
Deep learning is adaptive to intrinsic dimensionality of model smoothness in anisotropic Besov space

Taiji Suzuki

Department of Mathematical Informatics, The University of Tokyo, Tokyo, Japan
RIKEN Center for Advanced Intelligence Project, Tokyo, Japan
taiji@mist.i.u-tokyo.ac.jp

Atsushi Nitanda

Kyushu Institute of Technology, Fukuoka, Japan
RIKEN Center for Advanced Intelligence Project, Tokyo, Japan
nitanda@ai.kyutech.ac.jp

Abstract

Deep learning has exhibited superior performance for various tasks, especially for high-dimensional datasets, such as images. To understand this property, we investigate the approximation and estimation ability of deep learning on *anisotropic Besov spaces*. The anisotropic Besov space is characterized by direction-dependent smoothness and includes several function classes that have been investigated thus far. We demonstrate that the approximation error and estimation error of deep learning only depend on the average value of the smoothness parameters in all directions. Consequently, the curse of dimensionality can be avoided if the smoothness of the target function is highly anisotropic. Unlike existing studies, our analysis does not require a low-dimensional structure of the input data. We also investigate the minimax optimality of deep learning and compare its performance with that of the kernel method (more generally, linear estimators). The results show that deep learning has better dependence on the input dimensionality if the target function possesses anisotropic smoothness, and it achieves an adaptive rate for functions with spatially inhomogeneous smoothness.

1 Introduction

Based on the recent literature pertaining to machine learning, deep learning has exhibited superior performance in several tasks such as image recognition (Krizhevsky et al., 2012), natural language processing (Devlin et al., 2018), and image synthesis (Radford et al., 2015). In particular, its superiority is remarkable for complicated and high-dimensional data like images. This is mainly due to its high flexibility and superior feature-extraction ability for effectively extracting the intrinsic structure of data. Its theoretical analysis also has been extensively developed considering several aspects such as expressive ability, optimization, and generalization error.

Amongst representation ability analysis of deep neural networks such as universal approximation ability (Cybenko, 1989; Hornik, 1991; Sonoda & Murata, 2017), approximation theory of deep neural networks on typical function classes such as Hölder, Sobolev, and Besov spaces have been extensively studied. In particular, analyses of deep neural networks with the ReLU activation (Nair & Hinton, 2010; Glorot et al., 2011) have been recently developed. Schmidt-Hieber (2020) showed that the deep learning with ReLU activations can achieve the minimax optimal estimation accuracy to estimate composite functions in Hölder spaces by using the approximation theory of Yarotsky (2017). Suzuki (2019) generalized this analysis to those on the *Besov space* and the *mixed smooth*

Table 1: Relationship between existing research and our work. β indicates the smoothness of the target function, d is the dimensionality of input x , D is the dimensionality of a low-dimensional structure on which the data are distributed, and $\tilde{\beta}$ is the average smoothness of an anisotropic Besov space (Eq. (1)).

Function class	Hölder	Besov	mixed smooth Besov	Hölder on a low-dimensional set	anisotropic Besov
Author	Schmidt-Hieber (2020)	Suzuki (2019)	Suzuki (2019)	Nakada & Imaizumi (2020); Schmidt-Hieber (2019); Chen et al. (2019)	This work
Estimation error	$\tilde{O}(n^{-\frac{2\beta}{2\beta+d}})$	$\tilde{O}(n^{-\frac{2\beta}{2\beta+d}})$	$\tilde{O}\left(n^{-\frac{2\beta}{2\beta+1}} \log(n)^{\frac{2(d-1)(u+\beta)}{1+2\beta}}\right)$	$\tilde{O}(n^{-\frac{2\beta}{2\beta+D}})$	$\tilde{O}(n^{-\frac{2\beta}{2\beta+1}})$

Besov space by utilizing the techniques developed in approximation theories (Temlyakov, 1993; DeVore, 1998). It was shown that deep learning can achieve an *adaptive approximation* error rate that is faster than that of (non-adaptive) linear approximation methods (DeVore & Popov, 1988; DeVore et al., 1993; Düng, 2011), and it outperforms any linear estimators (including kernel ridge regression) in terms of the minimax optimal rate.

From these analyses, one can see that the approximation errors and estimation errors are strongly influenced by two factors, i.e., the *smoothness* of the target function and the *dimensionality* of the input (see Table 1). In particular, they suffer from the *curse of dimensionality*, which is unavoidable. However, these analyses are about the worst case errors and do not exploit specific intrinsic properties of the true distributions. For example, practically encountered data usually possess low intrinsic dimensionality, i.e., data are distributed on a low dimensional sub-manifold of the input space (Tenenbaum et al., 2000; Belkin & Niyogi, 2003). Recently, Nakada & Imaizumi (2020); Schmidt-Hieber (2019); Chen et al. (2019); Chen et al. (2019) have shown that deep ReLU network has adaptivity to the intrinsic dimensionality of data and can avoid curse of dimensionality if the intrinsic dimensionality is small. However, one drawback is that they assumed *exact* low dimensionality of the input data. This could be a strong assumption because practically observed data are always noisy, and injecting noise immediately destroys the low-dimensional structure. Therefore, we consider another direction in this paper. In terms of curse of dimensionality, Suzuki (2019) showed that deep learning can alleviate the curse of dimensionality to estimate functions in a so called mixed smooth Besov space (m-Besov). However, m-Besov space assumes strong smoothness toward *all* directions uniformly and does not include the ordinary Besov space as a special case. Moreover, the convergence rate includes heavy poly-log term which is not negligible (see Table 1).

In practice, one of the typically expected properties of a true function on high-dimensional data is that it is invariant against perturbations of an input in some specific directions (Figure 1). For example, in image-recognition tasks, the target function must be invariant against the spatial shift of an input image, which is utilized by data-augmentation techniques (Simard et al., 2003; Krizhevsky et al., 2012). In this paper, we investigate the approximation and estimation abilities of deep learning on *anisotropic Besov spaces* (Nikol’skii, 1975; Vybiral, 2006; Triebel, 2011) (also called dominated mixed-smooth Besov spaces). An anisotropic Besov space is a set of functions that have “direction-dependent” smoothness, whereas ordinary function spaces such as Hölder, Sobolev, and Besov spaces assume isotropic smoothness that is uniform in all directions. We consider a composition of functions included in an anisotropic Besov space, including several existing settings as special cases; it includes analyses of the Hölder space Schmidt-Hieber (2020) and Besov space Suzuki (2019), as well as the low-dimensional sub-manifold setting (Nakada & Imaizumi, 2020; Schmidt-Hieber, 2019; Chen et al., 2019; Chen et al., 2019)¹. By considering such a space, we can show that deep learning can alleviate curse of dimensionality if the smoothness in each direction is highly anisotropic. Interestingly, any linear estimator (including kernel ridge regression) has worse dependence on the dimensionality than deep learning. Our contributions can be summarized as follows:

- We consider a situation in which the target function is included in a class of anisotropic Besov spaces and show that deep learning can avoid the curse of dimensionality *even if the input data*

¹We would like to remark that the analysis of Nakada & Imaizumi (2020) does not require smoothness of the embedded manifold that is not covered in this paper.

do not lie on a low-dimensional manifold. Moreover, deep learning can achieve the optimal adaptive approximation error rate and minimax optimal estimation error rate.

