

Advantage of Deep Neural Networks for Estimating Functions with Singularity on Hypersurfaces

Masaaki Imaizumi

*The University of Tokyo
Meguro, Tokyo, Japan*

IMAIZUMI@G.ECC.U-TOKYO.AC.JP

Kenji Fukumizu

*The Institute of Statistical Mathematics
Tachikawa, Tokyo, Japan*

FUKUMIZU@ISM.AC.JP

Editor: Samory Kpotufe

Abstract

We develop a minimax rate analysis to describe the reason that deep neural networks (DNNs) perform better than other standard methods. For nonparametric regression problems, it is well known that many standard methods attain the minimax optimal rate of estimation errors for smooth functions, and thus, it is not straightforward to identify the theoretical advantages of DNNs. This study tries to fill this gap by considering the estimation for a class of non-smooth functions that have singularities on hypersurfaces. Our findings are as follows: (i) We derive the generalization error of a DNN estimator and prove that its convergence rate is almost optimal. (ii) We elucidate a phase diagram of estimation problems, which describes the situations where the DNNs outperform a general class of estimators, including kernel methods, Gaussian process methods, and others. We additionally show that DNNs outperform harmonic analysis based estimators. This advantage of DNNs comes from the fact that a shape of singularity can be successfully handled by their multi-layered structure.

Keywords: Nonparametric Regression, Minimax Optimal Rate, Singularity of Function

1. Introduction

Learning with deep neural networks (DNNs) has been applied extensively owing to their remarkable performance in many tasks. It has been observed that DNNs empirically achieve substantially higher accuracy in some tasks than many existing approaches (Schmidhuber, 2015; LeCun et al., 2015; Hinton et al., 2006; Le et al., 2011; Kingma and Ba, 2014). To understand such empirical successes, we investigate DNNs through nonparametric regression problems.

Suppose that we have n independently and identically distributed (i.i.d.) pairs $(X_i, Y_i) \in [0, 1]^D \times \mathbb{R}$ for $i = 1, \dots, n$ generated from the model

$$Y_i = f^*(X_i) + \xi_i, \quad i = 1, \dots, n, \quad (1)$$

where $f^* : [0, 1]^D \rightarrow \mathbb{R}$ is an unknown function and $(\xi_i)_{i=1}^n$ is an i.i.d. noise that is independent of $(X_i)_{i=1}^n$. For simplicity, we assume ξ_i is Gaussian. The aim of this study is to investigate the generalization error of the maximum likelihood estimator \hat{f}^{DL} given by DNNs; that is, we analyze $\|\hat{f}^{DL} - f^*\|_{L^2(P_X)}^2 := \mathbb{E}_{X \sim P_X}[(\hat{f}^{DL}(X) - f^*(X))^2]$.

This paper argues that deep learning has an advantage over other standard models in terms of the generalization error when f^* has singularities on a hypersurface in the domain. It will be also shown that the the existence of the advantage exhibits a phase transition in the phase space consisting of the *shape of the singularities*, which is defined by the smoothness of the hypersurfaces and the smoothness of the target function f^* . These results are based on the analysis of the minimax optimal rate of the generalization error.

Understanding the advantages of DNNs remains challenging when analyzing deep learning with the nonparametric regression problem. This is due to the well-known fact that some popular existing methods achieve the minimax optimal rate with a *smoothness assumption* for the regression model (1). Namely, when f^* is β -times (continuously) differentiable, the standard knowledge in the nonparametric statistics (Tsybakov, 2009; Wasserman, 2006) tells that a variety of existing methods, such as the kernel method, the Fourier series method, and Gaussian process methods, provide an estimator \tilde{f} , which satisfies

$$\|\tilde{f} - f^*\|_{L^2}^2 = O_{\mathbb{P}}\left(n^{-2\beta/(2\beta+D)}\right).$$

Since this convergence rate is known to be optimal in the minimax sense (Stone, 1982), it is almost impossible to present theoretical evidence for the empirical advantage of DNNs if the smoothness assumption is satisfied.

To overcome this limitation of theoretical understanding, we consider the estimation of a *piecewise smooth function*, which is a natural class of non-smooth functions. More specifically, the functions are singular (non-differentiable or discontinuous) only on smooth hypersurfaces in their multi-dimensional domain. Figure 1 presents an example of such a function, the domain of which is divided into three pieces with piecewise smooth boundaries. This class is flexible enough to express functions of singularities, while being broader than the usual spaces of smooth functions. As being suitable for representing edges in images, similar function classes have been studied in the areas of image analysis and harmonic analysis (Korostelev and Tsybakov, 2012; Candès and Donoho, 2004; Candès and Donoho, 2002; Kutyniok and Labate, 2012).

This study makes the following two contributions. The first one is to prove that the estimator by DNNs almost achieves the minimax optimal rate when estimating functions of singularities. Specifically, let $\mathcal{F}_{\alpha,\beta,M}^{PS}$ be the set of piecewise smooth functions such that their domain is divided into M pieces, they are β -times differentiable except on the boundaries of the pieces, and the boundary is piecewise α -times differentiable (Its rigorous definition will be provided in Section 3). We prove that the least-square estimator \hat{f}^{DL} by DNNs for $f^* \in \mathcal{F}_{\alpha,\beta,M}^{PS}$ satisfies

$$\mathbb{E}\left[\|\hat{f}^{DL} - f^*\|_{L^2(P_X)}^2\right] = \tilde{O}\left(\max\left\{n^{-2\beta/(2\beta+D)}, n^{-\alpha/(\alpha+D-1)}\right\}\right), \quad (2)$$

as $n \rightarrow \infty$ (Corollary 8). Here, \tilde{O} denotes the Big O notation ignoring logarithmic factors. It is interesting to note that this result holds even if a DNN does not contain any non-smooth elements, such as non-differentiable activation functions. That is, even smooth DNNs can estimate such a non-smooth function without being affected by singularities. We also evaluate the effect of the number of pieces in the domain.

As the second contribution, we develop a phase diagram of parameters that explains the superiority of DNNs over a certain class of nonparametric methods. Specifically, we consider




Figure 1: [Left] Example of piecewise smooth function $f(x_1, x_2)$ with two-dimensional input. The function is smooth within the pieces and singular on the boundaries of the pieces. [Right] Domain of example function, divided into three pieces. The bold red curve, which is a boundary of the pieces, is a hypersurface for the singularity of f .




Figure 2: The phase diagram of the parameters α (differentiability of the hypersurface for singularity) and β (differentiability of the function inside the pieces). The red dashed line shows $\alpha = 2\beta(D-1)/(D-2\beta)$ with $D = 5$. If a parameter configuration (α, β) is above the line, the minimax rate of DNNs is better than that of the linear estimator. Otherwise, the DNN and the linear estimator have the same minimax rate.

a class of existing nonparametric estimators known as linear estimators \hat{f}^{lin} , which includes commonly used methods such as estimators by kernel ridge regression, spline regression, and Gaussian process regression, and then derive a lower bound of its error rate as $n^{-\alpha/(2\alpha+D-1)}$ (Proposition 15). Combined with the result that the estimator by DNNs \hat{f}^{DL} has the

rate in (2), we can state that \hat{f}^{DL} has a faster convergence when $-\alpha/(2\alpha + D - 1) > \max\{-2\beta/(2\beta + D), -\alpha/(\alpha + D - 1)\}$. As a result, when $\alpha < 2\beta(D - 1)/(D - 2\beta)$ holds, the estimator by DNNs is superior to linear estimators. Figure 2 illustrates the configuration of the parameters of smoothness which retains the theoretical advantage of DNNs, that is, it describes a set of parameters such that

$$\inf_{\hat{f}^{\text{lin}}} \sup_{f^* \in \mathcal{F}_{\alpha, \beta, M}^{PS}} \mathbb{E}_{f^*} \left[\|\hat{f}^{\text{lin}} - f^*\|_{L^2} \right] \gtrsim \inf_{\hat{f}^{\text{DL}}} \sup_{f^* \in \mathcal{F}_{\alpha, \beta, M}^{PS}} \mathbb{E}_{f^*} \left[\|\hat{f}^{\text{DL}} - f^*\|_{L^2} \right]$$

holds. Here, we denote $a_n \gtrsim b_n$ by $|a_n/b_n| \rightarrow \infty$ holds as $n \rightarrow \infty$ with the sequences $\{a_n\}_n$ and $\{b_n\}_n$. The phase transition shows that the superiority of DNNs becomes apparent when f^* has a hypersurface of singularities that has a less smooth shape, i.e., the boundaries of the pieces are complicated. Otherwise, DNNs and the linear estimator achieve the same minimax rate. Additionally, we also consider other candidates used under the existence of singularities in the field of image analysis, such as the wavelet \hat{f}^{wav} and curvelet \hat{f}^{curve} methods, and then derive another phase diagram about the advantage of DNNs. These results indicate that DNNs certainly offer a theoretical advantage over the other methods under functions of singularities.

As intuitive reasons for these results, we discuss the following two roles of DNNs. First, a model by DNNs, which is a composition of several transforms, is suitable for decomposing non-smooth functions into simple elements. Let us begin with considering a simple example of the indicator function $\mathbf{1}_{S^{D-1}} : \mathbb{R}^D \rightarrow \mathbb{R}$ of the unit sphere $S^{D-1} \subset \mathbb{R}^D$; that is, $\mathbf{1}_{S^{D-1}}(x) = 1$ when $x \in S^{D-1}$, and $\mathbf{1}_{S^{D-1}}(x) = 0$ otherwise. Although $\mathbf{1}_{S^{D-1}}(x)$ is discontinuous, DNNs can approximate $\mathbf{1}_{S^{D-1}}(x)$ without a loss of efficiency from the discontinuity. Note that the function has the form $\mathbf{1}_{S^{D-1}}(x) = \mathbf{1}_{\{\cdot \geq 0\}} \circ h(x)$ with a certain smooth function $h(x)$ and a step function $\mathbf{1}_{\{\cdot \geq 0\}}$, such as $\mathbf{1}_{\{\cdot \geq 0\}}(x) = 1$ for $x \geq 0$ and $\mathbf{1}_{\{\cdot \geq 0\}}(x) = 0$ otherwise. The multi-transform structure of DNNs plays an important role: a first transform approximates $h(x)$, and a second one $\mathbf{1}_{\{\cdot \geq 0\}}$, following which the entire DNN model consists of a composite function of the two transforms. Owing to the composition, DNNs can approximate $\mathbf{1}_{S^{D-1}}$ as if it were a smooth function. Second, the activation function of DNNs plays a significant role. Several common activation functions, such as the sigmoid and rectified linear unit (ReLU) activations, can easily approximate a step function with an arbitrarily small error. The general conditions for achieving the approximation are explained in Assumption 1 in Section 2.

This paper is an extension of a conference proceeding Imaizumi and Fukumizu (2019). The contributions in the current paper differ from the results of the conference paper in the following ways. First, this study handles the general class of activation functions of DNNs, while the proceeding investigates the Rectified Linear Unit (ReLU) activation function. Hence, the proof of the first contribution with singularities is significantly different. Especially, this study develops a new proof for resolving singularities by possibly smooth activation functions, while the proof of the proceeding for singularity relies on the non-smoothness of the ReLU activation. Second, this paper develops the phase diagram by newly derived lower bounds on the minimax rate for the other estimators. Our new derivation of a specific lower bound on the rate allows us to describe such a phase transition, while Imaizumi and Fukumizu (2019) does not derive the lower rate. Further, this study additionally

examines several methods that are adept at handling singularities, such as the wavelet and curvelet approaches, and subsequently demonstrates that they do not achieve optimality.

It is important to mention that the superiority of DNNs shown in our results does not depend on the property of learning algorithms, but only on an expressive power and a degree of freedom of DNNs. Toward a very practical aspect of DNN, it is certainly important to investigate the error from learning algorithms. However, the environment surrounding practical DNNs is highly complicated that it is not very meaningful to consider all of these factors at the same time. For a more rigorous analysis and correct understanding, we follow the line of research in nonparametric regression, such as Schmidt-Hieber (2020) and Bauer and Kohler (2019), we focus on the advantage of DNNs on the expressive power and the degree of freedom of DNNs with singularities of data.

1.1 Related Studies

Several pioneering works have investigated deep learning in terms of the nonparametric regression problems. A recent study Schmidt-Hieber (2020) considered the case in which f^* is expressed as a composition of several smooth functions, and then derived the minimax optimal rate with this setting. Bauer and Kohler (2019) also derived the convergence rate of errors when f^* had the form of a generalized hierarchical interaction model, and revealed that the obtained rate was dependent on the lower dimensionality of the model. These studies focused on the composition structure of f^* , and they did not consider the non-smoothness or discontinuity of f^* . These works thus did not take into account singularities, which is the main focus of our study. We also mention that the comparison with the Harmonic and linear estimators is our unique result.

Suzuki (2019) and Hayakawa and Suzuki (2020) also demonstrated the superiority of DNNs. These works investigated the generalization error of DNNs when f^* belongs to the Besov space. Interestingly, the convergence rate of DNNs was faster than that of the linear estimator when the norm parameter of the Besov space was less than 1, following the theory of Donoho and Johnstone (1998). Although their motivation is similar to ours, the Besov space is not suitable for representing functions of singularities on a smooth hypersurface. This is because the wavelet decomposition, which is used to define the Besov space, loses its efficiency for handling the hypersurfaces for singularity, as explained in Section 5. Note also that the comparison in Suzuki (2019) and Hayakawa and Suzuki (2020) between DNNs and linear estimators was motivated by the proceeding version of this paper.

The current work has been technically inspired by Petersen and Voigtlaender (2018), which investigated the approximation power of DNNs with discontinuity. The main difference between the above paper and current work is the focus on the advantage of deep learning. Their study mainly investigated the approximation error of DNNs. Thus, a comparison with existing methods was not the main purpose. Another major difference is that their study focused on the approximation power, whereas we investigate the generalization error, including the variance control of DNNs. A further difference is that our study investigates a broader class of discontinuous functions, as our definition directly controls hypersurfaces in the domain, whereas Petersen and Voigtlaender (2018) defined the discontinuity by a transform of the Heaviside function.

1.2 Paper Organization

The remainder of this paper is organized as follows. Section 2 introduces a functional model by DNNs, following which an estimator for the regression problem with DNNs is defined. The notion of functions with singularities is explained in Section 3. Section 4 derives the convergence rate of the estimator by means of DNNs. Furthermore, the minimax optimality of the convergence rate is derived. Section 5 presents the non-optimal convergence rate obtained by a certain class of other estimators, and compares this rate with that of DNNs. Section 6 summarizes our work. Full proofs are deferred to the supplementary material.

1.3 Notation

Let $I := [0, 1]$ be the unit interval, \mathbb{N} be natural numbers, and $\mathbb{N}_0 := \mathbb{N} \cup \{0\}$. For $z \in \mathbb{N}$, $[z] := \{1, 2, \dots, z\}$ is the set of natural numbers that are no more than z . The d -th element of vector $b \in \mathbb{R}^D$ is denoted by b_d , and $b_{-d} = (b_1, \dots, b_{d-1}, b_{d+1}, \dots, b_D)$ for $d \in [D]$. $\|b\|_q := (\sum_d b_d^q)^{1/q}$ is the q -norm for $q \in (0, \infty)$, $\|b\|_\infty := \max_{j \in [J]} |b_j|$, and $\|b\|_0 := \sum_{j \in [J]} \mathbf{1}_{\{b_j \neq 0\}}$. For a measure space (A, \mathcal{B}, μ) and a measurable function $f : A \rightarrow \mathbb{R}$, let $\|f\|_{L^2(\mu)} := (\int_A |f(x)|^2 d\mu(x))^{1/2}$ denote the $L^2(\mu)$ -norm if the integral is finite. When μ is the Lebesgue measure on a measurable set A in \mathbb{R}^D , we omit μ and simply write $\|f\|_{L^2(A)}$. For a set $A \subset \mathbb{R}^D$, $\text{vol}(A)$ denotes the Lebesgue measure of A . The tensor product is denoted by \otimes . For a set $R \subset I^D$, let $\mathbf{1}_R : I^D \rightarrow \{0, 1\}$ denote the indicator function of R ; that is, $\mathbf{1}_R(x) = 1$ if $x \in R$, and $\mathbf{1}_R(x) = 0$ otherwise. For the sequences $\{a_n\}_n$ and $\{b_n\}_n$, $a_n \lesssim b_n$ means that there exists $C > 0$ such that $a_n \leq Cb_n$ holds for every $n \in \mathbb{N}$. $a_n \gtrsim b_n$ denotes the opposite of $a_n \lesssim b_n$. Furthermore, $a_n \asymp b_n$ denotes both $a_n \gtrsim b_n$ and $a_n \lesssim b_n$. $a_n \lesssim\lesssim b_n$ denotes $|a_n/b_n| \rightarrow 0$ as $n \rightarrow \infty$, and $a_n \gtrsim\gtrsim b_n$ is its opposite. For a set of parameters θ , $C_\theta > 0$ denotes an existing finite constant depending on θ . Let O_P and o_P be the Landau big O and small o in probability. \tilde{O} ignores every multiplicative polynomial of logarithmic factors.

2. Deep Neural Networks

A deep neural network (DNN) is a model of functions defined by a layered structure. Let $L \in \mathbb{N}$ be the number of layers in DNNs, and for $\ell \in [L + 1]$, let $D_\ell \in \mathbb{N}$ be the dimensionality of variables in the ℓ -th layer. DNNs have a matrix parameter $A_\ell \in \mathbb{R}^{D_{\ell+1} \times D_\ell}$ and a vector parameter $b_\ell \in \mathbb{R}^{D_\ell}$ for $\ell \in [L]$ to represent weights and biases, respectively. We introduce an activation function $\eta : \mathbb{R} \rightarrow \mathbb{R}$, which will be specified later. For a vector input $z \in \mathbb{R}^d$, $\eta(x) = (\eta(z_1), \dots, \eta(z_d))^\top$ denotes an element-wise operation. For $\ell \in [L - 1]$, with an input vector $z \in \mathbb{R}^{D_\ell}$, we define $g_\ell : \mathbb{R}^{D_\ell} \rightarrow \mathbb{R}^{D_{\ell+1}}$ as $g_\ell(z) = \eta(A_\ell z + b_\ell)$. We also define $g_L(z) = A_L z + b_L$ with $z \in \mathbb{R}^{D_L}$. Thereafter, we define a function $g : \mathbb{R}^{D_1} \rightarrow \mathbb{R}^{D_{L+1}}$ of DNNs with $(A_1, b_1), \dots, (A_L, b_L)$ by

$$g(x) = g_L \circ g_{L-1} \circ \dots \circ g_1(x). \quad (3)$$

Intuitively, $g(x)$ is constituted by compositions of L maps.

For each g with the form (3), we introduce several operators to extract information of g . Let $L(g) = L$ be the number of layers, $S(g) := \sum_{\ell \in [L]} \|\text{vec}(A_\ell)\|_0 + \|b_\ell\|_0$ as a number of non-zero elements in the parameter matrix and tensor in g , and $B(g) := \max_{\ell \in [L]} \|\text{vec}(A_\ell)\|_\infty \vee$

$\max_{\ell \in [L]} \|b_\ell\|_\infty$ be the largest absolute value of the parameters. Here, $\text{vec}(\cdot)$ is a vectorization operator for matrices.

We define the set of functions of DNNs. With a tuple $(L', S', B') \in \mathbb{N}^3$, we write it as

$$\mathcal{G}(L', S', B') := \left\{ g \in L^\infty(I^D) \mid g \text{ as (3)} : L(g) \leq L', S(g) \leq S', B(g) \leq B', \|g\|_{L^\infty(I^D)} \leq F \right\},$$

where $F > 0$ is a threshold. Since the form of DNNs is flexible, we control the size and complexity of it through the layers and parameters through the tuple (L', S', B') . Here, the internal dimensionality D_ℓ is implicitly regularized by the tuple.




Figure 3: Common activation functions; Sigmoid: $\eta(x) = 1/(1 + \exp(x))$, ReLU: $\eta(x) = \max\{x, 0\}$, LeakyReLU: $\eta(x) = \max\{x, 0\} + 0.2 \min\{x, 0\}$, SoftPlus: $\eta(x) = \log(1 + \exp(x))$, and Swish: $\eta(x) = x/(1 + \exp(x))$.

The explicit form of η plays a critical role in DNNs, and numerous variations of activation functions have been suggested. We select several of the most representative activation functions in Figure 3. To investigate a wide class of activation functions, we introduce the following assumption.

Assumption 1. *An activation function η satisfies either of the following conditions:*

- (i) *For $N \in \mathbb{N}$, there exists $q \geq 1$ and $k \in \{0, 1\}$ such that $\partial^j \eta$ exists at every point and is bounded for $j = 1, \dots, N + 1$. Furthermore, there exists $x' \in \mathbb{R}$ such that $\min_{j=1, \dots, N} |\partial^j \eta(x')| \geq c_\eta > 0$ with a constant $c_\eta > 0$, and the followings hold:*

$$|\eta(x) - \bar{c}x^k| = O(1/x^q), \quad (x \rightarrow \infty), \quad \text{and} \quad |\eta(x) - \underline{c}| = O(1/|x|^q), \quad (x \rightarrow -\infty),$$

with some constants $\bar{c} > \underline{c} \geq 0$. There also exists $C_K > 0$ such as $|\eta(x)| \leq C_K(1 + |x|^k)$ for any $x \in \mathbb{R}$.

- (ii) *There exist constants $c_1 > c_2 \geq 0$ such that*

$$\eta(x) = \begin{cases} c_1 x & \text{if } x \geq 0 \\ c_2 x & \text{if } x < 0. \end{cases}$$

The condition (i) describes smooth activation functions such as the sigmoid function, the softplus function, and the Swish function with $N = \infty$. The lower bound on $|\partial^j \eta(x')|$ describes a non-vanishing property of the derivatives. The condition (ii) indicates a piecewise linear function such as the rectified linear unit (ReLU) function and the leaky ReLU function, which require another technique to investigate.

2.1 Regression Problem and Estimator by DNNs

We consider the least square estimator by DNNs for the regression problem (1). Let the D -dimensional cube I^D ($D \geq 2$) be a space for input variables X_i . Suppose we have a set of observations $(X_i, Y_i) \in I^D \times \mathbb{R}$ for $i \in [n]$ which is independently and identically distributed with the data generating process (1) where $f^* : I^D \rightarrow \mathbb{R}$ is an unknown true function and ξ_i is Gaussian noise with mean 0 and variance $\sigma^2 > 0$ for $i \in [n]$. We also suppose that X_i follows a marginal distribution P_X on I^D and it has a density function p_X which is bounded away from zero and infinity. Then, we define an estimator by empirical risk minimization with DNNs as

$$\hat{f}^{DL} \in \operatorname{argmin}_{f \in \mathcal{G}(L, S, B)} \frac{1}{n} \sum_{i=1}^n (Y_i - f(X_i))^2. \quad (4)$$

The minimizer always exists since $\mathcal{G}(L, S, B)$ is a compact set in $L^\infty(I^D)$ due to the parameter bound and continuity of η . Note that we do not discuss the optimization issues from the non-convexity of the loss function, since we mainly focus on an estimation aspect.

3. Characterization of Functions with Singularity

In this section, we provide a rigorous formulation of functions with singularity on smooth hypersurfaces. To describe the singularity of functions, we introduce a notion of *piecewise smooth functions*, which have its domain divided into several pieces and smooth only within each of the pieces. Furthermore, piecewise smooth functions are singular (non-differentiable or discontinuous) on the boundaries of the pieces.

Smooth Functions (Hölder space): Let Ω be a closed subset of \mathbb{R}^D and $\beta, F > 0$ be parameters. For a multi-index $a \in \mathbb{N}_0^D$, $\partial^a = \partial^{a_1} \cdots \partial^{a_D}$ denotes a partial derivative operator. The Hölder space $H^\beta(\Omega)$ is defined as a set of functions $f : \Omega \rightarrow \mathbb{R}$ such as

$$H^\beta(\Omega) := \left\{ f \left| \max_{\|a\|_1 \leq \lfloor \beta \rfloor} \|\partial^a f\|_{L^\infty(\Omega)} + \max_{\|a\|_1 = \lfloor \beta \rfloor} \sup_{x, x' \in \Omega, x \neq x'} \frac{|\partial^a f(x) - \partial^a f(x')|}{|x - x'|^{\beta - \lfloor \beta \rfloor}} < \infty \right\}.$$

Also, let $H_F^\beta(\Omega)$ be a ball in $H^\beta(\Omega)$ with its radius $F > 0$ in terms of the norm $\|\cdot\|_{L^\infty}$.

