8 *The boundedness of the noises $\eta_{1:n}$ is only used to derive the dimension-free high probability bound (32) in the proof of Lemma 6 below. The proof relies on the concentration inequality for convex Lipschitz functions. This argument can also be applied to certain sub-Gaussian noises characterized by Gozlan et al. (2017, Corollary 5.11) (which of course includes Gaussian). Hence, similar results as Theorem 3 can be obtained for such kind of noises.*

3.2 Regularized Least Squares

The least squares estimate (7) with $\mathcal{F}_n = \mathcal{NN}(N_n, M_n)$ is a constrained optimization problem, where we make constraints on the norm $\kappa(f) \leq M_n$. One can also estimate the regression function h by solving the corresponding regularized optimization problem

$$\operatorname{argmin}_{f \in \mathcal{NN}(N_n)} \frac{1}{n} \sum_{i=1}^n (f(X_i) - Y_i)^2 + \lambda_n \kappa(f), \quad (14)$$

where $\lambda_n > 0$ is a tunable parameter. Usually, the regularization parameter λ_n is chosen to satisfy $\lim_{n \rightarrow \infty} \lambda_n = 0$. We show in the next theorem that, if λ_n is chosen properly, the regularized least squares has the same rate of convergence as the least squares estimator. The result is also true if the norm $\kappa(f)$ in (14) is replaced by $\gamma(f)$.

Theorem 9 *Let $\widehat{f}_{n, \lambda_n}$ be a solution of the optimization problem (14).*

(1) *If $\mathcal{H} = \mathcal{H}^\alpha$ with $\alpha < (d+3)/2$, we choose*

$$N_n \gtrsim n^{\frac{d}{d+2\alpha}}, \quad \lambda_n \asymp n^{-\frac{d+3+2\alpha}{2d+4\alpha}} \log n,$$

then the bound

$$\|\mathcal{T}_B \widehat{f}_{n, \lambda_n} - h\|_{L^2(\mu)}^2 \lesssim n^{-\frac{2\alpha}{d+2\alpha}} \log n$$

holds with probability at least $1 - c_1 \exp(-c_2 n^{\frac{d}{d+2\alpha}} \log n)$ for some constants $c_1, c_2 > 0$.

(2) *If $\mathcal{H} = \mathcal{F}_\sigma(1)$, we choose*

$$N_n \gtrsim n^{\frac{d}{2d+3}}, \quad \lambda_n \asymp n^{-\frac{d+3}{2d+3}} \log n,$$

then the bound

$$\|\mathcal{T}_B \widehat{f}_{n, \lambda_n} - h\|_{L^2(\mu)}^2 \lesssim n^{-\frac{d+3}{2d+3}} \log n$$

holds with probability at least $1 - c_1 \exp(-c_2 n^{\frac{d}{2d+3}} \log n)$ for some constants $c_1, c_2 > 0$.

As in the proof of Theorem 3, we first give a decomposition of the excess risk: for any $f \in \mathcal{NN}(N_n)$,

$$\|\mathcal{T}_B \widehat{f}_{n, \lambda_n} - h\|_{L^2(\mu)}^2 \leq 2\|\mathcal{T}_{2B}(\widehat{f}_{n, \lambda_n} - f)\|_{L^2(\mu)}^2 + 2\|f - h\|_{L^2(\mu)}^2. \quad (15)$$

The empirical counterpart of the first term on the right hand side can be bounded as

$$\|\mathcal{T}_{2B}(\widehat{f}_{n, \lambda_n} - f)\|_{L^2(\mu_n)}^2 \leq 2\|\widehat{f}_{n, \lambda_n} - h\|_{L^2(\mu_n)}^2 + 2\|f - h\|_{L^2(\mu_n)}^2. \quad (16)$$

By the definition of $\widehat{f}_{n, \lambda_n}$,

$$\frac{1}{n} \sum_{i=1}^n (\widehat{f}_{n, \lambda_n}(X_i) - Y_i)^2 + \lambda_n \kappa(\widehat{f}_{n, \lambda_n}) \leq \frac{1}{n} \sum_{i=1}^n (f(X_i) - Y_i)^2 + \lambda_n \kappa(f).$$

Using $Y_i = h(X_i) + \eta_i$, we get the base inequality

$$\|\widehat{f}_{n, \lambda_n} - h\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(\widehat{f}_{n, \lambda_n}) \leq \|f - h\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(f) + \frac{2}{n} \sum_{i=1}^n \eta_i \left(\widehat{f}_{n, \lambda_n}(X_i) - f(X_i) \right). \quad (17)$$

Similar to the proof of Theorem 3, our proof strategy contains three steps:

- Step 1.** Estimating $\|\widehat{f}_{n,\lambda_n} - h\|_{L^2(\mu_n)}^2$ and $\kappa(\widehat{f}_{n,\lambda_n})$ by using the base inequality (17).
- Step 2.** Bounding $\|\mathcal{T}_{2B}(\widehat{f}_{n,\lambda_n} - f)\|_{L^2(\mu)}^2$ by its empirical counterpart $\|\mathcal{T}_{2B}(\widehat{f}_{n,\lambda_n} - f)\|_{L^2(\mu_n)}^2$.
- Step 3.** Combining the above estimates with inequalities (15), (16) and the approximation result (Theorem 2).

The last two steps can be done in a similar way as Theorem 3. The main difficulty in proving Theorem 9 is that we do not have an explicit bound on the norm $\kappa(\widehat{f}_{n,\lambda_n})$. In order to apply Theorems 2 and 5, one needs to estimate the norm of the learned function from the optimization problem (14) or equivalently the base inequality (17). A trivial bound is to choose $f = 0$ in (14). Then, the definition of $\widehat{f}_{n,\lambda_n}$ implies

$$\kappa(\widehat{f}_{n,\lambda_n}) \leq \frac{1}{n\lambda_n} \sum_{i=1}^n Y_i^2 \leq \frac{B^2}{\lambda_n}.$$

Unfortunately, this bound is too loose to obtain the optimal rate. To see this, let f^* be a minimizer of the regularized approximation error

$$\mathcal{E}_{\lambda_n} := \inf_{f \in \mathcal{NN}(N_n)} \|f - h\|_{L^2(\mu)}^2 + \lambda_n \kappa(f),$$

which can be regarded as the expected version of the regularized least squares (14). Then, we know that $\kappa(f^*) \leq \mathcal{E}_{\lambda_n}/\lambda_n$ by definition. Yet, we expect $\widehat{f}_{n,\lambda_n}$ to be a good approximation of f^* . So, one would expect that $\kappa(\widehat{f}_{n,\lambda_n}) \lesssim \mathcal{E}_{\lambda_n}/\lambda_n$, which is always better than the bound B^2/λ_n , because $\mathcal{E}_{\lambda_n} \rightarrow 0$ if λ_n is chosen properly. This phenomenon is also discussed in detail by Wu et al. (2005) and Steinwart and Christmann (2008, Chapter 7) in the analysis of support vector machines. They also discussed how to use iteration technique to improve the trivial bound. Nonetheless, we overcome the difficulty by using a new oracle inequality for the regularized least squares estimator (see Lemma 10 below). Our proof shows that one can apply this oracle inequality to obtain a bound close to $\mathcal{E}_{\lambda_n}/\lambda_n$.

The following lemma can be viewed as a modification of Lemma 6 to regularized estimators. Note that the constant c_0 in the lemma will be useful in the proof of Corollary 12 below. Recall that $\partial\mathcal{F} := \{f_1 - f_2 : f_1, f_2 \in \mathcal{F}\}$.

Lemma 10 *For any fixed sample points $X_{1:n}$ and constant $c_0 \geq 1$, let $\widehat{f}_{n,\lambda_n} \in \mathcal{NN}(N_n)$ be any estimator (i.e. a function depending on the random noises $\eta_{1:n}$, where $\mathbb{E}[\eta_i] = 0$ and $|\eta_i| \leq 2B$) that satisfies the inequality*

$$\begin{aligned} & \|\widehat{f}_{n,\lambda_n} - h\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(\widehat{f}_{n,\lambda_n}) \\ & \leq c_0 \left(\|f - h\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(f) \right) + \frac{2}{n} \left| \sum_{i=1}^n \eta_i \left(\widehat{f}_{n,\lambda_n}(X_i) - f(X_i) \right) \right|, \quad \forall f \in \mathcal{NN}(N_n). \end{aligned} \quad (18)$$

For a user-defined parameter $R > 0$, let $\delta_n = \delta_n(R)$ be any positive solution to the inequality

$$\mathcal{G}_n(\text{star}(\partial\mathcal{F}_{n,R}); \delta_n, \eta_{1:n}) \leq \delta_n^2,$$

where $\mathcal{F}_{n,R} := \{f \in \mathcal{NN}(N_n) : \kappa(f) \leq R\}$. There are constants $c_1, c_2 > 0$ such that, if $\lambda_n \geq 8\delta_n^2/R$, then the bound

$$\|\widehat{f}_{n,\lambda_n} - h\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(\widehat{f}_{n,\lambda_n}) \leq (1 + 2c_0) \left(\|f - h\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(f) \right) + 64\delta_n^2, \quad \forall f \in \mathcal{F}_{n,R},$$

holds with probability at least $1 - c_1 \exp(-c_2 n \delta_n^2 / B^2)$.

The proof of Lemma 10 is given in Subsection 5.4. The inequality we obtained in this lemma is in a similar form as the bound (1), where the complexity is measured by δ_n^2 and the approximation in L^2 norm is replaced by the regularized approximation in the empirical norm.

Proof of Theorem 9

To apply Lemma 10, let us first compute $\delta_n = \delta_n(R)$, where $R > 0$ will be chosen later. Since $\text{star}(\partial\mathcal{F}_{n,R}) \subseteq \mathcal{F}_\sigma(2R)$, by Theorem 5,

$$\mathcal{G}_n(\text{star}(\partial\mathcal{F}_{n,R}); \delta_n, \eta_{1:n}) \lesssim \frac{B\delta_n^{\frac{3}{d+3}} R^{\frac{d}{d+3}}}{\sqrt{n}} \sqrt{\log(nR/\delta_n)},$$

holds for all sample points $X_{1:n}$ in \mathbb{B}^d . Hence, we can choose

$$\delta_n^2 \asymp n^{-\frac{d+3}{2d+3}} R^{\frac{2d}{2d+3}} \log(nR).$$

We denote the approximation error of $\mathcal{F}_{n,R}$ by

$$\mathcal{E}(R) := \inf_{f \in \mathcal{F}_{n,R}} \|f - h\|_{L^\infty(\mathbb{B}^d)}^2.$$

Note that we use the supremum norm instead of $L^2(\mu_n)$ norm in the definition. Lemma 10 with $c_0 = 1$ shows that, if $\lambda_n = 8\delta_n^2/R$, then with probability at least $1 - c_1 \exp(-c_2 n \delta_n^2)$, it holds

$$\|\widehat{f}_{n,\lambda_n} - h\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(\widehat{f}_{n,\lambda_n}) \lesssim \mathcal{E}(R) + \delta_n^2,$$

In particular, we have

$$\kappa(\widehat{f}_{n,\lambda_n}) \lesssim \frac{\mathcal{E}(R) + \delta_n^2}{\lambda_n} \lesssim \frac{\mathcal{E}(R)}{\lambda_n} + R,$$

where we have used $\delta_n^2 \lesssim \lambda_n R$ by the choice of λ_n .