- We compare deep learning with general linear estimators (including kernel methods) and show that deep learning has better dependence on the input dimensionality than linear estimators.

2 Problem setting and the model

In this section, we describe the problem setting considered in this work. We consider the following nonparametric regression model:

$$y_i = f^\circ(x_i) + \xi_i \quad (i = 1, \dots, n),$$

where x_i is generated from a probability distribution P_X on $[0, 1]^d$, $\xi_i \sim N(0, \sigma^2)$, and the data $D_n = (x_i, y_i)_{i=1}^n$ are independently identically distributed. f° is the true function that we want to estimate. We are interested in the mean squared estimation error of an estimator \hat{f} : $E_{D_n}[k\hat{f} - f^\circ]_{L^2(P_X)}^2$, where $E_{D_n}[\cdot]$ indicates the expectation with respect to the training data D_n .

We consider a least-squares estimator in the deep neural network model as \hat{f} (see Eq. (5)) and discuss its optimality. More specifically, we investigate how the ‘‘intrinsic dimensionality’’ of data affects the estimation accuracy of deep learning. For this purpose, we consider an *anisotropic Besov space* as a model of the target function.

2.1 Anisotropic Besov space

In this section, we introduce the anisotropic Besov which was investigated as the model of the true function in this paper. Throughout this paper, we set the domain of the input to $\Omega = [0, 1]^d$. For a function $f : \Omega \rightarrow \mathbb{R}$, let $kfk_p := kfk_{L^p(\Omega)} := (\int_\Omega |f|^p dx)^{1/p}$ for $0 < p < \infty$. For $p = \infty$, we define $kfk_\infty := kfk_{L^\infty(\Omega)} := \sup_{x \in \Omega} |f(x)|$. For $\beta \in \mathbb{R}_{++}^d$, let $j\beta j = \sum_{j=1}^d j\beta_j^2$.

For a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, we define the r th difference of f in the direction $h \in \mathbb{R}^d$ as

$$\Delta_h^r(f)(x) := \Delta_h^{r-1}(f)(x+h) - \Delta_h^{r-1}(f)(x), \quad \Delta_h^0(f)(x) := f(x),$$

for $x \in \Omega$ with $x + rh \in \Omega$, otherwise, let $\Delta_h^r(f)(x) = 0$.

Definition 1. For a function $f \in L^p(\Omega)$ where $p \in (0, \infty]$, the r -th modulus of smoothness of f is defined by $w_{r,p}(f, t) = \sup_{h \in \mathbb{R}^d: |h| \leq t} k\Delta_h^r(f)k_p$, for $t = (t_1, \dots, t_d)$, $t_i > 0$.

With this modulus of smoothness, we define the anisotropic Besov space $B_{p,q}^\beta(\Omega)$ for $\beta = (\beta_1, \dots, \beta_d) \in \mathbb{R}_{++}^d$ as follows.

Definition 2 (Anisotropic Besov space $(B_{p,q}^\beta(\Omega))$). For $0 < p, q \leq \infty$, $\beta = (\beta_1, \dots, \beta_d) \in \mathbb{R}_{++}^d$, $r := \max_i \beta_i \leq 1$, let the seminorm $j\cdot j_{B_{p,q}^\beta}$ be

$$j\cdot j_{B_{p,q}^\beta} := \begin{cases} \left(\sum_{k=0}^{\lfloor 1/t \rfloor} [2^k w_{r,p}(f, (2^{-k/\beta_1}, \dots, 2^{-k/\beta_d}))]^q \right)^{1/q} & (q < \infty), \\ \sup_k \sum_{k=0}^{\lfloor 1/t \rfloor} 2^k w_{r,p}(f, (2^{-k/\beta_1}, \dots, 2^{-k/\beta_d})) & (q = \infty). \end{cases}$$

The norm of the anisotropic Besov space $B_{p,q}^\beta(\Omega)$ is defined by $kfk_{B_{p,q}^\beta} := kfk_p + j\cdot j_{B_{p,q}^\beta}$, and $B_{p,q}^\beta(\Omega) = \{f \in L^p(\Omega) \mid kfk_{B_{p,q}^\beta} < \infty\}$.

Roughly speaking β represents the smoothness in each direction. If β_i is large, then a function in $B_{p,q}^\beta$ is smooth to the i th coordinate direction, otherwise, it is non-smooth to that direction. p is also an important quantity that controls the *spatial inhomogeneity* of the smoothness. If $\beta_1 = \beta_2 = \dots = \beta_d$, then the definition is equivalent to the usual Besov space (DeVore & Popov, 1988; DeVore et al., 1993). Suzuki (2019) analyzed curse of dimensionality of deep learning through a so-called *mixed smooth Besov* (m-Besov) space which imposes a stronger condition toward all directions uniformly.

²We let $\mathbb{N} := \{1, 2, 3, \dots\}$, $\mathbb{Z}_+ := \{0, 1, 2, 3, \dots\}$, $\mathbb{Z}_+^d := \{(z_1, \dots, z_d) \mid z_i \in \mathbb{Z}_+\}$, $\mathbb{R}_+ := \{x \in \mathbb{R} \mid x > 0\}$, and $\mathbb{R}_{++} := \{x \in \mathbb{R} \mid x > 0\}$. We let $[N] := \{1, \dots, N\}$ for $N \in \mathbb{N}$.

Particularly, it imposes stronger smoothness toward non-coordinate axis directions. Moreover, m -Besov space does *not* include the vanilla Besov space as a special case and thus cannot capture the situation that we consider in this paper.

