Asymptotic Properties for Bayesian Neural Network in Besov Space

Kyeongwon Lee
Department of Statistics
Seoul National University
Seoul, Republic of Korea 08826
lkw1718@snu.ac.kr

Jaeyong Lee
Department of Statistics
Seoul National University
Seoul, Republic of Korea 08826
leejyc@gmail.com

Abstract

Neural networks have shown great predictive power when dealing with various unstructured data such as images and natural languages. The Bayesian neural network captures the uncertainty of prediction by putting a prior distribution for the parameter of the model and computing the posterior distribution. In this paper, we show that the Bayesian neural network using spike-and-slab prior has consistency with nearly minimax convergence rate when the true regression function is in the Besov space. Even when the smoothness of the regression function is unknown the same posterior convergence rate holds and thus the spike-and-slab prior is adaptive to the smoothness of the regression function. We also consider the shrinkage prior, which is more feasible than other priors, and show that it has the same convergence rate. In other words, we propose a practical Bayesian neural network with guaranteed asymptotic properties.

1 Introduction

The neural network is a machine learning technique that can handle complex structures in data by combining linear transformations and non-linear transformations [Goodfellow et al., 2016]. The neural network can show great predictive power when dealing with unstructured data ranging from image data to natural language processing. Moreover, end-to-end learning is possible by grasping the inherent structure of data only with inputs and outputs without explicit intermediate steps. Such properties have played major roles in the rapid development of numerous artificial intelligence applications.

The Bayesian neural network (BNN) provides the uncertainty of the prediction from the probability distribution of the model parameters. Before observing the data, a prior distribution that reflects prior beliefs about the model is put over the parameter space of the neural network model. As observations are added, a posterior distribution is computed by updating the prior distribution via Bayes' rule. Neal [1996], Williams [1996] showed that a shallow BNN which has only one infinite-width hidden layer converges to a Gaussian process regression model. This means that the shallow BNN can be regarded as a finite approximation of the nonparametric Bayesian regression model. Lee et al. [2017], Matthews et al. [2018] extended this result to the deep BNN model.

However, it is difficult to compute the exact posterior distribution of the BNN since the normalizing constant of the posterior distribution is intractable. Therefore, various methods have been proposed to infer BNN through approximate Bayesian computation algorithms. The methods based on the Markov chain Monte Carlo (MCMC) algorithm such as the Gibbs sampler [Geman and Geman, 1984, Casella and George, 1992] and Hamiltonian Monte Carlo (HMC) [Neal et al., 2011] is effective to approximate true posterior. By constructing a Markov chain that has the desired distribution as its stationary distribution, one can obtain a sample of the desired distribution. Recently, scal-

36th Conference on Neural Information Processing Systems (NeurIPS 2022).

able MCMC algorithms such as stochastic gradient Langevin dynamics [Welling and Teh, 2011] and stochastic gradient HMC (SGHMC, Chen et al. [2014]) were proposed. The variational inference (VI) is an approximation method that finds the closest variational distribution from the posterior distribution. VI method changes the Bayesian inference problem to an optimization problem. Graves [2011] proposed to use of a Gaussian distribution as variational distribution for BNN and Kingma and Welling [2013] suggested the *reparameterization trick* for Gaussian variational distribution. Practically, VI methods for BNN are faster and more stable than MCMC methods. Blundell et al. [2015] suggested *Bayes by Backprop* algorithm for BNN via variational inference. There are also studies that approximate BNN with the dropout [Gal and Ghahramani, 2016] or the Gaussian process [Lee et al., 2017].

1.1 Related works

The performance of the neural network is justified by the *universal approximation* capability of the neural network. It is known that even a shallow neural network can approximate arbitrary continuous functions sufficiently precisely in L^p-sense [Cybenko, 1989, Hornik, 1991], and as the network deepens, the expressive power of the neural network gets stronger [Delalleau and Bengio, 2011, Bianchini and Scarselli, 2014, Telgarsky, 2015, 2016]. Lee III [1999] showed that the BNN has the universal approximation property.

As a measure of the expressive power of a model, consider the range of function spaces that a model can express. The choice of the activation function can affect the expressive power of a model, and in this paper, we consider the Rectified Linear Unit (ReLU) $ReLU(x) = \max\{0, x\}$ [Nair and Hinton, 2010], which is the most popular activation function in practice. Yarotsky [2017] evaluated the approximation error of the neural network using the ReLU activation function (ReLU network) for functions in the Sobolev space. Schmidt-Hieber [2020] showed that the regularized least squared estimator performed by the ReLU network in a nonparametric regression problem converges to the true regression function with nearly minimax rate up to logarithmic factors on the Hölder space. Suzuki [2018] extended the results in Schmidt-Hieber [2020] from the Hölder space to the Besov space which contains the Hölder and Sobolev spaces. Polson and Ročková [2018] proved a Bayesian version of the results in Schmidt-Hieber [2020]. They proved the Bayesian ReLU network has posterior consistency, the posterior probability is concentrated on the true parameter as collecting data, when the priors on the parameters of the ReLU network are spike-and-slab priors. Chérief-Abdellatif [2020] showed that the results in Polson and Ročková [2018] are still valid when variational methods are applied.

All of the above results assume a sparse structure in the neural networks. Sparsity is often assumed in high-dimensional models for asymptotic results. Castillo et al. [2015] showed theoretical properties including posterior consistency and rates of consistency of the high-dimensional linear regression model under the spike-and-slab prior. Song and Liang [2017] extended the results in Castillo et al. [2015] to the shrinkage priors which are more feasible than the spike-and-slab prior.

1.2 Main contribution

We show that the posterior of the Bayesian ReLU network with a spike-and-slab prior with appropriate width, depth, and sparsity level is consistent with a nearly minimax rate. We extend the result of Suzuki [2018] to the Bayesian framework and the results of Polson and Ročková [2018] to the Besov space. As in Polson and Ročková [2018], we show that the same convergence rate is achieved even when the smoothness parameter of the true regression function is not known. That is, the ReLU network has rich expressive power, and this ability is not constrained even without the smoothness information of the true function. We get similar results for the shrinkage prior.

We try to narrow the gap between the study of the theoretical properties and the practical application of the Bayesian ReLU networks. For instance, Gaussian distribution is frequently considered as a prior distribution in practical Bayesian ReLU networks. In this case, the posterior distribution may not be consistent since the posterior becomes heavy tail distribution over the parameter space [Vladimirova et al., 2019]. Approaches to interpreting the Bayesian neural networks as Gaussian Process [Lee et al., 2017] avoid this problem by setting the variance of the prior distribution to be proportional to the width of the network. We suggest similar sufficient conditions for the posterior consistency of the Bayesian ReLU networks. Compared to precedent studies, our setting is closer

to practical use. The results in Suzuki [2018] and Polson and Ročková [2018] require to solve the optimization problem under the l_0 constraint or compute posterior from spike-and-slab prior respectively. We suggest BNN model, which is easy to infer and consistent with the (nearly) optimal convergence rate.

The rest of the paper is organized as follows. We set up the problem and introduce the necessary concepts in Section 2. We present the obtained theoretical results in Section 3 and numerical examples in Section 4. A summary of the paper and a discussion are in Section 5.

2 Preliminaries

2.1 Notation

The floor function $\lfloor x \rfloor$ denotes the largest integer less than or equal to $x \in \mathbb{R}$, the ceiling function $\lceil x \rceil$ denotes the smallest integer greater than or equal to $x \in \mathbb{R}$, and $\|a\|_p$ is an l_p-norm of a real vector a. The u-th derivative of a d-dimensional function f is denoted by $D^u f$ with $u \in \{0, 1, \cdots\}^d$. The covering number $N(\epsilon, A, d)$ is the minimal number of d-balls of radius $\epsilon > 0$ necessary to cover a set A. We denote the dirac-delta function at $a \in \mathbb{R}$ by δ_a. We use $A \lesssim B$ as shorthand for the inequality $A \leq CB$ for some constant $C > 0$ and $A \vee B$ as shorthand for the $\max\{A, B\}$.

2.2 Model

Suppose that n input-output observations $\mathbb{D}_n = (X_i, y_i)_{i=1}^n \subset [0,1]^d \times \mathbb{R}$ are i.i.d. random sample from a regression model

$$y_i = f_0(X_i) + \xi_i \ (i = 1, 2, \cdots, n), \tag{1}$$

where $(\xi_i)_{i=1}^n$ is an i.i.d. sequence of Gaussian noises $\mathcal{N}\big(0,\ \sigma^2\big)$ with known variance $\sigma^2 > 0$ and f_0 is the true regression function belonging to the space $\mathscr{F}$.

In frequentist statistics, the maximum likelihood estimation is used to estimate the function f_0. From the machine learning viewpoint, it is equivalent to finding $\hat{f}$ that minimizes the square loss. Suzuki [2018] showed that f_0 can be estimated with a nearly optimal rate in this setting.

In practice, a sparse regularization term $r(f)$ such as dropout and l_1-regularizer is also considered together with a loss function. This is equivalent to estimating the MAP (maximum a posteriori) in the Bayesian modeling. The posterior distribution gives not only predictive values but also the uncertainty of the prediction. Thus, in the Bayesian statistics, it is important to find sufficient conditions for a prior distribution or a regularizer which guarantees the consistency or optimal convergence rate of the posterior.

2.3 Function spaces

In this section, we introduce function spaces that statisticians have studied for the *smoothness* of functions. Let $\Omega = [0,1]^d \subset \mathbb{R}^d$ be the domain of functions we consider.

The L^p-norm of a function is defined as follows:

$$\|f\|_{L^p} := \begin{cases} \left(\int_\Omega |f(x)|^p dx\right)^{1/p} & 0 < p < \infty, \\ \sup_{x \in \Omega} |f(x)| & p = \infty \end{cases}$$

where f is a real-valued function defined on Ω. The L^p space is the function space consisting of functions with bounded L^p-norm and is denoted by $L^p(\Omega) = \{f : \|f\|_{L^p} < \infty\}$.

Let $s > 0$ be the smoothness parameter and $m = \lfloor s \rfloor$. The s-Hölder norm of function f is defined by

$$\|f\|_{C^s} := \sup_{\|u\|_1 \leq m} \|D^u f\|_{L^\infty} + \sup_{\|u\|_1 = m} \sup_{x, y \in \Omega} \frac{|D^u f(x) - D^u f(y)|}{|x - y|^{s-m}}.$$

The s-Hölder space $C^s(\Omega)$ is defined as a set of functions with bounded s-Hölder norm. We call s the smoothness parameter of the s-Hölder space. Let $k \in \mathbb{N}$ and $1 \leq p \leq \infty$. The Sobolev norm of

a function f is defined by

$$\|f\|_{W^{k,p}} := \begin{cases} \left(\sum_{\|\alpha\|_1 \le k} \|D^\alpha f\|_{L^p}^p\right)^{1/p} & 1 \le p < \infty \\ \sup_{\|\alpha\|_1 \le k} \|D^\alpha f\|_{L^\infty} & p = \infty. \end{cases}$$

The Sobolev space $W^{k,p}(\Omega)$ is defined by the function space consisting of functions with bounded Sobolev norm.

