

Density Regression and Uncertainty Quantification with Bayesian Deep Noise Neural Networks

Daiwei Zhang¹, Tianci Liu², and Jian Kang^{3,*}

¹Department of Biostatistics, Epidemiology and Informatics, University of Pennsylvania

²School of Electrical and Computer Engineering, Purdue University

³Department of Biostatistics, University of Michigan

*Correspondence: jiankang@umich.edu, 1415 Washington Heights, Ann Arbor, MI 48109

Abstract

Deep neural network (DNN) models have achieved state-of-the-art predictive accuracy in a wide range of applications. However, it remains a challenging task to accurately quantify the uncertainty in DNN predictions, especially those of continuous outcomes. To this end, we propose the Bayesian Deep Noise Neural Network (B-DeepNoise), which generalizes standard Bayesian DNNs by extending the random noise variable from the output layer to all hidden layers. Our model is capable of approximating highly complex predictive density functions and fully learn the possible random variation in the outcome variables. For posterior computation, we provide a closed-form Gibbs sampling algorithm that circumvents tuning-intensive Metropolis-Hastings methods. We establish a recursive representation of the predictive density and perform theoretical analysis on the predictive variance. Through extensive experiments, we demonstrate the superiority of B-DeepNoise over existing methods in terms of density estimation and uncertainty quantification accuracy. A neuroimaging application is included to show our model's usefulness in scientific studies.

1 INTRODUCTION

Deep neural networks (DNNs) have achieved outstanding prediction performance in a wide range of artificial intelligence (AI) applications (Pouyanfar et al., 2018; Berner et al., 2021). Despite overwhelming cases of success, a major drawback of standard DNNs is the lack of reliable uncertainty quantification (UQ) (Begoli et al., 2019). UQ is an essential task in safety-critical AI applications (Amodei et al., 2016). For example, in medical diagnosis, an individualized risk assessment AI model should be able to report its confidence in its predictions. When the AI model is not sufficiently certain in its assessment of a patient, the patient should be referred to human physicians for further evaluation (Jiang et al., 2012; Leibig et al., 2017).

In this work, we seek to solve the problem of UQ in DNN regression tasks, where a representation of the outcome's total random variation cannot be achieved by a finite-length probability vector (unlike in classification tasks) but rather requires an infinite-dimensional predictive probability density function. In a standard DNN regression model, the outcome $y_i \in \mathbb{R}$ and the predictors $\mathbf{x}_i \in \mathbb{R}^P$ are assumed to follow the relation $y_i = f(\mathbf{x}_i) + \epsilon_i$, where the mean function $f(\mathbf{x}_i) = \mathbb{E}(y_i | \mathbf{x}_i)$ is constructed by a DNN, and the random noise ϵ_i follows a zero-mean homoscedastic Gaussian distribution $\epsilon_i \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma^2)$ for some unknown $\sigma^2 > 0$. This formulation implies the conditional variance of the outcome variable $\text{Var}(y_i | \mathbf{x}_i)$, given the predictor, to be constant. However, in real applications, the true conditional distribution could be heteroscedastic (Gaussian or non-Gaussian) or non-Gaussian at all, e.g., the distribution of ϵ_i is asymmetric or multimodal). In these cases, simple UQ statistics (e.g., prediction variance or prediction interval width) may fail to capture important patterns in the outcome variable and result in under- or over-confident UQ. In order to achieve accurate UQ in DNN regression, it is critical to learn the predictive density of the outcome given the predictors. We refer to the problem of estimating the predictive density function as the *density regression* (DR) problem (Dunson et al., 2007). This work focuses on DR tasks with DNNs. As uncertainty in the

outcome variable is quantified by summary statistics of the predictive density, a solution to DR trivially leads to a solution to UQ, regardless of the chosen UQ metric.

1.1 Related Work

To estimate predictive density function, several DNN-based frequentist DR methods have been proposed (Abdar et al., 2021; Ståhl et al., 2020; Caldeira and Nord, 2020; Zhu et al., 2019), such as mixture density networks (Bishop, 1994; Bishop and Nasrabadi, 2006), deep ensembles (Lakshminarayanan et al., 2016), distribution-free methods (Lei and Wasserman, 2014; Pearce et al., 2018), quantile-based models (Romano et al., 2019; Tagasovska and Lopez-Paz, 2019), categorization (Li et al., 2021), and estimation of the predictive cumulative distribution function (CDF) (Huberman et al., 2021).

Compared to ad hoc frequentist DR methods, Bayesian frameworks for DNNs, also known as Bayesian neural networks (BNNs) (MacKay, 1995; Neal, 2012; Xue et al., 2019), provide a more natural and systematic solution to the task of DR. In addition to the asymptotic well-calibratedness of the posterior prediction intervals (Hwang and Ding, 1997; Wang and Rocková, 2020; Sun et al., 2021), the Bayesian framework also improves the prediction accuracy of deterministic DNNs (Kendall and Gal, 2017; Izmailov et al., 2018). Due to the intractability of BNNs’ posterior distributions, variational inference (VI) or Markov Chain Monte Carlo (MCMC) simulation are required for posterior computation. VI methods (Blei et al., 2017; Kingma and Welling, 2013; Mandt et al., 2017) approximate the posterior distribution with simpler distributions (Graves, 2011; Louizos and Welling, 2016; Lee et al., 2020; Rezende and Mohamed, 2015; Louizos and Welling, 2017). Common randomness-based regularization techniques for DNNs, such as dropout (Srivastava et al., 2014; Gal and Ghahramani, 2016; Molchanov et al., 2017), batch normalization (Ioffe and Szegedy, 2015; Teye et al., 2018), and random weights (Blundell et al., 2015; Hernández-Lobato and Adams, 2015), can be interpreted as special cases of VI. However, although computationally efficient, VI methods induce extra approximation errors in posterior computation, which may lead to underestimated variance or oversimplified covariance structures (Blei et al., 2017).

In contrast to VI, MCMC methods simulate the exact posterior distribution. The most popular MCMC algorithm for modern Bayesian methods is arguably the Metropolis-Hastings (MH) algorithm (Chib and Greenberg, 1995; Hitchcock, 2003; Andrieu and Thoms, 2008). However, even with efficient techniques such as Hamiltonian dynamics (Wenzel et al., 2020; Wilson and Izmailov, 2020), Langevin dynamics (Welling and Teh, 2011), stochastic gradients (Chen et al., 2014, 2016), and mini-batches (Wu et al., 2020) that mitigates the high computation burden (Liang et al., 2016; Jospin et al., 2022), MH-based MCMC methods require intensive hyperparameter tuning in order to compute the posterior distribution efficiently. As an alternative MCMC simulation method, Gibbs sampling algorithms (Geman and Geman, 1984; Gelfand and Smith, 1990; Roberts and Smith, 1994; Gelfand, 2000) draw posterior samples for each model parameter (or a block of model parameters) conditioned on all the other parameters. Although Gibbs sampling methods have been developed for deep generative models such as sigmoid belief networks (Gan et al., 2015), Gibbs sampling cannot be applied to standard predictive BNNs, due to their lack of closed-form posterior full conditional distributions. Finally, most existing MCMC- and VI-based BNN methods focus on UQ exclusively, whereas the problem of DR in Bayesian frameworks has only been studied for linear models (Dunson et al., 2007).

In addition, for DR tasks with DNN, no existing method incorporates latent noise in hidden layers, to our best knowledge. Previous works on latent noise in DNNs primarily use it for regularization (You et al., 2019; Gulcehre et al., 2016). The potential of stochastic activation layers for UQ was briefly discussed in Lee et al. (2019), but the context of this work was classification tasks, where the predictive uncertainty could already be fully characterized by well-calibrated categorical distributions without using any latent noise. More recently, Sun and Liang (2022) formulated DNNs as latent variable models and included kernel maps in the input layer to avoid feature collinearity. Although the proposed model is capable of UQ, the more challenging problem of DR has not been studied.