Throughout this paper, for given $\beta = (\beta_1, \dots, \beta_d) \in \mathbb{R}_{++}^d$, we write $\underline{\beta} := \min_i \beta_i$ (smallest smoothness) and $\bar{\beta} := \max_i \beta_i$ (largest smoothness). The approximation error of a function in anisotropic Besov spaces is characterized by the harmonic mean of $(\beta_j)_{j=1}^d$, which corresponds to the average smoothness, and thus we define

$$\tilde{\beta} := \left(\sum_{j=1}^d 1/\beta_j \right)^{-1}. \quad (1)$$

The Besov space is closely related to other function spaces such as Hölder space. Let $\partial^\alpha f(x) = \frac{\partial^{|\alpha|} f}{\partial x_1^{\alpha_1} \dots \partial x_d^{\alpha_d}}(x)$.

Definition 3 (Hölder space $(C^\beta(\Omega))$). For a smoothness parameter $\beta \in \mathbb{R}_{++}$ with $\beta \notin \mathbb{N}$, consider an m times differentiable function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ where $m = \lfloor \beta \rfloor$ (the largest integer less than β), and let the norm of the Hölder space $C^\beta(\Omega)$ be $\|f\|_{C^\beta} := \max_{|\alpha| \leq m} \|\partial^\alpha f\| + \max_{|\alpha|=m} \sup_{x,y \in \Omega} \frac{|\partial^\alpha f(x) - \partial^\alpha f(y)|}{|x-y|^{m-\beta}}$. Then, (β) -Hölder space $C^\beta(\Omega)$ is defined as $C^\beta(\Omega) = \{f \mid \|f\|_{C^\beta} < \infty\}$.

Let $C^0(\Omega)$ be the set of continuous functions equipped with L^1 -norm: $C^0(\Omega) := \{f : \Omega \rightarrow \mathbb{R} \mid f \text{ is continuous and } \|f\|_1 < \infty\}$. These function spaces are closely related to each other.

Proposition 1 (Triebel (2011)). There exist the following relations between the spaces:

1. For $\beta = (\beta_0, \dots, \beta_0) \in \mathbb{R}^d$ with $\beta_0 \notin \mathbb{N}$, it holds that $C^{\beta_0}(\Omega) = B_{1,1}^{\beta_0}(\Omega)$.
2. For $0 < p_1, p_2, q < \infty$, $p_1 \leq p_2$ and $\beta \in \mathbb{R}_{++}^d$ with $\tilde{\beta} > (1/p_1 - 1/p_2)_+$ ³, it holds that⁴ $B_{p_1,q}^\beta(\Omega) \hookrightarrow B_{p_2,q}^\beta(\Omega)$ for $\gamma = 1 - (1/p_1 - 1/p_2)_+/\tilde{\beta}$.
3. For $0 < p, q_1, q_2 < \infty$, $q_1 < q_2$, and $\beta \in \mathbb{R}_{++}^d$, it holds that $B_{p,q_1}^\beta \hookrightarrow B_{p,q_2}^\beta$. In particular, with properties 1 and 2, if $\tilde{\beta} > 1/p$, it holds that $B_{p,q}^\beta(\Omega) \hookrightarrow C^\gamma(\Omega)$ where $\gamma = 1 - 1/(\tilde{\beta}p)$.
4. For $0 < p, q < \infty$ and $\beta \in \mathbb{R}_{++}^d$, if $\tilde{\beta} > 1/p$, then $B_{p,q}^\beta(\Omega) \hookrightarrow C^0(\Omega)$.

This result is basically proven by Triebel (2011). For completeness, we provide its derivation in Appendix D. If the average smoothness $\tilde{\beta}$ is sufficiently large ($\tilde{\beta} > 1/p$), then the functions in $B_{p,q}^\beta$ are continuous; however, if it is small ($\tilde{\beta} < 1/p$), then they are no longer continuous. Small p indicates spatially inhomogeneous smoothness; thus, spikes and jumps appear (see Donoho & Johnstone (1998) for this perspective, from the viewpoint of wavelet analysis).

2.2 Model of the true function

As a model of the true function f° , we consider two types of models: *Affine composition model* and *deep composition model*. For a Banach space H , we let $U(H)$ be the unit ball of H .

(a) Affine composition model: The first model we introduced is a very naive model which is just a composition of an affine transformation and a function in the anisotropic Besov space:

$$H_{\text{aff}} := \{f \circ h \mid h \in U(B_{p,q}^\beta([0,1]^{\tilde{d}})), A \in \mathbb{R}^{\tilde{d} \times d}, b \in \mathbb{R}^{\tilde{d}} \text{ s.t. } Ax + b \in [0,1]^{\tilde{d}} (\forall x \in \Omega)\},$$

where we assume $\tilde{d} \geq d$. Here, we assumed that the affine transformation has an appropriate scaling such that $Ax + b$ is included in the domain of h for all $x \in \Omega$. This is a quite naive model but provides an instructive example to understand how the estimation error of deep learning behaves under the anisotropic setting.

(b) Deep composition model: The *deep composition model* generalizes the affine composition model to a composition of nonlinear functions. Let $m_1 = d$, $m_{L+1} = 1$, m_ℓ be the dimension of the

³Here, we let $(x)_+ := \max\{x, 0\}$.

⁴The symbol \hookrightarrow means continuous embedding.




Figure 1: Near low dimensional data distribution with anisotropic smoothness of the target function. The target function has less smoothness (s_1, s_2) toward the first two coordinates on the manifold while it is almost constant toward the third coordinate (large s_3).

ℓ th layer, and let $\beta^{(\ell)} \in \mathbb{R}_{++}^{m_\ell}$ be the smoothness parameter in the ℓ th layer. The deep composition model is defined as

$$H_{\text{deep}} := f \circ h_H \quad h_{\ell,k} : [0, 1]^{m_\ell} \rightarrow [0, 1]^{m_{\ell+1}}, \quad h_{\ell,k} \in U(B_{p,q}^{\beta^{(\ell)}}([0, 1]^{m_\ell})) \quad (k \in [m_{\ell+1}])$$

Here, the interval $[0, 1]$ can be replaced by another compact interval, such as $[a_\ell, b_\ell]$, but this difference can be absorbed by changing a scaling factor. The assumption $\|h_{\ell,k}\|_{B_{p,q}^{\beta^{(\ell)}}} \leq 1$ can also be relaxed, but we do not pursue that direction due to presentation simplicity. This model includes the affine composition model as a special case. However, it requires a stronger assumption to properly evaluate the estimation error on this model.

Examples The model we have introduced includes some instructive examples as listed below:

(a) **Linear projection** Schmidt-Hieber (2020) analyzed estimation of the following model by deep learning: $f^\circ(x) = g(w^\top x)$ where $g \in \mathcal{C}^\beta([0, 1])$ and $w \in \mathbb{R}^d$. In this example, the function f° varies along only one direction, w . Apparently, this is an example of the affine composition model.