Note the notion of the smoothness of a function is related to the differentiability of the function. The Besov space extends the concept of smoothness. Before defining the Besov space, define the r-th modulus of smoothness of the function f as follows [Gine and Nickl, 2016]:

$$w_{r,p}(f,t) = \sup_{\|h\|_2 \le t} \|\Delta_h^r(f)\|_p, \tag{2}$$

$$\Delta_h^r(f)(x) = \begin{cases} \sum_{j=0}^r \binom{r}{j} (-1)^{r-j} f(x+jh) & x \in \Omega, x+rh \in \Omega, \\ 0 & \text{o.w.} \end{cases} \tag{3}$$

For $0 < p, q \le \infty,\ s > 0,\ r = \lfloor s \rfloor + 1$, define the Besov norm of a function f by

$$\|f\|_{B^s_{p,q}} := \|f\|_p + \begin{cases} \left(\int_0^\infty \left(t^{-s} w_{r,p}(f,t)\right)^q \frac{dt}{t}\right)^{1/q} & q < \infty, \\ \sup_{t>0}\{t^{-s} w_{r,p}(f,t)\} & q = \infty. \end{cases} \tag{4}$$

The Besov space $B^s_{p,q}(\Omega)$ is defined as the set of functions with finite Besov norms [Gine and Nickl, 2016], i.e.,

$$B^s_{p,q}(\Omega) = \left\{ f : \|f\|_{B^s_{p,q}} < \infty \right\}. \tag{5}$$

Note that the Besov spaces can be defined without the differentiability and continuity of functions, and are more general than the Hölder and Sobolev spaces, which are subspaces of the Besov spaces. The Besov space may contain complicated functions including discontinuous function if $d/p \ge s$. For instance, the Cantor function f_1 does not belong to any Sobolev space since it does not have a weak derivative but belongs to a Besov space [Sawano, 2018]. In fact, $f_1 \in C^{\log 2/\log 3}([0,1]) \subset B^{\log 2/\log 3}_{\infty,\infty}([0,1])$. Also,

$$f_2(x) = \begin{cases} 1/\log(x/2), & 0 < x \le 1 \\ 0, & x = 0 \end{cases} \tag{6}$$

is not in any Hölder space. Note when $p < 2$, the smoothness of the Besov space is inhomogeneous [Suzuki, 2018]. Donoho and Johnstone [1994] suggested spatially variable functions including Blocks

$$\begin{gathered} f_3(x) = \sum_j h_j K(x - x_j),\ K(x) = (1 + \text{sgn}(x))/2, \\ (x_j) = (0.1, 0.13, 0.15, 0.23, 0.25, 0.40, 0.44, 0.65, 0.76, 0.78, 0.81), \\ (h_j) = (4, -5, 3, -4, 5, -4.2, 2.1, 4.3, -3.1, 2.1, -4.2), \end{gathered} \tag{7}$$

where $\text{sgn}(x)$ is a function that extracts the sign of a real number x, and HeaviSine

$$f_4(x) = 4\sin(4\pi x) - \text{sgn}(x - 0.3) - \text{sgn}(0.72 - x). \tag{8}$$

Since $f_2,\ f_3$ and f_4 are functions of bounded variation, in $B^1_{1,\infty}([0,1])$ [Peetre, 1976, Suzuki, 2018]. We plot the functions f_1, f_2, f_3 and f_4 in Figure 1.

2.4 Sparse neural network

Define (deep) neural network space generated by a parameter space Θ as follows:

$$\begin{gathered} \Phi(\Theta) = \Big\{ f_\theta(x) = (W^{(L+1)}(\cdot) + b^{(L+1)}) \circ \zeta \circ \cdots \circ \zeta \circ (W^{(1)}x + b^{(1)}) : \\ \theta = (W^{(1)}, b^{(1)}, \cdots, W^{(L+1)}, b^{(L+1)}) \in \Theta \Big\}. \end{gathered} \tag{9}$$

Figure 1: Example functions in the Besov spaces

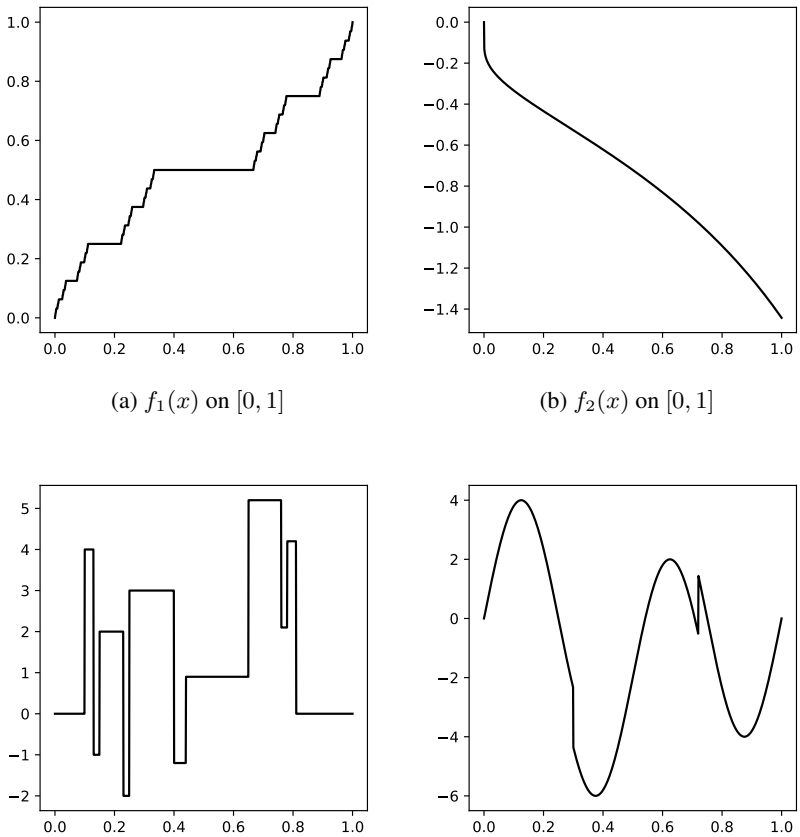

(a) $f_1(x)$ on $[0,1]$

(b) $f_2(x)$ on $[0,1]$

(c) $f_3(x)$ on $[0,1]$

(d) $f_4(x)$ on $[0,1]$

where ζ is an activation function. In this paper, we consider $\zeta(x) = ReLU(x)$ and the following B-bounded and S-sparse parameter space [Suzuki, 2018]

$$\begin{aligned}\Theta(L,W,S,B) := \Big\{ & W^{(l)} \in \mathbb{R}^{p_{l-1}\times p_l}, b^{(l)} \in \mathbb{R}^{p_l}, \\ & p_l = W, l = 1,2\cdots,L,\ p_0 = d, p_{L+1} = 1, \\ & \|\theta\|_0 = \sum_{l=1}^{L+1}\left(\left\|W^{(l)}\right\|_0 + \left\|b^{(l)}\right\|_0\right) \le S, \\ & \|\theta\|_\infty = \max_l \left\{\left\|W^{(l)}\right\|_\infty \vee \left\|b^{(l)}\right\|_\infty\right\} \le B\Big\}\end{aligned} \tag{10}$$

and denote S-sparse neural network space with B-bounded parameters $\Phi(\Theta(L,W,S,B))$ as $\Phi(L,W,S,B)$. We can define the entire neural network model space as follows

$$\Phi = \bigcup_{L=1}^{\infty}\bigcup_{W=1}^{\infty}\bigcup_{S=0}^{\infty}\bigcup_{B=0}^{\infty}\Phi(L,W,S,B). \tag{11}$$

3 Main results

Let $\mathcal{UB}(F) = \left\{f \in L^\infty([0,1]^d) : |f(x)| \le F,\ \forall x \in [0,1]^d\right\}$ be an uniformly bounded family of functions with bound $F > 0$. Consider $\mathscr{F} = B^s_{p,q}([0,1]^d) \cap \mathcal{UB}(F)$. For each $n \in \mathbb{N}$, and $f \in \mathcal{F}$,

let $P_f^{(n)}$ be the probability measure of the data $\mathbb{D}_n = (X_i, y_i)_{i=1}^n$ with a density function $p_f^{(n)}$ with respect to the measure $\mu^{(n)}$ and Π be a prior distribution of f. The posterior distribution is given by

$$\Pi(A|\mathbb{D}_n) = \frac{\int_A p_f^{(n)}(\mathbb{D}_n) d\Pi(f)}{\int_{\mathcal{F}} p_f^{(n)}(\mathbb{D}_n) d\Pi(f)}, \tag{12}$$

for $A \subset \Theta$. Let

$$\|f\|_n = \left(\frac{1}{n} \sum_{i=1}^n (f(X_i))^2 \right)^{1/2}$$

be the norm on empirical measure $\mathbb{P}_n^X = \frac{1}{n} \sum_{i=1}^n \delta_{X_i}$ and

$$\|f\|_{L^2(P_X)} = \left(\int (f(X))^2 \, dP_X \right)^{1/2},$$

where P_X is a marginal distribution of X. Suppose that P_X has a bounded density function $p_X(x)$ on $[0,1]^d$. Without loss of generality, we assume that $0 \le p_X(x) \le R \le 2$.

Suppose that $0 < F < \infty$, $0 < p, q \le \infty$, $\omega := d(1/p - 1/2)_+ < s < \infty$ and set $\nu = (s - \omega)/(2\omega)$. Assume that $m \in \mathbb{N}$ satisfies $0 < s < \min\{m, m - 1 + 1/p\}$. Let $N_n = \lceil n^{d/(2s+d)} \rceil$, $W_0 = 6dm(m+2) + 2d$ and

$$\begin{aligned} L_n = L(N_n), \quad W_n = N_n W_0, \\ S_n = (L_n - 1) W_0^2 N_n + N_n, \quad B_n = O(N_n^{\Xi}), \end{aligned} \tag{13}$$

where $c_{(d,m)} = \left(1 + 2de\frac{(2e)^m}{\sqrt{m}}\right)^{-1}$, $L(N_n) = 3 + 2\lceil \log_2 \left(\frac{3^{d \vee m}}{\tau_n c_{(d,m)}} \right) + 5 \rceil \lceil \log_2(d \vee m) \rceil$, $\tau_n = N_n^{-s/d - (\nu^{-1} + d^{-1})(d/p - s)_+} (\log N_n)^{-1}$ and $\Xi = (\nu^{-1} + d^{-1})(1 \vee (d/p - s)_+)$.

Note the theorems in this section are similar to the results in Suzuki [2018], which studied the asymptotic properties of the least squared estimator. They showed that for the least squared estimator

$$\begin{aligned} \hat{f} = \underset{\bar{f}: f \in \Phi(L_n, W_n, S_n, B_n)}{\arg\min} \sum_{i=1}^n (y_i - \bar{f}(x_i))^2, \\ \mathbb{E}_{\mathbb{D}_n} \left[\left\| f_0 - \hat{f} \right\|_{L^2(P_X)} \right] \lesssim n^{-\frac{2s}{2s+d}} (\log n)^3, \end{aligned} \tag{14}$$

where $\bar{f}$ is the clipping of f defined by $\bar{f} = \min\{\max\{f, -F\}, F\}$ for $f_0 \in \mathcal{UB}(F)$ and $\mathbb{E}_{\mathbb{D}_n}[A]$ is an expectation of A with respect to the training data $\mathbb{D}_n$. Note the minimax convergence rate of a Besov function f_0 in (1) has a lower bound $n^{-\frac{s}{2s+d}}$ [Gine and Nickl, 2016, Donoho and Johnstone, 1998]. Donoho and Johnstone [1998] also showed that the minimax rate of a linear estimator has a lower bound, $n^{-\frac{s-(1/p-1/2)_+}{2s+1-2(1/p-1/2)_+}}$ for $d = 1$, $s > 1/p$ and $1 \le p, q \le \infty$ or $s = p = q = 1$. Thus, any linear estimator cannot achieve the minimax optimal rate on the Besov space for $p < 2$ and is dominated by the neural network model.

Spike-and-slab prior

This result is an extension of Theorem 6.1 in Polson and Ročková [2018]. They showed that the Bayesian ReLU networks are consistent with nearly minimax rates in the Hölder space. We extend the function space from the Hölder space to the Besov space which contains the Hölder and Sobolev spaces. Assume the following spike-and-slab prior $\pi(\theta)$

$$\begin{aligned} \pi(\theta_j | \gamma_j, L, W, S, B) = \gamma_j \tilde{\pi}(\theta_j | L, W, S, B) + (1 - \gamma_j) \delta_0(\theta_j), \\ \pi(\boldsymbol{\gamma} | L, W, S, B) = \frac{1}{\binom{T}{S}}, \end{aligned} \tag{15}$$

$$\pi(L = L_n) = \pi(W = W_n) = \pi(S = S_n) = \pi(B = B_n) = 1, \tag{16}$$

where $\tilde{\pi}(\theta_j | L, W, S, B) = U(\theta_j; [-B, B])$ and $T = |\Theta(L, W, S, B)|$.