Pieces in the Domain: We describe pieces as subsets of the domain I^D by dividing I^D with several hypersurfaces. For $j = 1, \dots, J$, let $h_j \in H_F^\alpha(I^{D-1})$ be a function with input $x_{-d_j} \in I^{D-1}$ for some $d_j \in [D]$. We define a family of M pieces which are an intersection of one side of J hypersurfaces. Let $I_j^+ := \{x \in I^D \mid x_{d_j} \geq h_j(x_{-d_j})\}$ and $I_j^- := \{x \in I^D \mid x_{d_j} \leq h_j(x_{-d_j})\}$. For a J -tuple $t = (t_j)_{j=1}^J \in \{+, -\}^{[J]}$, a unit piece of I^D is defined by $I_t := \bigcap_{j \in [J]} I_j^{t_j}$. Let \mathcal{T} be a subset of $\{+, -\}^{[J]}$, and define *piece* $R_{\mathcal{T}}$ by

$$R_{\mathcal{T}} := \bigcup_{t \in \mathcal{T}} I_t.$$

Let $\{\mathcal{T}_1, \dots, \mathcal{T}_M\}$ be a partition of $\{+, -\}^{[J]}$. Then, it is easy to see $\bigcup_{m \in [M]} R_{\mathcal{T}_m} = I^D$ and $R_{\mathcal{T}_m} \cap R_{\mathcal{T}_{m'}}$ is of Lebesgue measure zero. The family of pieces that is of size M is given by

$$\mathcal{R}_{\alpha, M} := \left\{ \{R_{\mathcal{T}_m}\}_{m \in [M]} \mid \{\mathcal{T}_m\}_{m \in [M]} \text{ is a partition of } \{+, -\}^{[J]} \right\}.$$

Intuitively, $\{R_{\mathcal{T}_m}\}_{m \in [M]} \in \mathcal{R}_{\alpha, M}$ is a partition of I^D , allowing overlap on the piecewise α -smooth boundaries. Figure 1 presents an example. In the following, when the partition $\{\mathcal{T}_m\}_{m \in [M]}$ is fixed, we can also write $R_m := R_{\mathcal{T}_m}$ by slightly changing the notation.

Piecewise Smooth Functions: By using $H_F^\beta(I^D)$ and $\mathcal{R}_{\alpha, M}$, we introduce a space of piecewise smooth functions as

$$\mathcal{F}_{\alpha, \beta, M}^{PS} := \left\{ \sum_{m \in [M]} f_m \otimes \mathbf{1}_{R_m} : f_m \in H_F^\beta(I^D), \{R_m\}_{m \in [M]} \in \mathcal{R}_{\alpha, M} \right\}.$$

Since $f_m(x)$ realizes only when $x \in R_m$, the notion of $\mathcal{F}_{\alpha, \beta, M}^{PS}$ can express a combination of smooth functions on each piece R_m . Hence, functions in $\mathcal{F}_{\alpha, \beta, M}^{PS}$ are non-smooth (and even discontinuous) on boundaries of R_m . Obviously, $H^\beta(I^D) \subset \mathcal{F}_{\alpha, \beta, M}^{PS}$ holds for any M and α , since $f_1 = f_2 = \dots = f_M$ makes the function globally smooth.

Remark 1 (Similar definitions). *Several studies (Petersen and Voigtlaender, 2018; Imaizumi and Fukumizu, 2019) also define a class of piecewise smooth functions. There are mainly two differences of our definition in this study. First, our definition can describe a wider class of pieces. The definition utilizes a direct definition of a smooth hypersurface function h , while the other definition defines pieces by a transformation of the Heaviside function which is slightly restrictive. Second, our definition is less redundant. We do not allow the pieces to overlap with one another, whereas some of the other definitions allow pieces to overlap. When overlap exists, it may make approximation and estimation errors worse, which is a problem that our definition can avoid.*

Remark 2 (Comparison with a boundary by Mammen and Tsybakov (1995)). *We compare our formulation with $h_j \in H_\alpha^\beta(I^{D-1})$ with another way of making boundaries obtained by continuous transformation of a sphere by Mammen and Tsybakov (1995). In our approach, each h_j can only represent a singularity in one fixed dimensional direction. In contrast, the sphere-based boundary has the flexibility to create boundaries in various dimensional directions at once. However, since our formulation can provide several I_j^+ and I_j^- in different dimensional directions, we can reproduce the sphere-based boundary if we combine multiple I_j^+ and I_j^- by taking their intersections as the definition of I_t . We adopt the current formulation because this way of decomposing the sphere-based singularity into its parts is more general. Even if we adopt the sphere-based boundary, we can obtain the same result by using the approximation on step functions in Lemma 4.*

4. Generalization Error of Deep Neural Networks

We provide theoretical results regarding DNN performances for estimating piecewise smooth functions. To begin with, we decompose the estimator error into an approximation error and a complexity error, analogously to the bias-variance decomposition. By a simple calculation on (4), we obtain the following inequality:

$$\|\hat{f}^{DL} - f^*\|_n^2 \leq \underbrace{\inf_{f' \in \mathcal{G}(L, S, B)} \|f^* - f'\|_n^2}_{\mathcal{B}} + \underbrace{\frac{2}{n} \sum_{i=1}^n \xi_i(\hat{f}^{DL}(X_i) - f(X_i))}_{\mathcal{V}}, \quad (5)$$

with some $f \in \mathcal{G}(L, S, B)$ which will be specified later. Here, $\|f\|_n^2 := n^{-1} \sum_{i \in [n]} f(X_i)^2$ is an empirical (pseudo) norm. We note that ξ_i is the Gaussian noise displayed in (1). The term \mathcal{B} in the right hand side is the approximation error, and the term \mathcal{V} is the complexity error. In the following section, we bound \mathcal{B} and subsequently combine it with the bound for \mathcal{V} .

4.1 Approximation Result

We evaluate the approximation error with piecewise smooth functions according to the following three preparatory steps: approximating (i) smooth functions, (ii) step functions, and (iii) indicator functions on the pieces. Thereafter, we provide a theorem for the approximation of piecewise smooth functions.

As the first step, we state the approximation power of DNNs for smooth functions in the Hölder space. Although this topic has been studied extensively (Mhaskar, 1996; Yarotsky, 2017), we provide a formal statement because a condition on activation functions is slightly different.

Lemma 3 (Smooth function approximation). *Let $\beta > 0$ be a constant. Suppose Assumption 1 holds with $N > \beta$. Then, there exist constants $C_{\beta,D,F}, C_{\beta,D,F,q} > 0$ such that a tuple (L, S, B) such as $L \geq C_{\beta,D,F}(|\beta| + \log_2(1/\varepsilon) + 1)$, $S \geq C_{\beta,D,F}\varepsilon^{-D/\beta}(\log_2(1/\varepsilon))^2$, and $B \geq C_{\beta,D,F,q}\varepsilon^{-C_\beta}$, which satisfies*

$$\inf_{g \in \mathcal{G}(L,S,B)} \sup_{f \in H_F^\beta(I^D)} \|g - f\|_{L^2(R)} \leq \text{vol}(R)\varepsilon,$$

for any non-empty measurable set $R \subset I^D$ and $\varepsilon > 0$.

As we reform the result, the approximation error is written as $O(\text{vol}(R)S^{-\beta/D})$ up to logarithmic factors, with L and B satisfying the conditions.

As the second step, we investigate an approximation for a step function $\mathbf{1}_{\{\cdot \geq 0\}}$, which will play an important role in handling singularities of functions.

Lemma 4 (Step function approximation). *Suppose η satisfies Assumption 1. Then, for any $\varepsilon \in (0, 1)$ and $T > 0$, we obtain*

$$\inf_{g \in \mathcal{G}(2,6,C_{T,q}\varepsilon^{-8})} \|g - \mathbf{1}_{\{\cdot \geq 0\}}\|_{L^2([-T,T])} \leq \varepsilon.$$

The result states that any activation functions satisfying Assumption 1 can approximate indicator functions. Importantly, DNNs can achieve an arbitrary error ε with $O(1)$ parameters. The approximation with this constant number of parameters is very important in obtaining the desired rate, since the generalization error of DNNs is significantly influenced by the number of parameters.

As the third step, we investigate the approximation error for the indicator function of a piece. For approximation by DNNs, we reform the indicator function into a composition of a step function and a smooth function:

$$\mathbf{1}_{\{x_d \leq h_j(x_{-d})\}} = \mathbf{1}_{\{\cdot \geq 0\}} \circ (x \mapsto \mp(x_d - h_j(x_{-d}))), \quad x \in I^D.$$

Using this formulation, we obtain the following result:

Lemma 5 (Indicator functions for $\mathcal{R}_{\alpha,J}$). *Suppose that Assumption 1 holds with $N > \alpha$. Then, there exist constants $C_{\alpha,D,F,J}, C_{F,J,q} > 0$ such that for any $\{R_m\}_{m \in [M]} \in \mathcal{R}_{\alpha,M}$ and $\varepsilon > 0$ we can find a function $f = (f_1, \dots, f_M)^\top \in \mathcal{G}(L, S, B)$ such as $L \geq C_{\alpha,D,F,J}([\alpha] + \log_2(1/\varepsilon))$, $S \geq C_{\alpha,D,F,J}(J\varepsilon^{-2(D-1)/\alpha}(\log_2(1/\varepsilon))^2 + M(\log_2(1/\varepsilon))^2)$, and $B \geq C_{F,J,q}\varepsilon^{-C_\alpha}$, which satisfies*

$$\|\mathbf{1}_{R_m} - f_m\|_{L^2(I)} \leq \varepsilon, \quad \forall m \in [M].$$

Since the model of DNNs has a composition structure, DNNs can implicitly decompose the indicator function into a step function and a smooth function, which derives the convergence rate. Importantly, even though the function $\mathbf{1}_{\{x_d \leq h_j(x_{-d})\}}$ with $h_j \in H_F^\alpha(I^{D-1})$ has singularity on a set $\{x \in I^D \mid x_d = h_j(x_{-d})\}$, DNNs can achieve a fast approximation rate as if the boundary is a smooth function in $H_F^\alpha(I^{D-1})$.

Based on the above steps, we derive an approximation rate of DNNs for a piecewise smooth function $f \in \mathcal{F}_{\alpha,\beta,M}^{PS}$:

Theorem 6 (Approximation Error). *Suppose Assumption 1 holds with $N > \alpha \vee \beta$. Then, there exist constants $C_{\alpha,\beta,D,F}, C_{\alpha,\beta,D,F,J}, C_{F,M,q} > 0$ such that there exists a tuple (L, S, B) such as $L \geq C_{\alpha,\beta,D,F}([\alpha] + [\beta] + \log_2(1/\varepsilon_1) + \log_2(M/\varepsilon_2) + 1)$, $S \geq C_{\alpha,\beta,D,F,J}(M\varepsilon_1^{-D/\beta}(\log_2(1/\varepsilon_1))^2 + (\varepsilon_2/M)^{-2(D-1)/\alpha}(\log_2(M/\varepsilon_2))^2)$, and $B \geq C_{F,M,q}(\varepsilon_1 \wedge \varepsilon_2)^{-16 \wedge -C_{\alpha,\beta}}$, which satisfies*

$$\inf_{f \in \mathcal{G}(L,S,B)} \sup_{f^* \in \mathcal{F}_{\alpha,\beta,M}^{PS}} \|f - f^*\|_{L^2(I^D)} \leq \varepsilon_1 + \varepsilon_2,$$

for any $\varepsilon_1, \varepsilon_2 \in (0, 1)$.

The result states that the approximation error contains two main terms. A simple calculation yields that the error is reformulated as $O(S^{-\beta/D} + S^{-\alpha/2(D-1)})$ up to logarithmic factors, with L and B satisfying the conditions. The first rate $O(S^{-\beta/D})$ describes approximation for $f_m \in H_F^\beta(I^D)$, and the second rate $O(S^{-\alpha/2(D-1)})$ is for $\mathbf{1}_{R_m}$.

4.2 Generalization Result

We evaluate a generalization error of DNNs, based on the decomposition (5), associated with the bound on \mathcal{B} . To evaluate the remained term \mathcal{V} , we utilize the celebrated theory of the local Rademacher complexity (Bartlett, 1998; Koltchinskii, 2006). Then, we obtain one of our main results as follows. $\mathbb{E}_{f^*}[\cdot]$ denotes the expectation with respect to the true distribution of (X, Y) .

Theorem 7 (Generalization Error). *Suppose $f^* \in \mathcal{F}_{\alpha,\beta,M}^{PS}$ and Assumption 1 holds with $N > \alpha \vee \beta$. Then, there exists a sufficiently large F and a tuple (S, B, L) satisfying $L \geq C_{\alpha,\beta,D,F}(1 + [\alpha] + [\beta] + \log_2(n/M))$, $S = C_{\alpha,\beta,D,F,J}(Mn^{D/(2\beta+D)} + n^{(D-1)/(\alpha+D-1)})\log^2 n$, and $B \geq C_{F,M,q}n^{C_{\alpha,\beta,D}}$, such that there exist $C = C_{\sigma,\alpha,\beta,D,F,J,P_X} > 0$ and $c_1 > 0$ which satisfy*

$$\mathbb{E}_{f^*} \left[\|\widehat{f}^{DL} - f^*\|_{L^2(P_X)}^2 \right] \leq CM(n^{-2\beta/(2\beta+D)} + n^{-\alpha/(\alpha+D-1)})\log^2 n + \frac{C_{\sigma,F}\sqrt{\log n}}{n}.$$

The dominant term appears in the first term in the right hand, thus it mainly describes the error bound for \widehat{f}^{DL} . We note that the main term of the squared error increases linearly in the number of pieces M . To simplify the order of the bound, we provide the following:

Corollary 8. *With the settings in Theorem 7, we obtain*

$$\mathbb{E}_{f^*} \left[\|\hat{f}^{DL} - f^*\|_{L^2(P_X)}^2 \right] = \tilde{O} \left(\max\{n^{-2\beta/(2\beta+D)}, n^{-\alpha/(\alpha+D-1)}\} \right).$$

The order is interpreted as follows. The first term $n^{-2\beta/(2\beta+D)}$ describes an effect of estimating $f_m \in H_F^\beta(I^D)$ for $m \in [M]$. The rate corresponds to the minimax optimal convergence rate of generalization errors for estimating smooth functions in $H_F^\beta(I^D)$ (for a summary, see Tsybakov (2009)). The second term $n^{-\alpha/(\alpha+D-1)}$ reveals an effect from estimation of $\mathbf{1}_{R_m}$ for $m \in [M]$ through estimating the boundaries of $R_m \in \mathcal{R}_{\alpha,J}$. The same rate of convergence appears in a problem for estimating sets with smooth boundaries (Mammen and Tsybakov, 1995). Based on the result, we state that DNNs can divide a piecewise smooth function into its various smooth functions and indicators, and estimate them by parts. Thus, the overall convergence rate is the sum of the rates of the parts. We also note that we consider a sufficiently large F increasing in n such as $F = \Theta(\log n)$, whose effect is asymptotically negligible.

Remark 9 (Smoothness of η). *It is worth noting that the rate in Theorem 7 holds regardless of smoothness of the activation function η , because Assumption 1 allows both smooth and non-smooth activation functions. That is, even when \hat{f}^{DL} by DNNs is a smooth function with smooth activation, we can obtain the rate in Corollary 8 with non-smooth f^* .*

We can consider the error from optimization independently from the statistical generalization. The following proposition provides the statement.

Proposition 10 (Effect of Optimization). *If a learning algorithm outputs $\hat{f}^{\text{Algo}} \in \mathcal{G}(L, S, B)$ such as*

$$n^{-1} \sum_{i \in [n]} (Y_i - \hat{f}^{\text{Algo}}(X_i))^2 - \inf_{f \in \mathcal{G}(L, S, B)} n^{-1} \sum_{i \in [n]} (Y_i - f(X_i))^2 \leq \Delta,$$

with an existing constant $\Delta > 0$, then the following holds:

$$\mathbb{E}_{f^*} \left[\|\hat{f}^{\text{Algo}} - f^*\|_{L^2(P_X)}^2 \right] \leq \tilde{O} \left(\max\{n^{-2\beta/(2\beta+D)}, n^{-\alpha/(\alpha+D-1)}\} \right) + \Delta.$$

We can evaluate the generalization error, including the optimization effect, by combining several results on the magnitude of Δ (for example, Kawaguchi (2016) and Allen-Zhu et al. (2019)).

4.3 Minimax Lower Bound of Generalization Error

We investigate the efficiency of the convergence rate in Corollary 8. To this end, we consider the minimax generalization error for a functional class \mathcal{F} such as

$$\mathcal{R}(\mathcal{F}) := \inf_{\bar{f}} \sup_{f^* \in \mathcal{F}} \mathbb{E}_{f^*} \left[\|\bar{f} - f^*\|_{L^2(I^D)}^2 \right],$$

where $\bar{f} = \bar{f}(X_1, \dots, X_n, Y_1, \dots, Y_n)$ is taken from all possible estimators depending on the observations. In this section, we derive a lower bound of the generalization error of DNNs,

and then prove that it corresponds to the rate Theorem 7 up to logarithmic factors. By the result, we can claim that the estimation by DNNs is (almost) optimal in the minimax sense.

We introduce several new notations. For set A equipped with a norm $\|\cdot\|$, let $\mathcal{N}(\varepsilon, A, \|\cdot\|)$ be the covering number of A in terms of $\|\cdot\|$, and $\mathcal{M}(\varepsilon, A, \|\cdot\|)$ be the packing number of A , respectively. For sequences $\{a_n\}_n$ and $\{b_n\}_n$, we write $a_n = \Omega(b_n)$ for $\limsup_{n \rightarrow \infty} |a_n/b_n| > 0$. Also, $a_n = \Theta(b_n)$ means that both of $a_n = O(b_n)$ and $a_n = \Omega(b_n)$ hold.

To derive the lower bound, we apply the following information theoretic result by Yang and Barron (1999):

Theorem 11 (Theorem 6 in Yang and Barron (1999)). *Let \mathcal{F} be a set of functions, and ε_n be a sequence such that $\varepsilon_n^2 = \mathcal{M}(\varepsilon_n, \mathcal{F}, \|\cdot\|_{L^2})/n$ holds. Then, we obtain*

$$\mathcal{R}(\mathcal{F}) = \Theta(\varepsilon_n^2).$$

Since the minimax rate for $\mathcal{F}_{\alpha,\beta,M}^{PS}$ is bounded below by that of its subset, we will find a suitable subset of $\mathcal{F}_{\alpha,\beta,M}^{PS}$ and measure its packing number. In the rest of this section, we will take the following two steps. First, we define a subset of $\mathcal{F}_{\alpha,\beta,M}^{PS}$ by introducing a notion of basic pieces. Second, we measure a packing number of the subset of $\mathcal{F}_{\alpha,\beta,M}^{PS}$.

As the first step, we define a *basic piece indicator*, which is a set of piecewise functions whose piece is an embedding of D -dimensional balls, and then define a certain subset of $\mathcal{F}_{\alpha,\beta,M}^{PS}$. As a preparation, let $\mathbb{S}^{D-1} := \{x \in \mathbb{R}^D : \|x\|_2 = 1\}$ is the $D - 1$ dimensional sphere, and let $(V_j, F_j)_{j=1}^\ell$ be its coordinate system with some ℓ as a C^∞ -differentiable manifold such that $F_j : V_j \rightarrow \mathring{B}^{D-1} := \{x \in \mathbb{R}^{D-1} \mid \|x\| < 1\}$ is a diffeomorphism. Compactness of the domain guarantees that we can find finite ℓ . A function $f : \mathbb{S}^{D-1} \rightarrow \mathbb{R}$ is said to be in the Hölder class $H^\alpha(\mathbb{S}^{D-1})$ with $\alpha > 0$ if $f \circ F_j^{-1}$ is in $H^\alpha(\mathring{B}^{D-1})$.

Basic Piece Indicator: A subset $R \subset I^D$ is called an α -*basic piece*, if it satisfies two conditions: (i) there is a continuous embedding $g : \{x \in \mathbb{R}^D \mid \|x\| \leq 1\} \rightarrow \mathbb{R}^D$ such that its restriction to the boundary \mathbb{S}^{D-1} is in $H^\alpha(\mathbb{S}^{D-1})$ and $R = I^D \cap \text{Image}(g)$, (ii) there exist $d \in [D]$ and $h \in H_F^\alpha(I^{D-1})$ such that the indicator function of R is given by the graph

$$\mathbf{1}_R(x) = \Psi_d(x_1, \dots, x_{d-1}, x_d + h(x_{-d}), x_{d+1}, \dots, x_D), x \in I^D,$$

where $\Psi_d(x) := \mathbf{1}_{\{x_d \geq 0\}}$ is the Heaviside function. Then, we define a *basic piece indicator* as an indicator function $\mathbf{1}_R : I^D \rightarrow \mathbb{R}$ with an α -basic piece R . We also define a set of basic piece indicators as

$$\mathcal{I}_\alpha := \{\mathbf{1}_R \mid R \text{ is an } \alpha\text{-basic piece}\}.$$

The condition (i) tells that a basic piece belongs to the *boundary fragment class* which is developed by Dudley (1974) and Mammen and Tsybakov (1999). The condition (ii) means R is a set defined by a horizon function discussed in Petersen and Voigtlaender (2018).

In the following, we consider a functional class

$$H^\beta(I^D) \otimes \mathcal{I}_\alpha := \left\{ f \otimes \mathbf{1}_R \mid f \in H^\beta(I^D), \mathbf{1}_R \in \mathcal{I}_\alpha \right\},$$

and use this set as a key subset for the minimax lower bound. Obviously, $H^\beta(I^D) \otimes \mathcal{I}_\alpha \subset \mathcal{F}_{\alpha,\beta,M}^{PS}$ holds for any $M \in \mathbb{N}$.

At the second step of this section, we measure a packing number of $H^\beta(I^D) \otimes \mathcal{I}_\alpha$.

Proposition 12 (Packing Bound). *For any $D \geq 2$, we have*

$$\log \mathcal{M}(\varepsilon, H^\beta(I^D) \otimes \mathcal{I}_\alpha, \|\cdot\|_{L^2(I^D)}) = \Theta\left(\varepsilon^{-D/\beta} + \varepsilon^{-2(D-1)/\alpha}\right), \quad (\varepsilon \rightarrow 0).$$

With the bound, we apply Theorem 11 and thus obtain the minimax lower bound of the estimation with $H^\beta(I^D) \otimes \mathcal{I}_\alpha$. By using the relation $\mathcal{R}(\mathcal{F}_{\alpha,\beta,M}^{PS}) \geq \mathcal{R}(H^\beta(I^D) \otimes \mathcal{I}_\alpha)$, we obtain the following theorem.

Theorem 13 (Minimax Rate for $\mathcal{F}_{\alpha,\beta,M}^{PS}$). *For any $\alpha, \beta \geq 1$ and $M \in \mathbb{N}$ we obtain*

$$\mathcal{R}(\mathcal{F}_{\alpha,\beta,M}^{PS}) = \Omega\left(\max\left\{n^{-2\beta/(2\beta+D)}, n^{-\alpha/(\alpha+D-1)}\right\}\right), \quad (n \rightarrow \infty).$$

This result indicates that the convergence rate by \hat{f}^{DL} is almost optimal in the minimax sense, since the rates in Theorem 7 correspond to the lower bound of Theorem 13 up to a log factor. In other words, for estimating $f^* \in \mathcal{F}_{\alpha,\beta,M}^{PS}$, no other methods could achieve a better rate than the estimators by DNNs.

4.4 Singularity Control by Deep Neural Network

In this section, we present an intuition for the optimality of DNNs for the functions with singularities. In the following, we will present that the approximation error of a non-smooth function by DNNs is as if the function to be approximated is smooth.

We consider an example with $D = 2$ and $M = 2$. Let $f^S \in \mathcal{F}_{\alpha,\beta,2}^{PS}$ be

$$f^S(x_1, x_2) = \mathbf{1}_R(x_1, x_2), \quad R = \{(x_1, x_2) \in I^2 \mid x_2 \geq h(x_1)\},$$

with a function $h \in H_1^\alpha(I)$. The function f^S is singular on the set $\{(x_1, x_2) \in I^2 \mid x_2 = h(x_1)\}$. Moreover, the function f^S is rewritten as

$$f^S(x_1, x_2) = \mathbf{1}_{\{ \cdot \geq 0 \}} \circ \underbrace{((x_1, x_2) \mapsto (x_2 - h(x_1)))}_{=: f^H(x_1, x_2)},$$

where $\mathbf{1}_{\{ \cdot \geq 0 \}}$ is the step function and $f^H \in H^\alpha(I^2)$ is a smooth function induced by h . We can rewrite the function f^S with the singularities as a composition of the step function and the smooth function f^H .