(1) If $\mathcal{H} = \mathcal{H}^\alpha$ with $\alpha < (d+3)/2$, then by Theorem 2,

$$\mathcal{E}(R) \leq \sup_{h \in \mathcal{H}} \inf_{f \in \mathcal{F}_{n,R}} \|f - h\|_{L^\infty(\mathbb{B}^d)}^2 \lesssim N_n^{-\frac{2\alpha}{d}} \vee R^{-\frac{4\alpha}{d+3-2\alpha}}.$$

We choose

$$R \asymp n^{\frac{d+3-2\alpha}{2d+4\alpha}}, \quad N_n \gtrsim R^{\frac{2d}{d+3-2\alpha}} \gtrsim n^{\frac{d}{d+2\alpha}},$$

then

$$\mathcal{E}(R) \lesssim n^{-\frac{2\alpha}{d+2\alpha}}, \quad \delta_n^2 \asymp n^{-\frac{2\alpha}{d+2\alpha}} \log n, \quad \lambda_n = 8\delta_n^2/R \asymp n^{-\frac{d+3+2\alpha}{2d+4\alpha}} \log n. \quad (19)$$

(2) If $\mathcal{H} = \mathcal{F}_\sigma(1)$, then by Theorem 2, for any $R \geq 1$,

$$\mathcal{E}(R) \leq \sup_{h \in \widehat{\mathcal{H}}} \inf_{f \in \mathcal{F}_{n,R}} \|f - h\|_{L^\infty(\mathbb{B}^d)}^2 \lesssim N_n^{-\frac{d+3}{d}}.$$

We choose

$$R = 1, \quad N_n \gtrsim n^{\frac{d}{2d+3}},$$

then

$$\mathcal{E}(R) \lesssim n^{-\frac{d+3}{2d+3}}, \quad \delta_n^2 \asymp n^{-\frac{d+3}{2d+3}} \log n, \quad \lambda_n = 8\delta_n^2/R \asymp n^{-\frac{d+3}{2d+3}} \log n. \quad (20)$$

Notice that, in both cases, we always have $\mathcal{E}(R) \lesssim \lambda_n R$, which implies $M_n := \kappa(\widehat{f}_{n,\lambda_n}) \lesssim R$. Similar to step 2 in the proof of Theorem 3, we can apply Lemma 7 to the function class $\text{star}(\mathcal{T}_{2B}\mathcal{F}_\sigma(M_n + R))$ to show that

$$\|\mathcal{T}_{2B}(\widehat{f}_{n,\lambda_n} - f)\|_{L^2(\mu)}^2 \lesssim \|\mathcal{T}_{2B}(\widehat{f}_{n,\lambda_n} - f)\|_{L^2(\mu_n)}^2 + \epsilon_n^2, \quad \forall f \in \mathcal{F}_{n,R},$$

holds with probability at least $1 - c_1 \exp(-c_2 n \epsilon_n^2)$, where

$$\epsilon_n^2 \asymp n^{-\frac{d+3}{2d+3}} R^{\frac{2d}{2d+3}} \log(nR) \asymp \delta_n^2.$$

Combining this with inequalities (15) and (16), where we choose $f \in \mathcal{F}_{n,R}$ to satisfy the bound $\|f - h\|_{L^\infty(\mathbb{B}^d)}^2 \lesssim \mathcal{E}(R)$, we obtain

$$\|\mathcal{T}_B \widehat{f}_{n,\lambda_n} - h\|_{L^2(\mu)}^2 \lesssim \mathcal{E}(R) + \delta_n^2,$$

holds with probability at least $1 - c_1 \exp(-c_2 n \delta_n^2)$. The estimates (19) and (20) then give the desired results. \blacksquare

Next, we show that it is possible to simplify the optimization problem (14) without worsening the convergence rates. In practice, the neural network $f_\theta \in \mathcal{NN}(N_n)$ is parameterized by θ in the form (4). By Theorem 1, we can estimate $\kappa(f)$ by $\kappa(\theta)$, which reduces the optimization problem (14) to

$$\operatorname{argmin}_{\theta \in \mathbb{R}^{(d+2)N_n}} \frac{1}{n} \sum_{i=1}^n (f_\theta(X_i) - Y_i)^2 + \lambda_n \kappa(\theta). \quad (21)$$

One can further simplify the optimization problem by observing that the regularizer $\kappa(\theta) \leq \|\theta\|_2^2/2$. It may be easier to solve the following optimization problem

$$\operatorname{argmin}_{\theta \in \mathbb{R}^{(d+2)N_n}} \frac{1}{n} \sum_{i=1}^n (f_\theta(X_i) - Y_i)^2 + \frac{\lambda_n}{2} \|\theta\|_2^2. \quad (22)$$

The next proposition shows that the two optimization problems (21) and (22) essentially have the same solutions. To the best of our knowledge, this equivalence was first pointed out by Neyshabur et al. (2015b, Theorem 1).

Proposition 11 *If θ^* is a solution of the optimization problem (21), then there exists a solution $\tilde{\theta}^*$ of (22) such that $f_{\theta^*} = f_{\tilde{\theta}^*}$. Any solution of (22) is also a solution of (21).*

Proof Let us denote $N = N_n$ for simplicity. Due to the positive homogeneity of the ReLU function, we can always rescale the parameter $\theta = (a_1, w_1^\top, \dots, a_N, w_N^\top)^\top$ to the parameter $(c_1 a_1, c_1^{-1} w_1^\top, \dots, c_N a_N, c_N^{-1} w_N^\top)^\top$ with $c_i > 0$, without altering the function f_θ . We use $r(\theta) = (\tilde{a}_1, \tilde{w}_1^\top, \dots, \tilde{a}_N, \tilde{w}_N^\top)^\top$ to denote the rescaled parameter of θ that satisfies $|\tilde{a}_i| = \|\tilde{w}_i\|_2$ for all $1 \leq i \leq N$ (if $|a_i| \|w_i\|_2 = 0$, we let $|\tilde{a}_i| = \|\tilde{w}_i\|_2 = 0$). Note that we always have $\kappa(\theta) = \kappa(r(\theta)) = \|r(\theta)\|_2^2/2 \leq \|\theta\|_2^2/2$.

Let us denote the empirical square loss of f_θ by $\mathcal{L}_n(f_\theta) = \frac{1}{n} \sum_{i=1}^n (f_\theta(X_i) - Y_i)^2$ for convenience. If θ^* is a solution of (21), then $r(\theta^*)$ is also a solution of (21). Thus, for any θ , we have

$$\begin{aligned} \mathcal{L}_n(f_{r(\theta^*)}) + \frac{\lambda_n}{2} \|r(\theta^*)\|_2^2 &= \mathcal{L}_n(f_{r(\theta^*)}) + \lambda_n \kappa(r(\theta^*)) \\ &\leq \mathcal{L}_n(f_\theta) + \lambda_n \kappa(\theta) \\ &\leq \mathcal{L}_n(f_\theta) + \frac{\lambda_n}{2} \|\theta\|_2^2, \end{aligned}$$

which shows that $r(\theta^*)$ is a solution of (22). Conversely, if θ^* is a solution of (22), then $r(\theta^*) = \theta^*$ because $\|r(\theta)\|_2^2 \leq \|\theta\|_2^2$. Thus, for any θ ,

$$\begin{aligned} \mathcal{L}_n(f_{\theta^*}) + \lambda_n \kappa(\theta^*) &= \mathcal{L}_n(f_{\theta^*}) + \frac{\lambda_n}{2} \|\theta^*\|_2^2 \\ &\leq \mathcal{L}_n(f_{r(\theta)}) + \frac{\lambda_n}{2} \|r(\theta)\|_2^2 \\ &= \mathcal{L}_n(f_\theta) + \kappa(\theta), \end{aligned}$$

which shows that θ^* also minimizes (21). ■

Using the same proof idea as Theorem 9, one can show that the solutions of the regularized optimization problems (21) and (22) can also achieve the minimax optimal rates for learning functions in \mathcal{H}^α or $\mathcal{F}_\sigma(1)$. The fundamental reason is that these estimators satisfy the base inequality (17) up to constant factors.

Corollary 12 *Let $\hat{f}_{n,\lambda_n} = f_{\hat{\theta}_{n,\lambda_n}}$, where $\hat{\theta}_{n,\lambda_n}$ is a solution of the optimization problem (21) or (22). Then the conclusion of Theorem 9 also holds.*

Proof By Proposition 11, it is enough to consider the case that $\hat{\theta}_{n,\lambda_n}$ is a solution of the optimization problem (21). As in the proof of Theorem 9, the inequalities (15) and (16) still hold for \hat{f}_{n,λ_n} . The only difference is the base inequality

$$\|\hat{f}_{n,\lambda_n} - h\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(\hat{\theta}_{n,\lambda_n}) \leq \|f - h\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(\theta) + \frac{2}{n} \sum_{i=1}^n \eta_i \left(\hat{f}_{n,\lambda_n}(X_i) - f(X_i) \right),$$

where $\theta \in \mathbb{R}^{(d+2)N_n}$ satisfy $f_\theta = f \in \mathcal{NN}(N_n)$. By Theorem 1, we can choose some θ such that $\kappa(\theta) \leq 3\kappa(f)$. Hence,

$$\|\hat{f}_{n,\lambda_n} - h\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(\hat{f}_{n,\lambda_n}) \leq 3 \left(\|f - h\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(f) \right) + \frac{2}{n} \sum_{i=1}^n \eta_i \left(\hat{f}_{n,\lambda_n}(X_i) - f(X_i) \right),$$

where we also use $\kappa(\widehat{f}_{n,\lambda_n}) \leq \kappa(\widehat{\theta}_{n,\lambda_n})$. Therefore, we can apply Lemma 10 with $c_0 = 3$. The remained proof is the same as Theorem 9. \blacksquare

4. Discussions and Related Works

Infinitely wide neural networks. The seminal work of Barron (1993) obtained dimension independent rate of approximation by shallow neural networks (with sigmoidal activations) for functions h , whose Fourier transform \widehat{h} satisfies $\int_{\mathbb{R}^d} \|\omega\| |\widehat{h}(\omega)| d\omega < \infty$. Barron’s results have been improved and extended to functions of certain integral forms, such as (2), and ReLU activation in recent years (Klusowski and Barron, 2018; Makovoz, 1996; Mhaskar, 2004, 2020; Siegel and Xu, 2020). These results lead to the study of variation spaces corresponding to neural networks (Bach, 2017; Siegel and Xu, 2023; Siegel, 2023; Yang and Zhou, 2024), which are also called Barron spaces in E et al. (2019, 2022). We refer the reader to Siegel and Xu (2022) for detail discussions on these spaces from an approximation theory perspective.