(b) **Distribution on low dimensional smooth manifold** Assume that the input x is distributed on a low-dimensional smooth manifold embedded in Ω , and the smoothness of the true function f° is anisotropic along a coordinate direction on the manifold. We suppose that the low dimensional manifold is \tilde{d} -dimensional and $\tilde{d} \ll d$. In this situation, the true function can be written as $f^\circ(x) = h(\phi(x))$ where $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^{\tilde{d}}$ is a map that returns the coordinate of x on the manifold and h is an element in an anisotropic Besov space on $\mathbb{R}^{\tilde{d}}$. This situation appears if data is distributed on a low-dimensional sub-manifold of Ω and the target function is invariant against noise injection to some direction on the manifold at each input point x (Figure 1 illustrates this situation). One typical example of this situation is a function invariant with data augmentation (Simard et al., 2003; Krizhevsky et al., 2012). Even if the noise injection destroys low dimensionality of the data distribution (i.e., $\tilde{d} = d$), an anisotropic smoothness of the target function eases the curse of dimensionality as analyzed below, which is quite different from existing works (Yang & Dunson, 2016; Bickel & Li, 2007; Nakada & Imaizumi, 2020; Schmidt-Hieber, 2019; Chen et al., 2019; Chen et al., 2019).

Related work Here, we introduce some more related work and discuss their relation to our analysis. The statistical analysis on an anisotropic Besov space can be back to Ibragimov & Khas'minskii (1984) who considered density estimation, where the density is assumed to be included in an anisotropic Sobolev space with $p \geq 2$, and derived the minimax optimal rate $n^{-r\tilde{\beta}/(2\tilde{\beta}+1)}$ with respect to L^r -norm. Nyssbaum (1983, 1987) analyzed a nonparametric regression problem on an anisotropic Besov space. Following these results, several studies have been conducted in the literature pertaining to nonparametric statistics, such as nonlinear kernel estimator Kerkycharian et al. (2001), adaptive confidence band construction Hoffman & Lepski (2002), optimal aggregation Gaifas & Lecue (2011), Gaussian process estimator Bhattacharya et al. (2011, 2014), and kernel ridge regression Hang & Steinwart (2018). Basically, these studies investigated estimation problems in which the target function is in anisotropic Besov spaces, but the composition models considered in this paper have not been analyzed. Hoffman & Lepski (2002); Bhattacharya et al. (2011) considered a dimension reduction model; that is, the target function is dependent on only a few variables of x , but they did not deal with more general models, such as the affine/deep composition models. The nonparametric regression problems where the input data are distributed on a low-dimensional smooth manifold has been studied as a ‘‘manifold regression’’ Yang & Dunson (2016); Bickel & Li

(2007); Yang & Tokdar (2015). Such a model can be considered as a specific example of the deep composition model. In this sense, our analysis is a significant extension of these analyses.

3 Approximation error analysis

Here, we consider approximating the true function f^o via a deep neural network and derive the approximation error. As the activation function, we consider the ReLU activation denoted by $\eta(x) = \max\{x, 0\}$ ($x \in \mathbb{R}$). For a vector x , $\eta(x)$ is operated in an element-wise manner. The model of neural networks with height L , width W , sparsity constraint S , and norm constraint B as $\Phi(L, W, S, B) := \{f(W^{(L)}\eta(\dots + b^{(L)})) \mid (W^{(1)}x + b^{(1)}) \dots W^{(L)} \in \mathbb{R}^{1 \times W}, b^{(L)} \in \mathbb{R}, W^{(1)} \in \mathbb{R}^{W \times d}, b^{(1)} \in \mathbb{R}^W, W^{(\ell)} \in \mathbb{R}^{W \times W}, b^{(\ell)} \in \mathbb{R}^W (1 < \ell < L), \sum_{\ell=1}^L (kW^{(\ell)}k_0 + kb^{(\ell)}k_0) \leq S, \max_{\ell} kW^{(\ell)}k_1 \leq B\}$, where $k \cdot k_0$ is the ℓ_0 -norm of the matrix (the number of non-zero elements of the matrix), and $k \cdot k_1$ is the ℓ_1 -norm of the matrix (maximum of the absolute values of the elements). The sparsity constraint and norm bounds are required to obtain the near-optimal rate of the estimation error. To evaluate the accuracy of the deep neural network model in approximating target functions, we define the worst-case approximation error as

$$R_r(F, H) := \sup_{f \in F} \inf_{f \in H} \|f - f\|_{L^r(\Omega)},$$

where F is the set of functions used for approximation, and H is the set of target functions.

Proposition 2 (Approximation ability for anisotropic Besov space). *Suppose that $0 < p, q, r \leq 1$ and $\beta \in \mathbb{R}_{++}^d$ satisfy the following condition: $\tilde{\beta} > (1/p - 1/r)_+$. Assume that $m \in \mathbb{N}$ satisfies $0 < \tilde{\beta} < \min(m, m - 1 + 1/p)$. Let $\delta = (1/p - 1/r)_+$, $\nu = (\tilde{\beta} - \delta)/(2\delta)$ and $W_0(d) := 6dm(m+2) + 2d$. Then, for $N \in \mathbb{N}$, we can bound the approximation error as*

$$R_r(\Phi(L_1, W_1, S_1, B_1), U(B_{p,q}^{\tilde{\beta}}(\Omega))) \leq N^{-\tilde{\beta}},$$

by setting

$$L_1(d) := 3 + 2d \log_2 \left(\frac{3^{d-m}}{\epsilon c_{(d,m)}} \right) + 5ed \log_2(d - m)e, \quad W_1(d) := NW_0, \quad (2)$$

$$S_1(d) := [(L - 1)W_0^2 + 1]N, \quad B_1(d) := O(N^{d(1+\nu - 1)(1/p - \tilde{\beta})_+}), \quad (3)$$

for $\epsilon = N^{-\tilde{\beta} \log(N) - 1}$ and a constant $c_{(d,m)}$ depending only on d and m .