To approximate and estimate f^S , we consider an explicit function g_E by DNNs as $g_E(x_1, x_2) = g_s \circ g_H(x_1, x_2)$, where g_s is a DNN approximator for the step function by a DNN, and g_H is a DNN approximator for f^H . Subsequently, we can measure its approximation as

$$\|f^E - g_E\|_{L^2(I^2)} \lesssim \|\mathbf{1}_{\{ \cdot \geq 0 \}} - g_s\|_{L^2(I)} + \|f^H - g_H\|_{L^2(I^2)}.$$

For the right hand side, Lemma 4 indicates that the first term $\|\mathbf{1}_{\{ \cdot \geq 0 \}} - g_s\|_{L^2(I)}$ is negligible, because it is arbitrary small with a constant number of parameters. Hence, a dominant error appears in the second term $\|f^H - g_H\|_{L^2(I^2)}$, which is an approximation error of a smooth function $f^H \in H^\beta(I^2)$.

In summary, DNNs can approximate and estimate non-smooth f^E , as if f^E is a smooth function in $H^\beta(I^2)$. This is because DNNs can represent a composition of functions, which can eliminate the singularity of f^E .

5. Advantages of DNNs with Singularity

In this section, we compare the result of DNNs with several other methods.

5.1 Sub-optimality of Linear Estimators

We discuss sub-optimality of some of other standard methods in estimating piecewise smooth functions. To this end, we consider a class of *linear estimators*.

Definition 14 (Linear Estimator). *The class contains any estimators written as*

$$\hat{f}^{\text{lin}}(x) = \sum_{i \in [n]} \Upsilon_i(x; X_1, \dots, X_n) Y_i, \quad (6)$$

where Υ_i is an arbitrary measurable function which depends on X_1, \dots, X_n .

Linear estimators include various popular estimators such as kernel ridge regression, sieve regression, spline regression, and Gaussian process regression. If a regression model is a linear sum of (not necessarily orthogonal) basis functions and its parameter is a minimizer of the sum of square losses, the model is a linear estimator. Linear estimators have been studied extensively (e.g. Donoho and Johnstone (1998) and Korostelev and Tsybakov (2012), particularly Section 6 in Korostelev and Tsybakov (2012)), and the results can be adapted to our setting, thereby providing the following results.

In the following, we prove the sub-optimality of linear estimators:

Proposition 15 (Sub-Optimality of Linear Estimators). *For any $\alpha, \beta \geq 1$ and $D \in \mathbb{N}$, we obtain*

$$\sup_{f^* \in \mathcal{F}_{\alpha, \beta, M}^{PS}} \mathbb{E}_{f^*} \left[\|\hat{f}^{\text{lin}} - f^*\|_{L^2(P_X)}^2 \right] = \Omega \left(n^{-\alpha/(2\alpha+D-1)} \right).$$

This rate is slower than the minimax optimal rate $O \left(n^{-2\beta/(2\beta+D)} \vee n^{-\alpha/(\alpha+D-1)} \right)$ in Theorem 13 with the parameter configuration $\alpha < 2\beta(D-1)/(D-2\beta)$. This result implies that linear estimators perform worse than DNNs when it is relatively difficult to estimate a hypersurface with singularity.

Remark 16 (Linear Estimators with Cross-Validation). *We discuss a sub-optimality of a certain class of nonlinear estimators associated with cross-validation (CV). If we select the hyperparameters of a linear estimator by CV, the estimator is no longer a linear estimator. Typically, a bandwidth parameter of kernel methods and a number of basis functions for sieve estimators are examples of hyperparameters. Despite this fact, the generalization error of a CV-supported estimator is bounded below by that of the estimators with optimal hyperparameters, which are sometimes referred to as oracle estimators. Because oracle estimators are linear estimators in most cases, we can still claim that the CV-supported estimators are sub-optimal according to Proposition 15.*

Remark 17 (Relation to Other Non-Optimality of Linear Estimators). *The non-optimality of the linear estimator has been shown in several studies: Donoho and Johnstone (1998)*

showed the non-optimality of linear estimators in estimation for an element of a Besov space with a specific order. Proof by Donoho and Johnstone (1998) is based on the fact that a linear estimator cannot adapt to the optimal rate on a non-convex set of functions, and shows the non-convexity of a unit ball in the Besov space makes linear estimators sub-optimal.

In contrast, our result on the non-optimality is based on a different approach, which utilizes the shape of a hypersurface on which the singularity of f^* is placed. A linear estimator uses a level-set of its basis function Υ_i to approximate the hypersurface. However, due to the limited expressive power and flexibility of the utilization of the level-set, linear estimators cannot adapt to the higher-order smoothness of the hypersurface and lose their optimality. This proof is an application of the image analysis by Korostelev and Tsybakov (2012) to our setting with the basic pieces.

5.2 Sub-Optimality of Wavelet Estimator

We investigate the generalization error of a wavelet series estimator for $f^* \in \mathcal{F}_{\alpha, \beta, M}^{PS}$. The estimator by wavelets is one of the most common estimators for the nonparametric regression problem, and can attain the optimal rate in many settings (for a summary, see Giné and Nickl (2015)). Moreover, wavelets can handle discontinuous functions. Since this estimator is a linear estimator, the analysis with wavelets is strictly included in Proposition 15. However, we provide the results here, because it is useful for readers to understand how basis functions handle singularities through wavelets and other basis functions. In this section, we will prove the sub-optimality of the wavelet for the class of piecewise smooth functions. Since these methods are well known for dealing with singularities, their investigation can provide a better understanding of why linear estimators lose their optimality.

Intuitively, wavelets resolve singularity by fitting the support of their basis functions to the singularity's shape. That is, their approximation error decreases, when the supports fit the shape and fewer basis functions overlap the singularity hypersurface. As illustrated in the left panel of Figure 4, the wavelet divides I^D into cubes, with each basis is concentrated on each cube.




Figure 4: [Left] Hypersurface with singularity (black curve) in I^2 and cubes in which wavelet basis functions are concentrated. [Middle] Hypersurface with singularity (black curve) in I^2 and ellipses in which curvelet basis functions are concentrated. [Right] Domain on the frequency of the curvelet basis. Each basis function is concentrated on a colored shape.

To this end, we consider an orthogonal wavelet basis for $L^2(I)$. We define a set of indexes $\mathcal{H} := \{(j, k) : j \in \{-1, 0\} \cup \mathbb{N}, k \in K_j\}$ be the index set with a set $K_j = \{0\} \cup [2^j - 1]$ and

consider the wavelet basis $\{\phi_\kappa\}_{\kappa \in \mathcal{H}}$ of $L^2(\mathbb{R})$ with setting $\phi_{-1,k}$ be a shifted scaling function. Then, we restrict its domain to I and obtain the wavelet basis for $L^2(I)$.

For a wavelet analysis for multivariate functions, we consider a tensor product of the basis $\{\phi_\kappa\}_{\kappa \in \mathcal{H}}$. Namely, let us define $\Phi_{\kappa_1, \dots, \kappa_D}(x) := \prod_{d \in [D]} \phi_{\kappa_d}(x_d)$ with $x \in \mathbb{R}^D$. Then, consider an orthonormal basis $\{\Phi_{\kappa_1, \dots, \kappa_D}\}_{(\kappa_1, \dots, \kappa_D) \in \mathcal{H}^{\times D}}$ for $L^2(\mathbb{R}^D)$, then restrict it to $L^2(I^D)$. $\mathcal{H}^{\times D}$ is a D -times direct product of \mathcal{H} . Then, a decomposition of a restricted function $f \in L^2(I^D)$ is formulated as

$$f = \sum_{(\kappa_1, \dots, \kappa_D) \in \mathcal{H}^{\times D}} w_{\kappa_1, \dots, \kappa_D}(f) \Phi_{\kappa_1, \dots, \kappa_D},$$

where $w_{\kappa_1, \dots, \kappa_D}(f) = \langle f, \Phi_{\kappa_1, \dots, \kappa_D} \rangle$.

We define an estimator of f^* by the wavelet decomposition. For simplicity, let P_X be the uniform distribution on I^D . Using a truncation parameter $\tau \in \mathbb{N}$, we set $\mathcal{H}_\tau := \{(j, k) : j \in \{-1, 0\} \cup [\tau], k \in K_j\} \subset \mathcal{H}$ be a subset of indexes. Since the decomposition is a linear sum of orthogonal basis, an wavelet estimator \hat{f}^{wav} which minimizes an empirical squared loss has the following form

$$\hat{f}^{\text{wav}} = \sum_{(\kappa_1, \dots, \kappa_D) \in \mathcal{H}_\tau^{\times D}} \hat{w}_{\kappa_1, \dots, \kappa_D} \Phi_{\kappa_1, \dots, \kappa_D}.$$

Moreover, $\hat{w}_{\kappa_1, \dots, \kappa_D}$ is an empirical analogue version of the inner product as $\hat{w}_{\kappa_1, \dots, \kappa_D} = \frac{1}{n} \sum_{i \in [n]} Y_i \Phi_{\kappa_1, \dots, \kappa_D}(X_i)$. τ is selected to follow an order in n which balances a bias and variance of the error and minimizes its generalization error. With the wavelet estimator, we obtain the following result:

Proposition 18 (Sub-Optimality of Wavelets). *For any $\alpha, \beta \geq 1$ and $D \geq 2$, we obtain*

$$\sup_{f^* \in \mathcal{F}_{\alpha, \beta, M}^{PS}} \mathbb{E}_{f^*} \left[\|\hat{f}^{\text{wav}} - f^*\|_{L^2(P_X)}^2 \right] = \Omega \left(n^{-1/2} \right).$$

When comparing the derived rate with Theorem 13, we can find that the wavelet estimator cannot attain the minimax rate when both $\beta > D/2$ and $\alpha > (D-1)$ hold. This result describes wavelets cannot adapt a higher smoothness of f^* and the boundary of pieces.

Note that the truncation parameter τ in this estimator depends on unknown data distribution, and it should be chosen in a data-dependent way in practice. Even in this case, the same lower bound in Proposition 18 holds by the data-dependent choice, because the current data-independent choice is the optimal choice that minimizes the error.

Schmidt-Hieber (2020) shows that the wavelet estimator is sub-optimal in estimation of composite functions. Although our wavelet estimator is identical to that of Schmidt-Hieber (2020), the proof of the non-optimality is different. This is because our proof utilizes the difficulty of approximating the singularity of the true function, but the composite function by Schmidt-Hieber (2020) has no singularity. This argument is similar to Remark 17.

5.3 Sub-optimality of Harmonic Based Estimator

We investigate another estimator from the harmonic analysis and its optimality. The harmonic analysis provides several methods for non-smooth structures, such as curvelets

(Candes and Donoho, 2002; Candès and Donoho, 2004) and shearlets (Kutyniok and Lim, 2011). The methods are designed to approximate piecewise smooth functions on pieces with C^2 boundaries. As illustrated in Figure 4, each base of the curvelet is concentrated on an ellipse with different scales, locations, and angles in I^2 . The ellipses covering the hypersurface resolve the singularity. Each basis has a fan-shaped support with a different radius and angle in the frequency domain (see Candès and Donoho (2004) for details).

We focus on curvelets as one of the most common methods. For brevity, we study the case with $D = 2$, which is of primary concern for curvelets. Curvelets can be extended to higher dimensions $D \geq 3$, and we can study the case in a similar manner.

As preparation, we define curvelets. In this analysis, we consider the domain of f^* is $[-1, 1]^D$ for technical simplification. Furthermore, we set P_X as the uniform distribution on $[-1, 1]^D$, as similar to the wavelet case. Let $\mu := (j, \ell, k)$ be a tuple of scale index $j = 0, 1, 2, \dots$, rotation index $\ell = 0, 1, 2, \dots, 2^j$, and location parameter $k \in \mathbb{Z}^2$. For each μ , we define the parabolic scaling matrix $D_j = \begin{pmatrix} 2^{2j} & 0 \\ 0 & 2^j \end{pmatrix}$, the rotation angle $\theta_{j,\ell} = 2\pi 2^{-j}\ell$, and a location $k = k_\delta = (k_1\delta_1, k_2\delta_2)$ with hyper-parameters δ_1 and δ_2 , where $\delta_1, \delta_2 > 0$. Thereafter, we consider a curvelet $\gamma_\mu : \mathbb{R}^2 \rightarrow \mathbb{R}$ as $\gamma_\mu(x) = 2^{3j/2}\gamma(D_j R_{\theta_{j,\ell}}x - k_\delta)$. In this case, γ is defined by an inverse Fourier transform of a localized function in the Fourier domain. Figure 4 provides an illustration of its support and its rigorous definition is deferred to the appendix. According to (Candès and Donoho, 2004), it is shown that $\{\gamma_\mu\}_{\mu \in \mathcal{L}}$ is a tight frame in $L^2(\mathbb{R}^2)$, where \mathcal{L} is a set of μ . Hence, we obtain the following formulation $f(x) = \sum_{\mu \in \mathcal{L}} w_\mu(f)\gamma_\mu(x)$, where $w_\mu(f) = \langle f, \gamma_\mu \rangle$. For estimation, we consider the truncation parameter $\tau \in \mathbb{N}$ and define an index subset as $\mathcal{L}_\tau := \{\mu \mid j \in [\tau]\} \subset \mathcal{L}$. Thereafter, similar to the wavelet case, a curvelet estimator which minimizes an empirical squared loss is written as

$$\hat{f}^{\text{curve}}(x) = \sum_{\mu \in \mathcal{L}_\tau} \hat{w}_\mu \gamma_\mu(x), \text{ where } \hat{w}_\mu = \frac{1}{n} \sum_{i \in [n]} Y_i \gamma_\mu(X_i).$$

In this case, τ is selected to minimize the generalization error of the curvelet estimator. We then obtain the following statement.

Proposition 19 (Sub-Optimality of Curvelets). *For $D = 2$ and any $\alpha, \beta \geq 2$, we obtain*

$$\sup_{f^* \in \mathcal{F}_{\alpha, \beta, M}} \mathbb{E}_{f^*} \left[\|\hat{f}^{\text{curve}} - f^*\|_{L^2(P_X)}^2 \right] = \Omega \left(n^{-1/3} \right).$$

The result implies the sub-optimality of the curvelet; that is, the rate is slower than the minimax rate when $\beta > D/4 = 1/2$ and $\alpha > (D - 1)/2 = 1/2$. Similar to the wavelet estimator, the curvelet estimator does not adapt to the higher smoothness in the nonparametric regression setting.

5.4 Advantage of DNNs against the Other Methods

We summarize the sub-optimality of the other methods and compare them with DNNs. To this end, we provide a formal statement for comparing the estimator \hat{f}^{DL} by DNNs and the other estimators, namely, \hat{f}^{lin} , \hat{f}^{vac} , and \hat{f}^{curve} . The following corollary is the second main result of this study:




Figure 5: Parameter spaces of (α, β) . [Left] The red dashed line is $\alpha = 2\beta(D-1)/(D-2\beta)$. The red region presents $\{(\alpha, \beta) \mid \alpha < 2\beta(D-1)/(D-2\beta)\}$, which a set of parameter configurations such that (7) holds with $\tilde{f} = \hat{f}^{\text{lin}}$. [Right] The blue region presents $\{(\alpha, \beta) \mid \alpha > D-1, \beta > D/2\}$, which a set of parameter configurations such that (7) holds with $\tilde{f} = \hat{f}^{\text{wav}}$. The blue dashed line is its boundary.

Corollary 20 (Advantage of DNNs). *Fix $M \geq 2$. If $\alpha < 2\beta(D-1)/D$ holds with $\alpha, \beta \geq 1$ and $D \geq 2$, the estimator $\tilde{f} = \hat{f}^{\text{lin}}$ satisfies*

$$\sup_{f^* \in \mathcal{F}_{\alpha, \beta, M}^{PS}} \mathbb{E}_{f^*} \left[\|\tilde{f} - f^*\|_{L^2(P_X)}^2 \right] \gtrsim \sup_{f^* \in \mathcal{F}_{\alpha, \beta, M}^{PS}} \mathbb{E}_{f^*} \left[\|\hat{f}^{DL} - f^*\|_{L^2(P_X)}^2 \right]. \quad (7)$$

If $\beta > D/2$ and $\alpha > D-1$ hold, (7) holds with $\tilde{f} = \hat{f}^{\text{wav}}$.

If $\beta > D/4$ and $\alpha > (D-1)/2$ hold with $D = 2$, (7) holds with $\tilde{f} = \hat{f}^{\text{curve}}$.

These results are naturally derived from the discussion of the minimax optimal rate of DNN (Corollary 8 and Theorem 13), and the sub-optimality of the other methods (Propositions 15, 18, and 19). We can conclude that DNNs offer a theoretical advantage over the other methods in terms of estimation for functions with singularity; that is, with the parameter configurations in Corollary 20, there exist $f^* \in \mathcal{F}_{\alpha, \beta, M}^{PS}$, such that

$$\mathbb{E}_{f^*} \left[\|\hat{f}^{DL} - f^*\|_{L^2(P_X)}^2 \right] < \mathbb{E}_{f^*} \left[\|\tilde{f} - f^*\|_{L^2(P_X)}^2 \right]$$

holds for a sufficiently large n .

The parameter configurations introduced in Corollary 20 are classified into two cases, as illustrated in Figure 5. First, the configuration for the sub-optimality of the linear estimators appears when α is relatively smaller than β . With a small α , the rate $O(n^{\alpha/(\alpha+D-1)})$ in the minimax optimal rate dominates the generalization error, because it is more difficult to estimate the boundary of pieces than to estimate the function within the pieces. In this case, the linear estimators lose optimality, because their generalization error is sensitive to α . The second case is that in which both α and β are above certain thresholds. It is thought that the approximation of the singularity by the orthogonal bases has not been adapted to the higher smoothness of f^* .

We add a further explanation for the limitations of the other estimators by describing the difficulty of shape fitting to the hypersurface of singularities. The other estimators take

the form of sums of (not necessarily orthogonal) basis functions, and each base has a (nearly) compact support. When the other methods approximate a function with the hypersurface of singularities, they approximate the hypersurface by fitting its support. For example, consider the curvelet of which the basis function has an ellipsoid-shaped (nearly) compact support. The ellipsoid is fitted to the hypersurface of the singularity with rotation, as illustrated in Figure 6. The number of bases determines the magnitude of the error. However, when the hypersurface has higher-order smoothness, the fitting in the domain cannot adapt to the hypersurface with optimality.

DNNs do not use shape fitting to the hypersurface for singularity, but represent the hypersurface using a composition of functions, as mentioned in Section 4.4. Therefore, even if the hypersurface has larger smoothness, DNNs can handle this without losing efficiency.




Figure 6: [Left] An example of piecewise smooth functions f^C such as the indicator function of a disk. [Right] The singularities of f^C (black circle), and an illustration of approximation for the circle by a curvelet basis. The curvelet basis concentrates in the gray ellipses, that cannot have sufficient approximation ability when the curves are too smooth.

Remark 21 (Limitations of the advantage with learning algorithms). *We also need to reiterate the limitations of these advantages we claim to have. Needless to say, our analysis is based on the setting that a global optimum of the loss minimization problem has been found in (4). In other words, we only study the errors from an approximation power and a degree of freedom of DNN models, and ignore the errors from the training algorithm of DNNs in practical use. However, this algorithmic error requires a completely different analysis, and analyzing them all at the same time would be very difficult and not very meaningful. What we argue in this work is that DNNs can handle singularity more optimally than existing linear estimators, including those for singularity resolution, even if only in terms of approximation power and a degree of freedom derived error aspects.*

6. Conclusion

In this study, we have derived theoretical results that explain why DNNs outperform other methods. We considered the regression setting in the situation whereby the true function is singular on a smooth hypersurface in its domain. We derived the convergence rates of the estimator obtained by DNNs and proved that the rates were almost optimal in the minimax sense. We explained that the optimality of DNNs originates from their composition structure, which can resolve singularity. Furthermore, to analyze the advantage of DNNs, we investigated the sub-optimality of several other estimators, such as linear, wavelet, and

curvelet estimators. We proved the sub-optimality of each estimator with certain parameter configurations. This advantage of DNNs comes from the fact that the shape of smooth curves for the singularity can be handled by DNNs, while the other methods fail to capture the shape efficiently. Theoretically, this is a vital step for analyzing the mechanism of DNNs.

Acknowledgments

We have greatly benefited from insightful comments and suggestions by Alexandre Tsybakov, Taiji Suzuki, Bharath K Sriperumbudur, Johannes Schmidt-Hieber, and Motonobu Kanagawa. We also thank Samory Kpotufe and the anonymous reviewers for thoughtful and constructive comments. M.Imaizumi was supported by JSPS KAKENHI Grant Number 18K18114 and JST Presto Grant Number JPMJPR1852.

Appendix A. Proof Overview

A.1 Upper Bound on Generalization Error (Theorem 7)

We give an overview of the proof of Theorem 7, which derives the upper bound on the expected generalization error $\mathbb{E}_{f^*}[\|\hat{f}^{DL} - f^*\|_{L^2(P_X)}^2]$. As shown in (5), the generalization error is decomposed into the approximation error \mathcal{B} and the complexity error \mathcal{V} . In the proof, we evaluate these two terms separately. A flowchart of the overview is shown in Figure 7.




Figure 7: Flowchart showing the relationship between theorems/lemmas/propositions for Theorem 7 on the upper bound of the generalization error. Theorem 7 is mainly derived from the two groups of results bounding the approximation and complexity errors, and also is applied to the additional results. When multiple theorems/lemmas/propositions are specified within one box, it represents a more generalised result or a variant with different activation functions.

The evaluation of the approximation error \mathcal{B} , which is carried out in Section B.1, uses the evaluation of an approximation error of the piecewise smooth functions $\mathcal{F}_{\alpha,\beta,M}^{PS}$ by DNNs. According to the definition of piecewise smooth functions $f \in \mathcal{F}_{\alpha,\beta,M}^{PS}$ as

$$f = \sum_{m \in [M]} f_m \otimes \mathbf{1}_{R_m}, \text{ where } f_m \in H_F^\beta(I^D) \text{ and } \{R_m\}_{m \in [M]} \in \mathcal{R}_{\alpha,M}.$$

the approximation error is described by an approximation error of a simple smooth function $f_m \in H_F^\beta(I^D)$ and the approximation error of the indicator function on a piece $\mathbf{1}_{R_m}$. This basic idea is originally proposed by Petersen and Voigtlaender (2018). To approximate $f_m \in H_F^\beta(I^D)$, we apply an approximation technique based on an aggregation of local polynomial functions, which is originally developed by Yarotsky (2017). Since this technique

requires a step function $\mathbf{1}_{\{\geq 0\}}$, we show that a step function $\mathbf{1}_{\{\geq 0\}}$ can be well approximated by the wide class of admissible activation functions in Lemma 4 and 26. This approximation on a step function is a key part of our proof, and the approximation results on smooth functions are updated to adapt to this key part. Based on this result, we can approximate polynomials in Lemma 24 and 25, and a smooth function $f \in H_F^\beta$ in Lemma 3 and 27. In contrast, approximation of the indicator function $\mathbf{1}_{R_m}$ is evaluated by combining the approximation on both a step function $\mathbf{1}_{\{\geq 0\}}$ and a smooth function $f \in H_F^\beta(I^D)$. This approach takes advantage of the composition structure of DNNs described in Section 4.4. As a result, the final approximation error has the following rate:

$$\mathcal{B} = \tilde{O}(S^{-2\beta/D} + S^{-\alpha/(D-1)}),$$

which is a sum of the approximation errors of $f \in H_F^\beta(I^D)$ and $\mathbf{1}_{R_m}$.

The evaluation of the complexity error \mathcal{V} in Section B.2 is based on a covering number analysis on the function set by DNNs $\mathcal{G}(L, S, B)$. This makes use of the technique of peeling from the empirical process theory, which is developed, for example, by Koltchinskii (2006), and its application to DNNs is proposed in Schmidt-Hieber (2020) and others. Specifically, we define a localized subset of functions around a specific function $f \in \mathcal{G}(L, S, B)$, then use uniform convergence over a finite subset in the localized subset to bound the complexity error \mathcal{V} . Since we use a covering set to construct the finite subset, the covering number $\log \mathcal{N}(\varepsilon, \mathcal{G}(L, S, B), \|\cdot\|_{L^\infty(I^D)})$ of the DNN functions $\mathcal{G}(L, S, B)$ becomes essential. Briefly summarised, an expectation of the complexity error has the following order

$$O\left(\frac{\log \mathcal{N}(1/n, \mathcal{G}(L, S, B), \|\cdot\|_{L^\infty(I^D)})}{n}\right) = \tilde{O}\left(\frac{SL \log(SLB)}{n}\right),$$

with a number of parameters S and layers L and a parameter volume B . The equality follows the covering number bound in Lemma 23 from Nakada and Imaizumi (2020).