1.2 Our Contributions

To address the challenges for DR with DNNs, we propose the Bayesian Deep Noise Neural Network (B-DeepNoise). B-DeepNoise generalizes standard BNNs by adding latent random noise both before and after every activation layer. Although the latent random noise variables independently follow Gaussian distribu-

tions, their composition across multiple layers with non-linear activations generates highly complex predictive density functions. Moreover, the unique structure of B-DeepNoise induces closed-form posterior full conditional distributions for the model parameters, which eliminates the primary barrier for Gibbs sampling in DNN-based models and therefore makes it possible to simulate the exact posterior distribution without using tuning-intensive MH algorithms.

To our best knowledge, this is the first work on estimating complex predictive density functions by utilizing DNNs with latent random noise. Furthermore, no previous work has developed Gibbs sampling algorithms for DNN-based Bayesian predictive models. In short, our work contributes to the existing literature on DR and UQ with DNNs in the following ways: 1) We propose a Bayesian DNN model for learning complex, non-Gaussian predictive density functions. 2) We develop a virtually tuning-free Gibbs sampling algorithm for posterior computation that uses common samplers only, without the need of MH steps. 3) We perform theoretical analysis for analytic expressions of the predictive densities and variance propagation. 4) We evaluate our model on multiple benchmark datasets and demonstrate its usefulness in a neuroimaging study. The codes for our method is available at github.com/daviddaiweizhang/B-DeepNoise.

2 MODEL DESCRIPTION




Figure 1: Observed data (blue dots) and estimated predictive density (heatmap) by B-DeepNoise (top row), DE (middle row), and BNN (bottom row) for heteroscedastic (left column), asymmetric (middle column), and multimodal (right column) noise.

2.1 DNNs with Latent Noise Variables

Suppose the data consist of features $\mathbf{x}^{(n)} \in \mathbb{R}^P$ and outcomes $\mathbf{y}^{(n)} \in \mathbb{R}^Q$, with $n \in 1, \dots, N$. Let $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ be a Gaussian distribution with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$. To specify the nonlinear association between $\mathbf{y}^{(n)}$ and $\mathbf{x}^{(n)}$, a standard L -layer feed-forward DNN model with Gaussian noise can be represented as

$$\mathbf{y}^{(n)} \mid \mathbf{u}_L^{(n)} \stackrel{\text{iid}}{\sim} \mathcal{N}(\boldsymbol{\beta}_L \mathbf{u}_L^{(n)} + \boldsymbol{\gamma}_L, \mathbf{T}_L) \quad (1)$$

$$\mathbf{u}_{l+1}^{(n)} = h(\boldsymbol{\beta}_l \mathbf{u}_l^{(n)} + \boldsymbol{\gamma}_l), \quad l \in 0, \dots, L-1 \quad (2)$$

$$\mathbf{u}_0^{(n)} = \mathbf{x}^{(n)} \quad (3)$$

where $\boldsymbol{\beta}_l \in \mathbb{R}^{K_l \times K_{l-1}}$ and $\boldsymbol{\gamma}_l \in \mathbb{R}^{K_l}$ are unknown parameters, K_l is the number of units in the l^{th} layer, and $h(\cdot)$ is an element-wise nonlinear activation function. In this formulation, $\mathbf{u}_L^{(n)}$ is a deterministic function of $\mathbf{x}^{(n)}$, which implies that $\mathbf{y}^{(n)} \mid \mathbf{x}^{(n)}$ and $\mathbf{y}^{(n)} \mid \mathbf{u}_L^{(n)}$ follow the same homoscedastic Gaussian distribution

with constant covariance \mathbf{T}_L . To model more complex conditional distributions, we propose the Deep Noise Neural Network (DeepNoise), which generalizes Equation (2) of the standard DNN model into Equations (4) and (5) by including noise variables before and after every activation layer:

$$\mathbf{u}_{l+1}^{(n)} = h\left(\beta_l \mathbf{u}_l^{(n)} + \gamma_l + \epsilon_l^{(n)}\right) + \delta_l^{(n)} \quad (4)$$

$$\epsilon_l^{(n)} \stackrel{\text{iid}}{\sim} \mathcal{N}(\mathbf{0}, \mathbf{T}_l), \delta_l^{(n)} \stackrel{\text{iid}}{\sim} \mathcal{N}(\mathbf{0}, \mathbf{\Sigma}_l), l \in 0, \dots, L-1 \quad (5)$$

where $\mathbf{T}_l = \text{diag}\{\tau_l^2\}$ and $\mathbf{\Sigma}_l = \text{diag}\{\sigma_l^2\}$ with $\tau^2 = (\tau_{l,1}^2, \dots, \tau_{l,K_l}^2)^\top$ and $\sigma_l^2 = (\sigma_{l,1}^2, \dots, \sigma_{l,K_l}^2)^\top$. By composing the latent Gaussian noise variables with linear maps and non-linear activations, DeepNoise is capable of representing a wide range of heteroscedastic Gaussian and non-Gaussian conditional density functions (e.g. asymmetric, multimodal), as illustrated in Figure 1. Intuitively, as Gaussian mixtures are universal approximators of densities (Plataniotis and Hatzinakos, 2017; Calcaterra and Boldt, 2008; Goodfellow et al., 2016, Sec. 3.9.6) and DNNs are universal approximators of functions (Scarselli and Tsoi, 1998; Yarotsky, 2017; Lu and Lu, 2020), DeepNoise is designed to be a universal approximator of conditional densities. The nonparametric nature of DeepNoise enables it to approximate increasingly complex conditional density functions by increasing the number of hidden layers and the number of nodes per layer. In the special case where the variance of the noise variables is set to zero in all but the output layer, DeepNoise is reduced to a standard DNN.

2.2 Model Representation and Prior Specifications

DeepNoise transforms the predictor vector into the outcome vector by iteratively and stochastically applying the linear-noise-nonlinear-noise maps. Thus, the DeepNoise model defined by combining Equations (1) and (3) to (5) is equivalent to

$$\mathbf{v}_l^{(n)} \mid \mathbf{u}_l^{(n)} \stackrel{\text{iid}}{\sim} \mathcal{N}\left(\beta_l \mathbf{u}_l^{(n)} + \gamma_l, \mathbf{T}_l\right), l \in 0, \dots, L \quad (6)$$

$$\mathbf{u}_{l+1}^{(n)} \mid \mathbf{v}_l^{(n)} \stackrel{\text{iid}}{\sim} \mathcal{N}\left(h\left(\mathbf{v}_l^{(n)}\right), \mathbf{\Sigma}_l\right), l \in 0, \dots, L-1 \quad (7)$$

$$\mathbf{u}_0^{(n)} = \mathbf{x}^{(n)}, \quad \mathbf{v}_L^{(n)} = \mathbf{y}^{(n)} \quad (8)$$

Write $\beta_l = (\beta_{l,k,k'})$ and $\gamma_l = (\gamma_{l,k})$. To make fully Bayesian inference, we impose normal-inverse-gamma prior distributions on the weight-bias parameters:

$$\begin{aligned} \beta_{l,k,k'} &\stackrel{\text{iid}}{\sim} \mathcal{N}(0, \rho_{l,k,k'}^2), & \gamma_{l,k} &\stackrel{\text{iid}}{\sim} \mathcal{N}(0, \xi_{l,k}^2) \\ \rho_{l,k,k'}^2 &\stackrel{\text{iid}}{\sim} \mathcal{IG}(a, b) & \xi_{l,k}^2 &\stackrel{\text{iid}}{\sim} \mathcal{IG}(a, b) \end{aligned}$$

and assign inverse-gamma prior distributions to the pre-activation and post-activation noise variances:

$$\tau_{l,k}^2 \stackrel{\text{iid}}{\sim} \mathcal{IG}(a, b), \quad \sigma_{l,k}^2 \stackrel{\text{iid}}{\sim} \mathcal{IG}(a, b).$$

To use weakly informative priors, we set $a = b = 0.001$. We refer to Equations (6) to (8) along with the prior specifications above as the Bayesian Deep Noise Neural Network (B-DeepNoise) model. Although DeepNoise (i.e. Equations (6) to (8) without the prior distributions) is an effective frequentist model for DR, for the rest of the paper, we focus on B-DeepNoise and conduct theoretical, computational, and empirical analyses in the Bayesian framework.