There is another line of works (Bartolucci et al., 2023; Ongie et al., 2020; Parhi and Nowak, 2021, 2022, 2023; Savarese et al., 2019) trying to characterize infinitely wide neural networks from a functional analysis point of view. These papers derived representer theorems showing that (finitely wide) neural networks are sparse solutions to data fitting problems with total variation regularization in the Radon domain:

$$\operatorname{argmin}_f \frac{1}{n} \sum_{i=1}^n \ell(Y_i, f(X_i)) + \lambda \operatorname{RTV}(f), \tag{23}$$

where the loss function ℓ is convex and coercive in its second argument and RTV is a seminorm defined through the Radon transform. Note that the null space of the seminorm RTV is the space of affine functions (Parhi and Nowak, 2021, Lemma 19). For $f \in \mathcal{NN}(N)$ in the reduced form (5), $\operatorname{RTV}(f)$ is exactly the variation $\sum_{k=1}^K |c_k|$ by Parhi and Nowak (2021, Lemma 25), and hence $\kappa(f) \approx \operatorname{RTV}(f) + \|\text{linear part of } f\|$ by inequality (6). Thus, for the square loss ℓ , we can view the optimization problem (23) as the problem (14) with width $N_n = \infty$.

Infinitely wide neural networks are also used in the study of training dynamics of neural networks. The neural tangent kernel theory demonstrates that the evolution of a neural network during gradient descent training can be described by a kernel in the infinite width limit (Allen-Zhu et al., 2019; Du et al., 2019; Jacot et al., 2018; Oymak and Soltanolkotabi, 2020). The mean field analysis approximates the evolution of the network weights by a gradient flow, defined through a partial differential equation, in the Wasserstein space of probability distributions (Mei et al., 2018, 2019; Sirignano and Spiliopoulos, 2020). These theories established the convergences of gradient descent training of over-parameterized neural networks in certain scalings of network weights (Chizat et al., 2019). It would be interesting to see whether these techniques can be applied to study the optimization problems (7) and (21).

Nonparametric regression using neural networks. As discussed in the introduction, rates of convergence of neural network regression estimators have been analyzed by many authors. Minimax optimal rates have been established for sparse deep neural networks

(Schmidt-Hieber, 2020; Suzuki, 2019), for fully connected deep neural networks (Kohler and Langer, 2021) and for shallow neural networks (Yang and Zhou, 2024). It has also been shown that deep neural networks are able to circumvent the curse of dimensionality under certain conditions, for example, when the intrinsic dimension is low (Chen et al., 2022; Jiao et al., 2023a; Nakada and Imaizumi, 2020) or the regression function has certain hierarchical structures (Bauer and Kohler, 2019; Schmidt-Hieber, 2020; Kohler and Langer, 2021). In the over-parameterized regime, Yang and Zhou (2024) proved nearly optimal rates by using (deep or shallow) neural networks with weight constraints. The paper of Parhi and Nowak (2023) established the minimax optimal rate for the variation space \mathcal{F}_σ using infinitely wide shallow neural networks. Our results can be viewed as generalizations of this result to Hölder classes \mathcal{H}^α and finitely wide networks. Note that all these results are established for least squares, where either the network architecture or the network weights are properly constrained. For regularized least squares, the recent work (Zhang and Wang, 2023) proved convergence rates for deep neural networks with a regularization similar to this paper. The rates they obtained are asymptotically optimal as the depth goes to infinity. But they did not show how to choose the regularization parameter. Our results show that regularized least squares with shallow networks can achieve the minimax optimal rates, at least for function classes with low smoothness. Our proof method (see Lemma 10) may be used to study other regularized estimators. A particular interesting direction is to generalize our results to deep neural networks.

Interpolation and regularization. The recent progress in deep learning revealed that over-parameterized models can interpolate noisy data and yet generalize well (Belkin et al., 2019a; Zhang et al., 2017). Motivated by this phenomenon, which is called benign overfitting, there has recently been a line of works on understanding theoretical mechanisms for the good generalization performance of interpolating models (see Bartlett et al. 2021; Belkin 2021 for reviews). In particular, benign overfitting has been theoretically established for many models, such as Nadaraya-Watson smoothing estimator (Belkin et al., 2019b), linear regression (Bartlett et al., 2020; Hastie et al., 2022), kernel regression (Liang and Rakhlin, 2020) and random features models (Mei et al., 2022). However, this certainly does not mean that any interpolator can generalize well, simply because there are too many estimators that interpolate observed data. Hence, explicit or implicit regularization is necessary to obtain estimators with good generalization.

In fact, most works on benign overfitting studied the minimal norm interpolations. In our setting, this corresponds to the solutions of

$$\operatorname{argmin}_{f \in \mathcal{NN}(N_n)} \kappa(f), \quad \text{s.t.} \quad f(X_i) = Y_i, \quad i = 1, \dots, n, \quad (24)$$

which can be viewed as the limiting case of the optimization (14) with $\lambda_n = 0$. The setting of over-parameterized neural networks allows us to make an explicit comparison between interpolating and regularized solutions. We think it is a promising research direction to study rates of convergence for the minimal norm interpolation (24) and compare them with our results. Note that, for linear regression, Hastie et al. (2022) showed that optimally regularized ridge regression dominates the minimum ℓ^2 -norm interpolation in test performances. It is reasonable to expect that such a phenomenon also occurs for neural networks (since the regularized least squares (14) achieves the minimax rates). On the other hand, some

kernel interpolations have been proven to be inconsistent in low dimensions (Buchholz, 2022; Rakhlin and Zhai, 2019). It is an interesting problem to determine whether the minimal norm solutions of (24) are consistent and minimax optimal.

5. Proofs

This section gives the proofs of Theorem 1, Theorem 5, Lemma 6 and Lemma 10.

5.1 Proof of Theorem 1

Recall that S_- is defined by (3), $S_+ = -S_-$ and $S_0 = \mathbb{S}^d \setminus (S_- \cup S_+)$. Notice that S_- is the set of vectors $v \in \mathbb{S}^d$ such that $(x^\top, 1)v \leq 0$ for all $x \in \mathbb{B}^d$. Thus, $v \in \mathbb{S}^d$ is in S_+ if and only if $(x^\top, 1)v \geq 0$ holds for all $x \in \mathbb{B}^d$. And, for any $v \in S_0$, $\sigma((x^\top, 1)v)$ is a nonlinear function on \mathbb{B}^d . Let us first characterize the measures $\tau \in \mathcal{M}_{disc}(\mathbb{S}^d)$ such that $f_\tau = 0$.

Lemma 13 *Let $\tau = \sum_{i=1}^N a_i \delta_{u_i} \in \mathcal{M}_{disc}(\mathbb{S}^d)$, where $u_1, \dots, u_N \in \mathbb{S}^d$ are distinct. Then, $f_\tau(x) = \sum_{i=1}^N a_i \sigma((x^\top, 1)u_i) = 0$ for all $x \in \mathbb{B}^d$ if and only if the following two conditions hold*

- (1) *For any $u_i \in S_0$, there exists $j \neq i$ such that $u_j = -u_i$ and $a_j = -a_i$.*
- (2) $\frac{1}{2} \sum_{u_i \in S_0} a_i u_i + \sum_{u_i \in S_+} a_i u_i = 0$.

Proof Without loss of generality, we assume that $u_i \in S_0 \cup S_+$ and $a_i \neq 0$ for $i = 1, \dots, N$. Notice that $a_i(x^\top, 1)u_i = a_i \sigma((x^\top, 1)u_i) - a_i \sigma(-(x^\top, 1)u_i)$. If τ satisfies condition (1), then

$$\sum_{u_i \in S_0} a_i(x^\top, 1)u_i = 2 \sum_{u_i \in S_0} a_i \sigma((x^\top, 1)u_i).$$

If, in addition, τ satisfies condition (2), then

$$\begin{aligned} f_\tau(x) &= \sum_{u_i \in S_0} a_i \sigma((x^\top, 1)u_i) + \sum_{u_i \in S_+} a_i(x^\top, 1)u_i \\ &= \frac{1}{2} \sum_{u_i \in S_0} a_i(x^\top, 1)u_i + \sum_{u_i \in S_+} a_i(x^\top, 1)u_i = 0. \end{aligned}$$

Conversely, assume $f_\tau = 0$ on \mathbb{B}^d . Notice that $\sigma((x^\top, 1)u_i)$ is a piecewise linear function. We denote its ‘‘break points’’ by $A_i = \{x : \|x\|_2 < 1, (x^\top, 1)u_i = 0\}$. One can check that A_i is non-empty if and only if $u_i \in S_0$. For any $u_i \in S_0$, let j be the index such that $u_j = -u_i$ (it may not exist), then we can find a point $x_i \in A_i = A_j$ such that $x_i \notin A_k$ for any $k \neq i, j$. Consider the decomposition

$$0 = f_\tau(x) = a_i \sigma((x^\top, 1)u_i) + a_j \sigma((x^\top, 1)u_j) + \sum_{k \neq i, j} a_k \sigma((x^\top, 1)u_k).$$

If u_j does not exist, then the first term is nonlinear in a neighborhood of x_i and the remain terms are linear in this neighborhood, which is a contradiction. Hence, u_j exists and the

summation of the first two terms is linear, which implies $a_j = -a_i$ and condition (1) holds. As a consequence,

$$\begin{aligned} 0 = f_\tau(x) &= \frac{1}{2} \sum_{u_i \in S_0} a_i(x^\top, 1)u_i + \sum_{u_i \in S_+} a_i(x^\top, 1)u_i \\ &= (x^\top, 1) \left(\frac{1}{2} \sum_{u_i \in S_0} a_i u_i + \sum_{u_i \in S_+} a_i u_i \right) \end{aligned}$$

holds for all $x \in \mathbb{B}^d$. Thus, $\frac{1}{2} \sum_{u_i \in S_0} a_i u_i + \sum_{u_i \in S_+} a_i u_i = 0$, which proves condition (2). ■