The proof of this proposition is provided in Appendix B. The rate $N^{-\tilde{\beta}}$ is the optimal adaptive approximation error rate that can be achieved by a model with N parameters (the difference between adaptive and non-adaptive methods is explained in the discussion below). Note that this is an approximation error in an oracle setting and no sample complexity appears here. We notice that we can avoid the *curse of dimensionality* if the average smoothness $\tilde{\beta}$ is small. This means that if the target function is non-smooth in only a few directions and smooth in other directions, we can avoid the curse of dimensionality. In contrast, if we consider an isotropic Besov space where $\beta_1 = \dots = \beta_d (= \beta)$, then $\tilde{\beta} = \beta/d$, which directly depends on the dimensionality d , and we need an exponentially large number of parameters in this situation to achieve ϵ -accuracy. Therefore, the anisotropic smoothness has a significant impact on the approximation error rate. The assumption $\tilde{\beta} > (1/p - 1/r)_+$ ensures the L_r -integrability of the target function, and the inequality (without equality) admits a near-optimal wavelet approximation of the target function in terms of L_r -norm.

Using this evaluation as a basic tool, we can obtain the approximation error for the deep composition models. We can also obtain the approximation error for the affine composition models, but it is almost identical to Proposition 2. Therefore, we defer it to Appendix A.

Theorem 1 (Deep composition model). *Assume that $\tilde{\beta}^{(\ell)} > 1/p$ for all $\ell = 1, \dots, H$. Then, the estimation error on the deep composition model is bounded as*

$$R_1(\Phi(L, W, S, B), H_{\text{deep}}) \leq \max_{\ell \in [H]} N^{-\tilde{\beta}^{(\ell)}}, \quad (4)$$

where $\tilde{\beta}^{(\ell)} = \tilde{\beta}^{(\ell)} \prod_{k=\ell+1}^H [(\beta^{(k)} - 1/p) \wedge 1]$, and $L = \sum_{\ell=1}^H (L_1(m_\ell) + 1)$, $W = \max_{\ell} (W_1(m_\ell) - m_{\ell+1})$, $S = \sum_{\ell=1}^H (S_1(m_\ell) + 3m_{\ell+1})$, $B = \max_{\ell} B_1(m_\ell)$.

The proof can be found in Appendix B.1. Since the model is more general than the vanilla anisotropic Besov space, we require a stronger assumption $\tilde{\beta}^{(\ell)} > 1/p$ on $\tilde{\beta}^{(\ell)}$ than the condition in Proposition 2. This is because we need to bound the Hölder smoothness of the remaining layers to bound the influence of the approximation error in the internal layers to the entire function. Hölder smoothness is ensured according to the embedding property under this condition (Proposition 1). This Hölder smoothness assumption affects the approximation error rate. The convergence rate $\tilde{\beta}^{(\ell)}$ in Eq. (4) is different from that in Eq. (8). This is because the approximation error in the internal layers are propagated through the remaining layers with Hölder smoothness and its amplitude is controlled by the Hölder smoothness.

Approximation error by non-adaptive method The approximation error obtained in the previous section is called an adaptive error rate in the literature regarding approximation theory (DeVore, 1998). If we fix N bases beforehand and approximate the target function by a linear combination of the N bases (which is called the non-adaptive method), then we *cannot* achieve the adaptive error rate obtained in the previous section. Roughly speaking, the approximation error of non-adaptive methods is lower bounded by $N^{-\left(\frac{1}{p} - \frac{1}{\min\{2, r, g\} + 1}\right)}$ (Myronyuk, 2015, 2016, 2017), which is slower than the approximation error rate of deep neural networks especially for small p .

4 Estimation error analysis

In this section, we analyze the accuracy of deep learning in estimating a function in compositions of anisotropic Besov spaces. We consider a least-squares estimator in the deep neural network model:

$$\hat{f} = \operatorname{argmin}_{\bar{f}: f \in \mathcal{F}(L, W, S, B)} \sum_{i=1}^n (y_i - \bar{f}(x_i))^2 \quad (5)$$

where \bar{f} is the *clipping* of f defined by $\bar{f} = \min\{f, F\} \max\{f, -F\}$, F, G for a constant $F > 0$ which is realized by ReLU units. The network parameters (L, W, S, B) should be specified appropriately as indicated in Theorems 2 and 3. In practice, these parameters can be specified by cross validation. Indeed, we can theoretically show that cross validation can provide the appropriate choice of these parameters in compensation to an additional $\log(n)$ -factor in the estimation error bound. This estimator can be seen as a sparsely regularized estimator because there are constraints on S . In terms of optimization, this requires a combinatorial optimization, but we do not pursue the computational aspect. The estimation error that we derive in this section can involve the optimization error, but for simplicity, we only demonstrate the estimation error of the *ideal* situation where there is no optimization error.

Affine composition model The following theorem provides an upper bound of the estimation error for the affine composition model.

Theorem 2. *Assume the same condition as in Theorem 6; in particular, suppose $0 < p, q \leq 1$ and $\tilde{\beta} > (1/p - 1/2)_+$ for $\tilde{\beta} \geq \mathbb{R}_{++}^d$. Moreover, we assume that the distribution P_X has a density p_X such that $k p_X k_1 \leq R$ for a constant $R > 0$. If $f^\circ \in H_{\text{aff}} \setminus L^1(\Omega)$, and $k f^\circ k_1 \leq F$ for $F \geq 1$; then, letting $(W, L, S, B) = (L_1(\tilde{d}), W_1(\tilde{d}), S_1(\tilde{d}), (\tilde{d}C + 1)B_1(\tilde{d}))$ as in Theorem 6 with $N \leq n^{\frac{1}{2\tilde{\beta}+1}}$, we obtain*

$$\mathbb{E}_{D_n} [k f^\circ - \hat{f} k_{L^2(P_X)}^2] \leq n^{-\frac{2\tilde{\beta}}{2\tilde{\beta}+1}} \log(n)^3,$$

where $\mathbb{E}_{D_n}[\cdot]$ indicates the expectation with respect to the training data D_n .

The proof is provided in Appendix C. We will show that the convergence rate $n^{-2\tilde{\beta}/(2\tilde{\beta}+1)}$ is minimax optimal in Section 5 (see also Kerkycharian & Picard (1992); Donoho et al. (1996); Donoho & Johnstone (1998); Giné & Nickl (2015) for ordinary Besov spaces). The L^1 -norm constraint $k f^\circ k_1 \leq F$ is used to derive a uniform bound on the discrepancy between the population and the empirical L^2 -norm. Without this condition, the convergence rate could be slower.

Deep composition model For the deep composition model, we obtain the following convergence rate. This is an extension of Theorem 2 but requires a stronger assumption on the smoothness.