The combination of the above two evaluations bounds the generalization error $\mathbb{E}_{f^*}[\|\hat{f}^{DL} - f^*\|_{L^2(P_X)}^2]$, together with a choice of S , L , and B . By selecting L and B as a polynomial $\log n$, the generalisation error is bounded as

$$\mathbb{E}_{f^*}[\|\hat{f}^{DL} - f^*\|_{L^2(P_X)}^2] = \tilde{O}(S^{-2\beta/D} + S^{-\alpha/(D-1)}) + \tilde{O}(S/n) \quad (8)$$

In (8), the first item is an upper bound on the approximation error \mathcal{B} , and the second term is a bound on the complexity error \mathcal{V} . Since the approximation error \mathcal{B} is decreasing in S and the complexity error \mathcal{V} is increasing in S , we select the number of parameter $S = \Theta(n^{D/(2\beta+D)} + n^{(D-1)/(\alpha+D-1)})$ to handle the trade-off between \mathcal{B} and \mathcal{V} . Then, we obtain the statement of Theorem 7. By applying this result, the convergence rate is derived.

A.2 Lower Bound of Generalization Error (Theorem 13)

We provide an overview of the proof of Theorem 13 for the lower bound on the expected generalization error. The proof depends mainly on the information-theoretic analysis of the minimax rate, which is developed by Yang and Barron (1999). We cite Theorem 6 of Yang and Barron (1999) as Theorem 11 in this paper, then focus on a packing number $\mathcal{M}(\varepsilon_n, \mathcal{F}_{\alpha, \beta, M}^{PS}, \|\cdot\|_{L^2})$ to apply the theory. In Proposition 12, we investigate packing numbers of each element set that constitutes $\mathcal{F}_{\alpha, \beta, M}^{PS}$. We utilize several results from Dudley (2014) and derive the bound for the packing number. Figure 8 illustrates the relation.




Figure 8: Flowchart for Theorem 13 on the lower bound of the generalization error.

Appendix B. Proof of Theorem 7

We start with the basic inequality (5). As preparation, we introduce additional notation. Given an empirical measure, the empirical (pseudo) norm of a random variable is defined by $\|Y\|_n := (n^{-1} \sum_{i \in [n]} Y_i^2)^{1/2}$ and $\|\xi\|_n := (n^{-1} \sum_{i \in [n]} \xi_i^2)^{1/2}$. By the definition of \hat{f}^{DL} in (4), we obtain the following inequality

$$\|Y - \hat{f}^{DL}\|_n^2 \leq \|Y - f\|_n^2$$

for all $f \in \mathcal{G}(L, S, B)$. It follows from $Y_i = f^*(X_i) + \xi_i$ that

$$\|f^* + \xi - \hat{f}^{DL}\|_n^2 \leq \|f^* + \xi - f\|_n^2,$$

then, simple calculation yields

$$\|\hat{f}^{DL} - f^*\|_n^2 \leq \|f^* - f\|_n^2 + \frac{2}{n} \sum_{i \in [n]} \xi_i (\hat{f}^{DL}(X_i) - f(X_i)).$$

We set $f \in \mathcal{G}(L, S, B)$ as satisfying $\|f^* - f\|_n = \inf_{f' \in \mathcal{G}(L, S, B)} \|f^* - f'\|_n$, then we obtain

$$\begin{aligned} \|\hat{f}^{DL} - f^*\|_n^2 &\leq \inf_{f' \in \mathcal{G}(L, S, B)} \|f^* - f'\|_n^2 + \frac{2}{n} \sum_{i \in [n]} \xi_i (\hat{f}^{DL}(X_i) - f(X_i)) \\ &=: \mathcal{B} + \mathcal{V}. \end{aligned} \tag{9}$$

In the first subsection, we bound \mathcal{B} by investigating an approximation power of DNNs. In the second subsection, we evaluate \mathcal{V} by evaluating the complexity of the estimator. Afterward, we combine the results and derive an overall rate.

B.1 Approximate Piecewise Smooth Functions by DNNs

In this subsection, we provide proof of Theorem 6. The result follows the following proposition:

Proposition 22 (General Version of Theorem 6). *Suppose Assumption 1 holds. Then, for any $\varepsilon_1 \in (0, 1)$ and $\varepsilon_2 \in (0, 1)$, there exists a tuple (L, S, B) such as*

- $L = C_{\alpha, \beta, D, F}([\alpha] + [\beta] + \log_2(1/\varepsilon_1) + \log_2(M/\varepsilon_2) + 1)$,
- $S = C_{\alpha, \beta, D, F, J}(M\varepsilon_1^{-D/\beta}(\log_2(1/\varepsilon_1))^2 + (\varepsilon_2/M)^{-2(D-1)/\alpha} \log_2(M/\varepsilon_2)^2 + M(\log_2(M/\varepsilon_2))^2)$,
- $B = C_{F, M, q}(\varepsilon_1 \wedge \varepsilon_2)^{-16 \wedge -C_{\alpha, \beta}}$,

which satisfy

$$\inf_{f \in \mathcal{G}(L, S, B)} \sup_{f^* \in \mathcal{F}_{\alpha, \beta, M}^{PS}} \|f - f^*\|_{L^2} \leq \varepsilon_1 + \varepsilon_2.$$

Proof [Proof of Proposition 22] Fix $f^* \in \mathcal{F}_{\alpha, \beta, M}^{PS}$ such that $f^* = \sum_{m \in [M]} f_m^* \otimes \mathbf{1}_{R_m^*}$ with $f_m^* \in H_F^\beta(I^D)$ and $\{R_m\}_{m \in [M]} \in \mathcal{R}_{\alpha, M}$ for $m \in [M]$. By Lemma 27, for any $\delta_1 \in (0, 1)$, there exist a constant $c_1 > 0$ and functions $g_{f,1}, \dots, g_{f,M} \in \mathcal{G}(C_{\beta, D, F}([\beta] + \log_2(1/\delta_1) + 1), C_{\beta, D, F} \delta_1^{-D/\beta} (\log_2(1/\delta_1))^2, C_{F, q} \delta_1^{-16 \wedge -C_\beta})$ such that $\|g_{f,m} - f_m^*\|_{L^2(R_m)} \leq \text{vol}(R_m) \delta_1$ for $m \in [M]$. Similarly, by Lemma 29, we can find $g_R \in \mathcal{G}(C_{\alpha, D, F, J}([\alpha] + \log_2(1/\delta_2) + 1), C_{\alpha, D, F, J}(\delta_2^{-2(D-1)/\alpha} \log_2(1/\delta_2)^2 + 1 + M(\log_2(1/\delta_2))^2), C_{F, q} \delta_2^{-16 \wedge -C_\alpha})$ such that $\|g_{R,m} - \mathbf{1}_{R_m^*}\|_{L^2} \leq \delta_2$ for $\delta_2 \in (0, 1)$. For approximation, we follow (25) in the proof of Lemma 25 as $g_c \in \mathcal{G}(\log_2(1/\delta_1), C((\log_2(1/\delta_1))^2 + 1), C)$ which approximates a multiplication $\|((x, x') \mapsto xx') - g_c(x, x')\|_{L^\infty([-F, F]^2)} \leq F^2 \delta_1$.

With these components, we construct a function $\hat{g} \in \mathcal{G}(C_{\alpha, \beta, D, F}([\alpha] + [\beta] + \log_2(1/\delta_1) + \log_2(1/\delta_1) + 1), C_{\alpha, \beta, D, F, J}(M\delta_1^{-D/\beta}(\log_2(1/\delta_1))^2 + \delta_2^{-(D-1)/\alpha}(\log_2(1/\delta_2))^2 + 1 + M(\log_2(1/\delta_2))^2), C_{F, q}(\delta_1^{-16 \wedge C_\alpha} \wedge \delta_2^{-16 \wedge C_\beta}))$ as

$$\hat{g}(x) = \sum_{m \in [M]} g_c(g_{f,m}(x), g_{R,m}(x)), \quad (10)$$

by setting a parameter matrix $A_L = (1, 1, \dots, 1)^\top$ in the last layer of \hat{g} . Then, we evaluate the distance between f^* and the combined DNN:

$$\begin{aligned} & \|f^* - \hat{g}\|_{L^2(I^D)} \\ &= \left\| \sum_{m \in [M]} f_m^* \mathbf{1}_{R_m^*} - \sum_{m \in [M]} g_c(g_{f,m}(\cdot), g_{R,m}(\cdot)) \right\|_{L^2(I^D)} \\ &\leq \sum_{m \in [M]} \|f_m^* \otimes \mathbf{1}_{R_m^*} - g_{f,m} \otimes g_{R,m}\|_{L^2(I^D)} \\ &\quad + \sum_{m \in [M]} \|g_{f,m} \otimes g_{R,m} - g_c(g_{f,m}(\cdot), g_{R,m}(\cdot))\|_{L^2(I^D)} \\ &\leq \sum_{m \in [M]} \|(f_m^* - g_{f,m}) \otimes \mathbf{1}_{R_m^*}\|_{L^2(I^D)} + \sum_{m \in [M]} \|g_{f,m} \otimes (\mathbf{1}_{R_m^*} - g_{R,m})\|_{L^2(I^D)} \\ &\quad + \sum_{m \in [M]} \|g_{f,m} \otimes g_{R,m} - g_c(g_{f,m}(\cdot), g_{R,m}(\cdot))\|_{L^2(I^D)} \end{aligned}$$

$$=: \sum_{m \in [M]} B_{1,m} + \sum_{m \in [M]} B_{2,m} + \sum_{m \in [M]} B_{3,m}. \quad (11)$$

We will bound $B_{1,m}$, $B_{2,m}$ and $B_{3,m}$ for $m \in [M]$. About $B_{1,m}$, the choice of $g_{f,m}$ gives

$$B_{1,m} = \|(f_m^* - g_{f,m}) \otimes \mathbf{1}_{R_m^*}\|_{L^2(I^D)} = \|f_m^* - g_{f,m}\|_{L^2(R_m^*)} \leq \text{vol}(R_m^*) \delta_1.$$

About $B_{2,m}$, similarly, the Hölder inequality yields

$$\begin{aligned} B_{2,m} &= \|g_{f,m} \otimes (\mathbf{1}_{R_m^*} - g_{R,m})\|_{L^2(I^D)} \\ &\leq \|g_{f,m}\|_{L^\infty(I^D)} \|\mathbf{1}_{R_m^*} - g_{R,m}\|_{L^2(I^D)} \\ &\leq (1 + \delta_2) \delta_2. \end{aligned}$$

About $B_{3,m}$, since $g_{f,m}$ and $g_{R,m}$ is a bounded function by F , we obtain

$$B_{3,m} \leq \|(x, x' \mapsto xx') \rightarrow g_c\|_{L^\infty([-F, F]^2)} \leq F^2 \delta_2.$$

We combine the results about $B_{1,m}$, $B_{2,m}$ and $B_{3,m}$. Substituting the bounds for (11) yields

$$\begin{aligned} \|f^* - \hat{g}\|_{L^2(I^D)} &\leq \sum_{m \in [M]} \{\text{vol}(R_m^*) \delta_1 + \delta_2 + \delta_2^2 + F^2 \delta_2\} \\ &\leq \delta_1 + M \delta_2 + M \delta_2^2 + M F^2 \delta_2, \end{aligned}$$

where the second inequality follows $\sum_{m \in [M]} \text{vol}(R_m^*) = \text{vol}(I^D) = 1$. Set $\delta_1 = \varepsilon_1$ and $\delta_2 = C_F \varepsilon_2 / M$. Then we obtain

$$\|f^* - \hat{g}\|_{L^2(I^D)} \leq \varepsilon_1 + \varepsilon_2 / 2 + \varepsilon_2^2 / 2 \leq \varepsilon_1 + \varepsilon_2.$$

Adjusting the coefficients, we obtain the statement. ■

We are now ready to prove Theorem 6.

Proof [Proof of Theorem 6] We note that $\mathcal{G}(L, S, B)$ has an inclusion property, namely, for $L' \geq L, S' \geq S$ and $B' \geq B$, we obtain

$$\mathcal{G}(L', S', B') \supseteq \mathcal{G}(L, S, B).$$

Applying Proposition 22 and adjusting the coefficients yield the statement. ■

B.2 Combining the Bounds

Here, we evaluate \mathcal{V} in (9) by the empirical process and its applications (Koltchinskii, 2006; Giné and Nickl, 2015; Schmidt-Hieber, 2020). We then combine the result with Theorem 6, obtaining Theorem 7. Recall that F denotes an upper bound of $g \in \mathcal{G}(L, S, B)$ by its definition.

Proof [Proof of Theorem 7] The proof starts with the basis inequality (9) and follows the following two steps: (i) apply the covering number bound for \mathcal{V} in (9), and (ii) combine the results with the approximation result (Theorem 6) on \mathcal{B} .

Step (i). Covering bound for the cross term. We bound an expected term $|\mathbb{E}[\mathcal{V}]| = |\mathbb{E}[\frac{2}{n} \sum_{i \in [n]} \xi_i(\hat{f}^{DL}(X_i) - f(X_i))]|$ by the covering number of $\mathcal{G}(L, S, B)$. We fix $\delta \in (0, 1)$ and consider a covering set $\{g_j\}_{j=1}^N \subset \mathcal{G}(L, S, B)$ for $N = \mathcal{N}(\delta, \mathcal{G}(L, S, B), \|\cdot\|_{L^\infty})$, that is, for any $g \in \mathcal{G}(L, S, B)$, there exists g_j with $j \in [N]$ such that $\|g - g_j\|_{L^\infty} \leq \delta$. For \hat{f}^{DL} , let $\hat{j} \in [N]$ be such that $\|\hat{f}^{DL} - g_{\hat{j}}\|_{L^\infty} \leq \delta$ holds. Then, we bound the expected term as

$$\begin{aligned}
 & \left| \mathbb{E} \left[\frac{2}{n} \sum_{i \in [n]} \xi_i(\hat{f}^{DL}(X_i) - f(X_i)) \right] \right| \\
 & \leq \left| \mathbb{E} \left[\frac{2}{n} \sum_{i \in [n]} \xi_i(\hat{f}^{DL}(X_i) - g_{\hat{j}}(X_i)) \right] \right| + \left| \mathbb{E} \left[\frac{2}{n} \sum_{i \in [n]} \xi_i(g_{\hat{j}}(X_i) - f(X_i)) \right] \right| \\
 & \leq 2\delta \mathbb{E} \left[\frac{1}{n} \sum_{i \in [n]} |\xi_i| \right] + \mathbb{E} \left[\left(\frac{\|\hat{f}^{DL} - f\|_n + \delta}{\|g_{\hat{j}} - f\|_n} \right) \left| \frac{2}{n} \sum_{i \in [n]} \xi_i(g_{\hat{j}}(X_i) - f(X_i)) \right| \right] \\
 & \leq 2\delta(2\sigma^2/\pi)^{1/2} + \mathbb{E} \left[\left(\frac{\|\hat{f}^{DL} - f\|_n + \delta}{\|g_{\hat{j}} - f\|_n} \right) \left| \frac{2}{n} \sum_{i \in [n]} \xi_i(g_{\hat{j}}(X_i) - f(X_i)) \right| \right] \\
 & \leq 2\delta(2\sigma^2/\pi)^{1/2} + 2\mathbb{E} \left[\underbrace{\left(\frac{\|\hat{f}^{DL} - f\|_n + \delta}{\sqrt{n}} \right) \left| \frac{\sum_{i \in [n]} \xi_i(g_{\hat{j}}(X_i) - f(X_i))}{\sqrt{n}\|g_{\hat{j}} - f\|_n} \right|}_{=: \eta_{\hat{j}}} \right] \\
 & \leq 2\delta(2\sigma^2/\pi)^{1/2} + \frac{2(\mathbb{E}[\|\hat{f}^{DL} - f\|_n^2]^{1/2} + \delta)}{\sqrt{n}} \mathbb{E}[\eta_{\hat{j}}^2]^{1/2},
 \end{aligned}$$

where the second inequality follows $\|g_{\hat{j}} - f\|_n \leq \|\hat{f}^{DL} - f\|_n + \delta$, and the last inequality follows the Cauchy-Schwartz inequality. With conditional on the observed covariates X_1, \dots, X_n , η_j follows a centered Gaussian distribution with its variance σ^2 , hence $\mathbb{E}[\eta_j^2] \leq \sigma^2 \mathbb{E}[\max_{j \in [N]} \eta_j^2] \leq 3 \log N + 1$ by Lemma C.1 in (Schmidt-Hieber, 2020). Then, we have

$$\begin{aligned}
 |\mathbb{E}[\mathcal{V}]| & \leq 2\delta(2\sigma^2/\pi)^{1/2} + \frac{2\sigma^2(\mathbb{E}[\|\hat{f}^{DL} - f\|_n^2]^{1/2} + \delta)}{\sqrt{n}} (3 \log \mathcal{N}(\delta, \mathcal{G}(L, S, B), \|\cdot\|_{L^\infty}) + 1)^{1/2} \\
 & \leq c_N \delta + \frac{2\sigma^2 \mathbb{E}[\|\hat{f}^{DL} - f\|_n^2]^{1/2}}{\sqrt{n}} (3 \log \mathcal{N}(\delta, \mathcal{G}(L, S, B), \|\cdot\|_{L^\infty}) + 1)^{1/2}, \tag{12}
 \end{aligned}$$

where $c_N > 0$ is a constant. The last inequality with c_N follows $\mathcal{N}(\delta, \mathcal{G}(L, S, B), \|\cdot\|_{L^\infty})/\sqrt{n} = O(1)$.

Step (ii). Combine the results. We combine the results the bound for \mathcal{B} by Theorem 6 and $|\mathbb{E}[\mathcal{V}]|$ in the Step (i), then evaluate $\mathbb{E}[\|\hat{f}^{DL} - f^*\|_{L^2(P_X)}]$. Combining the bound (12)

with (9) yields that

$$\begin{aligned} & \mathbb{E} \left[\|\hat{f}^{DL} - f^*\|_n^2 \right] \\ & \leq \mathbb{E}[\mathcal{B}] + c_N \delta + \frac{2\sigma^2 \mathbb{E}[\|\hat{f}^{DL} - f\|_n^2]^{1/2}}{\sqrt{n}} (3 \log \mathcal{N}(\delta, \mathcal{G}(L, S, B), \|\cdot\|_{L^\infty}) + 1)^{1/2}. \end{aligned}$$

For any $a, b, c \in \mathbb{R}$, $a \leq b + c\sqrt{a}$ implies $a^2 \leq c^2 + 2b$. Hence, we obtain

$$\begin{aligned} & \mathbb{E} \left[\|\hat{f}^{DL} - f^*\|_n^2 \right] \\ & \leq 2\mathbb{E}[\mathcal{B}] + 2c_N \delta + \frac{12\sigma^2 \log \mathcal{N}(\delta, \mathcal{G}(L, S, B), \|\cdot\|_{L^\infty}) + 4}{n} \\ & = 2\mathbb{E} \left[\|\hat{g} - f^*\|_n^2 \right] + 2c_N \delta + \frac{12\sigma^2 \log \mathcal{N}(\delta, \mathcal{G}(L, S, B), \|\cdot\|_{L^\infty}) + 4}{n}, \end{aligned} \quad (13)$$

by following the definition of \mathcal{B} . Here, we apply the inequality (I) in the proof of Lemma 4 of Schmidt-Hieber (2020) with $\varepsilon = 1$ and apply (13) as

$$\begin{aligned} & \mathbb{E} \left[\|\hat{f}^{DL} - f^*\|_{L^2(P_X)}^2 \right] \\ & \leq 2 \left\{ \mathbb{E} \left[\|\hat{f}^{DL} - f^*\|_n^2 \right] + \frac{2F^2}{n} (12 \log \mathcal{N}(\delta, \mathcal{G}(L, S, B), \|\cdot\|_{L^\infty}) + 70) + 26\delta F \right\} \\ & \leq 4\mathbb{E} \left[\|\hat{g} - f^*\|_n^2 \right] + (52F + 4c_N)\delta \\ & \quad + \frac{24(\sigma^2 + 2F^2) \log \mathcal{N}(\delta, \mathcal{G}(L, S, B), \|\cdot\|_{L^\infty}) + 8 + 280F^2}{n}. \end{aligned}$$

Substituting the covering bound in Lemma 23 yields

$$\begin{aligned} & \mathbb{E} \left[\|\hat{f}^{DL} - f^*\|_{L^2(P_X)}^2 \right] \\ & \leq 4B_P \|\hat{g} - f^*\|_{L^2(I^D)}^2 + \frac{24(\sigma^2 + 2F^2)S}{n} \{(\log(nLB^L(S+1)^L) \vee 1)\}^2 \\ & \quad + \frac{8 + 2c_N + 26F + 280F^2}{n}, \end{aligned} \quad (14)$$

by setting $\delta = 1/(2n)$. Here, p_X is a density of P_X and $\sup_{x \in I^D} p_X(x) \leq B_P$ is finite by the assumption. About the last inequality, we apply the following

$$\mathbb{E} \left[\|\hat{g} - f^*\|_n^2 \right] = \int_{I^D} (\hat{g} - f^*)^2 d\lambda \frac{dP_X}{d\lambda} \leq \|\hat{g} - f^*\|_{L^2(I^D)}^2 \sup_{x \in I^D} p_X(x) \quad (15)$$

by the Hölder's inequality.

At last, we substitute the result of Theorem 6, and then adjust the coefficients. For ε in the statement of Theorem 6, we set

$$\varepsilon_1 = n^{-\beta/(2\beta+D)}, \text{ and } \varepsilon_2 = Mn^{-\alpha/2(\alpha+D-1)},$$

then, we rewrite the condition of Theorem 6 as

$$L \geq C_{\alpha, \beta, D, F} (1 + \lfloor \alpha \rfloor + \lfloor \beta \rfloor + \log_2 n),$$

$$S \geq C_{\alpha,\beta,D,F,J}(Mn^{D/(2\beta+D)} + n^{(D-1)/(\alpha+D-1)}) \log^2 n,$$

and $B \geq C_{F,M,q} n^{C_{\alpha,\beta,D}}$. Then, substitute them into (14) and obtain

$$\begin{aligned} & \mathbb{E} \left[\|\hat{f}^{DL} - f^*\|_{L^2(P_X)}^2 \right] \\ & \leq 4B_P(n^{-2\beta/(2\beta+D)} + n^{-\alpha/(\alpha+D-1)}) \\ & \quad + \frac{C_{\alpha,\beta,D,F,J}(\sigma^2 + F^2)}{n} (Mn^{D/(2\beta+D)} + Mn^{(D-1)/(\alpha+D-1)}) C_{\alpha,\beta,D,F} \log^2(n) + \frac{C_F}{n} \\ & \leq C_{\sigma,\alpha,\beta,D,F,J,P_X} (Mn^{-2\beta/(2\beta+D)} + Mn^{-\alpha/(\alpha+D-1)}) \log^2 n + \frac{C_F}{n}. \end{aligned}$$

Then, we adjust the coefficients and obtain the statement. \blacksquare

We provide a lemma which provides an upper bound for a covering number of $\mathcal{G}(L, S, B)$. Although similar results are well studied in several studies (Anthony and Bartlett, 2009; Schmidt-Hieber, 2020), we cite the following result in Nakada and Imaizumi (2020), which is more suitable for our result.

Lemma 23 (Covering Bound: Lemma 22 in Nakada and Imaizumi (2020)). *For any $\varepsilon > 0$, we have*

$$\log \mathcal{N}(\varepsilon, \mathcal{G}(L, S, B), \|\cdot\|_{L^\infty(I^D)}) \leq S \log \left(\frac{2LB^L(S+1)^L}{\varepsilon} \right).$$

Appendix C. Proof of Theorem 13

We first provide a proof of Proposition 12, and then prove Theorem 13 by applying Theorem 11.

Proof [Proof of Proposition 12] We give an upper bound and lower bound separately.

(i) **Upper bound:** First, we bound the packing number by a covering number as

$$\log \mathcal{M}(\varepsilon, H^\beta(I^D) \otimes \mathcal{I}_\alpha, \|\cdot\|_{L^2(I^D)}) \leq \log \mathcal{N}(\varepsilon/2, H^\beta(I^D) \otimes \mathcal{I}_\alpha, \|\cdot\|_{L^2(I^D)}),$$

by Section 2.2 in van der Vaart and Wellner (1996).