Theorem 1. *Under model (1) and prior distribution (15) and (16). Suppose that* $0 < F < \infty,\ 0 < p, q \leq \infty$ *and* $d(1/p - 1/2)_+ < s$*. Then the posterior distribution concentrates at the true function with a rate* $\epsilon_n = n^{-s/(2s+d)}(\log n)^{3/2}$*. That is,*

$$\Pi(f_\theta \in \Phi \cap \mathcal{UB}(F) : \|f_\theta - f_0\|_n > M_n \epsilon_n | \mathbb{D}_n) \to 0$$

in $P_{f_0}^{(n)}$*-probability as* $n \to \infty$ *for any* $M_n \to \infty$*.*

Proof. See appendix C.4. □

As mentioned in Suzuki [2018], the condition $d(1/p - 1/2)_+ < s$ indicates that $f_0 \in B^s_{p,q}([0,1]^d)$ is in $L^2([0,1]^d)$ and can be discontinuous. When $p = q = \infty$, $B^s_{p,q}([0,1]^d) = C^s([0,1]^d)$ and the result by Polson and Ročková [2018] is a special case of Theorem 1.

Adaptive estimation

This result is an extension of Theorem 6.2 in Polson and Ročková [2018] to the Besov space. They showed that the Bayesian ReLU networks adapt to smoothness on the Hölder space under the spike-and-slab prior.

Theorem 2. *Assume model (1). Let*

$$\tilde{L}_n(H) = \lceil H(\log n) \rceil \vee 1,\ \tilde{W}_n(H, N) = HN,\ \tilde{S}_n(H, N) = HN\tilde{L}_n(H),\ \tilde{B}_n(H, N) = N^H. \tag{17}$$

Consider the following prior

$$N \overset{d}{\equiv} 1 \vee \lceil Z/(\log n)^2 \rceil, \quad \pi_Z(Z) = \frac{\lambda_N^Z}{Z!(e^{\lambda_N} - 1)} \quad \text{for } Z = 1, 2, \cdots, \tag{18}$$

and prior distribution (15) of θ *given* $\tilde{L}_n(H_n),\ \tilde{W}_n(H_n, N),\ \tilde{S}_n(H_n, N),\ \tilde{B}_n(H_n, N)$ *on the function space*

$$\Phi = \bigcup_{n=1}^{\infty} \bigcup_{N=1}^{\infty} \Phi(\tilde{L}_n(H_n),\ \tilde{W}_n(H_n, N),\ \tilde{S}_n(H_n, N),\ \tilde{B}_n(H_n, N)) \tag{19}$$

for any $H_n \to \infty$ *and* $\lambda_N > 0$*. Suppose that* $0 < F < \infty,\ 0 < p, q \leq \infty$ *and* $d(1/p - 1/2)_+ < s < \min\{m, m - 1 + 1/p\}$*. The posterior distribution concentrates at the true function with a rate* $\epsilon_n = n^{-s/(2s+d)}(\log n)^{3/2}$*. That is,*

$$\Pi\left(f_\theta \in \Phi \cap \mathcal{UB}(F) : \|f_\theta - f_0\|_n > M_n \epsilon_n | \mathbb{D}_n\right) \to 0$$

in $P_{f_0}^{(n)}$*-probability as* $n \to \infty$ *for any* $M_n \to \infty$*.*

Proof. See appendix C.5. □

In Theorem 1, the prior distribution depends on the smoothness parameters p, q, and s. Theorem 2 shows that, as in Polson and Ročková [2018], a similar result can be obtained by considering the appropriate prior distribution even if the values of these parameters are unknown. Note it can be easily shown that BNNs do not overfit as in Polson and Ročková [2018].

Shrinkage prior

Asymptotic properties of the spike-and-slab prior can be obtained using Theorem 1, but its practical use for large neural network models is unsuitable due to the computational cost. The following prior distributions are more practical than a spike-and-slab prior in terms of computation.

The result to be introduced in this section shows that all the aforementioned results still hold even under the shrinkage prior. Assume the following prior

$$\pi(\theta | L, W, S, B) = \prod_{j=1}^{T} g(\theta_j | L, W, S, B), \tag{20}$$

where $T = |\Theta(L, W, S, B)|$, $g(t) := g(t|L, W, S, B)$ is a continuous, symmetric density function and decreasing on $t > 0$.

Theorem 3. *Assume model (1), prior distribution (16) and (20). Suppose that* $0 < F < \infty$, $0 < p, q \leq \infty$ *and* $d(1/p - 1/2)_+ < s < \min\{m, m - 1 + 1/p\}$. *Let* $\epsilon_n = n^{-s/(2s+d)}(\log n)^{3/2}$ *and* $g(t)$ *be a function such that*

$$\begin{aligned} a_n &\leq \frac{\epsilon_n}{72L_n(B_n \vee 1)^{L_n-1}(W_n+1)^{L_n}} \\ u_n &= \int_{[-a_n,a_n]} g(t|L_n, W_n, S_n, B_n)dt \end{aligned} \tag{21}$$

satisfies

$$\frac{S_n}{T_n} > 1 - u_n \geq \frac{S_n}{T_n}\eta_n, \tag{22}$$

$$-\log g(B_n|L_n, W_n, S_n, B_n) \lesssim (\log n)^2, \tag{23}$$

and

$$v_n = \int_{[-B_n,B_n]^c} g(t|L_n, W_n, S_n, B_n)dt = o\left(e^{-K_0 n\epsilon_n^2}\right), \tag{24}$$

where $\eta_n = \exp(-Kn\epsilon_n^2/S_n)$ *for some* $K,\ K_0 > 4$. *The posterior distribution concentrates at the true function with a rate* ϵ_n. *That is,*

$$\Pi(f_\theta \in \Phi \cap \mathcal{UB}(F) : \|f_\theta - f_0\|_n > M_n\epsilon_n|\mathbb{D}_n) \to 0$$

in $P_{f_0}^{(n)}$*-probability as* $n \to \infty$ *for any* $M_n \to \infty$.

Proof. See appendix C.6. □

Note that condition (22) is a continuous approximation for the spike part of (15). The next condition (23) means that the prior should have enough thick tail to sample true function. The last condition (24) restricts the thickness of the tail to prevent a function from being divergent.

Unlike the other priors, the shrinkage prior is relatively straightforward to implement. For example, to implement the model in Theorem 1, one can consider the MCMC algorithm [Sun et al., 2022] which combines Gibbs sampler, the Metropolis-Hastings algorithm, and the SGHMC algorithm. The shrinkage prior avoids the posterior computation with varying dimensions, and enables feasible computation.

Note that the conditions are sufficient conditions for the neural network model necessary to estimate the function in the worst case. In other words, we may infer true function f_0 sufficiently even for a prior distribution that satisfies only some conditions.

Example 1 (Gaussian prior)**.** *Let* $\epsilon_n = n^{-s/(2s+d)}(\log n)^{3/2}$ *and* ψ *be a density functions of the standard Gaussian distribution* $\mathcal{N}(0,\ 1)$.

$$g(t) = \frac{1}{\sigma_n}\psi\left(\frac{t}{\sigma_n}\right), \quad \sigma_n^2 = \frac{B_n^2}{2K_0 n\epsilon_n^2}$$

satisfies (24).

Example 2 (Gaussian mixture prior)**.** *Let* $\epsilon_n = n^{-s/(2s+d)}(\log n)^{3/2}$, ψ *and* Ψ *be density and inverse survival functions of the standard Gaussian distribution* $\mathcal{N}(0,\ 1)$ *respectively. When* $B_n \asymp N_n^{\min\{1,\Xi\}}$,

$$g(t) = \pi_{1n}\frac{1}{\sigma_{1n}}\psi\left(\frac{t}{\sigma_{1n}}\right) + \pi_{2n}\frac{1}{\sigma_{2n}}\psi\left(\frac{t}{\sigma_{2n}}\right)$$

with

$$\pi_{2n} = \frac{S_n}{T_n},\ \pi_{1n} = 1 - \pi_{2n},$$

$$a_n \leq \frac{\epsilon_n}{72L_n(B_n \vee 1)^{L_n-1}(W_n+1)^{L_n}},\ \sigma_{2n}^2 = \frac{B_n^2}{2K_0 n\epsilon_n^2}$$

and

$$a_n/\Psi\left(\frac{\pi_{2n}}{\pi_{1n}}\left[\frac{1}{2}\eta_n - \Psi^{-1}\left(\frac{a_n}{\sigma_{2n}}\right)\right]\right) \leq \sigma_{1n} < a_n/\Psi\left(\frac{\pi_{2n}}{\pi_{1n}}\left[\frac{1}{2} - \Psi^{-1}\left(\frac{a_n}{\sigma_{2n}}\right)\right]\right)$$

satisfies (22) and (24).

4 Numerical examples

However, these results are still not enough practical to be used. As the scale of the data n increases, the complexity of the neural network to obtain the theoretical optimality increases rapidly. For instance, consider a problem of estimating the Besov functions $f_1,\ f_2,\ f_3$ and f_4. Assume the Gaussian mixture prior distribution for each parameter as in Example 2,

$$B_n = 10N_n^{\min\{\Xi,1\}},\ K_0 = 5,\ \eta_n = \exp(-K_0 n\epsilon_n^2/S_n)$$
$$a_n = \frac{\epsilon_n}{72L_n(B_n \vee 1)^{L_n-1}(W_n+1)^{L_n}}$$
$$\sigma_{2n}^2 = \frac{B_n^2}{2(K_0+1)n\epsilon_n^2},\ \sigma_{1n} = a_n/\Psi\left(\frac{\pi_{2n}}{\pi_{1n}}\left[\frac{1}{2}\eta_n - \Psi^{-1}\left(\frac{a_n}{\sigma_{2n}}\right)\right]\right).$$

In Table 1 and Table 2, the hyperparameters (model parameters) to estimate each function were presented. These values are calculated numerically and may vary depending on the computer environment. Its speed is affected by the complexity of the true regression function. As can be seen from the tables, the smaller variance is almost zero. Considering the numerical precision, it is indistinguishable from spike-and-slab prior [Ishwaran and Rao, 2005].

Table 1: Hyperparameters of BNN to estimate $f_1 \in B_{\infty,\infty}^{\log 2/\log 3}([0,1])$.

n	L	W	σ_{1n}	σ_{2n}	π_{1n}	π_{2n}
100	39	400	3.747×10^{-179}	8.443×10^{-1}	8.719×10^{-1}	1.280×10^{-1}
1,000	43	1,100	1.505×10^{-234}	7.597×10^{-1}	9.534×10^{-1}	4.651×10^{-2}
10,000	45	2,950	1.508×10^{-283}	7.954×10^{-1}	9.826×10^{-1}	1.733×10^{-2}

Table 2: Hyperparameters of BNN to estimate $f_2,\ f_3$ and $f_4 \in B_{1,1}^{1}([0,1])$.

n	L	W	σ_{1n}	σ_{2n}	π_{1n}	π_{2n}
100	41	250	8.669×10^{-172}	6.779×10^{-1}	7.957×10^{-1}	2.042×10^{-1}
1,000	43	500	1.259×10^{-205}	5.028×10^{-1}	8.977×10^{-1}	1.022×10^{-1}
10,000	47	1,100	1.963×10^{-256}	4.894×10^{-1}	9.535×10^{-1}	4.642×10^{-2}

In practice, even the smaller models can approximate the true function with a sufficient precision. We present the numerical experiments with the four functions in Appendix D.

5 Conclusions

In this paper, we justify the theoretical properties of neural networks that are currently widely used. Specifically, we focused on the asymptotic properties of the Bayesian neural networks and extend results in Polson and Ročková [2018] to the Besov space from the Hölder space. Theorem 1 shows that the Bayesian neural network with a spike-and-slab prior has consistency posterior with a nearly minimax rate at Besov functions with known smoothness. Theorem 3 extends the result to the shrinkage prior. According to these theorems, the Bayesian neural networks have posterior consistency and (nearly) optimal rates by assuming proper prior. We found there are three sufficient conditions to achieve posterior consistency, (1) enough mass around zero, (2) thick tail to sample true function (3) while not too much. Theorem 2 shows that even if the smoothness of the function is not known, the same rate can be obtained by considering a prior on the architecture of the neural networks.

The results in this paper are an extension of the results of Suzuki [2018] on the Bayesian neural networks also. We suggest a practical method that can achieve the optimal rate while providing the uncertainty of the prediction. Note they showed the theoretical properties of the neural networks in the Besov space via nonparametric regression based on B-spline. In other words, similar results can be obtained using the Bayesian LARK B-spline model [Park et al., 2021].