2.3 Posterior Computation

Compared to standard DNNs and BNNs, the addition of latent random noise not only makes B-DeepNoise more flexible in approximating complex predictive density functions, but it also provides closed-form expressions of the posterior full conditional distributions of the model parameters. The latter advantage makes it possible to derive efficient Gibbs sampling algorithms for B-DeepNoise. In order for a Gibbs sampler to be computationally efficient, a key requirement is that all the posterior full conditional distributions are easy to simulate. To this end, we require the activation function to be piecewise linear, as described in Assumption 1.

Assumption 1. The element-wise activation function h can be expressed as $h(t) = \sum_{j=1}^J (b_j t + b'_j) \cdot \mathbb{I}\{t \in [c_{j-1}, c_j]\}$ for some $J \in \mathbb{N}_+$, $b_1, \dots, b_J, b'_1, \dots, b'_J \in \mathbb{R}$, and $-\infty = c_0 < c_1 < \dots < c_{J-1} < c_J = \infty$.

Remark. The family of functions defined in Assumption 1 includes many common activation functions, such as ReLU and leaky ReLU (Maas et al., 2013). Moreover, smooth activation functions can be approximated by piecewise linear functions. For example, the logistic, tanh, and softplus functions can be approximated by the hard sigmoid, hard tanh, and ReLU functions, respectively.

We now derive the posterior full conditional distributions of all the model parameters in B-DeepNoise. Let $\mathcal{TN}_{a,b}(\lambda, \omega^2)$ be a truncated normal distribution on interval $[a, b]$ with location λ and scale ω .

Theorem 1. Suppose the activation function h satisfies Assumption 1. The model parameters in B-DeepNoise have the following posterior full conditional distributions.

$$v_{l,k}^{(n)} \mid \text{rest} \sim \sum_{j=1}^J \pi_{l,k,j}^{(n)} \cdot \mathcal{TN}_{c_{j-1}, c_j} \left(\lambda_{l,k,j}^{(n)}, \omega_{l,k,j}^2 \right) \quad (9)$$

$$\mathbf{u}_l^{(n)} \mid \text{rest} \sim \mathcal{N} \left(\boldsymbol{\mu}_l^{(n)}, \mathbf{U}_l^{(n)} \right) \quad (10)$$

$$(\boldsymbol{\beta}_{l,k}, \gamma_{l,k}) \mid \text{rest} \sim \mathcal{N}(\boldsymbol{\eta}_{l,k}, \mathbf{B}_{l,k}) \quad (11)$$

$$\tau_{l,k}^2 \mid \text{rest} \sim \mathcal{IG} \left(a + 0.5N, b + 0.5 \|\hat{\boldsymbol{\epsilon}}_{l,k}\|_2^2 \right) \quad (12)$$

$$\sigma_{l,k}^2 \mid \text{rest} \sim \mathcal{IG} \left(a + 0.5N, b + 0.5 \|\hat{\boldsymbol{\delta}}_{l,k}\|_2^2 \right) \quad (13)$$

$$\rho_{l,k,k'}^2 \mid \text{rest} \sim \mathcal{IG} \left(a + 0.5, b + 0.5 \beta_{l,k,k'}^2 \right) \quad (14)$$

$$\xi_{l,k}^2 \mid \text{rest} \sim \mathcal{IG} \left(a + 0.5, b + 0.5 \gamma_{l,k}^2 \right), \quad (15)$$

where $\pi_{l,k,j}^{(n)}$, $\lambda_{l,k,j}^{(n)}$, $\boldsymbol{\mu}_l^{(n)}$, $\mathbf{U}_l^{(n)}$, $\boldsymbol{\eta}_{l,k}$, $\mathbf{B}_{l,k}$, $\hat{\boldsymbol{\epsilon}}_{l,k}$, $\hat{\boldsymbol{\delta}}_{l,k}$ are defined in Section S1.1.1.

Remark. Theorem 1 shows that the posterior full conditional distribution of every model parameter in B-DeepNoise is either inverse gamma, normal, or mixture of truncated normal. Equations (10) to (15) follow properties of Bayesian linear regression with conjugate priors (Bishop and Nasrabadi, 2006, Sec. 2.3.3). Equation (9) describes the most complicated block of parameters, $v_{l,k}^{(n)}$, since the nonlinear activation function is involved. Intuitively, the piecewise linear property of the activation function causes the posterior full conditional distribution of $v_{l,k}^{(n)}$ to be “piecewise normal”, i.e. a mixture of truncated normal distributions with adjacent truncation endpoints. Thus sampling $v_{l,k}^{(n)}$ only requires samplers for categorical distributions and truncated normal distributions, which are widely available in scientific computation libraries. In addition, the number of the mixing components is very small for common activation functions (e.g. 2 for ReLU and 3 for hard tanh).

The Gibbs sampler for computing the posterior distributions of B-DeepNoise parameters consist of applying Equations (9) to (15) iteratively, as described in Algorithm 1, Appendix A. Since the predictive distribution only depends on the weight-bias parameters and the latent noise variance parameters, the posterior samples of the latent variables do not need to be stored, which saves memory in practice. In addition, to reduce burn-in time, model parameters can be initialized using gradient-based optimizers or pre-trained weights.

B-DeepNoise and BNN have the same posterior computation time complexity: $\mathcal{O}(LK^2N)$, where L is the network depth, K is the maximal network width, and N is the training sample size. Although the introduction of latent random variables in B-DeepNoise increases the time complexity by a constant term, the extra computation price is paid for more information on the outcomes. Unlike BNNs, which only estimate the predictive mean and variance of every input, B-DeepNoise estimates the predictive density itself.

For the choice of posterior sampling algorithms, B-DeepNoise has the options of using Gibbs samplers or MH samplers, while BNN can use MH samplers only, due to the lack of closed-form posterior full conditional distributions. For hyperparameter tuning, our Gibbs sampler for B-DeepNoise is tuning-free, while gradient-based MH algorithms for BNN, such as Hamiltonian Monte Carlo (HMC), are sensitive to the choice of the

Table 1: Predictive density estimation error (in unit of 0.001) on synthetic datasets for various noise densities and training sizes.

Method	heteroscedastic noise			skewed noise			multimodal noise		
	N=1000	N=2000	N=4000	N=1000	N=2000	N=4000	N=1000	N=2000	N=4000
BP	52 \pm 6	40 \pm 6	33 \pm 2	100 \pm 1	99 \pm 2	98 \pm 2	82 \pm 8	78 \pm 2	78 \pm 2
VI	103 \pm 1	103 \pm 1	101 \pm 1	103 \pm 2	102 \pm 1	98 \pm 1	173 \pm 2	175 \pm 1	175 \pm 0
BNN	103 \pm 2	102 \pm 2	95 \pm 2	105 \pm 2	104 \pm 2	101 \pm 1	128 \pm 8	128 \pm 6	122 \pm 6
DE	49 \pm 4	39 \pm 4	32 \pm 2	101 \pm 1	100 \pm 3	97 \pm 1	82 \pm 1	78 \pm 1	78 \pm 1
B-DeepNoise	72 \pm 4	38 \pm 3	26 \pm 2	84 \pm 6	53 \pm 7	39 \pm 3	98 \pm 6	60 \pm 4	47 \pm 5

integration step size and the number of steps (Hoffman et al., 2014). For scalability, the Gibbs sampling algorithm for B-DeepNoise naturally allows the usage of mini-batches of the whole training dataset, since Gibbs samplers allow arbitrary partitions of the model parameters. That is, at every sampling step, when sampling the latent noise variables (Equations (9) and (10)), instead of updating $v_{l,k}^{(n)}$ and $u_l^{(n)}$ for the entire training set, we only update them for a mini-batch of the training samples. (See Algorithm 1 for details.) By contrast, although efforts have been made in developing mini-batch MH methods (Wu et al., 2020), computing the log-likelihood function using subsets of the training samples inevitably introduce extra noise and causes the algorithm’s stationary distribution to deviate from the target distribution. In addition, whereas the forward and backward propagation steps in MH algorithms for standard BNNs must be computed sequentially by network layers, our Gibbs sampler for B-DeepNoise allows for sampling model parameters parallelly across the layers, since given the parameters in the odd layers, the parameters in the even layers are conditionally independent (and vice versa), thus allowing for simultaneous, parallel updates. (See Algorithm 1 for details.)