We now give a proof of Theorem 1. For any $f \in \mathcal{NN}(N)$, we can always represent it as

$$f(x) = \sum_{\tilde{v}_i \in S_0} \tilde{c}_i \sigma((x^\top, 1)\tilde{v}_i) + \sum_{\tilde{v}_i \in S_+} \tilde{c}_i (x^\top, 1)\tilde{v}_i. \quad (25)$$

Without loss of generality, we assume that the \tilde{v}_i are distinct. If there exist \tilde{v}_i and \tilde{v}_j satisfying $\tilde{v}_j = -\tilde{v}_i \in S_0$, then

$$\tilde{c}_i \sigma((x^\top, 1)\tilde{v}_i) + \tilde{c}_j \sigma((x^\top, 1)\tilde{v}_j) = (\tilde{c}_i + \tilde{c}_j) \sigma((x^\top, 1)\tilde{v}_i) - \tilde{c}_j (x^\top, 1)\tilde{v}_i.$$

Using this equality, we can reduce (25) to the following form

$$f(x) = \sum_{k=1}^K c_k \sigma((x^\top, 1)v_k) + (x^\top, 1)w, \quad (26)$$

where $0 \leq K \leq N$, $c_1, \dots, c_K \neq 0$, $w \in \mathbb{R}^{d+1}$ and $v_1, \dots, v_K \in S_0$ satisfy $v_i \neq \pm v_k$ for $1 \leq i \neq k \leq K$. Note that, if $K = N$, then $w = 0$. We are going to construct a measure $\mu \in \mathcal{M}_{disc}(\mathbb{S}^d)$ supported on at most N points in $S_0 \cup S_+$ such that $f_\mu = f$ and estimate $\kappa(f)$ by $\|\mu\|$. By Lemma 13, any measure ν satisfies $f_\nu = f_\mu = f$ if and only if $\nu = \mu + \tau$ for some τ satisfying the two conditions in Lemma 13. Hence, $\kappa(f) = \inf_\tau \|\mu + \tau\|$. Without loss of generality, we also assume that τ is supported on $S_0 \cup S_+$. By Lemma 13, we can decompose τ as

$$\begin{aligned} \tau &= \sum_{k=1}^K a_k (\delta_{v_k} - \delta_{-v_k}) + \sum_{j=1}^{J_1} b_j \delta_{u_j} + \sum_{j=J_1+1}^J b_j (\delta_{u_j} - \delta_{-u_j}) \\ &=: \tau_1 + \tau_2 + \tau_3, \end{aligned}$$

where $u_j \in S_+$ for $1 \leq j \leq J_1$, and $u_j \in S_0 \setminus \{\pm v_1, \dots, \pm v_K\}$ for $J_1 + 1 \leq j \leq J$ satisfies $u_j \neq \pm u_i$ for any $i \neq j$ and

$$\sum_{k=1}^K a_k v_k + \sum_{j=1}^J b_j u_j = 0. \quad (27)$$

We continue the proof in four different cases.

Case 1: $w = 0$. We let $\mu = \sum_{k=1}^K c_k \delta_{v_k}$, then

$$\begin{aligned} \|\mu + \tau\| &= \|\mu + \tau_1\| + \|\tau_2 + \tau_3\| \geq \|\mu + \tau_1\| \\ &= \sum_{k=1}^K (|c_k + a_k| + |a_k|) \geq \sum_{k=1}^K |c_k| = \|\mu\|, \end{aligned}$$

which implies $\kappa(f) = \|\mu\|$.

Case 2: $w/\|w\|_2 \in S_+ \cup S_-$. We let $\mu = \sum_{k=1}^K c_k \delta_{v_k} + c_{K+1} \delta_{v_{K+1}} =: \mu_1 + \mu_2$, where, if $w/\|w\|_2 \in S_+$, we take $v_{K+1} = w/\|w\|_2 \in S_+$ and $c_{K+1} = \|w\|_2$; otherwise, we take $v_{K+1} = -w/\|w\|_2 \in S_+$ and $c_{K+1} = -\|w\|_2$. Since $w \neq 0$, $K \leq N-1$ and μ is supported on at most $K+1 \leq N$ points. Let $m \in \{1, \dots, J_1\}$ be the index such that $u_m = v_{K+1}$. (If such m does not exist, we set $m = 0$ and $b_m = 0$ in the following.) Similar to Case 1, we have

$$\begin{aligned} \|\mu + \tau\| &= \|\mu_1 + \tau_1\| + \|\mu_2 + b_m \delta_{u_m}\| + \|\tau_2 - b_m \delta_{u_m} + \tau_3\| \\ &\geq \sum_{k=1}^K (|c_k + a_k| + |a_k|) + |c_{K+1} + b_m| + \sum_{j \neq m} |b_j|. \end{aligned}$$

By equality (27),

$$w = c_{K+1} v_{K+1} = \sum_{k=1}^K a_k v_k + (c_{K+1} + b_m) u_m + \sum_{j \neq m} b_j u_j,$$

which implies

$$\|\mu + \tau\| \geq \sum_{k=1}^K |a_k| + |c_{K+1} + b_m| + \sum_{j \neq m} |b_j| \geq \|w\|_2.$$

Combining with the bound $\|\mu + \tau\| \geq \sum_{k=1}^K |c_k|$, we conclude that

$$\kappa(f) = \inf_{\tau} \|\mu + \tau\| \geq \|w\|_2 \vee \sum_{k=1}^K |c_k|. \quad (28)$$

Therefore,

$$\kappa(f) \geq \frac{1}{2} \sum_{k=1}^K |c_k| + \frac{1}{2} \|w\|_2 = \frac{1}{2} \|\mu\|.$$

Case 3: $w = cv_i$ for some $1 \leq i \leq K$ and $c \neq 0$. We let $\mu = \sum_{k=1}^K c_k \delta_{v_k} + c(\delta_{v_i} - \delta_{-v_i})$, which is supported on at most $K+1 \leq N$ points, then $\|\mu\| = \sum_{k \neq i} |c_k| + |c_i + c| + |c|$ and

$$\begin{aligned} \|\mu + \tau\| &= \|\mu + \tau_1\| + \|\tau_2 + \tau_3\| \\ &\geq \sum_{k \neq i} (|c_k + a_k| + |a_k|) + |c_i + c + a_i| + |c + a_i| + \sum_{j=1}^J |b_j|. \end{aligned}$$

Similar to Case 2, by equality (27),

$$w = cv_i = \sum_{k \neq i} a_k v_k + (c + a_i)v_i + \sum_{j=1}^J b_j u_j,$$

which implies

$$\|\mu + \tau\| \geq \sum_{k \neq i} |a_k| + |c + a_i| + \sum_{j=1}^J |b_j| \geq \|w\|_2.$$

Combining with $\|\mu + \tau\| \geq \sum_{k=1}^K |c_k|$, we can obtain the bound (28). Since $|c| = \|w\|_2$,

$$\|\mu\| = \sum_{k \neq i} |c_k| + |c_i + c| + |c| \leq \sum_{k=1}^K |c_k| + 2\|w\|_2 \leq 3\kappa(f).$$

Case 4: $v_{K+1} := w/\|w\|_2 \in S_0 \setminus \{\pm v_1, \dots, \pm v_K\}$. By the construction of (26), it holds that $K \leq N - 2$. We let $\mu = \sum_{k=1}^K c_k \delta_{v_k} + \|w\|_2(\delta_{v_{K+1}} - \delta_{-v_{K+1}}) =: \mu_1 + \mu_2$, which is supported on $K+2 \leq N$ points. Let $m \in \{J_1+1, \dots, J\}$ be the index such that $u_m = \pm v_{K+1}$. (If such m does not exist, we set $m = 0$ and $b_m = 0$ in the following.) Then,

$$\begin{aligned} \|\mu + \tau\| &= \|\mu_1 + \tau_1\| + \|\mu_2 + b_m(\delta_{u_m} - \delta_{-u_m})\| + \|\tau_2 + \tau_3 - b_m(\delta_{u_m} - \delta_{-u_m})\| \\ &\geq \sum_{k=1}^K (|c_k + a_k| + |a_k|) + \| \|w\|_2 \pm b_m \| + \sum_{j \neq m} |b_j|. \end{aligned}$$

Similar to Case 2, by using equality (27), we can obtain the bound (28). Therefore,

$$\|\mu\| = \sum_{k=1}^K |c_k| + 2\|w\|_2 \leq 3\kappa(f).$$

In summary, we have shown that, for the parameterization (26), there exists a measure μ supported on at most N points such that $f_\mu = f$ and

$$\|w\|_2 \vee \sum_{k=1}^K |c_k| \leq \kappa(f) \leq \|\mu\| \leq \sum_{k=1}^K |c_k| + 2\|w\|_2.$$

It is easy to construct $\theta \in \mathbb{R}^{(d+2)N}$ such that $f_\mu = f_\theta$ and $\kappa(\theta) = \|\mu\|$. Hence, $\kappa(f) \leq \inf\{\kappa(\theta) : f = f_\theta, \theta \in \mathbb{R}^{(d+2)N}\} \leq 3\kappa(f)$.

5.2 Proof of Theorem 5

Without loss of generality, we can assume that $\varsigma = 1$ by rescaling. Note that the sub-Gaussian assumption $\mathbb{E}[\exp(\lambda \xi_i)] \leq \exp(\lambda^2/2)$ implies $\mathbb{E}[\xi_i] = 0$ and $\mathbb{E}[\xi_i^2] \leq 1$. It is well known that the Gaussian complexity can be bounded by Dudley's entropy integral (Wainwright, 2019, Section 5.3.3). We will provide a bound for the entropy by using Theorem 2.

