Masked Bayesian Neural Networks : Computation and Optimality

Insung Kong
Department of Statistics
Seoul National University
ggong369@snu.ac.kr

Dongyoon Yang
Department of Statistics
Seoul National University
ydy0415@gmail.com

Jongjin Lee
Department of Statistics
Seoul National University
ga0408@snu.ac.kr

Ilsang Ohn
Department of Statistics
Inha University
ilsang.ohn@inha.ac.kr

Yongdai Kim *
Department of Statistics
Seoul National University
ydkim0903@gmail.com

Abstract

As data size and computing power increase, the architectures of deep neural networks (DNNs) have been getting more complex and huge, and thus there is a growing need to simplify such complex and huge DNNs. In this paper, we propose a novel sparse Bayesian neural network (BNN) which searches a good DNN with an appropriate complexity. We employ the masking variables at each node which can turn off some nodes according to the posterior distribution to yield a nodewise sparse DNN. We devise a prior distribution such that the posterior distribution has theoretical optimalities (i.e. minimax optimality and adaptiveness), and develop an efficient MCMC algorithm. By analyzing several benchmark datasets, we illustrate that the proposed BNN performs well compared to other existing methods in the sense that it discovers well condensed DNN architectures with similar prediction accuracy and uncertainty quantification compared to large DNNs.

1 Introduction

Over the past decade, Deep Learning (DL) has achieved tremendous success in many AI applications including computer vision and natural language processing. A key component of DL is complex architectures called Deep Neural Networks (DNN), which consist of multiple non-linear hidden layers. As the size of available data and computing power increase, the architectures of DNNs have been getting more and more complex and huge. As a result, recent DNNs become heavily over-parameterized. These over-parameterized networks result in expensive computation costs and requirements of excessive storage capacities when they are deployed to application systems. Moreover, over-parameterized DNNs frequently overfit the training data to result in worse performance in prediction.

Various methodologies have been proposed to alleviate over-parameterization. A popular approach is to sparsify a given DNN by iteratively pruning unnecessary edges [22, 13]. However, these edge sparsified DNNs may not be fully efficient in reducing inferential costs (i.e. computation costs in the prediction phase) because the dimensions of the weight matrices are not changed even if some weights are zero. In contrast, pruning nodes - neurons for multilayer perceptron (MLP) and channels for convolution neural network (CNN), can fully take advantage of sparsity. By pruning nodes in

* *Corresponding author*

Preprint. Under review.

DNNs [58, 25, 37, 54], the sparsified networks have the reduced widths to have low dimensional weight matrices which results in less inferential costs without any special techniques.

Bayesian approaches for learning DNN, which is called Bayesian Neural Networks (BNN) [40, 42], have been received much attention and successfully applied to various applications. Particularly, BNNs have the merit of having better generalization ability as well as better uncertainty quantification [60, 29]. Applications of BNNs range from recommender systems [55] to topic modeling [15], medical diagnosis [12], and astrophysics [8], to name just a few.

In this paper, we propose a new node-sparse BNN model called the masked BNN (mBNN) which is computationally feasible and has good theoretical properties. We develop a novel MCMC algorithm which makes the Bayesian inference of the node-sparse BNN model possible to provide better uncertainty quantification compared the variational inference (VI). Moreover, we prove that the posterior concentration rate to the true model is near minimax optimal and adaptive to the smoothness of the true model. The adaptiveness, which is the first of its kind for node-sparse BNNs, means that the mBNN selects optimal sparse architectures without knowing the complexity of the true model.

Our contributions are summarized as follows:

- We develop a node-sparse prior for DNNs such that Bayesian inference with a specially designed MCMC algorithm is possible and the posterior concentration rate to the true model is near minimax optimal adaptively to the smoothness of the true model.
- We implement an efficient MCMC algorithm for searching good node-sparse BNNs. In particular, we develop a local informed proposal distribution in the Metropolis-Hastings (MH) algorithm to search good node-sparse architectures efficiently.
- We provide sufficient conditions on the prior such that the posterior achieves the near minimax concentration rate and adaptive to the smoothness of the true model.
- By analyzing benchmark datasets, we illustrate that the mBNN is compared favorably to other non-sparse Bayesian approaches in terms of generalization and uncertainty quantification while inferential cost and model capacity are saved significantly. In addition, it is shown that the Bayesian inference by our MCMC algorithm outperforms the VI.

1.1 Related Work

BNN algorithms Various computation algorithms to approximate the posterior distribution of BNN have been developed. Stochastic Gradient Markov Chain Monte Carlo (SG-MCMC) methods [57, 6, 36, 64, 26, 59] provide a promising direction for sampling from the posterior distribution to allow efficient scalable inference. VI methods [20, 3, 27, 38, 50, 10], which have the merit of less computation cost, are also popularly used for approximating the posterior distribution. Deep ensembles [34, 45, 65, 47], which are often treated as a non-Bayesian alternative, are a promising way of mimicking the Bayesian inference without approximating the posterior distribution directly.

Sparse BNN algorithms BNNs also play an important role as a tool for compressing DNNs. Specifically, [41] uses variatonal dropout [31] to prune redundant weights. [9] and [56] develop edge-sparse learning algorithms via adaptive empirical Bayesian methods. Although their algorithms achieve promising compression rates, the reduction of inferential cost is much less than that of node-sparse BNNs. Using scale-mixture priors, [39] and [17] develop node-sparse BNNs based on the VI. However, it is well known that the VI is not good at approximating the true posterior [2, 62, 29] and thus yields suboptimal results.

2 Preliminaries

2.1 Notation

Let $\mathbb{R}$ and $\mathbb{N}$ be the sets of real numbers and natural numbers, respectively. For an integer $n \in \mathbb{N}$, we denote $[n] := \{1, \dots, n\}$. A capital letter denotes a random variable or matrix interchangeably whenever its meaning is clear, and a vector is denoted by a bold letter, e.g. $\boldsymbol{x} := (x_1, \dots, x_d)^\top$. For a d-dimensional vector $\boldsymbol{x} \in \mathbb{R}^d$, we denote $|\boldsymbol{x}|_p := (\sum_{j=1}^d |x_j|^p)^{1/p}$ for $1 \le p < \infty$, $|\boldsymbol{x}|_0 := \sum_{j=1}^d \mathbb{I}(x_j \neq 0)$ and $|\boldsymbol{x}|_\infty := \max_{j \in [d]} |x_j|$. For a real-valued function $f : \mathcal{X} \to \mathbb{R}$ and $1 \le p < \infty$,

we denote $||f||_{p,n} := (\sum_{i=1}^{n} f(\boldsymbol{x}_i)^p/n)^{1/p}$ and $||f||_{p,\mathrm{P}_{\boldsymbol{X}}} := \left(\int_{\boldsymbol{X}\in\mathcal{X}} f(\boldsymbol{X})^p d\mathrm{P}_{\boldsymbol{X}}\right)^{1/p}$ where $\mathrm{P}_{\boldsymbol{X}}$ is a probability measure defined on input space $\mathcal{X}$. Moreover, we define $||f||_\infty = \sup_{\boldsymbol{x}\in\mathcal{X}} |f(\boldsymbol{x})|$. For $b_1 \leq b_2$, we define $f_{[b_1,b_2]}(\cdot) := \min(\max(f(\cdot), b_1), b_2)$ which is a truncated version of f on b_1 and b_2. We denote $\circ$ and $\odot$ as the composition of functions and element-wise product of vectors or matrices, respectively. For a probability vector $\boldsymbol{q} \in \mathbb{R}^r$, we denote $\mathrm{Cat}(\boldsymbol{q})$ as the categorical distribution with the probabilities of each category being $\boldsymbol{q}$ and $\mathrm{Multi}_{(N,\boldsymbol{q})}$ as the distribution of N many selected balls without replacement from the jar containing r many balls whose selection probabilities are $\boldsymbol{q}$. For technical simplicity, we assume $\mathcal{X} = [-1, 1]^d$.

2.2 Data generating process

We consider two supervised learning problems : regression and classification. In regression problems, the input vector $\boldsymbol{X}$ and the response variable $Y \in \mathbb{R}$ are generated from the model

$$\begin{aligned} \boldsymbol{X} \sim & \mathrm{P}_{\boldsymbol{X}}, \\ Y|\boldsymbol{X} \sim & N(f_0(\boldsymbol{X}), \sigma_0^2), \end{aligned} \tag{1}$$

where $\mathrm{P}_{\boldsymbol{X}}$ is the probability measure defined on $\mathcal{X}$. Here, $f_0 : \mathcal{X} \to \mathbb{R}$ and $\sigma_0^2 > 0$ are the unknown true regression function and unknown variance of the noise, respectively.

For K-class classification problems, the input vector $\boldsymbol{X}$ and the response variable $Y \in [K]$ are generated from the model

$$\begin{aligned} \boldsymbol{X} \sim & \mathrm{P}_{\boldsymbol{X}}, \\ Y|\boldsymbol{X} \sim & \,\mathrm{Cat}\left(\mathrm{softmax}(\boldsymbol{f}_0(\boldsymbol{X}))\right), \end{aligned} \tag{2}$$

where $\mathrm{P}_{\boldsymbol{X}}$ is the probability measure defined on $\mathcal{X}$ and $\boldsymbol{f}_0 : \mathcal{X} \to \mathbb{R}^K$ is the logit of the unknown true conditional class probability function.

3 Masked Bayesian Neural Network

To construct a node-sparse BNN, we propose to use masking vectors which screen some nodes of the hidden layers. We first define the masked Deep Neural Network (mDNN) model, and then propose the mBNN on the top of the mDNN by specifying a prior appropriately. We focus the mBNN for MLP in the main manuscript and the mBNN for CNN is explained in Appendix B.1.

3.1 Deep Neural Network

For $L \in \mathbb{N}$ and $\boldsymbol{p} = (p^{(0)}, p^{(1)}, ..., p^{(L)}, p^{(L+1)})^\top \in \mathbb{N}^{L+2}$, DNN with the $(L, \boldsymbol{p})$ architecture is a DNN model which has L hidden layers and $p^{(l)}$ many nodes at the l-th hidden layer for $l \in [L]$. The input and output dimensions are $p^{(0)}$ and $p^{(L+1)}$, respectively. The output of the DNN model can be written as

$$f_{\boldsymbol{\theta}}^{\mathrm{DNN}}(\cdot) := A_{L+1} \circ \rho \circ A_L \cdots \circ \rho \circ A_1(\cdot), \tag{3}$$

where $A_l : \mathbb{R}^{p^{(l-1)}} \mapsto \mathbb{R}^{p^{(l)}}$ for $l \in [L+1]$ is an affine map defined as $A_l(\boldsymbol{x}) := W_l\boldsymbol{x} + \boldsymbol{b}_l$ with $W_l \in \mathbb{R}^{p^{(l)} \times p^{(l-1)}}$ and $\boldsymbol{b}_l \in \mathbb{R}^{p^{(l)}}$ and ρ is the RELU activation function. The DNN model is parameterized by $\boldsymbol{\theta}$ which is the concatenation of the weight matrices and bias vectors, that is

$$\boldsymbol{\theta} := (\mathrm{vec}(W_1)^\top, b_1^\top, \ldots, \mathrm{vec}(W_{L+1})^\top, b_{L+1}^\top)^\top.$$

3.2 Masked Deep Neural Network

For a given standard DNN, the corresponding masked DNN (mDNN) is constructed by simply adding masking parameters to the standard DNN model. For $l \in [L]$, the mDNN model screens the nodes at the l-th hidden layer using the binary masking vector $\boldsymbol{m}^{(l)} \in \{0, 1\}^{p^{(l)}}$. When $\left(\boldsymbol{m}^{(l)}\right)_j = 0$, the j-th node at the l-th hidden layer becomes inactive. The output of the mDNN model with the $(L, \boldsymbol{p})$ architecture can be written as

$$f_{\boldsymbol{M},\boldsymbol{\theta}}^{\mathrm{mDNN}}(\cdot) := A_{L+1} \circ \rho_{\boldsymbol{m}^{(L)}} \circ A_L \cdots \circ \rho_{\boldsymbol{m}^{(1)}} \circ A_1(\cdot),$$

where A_l for $l \in [L+1]$ is the affine map defined in (3) and $\rho_{\boldsymbol{m}^{(l)}} : \mathbb{R}^{p^{(l)}} \to \mathbb{R}^{p^{(l)}}$ for $l \in [L]$ is the masked-RELU activation function defined as

$$\rho_{\boldsymbol{m}^{(l)}} \begin{pmatrix} x_1 \\ \vdots \\ x_{p^{(l)}} \end{pmatrix} := \boldsymbol{m}^{(l)} \odot \begin{pmatrix} \max(x_1, 0) \\ \vdots \\ \max(x_{p^{(l)}}, 0) \end{pmatrix}.$$

The model is parameterized by $\boldsymbol{M}$ and $\boldsymbol{\theta}$, where $\boldsymbol{M} \in \{0,1\}^{\sum_{l=1}^{L} p^{(l)}}$ is the concatenate of the all masking vectors

$$\boldsymbol{M} := \left(\boldsymbol{m}^{(1)\top}, \ldots, \boldsymbol{m}^{(L)\top} \right)^\top$$

and $\boldsymbol{\theta}$ is the concatenation of the weight matrices and the bias vectors in the standard DNN model.

Note that the mDNN is nothing but a standard DNN with the architecture $(L, \tilde{\boldsymbol{p}})$ where $(\tilde{\boldsymbol{p}})_l = \left|\boldsymbol{m}^{(l)}\right|_0$. That is, the mDNN is a reparameterization of the standard DNN using the masking vectors. This reparameterization, however, allows us to develop an efficient MCMC algorithm, in particular for searching good architectures (i.e. good masking vectors).

3.3 Prior and posterior distribution

We adopt a data-dependent prior Π_n on the parameters $\boldsymbol{M}$ and $\boldsymbol{\theta}$ (as well as σ^2 for regression problems). We consider a data-dependent prior to ensure the optimal posterior concentration rate and adaptiveness. We assume that a priori $\boldsymbol{M}$ and $\boldsymbol{\theta}$ (as well as σ^2 for regression problems) are independent.

For $\boldsymbol{M}$, we use the following hierarchical prior. For each $\boldsymbol{m}^{(l)}$ with $l \in [L]$, let $s^{(l)} := |\boldsymbol{m}^{(l)}|_0$ be the sparsity of the masking vector $\boldsymbol{m}^{(l)}$. We put a prior mass on $s^{(l)}$ by

$$\Pi_n(s^{(l)}) \propto e^{-(\lambda \log n)^5 s^{(l)2}} \qquad \text{for } s^{(l)} \in [p^{(l)}], \tag{4}$$

where $\lambda > 0$ is a hyper-parameter. Note that the prior on $s^{(l)}$ regularizes the width of the network, and more strong regularization is enforced as more data are accumulated. Given the sparsity level $s^{(l)}$, the masking vector $\boldsymbol{m}^{(l)}$ is sampled from the set $\{\boldsymbol{m}^{(l)} \in \{0,1\}^{p^{(l)}} : |\boldsymbol{m}^{(l)}|_0 = s^{(l)}\}$ uniformly. In other words,

$$\Pi_n(\boldsymbol{m}^{(l)} | s^{(l)}) \overset{iid}{\propto} \frac{1}{\binom{p^{(l)}}{s^{(l)}}} \mathbb{I}(|\boldsymbol{m}^{(l)}|_0 = s^{(l)}). \tag{5}$$

For the prior of $\boldsymbol{\theta}$, we assume that

$$\theta_i \overset{iid}{\sim} \mathfrak{p}(\theta_i) \qquad i \in [T], \tag{6}$$

where $T := \sum_{l=0}^{L} (p^{(l)} + 1) p^{(l+1)}$ is the length of the vector $\boldsymbol{\theta}$, and choose $\mathfrak{p}$ carefully to ensure desirable theoretical properties. For high-dimensional linear regression problems, [4] and [5] notice that using a heavy-tailed distribution for the prior of the regression coefficients is essential for theoretical optimality. Motivated by these observations, we consider a heavy-tailed distribution for $\mathfrak{p}$. Let $\mathfrak{P}$ be the class of polynomial tail distributions on $\mathbb{R}$ defined as

$$\mathfrak{P} := \left\{ \mathfrak{p} : \lim_{x \to \infty} \frac{x^{-\log x}}{\mathfrak{p}(x)} \to 0 \text{ and } \lim_{x \to \infty} \frac{x^{-\log x}}{\mathfrak{p}(-x)} \to 0 \right\}.$$

Examples of polynomial tail distributions are the Cauchy distribution and Student's t-distribution. On the other hand, the Gaussian and Laplace distributions do not belong to $\mathfrak{P}$. We assume that $\mathfrak{p}$ belongs to $\mathfrak{P}$. We will show in Section 5 that any prior in $\mathfrak{P}$ yields the optimal posterior concentration rate.