Table 2: Experiment results for nine real datasets from the UCI Machine Learning Repository.

Dataset	BP	VI	BNN	DMC	DE	B-DeepNoise
Root Mean Squared Error (RMSE)						
Yacht Hydrodynamics	3.09 \pm 1.28	2.10 \pm 0.43	2.95 \pm 0.01	3.30 \pm 1.14	2.02 \pm 0.76	0.64 \pm 0.32
Boston Housing	3.36 \pm 1.05	2.75 \pm 0.67	2.32 \pm 0.00	3.10 \pm 0.88	3.16 \pm 1.11	2.84 \pm 0.69
Energy Efficiency	2.45 \pm 0.32	0.68 \pm 0.09	1.00 \pm 0.00	1.46 \pm 0.18	2.56 \pm 0.32	0.45 \pm 0.07
Concrete Strength	5.98 \pm 0.62	4.68 \pm 0.52	7.09 \pm 1.95	6.11 \pm 0.47	5.45 \pm 0.55	4.54 \pm 0.46
Wine Quality	0.64 \pm 0.05	0.66 \pm 0.07	0.62 \pm 0.03	0.62 \pm 0.04	0.62 \pm 0.04	0.63 \pm 0.04
Kin8nm	0.08 \pm 0.00	0.09 \pm 0.00	0.07 \pm 0.00	2.27 \pm 0.23	0.07 \pm 0.00	0.07 \pm 0.00
Power Plant	4.02 \pm 0.15	3.89 \pm 0.20	4.27 \pm 0.12	4.12 \pm 0.15	3.98 \pm 0.15	3.62 \pm 0.18
Naval Propulsion	0.00 \pm 0.00	0.00 \pm 0.00	0.02 \pm 0.00	509.93 \pm 0.00	0.00 \pm 0.00	0.00 \pm 0.00
Protein Structure	4.08 \pm 0.06	4.35 \pm 0.09	4.93 \pm 0.27	4.02 \pm 0.04	3.91 \pm 0.03	3.64 \pm 0.03
Negative Log Likelihood (NLL)						
Yacht Hydrodynamics	1.39 \pm 0.33	2.64 \pm 0.04	1.85 \pm 0.00	2.29 \pm 1.14	1.09 \pm 0.19	0.45 \pm 0.22
Boston Housing	3.01 \pm 0.90	2.40 \pm 0.12	2.28 \pm 0.00	2.42 \pm 0.88	2.37 \pm 0.27	2.28 \pm 0.17
Energy Efficiency	1.77 \pm 0.59	1.36 \pm 0.06	1.46 \pm 0.00	1.79 \pm 0.18	1.54 \pm 0.25	0.55 \pm 0.19
Concrete Strength	3.41 \pm 0.47	3.03 \pm 0.21	3.18 \pm 0.04	3.19 \pm 0.47	3.04 \pm 0.22	2.84 \pm 0.14
Wine Quality	2.07 \pm 0.86	10.05 \pm 4.03	1.03 \pm 0.43	0.92 \pm 0.04	1.03 \pm 0.25	0.95 \pm 0.09
Kin8nm	-1.01 \pm 0.21	1.67 \pm 0.44	-1.17 \pm 0.06	-0.92 \pm 0.02	-1.30 \pm 0.04	-1.31 \pm 0.04
Power Plant	2.81 \pm 0.06	2.92 \pm 0.10	2.86 \pm 0.03	2.80 \pm 0.03	2.77 \pm 0.05	2.70 \pm 0.06
Naval Propulsion	-4.67 \pm 1.29	-7.19 \pm 0.48	-5.71 \pm 0.37	-4.10 \pm 0.03	-5.15 \pm 0.21	-7.25 \pm 0.07
Protein Structure	2.80 \pm 0.19	3.25 \pm 0.03	2.95 \pm 0.09	2.77 \pm 0.01	2.54 \pm 0.05	2.65 \pm 0.09
Width of 95% Empirical Prediction Intervals (WCPI-95)						
Yacht Hydrodynamics	2.63 \pm 1.17	6.01 \pm 1.72	6.41 \pm 0.08	∞	2.34 \pm 0.91	1.26 \pm 0.90
Boston Housing	∞	9.81 \pm 1.60	11.20 \pm 0.11	10.51 \pm 1.90	∞	9.06 \pm 1.91
Energy Efficiency	3.70 \pm 1.05	2.60 \pm 0.46	3.55 \pm 0.03	6.28 \pm 1.14	3.55 \pm 0.79	1.65 \pm 0.32
Concrete Strength	∞	∞	27.96 \pm 8.00	25.45 \pm 2.68	17.40 \pm 3.02	16.00 \pm 2.29
Wine Quality	∞	∞	∞	∞	2.88 \pm 0.34	2.48 \pm 0.25
Kin8nm	0.30 \pm 0.02	∞	0.31 \pm 0.02	11.24 \pm 1.11	0.26 \pm 0.01	0.26 \pm 0.01
Power Plant	14.89 \pm 0.52	15.28 \pm 0.76	16.15 \pm 0.91	∞	14.76 \pm 0.56	13.90 \pm 0.62
Naval Propulsion	∞	0.00 \pm 0.00	0.00 \pm 0.00	2548.88 \pm 0.00	0.00 \pm 0.00	0.00 \pm 0.00
Protein Structure	16.32 \pm 0.63	∞	17.65 \pm 0.81	∞	14.06 \pm 0.37	11.90 \pm 0.19

2.4 Predictive Density

We further evaluate the properties of the predictive density function. Theorem 2 expresses the predictive density in a recursive formulation. $\phi_{\mathcal{MVN}}(\cdot | \boldsymbol{\mu}, \boldsymbol{\Sigma})$ be a multivariate normal density function with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$.

Theorem 2. For $l \in \{0, \dots, L\}$, let $\boldsymbol{\theta}_l = \{\boldsymbol{\beta}_l, \boldsymbol{\gamma}_l, \mathbf{T}_l, \boldsymbol{\Sigma}_l\}$ and $\boldsymbol{\Theta}_l = \{\boldsymbol{\theta}_{l'}\}_{l'=0}^l$. In a B-DeepNoise model, the conditional density of the output $\mathbf{y} = \mathbf{v}_L$ given input value \mathbf{x} and model parameters $\boldsymbol{\Theta}_l$ can be iteratively constructed by

$$f(\mathbf{v}_l | \mathbf{x}, \boldsymbol{\Theta}_l) = \int_{\mathbb{R}^{K_{l-1}}} \phi_{\mathcal{MVN}}(\mathbf{v}_l | \boldsymbol{\mu}_l, \mathbf{S}_l) f(\mathbf{v}_{l-1} | \mathbf{x}, \boldsymbol{\Theta}_{l-1}) d\mathbf{v}_{l-1},$$

where $\boldsymbol{\mu}_l = \boldsymbol{\gamma}_l + \boldsymbol{\beta}_l h(\mathbf{v}_{l-1})$ and $\mathbf{S}_l = \mathbf{T}_l + \boldsymbol{\beta}_l \boldsymbol{\Sigma}_l \boldsymbol{\beta}_l^\top$ for $l \in \{1, \dots, L\}$, and

$$f(\mathbf{v}_0 | \mathbf{x}, \boldsymbol{\theta}_0) = \phi_{\mathcal{MVN}}(\mathbf{v}_0 | \boldsymbol{\gamma}_0 + \boldsymbol{\beta}_0 \mathbf{x}, \mathbf{T}_0).$$

Remark. Theorem 2 shows that the predictive density given an input value can be expressed as a continuous mixture of multivariate normal density (CMMVN), where the mixing density is a CMMVN over the output values of the previous layer. Although this highly flexible predictive density does not have a closed-form expression, it can be simulated easily by adding normal noise to the intermediate values of the hidden layers, as stated in Equations (6) and (7).