In future works, performance comparison experiments with other models will also be of great help in establishing the properties of the Bayesian neural network. Empirically, we know that there are

fields in which decision-tree-based models such as random forest or neural network models perform better than traditional statistical models that find smooth functional relationships such as kernel regression. For these data, it is natural that there would be a non-smooth relationship between the explanatory variable and the response variable. For this reason, it is essential to compare Bayesian neural network models and statistical models for future studies.

We obtained results under the limited assumption, setting the networks fully connected and using the ReLU activation function. Also, as mentioned in Section 4, the results are not enough practical to be used. In addition, we do not suggest a result for the adaptive version of Theorem 3 yet. It is a necessary future study to propose a general shrinkage prior distribution that satisfies all conditions even when the parameters of the network are varying. The results of this study are at a basic level compared to the needs for the Bayesian neural networks in the industrial field, such as visual data analysis and natural language processing. It is important to research the theoretical properties of the Bayesian neural networks used in the practical area through ongoing studies.

Acknowledgments

Kyeongwon Lee and Jaeyong Lee were supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government(MSIT) (No. 2018R1A2A3074973 and 2020R1A4A1018207).

References

R. Arratia and L. Gordon. Tutorial on large deviations for the binomial distribution. *Bulletin of mathematical biology*, 51(1):125–131, 1989.

M. Bianchini and F. Scarselli. On the complexity of neural network classifiers: A comparison between shallow and deep architectures. *IEEE transactions on neural networks and learning systems*, 25(8):1553–1565, 2014.

E. Bingham, J. P. Chen, M. Jankowiak, F. Obermeyer, N. Pradhan, T. Karaletsos, R. Singh, P. A. Szerlip, P. Horsfall, and N. D. Goodman. Pyro: Deep universal probabilistic programming. *J. Mach. Learn. Res.*, 20:28:1–28:6, 2019. URL `http://jmlr.org/papers/v20/18-403.html`.

C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra. Weight uncertainty in neural network. In *International Conference on Machine Learning*, pages 1613–1622. PMLR, 2015.

G. Casella and E. I. George. Explaining the gibbs sampler. *The American Statistician*, 46(3):167–174, 1992.

I. Castillo, J. Schmidt-Hieber, and A. Van der Vaart. Bayesian linear regression with sparse priors. *The Annals of Statistics*, 43(5):1986–2018, 2015.

T. Chen, E. Fox, and C. Guestrin. Stochastic gradient hamiltonian monte carlo. In *International conference on machine learning*, pages 1683–1691. PMLR, 2014.

B.-E. Chérief-Abdellatif. Convergence rates of variational inference in sparse deep learning. In *International Conference on Machine Learning*, pages 1831–1842. PMLR, 2020.

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

O. Delalleau and Y. Bengio. Shallow vs. deep sum-product networks. *Advances in neural information processing systems*, 24:666–674, 2011.

D. L. Donoho and I. M. Johnstone. Minimax estimation via wavelet shrinkage. *The annals of Statistics*, 26(3):879–921, 1998.

D. L. Donoho and J. M. Johnstone. Ideal spatial adaptation by wavelet shrinkage. *biometrika*, 81(3): 425–455, 1994.

Y. Gal and Z. Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059. PMLR, 2016.

S. Geman and D. Geman. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE Transactions on pattern analysis and machine intelligence*, (6):721–741, 1984.

S. Ghosal and A. Van Der Vaart. Convergence rates of posterior distributions for noniid observations. *The Annals of Statistics*, 35(1):192–223, 2007.

E. Gine and R. Nickl. *Mathematical Foundations of Infinite-Dimensional Statistical Models*. Cambridge University Press, 2016. ISBN 978-1-107-33786-2. doi: 10.1017/CBO9781107337862. URL `http://ebooks.cambridge.org/ref/id/CBO9781107337862`.

I. Goodfellow, Y. Bengio, A. Courville, and Y. Bengio. *Deep learning*, volume 1. MIT press Cambridge, 2016.

A. Graves. Practical variational inference for neural networks. *Advances in neural information processing systems*, 24, 2011.

M. D. Hoffman, A. Gelman, et al. The no-u-turn sampler: adaptively setting path lengths in hamiltonian monte carlo. *J. Mach. Learn. Res.*, 15(1):1593–1623, 2014.

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

H. Ishwaran and J. S. Rao. Spike and slab variable selection: frequentist and bayesian strategies. *The Annals of Statistics*, 33(2):730–773, 2005.

D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

J. Lee, Y. Bahri, R. Novak, S. S. Schoenholz, J. Pennington, and J. Sohl-Dickstein. Deep neural networks as gaussian processes. *arXiv preprint arXiv:1711.00165*, 2017.

H. K. H. Lee III. *Model selection and model averaging for neural networks*. PhD thesis, Carnegie Mellon University, 1999.

A. G. d. G. Matthews, M. Rowland, J. Hron, R. E. Turner, and Z. Ghahramani. Gaussian process behaviour in wide deep neural networks. *arXiv preprint arXiv:1804.11271*, 2018.

V. Nair and G. E. Hinton. Rectified linear units improve restricted boltzmann machines. In *Icml*, 2010.

R. M. Neal. Priors for infinite networks. In *Bayesian Learning for Neural Networks*, pages 29–53. Springer, 1996.

R. M. Neal et al. Mcmc using hamiltonian dynamics. *Handbook of markov chain monte carlo*, 2 (11):2, 2011.

S. Park, H.-S. Oh, and J. Lee. Lévy adaptive b-spline regression via overcomplete systems. *arXiv preprint arXiv:2101.12179*, 2021.

A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. URL `http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf`.

J. Peetre. *New thoughts on Besov spaces*. Number 1. Mathematics Department, Duke University, 1976.

N. G. Polson and V. Ročková. Posterior concentration for sparse deep learning. *Advances in Neural Information Processing Systems*, 31, 2018.

Y. Sawano. *Theory of Besov Spaces*, volume 56 of *Developments in Mathematics*. Springer Singapore, 2018. ISBN 9789811308352. doi: 10.1007/978-981-13-0836-9. URL `http://link.springer.com/10.1007/978-981-13-0836-9`.

J. Schmidt-Hieber. Nonparametric regression using deep neural networks with relu activation function. *The Annals of Statistics*, 48(4):1875–1897, 2020.

Q. Song and F. Liang. Nearly optimal bayesian shrinkage for high dimensional regression. *arXiv preprint arXiv:1712.08964*, 2017.

Y. Sun, Q. Song, and F. Liang. Learning sparse deep neural networks with a spike-and-slab prior. *Statistics & Probability Letters*, 180:109246, 2022.

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*, 2018.

M. Telgarsky. Representation benefits of deep feedforward networks. *arXiv preprint arXiv:1509.08101*, 2015.

M. Telgarsky. Benefits of depth in neural networks. In *Conference on learning theory*, pages 1517–1539. PMLR, 2016.

M. Vladimirova, J. Verbeek, P. Mesejo, and J. Arbel. Understanding priors in bayesian neural networks at the unit level. In *International Conference on Machine Learning*, pages 6458–6467. PMLR, 2019.

M. Welling and Y. W. Teh. Bayesian learning via stochastic gradient langevin dynamics. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pages 681–688. Citeseer, 2011.

C. Williams. Computing with infinite networks. *Advances in neural information processing systems*, 9, 1996.

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

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] See Section 5.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes]
   (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] See Section 3.
   (b) Did you include complete proofs of all theoretical results? [Yes] See Section 3 and Appendix C.
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)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]

(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)? [Yes]

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? [Yes] See supplementary material.

   (b) Did you mention the license of the assets? [Yes]

   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]

   (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]

A Posterior consistency

Lemma 1 (Lemma 2 in Ghosal and Van Der Vaart [2007]). *Suppose that observation $X^{(n)} = (X_1, \cdots, X_n)$ of independent observations X_i. Assume that the distribution $P_{\theta,i}$ of the ith component X_i has a density $p_{\theta,i}$ relative to a σ-finite measure μ_i. Let $P_\theta^{(n)} = \bigotimes_{i=1}^n P_{\theta,i}$ be the product measures and average Hellinger distance*

$$d_n^2(\theta_0, \theta) = \frac{1}{n}\sum_{i=1}^n \int (\sqrt{p_{\theta_0,i}} - \sqrt{p_{\theta,i}})^2 d\mu_i.$$

Then there exists test ϕ_n such that $P_{\theta_0}^{(n)}\phi_n \le e^{-\frac{1}{2}nd_n^2(\theta_0,\theta_1)}$ and $P_\theta^{(n)}(1-\phi_n) \le e^{-\frac{1}{2}nd_n^2(\theta_0,\theta_1)}$ for all $\theta \in \Theta$ such that $d_n(\theta,\theta_1) \le \frac{1}{18}d_n(\theta_0,\theta_1)$.

Proof. See Ghosal and Van Der Vaart [2007]. □

Let define divergences

$$K(f,g) = \mathbb{E}^f\left[\log\frac{f(X)}{g(X)}\right] = \int f\log\frac{f}{g}d\mu, \tag{25}$$

$$V_k(f,g) = \mathbb{E}^f\left[\left|\log\frac{f(X)}{g(X)}\right|^k\right], \tag{26}$$

$$V_{k,0}(f,g) = \mathbb{E}^f\left[\left|\log\frac{f(X)}{g(X)} - K(f,g)\right|^k\right] \tag{27}$$

and set

$$B_n^*(\theta_0,\epsilon;k) = \left\{\theta\in\Theta : \frac{1}{n}\sum_{i=1}^n K(P_{\theta_0,i}, P_{\theta,i}) \le \epsilon^2,\ \frac{1}{n}\sum_{i=1}^n V_{k,0}(P_{\theta_0,i}, P_{\theta,i}) \le C_k\epsilon^k\right\}.$$

Here, the C_k is the constant satisfying

$$\mathbb{E}\Big[\big|\bar{X}_n - \mathbb{E}[\bar{X}_n]\big|^k\Big] \le C_k n^{-k/2}\frac{1}{n}\sum_{i=1}^n \mathbb{E}\Big[|X_i|^k\Big]$$

for $k \ge 2$. The existence of C_k is guaranteed by Marcinkiewiz-Zygmund inequality. The following lemma gives a sufficient condition for obtaining posterior consistency.

Lemma 2 (Theorem 4 of Ghosal and Van Der Vaart [2007]). *Let $P_\theta^{(n)}$ be product measures and $d_n(\theta_0,\theta) = \frac{1}{n}\sum_{i=1}^n\int(\sqrt{p_{\theta_0,i}} - \sqrt{p_{\theta,i}})^2 d\mu_i$. Suppose that for a sequence $\epsilon_n \to 0$ such that $n\epsilon_n^2$ is bounded away from zero, some $k > 1$, all sufficiently large j and sets $\Theta_n \subset \Theta$ which satisfies following conditions:*

$$\sup_{\epsilon>\epsilon_n}\log N(\epsilon/36, \{\theta\in\Theta_n : d_n(\theta,\theta_0) < \epsilon\}, d_n) \le n\epsilon_n^2, \tag{28}$$

$$\frac{\Pi(\Theta - \Theta_n)}{\Pi(B_n^*(\theta_0,\epsilon_n;k))} = o\left(e^{-2n\epsilon_n^2}\right), \tag{29}$$

$$\frac{\Pi(\theta\in\Theta_n : j\epsilon_n < d_n(\theta,\theta_0) \le 2j\epsilon_n)}{\Pi(B_n^*(\theta_0,\epsilon_n;k))} \le e^{n\epsilon_n^2 j^2/4} \tag{30}$$

Then $P_{\theta_0}^{(n)}\left(\Pi(\theta : d_n(\theta,\theta_0) \ge M_n\epsilon_n|\mathbb{D}_n)\right) \to 0$ for any sequence $M_n \to \infty$.