For the prior of σ^2 in regression problems, a standard distribution such as the inverse-gamma distribution can be used. Any distribution whose density is positive at the true σ_0^2 works for theoretical optimality.

3.4 Comparison with MC Dropout

The idea of masking nodes in DNN has been already used in various algorithms. Dropout [49] and MC-dropout [14] are two representative examples, where they randomly mask the nodes of DNN during the training phase. While Dropout abolishes the masking vectors and uses the scaled-down version of the trained weights in the prediction phase, MC-dropout uses the trained weights obtained in the training phase multiplied by a random masking vectors. Since the masking vectors is treated as a random vectors following its posterior distribution at the prediction phase, our mBNN is similar to MC-dropout. A key difference between the mBNN and MC-dropout, however, is that the mBNN learns the distribution of the masking vectors from data via the posterior distribution but MC-dropout does not. That is, the mBNN learns the architecture of DNN from data, which makes the mBNN have good theoretical and empirical properties.

4 Posterior inference

Let $\mathcal{D}^{(n)} := \{(\boldsymbol{X}_i, Y_i)\}_{i\in[n]}$ be training data. Let w denote all of the parameters in the mBNN, that is, $w := (\boldsymbol{M}, \boldsymbol{\theta}, \sigma^2)$ for the regression problem (1) and $w := (\boldsymbol{M}, \boldsymbol{\theta})$ for the classification problem (2). For a new test example $\boldsymbol{x} \in \mathcal{X}$, the prediction with the mBNN is done by the predictive distribution:

$$p(y \mid \boldsymbol{x}, \mathcal{D}^{(n)}) = \int_w p(y \mid \boldsymbol{x}, w)\Pi_n(w \mid \mathcal{D}^{(n)})dw,$$

where $\Pi_n(w \mid \mathcal{D}^{(n)})$ is the posterior posterior distribution of w. When the integral is difficult to be evaluated, it is common to approximate it by the Monte Carlo method

$$p(y \mid \boldsymbol{x}, \mathcal{D}^{(n)}) \approx \frac{1}{T}\sum_{t=1}^{T} p(y \mid \boldsymbol{x}, w^{(t)}),$$

where $w^{(t)} \sim \Pi_n(w \mid \mathcal{D}^{(n)})$. In this section, we develop a MCMC algorithm to sample w efficiently from $\Pi_n(w \mid \mathcal{D}^{(n)})$.

4.1 MCMC algorithm

The proposed MCMC algorithm samples $\boldsymbol{\theta}$ (and σ^2) given $\boldsymbol{M}$ and $\mathcal{D}^{(n)}$ and then samples $\boldsymbol{M}$ given $\boldsymbol{\theta}$ (and σ^2) and $\mathcal{D}^{(n)}$, and iterates these two samplings until convergence.

There are various efficient sampling algorithms for $\boldsymbol{\theta}$ (and σ^2) given $\boldsymbol{M}$ and $\mathcal{D}^{(n)}$ such as Hamiltonian Monte Carlo (HMC) [43], Stochastic Gradient Langevin Dynamics (SGLD) [57] and Stochastic Gradient HMC (SGHMC) [6]. In practice, we select a sampling algorithm among those depending on the sizes of data and model.

For generating $\boldsymbol{M}$ from its conditional posterior, we consider the Metropolis-Hastings (MH) algorithm. The hardest part is to design a good proposal distribution since the dimension of $\boldsymbol{M}$ is quite large and all entries are binary. In the next subsection, we propose an efficient proposal distribution for $\boldsymbol{M}$.

4.2 Proposal for the MH algorithm

Essentially, sampling $\boldsymbol{M}$ is equivalent to sampling a large dimensional binary vector. A well known strategy for sampling a large dimensional binary vector is to use the MH algorithm with the locally informed proposal [52, 63] given as

$$q\left(\boldsymbol{M}^\star \mid \boldsymbol{M}\right) \propto e^{\frac{1}{2}(l(\boldsymbol{M}^\star)-l(\boldsymbol{M}))}\mathbb{I}\left(\boldsymbol{M}^\star \in H(\boldsymbol{M})\right),$$

where $l(\boldsymbol{M})$ is the log-posterior of $\boldsymbol{M}$ and $H(\boldsymbol{M})$ is the Hamming ball of a certain size around $\boldsymbol{M}$. While powerful, this locally informed proposal requires to compute $l(\boldsymbol{M}^\star)$ for every $\boldsymbol{M}^\star \in H(\boldsymbol{M})$, which could be time consuming.

To resolve this problem, we could use a linear approximation of l using the gradient information at $\boldsymbol{M}$ as is suggested by [19]. Even though $l(\boldsymbol{M})$ is defined only on binary vectors, the gradient

Algorithm 1 The algorithm of the proposed MH algorithm

1: Sample $u \sim \text{Bernoulli}(0.5)$, $N \sim \text{Uniform}(\{1, \dots, N_{\max}\})$.
2: Calculate the selection probability $\boldsymbol{Q}_u$, where

$$\boldsymbol{Q}_u := \text{softmax}\left(-|\nabla l(\boldsymbol{M}_u)|/2\right) \tag{7}$$

3: Sample N many nodes $\{i_1, \dots, i_N\}$ from $\{(l, j) : \boldsymbol{m}_j^{(l)} = u\}$ by

$$\{i_1, \dots, i_N\} \sim \text{Multi}_{(N, \boldsymbol{Q}_u)}$$

4: $\boldsymbol{M}^\star = \text{flip}(\boldsymbol{M}^{\text{curr}}, \{i_1, \dots, i_N\})$, where $\boldsymbol{M}^{\text{curr}}$ is the current masking vectors.
5: Accept or reject $\boldsymbol{M}^\star$ with the acceptance probability (10).

$\nabla l(\boldsymbol{M}) = \partial l(\boldsymbol{M})/\partial \boldsymbol{M}$ of $l(\boldsymbol{M})$ can be defined by extending the domain of $l(\boldsymbol{M})$ appropriately. Using the standard approximation technique

$$l(\boldsymbol{M}^*) - l(\boldsymbol{M}) \approx (\boldsymbol{M}^* - \boldsymbol{M}) \odot \nabla l(\boldsymbol{M}), \tag{8}$$

we could devise the proposal distribution as

$$\begin{aligned} \boldsymbol{Q} &= \text{softmax}\left(-(2\boldsymbol{M} - 1) \odot \frac{\nabla l(\boldsymbol{M})}{2}\right), \\ i &\sim \text{Cat}\left(\boldsymbol{Q}\right), \quad \boldsymbol{M}^\star = \text{flip}(\boldsymbol{M}, i), \end{aligned} \tag{9}$$

where $\text{flip}(\boldsymbol{M}, i)$ denotes the operation of flipping the value of the ith node. That is, the proposal distribution first selects a node i with probability $\boldsymbol{Q}$ and changes its value to $1 - \boldsymbol{M}_i$.

For the mBNN, we need to modify the proposal distribution (9). First, the approximation (8) is not accurate for the mBNN because the corresponding DNN is not locally smooth enough. Instead, we select an entry with probability reciprocally proportional to $|\nabla l(\boldsymbol{M})|$. That is, we select nodes which do not affect much to the current DNN model when their values are changed. By doing so, the acceptance rate becomes larger. This modification is motivated by the pruning algorithm of [35]. The second modification is that instead of changing an entry, our proposal changes multiple entries (i.e. multiple nodes) simultaneously to make the mixing of the MCMC algorithm be faster.

To be more specific, we first select the move $u \in \{0, 1\}$ with probability 1/2 and an inteager N from $[N_{\max}]$ uniformly, where $N_{\max}$ is a prespecified positive integer. Then, we select randomly N nodes among the nodes whose masking values are u following $\text{Multi}_{(N, \boldsymbol{Q}_u)}$, where $\boldsymbol{Q}_u := \text{softmax}\left(-|\nabla l(\boldsymbol{M}_u)|/2\right)$ and $\nabla l(\boldsymbol{M}_u) = \left(\nabla l(\boldsymbol{M})_k : \boldsymbol{M}_k = u, k \in [|\boldsymbol{M}|]\right)$. Finally, we flip the masking values of the selected nodes to have a new proposal $\boldsymbol{M}^*$. Then, we accept the new proposal $\boldsymbol{M}^*$ with probability

$$\alpha(\boldsymbol{M}^\star, \boldsymbol{M}) := \min\left(1, \frac{\Pi_n(\boldsymbol{M}^\star)}{\Pi_n(\boldsymbol{M})} \frac{p(Y^{(n)} \mid X^{(n)}, \boldsymbol{M}^\star, \boldsymbol{\theta})}{p(Y^{(n)} \mid X^{(n)}, \boldsymbol{M}, \boldsymbol{\theta})} \frac{\text{Multi}_{(N, \boldsymbol{Q}^\star_{1-u})}(\{i_1, \dots, i_N\})}{\text{Multi}_{(N, \boldsymbol{Q}_u)}(\{i_1, \dots, i_N\})}\right), \tag{10}$$

where $\boldsymbol{Q}^\star_{1-u} := \text{softmax}\left(-|\nabla l(\boldsymbol{M}^\star_{1-u})|/2\right)$ and $\{i_1, \dots, i_N\}$ is the set of selected nodes.

Note that some inactive nodes become active when $u = 0$ (i.e. birth of nodes) and some active nodes become inactive when $u = 1$ (i.e. death of nodes). The proposed MH algorithm is summarized in Algorithm 1.

5 Theoretical results

In this section, we derive the posterior concentration rates of the mBNNs for regression and classification problems. For technical reasons, we only consider mBNN for MLP. We assume that $\mathcal{D}^{(n)} := \{(\boldsymbol{X}_i, Y_i)\}_{i \in [n]}$ are independent copies following the true distribution $\mathbb{P}_0$ specified by either the model (1) or model (2).

5.1 Nonparametric regression

We consider the nonparametric regression model (1). We assume the true regression function f_0 belongs to the β-Hölder class $\mathcal{H}_d^\beta$. For inference, we consider the probabilistic model

$$Y_i \overset{ind.}{\sim} N(f^{\text{mDNN}}_{\boldsymbol{M},\boldsymbol{\theta}[-F,F]}(\boldsymbol{X}_i), \sigma^2),$$

where $f^{\text{mDNN}}_{\boldsymbol{M},\boldsymbol{\theta}}(\boldsymbol{X}_i)$ has the $(L_n, \boldsymbol{p}_n)$ architecture with L_n and $\boldsymbol{p}_n$ given as

$$L_n := \lceil C_L \log n \rceil, \tag{11}$$
$$p_n := \lceil C_p \sqrt{n} \rceil,$$
$$\boldsymbol{p_n} := (d, p_n, \dots, p_n, 1)^\top \in \mathbb{N}^{L_n+2} \tag{12}$$

for positive constants C_L and C_p that are defined in Lemma A.1. Then, the likelihood of $w := (\boldsymbol{M}, \boldsymbol{\theta}, \sigma^2)$ is expressed as

$$\mathcal{L}(w|\mathcal{D}^{(n)}) = (2\pi\sigma^2)^{-\frac{n}{2}} \exp\left(-\frac{\sum_{i=1}^n (Y_i - f^{\text{mDNN}}_{\boldsymbol{M},\boldsymbol{\theta}\,[-F,F]}(\boldsymbol{X}_i))^2}{2\sigma^2} \right),$$

and the corresponding posterior distribution is given as

$$\Pi_n(w \mid \mathcal{D}^{(n)}) \propto \Pi_n(w)\mathcal{L}(w|\mathcal{D}^{(n)}),$$

where Π_n is the prior defined on Section 3.3.

In the following theorem, we show that the mBNN model achieves the optimal (up to a logarithmic factor) posterior concentration rate to the true regression function.

Theorem 5.1 (Posterior Concentration of the mBNN for regression problems)**.** *Assume $f_0 \in \mathcal{H}_d^\beta$, $\beta < d$, and there exist $F > 0$ and $\sigma^2_{max} > 0$ such that $\|f_0\|_\infty \le F$ and $\sigma_0^2 \le \sigma^2_{max}$. Consider the mDNN model with the $(L_n, \boldsymbol{p}_n)$ architecture, where L_n and $\boldsymbol{p}_n$ are given in (11) and (12). If we put the prior given as (4), (5) and (6) over $\boldsymbol{M}$, $\boldsymbol{\theta}$ and any prior on σ^2 whose density (with respect to Lebesgue measure) is positive on its support $(0, \sigma^2_{max}]$, the posterior distribution concentrates to the f_0 and σ_0^2 at the rate $\varepsilon_n = n^{-\beta/(2\beta+d)} \log^\gamma(n)$ for $\gamma > \frac{5}{2}$ in the sense that*

$$\Pi_n\Big((f, \sigma^2) : ||f - f_0||_{2,\mathrm{P}_X} + |\sigma^2 - \sigma_0^2| > M_n\varepsilon_n \;\Big|\; \mathcal{D}^{(n)}\Big) \overset{\mathbb{P}_0^n}{\to} 0$$

as $n \to \infty$ for any $M_n \to \infty$, where $\mathbb{P}_0^n$ is the probability measure of the training data $\mathcal{D}^{(n)}$.

The convergence rate $n^{-\beta/(2\beta+d)}$ is known to be minimax optimal when estimating the β-Hölder smooth function [51]. Our concentration rate is nearly optimal up to a logarithmic factor and adaptive to the smoothness of the true model.

Similar convergence rates are derived in the non-bayesian theoretical deep learning literature [48, 32]. An important advantage of our result is that the prior does not depend on the smoothness β of the true regression function and hence our Bayesian model achieves the minimax optimality adaptive to the smoothness. In contrast, the architectures considered in [48, 32] depend on the smoothness β and thus are not adaptive.

Similar concentration rates have been derived for edge-sparse BNNs by [46, 7, 1]. As we explained before, edge-sparse BNNs are not fully efficient for the compression of DNNs. Recently, [44] consider a mixure of BNN models and prove that the VI posterior achieves a near minimax convergence rate adaptively to the smoothness of the true model. However, this model is not suitable for compression. Theorem 5.1 is the first result for theoretical optimality of the Bayesian analysis for node-sparse DNNs.

5.2 Binary Classification

We consider the classification problem (2) with $K = 2$, and denote $f_0 := (\boldsymbol{f}_0)_2 - (\boldsymbol{f}_0)_1$. We assume that f_0 belongs to the β-Hölder class $\mathcal{H}_d^\beta$. For inference, we consider the probabilistic model

$$Y_i \overset{ind.}{\sim} \text{Bernoulli}(\phi \circ f^{\text{mDNN}}_{\boldsymbol{M},\boldsymbol{\theta}[-F,F]}(\boldsymbol{X}_i)),$$

where ϕ is the sigmoid function and $f^{\text{mDNN}}_{\boldsymbol{M},\boldsymbol{\theta}}$ has the $(L_n, \boldsymbol{p}_n)$ architecture with L_n and $\boldsymbol{p}_n$ given as (11) and (12). Then, the likelihood of $w := (\boldsymbol{M}, \boldsymbol{\theta})$ is expressed as

$$\mathcal{L}(w|\mathcal{D}^{(n)}) = \prod_{i=1}^{n} (\phi \circ f^{\text{mDNN}}_{\boldsymbol{M},\boldsymbol{\theta}[-F,F]}(\boldsymbol{X}_i))^{Y_i} (1 - \phi \circ f^{\text{mDNN}}_{\boldsymbol{M},\boldsymbol{\theta}[-F,F]}(\boldsymbol{X}_i))^{1-Y_i}.$$

In the following theorem, we prove that the mBNN model achieves the optimal (up to a logarithmic factor) posterior concentration rate to the true conditional class probability adaptive to the smoothness of the true model.

Theorem 5.2 (Posterior Concentration of mBNN for classification problems)**.** *Assume $f_0 \in \mathcal{H}^{\beta}_d$, $\beta < d$, and there exists $F > 0$ such that $\|f_0\|_\infty \leq F$. Consider the mDNN model with the $(L_n, \boldsymbol{p}_n)$ architecture where L_n and $\boldsymbol{p}_n$ are given in (11) and (12). If we put the prior given as (4), (5) and (6) over $\boldsymbol{M}$ and $\boldsymbol{\theta}$, the posterior distribution concentrates to the true conditional class probability at the rate $\varepsilon_n = n^{-\beta/(2\beta+d)} \log^{\gamma}(n)$ for $\gamma > \frac{5}{2}$ in the sense that*

$$\Pi_n\Big(f : ||\phi \circ f - \phi \circ f_0||_{2,\mathrm{P}_X} > M_n \varepsilon_n \;\Big|\; \mathcal{D}^{(n)}\Big) \overset{\mathbb{P}^n_0}{\to} 0$$

as $n \to \infty$ for any $M_n \to \infty$, where $\mathbb{P}^n_0$ is the probability measure of the training data $\mathcal{D}^{(n)}$.