In standard DNNs and BNNs, the outcome variable follow a normal distribution with the variance equal to by the variance of the outcome noise variable. As a more general model, B-DeepNoise propagates variations in the latent noise variables to produce a complex distribution in the outcome. When the variances of the latent random noise variables are all zero, B-DeepNoise is reduced to a standard BNN. A natural question is how the variance in the output variable can be decomposed by variances of the latent noises variables. To this end, Theorem 3 bounds the outcome variance by the other model parameters.

Theorem 3. Let $\mathbf{y} | \mathbf{x}, \boldsymbol{\Theta}$ be the output value of the B-DeepNoise model given input value \mathbf{x} and model parameters $\boldsymbol{\Theta} = \{\boldsymbol{\beta}_l, \boldsymbol{\gamma}_l, \mathbf{T}_l, \boldsymbol{\Sigma}_l\}_{l=0}^L$. Let $g(\mathbf{x}, \boldsymbol{\Gamma})$ be the output value of the standard DNN model with the same activation function and weight-bias parameters $\boldsymbol{\Gamma} = \{\boldsymbol{\beta}_l, \boldsymbol{\gamma}_l\}_{l=0}^L$. (Note that \mathbf{y} equals to $g(\mathbf{x}, \boldsymbol{\Gamma})$ with probability one when all the latent noise variances are zero.) Assume the activation function h is Lipschitz continuous with Lipschitz constant C_h , and define $d_l^2 = \|\boldsymbol{\beta}_l\|_2^2$. Then

$$\text{Var}(\mathbf{y} | \mathbf{x}, \boldsymbol{\Theta}) + [\mathbb{E}(\mathbf{y} | \mathbf{x}, \boldsymbol{\Theta}) - g(\mathbf{x}, \boldsymbol{\Gamma})]^2 \leq \sum_{l=0}^L \left[d_l^2 \sum_{k=1}^{K_{l-1}} \sigma_{l-1,k}^2 + \sum_{k=1}^{K_l} \tau_{l,k}^2 \right] \left[\prod_{l'=l+1}^L d_{l'}^2 \right] C_h^{2(L-l)}.$$

Remark. According to Theorem 3, given the model parameters, the predictive variance of B-DeepNoise is bounded by the latent noise variances, the spectrum norm of the weight matrices, and the Lipschitz constant of the activation function. In addition, the same upper bound holds for the squared distance of B-DeepNoise's predictive mean from the corresponding deterministic DNN's output value. The expression of this bound can be simplified for common activation functions using global bounds of model parameters, as shown in Corollary 4.

Corollary 4. Let $K = \max_{-1 \leq l \leq L} K_l$, $d^2 = \max_{0 \leq l \leq L} \|\boldsymbol{\beta}_l\|_2^2$, $\sigma^2 = \max_{0 \leq l \leq L-1} \|\boldsymbol{\sigma}_l^2\|_\infty$, $\tau^2 = \max_{0 \leq l \leq L} \|\boldsymbol{\tau}_l^2\|_\infty$. Suppose activation function h is ReLU, leaky ReLU, hard sigmoid, or hard tanh. Then

$$\text{Var}(\mathbf{y} | \mathbf{x}, \boldsymbol{\Theta}) + [\mathbb{E}(\mathbf{y} | \mathbf{x}, \boldsymbol{\Theta}) - g(\mathbf{x}, \boldsymbol{\Gamma})]^2 \leq 3KL (d^{2L+2} + 1) (\sigma^2 + \tau^2)$$

Proofs of all theoretical results are provided in Section S1.2.

3 METHOD COMPARISON

We applied B-DeepNoise and existing methods to synthetic and real data to evaluate their accuracy for predictive density estimation and uncertainty quantification.

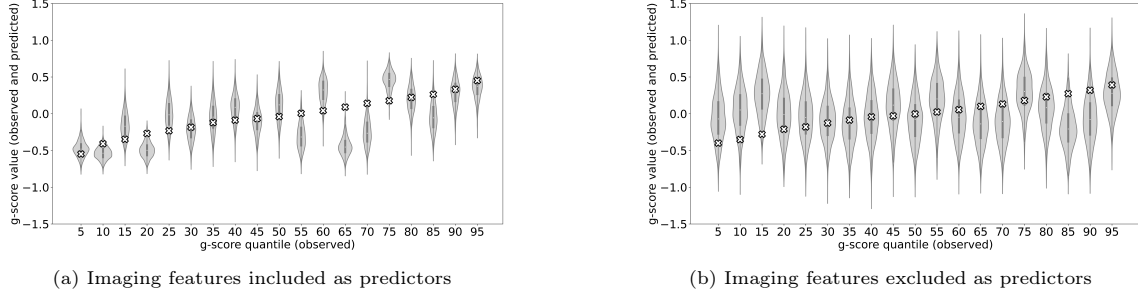


Figure 2: Observed g-scores (white crosses) and predictive densities (violin plots) estimated by B-DeepNoise for 19 testing subjects.

3.1 Experiments on Synthetic Continuous Data

We used synthetic datasets to compare the predictive density estimated by each method with the ground truth. The input variable x was one-dimensional and uniformly distributed on $[-1, 1]$. The output variable y was also one-dimensional, and its conditional median was a linear spline with respect to x . We designed the noise distribution to be heteroscedastic, asymmetric, or multimodal, as shown by the training data points (blue dots) in Figure 1. The training sample size varied among 1000, 2000, and 4000. Every experimental setting was repeated for 20 times.

We used a B-DeepNoise model with 4 hidden layers and 50 nodes in each layer, with the hard tanh function as the activation function. The prior distributions of the variance parameters were set to $\mathcal{IG}(0.001, 0.001)$. We used gradient descent to initialize the model parameters and drew 500 posterior samples. B-DeepNoise was compared against backpropagation (BP), variational inference (VI) (Ritter and Karaletsos, 2022), Bayesian neural networks (BNN) with HMC (Neal et al., 2011), and deep ensemble (DE) (Lakshminarayanan et al., 2016). The baseline methods used identical architecture as B-DeepNoise, and the hyperparameters were selected according to the original authors’ recommendations. Details of the experiments are described in Section S2.

As visualized in Figure 1, B-DeepNoise successfully captured key characteristics of the noise densities. The estimated predictive distributions identified variation in the output variance, opposite directions of skewedness, and abrupt changes between unimodal to bimodal distributions. In contrast, predictive densities estimated by the baseline methods were all unimodal and symmetric.

To quantitatively evaluate the accuracy of the estimated predictive density functions, we numerically computed the L_1 distance between the inverse CDFs of the true and estimated predictive distributions. This metric is equivalent to the simultaneous quantile loss (Tagasovska and Lopez-Paz, 2019), which measures the distance between two distributions by considering all quantiles jointly and is applicable to arbitrary distributions, including heteroscedastic, skewed, and multimodal ones. It is also generalizable to multivariate density functions by using multivariate CDFs. The estimation errors are reported in Table 1. Among all the methods, B-DeepNoise had the smallest error in all but two settings. Moreover, for all the three types of noise, B-DeepNoise’s accuracy improved much faster than the baseline methods as the training sample size increased, especially on the skewed and multimodal data. These simulations illustrate the accuracy of B-DeepNoise in learning complex predictive density functions. For uncertainty quantification, since uncertainty quantification is measured by summary statistics of the predictive density (e.g. variance, 95% prediction intervals), the superior performance of B-DeepNoise in predictive density estimation implies its superior performance in uncertainty quantification accuracy.