Further, we decompose the covering number of $H^\beta(I^D) \otimes \mathcal{I}_\alpha$. Let $\{f_j\}_{j=1}^{N_1} \subset H^\beta(I^D)$ be a set of centers of covering points with $N_1 = \mathcal{N}(\varepsilon/2, H^\beta(I^D), \|\cdot\|_{L^2(I^D)})$, and $\{\iota_j\}_{j=1}^{N_2} \subset \mathcal{I}_\alpha$ be centers of \mathcal{I}_α as $N_2 = \mathcal{N}(\varepsilon/2, \mathcal{I}_\alpha, \|\cdot\|_{L^2(I^D)})$. Then, we consider a set of points $\{f_j \otimes \iota_{j'}\}_{j,j'=1}^{N_1, N_2} \subset H^\beta(I^D) \otimes \mathcal{I}_\alpha$ whose cardinality is $N_1 N_2$. Then, for any element $f \otimes \mathbf{1}_R \in H^\beta(I^D) \otimes \mathcal{I}_\alpha$, there exists $f_j \in H^\beta(I^D)$ and $\iota_{j'}$ such as $\|f - f_j\|_{L^2(I^D)} \vee \|\mathbf{1}_R - \iota_{j'}\|_{L^2(I^D)} \leq \varepsilon/2$ due to the property of covering centers. Also, we can obtain

$$\begin{aligned} & \|f \otimes \mathbf{1}_R - f_j \otimes \iota_{j'}\|_{L^2(I^D)} \\ & \leq \|f \otimes (\mathbf{1}_R - \iota_{j'})\|_{L^2(I^D)} + \|(f - f_j) \otimes \iota_{j'}\|_{L^2(I^D)} \\ & \leq \|f\|_{L^\infty(I^D)} \|\mathbf{1}_R - \iota_{j'}\|_{L^2(I^D)} + \|f - f_j\|_{L^2(I^D)} \|\iota_{j'}\|_{L^\infty(I^D)} \\ & \leq \frac{(F+1)\varepsilon}{2}, \end{aligned}$$

where the second inequality follows the Hölder's inequality. By this result, we find that the set $\{f_j \otimes \iota_{j'}\}_{j,j'=1}^{N_1, N_2}$ is a $(F+1)\varepsilon/2$ covering set of $H^\beta(I^D) \otimes \mathcal{I}_\alpha$. Hence, we obtain

$$\begin{aligned} & \log \mathcal{M}((F+1)\varepsilon/2, H^\beta(I^D) \otimes \mathcal{I}_\alpha, \|\cdot\|_{L^2(I^D)}) \\ & \leq \log \mathcal{M}(\varepsilon/2, H^\beta(I^D), \|\cdot\|_{L^2(I^D)}) + \log \mathcal{M}(\varepsilon/2, \mathcal{I}_\alpha, \|\cdot\|_{L^2(I^D)}). \end{aligned} \quad (16)$$

We will bound the two entropy terms for $H^\beta(I^D)$ and \mathcal{I}_α . For $H^\beta(I^D)$, Theorem 8.4 in (Dudley, 2014) provides

$$\begin{aligned} \log \mathcal{N}(\varepsilon/2, H^\beta(I^D), \|\cdot\|_{L^2(I^D)}) & \leq \log \mathcal{N}(\varepsilon/2, H^\beta(I^D), \|\cdot\|_{L^\infty(I^D)}) \\ & = C_H(\varepsilon/2)^{-D/\beta}, \end{aligned} \quad (17)$$

with a constant $C_H > 0$. About the covering number of \mathcal{I}_α , we use the relation

$$\begin{aligned} \|\mathbf{1}_R - \mathbf{1}_{R'}\|_{L^2}^2 & = \int_{I^D} (\mathbf{1}_R(x) - \mathbf{1}_{R'}(x))^2 dx = \int |\mathbf{1}_R(x) - \mathbf{1}_{R'}(x)| dx \\ & = \int_{I^D} (\mathbf{1}_{R \cup R'}(x) - \mathbf{1}_{R' \cap R}(x)) dx =: d_1(R, R'), \end{aligned}$$

for basic sets $R, R' \subset I^D$, and d_1 is a difference distance with a Lebesgue measure for sets by Dudley (1974). Here, we consider a *boundary fragment class* $\tilde{\mathcal{R}}_\alpha$ defined by Dudley (1974), which is a set of subset of I^D with α -smooth boundaries. Since $R \subset I^D$ such that $\mathbf{1}_R \in \mathcal{I}_\alpha$ is a basis piece, we can see $R \in \tilde{\mathcal{R}}_\alpha$. Hence, we obtain

$$\begin{aligned} & \log \mathcal{M}(\varepsilon/2, \mathcal{I}_\alpha, \|\cdot\|_{L^2(I^D)}) \\ & = \log \mathcal{N}(\varepsilon^2/16, \{R \subset I^D \mid \mathbf{1}_R \in \mathcal{I}_\alpha\}, d_1) \\ & \leq \log \mathcal{N}(\varepsilon^2/16, \tilde{\mathcal{R}}_\alpha, d_1) \\ & = C_\lambda(\varepsilon/2)^{-2(D-1)/\alpha}, \end{aligned} \quad (18)$$

with a constant $C_\lambda > 0$. Here, the last equality follows Theorem 3.1 in (Dudley, 1974).

Combining (17) and (18) with (16), we obtain

$$\begin{aligned} & \log \mathcal{M}((F+1)\varepsilon/2, H^\beta(I^D) \otimes \mathcal{I}_\alpha, \|\cdot\|_{L^2(I^D)}) \\ & \leq C_H(\varepsilon/2)^{-D/\beta} + C_\lambda(\varepsilon/2)^{-2(D-1)/\alpha}. \end{aligned}$$

Adjusting the coefficients yields the upper bound.

(ii) Lower bound: We provide the lower bound by evaluating the packing number $\log \mathcal{M}(\varepsilon, H^\beta(I^D) \otimes \mathcal{I}_\alpha, \|\cdot\|_{L^2(I^D)})$ directory. Let $\mathbf{1}(x)$ be a constant function $\mathbf{1}(x) := 1, \forall x \in I^D$.

Here, we claim $\mathcal{I}_\alpha \subset H^\beta(I^D) \otimes \mathcal{I}_\alpha$, because $\mathbf{1} \in H^\beta(I^D)$ yields that $H^\beta(I^D) \otimes \mathcal{I}_\alpha \supset \{\mathbf{1}\} \otimes \mathcal{I}_\alpha = \mathcal{I}_\alpha$. Hence, we have

$$\begin{aligned} \log \mathcal{M}(\varepsilon, H^\beta(I^D) \otimes \mathcal{I}_\alpha, \|\cdot\|_{L^2(I^D)}) & \geq \log \mathcal{M}(\varepsilon, \mathcal{I}_\alpha, \|\cdot\|_{L^2(I^D)}) \\ & = C_\lambda \varepsilon^{-2(D-1)/\alpha}, \end{aligned} \quad (19)$$

by Theorem 3.1 in Dudley (1974). Similarly, $1 = \mathbf{1}_{I^D} \in \mathcal{I}_\alpha$ yields $H^\beta(I^D) \subset H^\beta(I^D) \otimes \mathcal{I}_\alpha$. Hence, we achieve

$$\begin{aligned} \log \mathcal{M}(\varepsilon, H^\beta(I^D) \otimes \mathcal{I}_\alpha, \|\cdot\|_{L^2(I^D)}) &\geq \log \mathcal{M}(\varepsilon, H^\beta(I^D), \|\cdot\|_{L^2(I^D)}) \\ &= C_H \varepsilon^{-D/\beta}, \end{aligned} \quad (20)$$

by Theorem 8.4 in Dudley (2014). Combining (19) and (20), we obtain

$$\log \mathcal{M}(\varepsilon, H^\beta(I^D) \otimes \mathcal{I}_\alpha, \|\cdot\|_{L^2(I^D)}) \geq \max\{C_H \varepsilon^{-D/\beta}, C_\lambda \varepsilon^{-2(D-1)/\alpha}\}.$$

Adjusting the coefficients, we obtain the statement. \blacksquare

Now, we are ready to prove Theorem 13.

Proof [Proof of Theorem 13] In this proof, we develop a lower bound of

$$\inf_{\bar{f}} \sup_{f^* \in H^\beta(I^D) \otimes \mathcal{I}_\alpha} \mathbb{E}_{f^*} \left[\|\bar{f} - f^*\|_{L^2(P_X)}^2 \right],$$

where \bar{f} is any estimator. Then the statement is immediate because of the following inequality

$$\sup_{f^* \in \mathcal{F}_{\alpha, \beta, M}^{PS}} \mathbb{E}_{f^*} \left[\|\bar{f} - f^*\|_{L^2(P_X)}^2 \right] \geq \sup_{f^* \in H^\beta(I^D) \otimes \mathcal{I}_\alpha} \mathbb{E}_{f^*} \left[\|\bar{f} - f^*\|_{L^2(P_X)}^2 \right] \quad (21)$$

due to $H^\beta(I^D) \otimes \mathcal{I}_\alpha \subset \mathcal{F}_{\alpha, \beta, M}^{PS}$.

To develop the lower bound with $H^\beta(I^D) \otimes \mathcal{I}_\alpha$, we apply Theorem 11. Let ε_n^* be a sequence for $n \in \mathbb{N}$ which satisfies

$$(\varepsilon_n^*)^2 = \log \mathcal{M}(\varepsilon_n^*, H^\beta(I^D) \otimes \mathcal{I}_\alpha, \|\cdot\|_{L^2(I^D)})/n.$$

From Proposition 12, we obtain

$$(\varepsilon_n^*)^2 = \Theta \left((\varepsilon_n^*)^{-D/\beta} + (\varepsilon_n^*)^{-2(D-1)/\alpha} \right) / n.$$

Solving this equation gives

$$(\varepsilon_n^*)^2 = \Theta \left(n^{-2\beta/(2\beta+D)} + n^{-\alpha/(\alpha+D-1)} \right). \quad (22)$$

Application of Theorem 11 with (22) derives

$$\begin{aligned} \inf_{\bar{f}} \sup_{f^* \in H^\beta(I^D) \otimes \mathcal{I}_\alpha} \mathbb{E}_{f^*} \left[\|\bar{f} - f^*\|_{L^2(P_X)}^2 \right] &= \Theta(\varepsilon_n^{*2}) \\ &= \Theta \left(n^{-2\beta/(2\beta+D)} + n^{-\alpha/(\alpha+D-1)} \right), \end{aligned}$$

which with (21) yields the claim. \blacksquare

Appendix D. Approximation Results of DNNs with General Admissible Activation

First, we show that the activation functions with Assumption 1 is suitable for DNNs to approximate polynomial functions, including an identity function.

Lemma 24. *Suppose η satisfies the condition (i) in Assumption 1. Then, for $\gamma \in \mathbb{N} \cup \{0\}$ with $\gamma \leq N + 1$, any $\varepsilon > 0$ and $T > 0$, there exists a tuple (L, S, B) such that*

- $L = 2$,
- $S = 3(\gamma + 1)$,
- $B = C_{\gamma, T} \varepsilon^{-C_{\gamma}}$,

and it satisfies

$$\inf_{g \in \mathcal{G}(L, S, B)} \|g - (x \mapsto x^\gamma)\|_{L^\infty([-T, T])} \leq \varepsilon.$$

Proof [Proof of Lemma 24] Consider the following neural network with one layer:

$$g_{sp}(x) := \sum_{j=1}^{\gamma+1} a_{2,j} \eta(a_{1,j}x + b_j). \quad (23)$$

Since η is N -times continuously differentiable by the condition (i) in Assumption 1, we set $b_j = x'$ for $j = 1, \dots, \gamma + 1$ and consider the Taylor expansion of η around $b_j = x'$. Then, for $j = 1, \dots, \gamma + 1$, we obtain

$$\eta(a_{1,j}x + b_j) = \sum_{k=0}^{\gamma} \frac{\partial^k \eta(x') a_{1,j}^k x^k}{k!} + \frac{\partial^{\gamma+1} \eta(\bar{x}) a_{1,j}^{\gamma+1} x^{\gamma+1}}{(\gamma + 1)!},$$

with some \bar{x} . We substitute it into (23) and obtain

$$\begin{aligned} & \sum_{j=1}^{\gamma+1} a_{2,j} \left(\sum_{k=0}^{\gamma} \frac{\partial^k \eta(x') a_{1,j}^k x^k}{k!} + \frac{\partial^{\gamma+1} \eta(\bar{x}) a_{1,j}^{\gamma+1} x^{\gamma+1}}{(\gamma + 1)!} \right) \\ &= \sum_{k=0}^{\gamma} \frac{\partial^k \eta(x') x^k}{k!} \sum_{j'=0}^{\gamma} a_{2,j'+1} a_{1,j'+1}^k + \sum_{j=1}^{\gamma+1} a_{2,j} \frac{\partial^{\gamma+1} \eta(\bar{x}) a_{1,j}^{\gamma+1} x^{\gamma+1}}{(\gamma + 1)!}. \end{aligned} \quad (24)$$

For each $j' = 0, \dots, \gamma$, we set $a_{1,j'+1} = j'/\bar{a}$ with $\bar{a} > 0$ and $a_{2,j'+1} = (-1)^{-\gamma+j'} \bar{a}^\gamma \binom{\gamma}{j'} / \partial^\gamma \eta(x')$. Note that $\partial^\gamma \eta(x') > 0$ holds by Assumption 1. Then, we obtain the following equality:

$$\begin{aligned} \sum_{j'=0}^{\gamma} a_{2,j'+1} a_{1,j'+1}^k &= \sum_{j'=0}^{\gamma} (-1)^{-\gamma+j'} \bar{a}^{\gamma-k} \binom{\gamma}{j'} \frac{j'^k}{\partial^\gamma \eta(x')} \\ &= \frac{\bar{a}^{\gamma-k}}{(-1)^\gamma \partial^\gamma \eta(x')} \sum_{j'=0}^{\gamma} (-1)^{j'} \binom{\gamma}{j'} j'^k \end{aligned}$$

$$= \begin{cases} \frac{\bar{a}^{\gamma-k}}{(-1)^\gamma \partial^\gamma \eta(x')} \gamma! (-1)^\gamma, & \text{if } k = \gamma \\ 0 & \text{otherwise,} \end{cases}$$

where the last equality follows the Stirling number of the second kind (described in (Graham et al., 1989)). Substituting it for (24) yields

$$\begin{aligned} \sum_{j=1}^{\gamma+1} a_{2,j} \eta(a_{1,j}x + b_j) &= x^\gamma + \sum_{j=1}^{\gamma+1} \frac{(-1)^{-\gamma+j-1} \partial^{\gamma+1} \eta(\bar{x})(j-1)^{\gamma+1} x^{\gamma+1}}{\partial^\gamma \eta(x') \bar{a}(\gamma+1)!} \binom{\gamma}{j-1} \\ &= x^\gamma + \frac{\partial^{\gamma+1} \eta(\bar{x}) x^{\gamma+1}}{\partial^\gamma \eta(x') \bar{a}(\gamma+1)!} \sum_{j=1}^{\gamma+1} (-1)^{-\gamma+j-1} (j-1)^{\gamma+1} \binom{\gamma}{j-1}. \end{aligned}$$

Regarding the second term, we obtain

$$\begin{aligned} &\left| \frac{\partial^{\gamma+1} \eta(\bar{x}) x^{\gamma+1}}{\partial^\gamma \eta(x') \bar{a}(\gamma+1)!} \sum_{j=1}^{\gamma+1} (-1)^{-\gamma+j-1} (j-1)^{\gamma+1} \binom{\gamma}{j-1} \right| \\ &\leq \frac{T^{\gamma+1}}{\bar{a} c_\eta} \|\partial^{\gamma+1} \eta\|_{L^\infty([-B,B])} \gamma^{\gamma+1} (1 + e^\gamma) =: \frac{C_{T,\gamma,\eta}}{\bar{a}}. \end{aligned}$$

As setting $\bar{a} = C_{T,\gamma,\eta}/\varepsilon$, we obtain the approximation with ε -error. From the result, we know that $g_{sp} \in \mathcal{G}(2, 3(\gamma+1), C_{T,\gamma}\varepsilon^{-C_\gamma})$. \blacksquare

Lemma 25. *Suppose η satisfies the condition (ii) in Assumption 1. Then, for $\gamma \in \mathbb{N} \cup \{0\}$ with $\gamma \leq N+1$, any $\varepsilon > 0$ and $T > 0$, there exists a tuple (L, S, B) such that*

- $L = (\gamma+1)(\log_2(C_{\gamma,T}/\varepsilon)/2 + 1)$,
- $S = C_{\gamma,T}((\log_2(1/\varepsilon))^2 + \log_2(1/\varepsilon))$,
- $B = C_{\gamma,T}\varepsilon^{-C_\gamma}$,

and it satisfies

$$\inf_{g \in \mathcal{G}(L,S,B)} \|g - (x \mapsto x^\gamma)\|_{L^\infty([-T,T])} \leq \varepsilon.$$

Proof [Proof of Lemma 25] As a preparation, we construct a *saw-tooth function* by Yarotsky (2017) with our Assumption 1. Let us define a *teeth function* $g_w : [0, 1] \rightarrow [0, 1]$ by a difference of two η as

$$\begin{aligned} g_w(x) &:= \frac{2c_2 + 2c_1}{c_2(c_2 - c_1)} \eta(x) - \frac{4}{c_2 - c_1} \eta(x - 1/2) - \frac{2\bar{c}}{c_2 - c_1} \\ &= \begin{cases} 2x, & \text{if } x \in [0, 1/2], \\ -2x + 2, & \text{if } x \in [1/2, 1]. \end{cases} \end{aligned}$$

Then, we consider the t -hold composition of $g_w \in \mathcal{G}(2, 6, c_w)$ with $c_w > 0$ as $g_t = g_w \circ \dots \circ g_w \in \mathcal{G}(t+1, 3(t+1), c_w)$ which satisfies

$$g_t(x) = \begin{cases} 2^t(x - 2k2^{-t}), & \text{if } x \in [2k2^{-t}, (2k+1)2^{-t}], k = 0, 1, \dots, 2^{t-1} - 1 \\ -2^t(x - 2k2^{-t}), & \text{if } x \in [(2k-1)2^{-t}, 2k2^{-t}], k = 1, \dots, 2^{t-1}. \end{cases}$$

Here, the domain $[0, 1]$ of g_t is divided into 2^{t+1} sub-intervals.

Then, we approximate a quadratic function by a linear sum of g_t . For $m \in \mathbb{N}$, We define a function $g_m : [0, 1] \rightarrow [0, 1]$ as

$$g_m(x) := x - \sum_{t=1}^m \frac{g_t(x)}{2^{2t}},$$

and it satisfies $\|g_m - (x \mapsto x^2)\|_{L^\infty([0,1])} \leq 2^{-2-2m}$ for $g_m \in \mathcal{G}(m+1, 3m^2/2 + 5m/2 + 1, c_m)$ with $c_m > 0$, by the similar way in Proposition 3 in (Yarotsky, 2017).

Further, we approximate a multiplicative function by f_m , intuitively, we represent a multiplication by a sum of a quadratic function as $xx' = \{|x+x'|^2 - |x|^2 - |x'|^2\}/2$. To the aim, we define an absolute function $g_a : [-1, 1] \rightarrow [0, 1]$ by DNNs as

$$g_a(x) := (\eta(x) - \eta(-x))/(c_2 + c_1) = |x|.$$

Then, we define $g_c : [-T, T] \times [-T, T] \rightarrow \mathbb{R}$ as

$$g_c(x, x') := \frac{T^2}{2} \{g_m \circ g_a((x+x')/T) - g_m \circ g_a(x/T) - g_m \circ g_a(x'/T)\}. \quad (25)$$

By the similar proof in Proposition 3 in Yarotsky (2017), we obtain $\|((x, x') \mapsto xx') - g_c(x, x')\|_{L^\infty([-T, T]^2)} \leq T^2 2^{-2m}$ with $g_c \in \mathcal{G}(m+1, (9m^2 + 15m)/2 + 10, c_c)$ with $c_c > 0$.

Finally, we approximate the polynomial function x^γ by an induction by g_c . When $\gamma = 1$, we consider

$$g_{p,1}(x) := (c_2 + c_1)^{-1}(\eta(x) - \eta(-x)) = x,$$

then obviously $\|g_{p,1} - (x \mapsto x)\|_{L^\infty([-T, T])} = 0$ holds with $g_{p,1} \in \mathcal{G}(2, 3, c_{p,1})$ with a constant $c_{p,1} > 0$.

Now, for the induction, assume that there exists a function by DNNs $g_{p,\gamma-1} \in \mathcal{G}(\gamma(m+1), c_{\gamma-1,T}((9m^2 + 15m)/2 + 10), c_p)$ with constants $c_p > 0$ and $C_{\gamma-1,T}$ depending $\gamma-1$ and T , and it satisfies $\|g_{p,\gamma-1} - (x \mapsto x^{\gamma-1})\|_{L^\infty([-T, T])} \leq C_{\gamma-1,T} 2^{-2m}$ with $C_{\gamma-1,T} > 0$. Also, we set g_m as (25). Then, we consider the following function by DNNs as

$$g_{p,\gamma}(x) := g_m(g_{p,\gamma-1}(x), g_{p,1}(x)).$$

Then, we consider the following approximation error

$$\begin{aligned} & \|g_{p,\gamma} - (x \mapsto x^\gamma)\|_{L^\infty([-T, T])} \\ &= \|g_m(g_{p,\gamma-1}(\cdot), g_{p,1}(\cdot)) - (x \mapsto x^{\gamma-1}) \otimes (x \mapsto x)\|_{L^\infty([-T, T])} \\ &\leq \|g_m(g_{p,\gamma-1}(\cdot), g_{p,1}(\cdot)) - g_{p,\gamma-1} \otimes g_{p,1}\|_{L^\infty([-T, T]^2)} \end{aligned}$$

$$\begin{aligned}
 & + \|g_{p,\gamma-1} \otimes g_{p,1} - (x \mapsto x^{\gamma-1}) \otimes g_{p,1}\|_{L^\infty([-T,T])} \\
 & + \|(x \mapsto x^{\gamma-1}) \otimes g_{p,1} - (x \mapsto x^{\gamma-1}) \otimes (x \mapsto x)\|_{L^\infty([-T,T])} \\
 & \leq 2T^2 2^{-2m} + 2T \|g_{p,\gamma-1} - (x \mapsto x^{\gamma-1})\|_{L^\infty([-T,T])} + 0 \\
 & \leq 2T^2 2^{-2m} + 2TC_{\gamma-1,T} 2^{-2m} \\
 & = (2T^2 + 2TC_{\gamma-1,T}) 2^{-2m} =: C_{\gamma,T} 2^{-2m}.
 \end{aligned}$$

Then, we obtain the statement with the condition (ii) in Assumption 1.

We combine the result with the conditions (i) and (ii), and obtain the statement with $\varepsilon = 2^{-2m}$. \blacksquare

Lemma 26 (General version of Lemma 4). *Suppose η satisfies Assumption 1. Then, for any $\varepsilon \in (0, 1)$ and $T > 0$, we obtain*

$$\inf_{g \in \mathcal{G}(2,6,C_T,q,\varepsilon^{-8 \vee (2q-1/2)})} \|g - \mathbf{1}_{\{\cdot \geq 0\}}\|_{L^2([-T,T])} \leq \varepsilon.$$

Proof [Proof of Lemma 26] First, we consider η that satisfies the condition (i) in Assumption 1. Without loss of generality, we set $\bar{c} = 1$ and $\underline{c} = 0$. We start with η with $k = 0$, where k is given in Assumption 1. Let us consider a shifted activation $\eta(ax)$ with $a > 0$. Then, the difference between $\eta(ax)$ and $\mathbf{1}_{\{\cdot \geq 0\}}$ is decomposed as

$$|\eta(ax) - \mathbf{1}_{\{\cdot \geq 0\}}| \leq \begin{cases} c_\eta(a^{-q}x^{-q} \wedge 1), & \text{if } x > 0, \\ c_\eta(a^{-q}|x|^{-q} \wedge 1), & \text{if } x < 0, \end{cases}$$

with an existing constant $c_\eta > 0$. The upper bound by 1 comes from the uniform bound by C_K in Assumption 1. Hence, the difference of them in terms of $L^2(\mathbb{R})$ norm is

$$\begin{aligned}
 \|\eta(a\cdot) - \mathbf{1}_{\{\cdot \geq 0\}}\|_{L^2(\mathbb{R})} & \leq \left(2c_\eta \int_0^\infty (a^{-2q}x^{-2q} \wedge 1) dx \right)^{1/2} \\
 & = \left(\frac{2c_\eta}{a} + \frac{2c_\eta}{(2q-1)a^{4q-1}} \right)^{1/2}.
 \end{aligned}$$

As we set $a \geq (4c_\eta/\varepsilon^2) \vee (4c_\eta(2q-1)/\varepsilon^2)^{4q-1}$, we obtain the statement with $\mathcal{G}(2, 1, C_q\varepsilon^{-(2q-1/2)})$.