Proof. See Ghosal and Van Der Vaart [2007]. □

In Ghosal and Van Der Vaart [2007], they mentioned that it can be replaced by any other distance d_n for which the conclusion of Lemma 1 holds. Moreover, for $k = 2$, Lemma 2 work with the smaller neighborhood

$$\overline{B}_n(\theta_0,\epsilon) = \left\{\theta\in\Theta : \frac{1}{n}\sum_{i=1}^n K(P_{\theta_0,i}, P_{\theta,i}) \le \epsilon^2,\ \frac{1}{n}\sum_{i=1}^n V_{2,0}(P_{\theta_0,i}, P_{\theta,i}) \le \epsilon^2\right\} \tag{31}$$

instead of $B_n^*(\theta_0, \epsilon_n; k)$.

The following is an adapted version of Lemma 2, which was introduced in Polson and Ročková [2018].

Lemma 3. *Assume model (1). Suppose that $\mathcal{F}$ is uniformly bounded. Let*

$$A_{\epsilon,M} := \{f \in \mathcal{F} : \|f - f_0\|_n \leq M\epsilon\}.$$

If there exist $C > 2/\sigma^2$ and $\mathcal{F}_n \subset \mathcal{F}$ such that

$$\sup_{\epsilon > \epsilon_n} \log N(\epsilon/36, A_{\epsilon,1} \cap \mathcal{F}_n, \|\cdot\|_n) \leq n\epsilon_n^2, \tag{32}$$

$$\Pi(A_{\epsilon_n,1}) \geq e^{-Cn\epsilon_n^2}, \tag{33}$$

$$\Pi(\mathcal{F} - \mathcal{F}_n) = o\left(e^{-(C\sigma^2+2)n\epsilon_n^2}\right) \tag{34}$$

for any $\epsilon_n \to 0$, $n\epsilon_n^2 \to \infty$,

$$\Pi(A^c_{\epsilon_n,M_n} | \mathbb{D}_n) \to 0$$

in $P_{f_0}^{(n)}$-probability as $n \to \infty$ for any $M_n \to \infty$.

Proof. See appendix C.1. □

B Consistency of neural network

The upper bound of the covering number of the neural network space is given by as follows.

Lemma 4 (Lemma 3 in Suzuki [2018])**.** $\forall\, \epsilon > 0$,

$$\log N(\epsilon, \Phi(L, W, S, B), \|\cdot\|_{L^\infty}) \leq (S+1)\log\left(2\epsilon^{-1}L(B \vee 1)^L (W+1)^{2L}\right) \tag{35}$$

for $W,\ L \geq 3$.

Proof. See appendix C.2. □

For $L,\ W,\ S \in \mathbb{N}$ and $B,\ a > 0$, define a function space

$$\Phi(L, W, S, B, a) = \{f_\theta : \theta \in \Theta(L, W, S, B, a)\}, \tag{36}$$

where

$$\Theta(L, W, S, B, a) = \left\{\theta : (\theta_i I(|\theta_i| > a))_{i=1}^{T_n} \in \Theta(L, W, S, B)\right\}.$$

Lemma 5. $\forall \epsilon \geq 2aL(B \vee 1)^{L-1}(W+1)^L$,

$$\log N(\epsilon, \Phi(L, W, S, B, a), \|\cdot\|_{L^\infty}) \leq (S+1)\log\left(2\epsilon^{-1}L(B \vee 1)^L (W+1)^{2L}\right) \tag{37}$$

for $W,\ L \geq 3$.

Proof. See appendix C.3. □

The following lemma shows that for any function f_0 in the Besov space, there exist neural networks which is closed enough to f_0.

Lemma 6 (Proposition 1 in Suzuki [2018])**.** *Suppose that $0 < p, q, r \leq \infty$, $\omega := d(1/p - 1/r)_+ < s < \infty$ and $\nu = (s - \omega)/(2\omega)$. Assume that $N \in \mathbb{N}$ is sufficiently large and $m \in \mathbb{N}$ satisfies $0 < s < \min\{m, m - 1 + 1/p\}$. Let $W_0 = 6dm(m+2) + 2d$. Then,*

$$\sup_{f_0 \in U(B^s_{p,q}([0,1]^d))} \inf_{f \in \Phi(L,W,S,B)} \|f_0 - f\|_{L^r} \lesssim N^{-s/d} \tag{38}$$

for

$$\begin{aligned} L &= 3 + 2\lceil \log_2\left(\frac{3^{d\vee m}}{\tau(N)c_{(d,m)}}\right) + 5\rceil \lceil \log_2(d \vee m)\rceil, \quad W = NW_0, \\ S &= (L-1)W_0^2 N + N, \quad B = O\left(N^{(\nu^{-1}+d^{-1})(1\vee(d/p-s)_+)}\right), \end{aligned} \tag{39}$$

where $U(\mathcal{H})$ is the unit ball of a quasi-Banach space $\mathcal{H}$, $c_{(d,m)} = \left(1 + 2de\frac{(2e)^m}{\sqrt{m}}\right)^{-1}$ and $\tau(N) = N^{-s/d-(\nu^{-1}+d^{-1})(d/p-s)_+}(\log N)^{-1}$.

Proof. See Suzuki [2018]. □

C Proof of theorems

C.1 Proof of Lemma 3

Proof. Let $\overline{d}_n(f_0, f) = \|f_0 - f\|_n$. As in Ghosal and Van Der Vaart [2007], page 214, we may use the norm $\overline{d}_n$ instead of the average Hellinger distance d_n and

$$K(P_{f_0,i}, P_{f,i}) = \frac{1}{2\sigma^2}|f_0(x_i) - f(x_i)|^2 = \frac{1}{2}V_{2,0}(P_{\theta_0,i}, P_{\theta,i}) \tag{40}$$

for all $i = 1, 2, \cdots, n$. Then,

$$\overline{B}_n(f_0, \epsilon) = \left\{ f \in \mathcal{F} : \overline{d}_n^2(f_0, f) \leq \sigma^2\epsilon^2 \right\} = A_{\sigma\epsilon_n,1} \tag{41}$$

and it is enough to show that

$$\frac{\Pi(\mathcal{F} - \mathcal{F}_n)}{\Pi(A_{\sigma\epsilon_n,1})} = o\left(e^{-2n\epsilon_n^2}\right), \tag{42}$$

$$\frac{\Pi(A_{2j\epsilon_n,1} - A_{j\epsilon_n,1})}{\Pi(A_{\sigma\epsilon_n,1})} \leq e^{n\epsilon_n^2 j^2/4} \tag{43}$$

for all sufficiently large j.

$$\frac{\Pi(\mathcal{F} - \mathcal{F}_n)}{\Pi(A_{\sigma\epsilon_n,1})} \leq e^{Cn\sigma^2\epsilon_n^2}\Pi(\mathcal{F} - \mathcal{F}_n) = o\left(e^{-2n\epsilon_n^2}\right) \tag{44}$$

and

$$\frac{\Pi(A_{\epsilon_n,2j} - A_{\epsilon_n,j})}{\Pi(A_{\sigma\epsilon_n,1})} \leq e^{Cn\sigma^2\epsilon_n^2}\Pi(A_{\epsilon_n,2j} - A_{\epsilon_n,j}) \leq e^{n\epsilon_n^2 j^2/4} \tag{45}$$

for all sufficiently large j. □

C.2 Proof of Lemma 4

Proof. Let denote $f \in \Phi(L, W, S, B)$ as

$$\begin{aligned} f(x) &= (W^{(L)}\zeta(\cdot) + b^{(L)}) \circ \cdots \circ (W^{(1)}x + b^{(1)}), \\ \mathcal{A}_k^+(f)(x) &= \zeta \circ (W^{(k-1)}\zeta(\cdot) + b^{(k-1)}) \circ \cdots \circ (W^{(1)}x + b^{(1)}), \\ \mathcal{A}_k^-(f)(x) &= (W^{(L)}\zeta(\cdot) + b^{(L)}) \circ \cdots \circ (W^{(k)}x + b^{(k)}) \end{aligned} \tag{46}$$

for $k = 2, \cdots, L$, and $\mathcal{A}_{L+1}^-(f)(x) = \mathcal{A}_1^+(f)(x) = x$. Then,

$$f(x) = \mathcal{A}_{k+1}^-(f) \circ (W^{(k)} \cdot + b^{(k)}) \circ \mathcal{A}_k^+(f)(x).$$

Since $f \in \Phi(L, W, S, B)$,

$$\begin{aligned} \left\|\mathcal{A}_k^+(f)(x)\right\|_\infty &\leq \max_j \left\|W_{j,:}^{(k-1)}\right\|_1 \left\|\mathcal{A}_{k-1}^+(f)(x)\right\|_\infty + \left\|b^{(k-1)}\right\|_\infty \\ &\leq WB\left\|\mathcal{A}_{k-1}^+(f)(x)\right\|_\infty + B \\ &\leq (W+1)(B \vee 1)\left\|\mathcal{A}_{k-1}^+(f)(x)\right\|_\infty \\ &\leq (W+1)^{k-1}(B \vee 1)^{k-1}, \end{aligned} \tag{47}$$

for all x, where $A_{j,:}$ is the j-th row of the matrix A. Similarly,

$$\left|\mathcal{A}_k^-(f)(x_1) - \mathcal{A}_k^-(f)(x_2)\right| \leq (BW)^{L-k+1}\|x_1 - x_2\|_\infty. \tag{48}$$

Fix $\epsilon > 0$ and $\theta \in \Theta(L, W, S, B)$. For any $\theta^* \in \Theta(L, W, S, B)$ which satisfies $\|\theta - \theta^*\|_\infty < \epsilon$,

$$
\begin{aligned}
&|f_\theta(x) - f_{\theta^*}(x)| \\
&= \left| \sum_{k=1}^{L} \mathcal{A}_{k+1}^{-}(f_{\theta^*}) \circ (W^{(k)} \cdot + b^{(k)}) \circ \mathcal{A}_k^{+}(f_\theta)(x) - \mathcal{A}_{k+1}^{-}(f_{\theta^*}) \circ (W^{(k)^*} \cdot + b^{(k)^*}) \circ \mathcal{A}_k^{+}(f_\theta)(x) \right| \\
&\leq \sum_{k=1}^{L} (BW)^{L-k} \left\| (W^{(k)} \cdot + b^{(k)}) \circ \mathcal{A}_k^{+}(f_\theta)(x) - (W^{(k)^*} \cdot + b^{(k)^*}) \circ \mathcal{A}_k^{+}(f_\theta)(x) \right\|_{L^\infty} \\
&\leq \sum_{k=1}^{L} (BW)^{L-k} \epsilon \left[W(B \vee 1)^{k-1}(W+1)^{k-1} + 1 \right] \\
&\leq \sum_{k=1}^{L} (BW)^{L-k} \epsilon (B \vee 1)^{k-1}(W+1)^{k} \\
&\leq \epsilon L (B \vee 1)^{L-1}(W+1)^{L}.
\end{aligned} \tag{49}
$$

Let s be the number of nonzero components of θ, $s \leq S$. Consider a subspace $\Theta_\theta(L, W, S, B)$ of the parameter space $\Theta(L, W, S, B)$ consists of parameters which have s nonzero component. Choose $\theta_1, \cdots, \theta_M$ from each grid divided with length $\epsilon^* = \frac{\epsilon}{L(B \vee 1)^{L-1}(W+1)^L}$ over $\Theta_\theta(L, W, S, B)$. Then

$$
f_\theta \in \bigcup_{m=1}^{M} \{ f : \|f - f_{\theta_m}\|_{L^\infty} < \epsilon \}.
$$

Note

$$
T = |\Theta(L, W, S, B)| \leq \sum_{l=1}^{L} (W+1)W \leq (L+1)(W+1)^2 \leq (W+1)^L
$$

for $W,\ L \geq 3$, and the number of cases of choose s nonzero components are

$$
\binom{T}{s} = \frac{T(T-1)\cdots(T-s+1)}{s!} \leq T^s \leq (W+1)^{Ls}.
$$

Thus,

$$
\begin{aligned}
N(\epsilon, \Phi(L, W, S, B), \|\cdot\|_{L^\infty}) &\leq \sum_{s^* \leq S} \binom{T}{s^*} (2B\epsilon^{-1}(L(B \vee 1)^{L-1}(W+1)^L)^{s^*} \\
&\leq \sum_{s^* \leq S} (2\epsilon^{-1} L (B \vee 1)^L (W+1)^{2L})^{s^*} \\
&\leq (2\epsilon^{-1} L (B \vee 1)^L (W+1)^{2L})^{S+1}
\end{aligned} \tag{50}
$$

and we get desired results by taking the logarithm of both sides. □

C.3 Proof of Lemma 5

Proof. Let $\tilde{\theta}(a) = (\theta_i I(|\theta_i| > a))$. Then $\left\| \theta - \tilde{\theta}(a) \right\|_\infty \leq a$ and

$$
\tilde{\Theta}(L, W, S, B, a) = \left\{ \tilde{\theta}(a) : \theta \in \Theta(L, W, S, B, a) \right\} = \Theta(L, W, S, B).
$$