Theorem 3. *Suppose that $0 < p, q \leq 1$ and $\tilde{\beta}^{(\ell)} > 1/p$ for all $\ell \geq [H]$. If $f^\circ \in H_{\text{deep}} \setminus L^1(\Omega)$, and $k f k_1 \leq F$ for $F \geq 1$, then we obtain*

$$\mathbb{E}_{D_n} [k f^\circ - \hat{f} k_{L^2(P_X)}^2] \leq \max_{\ell \geq [H]} n^{-2\tilde{\beta}^{(\ell)}/(2\tilde{\beta}^{(\ell)}+1)} \log(n)^3,$$

where $\tilde{\beta}^{(\ell)}$ is defined in Theorem 1, and (L, W, S, B) is as given in Theorem 1 with $N = \max_{\ell \in [L]} n^{\frac{1}{2\tilde{\beta}^{(\ell)} + 1}}$.

The proof is provided in Appendix C. We will show that this is also minimax optimal in Theorem 4. Because of the Hölder continuity, the convergence rate becomes slower than the affine composition model (that is, $\tilde{\beta}^{(\ell)} = \tilde{\beta}^{(\ell)}$). However, this slower rate is unavoidable in terms of the minimax optimal rate. Schmidt-Hieber (2020) analyzed the same situation for the Hölder class which corresponds to $\beta_1^{(\ell)} = \beta_d^{(\ell)}(\beta)$ and $p = q = 1$. Our analysis far extends their analysis to the setting of anisotropic Besov spaces in which the parameters $\beta^{(\ell)}, p, q$ have much more freedom.

From these two bounds (Theorems 2 and 3), we can see that as the smoothness $\tilde{\beta}$ becomes large, the convergence rates faster. If the target function is included in the isotropic Besov space with smoothness $\beta_1 = \beta_d(= \beta)$, then the estimation error becomes

$$\text{(Isotropic Besov)} \quad n^{-2\beta/(2\beta+d)}.$$

In the exponent, the dimensionality d appears, which causes the curse of dimensionality. In contrast, if the target function is in the anisotropic Besov space, and the smoothness in each direction is sufficiently imbalanced such that $\tilde{\beta}$ does not depend on d , our obtained rate

$$\text{(Anisotropic Besov)} \quad n^{-2\tilde{\beta}/(2\tilde{\beta}+1)}$$

avoids the curse of dimensionality. For high-dimensional settings, there would be several redundant directions in which the true function does not change. Deep learning is adaptive to this redundancy and achieves a better estimation error. However, in Section 6, we prove that linear estimators are affected by the dimensionality more strongly than deep learning. This indicates the superiority of deep learning.

5 Minimax optimal rate

Here, we show that the estimation error rate, that we have presented, of deep learning achieves the *minimax optimal rate*. Roughly speaking the minimax optimal risk on a model F of the true function is the smallest worst case error over all estimators:

$$R(F) := \inf_{\hat{f}} \sup_{f \in F} \mathbb{E}_{D_n} [k_{\hat{f}}(f) - f]_{L^2(P_X)}^2,$$

where \hat{f} runs over all estimators. The convergence rate of the minimax optimal risk is referred to as minimax optimal rate. We obtain the following minimax optimal rate for anisotropic Besov spaces.

Theorem 4. (a) Affine composition model: For $0 < p, q \leq 1$ and $\beta \in \mathbb{R}_{++}^d$, assume that $\tilde{\beta} > \max\{1/p - 1/2, 1/p - 1, 0\}$. Then, the minimax optimal risk of the affine composition model is lower bounded as $R(H_{\text{aff}}) \asymp n^{-\frac{2\tilde{\beta}}{2\tilde{\beta}+1}}$. **(b) Deep composition model:** For $0 < p, q \leq 1$ and $\beta^{(\ell)} \in \mathbb{R}_{++}^d$ ($\ell = 1, \dots, H$), assume that $\tilde{\beta}^{(\ell)} > 1/p$. Let $\epsilon > 0$ be arbitrarily small for $q < 1$, and let $\epsilon = 0$ for $q = 0$. Let $\tilde{\beta}^{(\ell)} = \tilde{\beta}^{(\ell)} \prod_{k=\ell+1}^H [(\beta^{(k)} - 1/p + \epsilon) \wedge 1]$, and $\tilde{\beta} := \min_{\ell} \tilde{\beta}^{(\ell)}$. Then, the minimax optimal risk of the deep composition model is lower bounded as $R(H_{\text{deep}}) \asymp n^{-\frac{2\tilde{\beta}}{2\tilde{\beta}+1}}$.

The proof is provided in Appendix E (see also Ibragimov & Khas'minskii (1984); Nyssbaum (1987)). From this theorem, we can see that the estimation error of deep learning shown in Theorems 2 and 3 indeed achieve the minimax optimal rate up to a poly-log(n) factor.

6 Suboptimality of linear estimators

In this section, we give the minimax optimal rate in the class of *linear estimators*. The linear estimator is a class of estimators that can be written as

$$\hat{f}(x) = \sum_{i=1}^n y_i \varphi_i(x; X^n),$$

where $X^n = (x_1, \dots, x_n)$ and $\varphi_i(x; X^n)$ ($i = 1, \dots, n$) are (measurable) functions that only depend on x and X^n . This is linearly dependent on $Y^n = (y_1, \dots, y_n)$. We notice that the kernel

ridge regression is included in this class because it can be written as $\widehat{f}(x) = k_{x, X^n}(k_{X^n, X^n} + \lambda I)^{-1} Y^n$, which linearly depends on Y^n . This class includes other important estimators, such as the Nadaraya–Watson estimator, the k -nearest neighbor estimator, and the sieve estimator. We compare deep learning with the linear estimators in terms of minimax risk. For this purpose, we define the minimax risk of the class of linear estimators:

$$R^{(\text{lin})}(F) := \inf_{\widehat{f}: \text{linear}} \sup_{f \in 2F} \mathbb{E}_{D_n} [k f \circ \widehat{f} k_{L^2(P_X)}^2],$$

where \widehat{f} runs over all *linear estimators*. We can see that linear estimators suffer from the sub-optimal rate because of the following two points: (i) they do not have adaptivity, and (ii) they significantly suffer from the curse of dimensionality.

Theorem 5. (i) *Suppose that the input distribution P_X is the uniform distribution on $\Omega = [0, 1]^d$ and assume that $\underline{\beta} > 1/p$ and $1 - p, q \leq 1$. Then, the minimax optimal rate of the linear estimators is lower bounded as*

$$R^{(\text{lin})}(U(B_{p,q}^\beta)) \asymp n^{-\frac{2\underline{\beta} - v}{2\underline{\beta} + 1 - v}}, \quad (6)$$

where $v = 2(1/p - 1/2)_+$.