6 Experiment

In this section, we perform experiments to empirically justify the usefulness of the mBNN. In section 6.1, by analyzing several benchmark datasets, we show that the mBNN is compared favorably to other non-sparse Bayesian approaches in terms of generalization and uncertainty quantification while inferential cost and model capacity are saved significantly. In addition, we confirm that the Bayesian inference based on our MCMC algorithm outperforms the VI. In section 6.2, we illustrate the efficiency of the proposal distribution (7) in the MH algorithm by comparing the random proposal. All the experimental details as well as the results of additional numerical experiments are given in Appendix B.

6.1 Performance of mBNN

We compare the mBNN with other non-sparse Bayesian methods in terms of generalization, uncertainty quantification as well as inferential cost and model capacity by analyzing several benchmark datasets. Along with the Bayesian inference through MCMC algorithms (BNN), we consider mean field VI (MFVI) [3] and Deep Ensemble [34] for non-sparse Bayesian approaches and a node-sparse VI algorithm with the group normal-Jeffreys (BC-GNJ) prior [39].

Two hidden layer MLP with the layer sizes (100,100) is used for UCI regression datasets (*Boston, Concrete, Energy, Yacht*) while ResNet18 [24] architecture is used for image datasets (CIFAR-10, CIFAR-100). When we use ResNet18 architecture for experiments, we mask the channels of the convolution layers. See Appendix B.1 for the detail of the masked Bayesian CNN. We construct 20 random 90-to-10 train-test splits for UCI datasets while we repeat the experiments 3 times with different seeds for image datasets.

In Table 1, we report the means and standard errors of the performance measures on the repeated experiments. For the performance measures, the root mean square error (RMSE) of the Bayes estimator and the negative log-likelihood (NLL) of test data are considered for regression models and the accuracy of the Bayes estimator and the expected calibration error (ECE) [33] are used in classification models. Note that the NLL and ECE are two popular measures for uncertainty quantification. In addition, we report inferential costs (FLOPs) and the numbers of nonzero parameters (Capacity) of the node-sparse models relative to the non-sparse models.

The results in Table 1 amply show that the mBNN saves the inferential cost and memory usage significantly compared to the non-sparse models while maintaining competitive and sometimes better prediction accuracy and uncertainty quantification. The BC-GNJ also takes advantage of less inferential cost but its generalization ability and uncertainty quantification are far below those of the mBNN, which suggests that the Bayesian inference by our MCMC algorithm is desirable. Performance based on different evaluation criteria and additional results are presented in Appendix B.3.

Table 1: **Benchmark datasets.** Performance of mBNN, BNN, MFVI [3], Deep Ensemble [34] and BC-GNJ [39]. When inferring the posterior of BNN and mBNN, we use HMC [43] for UCI datasets while we use SGLD [57] for image datasets.

	Dataset	Measure	BNN	MFVI	DeepEns	mBNN✓	BC-GNJ
Regression	Boston	RMSE	2.942(0.150)	3.295(0.183)	2.835(0.157)	**2.775(0.110)**	2.975(0.149)
		NLL	2.613(0.094)	2.689(0.084)	2.585(0.130)	**2.507(0.067)**	2.758(0.117)
		FLOPs ↓ (%)	·	·	·	**89.42(0.50)**	81.61(0.59)
	Concrete	RMSE	4.855(0.164)	5.416(0.184)	5.009(0.148)	**4.794(0.152)**	5.082(0.253)
		NLL	3.071(0.058)	3.172(0.055)	**3.028(0.040)**	3.036(0.049)	3.367(0.116)
		FLOPs ↓ (%)	·	·	·	**85.80(0.51)**	64.64(0.69)
	Energy	RMSE	0.454(0.013)	2.092(0.110)	0.485(0.013)	**0.451(0.018)**	0.642(0.041)
		NLL	0.638(0.034)	2.139(0.062)	0.744(0.035)	**0.626(0.038)**	0.956(0.052)
		FLOPs ↓ (%)	·	·	·	77.06(0.87)	**94.90(0.26)**
	Yacht	RMSE	0.638(0.066)	1.469(0.118)	**0.564(0.045)**	0.630(0.043)	1.176(0.104)
		NLL	**0.860(0.075)**	1.771(0.076)	0.942(0.234)	0.882(0.089)	1.584(0.070)
		FLOPs ↓ (%)	·	·	·	**81.26(0.72)**	69.77(0.70)
Classification	CIFAR-10	ACC	0.944(0.001)	0.905(0.002)	**0.945(0.001)**	0.938(0.002)	0.908(0.002)
		ECE	0.008(0.000)	0.018(0.003)	0.007(0.001)	**0.006(0.000)**	0.012(0.001)
		FLOPs ↓ (%)	·	·	·	**73.28(0.99)**	53.92(2.27)
		Capacity ↓ (%)	·	·	·	**92.23(0.09)**	78.88(1.51)
	CIFAR-100	ACC	**0.765(0.002)**	0.669(0.004)	0.751(0.002)	0.752(0.005)	0.603(0.003)
		ECE	**0.002(0.000)**	0.007(0.000)	0.003(0.000)	**0.002(0.000)**	0.006(0.000)
		FLOPs ↓ (%)	·	·	·	61.23(0.74)	**63.28(0.54)**
		Capacity ↓ (%)	·	·	·	**75.16(0.10)**	59.30(3.85)

6.2 Efficiency of the proposal distribution

To illustrate the efficiency of our proposal distribution in Algorithm 1, we conduct a comparative experiment. As an alternative to the proposal distribution with selection probability (7), we consider the random proposal which gives the same selection probability on each active (or inactive) node. In other words, we simply change (7) with

$$\boldsymbol{Q}_u^{\text{random}} \propto \mathbb{I}(\boldsymbol{m}_j^{(l)} = u).$$

Figure 1 presents how the numbers of the activated nodes in each layer decrease as the MCMC iteration proceeds for *Concrete* dataset. It is obvious that the numbers of activated nodes for our proposal distribution decrease much faster than those for the random proposal. This fast decrease is observed for the other datasets whose results are given in Appendix B.3.

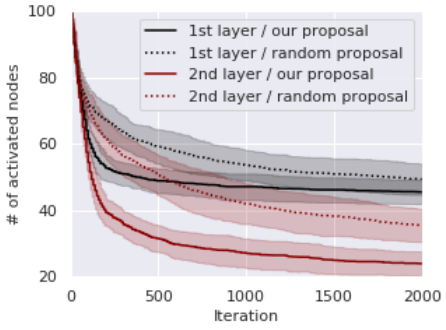

Figure 1: **Efficiency of the proposal.** The numbers of activated nodes in each layer (y-axis) as the MCMC iteration (x-axis) proceeds. We compare our proposal with the random proposal.

7 Discussion

As usual, computation cost for training the mBNN via the MCMC algorithm is relatively expensive compare to VI. Thus, reducing computation cost further to scale up the mBNN is still needed. A hybrid approach of combining the MCMC algorithm and VI would be a promising direction.

Even though theoretical properties of the mBNN for CNN is still unanswered, we have seen that the mBNN works well for CNN. For CIFAR-10, the capacity of the mBNN is less than 10% of the non-sparse model. Roughly speaking, 90% of filters used in the non-sparse model are redundant. It would be interesting to investigate which filters survive and which filters die in the mBNN, whose results would be amply used for interpretation of CNN. We will leave this issue as a future work.

References

[1] Jincheng Bai, Qifan Song, and Guang Cheng. Efficient variational inference for sparse deep learning with theoretical guarantee. *Advances in Neural Information Processing Systems*, 33: 466–476, 2020.

[2] David M Blei, Alp Kucukelbir, and Jon D McAuliffe. Variational inference: A review for statisticians. *Journal of the American statistical Association*, 112(518):859–877, 2017.

[3] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural network. In *International Conference on Machine Learning*, pages 1613–1622. PMLR, 2015.

[4] Ismaël Castillo and Aad van der Vaart. Needles and straw in a haystack: Posterior concentration for possibly sparse sequences. *The Annals of Statistics*, 40(4):2069–2101, 2012.

[5] Ismaël Castillo, Johannes Schmidt-Hieber, and Aad Van der Vaart. Bayesian linear regression with sparse priors. *The Annals of Statistics*, 43(5):1986–2018, 2015.

[6] Tianqi Chen, Emily Fox, and Carlos Guestrin. Stochastic gradient hamiltonian monte carlo. In *International conference on machine learning*, pages 1683–1691. PMLR, 2014.

[7] Badr-Eddine Chérief-Abdellatif. Convergence rates of variational inference in sparse deep learning. In *International Conference on Machine Learning*, pages 1831–1842. PMLR, 2020.

[8] Miles Cranmer, Daniel Tamayo, Hanno Rein, Peter Battaglia, Samuel Hadden, Philip J Armitage, Shirley Ho, and David N Spergel. A bayesian neural network predicts the dissolution of compact planetary systems. *Proceedings of the National Academy of Sciences*, 118(40), 2021.

[9] Wei Deng, Xiao Zhang, Faming Liang, and Guang Lin. An adaptive empirical bayesian method for sparse deep learning. *Advances in neural information processing systems*, 2019:5563, 2019.

[10] Michael Dusenberry, Ghassen Jerfel, Yeming Wen, Yian Ma, Jasper Snoek, Katherine Heller, Balaji Lakshminarayanan, and Dustin Tran. Efficient and scalable bayesian neural nets with rank-1 factors. In *International conference on machine learning*, pages 2782–2792. PMLR, 2020.

[11] Piero Esposito. Blitz - bayesian layers in torch zoo (a bayesian deep learing library for torch). `https://github.com/piEsposito/blitz-bayesian-deep-learning/`, 2020.

[12] Angelos Filos, Sebastian Farquhar, Aidan N Gomez, Tim GJ Rudner, Zachary Kenton, Lewis Smith, Milad Alizadeh, Arnoud De Kroon, and Yarin Gal. A systematic comparison of bayesian deep learning robustness in diabetic retinopathy tasks. *arXiv preprint arXiv:1912.10481*, 2019.

[13] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *International Conference on Learning Representations*, 2018.

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

[15] Zhe Gan, Changyou Chen, Ricardo Henao, David Carlson, and Lawrence Carin. Scalable deep poisson factor analysis for topic modeling. In *International Conference on Machine Learning*, pages 1823–1832. PMLR, 2015.

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

[17] Soumya Ghosh, Jiayu Yao, and Finale Doshi-Velez. Model selection in bayesian neural networks via horseshoe priors. *J. Mach. Learn. Res.*, 20(182):1–46, 2019.

[18] Tilmann Gneiting and Adrian E Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association*, 102(477):359–378, 2007.

[19] Will Grathwohl, Kevin Swersky, Milad Hashemi, David Duvenaud, and Chris Maddison. Oops i took a gradient: Scalable sampling for discrete distributions. In *International Conference on Machine Learning*, pages 3831–3841. PMLR, 2021.

[20] Alex Graves. Practical variational inference for neural networks. *Advances in neural information processing systems*, 24, 2011.

[21] László Györfi, Michael Kohler, Adam Krzyżak, and Harro Walk. *A distribution-free theory of nonparametric regression*, volume 1. Springer, 2002.

[22] Song Han, Jeff Pool, John Tran, and William Dally. Learning both weights and connections for efficient neural network. *Advances in neural information processing systems*, 28, 2015.

[23] Nick Harvey, Christopher Liaw, and Abbas Mehrabian. Nearly-tight vc-dimension bounds for piecewise linear neural networks. In *Conference on learning theory*, pages 1064–1068. PMLR, 2017.

[24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.

[25] Yihui He, Xiangyu Zhang, and Jian Sun. Channel pruning for accelerating very deep neural networks. In *Proceedings of the IEEE international conference on computer vision*, pages 1389–1397, 2017.

[26] Jonathan Heek and Nal Kalchbrenner. Bayesian inference for large scale image classification. *arXiv preprint arXiv:1908.03491*, 2019.

[27] José Miguel Hernández-Lobato and Ryan Adams. Probabilistic backpropagation for scalable learning of bayesian neural networks. In *International conference on machine learning*, pages 1861–1869. PMLR, 2015.

[28] Matthew D Hoffman, Andrew 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.

[29] Pavel Izmailov, Sharad Vikram, Matthew D Hoffman, and Andrew Gordon Gordon Wilson. What are bayesian neural network posteriors really like? In *International Conference on Machine Learning*, pages 4629–4640. PMLR, 2021.

[30] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[31] Durk P Kingma, Tim Salimans, and Max Welling. Variational dropout and the local reparameterization trick. *Advances in neural information processing systems*, 28:2575–2583, 2015.

[32] Michael Kohler and Sophie Langer. On the rate of convergence of fully connected deep neural network regression estimates. *The Annals of Statistics*, 49(4):2231–2249, 2021.

[33] Ananya Kumar, Percy S Liang, and Tengyu Ma. Verified uncertainty calibration. *Advances in Neural Information Processing Systems*, 32, 2019.

[34] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017.

[35] Namhoon Lee, Thalaiyasingam Ajanthan, and Philip Torr. Snip: Single-shot network pruning based on connection sensitivity. In *International Conference on Learning Representations*, 2018.

[36] Chunyuan Li, Changyou Chen, David Carlson, and Lawrence Carin. Preconditioned stochastic gradient langevin dynamics for deep neural networks. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.

[37] Zhuang Liu, Jianguo Li, Zhiqiang Shen, Gao Huang, Shoumeng Yan, and Changshui Zhang. Learning efficient convolutional networks through network slimming. In *Proceedings of the IEEE international conference on computer vision*, pages 2736–2744, 2017.

[38] Christos Louizos and Max Welling. Multiplicative normalizing flows for variational bayesian neural networks. In *International Conference on Machine Learning*, pages 2218–2227. PMLR, 2017.

[39] Christos Louizos, Karen Ullrich, and Max Welling. Bayesian compression for deep learning. *Advances in neural information processing systems*, 30, 2017.

[40] David JC MacKay. A practical bayesian framework for backpropagation networks. *Neural computation*, 4(3):448–472, 1992.

[41] Dmitry Molchanov, Arsenii Ashukha, and Dmitry Vetrov. Variational dropout sparsifies deep neural networks. In *International Conference on Machine Learning*, pages 2498–2507. PMLR, 2017.

[42] Radford M Neal. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012.

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

[44] Ilsang Ohn and Lizhen Lin. Adaptive variational bayes: Optimality, computation and applications. *arXiv preprint arXiv:2109.03204*, 2021.

[45] Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, David Sculley, Sebastian Nowozin, Joshua Dillon, Balaji Lakshminarayanan, and Jasper Snoek. Can you trust your model's uncertainty? evaluating predictive uncertainty under dataset shift. *Advances in neural information processing systems*, 32, 2019.

[46] Nicholas G Polson and Veronika Ročková. Posterior concentration for sparse deep learning. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 938–949, 2018.

[47] Rahul Rahaman et al. Uncertainty quantification and deep ensembles. *Advances in Neural Information Processing Systems*, 34, 2021.

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

[49] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.

[50] Jakub Swiatkowski, Kevin Roth, Bastiaan Veeling, Linh Tran, Joshua Dillon, Jasper Snoek, Stephan Mandt, Tim Salimans, Rodolphe Jenatton, and Sebastian Nowozin. The k-tied normal distribution: A compact parameterization of gaussian mean field posteriors in bayesian neural networks. In *International Conference on Machine Learning*, pages 9289–9299. PMLR, 2020.

[51] Alexandre B Tsybakov. Introduction to nonparametric estimation, 2009.

[52] CJ Umrigar. Accelerated metropolis method. *Physical review letters*, 71(3):408, 1993.

[53] Aad W van der Vaart and J Harry van Zanten. Rates of contraction of posterior distributions based on gaussian process priors. *The Annals of Statistics*, 36(3):1435–1463, 2008.

[54] Chaoqi Wang, Roger Grosse, Sanja Fidler, and Guodong Zhang. Eigendamage: Structured pruning in the kronecker-factored eigenbasis. In *International Conference on Machine Learning*, pages 6566–6575. PMLR, 2019.

[55] Hao Wang, Naiyan Wang, and Dit-Yan Yeung. Collaborative deep learning for recommender systems. In *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1235–1244, 2015.

[56] Yating Wang, Wei Deng, and Guang Lin. Bayesian sparse learning with preconditioned stochastic gradient mcmc and its applications. *Journal of Computational Physics*, 432:110134, 2021.

[57] Max Welling and Yee 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.

[58] Wei Wen, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li. Learning structured sparsity in deep neural networks. *Advances in neural information processing systems*, 29:2074–2082, 2016.

[59] Florian Wenzel, Kevin Roth, Bastiaan Veeling, Jakub Swiatkowski, Linh Tran, Stephan Mandt, Jasper Snoek, Tim Salimans, Rodolphe Jenatton, and Sebastian Nowozin. How good is the bayes posterior in deep neural networks really? In *International Conference on Machine Learning*, pages 10248–10259. PMLR, 2020.

[60] Andrew G Wilson and Pavel Izmailov. Bayesian deep learning and a probabilistic perspective of generalization. *Advances in neural information processing systems*, 33:4697–4708, 2020.