We consider η with $k = 1$. With a scale parameters a and a shift parameters $\delta/2 > 0$, we consider a difference of two η with difference shift as

$$g_s(x; a, \delta) = \eta(ax - \delta/2) - \eta(ax + \delta/2).$$

Let $t \in (0, B)$ be a parameter for threshold. When $x > t$, the property of η yields

$$g_s(x; a, \delta) \in [\delta \pm 2c_\eta|ax - \delta/2|^{-q}].$$

We set $\delta = 1$, hence we obtain

$$|g_s(x; a, 1) - 1| \leq 2c_\eta|ax - 1/2|^{-q}, \tag{26}$$

for $x > t$. Similarly, when $x < -t$, we have

$$g_s(x; a, \delta) \in [0 \pm 2c_\eta |ax - \delta/2|^{-q}]. \quad (27)$$

When $x \in [-t, t]$, the bounded property of η yields

$$g_s(x; a, \delta) \in [0 \pm C_K(2 + |ax - \delta/2| + |ax + \delta/2|)]. \quad (28)$$

Combining the inequalities (26), (27) and (28) with $\delta = 1$, we bound the difference between $\eta(ax - \delta/2) - \eta(ax + \delta/2)$ and $\mathbf{1}_{\{\cdot \geq 0\}}$ as

$$\begin{aligned} & \|g_s(\cdot; a, 1) - \mathbf{1}_{\{\cdot \geq 0\}}\|_{L^2([-B, B])}^2 \\ & \leq \|g_s(\cdot; a, 1)\|_{L^2([-B, -t])}^2 + \|g_s(\cdot; a, 1) - \mathbf{1}_{\{\cdot \geq 0\}}\|_{L^2((t, B])}^2 \\ & \quad + \|g_s(\cdot; a, 1) - \mathbf{1}_{\{\cdot \geq 0\}}\|_{L^2([-t, t])}^2 \\ & \leq 4c_\eta^2 \int_{[-B, -t]} |ax - 1/2|^{-2q} dx + 4c_\eta^2 \int_{(t, B]} |ax - 1/2|^{-2q} dx \\ & \quad + C_K^2 \int_{[-t, t]} \{2 + |ax - 1/2| + |ax + 1/2|\}^2 dx \\ & =: T_{s,1} + T_{s,2} + T_{s,3}. \end{aligned}$$

About $T_{s,2}$, simple calculation yields

$$T_{s,2} \leq 4c_\eta^2 (B - t) |at - 1/2|^{-2q} \leq 4c_\eta^2 B |at - 1/2|^{-2q}.$$

By the symmetric property, we can also obtain $T_{s,1} = 4c_\eta^2 B (at - 1/2)^{-2q}$. About $T_{s,3}$, the similar calculation yields

$$T_{s,3} \leq 8C_K^2 t (1 + |at + 1/2|)^2 \leq 8C_K^2 t (at + 3/2)^2.$$

Combining the results of the terms, we bound the norm $\|g_s(\cdot; a, 1) - \mathbf{1}_{\{\cdot \geq 0\}}\|_{L^2([-B, B])}^2$. Here, we set $t = a^{-3/4}$ and $a \geq 1/16$, then we obtain

$$\begin{aligned} \|g_s(\cdot; a, 1) - \mathbf{1}_{\{\cdot \geq 0\}}\|_{L^2([-B, B])}^2 & = 4c_\eta^2 B a^{-3q/4} + 8C_K^2 a^{-3/4} (a^{1/4} + 3/2)^2 \\ & = 4c_\eta^2 B a^{-3q/4} + 8C_K^2 (a^{-1/4} + 3a^{-1/2} + 9a^{-3/4}/4). \end{aligned}$$

Then, we set $a = 16c_\eta^2 B \varepsilon^{-8/3q} \vee 32C_K^2 \varepsilon^{-8} \vee 96C_K^2 \varepsilon^{-4} \vee 9\varepsilon^{-8/3} \vee 1/16$, then we obtain the statement, then we have $g_s \in \mathcal{G}(2, 6, C_{B,q} \varepsilon^{-8})$.

Second, we consider η satisfies the condition (ii) in Assumption 1. We consider a sum of two $\eta(x) = c_1 x + (c_2 - c_1)x_+$ with some scale change as

$$\begin{aligned} \eta(x) + \eta\left(-\frac{c_1}{c_2}x\right) & = (c_1 x + (c_2 - c_1)x_+) + \left(-c_1 x + \left(c_1 - \frac{c_1^2}{c_2}\right)x_+\right) \\ & = \left(c_2 - \frac{c_1^2}{c_2}\right)x_+ =: \eta_s(x; c_1, c_2). \end{aligned}$$

Then, we consider a difference of two $\eta_s(x : c_1, c_2)$ with shift change by $\delta/2 > 0$ and scale change $a > 0$ as

$$\begin{aligned} \eta_d(x) &= \eta_s(ax - \delta/2 : c_1, c_2) - \eta_s(ax + \delta/2 : c_1, c_2) \\ &= \begin{cases} \left(c_2 - \frac{c_1^2}{c_2}\right) \delta, & \text{if } x > \delta/2a, \\ \left(c_2 - \frac{c_1^2}{c_2}\right) (ax - \delta/2), & \text{if } x \in [-\delta/2a, \delta/2a], \\ 0, & \text{if } x < -\delta/2a. \end{cases} \end{aligned}$$

Here, we set $\delta = (c_2 - c_1^2/c_2)^{-1} < (c_2 - c_1)^{-1} < \infty$. Then, the $L^2([-B, B])$ -distance between $g_d(x)$ and $\mathbf{1}_{\{\cdot \geq 0\}}$ is written as

$$\|g_d - \mathbf{1}_{\{\cdot \geq 0\}}\|_{L^2([-B, B])} = \left(\int_{[-\delta/2a, \delta/2a]} \left| \frac{ax}{\delta} \right|^2 dx \right)^{1/2} = \frac{\delta^{1/2}}{12^{1/2}a}.$$

As we set $a = \delta^{1/2}/(12^{1/2}\varepsilon)$, we obtain the statement with $g_d \in \mathcal{G}(2, 6, C_{B,q}\varepsilon^{-1})$.

We combine the results with all the conditions and consider the largest functional set, and then we obtain the statement. \blacksquare

Lemma 27 (General Version of Lemma 3). *Suppose Assumption 1 holds with $N \geq \alpha \vee \beta$. For any non-empty measurable set $R \subset I^D$ and $\delta > 0$, a tuple (L, S, B) such as*

- $L = C_{\beta, D, F, J}(|\beta| + \log_2(1/\delta) + 1)$
- $S = C_{\beta, D, F, J} \delta^{-D/\beta} (\log_2(1/\delta))^2$
- $B = C_{\beta, D, F, J, q} \delta^{-16 \wedge -C_\beta}$,

satisfies

$$\inf_{g \in \mathcal{G}(L, S, B)} \sup_{f \in H_F^\beta(I^D)} \|g - f\|_{L^2(R)} \leq \text{vol}(R)\delta.$$

Proof [Proof of Lemma 27]

Before the central part of the proof, we provide some preparation. We divide the domain I^D into several hypercubes. Let $\ell \in \mathbb{N}$ and consider a D -dimensional multi-index $\lambda \in \{1, 2, \dots, \ell\}^D =: \Lambda$. For each λ , define a hypercube

$$I_\lambda := \prod_{d=1}^D \left[\frac{\lambda_d - 1}{\ell}, \frac{\lambda_d}{\ell} \right]$$

By the definition, I_λ is a D -dimensional hypercube whose side has a length $1/\ell$, and $\bigcup_{\lambda \in \Lambda} I_\lambda = I^D$. Also, the center of I_λ is denoted by x_λ ; i.e.,

$$x_\lambda := \left(\frac{\lambda_1 - 1/2}{\ell}, \frac{\lambda_2 - 1/2}{\ell}, \dots, \frac{\lambda_D - 1/2}{\ell} \right) \in I^D$$

We provide a Taylor polynomial for a smooth function. Fix $\lambda \in \Lambda$ and $f \in H_F^\beta(I^D)$ arbitrary. Let $a \in \mathbb{N}^D$ be a multi-index. Then, we consider the Taylor expansion of f in I_λ as

$$\begin{aligned}
 f(x) &= f(x_\lambda) + \sum_{a:|a|\leq\lfloor\beta\rfloor-1} \frac{\partial^a f(x_\lambda)}{a!} (x - x_\lambda)^a \\
 &\quad + \sum_{a:|a|=\lfloor\beta\rfloor} \frac{1}{a!} (x - x_\lambda)^a \int_0^1 (1-t)^{|\beta|-1} \partial^a f(x_\lambda + t(x - x_\lambda)) dt \\
 &= f(x_\lambda) + \sum_{a:|a|\leq\lfloor\beta\rfloor-1} \frac{\partial^a f(x_\lambda)}{a!} (x - x_\lambda)^a \\
 &\quad + \sum_{a:|a|=\lfloor\beta\rfloor} \frac{1}{a!} (x - x_\lambda)^a \int_0^1 (1-t)^{|\beta|-1} \partial^a f(x_\lambda) dt \\
 &\quad + \sum_{a:|a|=\lfloor\beta\rfloor} \frac{1}{a!} (x - x_\lambda)^a \int_0^1 (1-t)^{|\beta|-1} \{\partial^a f(x_\lambda + t(x - x_\lambda)) - \partial^a f(x_\lambda)\} dt \\
 &= f(x_\lambda) + \sum_{a:|a|\leq\lfloor\beta\rfloor} \frac{\partial^a f(x_\lambda)}{a!} (x - x_\lambda)^a \\
 &\quad + \sum_{a:|a|=\lfloor\beta\rfloor} \frac{1}{a!} (x - x_\lambda)^a \int_0^1 (1-t)^{|\beta|-1} \{\partial^a f(x_\lambda + t(x - x_\lambda)) - \partial^a f(x_\lambda)\} dt \\
 &=: f(x_\lambda) + f_{\lfloor\beta\rfloor}(x; x_\lambda) + R_\lambda(x),
 \end{aligned}$$

where $a! = \prod_{d \in [D]} a_d!$, $f_{\lfloor\beta\rfloor-1}(x; x_\lambda)$ is the Taylor polynomial with an order $\lfloor\beta\rfloor$, and $R_\lambda(x)$ is the remainder. By the Hölder continuity and the bounded property of $\partial^a f(x)$, the remainder $R_\lambda(x)$ is bounded as

$$\begin{aligned}
 |R_\lambda(x)| &\leq \sum_{a:|a|=\lfloor\beta\rfloor} \frac{F}{a!} |(x - x_\lambda)^a| \int_0^1 (1-t)^{|\beta|-1} |t(x - x_\lambda)|^{\beta-|\beta|} dx \\
 &\leq F \sum_{a:|a|=\lfloor\beta\rfloor} (a!)^{-1} |x - x_\lambda|^\beta \leq C_{D,\beta,F} \left(\frac{1}{\ell}\right)^\beta,
 \end{aligned}$$

for $x \in I_\lambda$. Here, $C_{F,\beta} > 0$ is a constant which depends on F and β . The last inequality follows since $\|x - x_\lambda\|_\infty \leq 1/\ell$ holds for any $x \in I_\lambda$.

Now, we approximate the Taylor polynomial $f_{\lfloor\beta\rfloor}(x; x_\lambda)$ by DNNs for each $\lambda \in \Lambda$. By the binomial theorem, we can rewrite the Taylor polynomial with a multi-index $b \in \mathbb{N}^D$ and $x \in I_\lambda$ as

$$\begin{aligned}
 f_{\lfloor\beta\rfloor}(x; x_\lambda) &= \sum_{a:|a|\leq\lfloor\beta\rfloor} \frac{\partial^a f(x_\lambda)}{a!} \sum_{b \leq a} \binom{a}{b} (-x_\lambda)^{a-b} x^b \\
 &= \sum_{b:|b|\leq\lfloor\beta\rfloor} x^b \sum_{a:a \geq b, |a|\leq\lfloor\beta\rfloor} \frac{\partial^a f(x_\lambda)}{a!} \binom{a}{b} (-x_\lambda)^{a-b} x^b
 \end{aligned}$$

$$=: \sum_{b:|b|\leq|\beta|} x^b c_b.$$

Since $\|x_\lambda\|_\infty \leq 1$ and $\partial^a f(x_\lambda) \leq F$ by their definition, we obtain $|c_b| \leq F$ for any b . Then, for each $d \in [D]$, we define a univariate function by DNNs $g_{\lambda,d} \in \mathcal{G}(C_b(\log_2(2/\varepsilon) + 1), C_\beta(\log_2(1/\varepsilon))^2, C_b \varepsilon^{-C_b})$ which satisfies $\|(x \mapsto x^{b_d}) - g_{\lambda,d}\|_{L^\infty([-1/\ell, 1/\ell])} \leq \varepsilon$ with any $\varepsilon > 0$ by following Lemma 24 or Lemma 25. Also, we set $g_{c,D} \in \mathcal{G}(C_D \log_2(1/\varepsilon), C_D(\log_2(1/\varepsilon))^2, c_c)$ which approximate D -variate multiplication as Lemma 28 by substituting $m = \log_2(1/\varepsilon)/2$. With the functions, we consider the following difference

$$\begin{aligned} & \|(x \mapsto x^b) - g_{c,D}(g_{\lambda,1}(\cdot), \dots, g_{\lambda,D}(\cdot))\|_{L^\infty([-1/\ell, 1/\ell]^D)} \\ & \leq \|(x \mapsto x^b) - g_{\lambda,1}(\cdot) \otimes \dots \otimes g_{\lambda,D}(\cdot)\|_{L^\infty([-1/\ell, 1/\ell]^D)} \\ & \quad + \|g_{\lambda,1}(\cdot) \otimes \dots \otimes g_{\lambda,D}(\cdot) - g_{c,D}(g_{\lambda,1}(\cdot), \dots, g_{\lambda,D}(\cdot))\|_{L^\infty([-1/\ell, 1/\ell]^D)} \\ & \leq D\ell^{-2}\varepsilon + D\ell^{-2}\varepsilon. \end{aligned}$$

We then define a function by DNNs as $g_\lambda(x) := f(x_\lambda) + \sum_{b:|b|\leq|\beta|} c_b g_{c,D}(g_{\lambda,1}(x + x_{\lambda,1}), \dots, g_{\lambda,D}(x + x_{\lambda,D})) \in \mathcal{G}(C_{\beta,D}[\beta](\log_2(1/\varepsilon) + 1), C_{\beta,D}(\log_2(1/\varepsilon))^2, C_\beta \varepsilon^{-C_\beta})$ with $C_\beta > 0$, which satisfies

$$\begin{aligned} & \|f - g_\lambda\|_{L^\infty(I_\lambda)} \\ & \leq \sum_{b:|b|\leq|\beta|} c_b \|(x \mapsto x^b) - g_{c,D}(g_{\lambda,1}(\cdot + x_{\lambda,1}), \dots, g_{\lambda,D}(\cdot + x_{\lambda,D}))\|_{L^\infty(I_\lambda)} \\ & \quad + \|R_\lambda\|_{L^\infty(I_\lambda)} \\ & \leq 2C_\beta D\ell^{-2}\varepsilon + C_{D,\beta,F}\ell^{-\beta}, \end{aligned} \tag{29}$$

for any $\lambda \in \Lambda$.

Finally, we approximate $f \in H^\beta(I^D)$ on R . For a preparation, we define an approximator for the indicator function $\mathbf{1}_{I_\lambda}$ for $\lambda \in \Lambda$. From Lemma 26, we define $g_s \in \mathcal{G}(3, 11, C_{B,q}\varepsilon^{-8})$ be an approximator for a step function $\mathbf{1}_{\{\cdot \geq 0\}}$. Then, we define

$$g_I(x) := (g_s(x + 1/2) + g_s(-x - 1/2) - 1),$$

which satisfies

$$\begin{aligned} & \|g_I - \mathbf{1}_{\{-1/2 \leq \cdot \leq 1/2\}}\|_{L^2(\mathbb{R})} \\ & \leq \|(g_I(\cdot + 1/2) + g_s(-\cdot - 1/2) - 1) - (\mathbf{1}_{\{\cdot \geq -1/2\}} + \mathbf{1}_{\{\cdot \leq 1/2\}} - 1)\|_{L^2(\mathbb{R})} \\ & \quad + \|(\mathbf{1}_{\{\cdot \geq -1/2\}} + \mathbf{1}_{\{\cdot \leq 1/2\}} - 1) - \mathbf{1}_{\{-1/2 \leq \cdot \leq 1/2\}}\|_{L^2(\mathbb{R})} \\ & \leq \|g_s(\cdot + 1/2) - \mathbf{1}_{\{\cdot \geq -1/2\}}\|_{L^2(\mathbb{R})} + \|g_s(-\cdot - 1/2) + \mathbf{1}_{\{\cdot \leq 1/2\}}\|_{L^2(\mathbb{R})} + 0 \\ & \leq 2\varepsilon. \end{aligned}$$

Further, for $\lambda \in \Lambda$ and $x \in I^D$, we define $g_{I_\lambda} \in \mathcal{G}(4 + (\log_2(1/\varepsilon)/2 + 1)(D - 1), 25D + (\log_2(1/\varepsilon)/2)^2 D, C_{B,q}\varepsilon^{-8})$ as

$$g_{I_\lambda}(x) = g_{c,D}(g_I((x_1 - x_{\lambda,1})\ell), \dots, g_I((x_D - x_{\lambda,D})\ell)),$$

which is analogous to $\mathbf{1}_{I_\lambda}(x) = \prod_{d \in [D]} \mathbf{1}_{\{x_{\lambda,d-1}/(2\ell) \leq x_d \leq x_{\lambda,d+1}/(2\ell)\}}(x)$. We bound the distance as

$$\begin{aligned}
 & \|\mathbf{1}_{I_\lambda} - g_{I_\lambda}\|_{L^2(I^D)} \\
 & \leq \|\prod_{d=1}^D \mathbf{1}_{\{x_{\lambda,d-1}/(2\ell) \leq x_d \leq x_{\lambda,d+1}/(2\ell)\}} - \prod_{d=1}^D g_c((x_d - x_{\lambda,d})\ell)\|_{L^2(I^D)} \\
 & \quad + \|\prod_{d=1}^D g_c((x_d - x_{\lambda,d})\ell) - g_\times(g_c((x_1 - x_{\lambda,1})\ell), \dots, g_c((x_D - x_{\lambda,D})\ell))\|_{L^2(I^D)} \\
 & \leq \sum_{d=1}^D \|\mathbf{1}_{\{x_{\lambda,d-1}/(2\ell) \leq x_d \leq x_{\lambda,d+1}/(2\ell)\}} - g_c((x_d - x_{\lambda,d})\ell)\| \\
 & \quad \times \prod_{d' \neq d}^{d-1} \mathbf{1}_{\{x_{\lambda,d'-1}/(2\ell) \leq x_{d'} \leq x_{\lambda,d'+1}/(2\ell)\}} \vee g_c((x_{d'} - x_{\lambda,d'})\ell)\|_{L^2(I^D)} + D\varepsilon \\
 & \leq \sum_{d=1}^D \|\mathbf{1}_{\{x_{\lambda,d-1}/(2\ell) \leq x_d \leq x_{\lambda,d+1}/(2\ell)\}} - g_c((x_d - x_{\lambda,d})\ell)\|_{L^2(I^D)} + D\varepsilon \\
 & \leq D\varepsilon/\ell + D\varepsilon.
 \end{aligned} \tag{30}$$

Here, the second last inequality follows a bounded property of the indicator functions and g_c , and the Hölder's inequality.

Finally, we unify the approximator on a set $R \subset I^D$. Let us define $\Lambda_R := \{\lambda : \text{vol}(R \cap I_\lambda) \neq 0\}$, and

$$I_R := \bigcup_{\lambda \in \Lambda_R} I_\lambda.$$

Now, we can find a constant C_Λ such that we have

$$|\text{vol}(R) - \text{vol}(I_R)| \leq \bigcup_{\lambda \in \Lambda: \partial R \cap I_\lambda \neq \emptyset} \text{vol}(I_\lambda) \leq C_{\Lambda,J} \ell^{-1}, \tag{31}$$

where ∂R is a boundary of R . The second inequality holds since ∂R is a $(D-1)$ -dimensional set in the sense of the box counting dimension. Then, we define an approximator $g_f \in \mathcal{G}(C_{\beta,D}(\lfloor \beta \rfloor + \log_2(1/\varepsilon) + 1), C_{\beta,D} \ell^D ((\log_2(1/\varepsilon))^2 + 1), C_{B,q} \varepsilon^{-8 \wedge -C_\beta})$ as

$$g_f(x) := \sum_{\lambda \in \Lambda_R} g_c(g_\lambda(x), g_{I_\lambda}(x)).$$

Then, the error between g_f and f is decomposed as

$$\begin{aligned}
 & \|f - g_f\|_{L^2(R)} \\
 & \leq \|\sum_{\lambda \in \Lambda_R} \mathbf{1}_{I_\lambda} \otimes f - \sum_{\lambda \in \Lambda_R} g_\times(g_\lambda(\cdot), g_{I_\lambda}(\cdot))\|_{L^2(I_\lambda)} \\
 & \leq \sum_{\lambda \in \Lambda} \|\mathbf{1}_{I_\lambda} \otimes f - g_\lambda \otimes g_{I_\lambda}\|_{L^2(I^D)} + |\Lambda_R| 2\ell^{-1} \varepsilon \\
 & \leq \sum_{\lambda \in \Lambda} \|\mathbf{1}_{I_\lambda} \otimes (f - g_{I_\lambda})\|_{L^2(I^D)} + \|(\mathbf{1}_{I_\lambda} - g_{I_\lambda}) \otimes g_\lambda\|_{L^2(I^D)} + |\Lambda_R| 2\ell^{-1} \varepsilon \\
 & \leq \sum_{\lambda \in \Lambda} \|f - g_{I_\lambda}\|_{L^2(I_\lambda)} + \|g_\lambda\|_{L^\infty(I^D)} \|(\mathbf{1}_{I_\lambda} - g_{I_\lambda})\|_{L^2(I^D)} + |\Lambda_R| 2\ell^{-1} \varepsilon \\
 & \leq \text{vol}(I_R) (2C_\beta D \ell^{-2} \varepsilon + D\varepsilon/\ell + C_{D,\beta,F} \ell^{-\beta}) \\
 & \quad + |\Lambda_R| (F(D\varepsilon + 2D\varepsilon) + 2\ell^{-1} \varepsilon) \\
 & \leq \text{vol}(R) (2C_\beta D \ell^{-2} \varepsilon + D\varepsilon/\ell + C_{D,\beta,F} \ell^{-\beta}) (1 + C_{\Lambda,J} \ell^{-1}) \\
 & \quad + \ell^{-D} (3FD\varepsilon + 2\ell^{-1} \varepsilon),
 \end{aligned}$$

by the Hölder's inequality and the result in (29), (30) and (31). We set $\ell = \lceil \delta^{-1/\beta} \rceil$ and $\varepsilon = \delta^2$ with $\delta > 0$, hence we have

$$\begin{aligned} & \|f - g_f\|_{L^2(R)} \\ & \leq C_{\beta,D,F,J} \text{vol}(R) (\delta + \delta^{2+2/\beta} + \delta^{2+1/\beta}) + C_{D,F} (\delta^2 + \delta^{2+(D+1)/\beta}). \end{aligned}$$

Then, adjusting the coefficients as we can ignore the smaller order terms than $\Theta(\delta^2)$ as $\delta \rightarrow 0$, we obtain the statement. \blacksquare

Lemma 28. *Suppose Assumption 1 holds. Then, for any $m \in \mathbb{N}$, $B > 0$, and $D' \geq 2$, we obtain*

$$\inf_{g \in \mathcal{G}((m+1)(D'-1), h_c(m)(D'-1), c_m)} \|(x \mapsto \Pi_{d \in [D']} x_d) - g\|_{L^\infty([-B, B]^{D'})} \leq D' B^2 2^{-2m},$$

where $h_c(m) := \frac{9m^2 + 15m}{2} + 10$.