(ii) *In addition to the above conditions, we assume that the true function is included in the affine composition model with $\widetilde{d} \leq d$, $\underline{\beta} = \beta_1 = \dots = \beta_{\widetilde{d}}$ and $0 < p \leq 2$. Let $\alpha_d = 1 + \kappa$ with arbitrary small $\kappa > 0$ when $\widetilde{d} < d/2$, and let $\alpha_d = 0$ when $\widetilde{d} \geq d/2$. Then, the minimax rate of the linear estimators on the affine composition model is lower bounded by*

$$R^{(\text{lin})}(H_{\text{aff}}) \asymp n^{-\frac{2(\underline{\beta} - \widetilde{d}/p + d/2 + \alpha_d)}{2(\underline{\beta} - \widetilde{d}/p + d/2 + \alpha_d) + d}}. \quad (7)$$

The proof is provided in Appendix F. (i) The lower bound (7) reveals the suboptimality of linear estimators in terms of input dimensionality. Actually, if we consider a particular case where $\widetilde{d} = 1$, $p = 1$ and $d \geq \widetilde{d}$, then the obtained minimax rate of linear estimators and the estimation error rate of deep learning can be summarized as

$$\text{linear} : n^{-\frac{2\underline{\beta} + d}{2\underline{\beta} + 2d}}, \quad \text{deep} : n^{-\frac{2\underline{\beta}}{2\underline{\beta} + 1}},$$

by Theorem 2 when $\underline{\beta} > 1$ (which can be checked by noticing $\widetilde{d} = \underline{\beta}/\widetilde{\beta} = 1$ in this situation). We can see that the dependence on the dimensionality of linear estimators is significantly worse than that of deep learning. This indicates poor adaptivity of linear estimators to the intrinsic dimensionality of data. Actually, as d becomes large, the rate for the linear estimator approaches to $1/\sqrt[n]{n}$ but that for the deep learning is not affected by d and still faster than $1/\sqrt[n]{n}$. To show the theorem, we used the ‘‘convex-hull argument’’ developed by Hayakawa & Suzuki (2019); Donoho & Johnstone (1998). We combined this technique with the so-called *Irie-Miyake’s integral representation* (Irie & Miyake, 1988; Hornik et al., 1990). Note that this difference appears because there is an affine transformation in the first layer of the affine composition model. Deep learning is flexible against such a coordinate transform so that it can find directions to which the target function is smooth. In contrast, kernel methods do not have such adaptivity because there is no feature extraction layer. (ii) The lower bound (6) states that when $p < 2$ (that is, $v > 0$), the minimax rate of the linear estimators is outperformed by that of deep learning (Theorem 2). This is due to the ‘‘adaptivity’’ of deep learning. When p is small, the smoothness of the target function is less homogeneous, and it requires an adaptive approximation scheme to achieve the best estimation error. Linear estimators do not have adaptivity and thus fail to achieve the minimax optimal rate. Our bound (6) extends the result by Zhang et al. (2002) to a multivariate anisotropic Besov space while Zhang et al. (2002) investigated the univariate space ($d = 1$).

7 Conclusion

We investigated the approximation error and estimation error of deep learning in the anisotropic Besov spaces. It was proved that the convergence rate is determined by the average of the anisotropic smoothness, which results in milder dependence on the input dimensionality. If the smoothness is

highly anisotropic, deep learning can avoid overfitting. We also compared the error rate of deep learning with that of linear estimators and showed that deep learning has better dependence on the input dimensionality. Moreover, it was shown that deep learning can achieve the adaptive rate and outperform non-adaptive approximation methods and linear estimators if the homogeneity p of smoothness is small. These analyses strongly support the practical success of deep learning from a theoretical perspective.

Limitations of this work Our work does not cover the optimization aspect of deep learning. It is assumed that the regularized least squares (5) can be executed. It would be nice to combine our study with recent developments of non-convex optimization techniques (Vempala & Wibisono, 2019; Suzuki & Akiyama, 2021).

Potential negative societal impact Since this is purely theoretical result, it is not expected that there is a direct negative societal impact. However, revealing detailed properties of the deep learning could promote an opportunity to pervert deep learning.

Acknowledgment

TS was partially supported by JSPS KAKENHI (18H03201), Japan Digital Design and JST CREST. AN was partially supported by JSPS Kakenhi (19K20337) and JST-PRESTO.

References

- M. Belkin and P. Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural computation*, 15(6):1373–1396, 2003.
- A. Bhattacharya, D. Pati, and D. B. Dunson. Adaptive dimension reduction with a gaussian process prior. *arXiv preprint arXiv:1111.1044*, 1445, 2011.
- A. Bhattacharya, D. Pati, and D. Dunson. Anisotropic function estimation using multi-bandwidth gaussian processes. *Annals of statistics*, 42(1):352, 2014.
- P. J. Bickel and B. Li. Local polynomial regression on unknown manifolds. In *Complex datasets and inverse problems*, pp. 177–186. Institute of Mathematical Statistics, 2007.
- M. Chen, H. Jiang, W. Liao, and T. Zhao. Nonparametric Regression on Low-Dimensional Manifolds using Deep ReLU Networks. *arXiv e-prints*, art. arXiv:1908.01842, Aug 2019.
- M. Chen, H. Jiang, W. Liao, and T. Zhao. Efficient approximation of deep relu networks for functions on low dimensional manifolds. In *Advances in Neural Information Processing Systems*, pp. 8172–8182, 2019.
- G. Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals, and Systems (MCSS)*, 2(4):303–314, 1989.
- J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *arXiv e-prints*, art. arXiv:1810.04805, Oct 2018.
- R. A. DeVore. Nonlinear approximation. *Acta Numerica*, 7:51–150, 1998.
- R. A. DeVore and V. A. Popov. Interpolation of Besov spaces. *Transactions of the American Mathematical Society*, 305(1):397–414, 1988.
- R. A. DeVore, G. Kyriazis, D. Leviatan, and V. M. Tikhomirov. Wavelet compression and nonlinear-widths. *Advances in Computational Mathematics*, 1(2):197–214, 1993.
- D. L. Donoho and I. M. Johnstone. Minimax estimation via wavelet shrinkage. *The Annals of Statistics*, 26(3):879–921, 1998.
- D. L. Donoho, I. M. Johnstone, G. Kerkycharian, and D. Picard. Density estimation by wavelet thresholding. *The Annals of Statistics*, 24(2):508–539, 1996.