[61] Fangzheng Xie and Yanxun Xu. Adaptive bayesian nonparametric regression using a kernel mixture of polynomials with application to partial linear models. *Bayesian Analysis*, 15(1): 159–186, 2020.

[62] Jiayu Yao, Weiwei Pan, Soumya Ghosh, and Finale Doshi-Velez. Quality of uncertainty quantification for bayesian neural network inference. *arXiv preprint arXiv:1906.09686*, 2019.

[63] Giacomo Zanella. Informed proposals for local mcmc in discrete spaces. *Journal of the American Statistical Association*, 115(530):852–865, 2020.

[64] Ruqi Zhang, Chunyuan Li, Jianyi Zhang, Changyou Chen, and Andrew Gordon Wilson. Cyclical stochastic gradient mcmc for bayesian deep learning. In *International Conference on Learning Representations*, 2019.

[65] Shaofeng Zhang, Meng Liu, and Junchi Yan. The diversified ensemble neural network. *Advances in Neural Information Processing Systems*, 33:16001–16011, 2020.

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 7.
   (c) Did you discuss any potential negative societal impacts of your work? [N/A] Our work does not have obvious direct negative social impacts.
   (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 5.
   (b) Did you include complete proofs of all theoretical results? [Yes] See Appendix A.
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] See Supplemental Material. It will be released as open-source after publication.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Section 6 and Appendix B.2.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See Section 6.
   (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] See Appendix B.2.
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]

Outline of the Supplementary Material

The Appendix is organized as follows.

- Appendix A provides the rigorous proofs of theoretical results in Sections 5.
  - Appendix A.1 provides additional notations.
  - Appendix A.2 states auxiliary lemmas for the proofs of the main theorems.
  - Appendix A.3 provides the proof of Theorem 5.1.
  - Appendix A.4 provides the proof of Theorem 5.2.
- Appendix B provides more experimental details.
  - Appendix B.1 provides the detailed description of mBNN.
  - Appendix B.2 provides the experimental setup.
  - Appendix B.3 provides additional experimental results.

A Proofs for Main Theorems

A.1 Additional notations

In this section, we describe additional notations not mentioned earlier.

We define $\mathcal{F}^{\mathrm{DNN}}(L, \boldsymbol{p}, F)$ as the function class of truncated DNNs with $(L, \boldsymbol{p})$ architecture :

$$\mathcal{F}^{\mathrm{DNN}}(L, \boldsymbol{p}, F) := \big\{ f : f = f^{\mathrm{DNN}}_{\boldsymbol{\theta}[-F,F]} \text{ is a DNN with } (L, \boldsymbol{p}) \\ \text{architecture truncated on } [-F, F] \big\}.$$

In a similar fashion, we define $\mathcal{F}^{\mathrm{mDNN}}(L, \boldsymbol{p}, F)$ as the function class of truncated mDNNs with $(L, \boldsymbol{p})$ architecture :

$$\mathcal{F}^{\mathrm{mDNN}}(L, \boldsymbol{p}, F) := \big\{ f : f = f^{\mathrm{mDNN}}_{\boldsymbol{M},\boldsymbol{\theta}[-F,F]} \text{ is a mDNN with } (L, \boldsymbol{p}) \\ \text{architecture truncated on } [-F, F] \big\}.$$

and $\mathcal{F}^{\mathrm{mDNN}}(L, \boldsymbol{p}, F, s)$ as

$$\mathcal{F}^{\mathrm{mDNN}}(L, \boldsymbol{p}, F, s) := \big\{ f : f = f^{\mathrm{mDNN}}_{\boldsymbol{M},\boldsymbol{\theta}[-F,F]} \in \mathcal{F}^{\mathrm{mDNN}}(L, \boldsymbol{p}, F), \ \max_{l \in [L]} |\boldsymbol{m}^{(l)}|_0 \le s \big\}.$$

For a real number $x \in \mathbb{R}$, we denote $\lceil x \rceil := \min\{z \in \mathbb{Z} : z \ge x\}$. For a vector $\boldsymbol{x} \in \mathbb{R}^d$ and m dimensional index vector $\boldsymbol{j} \subset [d]$, we denote $(\boldsymbol{x})_{\boldsymbol{j}} \in \mathbb{R}^m$ as the sub-vector whose elements are consist of $\boldsymbol{j}$th index of $\boldsymbol{x}$. Let $\boldsymbol{x}^{(n)} := (\boldsymbol{x}_i)_{i=1}^n$, $\boldsymbol{X}^{(n)} := (\boldsymbol{X}_i)_{i=1}^n$. For a real-valued function $f : \mathcal{X} \to \mathbb{R}$ and $1 \le p < \infty$, we denote $||f||_{p,n} := (\sum_{i=1}^n f(\boldsymbol{x}_i)^p / n)^{1/p}$ and $||f||_{p,\mathrm{P}_{\boldsymbol{X}}} := \left(\int_{\boldsymbol{X} \in \mathcal{X}} f(\boldsymbol{X})^p d\mathrm{P}_{\boldsymbol{X}}\right)^{1/p}$ where $\mathrm{P}_{\boldsymbol{X}}$ is a probability measure defined on input space $\mathcal{X}$.

For two positive sequences $\{a_n\}$ and $\{b_n\}$, we denote $a_n \lesssim b_n$ if there exists a positive sequence $C > 0$ such that $a_n \le C b_n$ for all $n \in \mathbb{N}$. We denote $a_n \asymp b_n$ if $a_n \lesssim b_n$ and $a_n \gtrsim b_n$ hold. We use the little o notation, that is, we write $a_n = o(b_n)$ if $\lim_{n\to\infty} a_n / b_n = 0$.

Let $\mathcal{F}$ be a set of functions $\mathcal{X} \to \mathbb{R}$ and d_n be a semimetric defined on $\mathcal{F}$. We denote $\mathcal{N}(\varepsilon, \mathcal{F}, d_n)$ and $\mathcal{M}(\varepsilon, \mathcal{F}, d_n)$ as the ε-covering number and ε-packing number of $\mathcal{F}$ w.r.t. d_n, respectively. Also, we denote $V^+_{\mathcal{F}}$ as the VC dimension of the set $\mathcal{F}^+ := \{\{(x,t) \in \mathcal{X} \times \mathbb{R}; t \le f(x)\}; f \in \mathcal{F}\}$.

A.2 Auxiliary lemmas

First, we describe a lemma that approximates Hölder smooth functions as DNN functions.

Lemma A.1 (Theorem 2 of [32])**.** *There exists $C_L > 0$ and $C_p > 0$ only depending on d such that for every $f_0 \in \mathcal{H}^\beta_d$ with $\|f_0\|_\infty \le F$, there exist $f^{\mathrm{DNN}}_{\hat{\boldsymbol{\theta}}\ [-F,F]} \in \mathcal{F}^{\mathrm{DNN}}(L_n, \boldsymbol{v}_n, F)$ with*

$$\begin{aligned} L_n &:= \lceil C_L \log n \rceil, \\ v_n &:= \left\lceil C_p \left(n^{\frac{d}{2\beta+d}} (\log n)^{-1} \right)^{1/2} \right\rceil, \\ \boldsymbol{v}_n &:= (d, v_n, ..., v_n, 1)^\top \in \mathbb{N}^{L_n+2}, \end{aligned}$$

such that

$$\left\| f^{\mathrm{DNN}}_{\hat{\boldsymbol{\theta}}\ [-F,F]} - f_0 \right\|_\infty \lesssim n^{-\frac{\beta}{2\beta+d}} \log n$$

and

$$|\hat{\boldsymbol{\theta}}|_\infty \le n \tag{A.1}$$

hold.

Note that the upper bound (A.1) is not mentioned in statement of [32], but it can be easily confirmed by following their proof. Next, we describe a standard tool for establishing concentration rates.

Lemma A.2 (Theorem 4 of [16]). *Let* $\left(\mathfrak{Y}^{(n)}, \mathcal{A}^{(n)}, P_\eta^{(n)} : \eta \in \mathcal{F}_n\right)$ *be a sequence statistical experiments with observations* $Y^{(n)}$. *We consider the case where the observation* $Y^{(n)}$ *is a vector* $Y^{(n)} = (Y_1, \ldots, Y_n)^\top$ *of independent observations* Y_i. *We assume that the distribution* $P_{\eta,i}$ *of the* i*th component* Y_i *possesses a density* $p_{\eta,i}$ *relative to Lebesgue measure for* $i \in [n]$. *We define*

$$K_i(\eta_0, \eta) = \int \log(p_{\eta_0,i}/p_{\eta,i}) dP_{\eta_0,i},$$
$$V_{2,0;i}(\eta_0, \eta) = \int \left(\log(p_{\eta_0,i}/p_{\eta,i}) - K_i(\eta_0, \eta)\right)^2 dP_{\eta_0,i},$$

and

$$B_n^*(\eta_0, \varepsilon_n; 2) = \left\{\eta \in \mathcal{F}_n : \frac{1}{n}\sum_{i=1}^n K_i(\eta_0, \eta) \le \varepsilon_n^2,\ \frac{1}{n}\sum_{i=1}^n V_{2,0;i}(\eta_0, \eta) \le \varepsilon_n^2\right\}.$$

Let h_n *be a semimetric on* $\mathcal{F}_n$ *with the property that there exist universal constants* $\xi > 0$ *and* $K > 0$ *such that for every* $\varepsilon > 0$ *and for each* $\eta_1 \in \mathcal{F}_n$ *with* $h_n(\eta_1, \eta_0) > \varepsilon$, *there exists a test* ϕ_n *such that*

$$P_{\eta_0}^{(n)}\phi_n \le e^{-Kn\varepsilon^2}, \quad \sup_{\eta_2 \in \mathcal{F}_n : h_n(\eta_2, \eta_1) < \varepsilon\xi} P_{\eta_2}^{(n)}(1 - \phi_n) \le e^{-Kn\varepsilon^2}.$$

Let $\varepsilon_n > 0, \varepsilon_n \to 0$ *and* $(n\varepsilon_n^2)^{-1} = O(1)$. *If for every sufficiently large* $j \in \mathbb{N}$,

$$\sup_{\varepsilon > \varepsilon_n} \log N\left(\frac{1}{2}\varepsilon\xi, \{\eta \in \mathcal{F}_n : h_n(\eta, \eta_0) < \varepsilon\}, h_n\right) \le n\varepsilon_n^2,$$
$$\frac{\Pi_n(\eta \in \mathcal{F}_n : j\varepsilon_n < h_n(\eta, \eta_0) \le 2j\varepsilon_n)}{\Pi_n(B_n(\eta_0, \varepsilon_n; 2))} \le e^{Kn\varepsilon_n^2 j^2/2}$$

for all but finite many n, *then we have that*

$$P_{\eta_0}^{(n)}\Pi_n\left(\eta \in \mathcal{F}_n : h_n(\eta, \eta_0) \ge M_n\varepsilon_n \mid X^{(n)}\right) \to 0$$

for every $M_n \to \infty$.

Next, we state the lemma which describes an upper bound of supremum norm distance of two DNN functions whose parameters are similar.

Lemma A.3. *Consider two DNN models* $f_{\boldsymbol{\theta}_1}^{\mathrm{DNN}} : [-1,1]^d \to \mathbb{R}$, $f_{\boldsymbol{\theta}_2}^{\mathrm{DNN}} : [-1,1]^d \to \mathbb{R}$ *with* $(L, \boldsymbol{p})$ *architecture, where* $L \in \mathbb{N}$ *and* $\boldsymbol{p} = (d, p, p, ..., p, 1)^\top \in \mathbb{N}^{L+2}$ *for some* $p \in \mathbb{N}$. *If* $|\boldsymbol{\theta}_1|_\infty \le B$, $|\boldsymbol{\theta}_2|_\infty \le B$ *and* $|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2|_\infty \le \delta$ *holds for some* $B > 0$ *and* $\delta > 0$, *then*

$$\|f_{\boldsymbol{\theta}_1}^{\mathrm{DNN}} - f_{\boldsymbol{\theta}_2}^{\mathrm{DNN}}\|_\infty \le dp^L B^{L+1}(L+1)\delta$$

holds.

proof of Lemma A.3. Define

$$f_{\boldsymbol{\theta}_1}^{\mathrm{DNN}}(\cdot) = A_{L+1,1} \circ \rho \circ A_{L,1} \cdots \circ \rho \circ A_{1,1}(\cdot),$$
$$f_{\boldsymbol{\theta}_2}^{\mathrm{DNN}}(\cdot) = A_{L+1,2} \circ \rho \circ A_{L,2} \cdots \circ \rho \circ A_{1,2}(\cdot).$$

Also, we define $\boldsymbol{h}_{\boldsymbol{\theta},L'} : [-1,1]^d \to \mathbb{R}^{p_{L'}}$ for $L' \in [L]$ as the DNN model whose output is L'-th hidden layer of $f_{\boldsymbol{\theta}}^{\mathrm{DNN}}$. In other words,

$$\boldsymbol{h}_{\boldsymbol{\theta}_1,L'}(\cdot) = A_{L',1} \circ \rho \circ A_{L'-1,1} \cdots \circ \rho \circ A_{1,1}(\cdot),$$
$$\boldsymbol{h}_{\boldsymbol{\theta}_2,L'}(\cdot) = A_{L',2} \circ \rho \circ A_{L'-1,2} \cdots \circ \rho \circ A_{1,2}(\cdot).$$

We let $\boldsymbol{h}_{\boldsymbol{\theta},L+1}(\cdot) = f_{\boldsymbol{\theta}}^{\mathrm{DNN}}(\cdot)$. Since

$$\left\||\boldsymbol{h}_{\boldsymbol{\theta}_1,L'+1} - \boldsymbol{h}_{\boldsymbol{\theta}_2,L'+1}|_\infty\right\|_\infty \le p|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2|_\infty \left\||\boldsymbol{h}_{\boldsymbol{\theta}_1,L'}|_\infty\right\|_\infty + p|\boldsymbol{\theta}_2|_\infty \left\||\boldsymbol{h}_{\boldsymbol{\theta}_1,L'} - \boldsymbol{h}_{\boldsymbol{\theta}_2,L'}|_\infty\right\|_\infty$$

and

$$\left\||\boldsymbol{h}_{\boldsymbol{\theta}_1,L'}|_\infty\right\|_\infty \le dp^{L'-1}B^{L'}$$

hold, we can show that $|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2|_\infty \le \delta$ implies

$$\left\||\boldsymbol{h}_{\boldsymbol{\theta}_1,L'}(x) - \boldsymbol{h}_{\boldsymbol{\theta}_2,L'}(x)|_\infty\right\|_\infty \le dp^{L'-1}B^{L'}L'\delta$$

for every $L' \in [L+1]$ by recursion. □

Lastly, we state the lemma about empirical process theory.