The derivation of the entropy integral bound (30) below uses essentially the same argument as in Srebro et al. (2010, Lemma A.3), which we repeat for the convenience of the

reader. For any $\epsilon > 0$, let us denote the ϵ -covering number of the set \mathcal{F} in $L^2(\mu_n)$ norm by $\mathcal{N}(\epsilon, \mathcal{F}, \|\cdot\|_{L^2(\mu_n)})$. In other words, there exists a subset $\{f_1, \dots, f_m\} \subseteq \mathcal{F}$ of minimal size $m = \mathcal{N}(\epsilon, \mathcal{F}, \|\cdot\|_{L^2(\mu_n)})$ such that for any $f \in \mathcal{F}$, there exists f_j satisfying $\|f - f_j\|_{L^2(\mu_n)} \leq \epsilon$, which implies

$$\begin{aligned} \left| \frac{1}{n} \sum_{i=1}^n \xi_i f(X_i) \right| &\leq \left| \frac{1}{n} \sum_{i=1}^n \xi_i f_j(X_i) \right| + \left| \frac{1}{n} \sum_{i=1}^n \xi_i (f(X_i) - f_j(X_i)) \right| \\ &\leq \max_{j=1, \dots, m} \left| \frac{1}{n} \sum_{i=1}^n \xi_i f_j(X_i) \right| + \sqrt{\frac{\sum_{i=1}^n \xi_i^2}{n}} \epsilon, \end{aligned}$$

where we use the Cauchy-Schwarz inequality in the second step. Taking the supremum over $f \in \mathcal{F}$ and then taking expectations over the noise, we obtain

$$\mathbb{E}_{\xi_{1:n}} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n \xi_i f(X_i) \right| \right] \leq \mathbb{E}_{\xi_{1:n}} \left[\max_{j=1, \dots, m} \left| \frac{1}{n} \sum_{i=1}^n \xi_i f_j(X_i) \right| \right] + \epsilon, \quad (29)$$

where we have used $\mathbb{E}[(\sum_{i=1}^n \xi_i^2)^{1/2}] \leq (\sum_{i=1}^n \mathbb{E}[\xi_i^2])^{1/2} \leq \sqrt{n}$. Define a family of zero-mean random variables indexed by $f \in \mathcal{F}$ as $Z(f) := \frac{1}{\sqrt{n}} \sum_{i=1}^n \xi_i f(X_i)$. Then, $Z(f)$ is a sub-Gaussian process with respect to the metric $\rho_Z(f, f') := \|f - f'\|_{L^2(\mu_n)}$, meaning that

$$\mathbb{E}[\exp(\lambda(Z(f) - Z(f')))] \leq \exp(\lambda^2 \rho_Z(f, f')^2 / 2), \quad \forall f, f' \in \mathcal{F}, \lambda \in \mathbb{R},$$

which can be proven by basic properties of sub-Gaussian variables (Wainwright, 2019, Exercise 2.13). Observe that the first term on the right-hand side of (29) is an expected supremum of the sub-Gaussian process $Z(f)$. We can apply a chaining argument as Wainwright (2019, Proof of Theorem 5.22) to show that

$$\mathbb{E}_{\xi_{1:n}} \left[\max_{j=1, \dots, m} |Z(f_j)| \right] \leq 16 \int_{\epsilon/4}^{D/2} \sqrt{\log \mathcal{N}(t, \mathcal{F}, \|\cdot\|_{L^2(\mu_n)})} dt,$$

where $D := \sup_{f, f' \in \mathcal{F}} \rho_Z(f, f')$ denotes the diameter. Therefore,

$$\mathbb{E}_{\xi_{1:n}} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n \xi_i f(X_i) \right| \right] \leq \inf_{\epsilon \geq 0} \left\{ 4\epsilon + \frac{16}{\sqrt{n}} \int_{\epsilon}^{D/2} \sqrt{\log \mathcal{N}(t, \mathcal{F}, \|\cdot\|_{L^2(\mu_n)})} dt \right\}. \quad (30)$$

In order to apply this bound to the variation space $\mathcal{F}_\sigma(M)$, we need to estimate its covering number. Since $\|f - f'\|_{L^2(\mu_n)} \leq \|f - f'\|_{L^\infty(\mathbb{B}^d)}$, it is enough to provide a bound for the entropy $\log \mathcal{N}(\epsilon, \mathcal{F}_\sigma(M), \|\cdot\|_{L^\infty(\mathbb{B}^d)})$. Note that this bound should be independent of the network width. But we will need to first bound the entropy of $\mathcal{NN}(N, M)$ by N and M . Recall that any $f \in \mathcal{NN}(N, M)$ can be represented in the form

$$f(x) = \sum_{i=1}^N a_i \sigma((x^\top, 1)v_i), \quad v_i \in \mathbb{S}^d, \quad \sum_{i=1}^N |a_i| \leq M.$$

It is easy to see that there exists $\epsilon/(2\sqrt{2}M)$ -cover V_ϵ of \mathbb{S}^d in the metric $\|\cdot\|_2$ with cardinality $|V_\epsilon| \lesssim (1 + M/\epsilon)^d$. Similarly, there exists $\epsilon/(2\sqrt{2})$ -cover A_ϵ of the set $A = \{a = (a_1, \dots, a_N) :$

$\|a\|_1 \leq M\}$ in the metric $\|\cdot\|_1$ such that $|A_\epsilon| \lesssim (1 + M/\epsilon)^N$. Hence, for any $v_i \in \mathbb{S}^d$ and $a = (a_1, \dots, a_N) \in A$, we can choose $\tilde{v}_i \in V_\epsilon$ and $\tilde{a} = (\tilde{a}_1, \dots, \tilde{a}_N) \in A_\epsilon$ such that $\|v_i - \tilde{v}_i\|_2 \leq \epsilon/(2\sqrt{2}M)$ and $\|a - \tilde{a}\|_1 \leq \epsilon/(2\sqrt{2})$. Therefore, letting $f(x) = \sum_{i=1}^N \tilde{a}_i \sigma((x^\top, 1)\tilde{v}_i)$, we have for $x \in \mathbb{B}^d$,

$$\begin{aligned} |f(x) - \tilde{f}(x)| &\leq \|a - \tilde{a}\|_1 \sup_{1 \leq i \leq N} |\sigma((x^\top, 1)v_i)| + \|\tilde{a}\|_1 \sup_{1 \leq i \leq N} |\sigma((x^\top, 1)v_i) - \sigma((x^\top, 1)\tilde{v}_i)| \\ &\leq \epsilon/2 + \sqrt{2}M \sup_{1 \leq i \leq N} \|v_i - \tilde{v}_i\|_2 \leq \epsilon, \end{aligned}$$

where we use the Lipschitz continuity of ReLU in the second inequality. Thus,

$$\log \mathcal{N}(\epsilon, \mathcal{NN}(N, M), \|\cdot\|_{L^\infty(\mathbb{B}^d)}) \lesssim N \log(1 + M/\epsilon). \quad (31)$$

Now, we estimate the entropy of $\mathcal{F}_\sigma(M)$. By Theorem 2,

$$\sup_{f \in \mathcal{F}_\sigma(M)} \inf_{f_N \in \mathcal{NN}(N, M)} \|f - f_N\|_{L^\infty(\mathbb{B}^d)} \lesssim MN^{-\frac{d+3}{2d}}.$$

Hence, for any $0 < \epsilon \leq M$, we can choose $N \asymp (\epsilon/M)^{-2d/(d+3)}$ such that for any $f \in \mathcal{F}_\sigma(M)$, there exists $f_N \in \mathcal{NN}(N, M)$ satisfying $\|f - f_N\|_{L^\infty(\mathbb{B}^d)} < \epsilon/2$. By the triangle inequality, any $\epsilon/2$ -cover of $\mathcal{NN}(N, M)$ is an ϵ -cover of $\mathcal{F}_\sigma(M)$. Then, the estimate (31) shows

$$\begin{aligned} \log \mathcal{N}(\epsilon, \mathcal{F}_\sigma(M), \|\cdot\|_{L^\infty(\mathbb{B}^d)}) &\leq \log \mathcal{N}(\epsilon/2, \mathcal{NN}(N, M), \|\cdot\|_{L^\infty(\mathbb{B}^d)}) \\ &\lesssim N \log(1 + 2M/\epsilon) \\ &\lesssim (\epsilon/M)^{-2d/(d+3)} \log(1 + M/\epsilon). \end{aligned}$$

Applying (30) to $\mathcal{F} = \{f \in \mathcal{F}_\sigma(M) : \|f\|_{L^2(\mu_n)} \leq \delta\}$, we have $D = 2\delta$ and

$$\mathcal{G}_n(\mathcal{F}_\sigma(M); \delta, \xi_{1:n}) \lesssim \inf_{\epsilon \geq 0} \left\{ 4\epsilon + \frac{16}{\sqrt{n}} \int_\epsilon^\delta (t/M)^{-d/(d+3)} \sqrt{\log(1 + M/t)} dt \right\}.$$

If we choose $\epsilon \asymp \delta^{\frac{3}{d+3}} M^{\frac{d}{d+3}} n^{-1/2}$, then for $n \geq 2$,

$$\begin{aligned} \mathcal{G}_n(\mathcal{F}_\sigma(M); \delta, \xi_{1:n}) &\lesssim \frac{\delta^{\frac{3}{d+3}} M^{\frac{d}{d+3}}}{\sqrt{n}} + \frac{M^{\frac{d}{d+3}}}{\sqrt{n}} \sqrt{\log(nM/\delta)} \int_0^\delta t^{-d/(d+3)} dt \\ &\lesssim \frac{\delta^{\frac{3}{d+3}} M^{\frac{d}{d+3}}}{\sqrt{n}} \sqrt{\log(nM/\delta)}. \end{aligned}$$

Finally, for the star hull $\text{star}(\mathcal{T}_B \mathcal{F}_\sigma(M))$, we can similarly derive the bound by estimating the covering number. It is easy to see that the covering number of $\mathcal{T}_B \mathcal{F}_\sigma(M)$ is not larger than that of $\mathcal{F}_\sigma(M)$. For any $\epsilon > 0$, let $\{f_1, \dots, f_m\} \subseteq \mathcal{T}_B \mathcal{F}_\sigma(M)$ be an $\epsilon/2$ -cover of $\mathcal{T}_B \mathcal{F}_\sigma(M)$ and $\{a_1, \dots, a_k\} \subseteq [0, 1]$ be an $\epsilon/(2B)$ -cover of $[0, 1]$. Then, for any $af \in \text{star}(\mathcal{T}_B \mathcal{F}_\sigma(M))$, we can choose $a_i f_j$ such that

$$\|af - a_i f_j\|_{L^\infty(\mathbb{B}^d)} \leq |a - a_i| \|f\|_{L^\infty(\mathbb{B}^d)} + |a_i| \|f - f_j\|_{L^\infty(\mathbb{B}^d)} \leq \epsilon.$$

This shows that $\{a_i f_j : 1 \leq i \leq k, 1 \leq j \leq m\}$ is ϵ -cover of $\text{star}(\mathcal{T}_B \mathcal{F}_\sigma(M))$, which implies

$$\begin{aligned} & \log \mathcal{N}(\epsilon, \text{star}(\mathcal{T}_B \mathcal{F}_\sigma(M)), \|\cdot\|_{L^\infty(\mathbb{B}^d)}) \\ & \leq \log \mathcal{N}(\epsilon, \mathcal{F}_\sigma(M), \|\cdot\|_{L^\infty(\mathbb{B}^d)}) + \log(1 + B/\epsilon) \\ & \lesssim (\epsilon/M)^{-2d/(d+3)} \log(1 + M/\epsilon). \end{aligned}$$

Hence, we can obtain the same bound for the local complexity.