As in the proof of Lemma 4,

$$
\left\| f_\theta - f_{\tilde{\theta}(a)} \right\|_{L^\infty} \leq aL(B \vee 1)^{L-1}(W+1)^L \leq \epsilon/2 \tag{51}
$$

for any $f_\theta,\ f_{\tilde{\theta}(a)} \in \Phi(L, W, S, B, a, F)$. Let $s \leq S$ be the number of nonzero components of $\tilde{\theta}(a)$. Consider a subspace $\tilde{\Theta}_\theta(L, W, S, B, a)$ of $\tilde{\Theta}(L, W, S, B, a)$ consisting of parameters which have s nonzero component. Choose $\theta_1, \cdots, \theta_M$ from each grid divided with length $\epsilon^* = \frac{\epsilon/2}{L(B \vee 1)^{L-1}(W+1)^L}$ over $\tilde{\Theta}_\theta(L, W, S, B, a)$. Then

$$
f_\theta \in \bigcup_{m=1}^{M} \{ f : \|f - f_{\theta_m}\|_{L^\infty} < \epsilon \}
$$

from triangular inequality. We get desired results in a similar way as in the proof of Lemma 4. □

C.4 Proof of Theorem 1

Proof. Let $\mathcal{F} = \Phi \cap \mathcal{UB}(F)$ and

$$A_{\epsilon,M} := \{f \in \mathcal{F} : \|f - f_0\|_n \le M\epsilon\}.$$

From Lemma 3, it is enough to show that there exist a constant $C > 2/\sigma^2$ and $\mathcal{F}_n \subset \mathcal{F}$ which satisfy

(a) $\sup_{\epsilon > \epsilon_n} \log N(\epsilon/36, A_{\epsilon,1} \cap \mathcal{F}_n, \|\cdot\|_n) \le n\epsilon_n^2$

(b) $-\log \Pi(A_{\epsilon_n,1}) \le Cn\epsilon_n^2$

(c) $\Pi(\mathcal{F} - \mathcal{F}_n) = o\left(e^{-(C\sigma^2+2)n\epsilon_n^2}\right)$

for sufficiently large n. Let $\mathcal{F}_n = \Phi(L_n, W_n, S_n, B_n) \cap \mathcal{UB}(F)$. First, (c) is trivial from (16). From

$$\{f \in \mathcal{F}_n : \|f\|_{L^\infty} \le \epsilon\} \subset \{f \in \mathcal{F}_n : \|f\|_n \le \epsilon\}$$

and Lemma 4,

$$\begin{aligned}
&\sup_{\epsilon > \epsilon_n} \log N\left(\frac{\epsilon}{36}, A_{\epsilon,1} \cap \mathcal{F}_n, \|\cdot\|_n\right) \\
&\le \sup_{\epsilon > \epsilon_n} \log N\left(\frac{\epsilon}{36}, A_{\epsilon,1} \cap \mathcal{F}_n, \|\cdot\|_{L^\infty}\right) \\
&\le \sup_{\epsilon > \epsilon_n} \log N\left(\frac{\epsilon}{36}, \mathcal{F}_n, \|\cdot\|_{L^\infty}\right) \\
&\le \log N\left(\frac{\epsilon_n}{36}, \mathcal{F}_n, \|\cdot\|_{L^\infty}\right) \\
&\le (S_n + 1)\left[\log L_n + L_n \log\big((B_n \vee 1)(W_n + 1)^2\big) - \log \frac{\epsilon_n}{72}\right] \\
&\lesssim N_n (\log n)^3 \\
&= n\epsilon_n^2
\end{aligned} \tag{52}$$

for sufficiently large n. Thus, (a) holds. Here, the last inequality holds from

$$L_n = O(\log n),\ W_n = O(N_n),\ S_n = O(N_n \log n) \tag{53}$$

Next, from Lemma 6, there is a constant $C_1 > 0$ and $\hat{f}_n = f_{\hat{\theta}} \in \mathcal{F}$ such that

$$\left\|\hat{f}_n - f_0\right\|_{L^2} \le C_1 \|f_0\|_{B^s_{p,q}([0,1]^d)} N_n^{-s/d} \le \epsilon_n/4. \tag{54}$$

for sufficiently large n. Moreover, by the assumptions and the strong law of large numbers,

$$\|f - f_0\|_n^2 \le 2\|f - f_0\|^2_{L^2(P_X)} \le 4\|f - f_0\|^2_{L^2} \tag{55}$$

for sufficiently large n $P_{f_0}^{(n)}$-almost surely. Let $\hat{\gamma}$ and $\hat{\theta}_{\hat{\gamma}}$ be index and value of nonzero components of $\hat{\theta}$ respectively. Let $\Theta(\hat{\gamma}; L_n, W_n, S_n, B_n)$ be a subset of parameter space $\Theta(L_n, W_n, S_n, B_n)$ consists of parameters which have nonzero components at $\hat{\gamma}$ only and $\mathcal{F}_n(\hat{\gamma}) = \Phi(\hat{\gamma}; L_n, W_n, S_n, B_n) \cap \mathcal{UB}(F)$ be an uniformly bounded neural network space generated by $\Theta(\hat{\gamma}; L_n, W_n, S_n, B_n)$. Note

$$\begin{aligned}
\Pi(A_{\epsilon_n,1}) &= \Pi(f \in \mathcal{F}_n : \|f - f_0\|_n \le \epsilon_n) \\
&\ge \Pi(f \in \mathcal{F}_n : \|f - f_0\|_{L^2} \le \epsilon_n/2) \\
&\ge \Pi\left(f \in \mathcal{F}_n : \left\|f - \hat{f}_n\right\|_{L^2} \le \epsilon_n/4\right) \\
&\ge \Pi\left(f \in \mathcal{F}_n : \left\|f - \hat{f}_n\right\|_{L^\infty} \le \epsilon_n/4\right) \\
&\ge \Pi\left(f \in \mathcal{F}_n(\hat{\gamma}) : \left\|f - \hat{f}_n\right\|_{L^\infty} \le \epsilon_n/4\right)
\end{aligned} \tag{56}$$

for sufficiently large n. As in the proof of Lemma 4,

$$
\begin{aligned}
&\Pi\left(f \in \mathcal{F}_n(\hat{\gamma}) : \left\|f - \hat{f}_n\right\|_{L^\infty} \le \epsilon_n/4\right) \\
&\ge \Pi\left(\theta \in \mathbb{R}^{T_n} : \theta_{\hat{\gamma}^c} = 0,\ \|\theta\|_\infty \le B_n,\ \left\|\hat{\theta} - \theta\right\|_\infty \le \frac{\epsilon_n}{4(W_n+1)^{L_n} L_n (B_n \vee 1)^{L_n - 1}}\right) \\
&\ge \left(\frac{\epsilon_n}{4B_n(W_n+1)^{L_n} L_n (B_n \vee 1)^{L_n-1}}\right)^{S_n} \binom{T_n}{S_n}^{-1} \\
&\ge \left(\frac{\epsilon_n}{4B_n(W_n+1)^{2L_n} L_n (B_n \vee 1)^{L_n-1}}\right)^{S_n} \\
&= \exp\left(-S_n \log\left(\frac{4B_n(W_n+1)^{2L_n} L_n (B_n \vee 1)^{L_n-1}}{\epsilon_n}\right)\right)
\end{aligned}
\tag{57}
$$

Thus,

$$
\begin{aligned}
-\log \Pi(A_{\epsilon_n,1}) &\le S_n \log\left(\frac{4B_n(W_n+1)^{2L_n} L_n (B_n \vee 1)^{L_n-1}}{\epsilon_n}\right) \\
&\le S_n \left[L_n \log\big((W_n+1)^2 (B_n \vee 1)\big) + \log 2L_n - \log \epsilon_n\right] \\
&\lesssim N_n (\log n)^3 \\
&= n\epsilon_n^2
\end{aligned}
\tag{58}
$$

for sufficiently large n. □

C.5 Proof of Theorem 2

Proof. Let $\mathcal{F} = \Phi \cap \mathcal{UB}(F)$. From Lemma 3, it is enough to show that there exist a constant $C' > 2/\sigma^2$ and $\mathcal{F}_n \subset \mathcal{F}$ which satisfy

(a) $\sup_{\epsilon > \epsilon_n} \log N(\epsilon/36, A_{\epsilon,1} \cap \mathcal{F}_n, \|\cdot\|_n) \le n\epsilon_n^2$

(b) $-\log \Pi(A_{\epsilon_n,1}) \le C' n \epsilon_n^2$

(c) $\Pi(\mathcal{F} - \mathcal{F}_n) = o\left(e^{-(C'\sigma^2+2)n\epsilon_n^2}\right)$

for sufficiently large n. From (53), we can choose

$$
H_0 > \sup_n \{L_n/\log n,\ W_n/N_n,\ S_n/(N_n \log n),\ \Xi\}. \tag{59}
$$

Let $\tilde{N}_n = C_N N_n$,

$$
\mathcal{F}_n = \mathcal{UB}(F) \cap \left(\bigcup_{N=1}^{\tilde{N}_n} \Phi(\tilde{L}_n(H_0),\ \tilde{W}_n(H_0,N),\ \tilde{S}_n(H_0,N),\ \tilde{B}_n(H_0,N))\right)
$$

for sufficiently large $C_N > 0$ and $\pi_N(N)$ be a density function of N. First, show that (a) holds. From Lemma 4,

$$
\begin{aligned}
&N\left(\frac{\epsilon_n}{36}, \mathcal{F}_n, \|\cdot\|_{L^\infty}\right) \\
&\le \sum_{N=1}^{\tilde{N}_n} \left(\frac{72}{\epsilon_n} \tilde{L}_n(H_0) (\tilde{B}_n(H_0,N) \vee 1)^{\tilde{L}_n(H_0)} (\tilde{W}_n(H_0,N)+1)^{2\tilde{L}_n(H_0)}\right)^{\tilde{S}_n(H,N)+1} \\
&\le \tilde{N}_n \left(\frac{72}{\epsilon_n} \tilde{L}_n(H_0) (\tilde{B}_n(H_0,\tilde{N}_n) \vee 1)^{\tilde{L}_n(H_0)} (\tilde{W}_n(H_0,\tilde{N}_n)+1)^{2\tilde{L}_n(H_0)}\right)^{\tilde{S}_n(H_0,\tilde{N}_n)+1}
\end{aligned}
\tag{60}
$$

and

$$
\begin{aligned}
&\sup_{\epsilon>\epsilon_n} \log N\left(\frac{\epsilon}{36}, A_{\epsilon,1}\cap\mathcal{F}_n, \|\cdot\|_n\right)\\
&\le \sup_{\epsilon>\epsilon_n} \log N\left(\frac{\epsilon}{36}, A_{\epsilon,1}\cap\mathcal{F}_n, \|\cdot\|_{L^\infty}\right)\\
&\le \sup_{\epsilon>\epsilon_n} \log N\left(\frac{\epsilon}{36}, \mathcal{F}_n, \|\cdot\|_{L^\infty}\right)\\
&\le \log N\left(\frac{\epsilon_n}{36}, \mathcal{F}_n, \|\cdot\|_{L^\infty}\right)\\
&\le \log \tilde{N}_n\\
&\quad + \left[\tilde{S}_n(H_0,\tilde{N}_n)+1\right]\log\left(\frac{72}{\epsilon_n}\tilde{L}_n(H_0)(\tilde{B}_n(H_0,\tilde{N}_n)\vee 1)^{\tilde{L}_n(H_0)}(\tilde{W}_n(H_0,\tilde{N}_n)+1)^{2\tilde{L}_n(H_0)}\right)\\
&\lesssim \tilde{N}_n(\log n)^3\\
&\lesssim n\epsilon_n^2.
\end{aligned}
\tag{61}
$$