Lemma A.4 (Theorem 19.3 of [21]). *Let $\boldsymbol{X}, \boldsymbol{X}_1, \ldots, \boldsymbol{X}_n$ be independent and identically distributed random vectors with values in $\mathbb{R}^d$. Let $K_1, K_2 \geq 1$ be constants and let $\mathcal{G}$ be a class of functions $g : \mathbb{R}^d \to \mathbb{R}$ with*

$$|g(\boldsymbol{x})| \leq K_1, \quad \mathbb{E}(g(\boldsymbol{X})^2) \leq K_2 \mathbb{E}(g(\boldsymbol{X})).$$

Let $0 < \kappa < 1$ and $\alpha > 0$. Assume that

$$\sqrt{n}\kappa\sqrt{1-\kappa}\sqrt{\alpha} \geq 288 \max\left\{2K_1, \sqrt{2K_2}\right\}$$

and that, for all $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n \in \mathbb{R}^d$ and for all $t \geq \frac{\alpha}{8}$,

$$\frac{\sqrt{n}\kappa(1-\kappa)t}{96\sqrt{2}\max\{K_1, 2K_2\}} \geq \int_{\frac{\kappa(1-\kappa)t}{16\max\{K_1,2K_2\}}}^{\sqrt{t}} \sqrt{\log \mathcal{N}\left(u, \left\{g \in \mathcal{G} : \frac{1}{n}\sum_{i=1}^n g\left(\boldsymbol{x}_i\right)^2 \leq 16t\right\}, ||\cdot||_{1,n}\right)} du.$$

Then,

$$\mathbf{P}\left\{\sup_{g\in\mathcal{G}} \frac{\left|\mathbb{E}\{g(\boldsymbol{X})\} - \frac{1}{n}\sum_{i=1}^n g\left(\boldsymbol{X}_i\right)\right|}{\alpha + \mathbb{E}\{g(\boldsymbol{X})\}} > \kappa\right\} \leq 60 \exp\left(-\frac{n\alpha\kappa^2(1-\kappa)}{128 \cdot 2304 \max\{K_1^2, K_2\}}\right).$$

A.3 Proof of Theorem 5.1

Let $\tau := \gamma - \frac{5}{2}$. It is enough to show the main statement for $0 < \tau < 1$. Note that $\varepsilon_n = n^{-\frac{\beta}{2\beta+d}}(\log n)^\gamma = n^{-\frac{\beta}{2\beta+d}}(\log n)^{\tau+\frac{5}{2}}$. For C_p defined in Lemma A.1, we define s_n as

$$s_n := \left\lceil C_p \left(n^{\frac{d}{2\beta+d}}(\log n)^{3\tau}\right)^{1/2} \right\rceil$$

and $\boldsymbol{s}_n = (d, s_n, \ldots, s_n, 1)^\top \in \mathbb{N}^{L_n+2}$. We define

$$T_n := (d+1)p_n + (L_n - 1)p_n(p_n+1) + (p_n+1)$$

and

$$S_n := (d+1)s_n + (L_n - 1)s_n(s_n+1) + (s_n+1)$$

as the numbers of parameters in the DNNs with $(L_n, \boldsymbol{p}_n)$ and $(L_n, \boldsymbol{s}_n)$ architectures, respectively. Let $\mathcal{F}_n$ be the set of pairs of truncated mDNN with $(L_n, \boldsymbol{p}_n)$ architecture and variances of the Gaussian noise,

$$\mathcal{F}_n := \left\{(f, \sigma^2)^\top \ : \ f \in \mathcal{F}^{\mathrm{mDNN}}(L_n, \boldsymbol{p}_n, F),\ 0 < \sigma^2 \leq \sigma^2_{\max}\right\}.$$

Also, we let $\mathcal{F}'_n \subset \mathcal{F}_n$ by

$$\mathcal{F}'_n := \left\{(f, \sigma^2)^\top \ : \ f \in \mathcal{F}^{\mathrm{mDNN}}(L_n, \boldsymbol{p}_n, F, s_n),\ 0 < \sigma^2 \leq \sigma^2_{\max}\right\}.$$

In the first step of the proof, we fix $\{\boldsymbol{x}^{(n)}\}_{n=1}^\infty$ and show

$$\mathbb{E}_0\left[\Pi_n\left((f, \sigma^2)^\top \in \mathcal{F}'_n : ||f - f_0||_{2,n} + |\sigma^2 - \sigma_0^2| > M_n\varepsilon_n \middle| \mathcal{D}^{(n)}\right) \middle| \boldsymbol{X}^{(n)} = \boldsymbol{x}^{(n)}\right] \to 0 \quad \text{(A.2)}$$

as $n \to \infty$ for any $M_n \to \infty$.

In the second step of proof, we extend empirical L_2 error to expected L_2 error. In other words, we show

$$\mathbb{E}_0\left[\Pi_n\left((f, \sigma^2)^\top \in \mathcal{F}'_n : ||f - f_0||_{2,\mathrm{P}_X} + |\sigma^2 - \sigma_0^2| > M_n\varepsilon_n \middle| \mathcal{D}^{(n)}\right)\right] \to 0 \quad \text{(A.3)}$$

as $n \to \infty$ for any $M_n \to \infty$.

In the last step of proof, we show

$$\mathbb{E}_0\left[\Pi_n\left((f, \sigma^2)^\top \in (\mathcal{F}_n \setminus \mathcal{F}'_n) \middle| \mathcal{D}^{(n)}\right)\right] \to 0. \quad \text{(A.4)}$$

and the proof of Theorem 5.1 is done by (A.3) and (A.4).

Step 1 For fixed $\{\boldsymbol{x}^{(n)}\}_{n=1}^{\infty}$, let $P_{(f,\sigma^2),i}$ and $p_{(f,\sigma^2),i}$ be the probability measure and density corresponding to Gaussian distribution $N(f(\boldsymbol{x}_i),\sigma^2)$, respectively. We define the semimetric h_n^2 on $\mathcal{F}_n'$ as the average of the squares of the Hellinger distances for the distributions of the n individual observations. In other words, for $(f_1,\sigma_1^2),(f_2,\sigma_2^2)\in\mathcal{F}_n'$,

$$h_n^2\left((f_1,\sigma_1^2),(f_2,\sigma_2^2)\right):=\frac{1}{n}\sum_{i=1}^{n}\int\left(\sqrt{p_{(f_1,\sigma_1^2),i}}-\sqrt{p_{(f_2,\sigma_2^2),i}}\right)^2 dP_{(f_1,\sigma_1^2),i}.$$

Note that h_n^2 satisfies

$$\begin{aligned}(||f_1-f_2||_{2,n}+|\sigma_1^2-\sigma_2^2|)^2&\le 2\left(||f_1-f_2||_{2,n}^2+|\sigma_1^2-\sigma_2^2|^2\right)\\&\lesssim h_n^2\left((f_1,\sigma_1^2),(f_2,\sigma_2^2)\right).\end{aligned}$$

by Lemma B.1 of [61]. Hence, to prove (A.2), it is suffices to show

$$\mathbb{E}_0\left[\Pi_n\left((f,\sigma^2)^\top\in\mathcal{F}_n':h_n\left((f,\sigma^2),(f_0,\sigma_0^2)\right)>M_n\varepsilon_n\middle|\mathcal{D}^{(n)}\right)\middle|\boldsymbol{X}^{(n)}=\boldsymbol{x}^{(n)}\right]\to 0. \tag{A.5}$$

Since Hellinger distance possesses an exponentially powerful local test with respect to both the type-I and type-II errors (Lemma 2 of [16]), we can use the standard tool to establish concentration rates that we restate in Lemma A.2 for the convenience of the reader.

If we define semimetric d_n on $\mathcal{F}_n'$ as

$$d_n^2\left((f_1,\sigma_1^2),(f_2,\sigma_2^2)\right):=||f_1-f_2||_{1,n}+|\sigma_1^2-\sigma_2^2|^2,$$

then $h_n^2(\cdot)\lesssim d_n^2(\cdot)$ holds by by Lemma B.1 of [61] and hence $\mathcal{N}\left(\varepsilon,\mathcal{F}_n',h_n\right)\le\mathcal{N}\left(\varepsilon^2,\mathcal{F}_n',d_n^2\right)$. Also, by the fact that $||f_1-f_2||_{1,n}\le\frac{\varepsilon^2}{2}$ and $|\sigma_1^2-\sigma_2^2|^2\le\frac{\varepsilon^2}{2}$ implies $||f_1-f_2||_{1,n}+|\sigma_1^2-\sigma_2^2|^2\le\varepsilon^2$ and for every $f^{\mathrm{mDNN}}_{\boldsymbol{M},\boldsymbol{\theta}\,[-F,F]}\in\mathcal{F}^{\mathrm{mDNN}}(L_n,\boldsymbol{p}_n,F,s_n)$ there exist $\boldsymbol{\psi}_{\boldsymbol{M},\boldsymbol{\theta}}\in\mathbb{R}^{S_n}$ and $f^{\mathrm{DNN}}_{\boldsymbol{\psi}_{\boldsymbol{M},\boldsymbol{\theta}}\,[-F,F]}\in\mathcal{F}^{\mathrm{DNN}}(L_n,\boldsymbol{s}_n,F)$ such that $f^{\mathrm{mDNN}}_{\boldsymbol{M},\boldsymbol{\theta}\,[-F,F]}=f^{\mathrm{DNN}}_{\boldsymbol{\psi}_{\boldsymbol{M},\boldsymbol{\theta}}\,[-F,F]}$ holds, we get

$$\begin{aligned}\mathcal{N}\left(\varepsilon,\mathcal{F}_n',h_n\right)\le&\mathcal{N}\left(\varepsilon^2,\mathcal{F}_n',d_n^2\right)\\\le&\mathcal{N}\left(\frac{\varepsilon^2}{2},\mathcal{F}^{\mathrm{mDNN}}(L_n,\boldsymbol{p}_n,F,s_n),||\cdot||_{1,n}\right)\frac{\sqrt{2}\sigma_{max}^2}{\varepsilon}\\\le&\mathcal{N}\left(\frac{\varepsilon^2}{2},\mathcal{F}^{\mathrm{DNN}}(L_n,\boldsymbol{s}_n,F),||\cdot||_{1,n}\right)\frac{\sqrt{2}\sigma_{max}^2}{\varepsilon}.\end{aligned}$$

Since functions in $\mathcal{F}^{DNN}(L_n,\boldsymbol{s}_n,F)$ are bounded by $[-F,F]$, there exists $c_1>0$ such that

$$\begin{aligned}\mathcal{N}\left(\frac{\varepsilon^2}{2},\mathcal{F}^{\mathrm{DNN}}(L_n,\boldsymbol{s}_n,F),||\cdot||_{1,n}\right)\le&\mathcal{M}\left(\frac{\varepsilon^2}{2},\mathcal{F}^{DNN}(L_n,\boldsymbol{s}_n,F),||\cdot||_{1,n}\right)\\\le&3\left(\frac{8eF}{\epsilon^2}\log\frac{12eF}{\epsilon^2}\right)^{V^+_{\mathcal{F}^{\mathrm{DNN}}(L_n,\boldsymbol{s}_n,F)}}\\\le&3\left(\frac{8eF}{\epsilon^2}\log\frac{12eF}{\epsilon^2}\right)^{c_1L_nS_n\log S_n}\end{aligned}$$

holds for every $\varepsilon>0$ by Theorem 9.4 of [21] and Theorem 6 of [23]. Hence,

$$\begin{aligned}\sup_{\varepsilon>\varepsilon_n}\log\mathcal{N}\left(\varepsilon,\mathcal{F}_n',h_n\right)\lesssim&L_nS_n\log S_n\log n\\\asymp&n^{\frac{d}{2\beta+d}}(\log n)^{4+3\tau}\\\le&n\varepsilon_n^2\end{aligned}\tag{A.6}$$

holds.

Now, we define

$$\begin{aligned}K_i((f_0,\sigma_0^2),(f,\sigma^2))&=\int\log(p_{(f_0,\sigma_0^2),i}/p_{(f,\sigma^2),i})dP_{(f_0,\sigma_0^2),i},\\V_{2,0;i}((f_0,\sigma_0^2),(f,\sigma^2))&=\int\left(\log(p_{(f_0,\sigma_0^2),i}/p_{(f,\sigma^2),i}-K_i((f_0,\sigma_0^2),(f,\sigma^2))\right)^2dP_{(f_0,\sigma_0^2),i}\end{aligned}$$

and

$$B_n^*\left((f_0,\sigma_0^2),\varepsilon;2\right)=\big\{(f,\sigma^2)\in\mathcal{F}_n':\frac{1}{n}\sum_{i=1}^n K_i\left((f_0,\sigma_0^2),(f,\sigma^2)\right)\le\varepsilon^2,\\ \frac{1}{n}\sum_{i=1}^n V_{2,0;i}\left((f_0,\sigma_0^2),(f,\sigma^2)\right)\le\varepsilon^2\big\}.$$

with notations on Lemma A.2. In addition, for $\varepsilon>0$, define

$$A_n^*\left((f_0,\sigma_0^2),\varepsilon;2\right):=\Big\{(f,\sigma^2)\in\mathcal{F}_n':\max_i|f(\boldsymbol{x}_i)-f_0(\boldsymbol{x}_i)|\le\frac{\sigma_0\varepsilon}{2},\\ \sigma^2\in[\sigma_0^2,(1+\varepsilon^2)\sigma_0^2]\Big\}.$$

Then for every $f\in A_n^*\left((f_0,\sigma_0^2),\varepsilon;2\right)$ and $i\in[n]$,

$$K_i((f_0,\sigma_0^2),(f,\sigma^2))=\frac{1}{2}\log\frac{\sigma^2}{\sigma_0^2}+\frac{\sigma_0^2+(f_0(\boldsymbol{x}_i)-f(\boldsymbol{x}_i))^2}{2\sigma^2}-\frac{1}{2}\le\varepsilon^2$$

and

$$\begin{aligned}V_{2,0;i}((f_0,\sigma_0^2),(f,\sigma^2))&=\mathrm{Var}_{f_0,\sigma_0^2}\left(-\frac{(Y_i-f_0(\boldsymbol{x}_i))^2}{2\sigma_0^2}+\frac{(Y_i-f(\boldsymbol{x}_i))^2}{2\sigma^2}\right)\\&=\mathrm{Var}_{f_0,\sigma_0^2}\left(-\frac{1}{2}(1-\frac{\sigma_0^2}{\sigma^2})Z_i^2+\frac{\sigma_0(f_0(\boldsymbol{x}_i)-f(\boldsymbol{x}_i))Z_i}{\sigma^2}\right)\le\varepsilon^2\end{aligned}$$

where $Y_i\sim N(f_0(\boldsymbol{x}_i),\sigma_0^2)$ and $Z_i:=\frac{Y_i-f_0(\boldsymbol{x}_i)}{\sigma_0}\sim N(0,1)$. Hence, we can conclude that

$$A_n^*\left((f_0,\sigma_0^2),\varepsilon_n;2\right)\subset B_n^*\left((f_0,\sigma_0^2),\varepsilon_n;2\right). \tag{A.7}$$

Now we define v_n as

$$v_n:=\left\lceil C_p\left(n^{\frac{d}{2\beta+d}}(\log n)^{-1}\right)^{1/2}\right\rceil,$$

which is the width of the network in Lemma A.1. Let $\boldsymbol{v}_n=(d,v_n,\dots,v_n,1)^\top\in\mathbb{N}^{L_n+2}$, and let

$$V_n:=(d+1)v_n+(L_n-1)v_n(v_n+1)+(v_n+1),$$

which is the number of parameters in DNN with $(L_n,\boldsymbol{v}_n)$ architecture. For any $\boldsymbol{\theta}\in\mathbb{R}^{T_n}$ and any $\boldsymbol{M}$ with $|\boldsymbol{m}^{(l)}|_0=v_n$ for $l\in[L]$, there exists V_n dimension index vector $\boldsymbol{j}_{\boldsymbol{M}}\subset[T_n]$ such that $f^{\mathrm{DNN}}_{\boldsymbol{\psi}_{\boldsymbol{M},\boldsymbol{\theta}}}=f^{\mathrm{mDNN}}_{\boldsymbol{M},\boldsymbol{\theta}}$ holds for $\boldsymbol{\psi}_{\boldsymbol{M},\boldsymbol{\theta}}:=(\boldsymbol{\theta})_{\boldsymbol{j}_{\boldsymbol{M}}}\in\mathbb{R}^{V_n}$. In other words, $f^{\mathrm{DNN}}_{\boldsymbol{\psi}_{\boldsymbol{M},\boldsymbol{\theta}}}$ is the sub-network of $f^{\mathrm{DNN}}_{\boldsymbol{M},\boldsymbol{\theta}}$ consisting of the unmasked nodes. Also, there exists $\hat{\boldsymbol{\psi}}\in[-n,n]^{V_n}$ such that

$$\left\|f^{\mathrm{DNN}}_{\hat{\boldsymbol{\psi}}\ [-F,F]}-f_0\right\|_\infty<\frac{\sigma_0\varepsilon_n}{4} \tag{A.8}$$

satisfies for large n by Lemma A.1.