5.3 Proof of Lemma 6

If $f \in \mathcal{F}_n$ satisfies $\|\widehat{f}_n - f\|_{L^2(\mu_n)} \leq \delta_n$, then

$$\begin{aligned} \|\widehat{f}_n - h\|_{L^2(\mu_n)}^2 & \leq 2\|f - h\|_{L^2(\mu_n)}^2 + 2\|\widehat{f}_n - f\|_{L^2(\mu_n)}^2 \\ & \leq 2\|f - h\|_{L^2(\mu_n)}^2 + 2\delta_n^2. \end{aligned}$$

Thus, we can assume $\|\widehat{f}_n - f\|_{L^2(\mu_n)} > \delta_n$.

Observing that $\widehat{f}_n - f \in \partial \mathcal{F}_n$, we consider the event

$$\mathcal{A} := \left\{ \exists g \in \partial \mathcal{F}_n : \|g\|_{L^2(\mu_n)} > \delta_n, \left| \frac{1}{n} \sum_{i=1}^n \eta_i g(X_i) \right| > 2\delta_n \|g\|_{L^2(\mu_n)} \right\}.$$

If \mathcal{A} is true, then there exists $\tilde{g} := \frac{\delta_n}{\|g\|_{L^2(\mu_n)}} g \in \text{star}(\partial \mathcal{F}_n)$ such that $\|\tilde{g}\|_{L^2(\mu_n)} = \delta_n$ and

$$\left| \frac{1}{n} \sum_{i=1}^n \eta_i \tilde{g}(X_i) \right| = \frac{\delta_n}{\|g\|_{L^2(\mu_n)}} \left| \frac{1}{n} \sum_{i=1}^n \eta_i g(X_i) \right| > 2\delta_n^2.$$

This implies that $\mathbb{P}(\mathcal{A}) \leq \mathbb{P}(Z_n > 2\delta_n^2)$, where Z_n is the random variable defined by

$$Z_n := \sup_{\substack{g \in \text{star}(\partial \mathcal{F}_n) \\ \|g\|_{L^2(\mu_n)} \leq \delta_n}} \left| \frac{1}{n} \sum_{i=1}^n \eta_i g(X_i) \right|.$$

Observe that we can view $Z_n(\eta_{1:n})$ as a convex Lipschitz function of the random vector $\eta_{1:n}$ supported on $[-2B, 2B]^n$ with Lipschitz constant δ_n/\sqrt{n} . By the concentration for convex Lipschitz functions on bounded random variables (Boucheron et al., 2013, Theorem 6.10), we conclude that, for any $t > 0$,

$$\mathbb{P}(Z_n > \mathbb{E}[Z_n] + t) \leq \exp\left(-\frac{nt^2}{32B^2\delta_n^2}\right). \quad (32)$$

Taking $t = \delta_n^2$ and noticing that $\mathbb{E}[Z_n] = \mathcal{G}_n(\text{star}(\partial \mathcal{F}_n); \delta_n, \eta_{1:n}) \leq \delta_n^2$ by assumption, we get

$$\mathbb{P}(\mathcal{A}) \leq \mathbb{P}(Z_n > 2\delta_n^2) \leq \mathbb{P}(Z_n > \mathbb{E}[Z_n] + \delta_n^2) \leq \exp\left(-\frac{n\delta_n^2}{32B^2}\right).$$

By the definition of the event \mathcal{A} and the assumption $\|\widehat{f}_n - f\|_{L^2(\mu_n)} > \delta_n$, the inequality

$$\left| \frac{1}{n} \sum_{i=1}^n \eta_i \left(\widehat{f}_n(X_i) - f(X_i) \right) \right| \leq 2\delta_n \|\widehat{f}_n - f\|_{L^2(\mu_n)}$$

holds with probability at least $1 - \exp(-n\delta_n^2/(32B^2))$. Combining this bound with the base inequality (10), we have

$$\begin{aligned} \|\widehat{f}_n - h\|_{L^2(\mu_n)}^2 &\leq \|f - h\|_{L^2(\mu_n)}^2 + 4\delta_n \|\widehat{f}_n - f\|_{L^2(\mu_n)} \\ &\leq \|f - h\|_{L^2(\mu_n)}^2 + 16\delta_n^2 + \frac{1}{4} \|\widehat{f}_n - f\|_{L^2(\mu_n)}^2 \\ &\leq \|f - h\|_{L^2(\mu_n)}^2 + 16\delta_n^2 + \frac{1}{2} \|\widehat{f}_n - h\|_{L^2(\mu_n)}^2 + \frac{1}{2} \|f - h\|_{L^2(\mu_n)}^2 \\ &= \frac{3}{2} \|f - h\|_{L^2(\mu_n)}^2 + 16\delta_n^2 + \frac{1}{2} \|\widehat{f}_n - h\|_{L^2(\mu_n)}^2. \end{aligned}$$

Rearranging yields the desired bound.

5.4 Proof of Lemma 10

We modify the proofs of Lemma 6 and Wainwright (2019, Theorem 13.17). For any $f \in \mathcal{F}_{n,R}$, let us denote $g_f := \widehat{f}_{n,\lambda_n} - f$ for convenience. We divide the proof into four cases.

Case 1: $\|g_f\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(g_f) \leq \delta_n^2$. We have

$$\begin{aligned} &\|\widehat{f}_{n,\lambda_n} - h\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(\widehat{f}_{n,\lambda_n}) \\ &\leq 2\|f - h\|_{L^2(\mu_n)}^2 + 2\|g_f\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(g_f) + \lambda_n \kappa(f) \\ &\leq 2\|f - h\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(f) + 2\delta_n^2. \end{aligned}$$

Case 2: $\|g_f\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(g_f) > \delta_n^2$ and $\kappa(\widehat{f}_{n,\lambda_n}) \leq R$. In this case, $g_f = \widehat{f}_{n,\lambda_n} - f \in \partial\mathcal{F}_{n,R}$. Consider the event

$$\mathcal{A}_1 := \left\{ \exists g \in \partial\mathcal{F}_{n,R} : \|g\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(g) > \delta_n^2, \left| \frac{1}{n} \sum_{i=1}^n \eta_i g(X_i) \right| > 2\delta_n \sqrt{\|g\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(g)} \right\}.$$

If \mathcal{A}_1 is true, then there exists

$$\widetilde{g} := \frac{\delta_n}{\sqrt{\|g\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(g)}} g \in \text{star}(\partial\mathcal{F}_{n,R}),$$

such that $\|\widetilde{g}\|_{L^2(\mu_n)} \leq \delta_n$ and

$$\left| \frac{1}{n} \sum_{i=1}^n \eta_i \widetilde{g}(X_i) \right| = \frac{\delta_n}{\sqrt{\|g\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(g)}} \left| \frac{1}{n} \sum_{i=1}^n \eta_i g(X_i) \right| > 2\delta_n^2.$$

This implies that $\mathbb{P}(\mathcal{A}_1) \leq \mathbb{P}(Z_n(\delta_n) > 2\delta_n^2)$, where $Z_n(\delta_n)$ is the random variable defined by

$$Z_n(\delta_n) := \sup_{\substack{g \in \text{star}(\partial\mathcal{F}_{n,R}) \\ \|g\|_{L^2(\mu_n)} \leq \delta_n}} \left| \frac{1}{n} \sum_{i=1}^n \eta_i g(X_i) \right|.$$

As in the proof of Lemma 6, by viewing $Z_n(\delta_n)$ as a convex Lipschitz function of $\eta_{1:n}$, we get, for any $u > 0$,

$$\mathbb{P}(Z_n(\delta_n) > \mathbb{E}[Z_n(\delta_n)] + u) \leq \exp\left(-\frac{nu^2}{32B^2\delta_n^2}\right).$$

Taking $u = \delta_n^2$ and noticing that $\mathbb{E}[Z_n(\delta_n)] = \mathcal{G}_n(\text{star}(\partial\mathcal{F}_{n,R}; \delta_n, \eta_{1:n}) \leq \delta_n^2$, we conclude

$$\mathbb{P}(\mathcal{A}_1) \leq \mathbb{P}(Z_n(\delta_n) > 2\delta_n^2) \leq \mathbb{P}(Z_n(\delta_n) > \mathbb{E}[Z_n(\delta_n)] + \delta_n^2) \leq \exp\left(-\frac{n\delta_n^2}{32B^2}\right).$$

This implies that the inequality

$$\left| \frac{1}{n} \sum_{i=1}^n \eta_i g_f(X_i) \right| \leq 2\delta_n \sqrt{\|g_f\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(g_f)}$$

holds with probability at least $1 - \exp(-n\delta_n^2/(32B^2))$. Combining this with inequality (18), we have

$$\begin{aligned} & \|\widehat{f}_{n,\lambda_n} - h\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(\widehat{f}_{n,\lambda_n}) \\ & \leq c_0 \left(\|f - h\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(f) \right) + 4\delta_n \sqrt{\|g_f\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(g_f)} \\ & \leq c_0 \left(\|f - h\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(f) \right) + 16\delta_n^2 + \frac{1}{4} \|\widehat{f}_{n,\lambda_n} - f\|_{L^2(\mu_n)}^2 + \frac{1}{4} \lambda_n \kappa(\widehat{f}_{n,\lambda_n} - f) \\ & \leq c_0 \left(\|f - h\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(f) \right) + 16\delta_n^2 + \frac{1}{2} \|\widehat{f}_{n,\lambda_n} - h\|_{L^2(\mu_n)}^2 + \frac{1}{2} \|f - h\|_{L^2(\mu_n)}^2 \\ & \quad + \frac{1}{4} \lambda_n \kappa(\widehat{f}_{n,\lambda_n}) + \frac{1}{4} \lambda_n \kappa(f) \\ & \leq \frac{1+2c_0}{2} \left(\|f - h\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(f) \right) + 16\delta_n^2 + \frac{1}{2} \left(\|\widehat{f}_{n,\lambda_n} - h\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(\widehat{f}_{n,\lambda_n}) \right). \end{aligned}$$

Rearranging shows that the desired bound holds with probability at least $1 - \exp(-n\delta_n^2/(32B^2))$.