Proof [Proof of Lemma 28] Let $g_c \in \mathcal{G}(m+1, h_c(m), c_c)$ from the proof of Lemma 25. We prove it by induction. When $D' = 2$, the statement holds by the property of g_c . Then, consider $D' = \bar{D} - 1$ case with \bar{D} , and suppose that $g_{c, \bar{D}} \in \mathcal{G}((m+1)(\bar{D}-2), h_c(m)(\bar{D}-2), c_m)$ satisfies $\|(x \mapsto \Pi_{d \in [D']} x_d) - g\|_{L^\infty([-B, B]^{D'})} \leq (D' - 1) B^2 2^{-2m}$. Then, we define $g_{c, \bar{D}+1} \in \mathcal{G}((m+1)\bar{D}, h_c(m)\bar{D}, c_m)$ such as

$$g_{c, \bar{D}}(x_1, \dots, x_{\bar{D}}) = g_c(g_{c, \bar{D}-1}(x_1, \dots, x_{\bar{D}-1}), x_{\bar{D}}).$$

Then, we bound the difference as

$$\begin{aligned} & \|(x \mapsto \Pi_{d \in [\bar{D}]} x_d) - g_{c, \bar{D}}\|_{L^\infty([-B, B]^{\bar{D}})} \\ & \leq B^2 2^{-2m} + \|(x \mapsto \Pi_{d \in [\bar{D}-1]} x_d) - g_{c, \bar{D}-1}\|_{L^\infty([-B, B]^{\bar{D}-1})} \\ & \leq B^2 2^{-2m} + (\bar{D} - 1) B^2 2^{-2m} = \bar{D} B^2 2^{-2m}. \end{aligned}$$

Then, by the induction, we obtain the statement for any $D' \geq 2$. \blacksquare

Lemma 29 (General Version of Lemma 5). *Suppose Assumption 1 holds with $N > \alpha$. Then, for $\{R_m\}_{m \in [M]} \in \mathcal{R}_{\alpha, M}$ and any $\varepsilon > 0$ and $m' \in \mathbb{N}$, there exists $f \in \mathcal{G}(C_{\alpha, D, F, J}(\lfloor \alpha \rfloor + \log_2(1/\varepsilon) + 1), C_{\alpha, D, F, J}(\varepsilon^{-2(D-1)/\alpha}(\log_2(1/\varepsilon))^2 + M(\log_2(1/\varepsilon))^2 + 1), C_{F, J, q} \varepsilon^{-16 \wedge -C_\alpha})$ with a M -dimensional output $f(x) = (f_1(x), \dots, f_M(x))^\top$ such that*

$$\|\mathbf{1}_{R_m} - f_m\|_{L^2(I)} \leq \varepsilon,$$

and

$$\|f_m\|_{L^\infty(I^D)} \leq 1 + \varepsilon,$$

for all $m \in [M]$.

Proof [Proof of Lemma 29] We define a function by DNNs $g_{h,j} \in \mathcal{G}(C_{\alpha,D,F}([\alpha] + \log_2(1/\delta) + 1), C_{\alpha,D,F}\delta^{-(D-1)/\alpha}(\log_2(1/\delta))^2, C_{\alpha,D,F}\delta^{-16 \wedge -C_\alpha})$ such that $\|h_j - g_{h,j}\|_{L^2(I^{D-1})} \leq \delta$ by Lemma 27, for $\delta > 0$. Also, we define $g_{c,J} \in \mathcal{G}(J(\log_2(1/\delta)), C(\log_2(1/\delta))^2, C)$ with $m' \geq 1$ as Lemma 28, and $g_s \in \mathcal{G}(2, 6, C_{F,q}\delta^{-8})$ such that $\|g_s - \mathbf{1}_{\{\cdot \geq 0\}}\|_{L^2([-F,F])} \leq \delta$ as Lemma 26. Then, for $m \in [M]$, we define a function $g_{h,j,+}, g_{h,j,-} \in \mathcal{G}(C_{\alpha,D,F}([\alpha] + \log_2(1/\delta) + 1), C_{\alpha,D,F}(\delta^{-(D-1)/\alpha}(\log_2(1/\delta))^2 + 1), C_{B,q}\delta^{-16 \wedge -C_\alpha})$ to approximate $\mathbf{1}_{\{x_{d_j} \leq h_j(x_{-d_j})\}}$ as

$$g_{h,j,+}(x) = g_s(x_{d_j} - h_j(x_{-d_j})), \text{ and } g_{h,j,-}(x) = g_s(-x_{d_j} + h_j(x_{-d_j})).$$

To approximate $\mathbf{1}_{R_m}(x) = \prod_{j \in [J]} \mathbf{1}_{\{x_{d_j} \leq h_j(x_{-d_j})\}}$, we define $g_{R,m} \in \mathcal{G}(C_{\alpha,D,F}([\alpha] + \log_2(1/\delta) + 1), C_{\alpha,D,F}(J\delta^{-(D-1)/\alpha}(\log_2(1/\delta))^2 + \log_2(1/\delta) + 1), C_{F,q}\delta^{-16})$ as $g_{R,m} : [0, 1]^D \rightarrow [0, 1]$ such that

$$g_{R,m}(x) = g_{c,J}(g_{h,1,\pm}(x), \dots, g_{h,J,\pm}(x)),$$

and define $g_R \in \mathcal{G}(C_{\alpha,D,F}([\alpha] + \log_2(1/\delta) + 1), C_{\alpha,D,F}(J\delta^{-(D-1)/\alpha}(\log_2(1/\delta))^2 + M(\log_2(1/\delta))^2 + 1), C_{F,q}\delta^{-16 \wedge -C_\alpha})$ as $g_R : [0, 1]^D \rightarrow [0, 1]^M$ as

$$g_R(x) = (g_{R_1}(x), \dots, g_{R_M}(x))^\top.$$

Then, its approximation error is bounded as

$$\begin{aligned} & \|\mathbf{1}_{R_m} - g_{R,m}\|_{L^2(I^D)} \\ & \leq \left\| \left(x \mapsto \prod_{j \in [J]} \mathbf{1}_{\{x_{d_j} \leq h_j(x_{-d_j})\}} \right) - \prod_{j \in [J]} g_{h,j,\pm} \right\|_{L^2(I^D)} + J\delta \\ & \leq \sum_{j \in [J]} \left\| \left(x \mapsto \mathbf{1}_{\{x_{d_j} \leq h_j(x_{-d_j})\}} \right) - g_{h,j,\pm} \right\|_{L^2(I^D)} \\ & \quad \times \prod_{j' \in [j]} \left\| \left(x \mapsto \mathbf{1}_{\{x_{d_{j'}} \leq h_{j'}(x_{-d_{j'}})\}} \right) \right\|_{L^\infty(I^D)} \prod_{j'' \in [J] \setminus [j]} \|g_{h,j'',\pm}\|_{L^\infty(I^D)} + J\delta \\ & \leq \sum_{j \in [J]} \left\| \left(x \mapsto \mathbf{1}_{\{x_{d_j} \leq h_j(x_{-d_j})\}} \right) - g_{h,j,\pm} \right\|_{L^2(I^D)} + J\delta \\ & \leq \sum_{j \in [J]} \left\| \left(x \mapsto \mathbf{1}_{\{x_{d_j} \leq h_j(x_{-d_j})\}} \right) - \left(x \mapsto \mathbf{1}_{\{x_{d_j} \leq g_{h,j}(x_{-d_j})\}} \right) \right\|_{L^2(I^D)} \\ & \quad + \sum_{j \in [J]} \left\| \left(x \mapsto \mathbf{1}_{\{x_{d_j} \leq g_{h,j}(x_{-d_j})\}} \right) - g_{h,j,\pm} \right\|_{L^2(I^D)} + J\delta \\ & =: \sum_{j \in [J]} T_{h,1,j} + \sum_{j \in [J]} T_{h,2,j} + J\delta, \end{aligned}$$

where $\|g_{h,j,\pm}\|_{L^\infty(I^D)} \leq 1$ is used in the last inequality.

We evaluate each of the two terms $T_{h,1,j}$ and $T_{h,2,j}$. As preparation, for sets $\Omega, \Omega' \subset I^{D-1}$, we define $\Omega \Delta \Omega' := (\Omega \cup \Omega') \setminus (\Omega \cap \Omega')$. For each $j \in [J]$, we obtain

$$T_{h,1,j} = \left\| \left(x \mapsto \mathbf{1}_{\{x_{d_j} \leq h_j(x_{-d_j})\}} \right) - \left(x \mapsto \mathbf{1}_{\{x_{d_j} \leq g_{h,j}(x_{-d_j})\}} \right) \right\|_{L^2(I^D)}$$

$$\begin{aligned}
 &= \lambda (\{x \in I^D \mid x_{d_j} \lesseqgtr h_j(x_{-d_j})\} \Delta \{x \in I^D \mid x_{d_j} \lesseqgtr g_{h,j}(x_{-d_j})\})^{1/2} \\
 &= \|h_j - g_{h,j}\|_{L^1(I^{D-1})}^{1/2} \leq \|h_j - g_{h,j}\|_{L^2(I^{D-1})}^{1/2} \leq \delta^{1/2},
 \end{aligned}$$

where the second last inequality follows the Cauchy-Schwartz inequality. About $T_{h,2,j}$, we obtain

$$\begin{aligned}
 T_{h,2,j} &= \left\| (x \mapsto \mathbf{1}_{\{x_{d_j} \lesseqgtr g_{h,j}(x_{-d_j})\}}) - g_{h,j,\pm} \right\|_{L^2(I^D)} \\
 &= \left\| \mathbf{1}_{\{\cdot \geq 0\}} \circ (x \mapsto x_{d_j} - g_{h,j}(x_{-d_j})) - g_s \circ (x \mapsto x_{d_j} - g_{h,j}(x_{-d_j})) \right\|_{L^2(I^D)} \\
 &\leq C_F \left\| \mathbf{1}_{\{\cdot \geq 0\}} - g_s \right\|_{L^1([-F,F])} \leq C_F \left\| \mathbf{1}_{\{\cdot \geq 0\}} - g_s \right\|_{L^2([-F,F])} \leq C_F \delta,
 \end{aligned}$$

by the setting of g_s and the second last inequality follows the Cauchy-Schwartz inequality. Combining the results on $T_{h,1,j}$ and $T_{h,2,j}$, we obtain

$$\| \mathbf{1}_{R_m} - g_{R,m} \|_{L^2(I^D)} \leq J(\delta^{1/2} + C_F \delta + \delta).$$

For the second inequality of the statement, we apply the following inequality:

$$\begin{aligned}
 \|g_{R,m}\|_{L^\infty(I^D)} &\leq \|g_{R,m} - (x \mapsto \prod_{j \in [J]} \mathbf{1}_{\{x_{d_j} \lesseqgtr h_j(x_{-d_j})\}})\|_{L^\infty(I^D)} \\
 &\quad + \|(x \mapsto \prod_{j \in [J]} \mathbf{1}_{\{x_{d_j} \lesseqgtr h_j(x_{-d_j})\}})\|_{L^\infty(I^D)} \\
 &\leq \|g_{c,J} - (x \mapsto \prod_{j \in [J]} x_j)\|_{L^\infty([0,1]^J)} + 1 \\
 &\leq J\delta + 1,
 \end{aligned}$$

We set $\varepsilon = C_F J \delta^{1/2}$, we obtain the statement. ■

Appendix E. Proof of Proposition 15

The sub-optimality is well studied by Section 6 in Korostelev and Tsybakov (2012). We slightly adapt the result to our setting, and obtain the following proof.

Proof [Proof of Proposition 15] We divide this proof into the following five steps: (i) preparation, (ii) define a sub-class of functions, (iii) reparametrize a lower bound of errors, (iv) define a subset of parameters, and (v) combine all the results.

Step (i). Preparation. First, we decompose the distance $\|f^* - \hat{f}^{\text{lin}}\|_{L^2(P_X)}^2$. Let us define $\Upsilon_i(\cdot) := \Upsilon_i(\cdot; X_1, \dots, X_n)$. By the definition of linear estimators, we obtain

$$\begin{aligned}
 \|f^* - \hat{f}^{\text{lin}}\|_{L^2(P_X)}^2 &= \left\| f^* - \sum_{i=1}^n (f^*(X_i) + \xi_i) \Upsilon_i \right\|_{L^2(P_X)}^2 \\
 &= \left\| f^* - \sum_{i=1}^n f^*(X_i) \Upsilon_i \right\|_{L^2(P_X)}^2 + \left\| \sum_{i=1}^n \xi_i \Upsilon_i \right\|_{L^2(P_X)}^2 \\
 &\quad + 2 \left\langle f^* - \sum_{i=1}^n f^*(X_i) \Upsilon_i, \sum_{i=1}^n \xi_i \Upsilon_i \right\rangle_{L^2(P_X)}
 \end{aligned}$$

$$=: T_1^{(L)} + T_2^{(L)} + T_3^{(L)},$$

where $\langle f, f' \rangle_{L^2(P_X)} := \int f \otimes f' dP_X$ is an inner product with respect to P_X . Since ξ_i is a noise variable which is independent to X_i , we can simplify the expectations of the terms as

$$\mathbb{E}_{f^*}[T_3^{(L)}] = 2 \sum_{i=1}^m \mathbb{E}_{f^*}[\xi_i] \left\langle f^* - \sum_{i=1}^n f^*(X_i) \Upsilon_i, \Upsilon_i \right\rangle_{L^2(P_X)} = 0,$$

and

$$\mathbb{E}_{f^*}[T_2^{(L)}] = \sum_{i,i'=1}^n \mathbb{E}_{f^*}[\xi_i \xi_{i'}] \langle \Upsilon_i, \Upsilon_{i'} \rangle_{L^2(P_X)} = \sigma^2 \sum_{i=1}^n \|\Upsilon_i\|_{L^2(P_X)}^2.$$

Since $T_1^{(L)}$ is a deterministic term with fixed X_1, \dots, X_n , we obtain

$$\begin{aligned} \mathbb{E}_{f^*} \left[\|f^* - \hat{f}^{\text{lin}}\|_{L^2(P_X)}^2 \right] &= \left\| f^* - \sum_{i=1}^n f^*(X_i) \Upsilon_i \right\|_{L^2(P_X)}^2 + \sigma^2 \sum_{i=1}^n \|\Upsilon_i\|_{L^2(P_X)}^2 \\ &\geq \left\| f^* - \sum_{i=1}^n f^*(X_i) \Upsilon_i \right\|_{L^2(P_X)}^2 \vee \sigma^2 \sum_{i=1}^n \|\Upsilon_i\|_{L^2(P_X)}^2. \end{aligned} \quad (32)$$

Step (ii). Define a class of functions. We investigate a lower bound of the term $\sup_{f^* \in \mathcal{F}_{\alpha, \beta, M}^{PS}} \mathbb{E}_{f^*}[\|f^* - \sum_{i=1}^n f^*(X_i) \Upsilon_i\|^2]$ by considering an explicit class of piecewise smooth functions by dividing the domain I^D . For $m = 1, \dots, M-1$, we will consider a smooth boundary function $B_m : I^{D-1} \ni (x_1, \dots, x_{D-1}) \mapsto x_D \in I$, then define pieces

$$R_m = \begin{cases} \{x \in I^D \mid 0 \leq x_D < B_1(x_{-D})\}, & \text{if } m = 1, \\ \{x \in I^D \mid B_{m-1}(x_{-D}) \leq x_D < B_m(x_{-D})\}, & \text{if } m = 2, \dots, M-1, \\ \{x \in I^D \mid B_{m-1}(x_{-D}) \leq x_D \leq 1\}, & \text{if } m = M. \end{cases}$$

An explicit form of B_m is provided below. Let $N \in \mathbb{N}$ be a parameter, and consider a grid for I^{D-1} such that $q_j := ((j_d - 0.5)/N)_{d=1, \dots, D-1}$ for $j \in \{1, \dots, N\}^{D-1} =: \mathcal{J}$. We also define another index set $\mathcal{J}^+ := \{N+1, \dots, 2N\}^{D-1}$. Also, let $\phi \in H_1^\alpha(\mathbb{R}^{D-1})$ be a function such that $\phi(x) = 0$ for $x \notin I^{D-1}$, and $\phi(x) = 1$ for $x \in [0.1, 0.9]^{D-1}$. Then, we define a boundary function for $j \in \mathcal{J} \cup \mathcal{J}^+$ as

$$B_m(x_{-D}; j, r) = \begin{cases} \frac{m-1}{M} + \frac{r}{MN^\alpha} \phi(N(x_{-D} - q_j)), & \text{if } j \in \mathcal{J}, \\ \frac{m-1}{M}, & \text{if } j \in \mathcal{J}^+. \end{cases}$$

for $r \in \{r' \in \mathbb{N} \mid 0 < r' < N^\alpha - 1\} =: D_R$ and $m = 1, \dots, M-1$. Here, $\phi(N(x - q_j))$ is a smooth approximator for the indicator function of a hyper-cube with a center q_j , and B_m is constructed by the approximated indicator functions. Also, we define the following subset of I^D by B_m as

$$\hat{R}_{j,r,m} := \left\{ x \in I^D \mid \frac{m-1}{M} \leq x_D < B_m(x_{-D}; j, r) \right\},$$




Figure 9: Illustration of I^D with $D = 2$ case. The red curve denotes $B_m(x_{-D}; j_m, r)$. Also, \bar{R}_m (gray region) and $\hat{R}_{j_m, r, m}$ (red region) are illustrated.

for $m = 1, \dots, M - 1$. Moreover, we define

$$\bar{R}_m := \left\{ x \in I^D \mid \frac{m-1}{M} \leq x_D < \frac{m}{M} \right\}.$$

for $m = 1, \dots, M$. Obviously, $\hat{R}_{j, r, m} \subset \bar{R}_m$ for any j, r , and m . Figure 9 provides its illustration.

We provide a specific functional form characterized by $B_m(\cdot; j, r)$ with $r \in D_R$ and $j \in \mathcal{J} \cup \mathcal{J}^+$ for each $m = 1, \dots, M - 1$. Let $\mathbf{j} := (j_1, \dots, j_{M-1}) \in (\mathcal{J} \cup \mathcal{J}^+)^{\otimes M-1}$, and c_m be a fixed coefficient for $m = 1, \dots, M - 1$ such that $c_{m+1} - c_m = c > 0$ holds. For $r \in D_R$ and $\mathbf{j} \in (\mathcal{J} \cup \mathcal{J}^+)^{\otimes M-1}$, we define the following function

$$\check{f}(x; r, \mathbf{j}) := \sum_{m=1}^{M-1} c_m \mathbf{1}_{\hat{R}_{j_m, r, m}}(x).$$

$\check{f} \in \mathcal{F}_{\alpha, \beta, M}^{PS}$ holds by its construction.

Step (iii). Reparametrize a lower bound by the parameters of B_m . We develop a lower bound of the minimax risk by parameters r, j, m of the boundary function B_m , with fixed c_1, \dots, c_M and N . Now, we provide a lower bound of the minimax risk as

$$\begin{aligned} & \sup_{f^* \in \mathcal{F}_{\alpha, \beta, M}} \mathbb{E}_{f^*} \left[\|f^* - \hat{f}^{\text{lin}}\|_{L^2(P_X)}^2 \right] \\ &= \sup_{f^* \in \mathcal{F}_{\alpha, \beta, M}} \mathbb{E}_{f^*} \left[\left\| f^* - \sum_{i=1}^n f^*(X_i) \Upsilon_i \right\|_{L^2(P_X)}^2 \right] \\ &\geq \sup_{r \in D_R} \sup_{\mathbf{j} \in (\mathcal{J} \cup \mathcal{J}^+)^{\otimes M-1}} \mathbb{E}_{\check{f}} \left[\left\| \check{f} - \sum_{i=1}^n \check{f}(X_i; r, \mathbf{j}) \Upsilon_i \right\|_{L^2(P_X)}^2 \right] \end{aligned}$$

$$\begin{aligned}
 &\geq \frac{1}{|D_R|(|\mathcal{J}| + |\mathcal{J}^+|)^{M-1}} \\
 &\times \sum_{r \in D_R} \sum_{\mathbf{j} \in (\mathcal{J} \cup \mathcal{J}^+)^{\otimes M-1}} \mathbb{E}_{\check{f}} \left[\left\| \check{f}(\cdot; r, \mathbf{j}) - \sum_{i=1}^n \check{f}(X_i; r, \mathbf{j}) \Upsilon_i \right\|_{L^2(P_X)}^2 \right]. \\
 &\gtrsim \frac{1}{|D_R|(|\mathcal{J}| + |\mathcal{J}^+|)^{M-1}} \\
 &\times \sum_{r \in D_R} \sum_{\mathbf{j} \in (\mathcal{J} \cup \mathcal{J}^+)^{\otimes M-1}} \sum_{m=1}^{M-1} \sum_{j \in \mathcal{J}} \mathbb{E}_{\check{f}} \left[\left\| \check{f}(\cdot; r, \mathbf{j}) - \sum_{i=1}^n \check{f}(X_i; r, \mathbf{j}) \Upsilon_i \right\|_{L^2(\hat{R}_{j,r,m})}^2 \right]. \tag{33}
 \end{aligned}$$

To derive the second inequality, we consider all possible configurations. The last inequality holds since P_X has a finite and positive density by its definition. Also, for any $r \in D_R$, $\hat{R}_{j,r,m} \cap \hat{R}_{j',r,m} = \emptyset$ for $j \neq j' \in \mathcal{J}$, and $\cup_{j \in \mathcal{J}} \hat{R}_{j,r,m} \subset \bar{R}_m$ yield the last inequality.

Afterwards, we provide a lower bound of $\left\| \check{f}(\cdot; r, \mathbf{j}) - \sum_{i=1}^n \check{f}(X_i; r, \mathbf{j}) \Upsilon_i \right\|_{L^2(\hat{R}_{j,r,m})}^2$. For $m = 1, \dots, M-1$ and $r \in D_R$, we can achieve

$$\begin{aligned}
 &\sum_{j_m \in \mathcal{J} \cup \mathcal{J}^+} \left\| \check{f}(\cdot; r, \mathbf{j}) - \sum_{i=1}^n \check{f}(X_i; r, \mathbf{j}) \Upsilon_i \right\|_{L^2(\hat{R}_{j,r,m})}^2 \\
 &= \sum_{j_m \in \mathcal{J}} \left[\int_{\hat{R}_{j,r,m}} \left(c_m + c \mathbf{1}_{\{x_D \geq B_m(x_{-D}; j_m, r)\}}(x) \right. \right. \\
 &\quad \left. \left. - \sum_{i=1}^n (c_m + c \mathbf{1}_{\{x_D \geq B_m(x_{-D}; j_m, r)\}}(X_i)) \Upsilon_i(x) \right)^2 \right. \\
 &\quad \left. + \left(c_m + c \mathbf{1}_{\{x_D \geq \frac{m-1}{M}\}}(x) - \sum_{i=1}^n (c_m + c \mathbf{1}_{\{x_D \geq \frac{m-1}{M}\}}(X_i)) \Upsilon_i(x) \right)^2 dx \right] \\
 &\geq \sum_{j_m \in \mathcal{J}} \frac{c^2}{2} \int_{\hat{R}_{j,r,m}} \left(\sum_{i: X_i \in \hat{R}_{j_m, r, m}} \Upsilon_i(x) - \mathbf{1}_{\hat{R}_{j_m, r, m}}(x) \right)^2 dx \\
 &\gtrsim |\mathcal{J}| c^2 \left\| \sum_{i: X_i \in \hat{R}_{j,r,m}} \Upsilon_i - 1 \right\|_{L^2(\hat{R}_{j,r,m})}^2.
 \end{aligned}$$

The second last inequality follows $x^2 + y^2 \geq (x - y)^2/2$. Substituting it into (33) yield

$$\begin{aligned}
 &\sup_{f^* \in \mathcal{F}_{\alpha, \beta, M}} \mathbb{E}_{f^*} \left[\|f^* - \hat{f}^{\text{lin}}\|_{L^2(P_X)}^2 \right] \\
 &\gtrsim \frac{c^2}{|D_R| 2^{M-1}} \sum_{r \in D_R} \sum_{m=1}^{M-1} \sum_{j \in \mathcal{J}} \mathbb{E}_{f^*} \left[\left\| \sum_{i: X_i \in \hat{R}_{j,r,m}} \Upsilon_i - 1 \right\|_{L^2(\hat{R}_{j,r,m})}^2 \right]. \tag{34}
 \end{aligned}$$

Step (iv). Define subsets of parameters. Here, we will consider suitable subsets $\tilde{D}_R \subset D_R$ and $\tilde{\mathcal{J}} \subset \mathcal{J}$ for a tight lower bound. Let us define $n_m := |\{X_i \in \bar{R}_m\}|$ and

$\tau_m^2 := \sum_{i: X_i \in \tilde{R}_m} \|\Upsilon_i\|_{L^2(\tilde{R}_m)}^2$ for each $m \in [M]$. For $m \in [M]$, let $S_{m,r} := \{x \in I^D \mid x_D \in [m/M + r/N^\alpha, m/M + (r+1)/N^\alpha]\}$, for $r = 0, \dots, N^\alpha - 1$. Then, let \tilde{D}_R be a set of integers which satisfies

$$|\{i : X_i \in S_{m,r}\}| \leq \frac{cn_m}{N^\alpha}, \text{ and } \sum_{i: X_i \in S_{m,r}} \|\Upsilon_i\|_{L^2(I^D)} \leq \frac{c'\tau_m^2}{N^\alpha},$$

where $c, c' > 2$ are coefficients. We can claim that at least $(1 - 1/c)N^\alpha$ and $(1 - 1/c')N^\alpha$ integers from D_R satisfy each of the conditions, because the rest N^α/c and N^α/c' integers from D_R should sum to N^α . Then, at least $(1 - 1/c)N^\alpha + (1 - 1/c')N^\alpha - N^\alpha = (1 - 1/c - 1/c')N^\alpha$ integers from D_R satisfies the both conditions simultaneously. We set \tilde{D}_R as a set of such the integers, then $|\tilde{D}_R| \geq (1 - 1/c - 1/c')N^\alpha \gtrsim N^\alpha$ holds since $c, c' > 2$.