The computation efficiency of posterior sampling methods hinges crucially on fast convergence rates and well-mixed behaviors. It is worth noting that the optimal choice of such a method can be both model-specific and data-dependent (Papaspiliopoulos and Roberts, 2008). In this regard, B-DeepNoise has an edge over conventional BNNs thanks to its compatibility with both Metropolis Hastings (MH) algorithms and Gibbs samplers, thereby offering a broader range of options for choosing the most suitable sampling method. In our simulations, we compared the convergence of the proposed Gibbs sampler on the B-DeepNoise model with HMC on the B-DeepNoise model. The trace plots of the negative log-likelihood (Figure S2) highlight that, compared to B-DeepNoise+HMC, B-DeepNoise+Gibbs converged within fewer steps and

reached lower negative log-likelihood values, which indicates a higher per-step efficiency and more accurate posterior predictive densities. Despite the superior empirical convergence performance of the Gibbs sampler, it remains essential to be mindful of the potential theoretical challenges with Gibbs samplers that can lead to inadequate mixing behaviors (Papaspiliopoulos and Roberts, 2008; Román and Hobert, 2012; Johnson and Jones, 2015). For example, when the model parameters are highly correlated or when the noise variance is small, Gibbs samplers may be less efficient than MH algorithms. While a comprehensive discussion on the theoretical issues surrounding Gibbs samplers are beyond the scope of this paper, we recommend in practice the identification of the most effective sampling methods by examining the effective sample sizes and inspecting the trace plots of the negative log-likelihood.

3.2 Experiments on Real Continuous Data




Figure 3: Influence of brain regions on the predictive mean of the g-score.

We applied B-DeepNoise and baseline methods to nine regression datasets on the UCI Machine Learning Repository (Dua and Graff, 2017). Experiment setup was similar to Gal and Ghahramani (2016); Lakshminarayanan et al. (2016), where each dataset was randomly split into training and testing sets for 5 times on Protein Structure and 20 times in the other datasets. Prediction accuracy was measured by the testing root mean squared error (RMSE) of the predictive mean, and uncertainty quantification accuracy was measured by the average negative log likelihood (NLL) of the predictive density on the testing data, as in Hernández-Lobato and Adams (2015); Gal and Ghahramani (2016); Lakshminarayanan et al. (2016). The NLL is equivalent to the empirical cross-entropy between the estimated and the true predictive density functions, which penalizes both over- and under-confident predictions. See Quinero-Candela et al. (2005) for NLL as a valid metric and a proper scoring rule for evaluating predictive uncertainty. Another common UQ accuracy metric is the width of the 95% uncalibrated prediction interval (UPI) (Pearce et al., 2018). However, due to model miscalibration, the 95% UPI might over- or under-cover 95% of the observed outcomes. To better measure UQ accuracy, we computed the 95% calibrated prediction intervals (CPI), defined as the minimal $x\%$ UPI that covers at least 95% of the observed outcomes. In other words, the 95% CPI is the miscalibration-adjusted version of the 95% UPI. (See Section S3.1.3 for details.) Using the three aforementioned metrics, we compared B-DeepNoise against BP, VI, BNN, DE, and Dropout Monte Carlo (DMC) (Gal and Ghahramani, 2016). Network architectures and experiment setups are similar to those in Section 3.1. (See Section S3 for details.)

Table 2 shows the experiment results. Compared to the baseline methods, B-DeepNoise had the least RMSE on all except two datasets, which indicates a superior prediction accuracy. Moreover, B-DeepNoise’s predictive distribution was also overall more accurate than the other methods, since the NLL of B-DeepNoise was the smallest on all except two datasets. The high UQ accuracy of B-DeepNoise was further demonstrated by its uniformly narrowest WCPI-95. In comparison, the baseline methods not only had significantly larger WCPI-95s, but some methods could not produce valid 95% CPIs at all on some of the datasets (as indicated by ∞ in Table 2), since even their 100% prediction intervals could not cover at least 95% of the observed testing outcomes. In addition, B-DeepNoise’s performance was overall more stable than the baseline methods, as reflected by its generally smaller standard errors with respect to all the metrics. In short, the experiments on the UCI datasets demonstrated the superior DR and UQ accuracy of B-DeepNoise compared to existing methods.

For computation efficiency, the total runtimes of all methods are reported in Table S2. Overall, the computation cost of B-DeepNoise is comparable with BNN.

4 NEUROIMAGING-BASED PREDICTION OF GENERAL INTELLIGENCE FOR ADOLESCENTS

We demonstrate the usefulness of B-DeepNoise in scientific studies by applying it to the neuroimaging data in the Adolescent Brain Cognitive Development (ABCD) Study (Casey et al., 2018). The dataset contains 1191 subjects recruited from multiple study sites in the United States. For every subject, we use B-DeepNoise to predict the density of the general intelligence score (g-score) (Dubois et al., 2018; O’Shea et al., 2016) using the 2-back task score (Cohen et al., 2016), general psychopathology factor (Caspi et al., 2014; Carver et al., 2017; Murray et al., 2016), demographic information (age, sex, parental education level, household marital status, household income, and ethnic backgrounds), and brain functional magnetic resonance imaging (fMRIs). (See Casey et al. (2018); Zhang et al. (2020) for details of the ABCD data.) We randomly selected 90% of the samples for training and the rest for testing. The model hyperparameters are similar to those in Sections 3.1 and 3.2. See Section S4 for details of the experiment setup.

To illustrate the predictive density functions learned by B-DeepNoise, we selected 19 testing subjects that correspond to the 5%, 10%, ..., 90%, 95% quantiles of the observed g-score. The results are shown in Figure 2a. The prediction distributions estimated by B-DeepNoise have successfully covered the observed outcomes, with only a couple of samples located near the tails of the predictive distributions. To further assess B-DeepNoise’s UQ accuracy, we removed the imaging predictors and refit the model with the non-imaging predictors only. As shown in Figure 2b, when the imaging information was not available, B-DeepNoise widened the prediction intervals to account for the higher degree of uncertainty. In contrast, the predictive densities in the imaging-included model are not only more concentrated but also exhibited greater magnitude of heteroscedasticity and skewedness. These results indicate that B-DeepNoise is able to appropriately adjust the predictive density to reflect its subject-level prediction confidence.

Furthermore, we investigated the most influential neuroimaging features on the predictive mean of the g-score, where influence is measured by the average absolute value of the gradient of the predictive mean with respect to the feature in the B-DeepNoise model. As shown in Table S3, the 2-back task score had the highest influence (1.32) on the predicted mean g-score, and the magnitude of influence was much higher than the other features (0.63 or less). This result is consistent with the current understanding that memory is a one of the major components that encompass cognitive abilities (Thompson et al., 2019). The rest of the most influential features were primarily imaging features, and Figure 3 visualizes the influence of brain regions (Tzourio-Mazoyer et al., 2002) on predicted g-scores. Many of the corresponding brain regions have been well studied for their associations with general intelligence in existing studies. To name a few, bilateral calcarine is associated with the intelligence quotient (IQ) of children and adolescents (Kilroy et al., 2011); putamen has been identified with verbal IQ in healthy adults (Grazioplene et al., 2015); the right paracentral lobule has been found to be associated with functioning decline in at-risk mental state patients (Sasabayashi et al., 2021). Overall, neuroimaging regions that are the most influential for B-DeepNoise’s predicted mean g-score are supported by existing findings in the literature. Although causes for individual differences in cognitive abilities are multifaceted (Dubois et al., 2018), and our analysis does not make causal claims on the relationships between the features and the outcomes, B-DeepNoise’s ability of providing predictive densities of the g-score makes the model useful for studies that involve quantification of uncertainty in general intellectual capacity.

5 CONCLUSION

In this work, we have presented B-DeepNoise, a novel Bayesian model based on deep neural networks for density regression, which generalizes the task of uncertainty quantification. We have demonstrated our model’s theoretical and computational properties and evaluated its performance on synthetic and real regression data.