for sufficiently large n. Next, show that (b) holds. Note $N_n(\log n)^3 \lesssim n\epsilon_n^2$ and

$$
L_n \le \tilde{L}_n(H_n),\ W_n \le \tilde{W}_n(H_n,N_n),\ S_n \le \tilde{S}_n(H_n,N_n),\ B_n \le \tilde{B}_n(H_n,\tilde{N}_n) \tag{62}
$$

for $N_n,\ L_n,\ W_n,\ S_n,\ B_n$ in Theorem 1 and sufficiently large n. Thus, there exists a constant $D>0$ such that

$$
\pi_N(N_n) \gtrsim \exp\left(-N_n(\log n)^2\log\frac{N_n}{\lambda_N}\right) \gtrsim \exp\left(-Dn\epsilon_n^2\right) \tag{63}
$$

and

$$
\begin{aligned}
&\Pi(f_\theta\in\mathcal{F}_n : \|f-f_0\|_n \le \epsilon_n)\\
&\ge \pi_N(N_n)\Pi(f_\theta\in\Phi(L_n,W_n,S_n,B_n) : \|f_\theta-f_0\|_n\le\epsilon_n|N_n)\\
&\gtrsim \exp\left(-(C+D)n\epsilon_n^2\right)
\end{aligned}
\tag{64}
$$

holds for sufficiently large n. (b) holds for $C' = \max\left\{C+D, 1+2/\sigma^2\right\}$. From

$$
\Pi(\mathcal{F}-\mathcal{F}_n) \le \pi_N(N>\tilde{N}_n)
$$

and *Chernoff bound*, for any positive number $t,\ Z_0>0$,

$$
P(Z>Z_0) < e^{-t(Z_0+1)}\mathbb{E}\left[e^{tZ}\right] \lesssim e^{-t(Z_0+1)}\left(\exp\left(e^t\lambda_N\right)-1\right). \tag{65}
$$

Letting $t=\log Z_0$, we get

$$
P(Z>Z_0) \lesssim e^{-(Z_0+1)\log Z_0}\left(\exp(Z_0\lambda_N)-1\right). \tag{66}
$$

Thus,

$$
\begin{aligned}
\pi_N(N>\tilde{N}_n) &\lesssim e^{-\left[(\tilde{N}_n+1)\log\tilde{N}_n+\tilde{N}_n\lambda_N\right](\log n)^2},\\
(C'\sigma^2+2)n\epsilon_n^2 + \lambda_N\tilde{N}_n(\log n)^2 - (\tilde{N}_n+1)\log\tilde{N}_n(\log n)^2 &\to -\infty
\end{aligned}
\tag{67}
$$

for sufficiently large $C_N>0$. (c) holds. □

C.6 Proof of Theorem 3

Proof. Let $\mathcal{F} = \Phi\cap\mathcal{UB}(F)$. From Lemma 3, it is enough to show that there exist a constant $C''>2/\sigma^2$ and $\mathcal{F}_n\subset\mathcal{F}$ which satisfy

(a) $\sup_{\epsilon>\epsilon_n}\log N(\epsilon/36, A_{\epsilon,1}\cap\mathcal{F}_n, \|\cdot\|_n) \le n\epsilon_n^2$

(b) $-\log\Pi(A_{\epsilon_n,1}) \le C''n\epsilon_n^2$

(c) $\Pi(\mathcal{F}-\mathcal{F}_n) = o\left(e^{-(C''\sigma^2+2)n\epsilon_n^2}\right)$

for sufficiently large n. Let $\mathcal{F}_n = \Phi(L_n, W_n, S_n, B_n, a_n) \cap \mathcal{UB}(F)$. It is easy to show that (a) holds from Lemma 5 as in the proof of Theorem 1. By the assumption,

$$
\begin{aligned}
&\Pi(\mathcal{F} - \mathcal{F}_n) \\
&\leq \pi\left(\exists |\theta_i| > B_n \middle| L_n, W_n, S_n, B_n\right) + \pi\left(\sum_{i=1}^{T_n} I(|\theta_i| > a_n) > S_n | L_n, W_n, S_n, B_n\right) \\
&= \left(1 - (1 - v_n)^{T_n}\right) + P\left(S > S_n | S \sim B(T_n, 1 - u_n)\right) \\
&\leq T_n v_n + \exp\left(-T_n \left\{\left(1 - \frac{S_n}{T_n}\right) \log \frac{1 - S_n/T_n}{u_n} + \frac{S_n}{T_n} \log \frac{S_n/T_n}{1 - u_n}\right\}\right) \\
&= o\left(e^{-K_0 n \epsilon_n^2 + \log T_n}\right) + \exp\left(T_n \left(1 - \frac{S_n}{T_n}\right) \log \frac{u_n}{1 - S_n/T_n} - S_n \log \frac{S_n/T_n}{1 - u_n}\right) \\
&\leq o\left(e^{-K_1 n \epsilon_n^2}\right) + \exp\left(T_n \left(1 - \frac{S_n}{T_n}\right) \log \frac{1 - \eta_n S_n/T_n}{1 - S_n/T_n} - S_n \log \frac{S_n/T_n}{1 - u_n}\right) \\
&\leq o\left(e^{-K_1 n \epsilon_n^2}\right) + \exp\left(T_n \left(1 - \frac{S_n}{T_n}\right) \left(\frac{(1 - \eta_n) S_n/T_n}{1 - S_n/T_n} + o\left(\frac{(1 - \eta_n) S_n/T_n}{1 - S_n/T_n}\right)\right) - K n \epsilon_n^2\right) \\
&= o\left(e^{-K_1 n \epsilon_n^2}\right) + \exp\left(S_n(1 - \eta_n) + o(S_n(1 - \eta_n)) - K n \epsilon_n^2\right) \\
&= o\left(e^{-K_1 n \epsilon_n^2}\right) + o\left(e^{-K_2 n \epsilon_n^2}\right) \\
&= o\left(e^{-\min\{K_1, K_2\} n \epsilon_n^2}\right).
\end{aligned}
\tag{68}
$$

for some $4 < K_1 < K$ and $4 < K_2 < K_0$. Letting $C'' = (\min\{K_1, K_2\} - 2)/\sigma^2$, (c) holds. We use Bernoulli's inequality and a tail bound for binomial distribution in Arratia and Gordon [1989] for the second inequality. Next, as in the proof of Theorem 1, there is a constant $C_1 > 0$ and $\hat{f}_n = f_{\hat{\theta}} \in \mathcal{F}_n$ such that

$$
\begin{gathered}
\left\|\hat{f}_n - f_0\right\|_{L^2} \leq C_1 \|f_0\|_{B^s_{p,q}([0,1]^d)} N_n^{-s/d} \leq \epsilon_n/4, \\
\|f - f_0\|_n^2 \leq 4 \|f - f_0\|_{L^2}^2
\end{gathered}
\tag{69}
$$

for sufficiently large n almost surely. Let $\hat{\gamma}$ and $\hat{\theta}_{\hat{\gamma}}$ be index and value of nonzero components of $\hat{\theta}$ respectively. Let

$$
\tilde{\Theta}(L, W, S, B, a) = \left\{\tilde{\theta} : \theta \in \Theta(L, W, S, B, a)\right\} = \Theta(L, W, S, B)
$$

and $\tilde{\Theta}(\hat{\gamma}; L_n, W_n, S_n, B_n, a_n)$ be a subset of parameter space $\tilde{\Theta}(L_n, W_n, S_n, B_n, a_n)$ consists of parameters which have nonzero components at $\hat{\gamma}$ only and $\mathcal{F}_n(\hat{\gamma}) = \tilde{\Phi}(\hat{\gamma}; L_n, W_n, S_n, B_n, a_n) \cap \mathcal{UB}(F)$ be an uniformly bounded neural network space generated by $\tilde{\Theta}(\hat{\gamma}; L_n, W_n, S_n, B_n, a_n)$. Note

$$
\begin{aligned}
\Pi(A_{\epsilon_n, 1}) &= \Pi(f \in \mathcal{F} : \|f - f_0\|_n \leq \epsilon_n) \\
&\geq \Pi\left(f \in \mathcal{F} : \left\|f - \hat{f}_n\right\|_{L^2} \leq \epsilon_n/4\right) \\
&\geq \Pi\left(f \in \mathcal{F} : \left\|f - \hat{f}_n\right\|_{L^\infty} \leq \epsilon_n/4\right) \\
&\geq \Pi\left(f \in \mathcal{F}_n(\hat{\gamma}) : \left\|f - \hat{f}_n\right\|_{L^\infty} \leq \epsilon_n/4\right)
\end{aligned}
\tag{70}
$$

for sufficiently large n. As in the proof of Lemma 4,

$$\begin{aligned}
&\Pi\left(f \in \mathcal{F}_n(\hat{\gamma}) : \left\|f - \hat{f}_n\right\|_{L^\infty} \leq \epsilon_n/4\right) \\
&\geq \Pi\Bigg(\theta \in \mathbb{R}^{T_n} : \theta_{\hat{\gamma}^c} \in [-a_n, a_n]^{T_n - S_n},\ \|\theta_{\hat{\gamma}}\|_\infty \leq B_n, \\
&\qquad \left\|\hat{\theta}_{\hat{\gamma}} - \theta_{\hat{\gamma}}\right\|_\infty \leq \frac{\epsilon_n}{4(W_n+1)^{L_n} L_n (B_n \vee 1)^{L_n - 1}}\Bigg) \\
&\geq u_n^{T_n - S_n} \left(\int_{B_n - t_n}^{B_n} g(t)dt\right)^{S_n},
\end{aligned} \tag{71}$$

where $t_n = \frac{\epsilon_n}{4(W_n+1)^{L_n} L_n (B_n \vee 1)^{L_n - 1}}$. Letting

$$\begin{aligned}
y_n &= \int_{B_n - t_n}^{B_n} g(t)dt \geq t_n g(B_n), \\
-\log \Pi(A_{\epsilon_n,1}) &\leq -S_n \log y_n - (T_n - S_n)\log u_n \\
&\leq -S_n \log(t_n g(B_n)) - T_n \left(1 - \frac{S_n}{T_n}\right)\log(1 - S_n/T_n) \\
&= -S_n \log t_n - S_n \log g(B_n) + T_n \left(1 - \frac{S_n}{T_n}\right)(S_n/T_n + o(S_n/T_n)) \\
&\lesssim S_n (\log n)^2 + S_n + o(S_n) \\
&\lesssim n\epsilon_n^2.
\end{aligned} \tag{72}$$

□

D Numerical Examples

D.1 Gaussian Mixture Prior

Consider a problem of estimating the Besov functions $f_1,\ f_2,\ f_3$ and f_4. Assume the Gaussian mixture prior distribution for each parameter as in Example 2. We sampled n points from $x \sim U(0,1)$ and y from $\mathcal{N}\big(f_i(x),\ \sigma_i^2\big)$ for $i = 1, 2, 3$ and 4. We set $\sigma_1^2 = \sigma_2^2 = 0.01^2$ and $\sigma_3^2 = \sigma_4^2 = 0.1^2$.

As mentioned in Section 4, we used a smaller model that was weaker than the conditions of the theorem. We fix the depth $L_n = 5$, the width $W_n = 200$ and set $N_n \leftarrow W_n/W_0$ of the model. We consider the prior in Example 2 (Gaussian mixture BNN) and set the smaller variance $\sigma_{1n} \leftarrow \max\{0.001,\ \sigma_{1n}\}$ to avoid numerical precision issues.

We fit the model using the NUTS algorithm [Hoffman et al., 2014] with the python `Pyro` [Bingham et al., 2019] and `PyTorch` [Paszke et al., 2019] packages. Experiments were run on a GPU server with Nvidia GeForce GTX TITAN X and RTX 3090. The code and instructions for the experiment are provided in the supplementary material. Figure 2 shows the results. Overall, as the number of the data n increase, the mean functions get closer to the true regression function.