With (A.7), (A.8) and Lemma A.3, we can obtain the lower bound of $\Pi_n\left(B_n^*\left((f_0,\sigma_0^2),\varepsilon_n;2\right)\right)$ by

$$
\begin{aligned}
&\Pi_n\left(B_n^*\left((f_0,\sigma_0^2),\varepsilon_n;2\right)\right)\\
\geq&\Pi_n\left(A_n^*\left((f_0,\sigma_0^2),\varepsilon_n;2\right)\right)\\
=&\Pi_n\Big(\Big\{(\boldsymbol{M},\boldsymbol{\theta}):\max_i|f^{\mathrm{mDNN}}_{\boldsymbol{M},\boldsymbol{\theta}\,[-F,F]}(\boldsymbol{x}_i)-f_0(\boldsymbol{x}_i)|\leq\frac{\sigma_0\varepsilon_n}{2}\Big\}\Big)\Pi_n\Big(\sigma^2\in[\sigma_0^2,(1+\varepsilon_n^2)\sigma_0^2]\Big)\\
\geq&\Pi_n\left(\left\{(\boldsymbol{M},\boldsymbol{\theta}):|\boldsymbol{m}^{(1)}|_0=\cdots=|\boldsymbol{m}^{(L)}|_0=v_n,\ \max_i\left|f^{\mathrm{DNN}}_{\boldsymbol{\psi}_{\boldsymbol{M},\boldsymbol{\theta}}\,[-F,F]}(\boldsymbol{x}_i)-f_0(\boldsymbol{x}_i)\right|\leq\frac{\sigma_0\varepsilon_n}{2}\right\}\right)\\
&\quad\times\Pi_n\Big(\sigma^2\in[\sigma_0^2,(1+\varepsilon_n^2)\sigma_0^2]\Big)\\
\geq&\Pi_n\left(\left\{(\boldsymbol{M},\boldsymbol{\theta}):|\boldsymbol{m}^{(1)}|_0=\cdots=|\boldsymbol{m}^{(L)}|_0=v_n,\ \max_i\left|f^{\mathrm{DNN}}_{\boldsymbol{\psi}_{\boldsymbol{M},\boldsymbol{\theta}}\,[-F,F]}(\boldsymbol{x}_i)-f^{\mathrm{DNN}}_{\hat{\boldsymbol{\psi}}\,[-F,F]}(\boldsymbol{x}_i)\right|\leq\frac{\sigma_0\varepsilon_n}{4}\right\}\right)\\
&\quad\times\Pi_n\Big(\sigma^2\in[\sigma_0^2,(1+\varepsilon_n^2)\sigma_0^2]\Big)\\
\geq&\Pi_n\left(\left\{(\boldsymbol{M},\boldsymbol{\theta}):|\boldsymbol{m}^{(1)}|_0=\cdots=|\boldsymbol{m}^{(L)}|_0=v_n,\ \left|\boldsymbol{\psi}_{\boldsymbol{M},\boldsymbol{\theta}}-\hat{\boldsymbol{\psi}}\right|_\infty\leq\frac{\sigma_0\varepsilon_n}{4dv_n^{L_n}n^{L_n+1}(L_n+1)}\right\}\right)\\
&\quad\times\Pi_n\Big(\sigma^2\in[\sigma_0^2,(1+\varepsilon_n^2)\sigma_0^2]\Big)\\
\gtrsim&\exp\left(-(\lambda\log n)^5v_n^2L_n\right)\exp\left(-V_n(\log n)^2\right)\varepsilon_n^2\\
\gtrsim&\exp\left(-\lambda^5C_p^2C_L(\log n)^5n^{\frac{d}{2\beta+d}}\right)\exp\left(-C_p^2C_L(\log n)^2n^{\frac{d}{2\beta+d}}\right)n^{-1}.
\end{aligned}
\tag{A.9}
$$

Hence we have that

$$
\Pi_n\left(B_n^*\left(\eta_0,\varepsilon_n;2\right)\right)\geq e^{-n\varepsilon_n^2}
\tag{A.10}
$$

for all but finite many n. Hence by (A.6), (A.10) and Lemma A.2, the proof of (A.5) is done. □

Step 2. Since (A.2) holds for arbitrary $\{\boldsymbol{x}^{(n)}\}_{n=1}^\infty$,

$$
\mathbb{E}_0\left[\Pi_n\left((f,\sigma^2)^\top\in\mathcal{F}_n':||f-f_0||_{2,n}+|\sigma^2-\sigma_0^2|>M_n\varepsilon_n\middle|\mathcal{D}^{(n)}\right)\right]\to0
\tag{A.11}
$$

also holds. Next, we will check the conditions in Lemma A.4 for

$$
\begin{aligned}
&\mathcal{G}:=\left\{g\ :\ g=(f^{\mathrm{mDNN}}_{\boldsymbol{M},\boldsymbol{\theta}\,[-F,F]}-f_0)^2,f^{\mathrm{mDNN}}_{\boldsymbol{M},\boldsymbol{\theta}\,[-F,F]}\in\mathcal{F}^{\mathrm{mDNN}}(L_n,\boldsymbol{p}_n,F,s_n)\right\},\\
&\kappa:=\frac{1}{2},\ \alpha:=\varepsilon_n^2,\ K_1=K_2=4F^2.
\end{aligned}
$$

First, it is easy to check $||g(\boldsymbol{x})||_\infty\leq4F^2$ and $\mathbb{E}(g(\boldsymbol{X})^2)\leq4F^2\mathbb{E}(g(\boldsymbol{X}))$ for $g\in\mathcal{G}$. Also, for every $f^{\mathrm{mDNN}}_{\boldsymbol{M},\boldsymbol{\theta}\,[-F,F]}\in\mathcal{F}^{\mathrm{mDNN}}(L_n,\boldsymbol{p}_n,F,s_n)$, there exist $\boldsymbol{\psi}_{\boldsymbol{M},\boldsymbol{\theta}}\in\mathbb{R}^{S_n}$ and $f^{\mathrm{DNN}}_{\boldsymbol{\psi}_{\boldsymbol{M},\boldsymbol{\theta}}\,[-F,F]}\in\mathcal{F}^{\mathrm{DNN}}(L_n,\boldsymbol{s}_n,F)$ such that $f^{\mathrm{mDNN}}_{\boldsymbol{M},\boldsymbol{\theta}\,[-F,F]}=f^{\mathrm{DNN}}_{\boldsymbol{\psi}_{\boldsymbol{M},\boldsymbol{\theta}}\,[-F,F]}$ holds. Since

$$
\left\|(f^{\mathrm{DNN}}_{\boldsymbol{\psi}_1\,[-F,F]}-f_0)^2-(f^{\mathrm{DNN}}_{\boldsymbol{\psi}_2\,[-F,F]}-f_0)^2\right\|_{n,1}\leq4F\left\|f^{\mathrm{DNN}}_{\boldsymbol{\psi}_1\,[-F,F]}-f^{\mathrm{DNN}}_{\boldsymbol{\psi}_2\,[-F,F]}\right\|_{n,1}
$$

holds for $\boldsymbol{\psi}_1,\boldsymbol{\psi}_2\in\mathbb{R}^{S_n}$, there exists $c_2>0$ such that

$$
\begin{aligned}
\mathcal{N}\left(u,\mathcal{G},||\cdot||_{n,1}\right)\leq&\mathcal{N}\left(\frac{u}{4F},\mathcal{F}^{\mathrm{DNN}}(L_n,\boldsymbol{s}_n,F),||\cdot||_{n,1}\right)\\
\leq&\mathcal{M}\left(\frac{u}{4F},\mathcal{F}^{\mathrm{DNN}}(L_n,\boldsymbol{s}_n,F),||\cdot||_{n,1}\right)\\
\leq&3\left(\frac{16eF^2}{u}\log\frac{24eF^2}{u}\right)^{V^+_{\mathcal{F}^{\mathrm{DNN}}(L_n,\boldsymbol{s}_n,F)}}\\
\lesssim&n^{c_2S_nL_n\log S_n}
\end{aligned}
$$

for $u\geq n^{-1}$ by Theorem 9.4 of [21] and Theorem 6 of [23]. Hence for all $t\geq\frac{\varepsilon_n^2}{8}$,

$$
\begin{aligned}
\int_{\frac{\kappa(1-\kappa)t}{16\max\{K_1,2K_2\}}}^{\sqrt{t}}\sqrt{\log\mathcal{N}\left(u,\mathcal{G},||\cdot||_{n,1}\right)}du\lesssim&\sqrt{t}\left(n^{\frac{d}{2\beta+d}}(\log n)^{4+3\tau}\right)^{\frac{1}{2}}\\
=&o\left(\frac{\sqrt{n}t/4}{96\sqrt{2}\max\{K_1,2K_2\}}\right)
\end{aligned}
$$

holds. To sum up, we conclude that

$$\mathbf{P}\left\{\sup_{f\in\mathcal{F}^{\mathrm{DNN}}(L_n,\boldsymbol{p}'_n)}\frac{\left|||f-f_0||^2_{2,\mathrm{P}_X}-||f-f_0||^2_{2,n}\right|}{\varepsilon_n^2+||f-f_0||^2_{2,\mathrm{P}_X}}>\frac{1}{2}\right\}\leq 60\exp\left(-\frac{n\varepsilon_n^2/8}{128\cdot 2304\cdot 16F^4}\right) \tag{A.12}$$

holds for all but finite many n by Lemma A.4. Hence by (A.11) and (A.12), the proof of (A.3) is done. □

Step 3. Since

$$\left(\frac{1}{2ks}-\frac{1}{4k^2s^3}\right)e^{-ks^2}\leq\int_s^\infty e^{-kt^2}dt\leq\frac{1}{2ks}e^{-ks^2}$$

for any $k>0$ and $s>0$,

$$\begin{aligned}\Pi_n(|\boldsymbol{m}^{(l)}|_0>s_n)\leq&\frac{\sum_{s=s_n+1}^{p_n}e^{-(\lambda\log n)^5s^2}}{e^{-(\lambda\log n)^5}}\\ \lesssim&e^{-(\lambda\log n)^5{s_n}^2}e^{(\lambda\log n)^5}(\log n)^{-5}\end{aligned}$$

holds for every $l\in[L]$. Then we have that

$$\begin{aligned}\Pi_n\left(\mathcal{F}_n\setminus\mathcal{F}'_n\right)=&\Pi_n\left(\bigcup_{l\in[L]}\{|\boldsymbol{m}^{(l)}|_0>s_n\}\right)\\ \leq&L_n\cdot\Pi_n\left(\{|\boldsymbol{m}^{(1)}|_0>s_n\}\right)\\ \lesssim&\exp\left(-\lambda^5C_p^2(\log n)^{5+3\tau}n^{\frac{d}{2\beta+d}}\right)\exp\left((\lambda\log n)^5\right)(\log n)^{-4}.\end{aligned} \tag{A.13}$$

By (A.9) and (A.13), we obtain

$$\begin{aligned}&\frac{\Pi_n\left(\mathcal{F}_n\setminus\mathcal{F}'_n\right)}{\Pi_n\left(B_n^\star\left(\eta_0,\varepsilon_n;2\right)\right)e^{-2n\varepsilon_n^2}}\\ &\lesssim\frac{\exp\left(-\lambda^5C_p^2(\log n)^{5+3\tau}n^{\frac{d}{2\beta+d}}\right)\exp\left((\lambda\log n)^5\right)}{\exp(-\lambda^5C_p^2C_L(\log n)^5n^{\frac{d}{2\beta+d}})\exp(-C_p^2C_L(\log n)^2n^{\frac{d}{2\beta+d}})n^{-1}\exp(-2(\log n)^{5+2\tau}n^{\frac{d}{2\beta+d}})}\\ &=o(1),\end{aligned}$$

and thus we have

$$\mathbb{E}_0\left[\Pi_n\left((f,\sigma^2)^\top\in\left(\mathcal{F}_n\setminus\mathcal{F}'_n\right)\middle|\mathcal{D}^{(n)}\right)\right]\to 0.$$

by apply Lemma 1 of [16], which completes the proof of (A.4). □

A.4 Proof of Theorem 5.2

The proof is a slight modification of the proof of Theorem 5.1. Let $\tau:=\gamma-\frac{5}{2}$. It suffices to show the main statement for $0<\tau<1$. Note that $\varepsilon_n=n^{-\frac{\beta}{2\beta+d}}(\log n)^\gamma=n^{-\frac{\beta}{2\beta+d}}(\log n)^{\tau+\frac{5}{2}}$. For C_p defined in Lemma A.1, we define s_n as

$$s_n:=\left\lceil C_p\left(n^{\frac{d}{2\beta+d}}(\log n)^{3\tau}\right)^{1/2}\right\rceil$$

and $\boldsymbol{s}_n=(d,s_n,\ldots,s_n,1)^\top\in\mathbb{N}^{L_n+2}$. We define

$$T_n:=(d+1)p_n+(L_n-1)p_n(p_n+1)+(p_n+1)$$

and

$$S_n:=(d+1)s_n+(L_n-1)s_n(s_n+1)+(s_n+1)$$

as the numbers of parameters in the DNNs with $(L_n, \boldsymbol{p}_n)$ and $(L_n, \boldsymbol{s}_n)$ architectures, respectively. Let $\mathcal{F}_n$ be the set of truncated mDNN with the $(L_n, \boldsymbol{p}_n)$ architecture,

$$\mathcal{F}_n := \mathcal{F}^{\text{mDNN}}(L_n, \boldsymbol{p}_n, F).$$

Also, we let $\mathcal{F}'_n \subset \mathcal{F}_n$ by

$$\mathcal{F}'_n := \mathcal{F}^{\text{mDNN}}(L_n, \boldsymbol{p}_n, F, s_n)$$

In the first step of proof, we fix $\{\boldsymbol{x}^{(n)}\}_{n=1}^{\infty}$ and show

$$\mathbb{E}_0\left[\Pi_n\left(f \in \mathcal{F}'_n : ||\phi \circ f - \phi \circ f_0||_{2,n} > M_n\varepsilon_n \middle| \mathcal{D}^{(n)}\right) \middle| \boldsymbol{X}^{(n)} = \boldsymbol{x}^{(n)}\right] \to 0 \tag{A.14}$$

as $n \to \infty$ for any $M_n \to \infty$.

In the second step of proof, we extend empirical L_2 error to expected L_2 error. In other words, we show

$$\mathbb{E}_0\left[\Pi_n\left(f \in \mathcal{F}'_n : ||\phi \circ f - \phi \circ f_0||_{2,\mathrm{P}_X} > M_n\varepsilon_n \middle| \mathcal{D}^{(n)}\right)\right] \to 0 \tag{A.15}$$

as $n \to \infty$ for any $M_n \to \infty$.

Then, since we already showed

$$\mathbb{E}_0\left[\Pi_n\left(f \in (\mathcal{F}_n \setminus \mathcal{F}'_n) \middle| \mathcal{D}^{(n)}\right)\right] \to 0 \tag{A.16}$$

in the last step of the proof of Theorem 5.1, the proof of Theorem 5.2 is done by (A.15) and (A.16).

Step 1 For fixed $\{\boldsymbol{x}^{(n)}\}_{n=1}^{\infty}$, let $P_{f,i}$ and $p_{f,i}$ be the probability measure and density corresponding to the Bernoulli distribution $\text{Ber}(\phi \circ f(\boldsymbol{x}_i))$, respectively. We define the semimetric h_n^2 on $\mathcal{F}'_n$ as the average of the squares of the Hellinger distances for the distributions of the n individual observations. In other words, for $f_1, f_2 \in \mathcal{F}'_n$,

$$h_n^2(f_1, f_2) := \frac{1}{n}\sum_{i=1}^{n}\int \left(\sqrt{p_{f_1,i}} - \sqrt{p_{f_2,i}}\right)^2 dP_{f_1,i}.$$

Also, we define semimetric d_n on $\mathcal{F}'_n$ as

$$d_n(f_1, f_2) := ||\phi \circ f_1 - \phi \circ f_2||_{2,n}.$$

Note that since $f_1, f_2 \in \mathcal{F}'_n$ are bounded,

$$\begin{aligned} d_n^2(f_1, f_2) =& \frac{1}{n}\sum_{i=1}^{n}(\phi \circ f_1(\boldsymbol{x}_i) - \phi \circ f_2(\boldsymbol{x}_i))^2 \\ =& \frac{1}{n}\sum_{i=1}^{n}\left(\sqrt{\phi \circ f_1(\boldsymbol{x}_i)} - \sqrt{\phi \circ f_2(\boldsymbol{x}_i)}\right)^2\left(\sqrt{\phi \circ f_1(\boldsymbol{x}_i)} + \sqrt{\phi \circ f_2(\boldsymbol{x}_i)}\right)^2 \\ \lesssim& h_n^2(f_1, f_2) \end{aligned}$$

and

$$\begin{aligned} d_n^2(f_1, f_2) =& \frac{1}{2n}\sum_{i=1}^{n}\left\{\left(\sqrt{\phi \circ f_1(\boldsymbol{x}_i)} - \sqrt{\phi \circ f_2(\boldsymbol{x}_i)}\right)^2\left(\sqrt{\phi \circ f_1(\boldsymbol{x}_i)} + \sqrt{\phi \circ f_2(\boldsymbol{x}_i)}\right)^2\right. \\ &\left. + \left(\sqrt{1-\phi \circ f_1(\boldsymbol{x}_i)} - \sqrt{1-\phi \circ f_2(\boldsymbol{x}_i)}\right)^2\left(\sqrt{1-\phi \circ f_1(\boldsymbol{x}_i)} + \sqrt{1-\phi \circ f_2(\boldsymbol{x}_i)}\right)^2\right\} \\ \gtrsim& h_n^2(f_1, f_2) \end{aligned}$$

holds. Hence, to prove (A.14), it suffices to show

$$\mathbb{E}_0\left[\Pi_n\left((f, \sigma^2)^\top \in \mathcal{F}'_n : h_n(f, f_0) > M_n\varepsilon_n \middle| \mathcal{D}^{(n)}\right) \middle| \boldsymbol{X}^{(n)} = \boldsymbol{x}^{(n)}\right] \to 0. \tag{A.17}$$

Since Hellinger distance possesses an exponentially powerful local test with respect to both the type-I and type-II errors (Lemma 2 of [16]), we can use standard tools to establish concentration rates that we restate in Lemma A.2 for the convenience of the reader.