- D. Dũng. Optimal adaptive sampling recovery. *Advances in Computational Mathematics*, 34(1): 1–41, 2011.
- S. Gaïffas and G. Lecue. Hyper-sparse optimal aggregation. *Journal of Machine Learning Research*, 12(Jun):1813–1833, 2011.
- E. Giné and R. Nickl. *Mathematical Foundations of Infinite-Dimensional Statistical Models*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2015.
- X. Glorot, A. Bordes, and Y. Bengio. Deep sparse rectifier neural networks. In *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics*, volume 15 of *Proceedings of Machine Learning Research*, pp. 315–323, 2011.
- H. Hang and I. Steinwart. Optimal learning with anisotropic gaussian svms. *arXiv preprint arXiv:1810.02321*, 2018.
- S. Hayakawa and T. Suzuki. On the minimax optimality and superiority of deep neural network learning over sparse parameter spaces. *arXiv preprint arXiv:1905.09195*, 2019.
- M. Hoffman and O. Lepski. Random rates in anisotropic regression (with a discussion and a rejoinder by the authors). *The Annals of Statistics*, 30(2):325–396, 04 2002.
- K. Hornik. Approximation capabilities of multilayer feedforward networks. *Neural Networks*, 4(2): 251–257, 1991.
- K. Hornik, M. Stinchcombe, and H. White. Universal approximation of an unknown mapping and its derivatives using multilayer feedforward networks. *Neural Networks*, 3(5):551–560, 1990.
- I. Ibragimov and R. Khas’minskii. More on the estimation of distribution densities. *Journal of Soviet Mathematics*, 25(3):1155–1165, 1984.
- B. Irie and S. Miyake. Capabilities of three-layered perceptrons. In *IEEE 1988 International Conference on Neural Networks*, pp. 641–648, 1988.
- G. Kerkycharian and D. Picard. Density estimation in Besov spaces. *Statistics & Probability Letters*, 13:15–24, 1992.
- G. Kerkycharian, O. Lepski, and D. Picard. Nonlinear estimation in anisotropic multi-index denoising. *Probability Theory and Related Fields*, 121(2):137–170, Oct 2001.
- A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.
- V. Myronyuk. Trigonometric approximations and kolmogorov widths of anisotropic besov classes of periodic functions of several variables. *Ukrainian Mathematical Journal*, 66(8), 2015.
- V. V. Myronyuk. Kolmogorov widths of the anisotropic besov classes of periodic functions of many variables. *Ukrainian Mathematical Journal*, 68(5):718–727, Oct 2016.
- V. V. Myronyuk. Widths of the anisotropic besov classes of periodic functions of several variables. *Ukrainian Mathematical Journal*, 68(8):1238–1251, Jan 2017. ISSN 1573-9376. doi: 10.1007/s11253-017-1290-1. URL <https://doi.org/10.1007/s11253-017-1290-1>.
- V. Nair and G. E. Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th International Conference on Machine Learning*, pp. 807–814, 2010.
- R. Nakada and M. Imaizumi. Adaptive approximation and generalization of deep neural network with intrinsic dimensionality. *Journal of Machine Learning Research*, 21(174):1–38, 2020. URL <http://jmlr.org/papers/v21/20-002.html>.
- S. M. Nikol’skii. *Approximation of functions of several variables and imbedding theorems*, volume 205. Springer-Verlag Berlin Heidelberg, 1975.
- M. Nyssbaum. Optimal filtration of a function of many variables in white gaussian noise. *Problems of Information Transmission*, 19:23–29, 1983.

- M. Nyssbaum. Nonparametric estimation of a regression function that is smooth in a domain in \mathbb{R}^k . *Theory of Probability & Its Applications*, 31(1):108–115, 1987.
- A. Radford, L. Metz, and S. Chintala. Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. *arXiv e-prints*, art. arXiv:1511.06434, Nov 2015.
- J. Schmidt-Hieber. Deep ReLU network approximation of functions on a manifold. *arXiv preprint arXiv:1908.00695*, 2019.
- J. Schmidt-Hieber. Nonparametric regression using deep neural networks with relu activation function. *The Annals of Statistics*, 48(4):1875–1897, 2020.
- P. Y. Simard, D. Steinkraus, and J. C. Platt. Best practices for convolutional neural networks applied to visual document analysis. In *Proceedings of the Seventh International Conference on Document Analysis and Recognition-Volume 2*, pp. 958. IEEE Computer Society, 2003.
- S. Sonoda and N. Murata. Neural network with unbounded activation functions is universal approximator. *Applied and Computational Harmonic Analysis*, 43(2):233–268, 2017.
- T. Suzuki. Adaptivity of deep ReLU network for learning in Besov and mixed smooth Besov spaces: optimal rate and curse of dimensionality. In *International Conference on Learning Representations (ICLR2019)*, 2019. URL <https://openreview.net/forum?id=H1ebTsActm>.
- T. Suzuki and S. Akiyama. Benefit of deep learning with non-convex noisy gradient descent: Provable excess risk bound and superiority to kernel methods. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=2m0g1wEafh>.
- V. Temlyakov. *Approximation of Periodic Functions*. Nova Science Publishers, 1993.
- J. B. Tenenbaum, V. De Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *science*, 290(5500):2319–2323, 2000.
- H. Triebel. Entropy numbers in function spaces with mixed integrability. *Revista matemática complutense*, 24(1):169–188, 2011.
- S. Vempala and A. Wibisono. Rapid convergence of the unadjusted langevin algorithm: Isoperimetry suffices. In *Advances in Neural Information Processing Systems*, pp. 8094–8106, 2019.
- J. Vybiral. Function spaces with dominating mixed smoothness. *Dissertationes Math. (Rozprawy Mat.)*, 436:3–73, 2006.
- Y. Yang and D. B. Dunson. Bayesian manifold regression. *The Annals of Statistics*, 44(2):876–905, 2016.
- Y. Yang and S. T. Tokdar. Minimax-optimal nonparametric regression in high dimensions. *The Annals of Statistics*, 43(2):652–674, 2015.
- D. Yarotsky. Error bounds for approximations with deep relu networks. *Neural Networks*, 94: 103–114, 2017.
- S. Zhang, M.-Y. Wong, and Z. Zheng. Wavelet threshold estimation of a regression function with random design. *Journal of Multivariate Analysis*, 80(2):256–284, 2002.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [\[Yes\]](#)
 - (b) Did you describe the limitations of your work? [\[Yes\]](#)
 - (c) Did you discuss any potential negative societal impacts of your work? [\[No\]](#)
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [\[Yes\]](#)

2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes] Each theorem explicitly describes assumptions on which the theorem is established.
 - (b) Did you include complete proofs of all theoretical results? [Yes] See Appendix in the supplementary material. All proofs are included there.
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [N/A]
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [N/A]
 - (b) Did you mention the license of the assets? [N/A]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]