A limitation of B-DeepNoise is the lack of advantage on classification tasks. Although B-DeepNoise combined with a softmax activation in the output layer is capable of handling categorical outcomes (Section S5.2), it (or any other stochastic models) is not expected theoretically to have superior uncertainty quantification accuracy than a standard DNN classifier. See Section S5.1 for a demonstration.

For future works, we are interested in developing frequentist versions of our method and exploring the

possibility of using them for outcome selection.

Acknowledgements

This research was partially supported by NIH grants NIH R01MH105561 (D.Z and J.K.), NIH R01DA048993 (J.K.), NIH R01GM124061 (J.K.), and NSF grant IIS2123777 (J.K.).

A Gibbs Sampling Algorithm for B-DeepNoise

Based on the posterior full conditional distributions in Theorem 1, we describe the Gibbs sampling algorithm for B-DeepNoise in Algorithm 1. Partition the training sample indices $\{1, \dots, N_{\text{train}}\}$ into T mini-batches $\{\mathcal{T}_t : t = 1, \dots, T\}$ (i.e. $\bigcup_{t=1}^T \mathcal{T}_t = \{1, \dots, N_{\text{train}}\}$ and $\mathcal{T}_{t_1} \cap \mathcal{T}_{t_2} = \emptyset \forall t_1 \neq t_2$). Without loss of generality, assume L is an even number.

Algorithm 1: Posterior Sampler for Bayesian Deep Noise Neural Network (B-DeepNoise)

Input: training features and targets $[\mathbf{x}^{(n)}, \mathbf{y}^{(n)}]_{n=1}^{N_{\text{train}}}$, testing features $[\mathbf{x}^{(n)}]_{n=1}^{N_{\text{test}}}$, number of posterior samples M , number of realizations of the predictive distribution R , training sample mini-batches $[\mathcal{T}_t]_{t=1}^T$.

Output: posterior predictive samples for testing targets $[[[\hat{\mathbf{y}}_{m,r}^{(n)}]_{r=1}^R]_{m=1}^M]_{n=1}^{N_{\text{test}}}$
 Randomly initialize $[\beta_l]_{l=0}^L, [\gamma_l]_{l=0}^L, [\tau_l^2]_{l=0}^L, [\sigma_l^2]_{l=0}^L, [[\mathbf{u}_l^{(n)}]_{l=1}^L]_{n=1}^{N_{\text{train}}}, [[\mathbf{v}_l^{(n)}]_{l=0}^{L-1}]_{n=1}^{N_{\text{train}}}$;
 $[\mathbf{u}_0^{(n)}]_{n=1}^{N_{\text{train}}} \leftarrow [\mathbf{x}^{(n)}]_{n=1}^{N_{\text{train}}}, [\mathbf{v}_L^{(n)}]_{n=1}^{N_{\text{train}}} \leftarrow [\mathbf{y}^{(n)}]_{n=1}^{N_{\text{train}}}$;

```

for  $m \leftarrow 1, \dots, M$  do
  for  $t \leftarrow 1, \dots, T$  do
    for  $l' \leftarrow \{0, 1\}$  do
      for  $l'' \leftarrow 0, \dots, L/2 - 1$  do
         $l \leftarrow l' + 2l''$ ;
        for  $n \in \mathcal{T}_t$  do
          Sample  $\mathbf{v}_l^{(n)} \mid \mathbf{u}_l^{(n)}, \mathbf{u}_{l+1}^{(n)}, \beta_l, \gamma_l, \sigma_l^2, \tau_l^2$  by Equation (9);
          Sample  $\mathbf{u}_{l+1}^{(n)} \mid \mathbf{v}_l^{(n)}, \mathbf{v}_{l+1}^{(n)}, \beta_{l+1}, \gamma_{l+1}, \sigma_{l+1}^2, \tau_{l+1}^2$  by Equation (10);
        end
        Sample  $\beta_l, \gamma_l \mid \mathbf{u}_l, \mathbf{v}_l, \tau_l^2, \rho_l^2, \xi_l^2$  by Equation (11);
        Sample  $\tau_l^2, \sigma_l^2, \rho_l^2, \xi_l^2 \mid \mathbf{v}_l, \mathbf{u}_l, \beta_l, \gamma_l, \mathbf{u}_{l+1}$  by Equations (12) to (15);
      end
    end
     $\hat{\beta}_m \leftarrow [\beta_l]_{l=0}^L, \hat{\gamma}_m \leftarrow [\gamma_l]_{l=0}^L, \hat{\tau}_m^2 \leftarrow [\tau_l^2]_{l=0}^L, \hat{\sigma}_m^2 \leftarrow [\sigma_l^2]_{l=0}^L$ ;
  end
end
for  $n \leftarrow 1, \dots, N_{\text{test}}$  do
  for  $m \leftarrow 1, \dots, M$  do
    for  $r \leftarrow 1, \dots, R$  do
      Sample  $\mathbf{y}_m^{(n)} \mid \mathbf{x}^{(n)}, \hat{\beta}_m, \hat{\gamma}_m, \hat{\tau}_m^2, \hat{\sigma}_m^2$  by Equations (6) and (7);
       $\hat{\mathbf{y}}_{m,r}^{(n)} \leftarrow \mathbf{y}_m^{(n)}$ 
    end
  end
end

```

References


- Abdar, M., Pourpanah, F., Hussain, S., Rezazadegan, D., Liu, L., Ghavamzadeh, M., Fieguth, P., Cao, X., Khosravi, A., Acharya, U. R., et al. (2021). A review of uncertainty quantification in deep learning: Techniques, applications and challenges. *Information Fusion*, 76:243–297.
- Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., and Mané, D. (2016). Concrete problems in ai safety. *arXiv preprint arXiv:1606.06565*.
- Andrieu, C. and Thoms, J. (2008). A tutorial on adaptive MCMC. *Statistics and computing*, 18(4):343–373.
- Begoli, E., Bhattacharya, T., and Kusnezov, D. (2019). The need for uncertainty quantification in machine-assisted medical decision making. *Nature Machine Intelligence*, 1(1):20–23.
- Berner, J., Grohs, P., Kutyniok, G., and Petersen, P. (2021). The modern mathematics of deep learning. *arXiv preprint arXiv:2105.04026*.
- Bishop, C. M. (1994). Mixture density networks.
- Bishop, C. M. and Nasrabadi, N. M. (2006). *Pattern recognition and machine learning*, volume 4. Springer.
- Blei, D. M., Kucukelbir, A., and McAuliffe, J. D. (2017). Variational inference: A review for statisticians. *Journal of the American statistical Association*, 112(518):859–877.
- Blundell, C., Cornebise, J., Kavukcuoglu, K., and Wierstra, D. (2015). Weight uncertainty in neural network. In *International Conference on Machine Learning*, pages 1613–1622. PMLR.
- Calcaterra, C. and Boldt, A. (2008). Approximating with gaussians. *arXiv preprint arXiv:0805.3795*.
- Caldeira, J. and Nord, B. (2020). Deeply uncertain: comparing methods of uncertainty quantification in deep learning algorithms. *Machine Learning: Science and Technology*, 2(1):015002.
- Carver, C. S., Johnson, S. L., and Timpano, K. R. (2017). Toward a functional view of the p factor in psychopathology. *Clinical Psychological Science*, 5(5):880–889.
- Casey, B., Cannonier, T., Conley, M. I., Cohen, A. O., Barch, D. M., Heitzeg, M. M., Soules, M. E., Teslovich, T., Dellarco, D. V., Garavan, H., et al. (2018). The adolescent brain cognitive development (abcd) study: imaging acquisition across 21 sites. *Developmental cognitive neuroscience*, 32:43–54.
- Caspi, A., Houts, R. M., Belsky, D. W., Goldman-Mellor, S. J., Harrington, H., Israel, S., Meier, M. H., Ramrakha, S., Shalev, I., Poulton, R., et al. (2014). The p factor: one general psychopathology factor in the structure of psychiatric disorders? *Clinical psychological science*, 2(2):119–137.
- Chen, C., Carlson, D., Gan, Z., Li, C., and Carin, L. (2016). Bridging the gap between stochastic gradient MCMC and stochastic optimization. In *Artificial Intelligence and Statistics*, pages 1051–1060. PMLR.
- Chen, T., Fox, E., and Guestrin, C. (2014). Stochastic gradient hamiltonian monte carlo. In *International conference on machine learning*, pages 1683–1691. PMLR.
- Chib, S. and Greenberg, E. (1995). Understanding the metropolis-hastings algorithm. *The american statistician*, 49(4):327–335.
- Cohen, A. O., Breiner, K., Steinberg, L., Bonnie, R. J., Scott, E. S., Taylor-Thompson, K., Rudolph, M. D., Chein, J., Richeson, J. A., Heller, A. S., et al. (2016). When is an adolescent an adult? assessing cognitive control in emotional and nonemotional contexts. *Psychological Science*, 27(4):549–562.
- Dua, D. and Graff, C. (2017). UCI machine learning repository.
- Dubois, J., Galdi, P., Paul, L. K., and Adolphs, R. (2018). A distributed brain network predicts general intelligence from resting-state human neuroimaging data. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1756):20170284.