D.2 Gaussian Prior

We performed the same experiment as before using Gaussian prior (Gaussian BNN). The standard deviation of the Gaussian prior is the same as σ_{2n} in the Gaussian mixture BNN (see Example 1). Figure 3 shows the results. Overall, as the number of the data n increase, the mean functions get closer to the true regression function.

D.3 Comparison

We compare two model in terms of empirical errors (training error) $\frac{1}{n}\sum_{i=1}^n |y_i - f(x_i)|^2$ and L^2 norms (test error) $e_f = \|f - f_0\|_{L^2}$ between true function f_0 and sampled function f. We approxi-

Figure 2: The results of estimating four functions $f_1,\ f_2,\ f_3$ and f_4 using Gaussian mixture BNN model with $n = 100$ (left), $n = 1,000$ (center) and $n = 10,000$ (right) samples. We construct 1,000 functions from the MCMC samples and plot the mean function with training data. The blue lines are the mean functions and the yellow intervals are the prediction intervals.

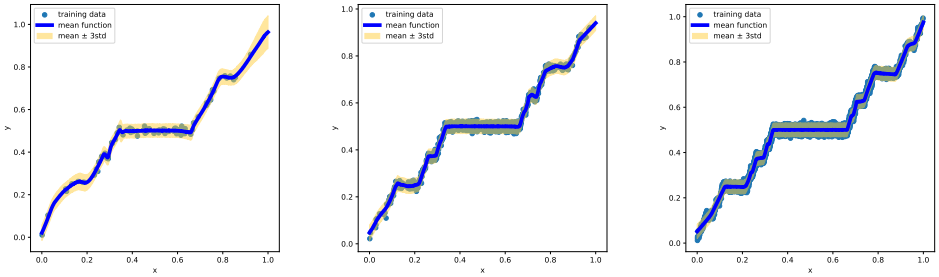

(a) The results of estimating f_1.

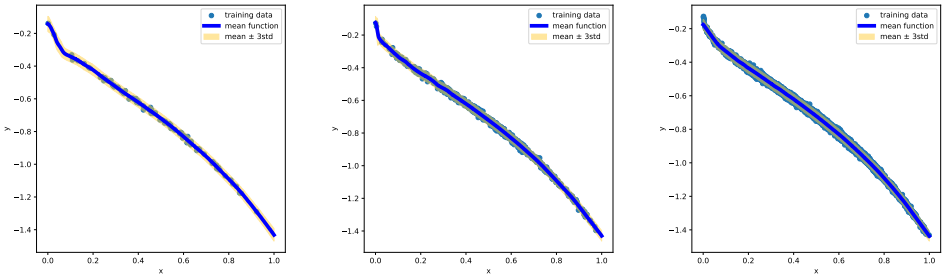

(b) The results of estimating f_2.

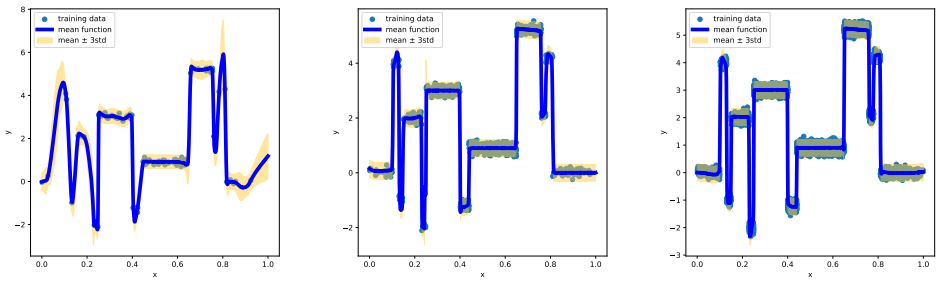

(c) The results of estimating f_3.

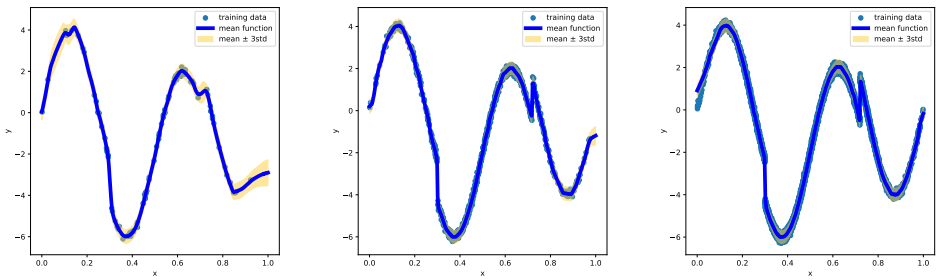

(d) The results of estimating f_4.

Figure 3: The results of estimating four functions $f_1,\ f_2,\ f_3$ and f_4 using Gaussian BNN models with $n = 100$ (left), $n = 1,000$ (center) and $n = 10,000$ (right) samples. We construct 1,000 functions from the MCMC samples and plot the mean function with training data. The blue lines are the mean functions and the yellow intervals are the prediction intervals.

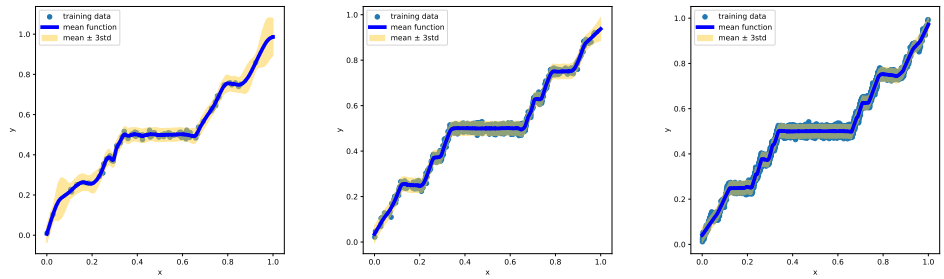

(a) The results of estimating f_1.

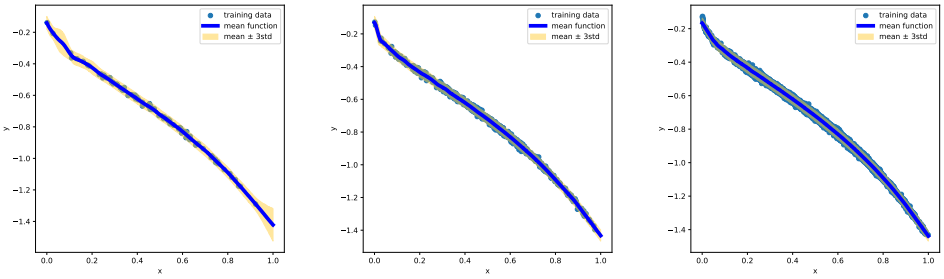

(b) The results of estimating f_2.

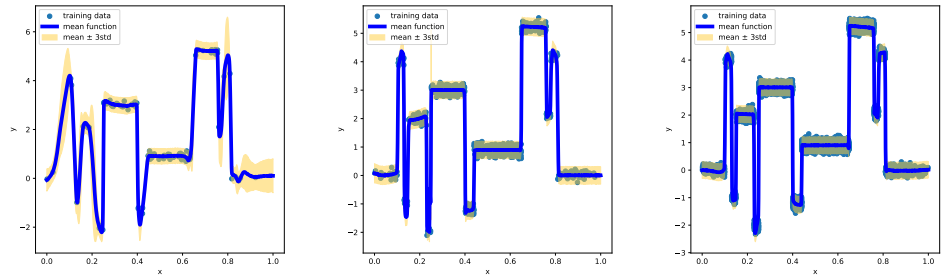

(c) The results of estimating f_3.

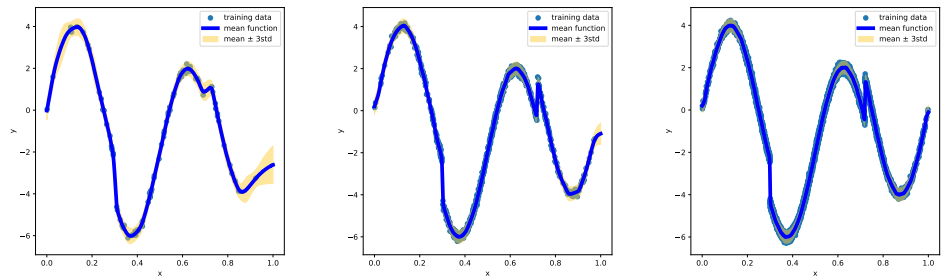

(d) The results of estimating f_4.

mate e_f as $\hat{e}_f$ using Riemann sum,

$$\hat{e}_f = \left(\frac{1}{n^*} \sum_{k=1}^{n^*} (f(x_k^*) - f_0(x_k^*))^2 \right)^{1/2}, \quad x_k^* = \frac{k}{n^*},\ 1 \le k \le n^*.$$

As in Table 3 and Table 4, both models give similar results. For the functions f_1 and f_2, Gaussian mixture BNN shows better results for both errors. For the functions f_3 and f_4, Gaussian BNN gives slightly better results for the test error. However, as shown in Figure 2 and Figure 3, the differences are negligible.

Table 3: Summary statistics of the empirical errors between true functions and sampled functions, where the numbers in the parentheses denote the standard deviations.

f_0	n	Gaussian mixture BNN	Gaussian BNN
f_1	100	1.398×10^{-2} (8.896×10^{-4})	1.386×10^{-2} (8.774×10^{-4})
	1,000	1.451×10^{-2} (2.899×10^{-4})	1.455×10^{-2} (2.840×10^{-4})
	10,000	1.455×10^{-2} (9.233×10^{-5})	1.461×10^{-2} (8.763×10^{-5})
f_2	100	1.336×10^{-2} (8.119×10^{-4})	1.359×10^{-2} (9.023×10^{-4})
	1,000	1.424×10^{-2} (2.781×10^{-4})	1.427×10^{-2} (2.741×10^{-4})
	10,000	1.420×10^{-2} (8.728×10^{-5})	1.419×10^{-2} (8.651×10^{-5})
f_3	100	1.428×10^{-1} (9.301×10^{-3})	1.497×10^{-1} (9.356×10^{-3})
	1,000	2.297×10^{-1} (4.256×10^{-3})	2.289×10^{-1} (4.632×10^{-3})
	10,000	2.200×10^{-1} (3.228×10^{-3})	2.204×10^{-1} (2.668×10^{-3})
f_4	100	1.389×10^{-1} (9.279×10^{-3})	1.390×10^{-1} (9.085×10^{-3})
	1,000	1.430×10^{-1} (2.898×10^{-3})	1.460×10^{-1} (2.921×10^{-3})
	10,000	1.543×10^{-1} (8.986×10^{-4})	1.500×10^{-1} (8.825×10^{-4})

Table 4: Summary statistics of the test errors between true functions and sampled functions, where the numbers in the parentheses denotes the standard deviations.

f_0	n	Gaussian mixture BNN	Gaussian BNN
f_1	100	1.921×10^{-2} (1.441×10^{-3})	2.359×10^{-2} (3.646×10^{-3})
	1,000	1.327×10^{-2} (3.172×10^{-4})	1.335×10^{-2} (5.983×10^{-4})
	10,000	1.212×10^{-2} (2.808×10^{-4})	1.201×10^{-2} (2.643×10^{-4})
f_2	100	1.363×10^{-2} (6.582×10^{-4})	1.652×10^{-2} (2.293×10^{-3})
	1,000	1.151×10^{-2} (3.835×10^{-4})	1.164×10^{-2} (4.203×10^{-4})
	10,000	1.177×10^{-2} (2.626×10^{-4})	1.155×10^{-2} (2.529×10^{-4})
f_3	100	1.120×10^{0} (5.175×10^{-2})	1.041×10^{0} (6.923×10^{-2})
	1,000	3.750×10^{-1} (5.497×10^{-3})	3.699×10^{-1} (6.699×10^{-3})
	10,000	2.478×10^{-1} (4.589×10^{-3})	2.479×10^{-1} (4.036×10^{-3})
f_4	100	4.942×10^{-1} (2.930×10^{-2})	4.493×10^{-1} (4.450×10^{-2})
	1,000	1.715×10^{-1} (8.959×10^{-3})	1.666×10^{-1} (1.029×10^{-2})
	10,000	1.550×10^{-1} (3.242×10^{-3})	1.128×10^{-1} (2.454×10^{-3})