Case 3: $\kappa(\widehat{f}_{n,\lambda_n}) > R$ and $\|g_f\|_{L^2(\mu_n)} \leq \delta_n \kappa(\widehat{f}_{n,\lambda_n})/R$. Observe that

$$\begin{aligned} \widetilde{g}_f &:= \frac{R}{\kappa(\widehat{f}_{n,\lambda_n})} g_f = \frac{R}{\kappa(\widehat{f}_{n,\lambda_n})} (\widehat{f}_{n,\lambda_n} - f) \in \partial\mathcal{F}_{n,R}, \\ \|\widetilde{g}_f\|_{L^2(\mu_n)} &= \frac{R}{\kappa(\widehat{f}_{n,\lambda_n})} \|g_f\|_{L^2(\mu_n)} \leq \delta_n. \end{aligned}$$

We have shown in Case 2 that $\mathbb{P}(Z_n(\delta_n) > 2\delta_n^2) \leq \exp(-n\delta_n^2/(32B^2))$. This implies that the inequality

$$\left| \frac{1}{n} \sum_{i=1}^n \eta_i g_f(X_i) \right| = \frac{\kappa(\widehat{f}_{n,\lambda_n})}{R} \left| \frac{1}{n} \sum_{i=1}^n \eta_i \widetilde{g}_f(X_i) \right| \leq 2\delta_n^2 \frac{\kappa(\widehat{f}_{n,\lambda_n})}{R}$$

holds with probability at least $1 - \exp(-n\delta_n^2/(32B^2))$. By inequality (18),

$$\begin{aligned} \|\widehat{f}_{n,\lambda_n} - h\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(\widehat{f}_{n,\lambda_n}) &\leq c_0 \left(\|f - h\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(f) \right) + \frac{4\delta_n^2}{R} \kappa(\widehat{f}_{n,\lambda_n}) \\ &\leq c_0 \left(\|f - h\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(f) \right) + \frac{1}{2} \lambda_n \kappa(\widehat{f}_{n,\lambda_n}), \end{aligned}$$

where we use the assumption $\lambda_n \geq 8\delta_n^2/R$ in the last inequality. Therefore,

$$\|\widehat{f}_{n,\lambda_n} - h\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(\widehat{f}_{n,\lambda_n}) \leq 2c_0 \left(\|f - h\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(f) \right),$$

holds with probability at least $1 - \exp(-n\delta_n^2/(32B^2))$.

Case 4: $\kappa(\widehat{f}_{n,\lambda_n}) > R$ and $\|g_f\|_{L^2(\mu_n)} > \delta_n \kappa(\widehat{f}_{n,\lambda_n})/R$. In this case,

$$\widetilde{g}_f = \frac{R}{\kappa(\widehat{f}_{n,\lambda_n})} g_f \in \partial\mathcal{F}_{n,R}, \quad \|\widetilde{g}_f\|_{L^2(\mu_n)} > \delta_n.$$

We are going to show that the event

$$\mathcal{A}_2 := \left\{ \exists g \in \partial\mathcal{F}_{n,R} : \|g\|_{L^2(\mu_n)} > \delta_n, \left| \frac{1}{n} \sum_{i=1}^n \eta_i g(X_i) \right| > 2\delta_n \|g\|_{L^2(\mu_n)} + \frac{1}{16} \|g\|_{L^2(\mu_n)}^2 \right\}$$

holds with small probability. We will prove it by a ‘‘peeling’’ argument. To do this, let us denote $t_m = 2^m \delta_n$ for $m = 0, 1, \dots$, and

$$\begin{aligned} \mathcal{A}_2(t_m) := &\left\{ \exists g \in \partial\mathcal{F}_{n,R} : t_m < \|g\|_{L^2(\mu_n)} \leq 2t_m, \right. \\ &\left. \left| \frac{1}{n} \sum_{i=1}^n \eta_i g(X_i) \right| > 2\delta_n \|g\|_{L^2(\mu_n)} + \frac{1}{16} \|g\|_{L^2(\mu_n)}^2 \right\}, \end{aligned}$$

then we have the decomposition $\mathcal{A}_2 = \cup_{m=0}^{\infty} \mathcal{A}_2(t_m)$. To estimate $\mathbb{P}(\mathcal{A}_2(t_m))$, we consider the random variables

$$Z_n(t) := \sup_{\substack{g \in \text{star}(\partial\mathcal{F}_{n,R}) \\ \|g\|_{L^2(\mu_n)} \leq t}} \left| \frac{1}{n} \sum_{i=1}^n \eta_i g(X_i) \right|, \quad t \geq \delta_n.$$

Notice that, if the event $\mathcal{A}_2(t_m)$ holds, then there exists some function $g \in \partial\mathcal{F}_{n,R}$ with $\|g\|_{L^2(\mu_n)} \in (t_m, t_{m+1}]$ such that

$$\begin{aligned} \left| \frac{1}{n} \sum_{i=1}^n \eta_i g(X_i) \right| &> 2\delta_n \|g\|_{L^2(\mu_n)} + \frac{1}{16} \|g\|_{L^2(\mu_n)}^2 \\ &\geq 2t_m \delta_n + \frac{1}{16} t_m^2 = t_{m+1} \delta_n + \frac{1}{64} t_{m+1}^2, \end{aligned}$$

where we use $t_{m+1} = 2t_m$. This lower bound implies

$$\mathbb{P}(\mathcal{A}_2(t_m)) \leq \mathbb{P} \left(Z_n(t_{m+1}) > t_{m+1} \delta_n + \frac{1}{64} t_{m+1}^2 \right).$$

As in Case 2, by viewing $Z_n(t)$ as a convex Lipschitz function of $\eta_{1:n}$, one obtain

$$\mathbb{P}(Z_n(t) > \mathbb{E}[Z_n(t)] + u) \leq \exp\left(-\frac{nu^2}{32B^2t^2}\right), \quad \forall u > 0.$$

Since $\text{star}(\partial\mathcal{F}_{n,R})$ is star-shaped, the function $\delta \mapsto \mathcal{G}_n(\text{star}(\partial\mathcal{F}_{n,R}); \delta, \eta_{1:n})/\delta$ is non-increasing. Hence, for any $t \geq \delta_n$,

$$\mathbb{E}[Z_n(t)] = \mathcal{G}_n(\text{star}(\partial\mathcal{F}_{n,R}); t, \eta_{1:n}) \leq t \frac{\mathcal{G}_n(\text{star}(\partial\mathcal{F}_{n,R}); \delta_n, \eta_{1:n})}{\delta_n} \leq t\delta_n.$$

Using this upper bound on the mean and setting $u = t^2/64$, we get

$$\mathbb{P}\left(Z_n(t) > t\delta_n + \frac{1}{64}t^2\right) \leq \exp\left(-\frac{nt^2}{2^{17}B^2}\right).$$

As a consequence,

$$\begin{aligned} \mathbb{P}(\mathcal{A}_2) &\leq \sum_{m=0}^{\infty} \mathbb{P}(\mathcal{A}_2(t_m)) \leq \sum_{m=1}^{\infty} \mathbb{P}\left(Z_n(t_m) > t_m\delta_n + \frac{1}{64}t_m^2\right) \\ &\leq \sum_{m=1}^{\infty} \exp\left(-\frac{nt_m^2}{2^{17}B^2}\right) = \sum_{m=1}^{\infty} \exp\left(-\frac{2^{2m}n\delta_n^2}{2^{17}B^2}\right) \\ &\leq c_1 \exp\left(-\frac{c_2n\delta_n^2}{B^2}\right), \end{aligned}$$

for some constant $c_1, c_2 > 0$.

By the definition of the event \mathcal{A}_2 and the assumption that $\tilde{g}_f \in \partial\mathcal{F}_{n,R}$ and $\|\tilde{g}_f\|_{L^2(\mu_n)} > \delta_n$, we have shown

$$\left|\frac{1}{n} \sum_{i=1}^n \eta_i \tilde{g}_f(X_i)\right| \leq 2\delta_n \|\tilde{g}_f\|_{L^2(\mu_n)} + \frac{1}{16} \|\tilde{g}_f\|_{L^2(\mu_n)}^2$$

holds with probability at least $1 - c_1 \exp(-c_2n\delta_n^2/B^2)$. Multiplying both sides by $\kappa(\hat{f}_{n,\lambda_n})/R$, we obtain

$$\begin{aligned} \left|\frac{1}{n} \sum_{i=1}^n \eta_i g_f(X_i)\right| &\leq 2\delta_n \|g_f\|_{L^2(\mu_n)} + \frac{1}{16} \frac{R}{\kappa(\hat{f}_{n,\lambda_n})} \|g_f\|_{L^2(\mu_n)}^2 \\ &\leq 2\delta_n \|g_f\|_{L^2(\mu_n)} + \frac{1}{16} \|g_f\|_{L^2(\mu_n)}^2, \end{aligned}$$

where we use $\kappa(\hat{f}_{n,\lambda_n}) > R$ in the last inequality. Combining this bound with inequality (18), we have

$$\begin{aligned} &\|\hat{f}_{n,\lambda_n} - h\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(\hat{f}_{n,\lambda_n}) \\ &\leq c_0 \left(\|f - h\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(f)\right) + 4\delta_n \|g_f\|_{L^2(\mu_n)} + \frac{1}{8} \|g_f\|_{L^2(\mu_n)}^2 \\ &\leq c_0 \left(\|f - h\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(f)\right) + 32\delta_n^2 + \frac{1}{4} \|g_f\|_{L^2(\mu_n)}^2 \\ &\leq c_0 \left(\|f - h\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(f)\right) + 32\delta_n^2 + \frac{1}{2} \|\hat{f}_{n,\lambda_n} - h\|_{L^2(\mu_n)}^2 + \frac{1}{2} \|f - h\|_{L^2(\mu_n)}^2 \\ &\leq \frac{1+2c_0}{2} \left(\|f - h\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(f)\right) + 32\delta_n^2 + \frac{1}{2} \left(\|\hat{f}_{n,\lambda_n} - h\|_{L^2(\mu_n)}^2 + \lambda_n \kappa(\hat{f}_{n,\lambda_n})\right), \end{aligned}$$

which yields the desired bound by rearrangements.

Acknowledgments

The work described in this paper was partially supported by InnoHK initiative, The Government of the HKSAR, Laboratory for AI-Powered Financial Technologies, the Research Grants Council of Hong Kong [Projects No. CityU 11306220 and 11308020] and National Natural Science Foundation of China [Project No. 12371103] when the second author worked at City University of Hong Kong. We thank the referees for their helpful comments and suggestions on the paper.

References

- Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via over-parameterization. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pages 242–252. PMLR, 2019.
- Francis Bach. Breaking the curse of dimensionality with convex neural networks. *Journal of Machine Learning Research*, 18(19):1–53, 2017.
- Andrew R. Barron. Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information Theory*, 39(3):930–945, 1993.
- Peter L. Bartlett and Shahar Mendelson. Rademacher and Gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3:463–482, 2002.
- Peter L. Bartlett, Olivier Bousquet, and Shahar Mendelson. Local Rademacher complexities. *The Annals of Statistics*, 33(4), 2005.
- Peter L. Bartlett, Philip M. Long, Gábor Lugosi, and Alexander Tsigler. Benign overfitting in linear regression. *Proceedings of the National Academy of Sciences*, 117(48):30063–30070, 2020.
- Peter L. Bartlett, Andrea Montanari, and Alexander Rakhlin. Deep learning: a statistical viewpoint. *Acta Numerica*, 30:87–201, 2021.
- Francesca Bartolucci, Ernesto De Vito, Lorenzo Rosasco, and Stefano Vigogna. Understanding neural networks with reproducing kernel Banach spaces. *Applied and Computational Harmonic Analysis*, 62:194–236, 2023.
- Benedikt Bauer and Michael Kohler. On deep learning as a remedy for the curse of dimensionality in nonparametric regression. *The Annals of Statistics*, 47(4):2261–2285, 2019.
- Mikhail Belkin. Fit without fear: remarkable mathematical phenomena of deep learning through the prism of interpolation. *Acta Numerica*, 30:203–248, 2021.

- Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32):15849–15854, 2019a.
- Mikhail Belkin, Alexander Rakhlin, and Alexandre B Tsybakov. Does data interpolation contradict statistical optimality? In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1611–1619. PMLR, 2019b.
- Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 2013.
- Simon Buchholz. Kernel interpolation in Sobolev spaces is not consistent in low dimensions. In *Proceedings of the 35th Conference on Learning Theory*, volume 178, pages 3410–3440. PMLR, 2022.
- Minshuo Chen, Haoming Jiang, Wenjing Liao, and Tuo Zhao. Nonparametric regression on low-dimensional manifolds using deep ReLU networks: function approximation and statistical recovery. *Information and Inference: A Journal of the IMA*, 11(4):1203–1253, 2022.
- Lénaïc Chizat, Edouard Oyallon, and Francis R. Bach. On lazy training in differentiable programming. In *Advances in Neural Information Processing Systems*, pages 2933–2943, 2019.
- Simon Du, Jason Lee, Haochuan Li, Liwei Wang, and Xiyu Zhai. Gradient descent finds global minima of deep neural networks. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pages 1675–1685. PMLR, 2019.
- Weinan E, Chao Ma, and Lei Wu. A priori estimates of the population risk for two-layer neural networks. *Communications in Mathematical Sciences*, 17(5):1407–1425, 2019.
- Weinan E, Chao Ma, and Lei Wu. The Barron space and the flow-induced function spaces for neural network models. *Constructive Approximation*, 55(1):369–406, 2022.
- Gerald B. Folland. *Real Analysis: Modern Techniques and Their Applications*. John Wiley & Sons, second edition, 1999.
- Noah Golowich, Alexander Rakhlin, and Ohad Shamir. Size-independent sample complexity of neural networks. *Information and Inference: A Journal of the IMA*, 9(2):473–504, 2020.
- Nathael Gozlan, Cyril Roberto, Paul-Marie Samson, and Prasad Tetali. Kantorovich duality for general transport costs and applications. *Journal of Functional Analysis*, 273(11):3327–3405, 2017.
- László Györfi, Michael Köhler, Adam Krzyżak, and Harro Walk. *A Distribution-free Theory of Nonparametric Regression*. Springer, 2002.
- Trevor Hastie, Andrea Montanari, Saharon Rosset, and Ryan J. Tibshirani. Surprises in high-dimensional ridgeless least squares interpolation. *The Annals of Statistics*, 50(2), 2022.

- Arthur Jacot, Franck Gabriel, and Clement Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in Neural Information Processing Systems*, volume 31, pages 8580–8589. Curran Associates, Inc., 2018.
- Yuling Jiao, Guohao Shen, Yuanyuan Lin, and Jian Huang. Deep nonparametric regression on approximate manifolds: Nonasymptotic error bounds with polynomial prefactors. *The Annals of Statistics*, 51(2), 2023a.
- Yuling Jiao, Yang Wang, and Yunfei Yang. Approximation bounds for norm constrained neural networks with applications to regression and GANs. *Applied and Computational Harmonic Analysis*, 65:249–278, 2023b.
- Jason M. Klusowski and Andrew R. Barron. Approximation by combinations of ReLU and squared ReLU ridge functions with l^1 and l^0 controls. *IEEE Transactions on Information Theory*, 64(12):7649–7656, 2018.
- Michael Kohler and Sophie Langer. On the rate of convergence of fully connected deep neural network regression estimates. *The Annals of Statistics*, 49(4):2231–2249, 2021.
- Vladimir Koltchinskii. Local Rademacher complexities and oracle inequalities in risk minimization. *The Annals of Statistics*, 34(6), 2006.
- Tengyuan Liang and Alexander Rakhlin. Just interpolate: Kernel “Ridgeless” regression can generalize. *The Annals of Statistics*, 48(3), 2020.
- Yuly Makovoz. Random approximants and neural networks. *Journal of Approximation Theory*, 85(1):98–109, 1996.
- Tong Mao and Ding-Xuan Zhou. Approximation of functions from Korobov spaces by deep convolutional neural networks. *Advances in Computational Mathematics*, 48(6):84, 2022.
- Tong Mao and Ding-Xuan Zhou. Rates of approximation by ReLU shallow neural networks. *Journal of Complexity*, 79:101784, 2023.
- Tong Mao, Zhongjie Shi, and Ding-Xuan Zhou. Theory of deep convolutional neural networks III: Approximating radial functions. *Neural Networks*, 144:778–790, 2021.
- Song Mei, Andrea Montanari, and Phan-Minh Nguyen. A mean field view of the landscape of two-layer neural networks. *Proceedings of the National Academy of Sciences*, 115(33), 2018.
- Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Mean-field theory of two-layers neural networks: dimension-free bounds and kernel limit. In *Proceedings of the 32nd Conference on Learning Theory*, volume 99, pages 2388–2464. PMLR, 2019.
- Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Generalization error of random feature and kernel methods: Hypercontractivity and kernel matrix concentration. *Applied and Computational Harmonic Analysis*, 59:3–84, 2022.
- Hrushikesh N. Mhaskar. On the tractability of multivariate integration and approximation by neural networks. *Journal of Complexity*, 20(4):561–590, 2004.

- Hrushikesh N. Mhaskar. Dimension independent bounds for general shallow networks. *Neural Networks*, 123:142–152, 2020.
- Ryumei Nakada and Masaaki Imaizumi. Adaptive approximation and generalization of deep neural network with intrinsic dimensionality. *Journal of Machine Learning Research*, 21(174):1–38, 2020.
- Behnam Neyshabur, Ruslan Salakhutdinov, and Nathan Srebro. Path-SGD: Path-normalized optimization in deep neural networks. In *Advances in Neural Information Processing Systems*, volume 28, pages 2422–2430. Curran Associates, Inc., 2015a.
- Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. In search of the real inductive bias: On the role of implicit regularization in deep learning. In *International Conference on Learning Representations*, 2015b.
- Greg Ongie, Rebecca Willett, Daniel Soudry, and Nathan Srebro. A function space view of bounded norm infinite width ReLU nets: The multivariate case. In *International Conference on Learning Representations*, 2020.
- Samet Oymak and Mahdi Soltanolkotabi. Towards moderate overparameterization: global convergence guarantees for training shallow neural networks. *IEEE Journal on Selected Areas in Information Theory*, 1(1):84–105, 2020.
- Rahul Parhi and Robert D. Nowak. Banach space representer theorems for neural networks and ridge splines. *Journal of Machine Learning Research*, 22(43):1–40, 2021.
- Rahul Parhi and Robert D. Nowak. What kinds of functions do deep neural networks learn? insights from variational spline theory. *SIAM Journal on Mathematics of Data Science*, 4(2):464–489, 2022.
- Rahul Parhi and Robert D. Nowak. Near-minimax optimal estimation with shallow ReLU neural networks. *IEEE Transactions on Information Theory*, 69(2):1125–1140, 2023.
- Alexander Rakhlin and Xiyu Zhai. Consistency of interpolation with Laplace kernels is a high-dimensional phenomenon. In *Proceedings of the 32nd Conference on Learning Theory*, volume 99, pages 2595–2623. PMLR, 2019.
- Pedro Savarese, Itay Evron, Daniel Soudry, and Nathan Srebro. How do infinite width bounded norm networks look in function space? In *Proceedings of the 32nd Conference on Learning Theory*, volume 99, pages 2667–2690. PMLR, 2019.
- Johannes Schmidt-Hieber. Nonparametric regression using deep neural networks with ReLU activation function. *The Annals of Statistics*, 48(4):1875–1897, 2020.
- Jonathan W. Siegel. Optimal approximation of zonoids and uniform approximation by shallow neural networks. *arXiv: 2307.15285*, 2023.
- Jonathan W. Siegel and Jinchao Xu. Approximation rates for neural networks with general activation functions. *Neural Networks*, 128:313–321, 2020.

- Jonathan W. Siegel and Jinchao Xu. Sharp bounds on the approximation rates, metric entropy, and n -widths of shallow neural networks. *Foundations of Computational Mathematics*, pages 1–57, 2022.
- Jonathan W. Siegel and Jinchao Xu. Characterization of the variation spaces corresponding to shallow neural networks. *Constructive Approximation*, 57(3):1109–1132, 2023.
- Justin Sirignano and Konstantinos Spiliopoulos. Mean field analysis of neural networks: A law of large numbers. *SIAM Journal on Applied Mathematics*, 80(2):725–752, 2020.
- Nathan Srebro, Karthik Sridharan, and Ambuj Tewari. Smoothness, low noise and fast rates. In *Advances in Neural Information Processing Systems*, volume 23, pages 2199–2207. Curran Associates, Inc., 2010.
- Ingo Steinwart and Andreas Christmann. *Support Vector Machines*. Springer Science & Business Media, 2008.
- Charles J. Stone. Optimal global rates of convergence for nonparametric regression. *The Annals of Statistics*, 10(4):1040–1053, 1982.
- Taiji Suzuki. Adaptivity of deep ReLU network for learning in Besov and mixed smooth Besov spaces: optimal rate and curse of dimensionality. In *International Conference on Learning Representations*, 2019.
- Vladimir N. Vapnik and Alexey Ya. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and Its Applications*, 16(2):264–280, 1971.
- Martin J. Wainwright. *High-dimensional Statistics: A Non-asymptotic Viewpoint*, volume 48. Cambridge University Press, 2019.
- Qiang Wu, Yiming Ying, and Ding-Xuan Zhou. Learning rates of least-square regularized regression. *Foundations of Computational Mathematics*, 6(2):171–192, 2005.
- Yunfei Yang and Ding-Xuan Zhou. Optimal rates of approximation by shallow ReLU^k neural networks and applications to nonparametric regression. *Constructive Approximation*, 2024.
- Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations*, 2017.
- Kaiqi Zhang and Yu-Xiang Wang. Deep learning meets nonparametric regression: Are weight-decayed DNNs locally adaptive? In *International Conference on Learning Representations*, 2023.