Further, for each $r \in \tilde{D}_R$, we consider the subset $\tilde{\mathcal{J}} \subset \mathcal{J}$. Similarly, we consider $\tilde{\mathcal{J}}$ is a set of indexes $j \in \mathcal{J}$ such as

$$|\{i : X_i \in \hat{R}_{j,r,m}\}| \leq \frac{c''n_m}{N^\alpha N^{D-1}}, \text{ and } \sum_{i: X_i \in \hat{R}_{j,r,m}} \|\Upsilon_i\|_{L^2(I^D)} \leq \frac{c'''\tau_m^2}{N^\alpha N^{D-1}},$$

with coefficients $c'', c''' > 2$. Repeating the argument for \tilde{D}_R , we can claim that there exist at least $(1 - 1/c'' - 1/c''')N^\alpha N^{D-1}$ indexes which satisfy the conditions simultaneously, then we have $\tilde{\mathcal{J}} \geq (1 - 1/c'' - 1/c''')N^\alpha N^{D-1} \gtrsim N^{\alpha+D-1}$ since $c'', c''' > 2$.

Step (v). Combining all the results. With the subsets \tilde{D}_R and $\tilde{\mathcal{J}}$, we derive a lower bound of the norm $\left\| \sum_{i: X_i \in \hat{R}_{j,r,m}} \Upsilon_i - 1 \right\|_{L^2(\hat{R}_{j,r,m})}$. For $m \in [M]$, $r \in \tilde{D}_R$, and $j \in \tilde{\mathcal{J}}$, we obtain

$$\begin{aligned} & \left\| \sum_{i: X_i \in \hat{R}_{j,r,m}} \Upsilon_i - 1 \right\|_{L^2(\hat{R}_{j,r,m})} \\ & \geq \|1\|_{L^2(\hat{R}_{j,r,m})} - \left\| \sum_{i: X_i \in \hat{R}_{j,r,m}} \Upsilon_i \right\|_{L^2(\hat{R}_{j,r,m})} \\ & = \text{vol}(\hat{R}_{j,r,m})^{1/2} - |\{i : X_i \in \hat{R}_{j,r,m}\}|^{1/2} \left(\int_{\hat{R}_{j,r,m}} \sum_{i: X_i \in \hat{R}_{j,r,m}} \Upsilon_i(x)^2 dx \right)^{1/2} \\ & \gtrsim \left(\frac{1}{MN^{D-1}N^\alpha} \right)^{1/2} - \frac{n_m^{1/2} \tau_m}{N^{D-1}N^\alpha}. \end{aligned}$$

Then, we substitute the result into (34). Here, note that $n_m \leq n$. Also, by (32), we have $\sum_{m \in [M]} \tau_m^2 \leq \sigma^{-2} \|f^* - \hat{f}^{\text{lin}}\|_{L^2(P_X)}^2$ and its integrable conditions, $\sum_{m \in [M]} \tau_m^2 = O(1)$ with probability at least $1/2$. Then, we continue the inequality as

$$\sup_{f^* \in \mathcal{F}_{\alpha, \beta, M}} \mathbb{E}_{f^*} \left[\|f^* - \hat{f}^{\text{lin}}\|_{L^2(P_X)}^2 \right]$$

$$\begin{aligned}
 &\gtrsim \frac{c^2}{|D_R|2^{M-1}} \sum_{m=2}^M \sum_{r \in \tilde{D}_R} \sum_{j \in \tilde{\mathcal{J}}} \mathbb{E}_{f^*} \left[\left\| \sum_{i: X_i \in \hat{R}_{j,r,m}} \Upsilon_i - 1 \right\|_{L^2(\hat{R}_{j,r,m'})}^2 \right] \\
 &\gtrsim N^{D-1} \sum_{m=2}^M \left\{ \left(\frac{1}{MN^{D-1}N^\alpha} \right)^{1/2} - \frac{n_m^{1/2} \tau_m}{N^{D-1}N^\alpha} \right\}^2.
 \end{aligned}$$

We substitute $N = \lceil n^{1/(2\alpha+D-1)} \rceil$ and ignore negligible terms, then obtain the statement. \blacksquare

Appendix F. Sub-optimality of Wavelet Estimators

Proof [Proof of Proposition 18] As preparation, we derive a lower bound of the minimax risk. Note that P_X is a uniform distribution on I^D , we have $\|f\|_{L^2(P_X)} = \|f\|_{L^2}$ for any measurable $f : I^D \rightarrow \mathbb{R}$. By the definition of \hat{f}^{wav} and the Parseval's equality, we obtain

$$\begin{aligned}
 \|f^* - \hat{f}^{\text{wav}}\|_{L^2(P_X)}^2 &= \sum_{(\kappa_1, \dots, \kappa_D) \in \mathcal{H}_\tau^{\times D}} (\hat{w}_{\kappa_1, \dots, \kappa_D} - w_{\kappa_1, \dots, \kappa_D}(f^*))^2 \\
 &\quad + \sum_{(\kappa_1, \dots, \kappa_D) \in \mathcal{H}^{\times D} \setminus \mathcal{H}_\tau^{\times D}} w_{\kappa_1, \dots, \kappa_D}(f^*)^2. \tag{35}
 \end{aligned}$$

For the first term in the right hand side, we evaluate the expectation of $(\hat{w}_{\kappa_1, \dots, \kappa_D} - w_{\kappa_1, \dots, \kappa_D}(f^*))^2$. Since $\mathbb{E}_{f^*}[\hat{w}_{\kappa_1, \dots, \kappa_D}] = \langle f^*, \Phi_{\kappa_1, \dots, \kappa_D} \rangle = w_{\kappa_1, \dots, \kappa_D}(f^*)$, we rewrite the expectation as

$$\begin{aligned}
 &\mathbb{E}_{f^*}[(\hat{w}_{\kappa_1, \dots, \kappa_D} - w_{\kappa_1, \dots, \kappa_D}(f^*))^2] \\
 &= \text{Var} \left(\frac{1}{n} \sum_{i \in [n]} Y_i \Phi_{\kappa_1, \dots, \kappa_D}(X_i) \right) \\
 &= \frac{1}{n} \text{Var}_{f^*} ((f^*(X_i) + \xi_i) \Phi_{\kappa_1, \dots, \kappa_D}(X_i)) \\
 &= \frac{1}{n} \mathbb{E}_X [\text{Var}_\xi ((f^*(X_i) + \xi_i) \Phi_{\kappa_1, \dots, \kappa_D}(X_i) \mid X_i)] \\
 &\geq \frac{\sigma^2}{n} \mathbb{E}_X [\Phi_{\kappa_1, \dots, \kappa_D}(X_i)^2] \\
 &= \frac{\sigma^2}{n},
 \end{aligned}$$

where the third equality follows the iterated law of expectation, and the last equality follows $\Phi_{\kappa_1, \dots, \kappa_D}$ is an orthonormal function. Substituting the result into (35) yields

$$\mathbb{E}_{f^*} \left[\|f^* - \hat{f}^{\text{wav}}\|_{L^2(P_X)}^2 \right] \geq \frac{\sigma^2 |\mathcal{H}_\tau|^D}{n} + \sum_{(\kappa_1, \dots, \kappa_D) \in \mathcal{H}^{\times D} \setminus \mathcal{H}_\tau^{\times D}} w_{\kappa_1, \dots, \kappa_D}(f^*)^2. \tag{36}$$

We will prove the statement by providing a specific configuration of f^* . Let $\dot{R} \subset I^D$ be a hyper-rectangle as

$$\dot{R} := \left\{ x \in I^D \mid 0 \leq x_d \leq \frac{2}{3}, d \in [D] \right\}.$$

Then, we define $f^* \in \mathcal{F}_{\alpha, \beta, M}^{PS}$ as

$$f^* := \mathbf{1}_{\dot{R}}.$$

Since \dot{R} is regarded as $\cap_{d \in [D]} \{x \in I^D \mid x_d \leq 2/3\}$, $\mathbf{1}_{\dot{R}}$ is a piecewise smooth function for any $M \geq 2, \alpha \geq 1$ and $\beta \geq 1$.

Then, we define the coefficient $w_{\kappa_1, \dots, \kappa_D}(f^*)$. Fix $j_d \in \{-1, 0, 1, 2, \dots\}$ for all $d \in [D]$. Since $\mathbf{1}_{\dot{R}}(x) = \prod_{d \in [D]} \mathbf{1}_{\{\cdot \leq 2/3\}}(x_d)$, we can decompose the coefficient as

$$\begin{aligned} \langle f^*, \Phi_{\kappa_1, \dots, \kappa_D} \rangle &= \int_{I^D} \prod_{d \in [D]} \mathbf{1}_{\{\cdot \leq 2/3\}}(x_d) \prod_{d \in [D]} \phi_{\kappa_d}(x_d) d(x_1, \dots, x_D) \\ &= \prod_{d \in [D]} \int_I \mathbf{1}_{\{\cdot \leq 2/3\}}(x_d) \phi_{\kappa_d}(x_d) dx_d. \end{aligned}$$

For each $d \in [D]$, a simple calculation yields

$$\int_I \mathbf{1}_{\{\cdot \leq 2/3\}}(x_d) \phi_{\kappa_d}(x_d) dx_d = \begin{cases} 2^{-j_d/2} c_{\kappa_d} & \text{if } 2/3 \in [k_d, k_d + 2^{-j_d}) \\ 0 & \text{if } 2/3 \notin [k_d, k_d + 2^{-j_d}), \end{cases}$$

with $c_{\kappa_d} > c > 0$ with a constant $c > 0$. When $2/3 \notin [k_d, k_d + 2^{-j_d})$, $\mathbf{1}_{\{\cdot \leq 2/3\}}(x_d)$ is a constant, then the integration is zero by the definition of ϕ_{κ_d} . Let k_d^* be a $k_d \in K_{j_d}$ such that $2/3 \in [k_d, k_d + 2^{-j_d})$. By the result of integration, we can rewrite $w_{\kappa_1, \dots, \kappa_D}$ as

$$w_{\kappa_1, \dots, \kappa_D} \geq \prod_{d \in [D]} c 2^{-j_d/2} \mathbf{1}_{\{k_d = k_d^*\}}.$$

Finally, we will select the truncation parameter τ for \mathcal{H}_τ and update the inequality (36). By its definition, we obtain $|\mathcal{H}_\tau| = \sum_{\ell=-1}^{\tau} 2^\ell = 2^{\tau+1}$. Also, about the second term of (36), we have

$$\begin{aligned} \sum_{(\kappa_1, \dots, \kappa_D) \in \mathcal{H}^{\times D} \setminus \mathcal{H}_\tau^{\times D}} w_{\kappa_1, \dots, \kappa_D}(f^*)^2 &\geq \sum_{j_1 > \tau} \sum_{j_2 > \tau} \dots \sum_{j_D > \tau} \prod_{d \in [D]} c^2 2^{-j_d} \\ &= c^{2D} \prod_{d \in [D]} \left(\sum_{j_d > \tau} 2^{-j_d} \right) \\ &= c^{2D} 2^{-\tau D}. \end{aligned}$$

Substituting the results into (36), we obtain

$$\mathbb{E}_{f^*} \left[\|f^* - \hat{f}^{\text{wav}}\|_{L^2(P_X)}^2 \right] \geq \frac{2\sigma^2 2^{\tau D}}{n} + c^{2D} 2^{-\tau D}. \quad (37)$$

By setting $\tau = \lceil (2D)^{-1} \log_2 n \rceil$ to minimize the right hand side of (37), we obtain the statement. \blacksquare

Appendix G. Sub-Optimality of Other Harmonic Estimators

Proof [Proof of Proposition 19]

In this proof, we provide an explicit example of $f^* \in \mathcal{F}_{\alpha,\beta,M}^{PS}$, and then derive a lower bound of a risk of \hat{f}^{curve} . Note that P_X is the uniform distribution on $[-1, 1]^2$. Also, we consider f^* to be the following non-smooth function

$$f^*(x_1, x_2) = \mathbf{1}_{\{x_1 \geq 0\}} \cdot \mathbf{1}_{\{x_2 \geq 0\}}, \quad (38)$$

and restrict it to $[-1, 1]^2$.

For analysis of curvelets, we consider a Fourier transformed curvelets $\hat{\gamma}_\mu$. As shown in (2.9) in Candès and Donoho (2004), with $\xi = (\xi_1, \xi_2)$, $\hat{\gamma}_\mu$ is written as

$$\hat{\gamma}_\mu(\xi) = 2\pi \cdot \chi_{j,\ell}(\xi) \cdot u_{j,k}(R_{\theta_j}^* \xi).$$

Here, $\chi_{j,\ell}$ is a polar symmetric window function

$$\chi_{j,\ell}(\xi) = \omega(2^{-2j} \|\xi\|_2) (\nu_{j,\ell}(\theta(\xi)) + \nu_{j,\ell}(\theta(\xi) + \pi)),$$

where $\theta(\xi) = \arcsin(\xi_2/\xi_1)$. Here, $\omega : \mathbb{R} \rightarrow \mathbb{R}$ is a compactly supported function, such as the Meyer wavelet, and we introduce $\nu_{j,\ell}(z) = \nu(2^j z - \pi\ell)$ where $\nu : \mathbb{R} \rightarrow \mathbb{R}$ is a function with a support $[-\pi, \pi]$ and satisfies $|\nu(\theta)|^2 + |\nu(\theta - \pi)|^2 = 1$ for $\theta \in [0, 2\pi)$. Without loss of generality, we assume that there exists a constant $c > 0$ such as $\text{vol}(\{z \mid \omega(z) \geq 0\})/2 \leq \text{vol}(\{z \mid \omega(z) \geq c\})$ and $\text{vol}(\{z \mid \nu(z) \geq 0\})/2 \leq \text{vol}(\{z \mid \nu(z) \geq c\})$, and the support of ω is $[1, 2]$. Also, $u_{j,k} : \mathbb{R}^2 \rightarrow \mathbb{R}$ is defined as

$$u_{j,k}(\xi) = \frac{2^{-3j/2}}{2\pi\sqrt{\delta_1\delta_2}} \exp(i(k_1 + 1/2)2^{-2j}\xi_1/\delta_1) \exp(ik_2 2^{-j}\xi_2/\delta_2),$$

and $\{u_{j,k}\}_k$ is an orthonormal basis for an L^2 -space on a rectangle which covers the support of $\chi_{j,\ell}$, with fixed j and ℓ .

Further, we provide a Fourier transform of f^* . Since a Fourier transform of $\mathbf{1}_{\{x \geq 0\}}$ is $\frac{1}{i\xi}$ and f^* a product of two step functions, its Fourier transform \hat{f}^* is written as $\frac{-1}{\xi_1\xi_2}$.

To obtain the statement of Proposition 19, we repeat the argument for (36), and obtain

$$\mathbb{E}_{f^*} \left[\|f^* - \hat{f}^{\text{curve}}\|_{L^2(P_X)}^2 \right] \geq \frac{\sigma^2 |\mathcal{L}_\tau|}{n} + \sum_{\mu \in \mathcal{L} \setminus \mathcal{L}_\tau} w_\gamma(f^*)^2. \quad (39)$$

Let us consider a partial sum of the coefficients $\{\gamma_\mu\}_\mu$. Here, fix j and ℓ , then consider a subset of indexes $\mathcal{L}_{j,\ell} := \{\mu \mid j = j', \ell = \ell'\}$. Then, since $\{u_{j,k}\}_k$ is an orthonormal basis, we obtain

$$\sum_{\mu \in \mathcal{L}_{j,\ell}} |w_\mu(f^*)|^2 = \int |\hat{f}^*(\xi)|^2 |\chi_{j,\ell}(\xi)|^2 d\xi.$$

Then, we utilize the form of \hat{f}^* and a compact support of $\chi_{j,\ell}$, hence obtain

$$\int |\hat{f}^*(\xi)|^2 |\chi_{j,\ell}(\xi)|^2 d\xi = \int \frac{1}{\xi_1^2 \xi_2^2} |\chi_{j,\ell}(\xi)|^2 d\xi$$

$$\geq \frac{c}{2} \text{vol}(\text{Supp}(\chi_{j,\ell})) \inf_{\xi \in \text{Supp}(\chi_{j,\ell})} \frac{1}{\xi_1^2 \xi_2^2}. \quad (40)$$

Since $\text{Supp}(\chi_{j,\ell})$ is a set

$$\{\xi \in \mathbb{R}^2 \mid 2^{2j} \leq \|\xi\|_2 \leq 2^{2j+1}, |\theta(\xi) - \pi \ell 2^{-j}| \leq \pi 2^{-j-1}\}.$$

Hence, simply we obtain

$$\text{vol}(\text{Supp}(\chi_{j,\ell})) = \frac{3\pi}{4} 2^{4j+2} = 3\pi 2^{4j}. \quad (41)$$

Also, about the infimum term, we obtain

$$\inf_{\mu \in \text{Supp}(\chi_{j,\ell})} \frac{1}{\xi_1^2 \xi_2^2} \geq \inf_{\xi: \|\xi\|_2 \leq 2^{2j+1}} \frac{1}{\xi_1^2 \xi_2^2} = \frac{1}{(2^{2j+1}/\sqrt{2})^4} = 2^{-8j-2}. \quad (42)$$

Substituting (41) and (42) into (40), then we obtain

$$\sum_{\mu \in \mathcal{L}_{j,\ell}} |w_\mu(f^*)|^2 = \int |\hat{f}^*(\xi)|^2 |\chi_{j,\ell}(\xi)|^2 d\xi \geq 3\pi 2^{-4j-2}. \quad (43)$$

Now, we will establish a lower bound by specifying \mathcal{L}_τ and associate it with (39). Let us define $\mathcal{L}_{j'} := \{\mu \mid j = j'\}$. By the setting of ℓ and k , we can obtain $|\mathcal{L}_j| = 2^j(1 + 2^{j/2}) = 2^j + 2^{3j/2}$. Also, with the truncation parameter τ , we have $|\mathcal{L}_\tau| = c_\tau(2^\tau + 2^{3\tau/2} - 1)$ with a coefficient $c_\tau > 0$.

For an approximation error, we consider

$$\begin{aligned} \sum_{\mu \in \mathcal{L} \setminus \mathcal{L}_\tau} w_\mu(f^*)^2 &= \sum_{j \in \mathbb{N}_0 \setminus \{[\tau] \cup \{0\}\}} \sum_{\ell} \sum_k w_{(j,\ell,k)}(f^*)^2 \\ &\geq \sum_{j \in \mathbb{N}_0 \setminus \{[\tau] \cup \{0\}\}} \sum_{\ell} 3\pi 2^{-4j-2} \\ &= \frac{3\pi}{4} \sum_{j \in \mathbb{N}_0 \setminus \{[\tau] \cup \{0\}\}} 2^{-3j} \\ &= \frac{3\pi}{28} 2^{-3\tau}. \end{aligned}$$

where the inequality follows (43) and the second equality follows $\ell = 0, 1, \dots, 2^j - 1$.

Combining the results with (39) by setting $\mathcal{L}_\tau = \mathcal{L}_\tau$, we obtain

$$\mathbb{E}_{f^*} \left[\|f^* - \hat{f}^{\text{curve}}\|_{L^2(P_X)}^2 \right] \geq \frac{\sigma^2 c_\tau (2^\tau + 2^{3\tau/2} - 1)}{n} + \frac{3\pi}{28} 2^{-3\tau}. \quad (44)$$

As we set τ as $2^{3\tau/2} = \Theta(n^{1/3})$, then obtain the statement. \blacksquare

References

- Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via over-parameterization. *Proceedings of Machine Learning Research (International Conference on Machine Learning)*, 97:242–252, 2019.
- Martin Anthony and Peter L Bartlett. *Neural network learning: Theoretical foundations*. cambridge university press, 2009.
- Peter L Bartlett. The sample complexity of pattern classification with neural networks: the size of the weights is more important than the size of the network. *IEEE transactions on Information Theory*, 44(2):525–536, 1998.
- 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.
- Emmanuel J Candes and David L Donoho. Recovering edges in ill-posed inverse problems: Optimality of curvelet frames. *Annals of statistics*, pages 784–842, 2002.
- Emmanuel J Candès and David L Donoho. New tight frames of curvelets and optimal representations of objects with piecewise c_2 singularities. *Communications on pure and applied mathematics*, 57(2):219–266, 2004.
- David L Donoho and Iain M Johnstone. Minimax estimation via wavelet shrinkage. *The Annals of Statistics*, 26(3):879–921, 1998.
- Richard M Dudley. Metric entropy of some classes of sets with differentiable boundaries. *Journal of Approximation Theory*, 10(3):227–236, 1974.
- Richard M Dudley. *Uniform central limit theorems*, volume 142. Cambridge university press, 2014.
- Evarist Giné and Richard Nickl. *Mathematical foundations of infinite-dimensional statistical models*, volume 40. Cambridge University Press, 2015.
- Ronald L Graham, Donald E Knuth, Oren Patashnik, and Stanley Liu. Concrete mathematics: a foundation for computer science. *Computers in Physics*, 3(5):106–107, 1989.
- Satoshi Hayakawa and Taiji Suzuki. On the minimax optimality and superiority of deep neural network learning over sparse parameter spaces. *Neural Networks*, 123:343–361, 2020.
- Geoffrey E Hinton, Simon Osindero, and Yee-Whye Teh. A fast learning algorithm for deep belief nets. *Neural computation*, 18(7):1527–1554, 2006.
- Masaaki Imaizumi and Kenji Fukumizu. Deep neural networks learn non-smooth functions effectively. *Proceedings of Machine Learning Research (International Conference on Artificial Intelligence and Statistics)*, 2019.

- Kenji Kawaguchi. Deep learning without poor local minima. In *Advances in Neural Information Processing Systems*, pages 586–594, 2016.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014.
- Vladimir Koltchinskii. Local rademacher complexities and oracle inequalities in risk minimization. *The Annals of Statistics*, 34(6):2593–2656, 2006.
- Aleksandr Petrovich Korostelev and Alexandre B Tsybakov. *Minimax theory of image reconstruction*, volume 82. Springer Science & Business Media, 2012.
- Gitta Kutyniok and Demetrio Labate. *Shearlets: Multiscale analysis for multivariate data*. Springer Science & Business Media, 2012.
- Gitta Kutyniok and Wang-Q Lim. Compactly supported shearlets are optimally sparse. *Journal of Approximation Theory*, 163(11):1564–1589, 2011.
- Quoc V Le, Jiquan Ngiam, Adam Coates, Abhik Lahiri, Bobby Prochnow, and Andrew Y Ng. On optimization methods for deep learning. In *Proceedings of Machine Learning Research (International Conference on Machine Learning)*, pages 265–272, 2011.
- Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.
- E Mammen and AB Tsybakov. Asymptotical minimax recovery of sets with smooth boundaries. *The Annals of Statistics*, 23(2):502–524, 1995.
- Enno Mammen and Alexandre B Tsybakov. Smooth discrimination analysis. *The Annals of Statistics*, 27(6):1808–1829, 1999.
- HN Mhaskar. Neural networks for optimal approximation of smooth and analytic functions. *Neural computation*, 8(1):164–177, 1996.
- Ryumei Nakada and Masaaki Imaizumi. Adaptive approximation and generalization of deep neural network with intrinsic dimensionality. *The Journal of Machine Learning Research*, 2020.
- Philipp Petersen and Felix Voigtlaender. Optimal approximation of piecewise smooth functions using deep relu neural networks. *Neural Networks*, 108:296–330, 2018.
- Jürgen Schmidhuber. Deep learning in neural networks: An overview. *Neural networks*, 61:85–117, 2015.
- Johannes Schmidt-Hieber. Nonparametric regression using deep neural networks with relu activation function. *Annals of Statistics*, 48(4):1875–1897, 2020.
- CJ Stone. Optimal global rates of convergence for nonparametric regression. *The Annals of Statistics*, 10:1040–1053, 1982.

Taiji Suzuki. Adaptivity of deep relu network for learning in besov and mixed smooth besov spaces: optimal rate and curse of dimensionality. *International Conference on Learning Representation*, 2019.

Alexandre B Tsybakov. Introduction to nonparametric estimation, 2009.

AW van der Vaart and Jon Wellner. *Weak Convergence and Empirical Processes: With Applications to Statistics*. Springer Science & Business Media, 1996.

Larry Alan Wasserman. *All of nonparametric statistics: with 52 illustrations*. Springer, 2006.

Yuhong Yang and Andrew Barron. Information-theoretic determination of minimax rates of convergence. *The Annals of Statistics*, 27(5):1564–1599, 1999.

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