Since ϕ is L1-Lipschitz function,

$$\begin{aligned}\mathcal{N}\left(\varepsilon, \mathcal{F}_n', h_n\right) \lesssim & \mathcal{N}\left(\varepsilon, \mathcal{F}_n', d_n\right) \\ \leq & \mathcal{N}\left(\varepsilon, \mathcal{F}_n', ||\cdot||_{2,n}\right).\end{aligned}$$

For every $f^{\text{mDNN}}_{\boldsymbol{M},\boldsymbol{\theta}\,[-F,F]} \in \mathcal{F}^{\text{mDNN}}(L_n, \boldsymbol{p}_n, F, s_n)$ there exist $\boldsymbol{\psi}_{\boldsymbol{M},\boldsymbol{\theta}} \in \mathbb{R}^{S_n}$ and $f^{\text{DNN}}_{\boldsymbol{\psi}_{\boldsymbol{M},\boldsymbol{\theta}}\,[-F,F]} \in \mathcal{F}^{\text{DNN}}(L_n, \boldsymbol{s}_n, F)$ such that $f^{\text{mDNN}}_{\boldsymbol{M},\boldsymbol{\theta}\,[-F,F]} = f^{\text{DNN}}_{\boldsymbol{\psi}_{\boldsymbol{M},\boldsymbol{\theta}}\,[-F,F]}$ holds. So we get

$$\begin{aligned}\mathcal{N}\left(\varepsilon, \mathcal{F}_n', ||\cdot||_{2,n}\right) = & \mathcal{N}\left(\varepsilon, \mathcal{F}^{\text{DNN}}(L_n, \boldsymbol{s}_n, F), ||\cdot||_{2,n}\right) \\ \leq & \mathcal{M}\left(\varepsilon, \mathcal{F}^{\text{DNN}}(L_n, \boldsymbol{s}_n, F), ||\cdot||_{2,n}\right).\end{aligned}$$

Since functions in $\mathcal{F}^{DNN}(L_n, \boldsymbol{s}_n, F)$ are bounded by $[-F, F]$, there exists $c_3 > 0$ such that

$$\begin{aligned}\mathcal{M}\left(\varepsilon, \mathcal{F}^{\text{DNN}}(L_n, \boldsymbol{s}_n, F), ||\cdot||_{2,n}\right) \leq & 3\left(\frac{8eF^2}{\epsilon^2}\log\frac{12eF^2}{\epsilon^2}\right)^{V^+_{\mathcal{F}^{\text{DNN}}(L_n,\boldsymbol{s}_n,F)}} \\ \leq & 3\left(\frac{8eF^2}{\epsilon^2}\log\frac{12eF^2}{\epsilon^2}\right)^{c_3 L_n S_n \log S_n}\end{aligned}$$

holds for every $\varepsilon > 0$ by Theorem 9.4 of [21] and Theorem 6 of [23]. To sum up,

$$\begin{aligned}\sup_{\varepsilon > \varepsilon_n} \log \mathcal{N}\Big(\varepsilon, \mathcal{F}_n', h_n\Big) \lesssim & L_n S_n \log S_n \log n \\ \asymp & n^{\frac{d}{2\beta+d}} (\log n)^{4+3\tau} \\ \leq & n\varepsilon_n^2 \end{aligned} \tag{A.18}$$

holds.

Now, we define

$$\begin{aligned}K_i(f_0, f) &= \int \log(p_{f_0,i}/p_{f,i}) dP_{f_0,i}, \\ V_{2,0;i}(f_0, f) &= \int \left(\log(p_{f_0,i}/p_{f,i}) - K_i(f_0, f)\right)^2 dP_{f_0,i},\end{aligned}$$

and

$$\begin{aligned}B_n^*\left(f_0, \varepsilon; 2\right) = \big\{ f \in \mathcal{F}_n' : & \frac{1}{n}\sum_{i=1}^n K_i\left(f_0, f\right) \leq \varepsilon^2, \\ & \frac{1}{n}\sum_{i=1}^n V_{2,0;i}\left(f_0, f\right) \leq \varepsilon^2 \big\}.\end{aligned}$$

For $\varepsilon > 0$, define

$$A_n^*\left(f_0, \varepsilon; 2\right) := \Big\{ f \in \mathcal{F}_n' : \max_i |f(\boldsymbol{x}_i) - f_0(\boldsymbol{x}_i)| \leq \varepsilon \Big\}.$$

Then by Lemma 3.2 of [53], we can get

$$A_n^*\left(f_0, \varepsilon_n; 2\right) \subset B_n^*\left(f_0, \varepsilon_n; 2\right).$$

Now by following the proof of (A.9), we obtain

$$\Pi_n\left(B_n^*\left(\eta_0, \varepsilon_n; 2\right)\right) \geq e^{-n\varepsilon_n^2} \tag{A.19}$$

for all but finite many n. Hence by (A.18), (A.19) and Lemma A.2, the proof of (A.17) is done. □

Step 2. Since (A.14) holds for arbitrary $\{\boldsymbol{x}^{(n)}\}_{n=1}^{\infty}$,

$$\mathbb{E}_0\left[\Pi_n\left(f \in \mathcal{F}_n' : ||\phi \circ f - \phi \circ f_0||_{2,n} > M_n\varepsilon_n \Big| \mathcal{D}^{(n)}\right)\right] \to 0 \tag{A.20}$$

also holds. Next, we will check the conditions in Lemma A.4 for

$$\mathcal{G} := \left\{ g \ : \ g = (\phi \circ f^{\mathrm{mDNN}}_{\boldsymbol{M},\boldsymbol{\theta}\,[-F,F]} - \phi \circ f_0)^2, f^{\mathrm{mDNN}}_{\boldsymbol{M},\boldsymbol{\theta}\,[-F,F]} \in \mathcal{F}^{\mathrm{mDNN}}(L_n, \boldsymbol{p}_n, F, s_n) \right\}$$
$$\kappa := \frac{1}{2}, \ \alpha := \varepsilon_n^2, \ K_1 = K_2 = 1.$$

First, it is easy to check $||g(\boldsymbol{x})||_\infty \leq 1$ and $E(g(\boldsymbol{X})^2) \leq E(g(\boldsymbol{X}))$ for $g \in \mathcal{G}$. Also, for every $f^{\mathrm{mDNN}}_{\boldsymbol{M},\boldsymbol{\theta}\,[-F,F]} \in \mathcal{F}^{\mathrm{mDNN}}(L_n, \boldsymbol{p}_n, F, s_n)$, there exist $\boldsymbol{\psi}_{\boldsymbol{M},\boldsymbol{\theta}} \in \mathbb{R}^{S_n}$ and $f^{\mathrm{DNN}}_{\boldsymbol{\psi}_{\boldsymbol{M},\boldsymbol{\theta}}\,[-F,F]} \in \mathcal{F}^{\mathrm{DNN}}(L_n, \boldsymbol{s}_n, F)$ such that $f^{\mathrm{mDNN}}_{\boldsymbol{M},\boldsymbol{\theta}\,[-F,F]} = f^{\mathrm{DNN}}_{\boldsymbol{\psi}_{\boldsymbol{M},\boldsymbol{\theta}}\,[-F,F]}$ holds. Since

$$\begin{aligned}
&\left\| (\phi \circ f^{\mathrm{DNN}}_{\boldsymbol{\psi}_1\,[-F,F]} - \phi \circ f_0)^2 - (\phi \circ f^{\mathrm{DNN}}_{\boldsymbol{\psi}_2\,[-F,F]} - \phi \circ f_0)^2 \right\|_{n,1} \\
&\qquad \leq 4 \left\| \phi \circ f^{\mathrm{DNN}}_{\boldsymbol{\psi}_1\,[-F,F]} - \phi \circ f^{\mathrm{DNN}}_{\boldsymbol{\psi}_2\,[-F,F]} \right\|_{n,1} \\
&\qquad \leq 4 \left\| f^{\mathrm{DNN}}_{\boldsymbol{\psi}_1\,[-F,F]} - f^{\mathrm{DNN}}_{\boldsymbol{\psi}_2\,[-F,F]} \right\|_{n,1}
\end{aligned}$$

holds for $\boldsymbol{\psi}_1, \boldsymbol{\psi}_2 \in \mathbb{R}^{S_n}$, there exists $c_4 > 0$ such that

$$\begin{aligned}
\mathcal{N}(u, \mathcal{G}, ||\cdot||_{n,1}) \leq & \mathcal{N}\left(\frac{u}{4}, \mathcal{F}^{\mathrm{DNN}}(L_n, \boldsymbol{s}_n, F), ||\cdot||_{n,1}\right) \\
\leq & \mathcal{M}\left(\frac{u}{4}, \mathcal{F}^{\mathrm{DNN}}(L_n, \boldsymbol{s}_n, F), ||\cdot||_{n,1}\right) \\
\leq & 3\left(\frac{16eF}{u} \log \frac{24eF}{u}\right)^{V^+_{\mathcal{F}^{\mathrm{DNN}}(L_n, \boldsymbol{s}_n, F)}} \\
\lesssim & n^{c_4 S_n L_n \log S_n}
\end{aligned}$$

for $u \geq n^{-1}$ by Theorem 9.4 in [21] and Theorem 6 of [23]. Hence for all $t \geq \frac{\varepsilon_n^2}{8}$,

$$\begin{aligned}
\int_{\frac{\kappa(1-\kappa)t}{16\max\{K_1, 2K_2\}}}^{\sqrt{t}} \sqrt{\log \mathcal{N}(u, \mathcal{G}, ||\cdot||_{n,1})} du \lesssim & \sqrt{t}\left(n^{\frac{d}{2\beta+d}} (\log n)^{4+3\tau}\right)^{\frac{1}{2}} \\
= & o\left(\frac{\sqrt{n}t/4}{96\sqrt{2}\max\{K_1, 2K_2\}}\right)
\end{aligned}$$

holds. To sum up, we conclude that

$$\mathbf{P}\left\{ \sup_{f \in \mathcal{F}^{\mathrm{DNN}}(L_n, \boldsymbol{p}'_n)} \frac{\left| ||f - f_0||^2_{2,\mathrm{P}_X} - ||f - f_0||^2_{2,n} \right|}{\varepsilon_n^2 + ||f - f_0||^2_{2,\mathrm{P}_X}} > \frac{1}{2} \right\} \leq 60 \exp\left(-\frac{n\varepsilon_n/8}{128 \cdot 2304}\right). \quad \text{(A.21)}$$

holds for all but finite many n by Lemma A.4. Hence by (A.20) and (A.21), the proof of (A.15) is done. □

B Experiment

B.1 More detail for mBNN

Masked Bayesian CNN First, we introduce how to apply masking variables for masked Bayesian CNN. Most of CNN architectures consist of a mixture of sequences of convolution layer and RELU activation function. For a given CNN, the corresponding masked CNN is constructed by simply adding masking parameters to the CNN model. For the l-th convolution layer, the masked CNN screens the channels using binary masking vector whose dimension is equal to the number of channels in l-th layer. Figure 2 is a illustration of the masked convolution layer.

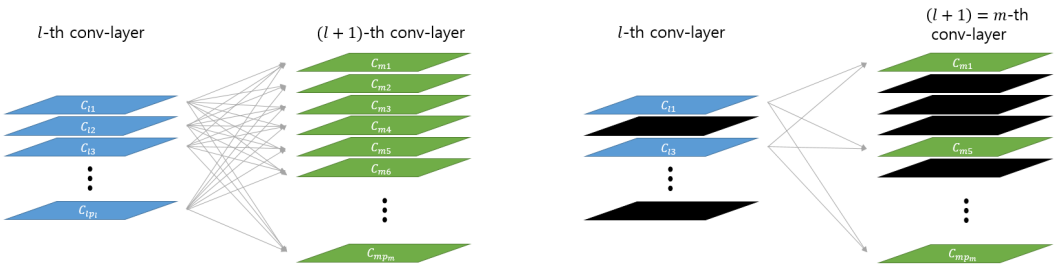

(a) Original convolution layer (b) Convolution layer with masking vector

Figure 2: **Convolution layer and masked convolution layer.** Masking vector $\boldsymbol{m}^{(l)}$ screens channels of the l-th layer. When $\left(\boldsymbol{m}^{(l)}\right)_j = 0$, the j-th channel of the l-th layer becomes inactive. Blue and green rectangles represent input and output channels, and black rectangles represent masked channels.

As we did in Section 3.2, we can simply implement this by converting RELU activation functions into masked-RELU activation functions. For instance, since Resnet18 structure has 17 RELU activation functions, we use 17 masking vectors whose dimensions are equal to the number of channels of the corresponding hidden layer.

Then, the prior Π_n given as (4), (5) and (6) is adopted on parameters $\boldsymbol{M}$ and $\boldsymbol{\theta}$, where $\boldsymbol{M}$ is the concatenate of the all masking vectors and $\boldsymbol{\theta}$ is the concatenation of the filters, weight matrices and bias vectors in the CNN model.

MCMC algorithm First we keep $\boldsymbol{M}^{(t)}$ fixed and update $\boldsymbol{\theta}^{(t)}$ (and ${\sigma^2}^{(t)}$) using existing MCMC algorithm, then we update $\boldsymbol{M}^{(t)}$ using MH algorithm. In practical, we can update $\boldsymbol{M}^{(t)}$ using only the data in single batch for large-scale dataset. Algorithm 2 is a brief summary of our algorithm.

Algorithm 2 Proposed MCMC algorithm

INPUT: $T_{\text{total}}, T_{\text{MH}}, n_{\text{MH}} \in \mathbb{N}$

for $t = 1$ to T_{total} **do**
  Update $\boldsymbol{\theta}^{(t)}$ (and ${\sigma^2}^{(t)}$) using existing MCMC algorithm (e.g. HMC, SGLD).
  if $t \,\%\, T_{\text{MH}} = 0$ **then**
    Update $\boldsymbol{M}^{(t)}$ using Algorithm 1, n_{MH} times.
  end if
end for

B.2 Experimental setup

We report the experimental setup in Table 2. We use Cauchy prior on the weights and biases in both BNN and mBNN. Comparisons for different priors are in Section B.3. Also, we use HMC [43] with 20 leapfrog steps and SGLD [57] with temperature $1/\sqrt{n}$, which is common in SG-MCMC literature [64, 59]. We use the package of [11] for implementing MFVI. For experiments, we used NVIDIA TITAN Xp GPUs.

Table 2: **Hyper-parameters.** Most of the hyper-parameters are set to be the ones given in the corresponding articles.

Method	Hyper-parameters	Experiments	
		UCI regression datasets	Image datasets
mBNN	Prior scale	1.0	0.1
	λ	0.1	0.04
	$N_{\max}$	3	3
	MCMC method	HMC	SGLD
	Step size	10^{-2}	10^{-3}
	Step size schedule	NUTs [28]	Cosine + Constant
	Batch size	Full batch	100
	Thinning interval	200	20
	Burn-in sample	5	5
	T_{MH}	10 (epoch)	10 (batch)
	n_{MH}	10	1
	Number of samples	5	5
	Total epochs	2000	200
BNN	Prior scale	1.0	0.1
	MCMC method	HMC	SGLD
	Step size	10^{-2}	10^{-3}
	Step size schedule	NUTs	Cosine + Constant
	Batch size	Full batch	100
	Thinning interval	200	20
	Burn-in sample	5	5
	Number of samples	5	5
	Total epochs	2000	200
MFVI [3]	Pretrain	$\times$	100 epochs with step size 10^{-3}
	Prior scale	0.1	0.1
	Step size	10^{-2}	10^{-4}
	Optimizer	Adam [30]	Adam
	Batch size	100	100
	Epochs	500	100
	Number of samples	5	5
	Total epochs	500	200
Deep Ensemble [34]	Step size	10^{-2}	10^{-3}
	Optimizer	Adam	Adam
	Batch size	100	100
	Epochs per sample	400	100
	Number of samples	5	5
	Total epochs	2000	500
BCGNJ [39]	Pretrain	$\times$	100 epochs with step size 10^{-3}
	Step size	10^{-2}	10^{-4}
	Optimizer	Adam	Adam
	batch size	100	100
	Epochs	500	100
	Drop criterion	Dropout rate $p > 0.95$	Dropout rate $p > 0.95$
	Number of samples	5	5
	Total epochs	500	200

B.3 Additional results

In Section 6, we reported the results of the experiments to empirically justify the efficiency of our mBNN. In this section, we present results of additional experiments. Unless otherwise noted, all the experimental settings are the same as before.

Extended results for Section 6.1 In Table 1, we compared mBNN with other Bayesian methods in terms of generalization and uncertainty quantification as well as inferential cost and model capacity by analyzing several benchmark datasets. Due to space constraints, we did not report all the results in Table 1. We report the extended results of the experiments in Table 3 and Table 4.

Table 3: **Performance on benchmark datasets.** Performance of mBNN, BNN, MFVI, Deep Ensemble and BC-GNJ. When inferring the posterior of BNN and mBNN, we use HMC for UCI datasets while we use SGLD for image datasets. We report the averages (standard errors) of the performance measures on the repeated experiments.

	Dataset	Measure	BNN	MFVI	DeepEns	mBNN✓	BC-GNJ
Regression	Boston	RMSE	2.942(0.150)	3.295(0.183)	2.835(0.157)	**2.775(0.110)**	2.975(0.149)
		NLL	2.613(0.094)	2.689(0.084)	2.585(0.130)	**2.507(0.067)**	2.758(0.117)
		CRPS	1.485(0.056)	1.713(0.062)	1.444(0.054)	**1.430(0.042)**	1.566(0.056)
		FLOPs ↓ in %	·	·	·	**89.42(0.50)**	81.61(0.59)
		Capacity ↓ in %	·	·	·	**89.11(0.51)**	81.24(0.59)
		# Nodes (average)	[13,100,100,1]	[13,100,100,1]	[13,100,100,1]	[13,38,18,1]	[13,39,40,1]
	Concrete	RMSE	4.855(0.164)	5.416(0.184)	5.009(0.148)	**4.794(0.152)**	5.082(0.253)
		NLL	3.071(0.058)	3.172(0.055)	**3.028(0.040)**	3.036(0.049)	3.367(0.116)
		CRPS	**2.558(0.072)**	2.931(0.090)	2.648(0.064)	2.563(0.061)	2.640(0.106)
		FLOPs ↓ in %	·	·	·	**85.80(0.51)**	64.64(0.69)
		Capacity ↓ in %	·	·	·	**85.41(0.52)**	64.22(0.69)
		# Nodes (average)	[8,100,100,1]	[8,100,100,1]	[8,100,100,1]	[8,46,25,1]	[8,67,49,1]
	Energy	RMSE	0.454(0.013)	2.092(0.110)	0.485(0.013)	**0.451(0.018)**	0.642(0.041)
		NLL	0.638(0.034)	2.139(0.062)	0.744(0.035)	**0.626(0.038)**	0.956(0.052)
		CRPS	0.244(0.005)	1.138(0.057)	0.268(0.006)	**0.241(0.006)**	0.355(0.022)
		FLOPs ↓ in %	·	·	·	77.06(0.87)	**94.90(0.26)**
		Capacity ↓ in %	·	·	·	76.64(0.87)	**94.62(0.27)**
		# Nodes (average)	[8,100,100,1]	[8,100,100,1]	[8,100,100,1]	[8,53,39,1]	[6,21,20,1]
	Yacht	RMSE	0.638(0.066)	1.469(0.118)	**0.564(0.045)**	0.630(0.043)	1.176(0.104)
		NLL	**0.860(0.075)**	1.771(0.076)	0.942(0.234)	0.882(0.089)	1.584(0.070)
		CRPS	0.323(0.036)	0.747(0.054)	**0.243(0.018)**	0.292(0.014)	0.614(0.038)
		FLOPs ↓ in %	·	·	·	**81.26(0.72)**	69.77(0.70)
		Capacity ↓ in %	·	·	·	**80.83(0.72)**	69.33(0.70)
		# Nodes (average)	[6,100,100,1]	[6,100,100,1]	[6,100,100,1]	[6,49,34,1]	[6,60,47,1]
Classification	CIFAR10	ACC	0.944(0.001)	0.905(0.002)	**0.945(0.001)**	0.938(0.002)	0.908(0.002)
		NLL	0.199(0.003)	0.301(0.005)	**0.196(0.002)**	0.205(0.002)	0.305(0.002)
		ECE	0.008(0.000)	0.018(0.003)	0.007(0.001)	**0.006(0.000)**	0.012(0.001)
		FLOPs ↓ in %	·	·	·	**73.28(0.99)**	53.92(2.27)
		Capacity ↓ in %	·	·	·	**92.23(0.09)**	78.88(1.51)
	CIFAR100	ACC	**0.765(0.002)**	0.669(0.004)	0.751(0.002)	0.752(0.005)	0.603(0.003)
		NLL	**0.981(0.012)**	1.322(0.007)	1.200(0.006)	1.014(0.013)	2.000(0.035)
		ECE	**0.002(0.000)**	0.007(0.000)	0.003(0.000)	**0.002(0.000)**	0.006(0.000)
		FLOPs ↓ in %	·	·	·	61.23(0.74)	**63.28(0.54)**
		Capacity ↓ in %	·	·	·	**75.16(0.10)**	59.30(3.85)

For regression problems, we report the root-mean-squared error (RMSE) of the Bayes estimator, negative log-likelihood (NLL) and continuous ranked probability score (CRPS) [18] as well as the reduced inferential costs (FLOPs ↓) and reduced numbers of nonzero parameters (Capacity ↓) of the node-sparse models. Also, we report the average numbers of nodes in each layer of MLP.

For classification problems, we report the accuracy (ACC), NLL, expected calibration error (ECE) as well as the reduced inferential costs (FLOPs ↓) and reduced numbers of nonzero parameters (Capacity ↓) of the node-sparse models.

The results in Table 3 show similar results as before. The mBNN saves the inferential cost and memory usage significantly compared to the non-sparse models while maintaining competitive and sometimes better prediction accuracy and uncertainty quantification.

Table 4: **Selected structures on image datasets** For image dataset experiments with Resnet18 architecture, we report the average numbers of channels of each layer. Compressed models have the same structure with original model in [24] (i.e. basic-blocks with shortcuts), with only the numbers of channels decreasing.

Dataset	Method	# channels (mean)
CIFAR10	Non-sparse	[3,64, 64, 64, 64, 64, 128, 128, 128, 128, 256, 256, 256, 256, 512, 512, 512, 512, 10]
	mBNN✓	[3, 33, 38, 46, 33, 53, 66, 94, 76, 105, 127, 126, 86, 100, 67, 55, 30, 50, 10]
	BCGNJ	[3, 63, 28, 64, 48, 64, 121, 108, 86, 128, 199, 230, 29, 151, 57, 420, 96, 429, 10]
CIFAR100	Non-sparse	[3, 64, 64, 64, 64, 64, 128, 128, 128, 128, 256, 256, 256, 256, 512, 512, 512, 512, 100]
	mBNN✓	[3, 37, 32, 48, 41, 48, 63, 95, 74, 119, 181, 218, 143, 221, 246, 188, 131, 143, 100]
	BCGNJ	[3, 62, 17, 63, 23, 61, 73, 92, 48, 114, 92, 176, 52, 162, 257, 416, 299, 511, 100]

In Table 4, we report the average numbers of channels in each layer of selected models by mBNN and BCGNJ. In mBNN, it is observed that the number of channels increases initially and then decreases as the layer increases. In contrast, relatively large numbers of channels are preserved in odd-numbered hidden layers for BCGNJ. This would be partly due to skip connections in the Resnet architecture. The selected the mBNN architecture would give a useful information for designing DNN architectures efficiently. We will pursue this issue in near future.

Extended results for Section 6.2 In Figure 1, we reported the results of the experiment on *Concrete* dataset to illustrate the efficiency of our proposal distribution in Algorithm 1. Here, we report the full results for the other UCI datasets in Figure 3.

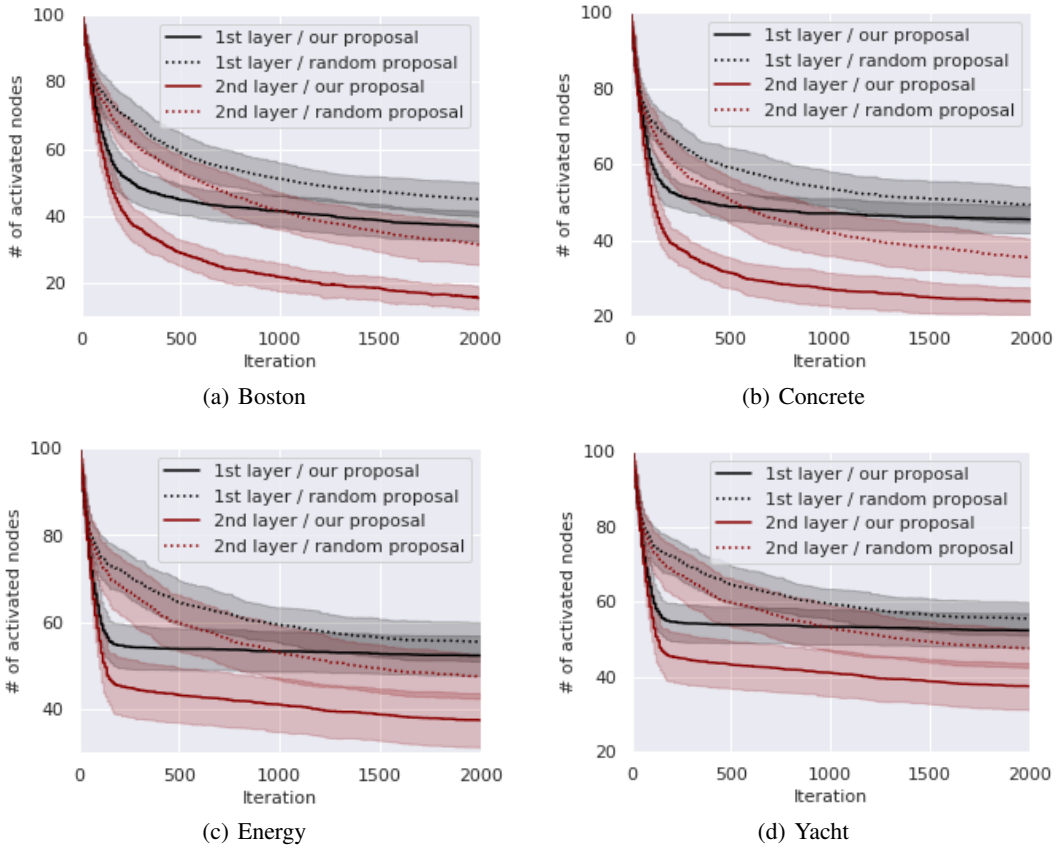

(a) Boston (b) Concrete

(c) Energy (d) Yacht

Figure 3: **Efficiency of the proposal** The numbers of activated nodes in each layer (y-axis) as the MCMC iteration (x-axis) proceeds. We compare our proposal (7) with the random proposal on Section 6.2. It is obvious that the numbers of activated nodes for our proposal distribution decrease much faster than those for the random proposal.

Performance of mBNN on different prior on θ Gaussian prior is usually used for θ in other BNN literature. However, as we report on Section B.2, we used Cauchy prior for θ to ensure theoretical property. To check performances of mBNN on different priors, we consider the Gaussian prior

$$\Pi_n(\theta_i) = \frac{1}{\sqrt{2\pi\varsigma^2}} \exp\left(-\frac{\theta_i^2}{2\varsigma^2}\right)$$

with various prior scale ς and the Cauchy distribution

$$\Pi_n(\theta_i) = \frac{1}{\pi\varsigma\left[1 + \left(\frac{\theta_i}{\varsigma}\right)^2\right]}$$

with various prior scale ς.

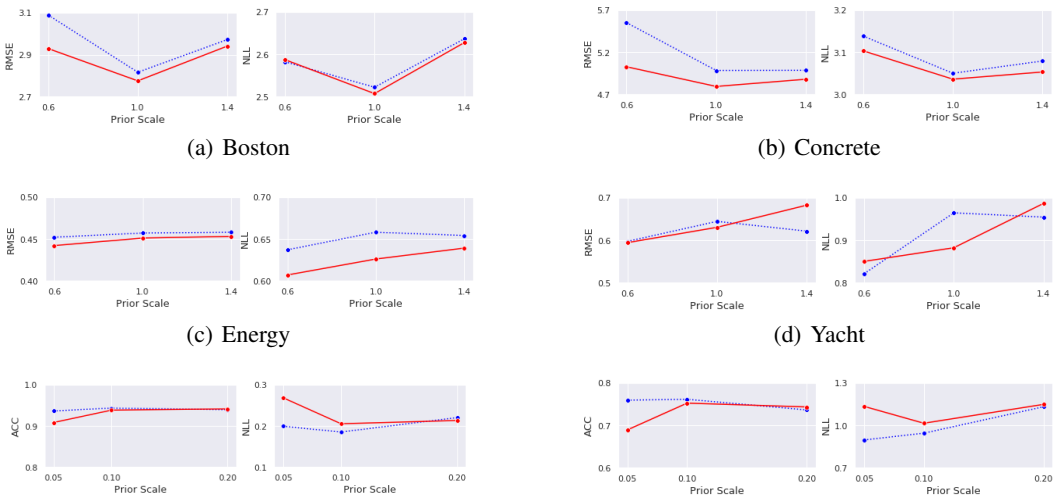

(a) Boston (b) Concrete

(c) Energy (d) Yacht

(e) CIFAR10 (f) CIFAR100

Figure 4: **mBNN with different priors** Solid red lines represent Cauchy prior, and dotted blue lines represent Gaussian prior. We use the prior scales $\{0.6, 1.0, 1.4\}$ and report the averages of RMSE and NLL for UCI regression datasets while we use the prior scales $\{0.05, 0.1, 0.2\}$ and report the averages of ACC and NLL for image datasets.

In Figure 4, we can see that the Cauchy prior yields improved results for UCI datasets compared to the Gaussian prior, which is well matched with the theoretical results. However, for the image datasets, the Gaussian prior seems to work better. Note that our theoretical results do not cover Bayesian CNN. Theoretical investigation of Bayesian CNN would be an interesting research topic.

Performance of mBNN on different λ The λ is the hyper-parameter used in the prior of masking vectors in (4). As mentioned, λ regularizes the widths of the network, and hence larger λ makes the model smaller. In Figures 5 and 6, we report the performances of mBNN for various choices of the regularization parameter λ.

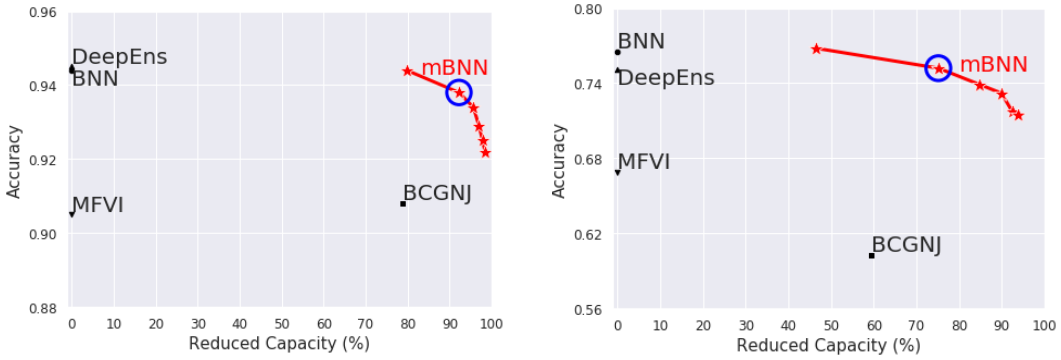

(a) CIFAR10 (b) CIFAR100

Figure 5: **Reduced Capacity vs ACC** Reduced ratio of the number of nonzero parameters (x-axis) and accuracy (y-axis) according to $\lambda \in \{0.03, 0.04, 0.05, 0.06, 0.07, 0.08\}$. We reported the results with $\lambda = 0.04$ (blue circle) in Tables 1 and 3.

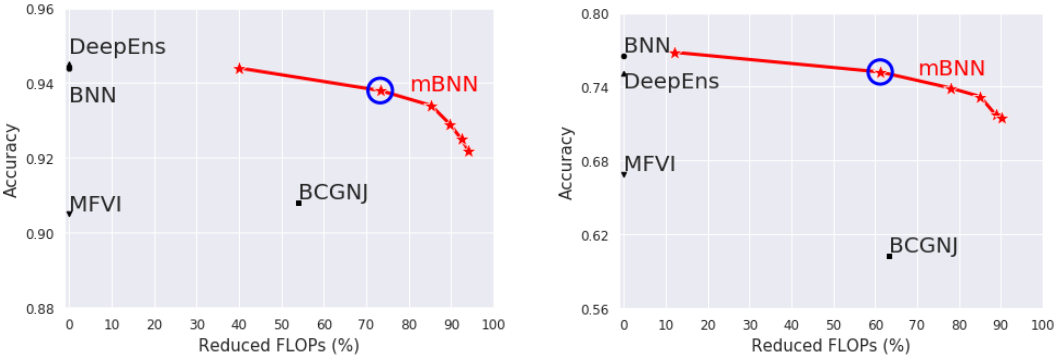

(a) CIFAR10 (b) CIFAR100

Figure 6: **Reduced FLOPs vs ACC** Reduced inferential costs (x-axis) and accuracy (y-axis) according to $\lambda \in \{0.03, 0.04, 0.05, 0.06, 0.07, 0.08\}$. We reported the results with $\lambda = 0.04$ (blue circle) in Tables 1 and 3.

There is a clear trade-off between the model size and accuracy, which appears in common in other sparse deep learning literature. By choosing λ appropriately, we can reduce the model capacity and inferential cost much while not sacrificing the accuracy much (sometimes better) compared to non-sparse methods. Choosing too large λ result in poor accuracy compared to non-sparse methods, but mBNN is more accurate than BCGNJ even though it compresses much more, which supports that mBNN is an efficient tool for compression of DNNs.