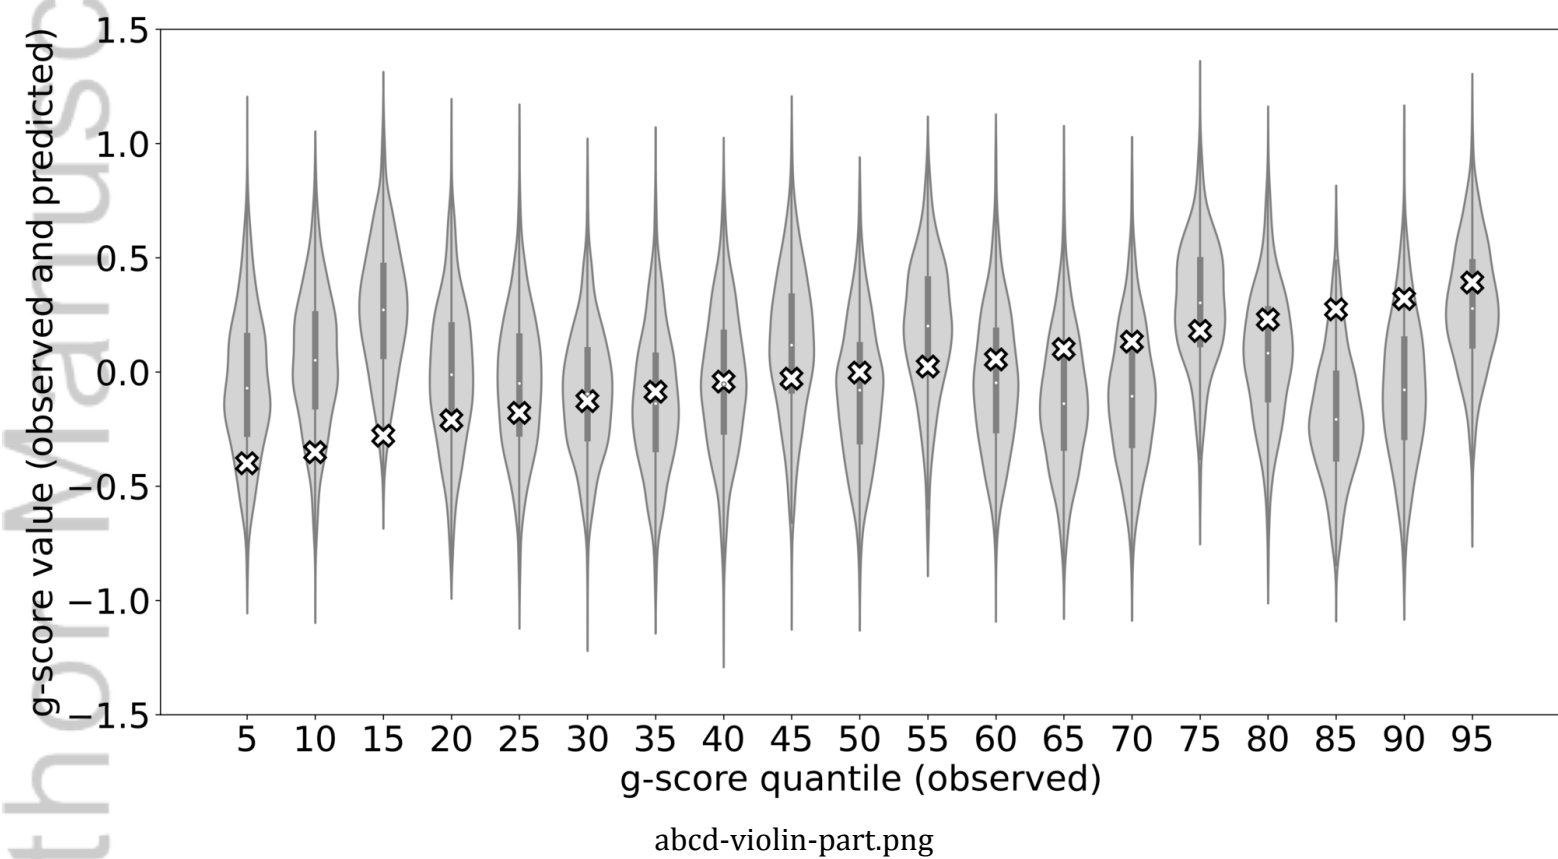
- Dunson, D. B., Pillai, N., and Park, J.-H. (2007). Bayesian density regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 69(2):163–183.
- Gal, Y. and Ghahramani, Z. (2016). Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059. PMLR.
- Gan, Z., Henao, R., Carlson, D., and Carin, L. (2015). Learning deep sigmoid belief networks with data augmentation. In *Artificial Intelligence and Statistics*, pages 268–276. PMLR.
- Gelfand, A. E. (2000). Gibbs sampling. *Journal of the American statistical Association*, 95(452):1300–1304.
- Gelfand, A. E. and Smith, A. F. (1990). Sampling-based approaches to calculating marginal densities. *Journal of the American statistical association*, 85(410):398–409.
- Geman, S. and Geman, D. (1984). Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE Transactions on pattern analysis and machine intelligence*, (6):721–741.
- Goodfellow, I., Bengio, Y., Courville, A., and Bengio, Y. (2016). *Deep learning*, volume 1. MIT press Cambridge.
- Graves, A. (2011). Practical variational inference for neural networks. *Advances in neural information processing systems*, 24.
- Grazioplene, R. G., Ryman, S., Gray, J. R., Rustichini, A., Jung, R. E., and DeYoung, C. G. (2015). Subcortical intelligence: Caudate volume predicts iq in healthy adults. *Human brain mapping*, 36(4):1407–1416.
- Gulcehre, C., Moczulski, M., Denil, M., and Bengio, Y. (2016). Noisy activation functions. In *International conference on machine learning*, pages 3059–3068. PMLR.
- Hernández-Lobato, J. M. and Adams, R. (2015). Probabilistic backpropagation for scalable learning of bayesian neural networks. In *International conference on machine learning*, pages 1861–1869. PMLR.
- Hitchcock, D. B. (2003). A history of the metropolis–hastings algorithm. *The American Statistician*, 57(4):254–257.
- Hoffman, M. D., Gelman, A., et al. (2014). The no-u-turn sampler: adaptively setting path lengths in hamiltonian monte carlo. *J. Mach. Learn. Res.*, 15(1):1593–1623.
- Huberman, D. B., Reich, B. J., and Bondell, H. D. (2021). Nonparametric conditional density estimation in a deep learning framework for short-term forecasting. *Environmental and Ecological Statistics*, pages 1–15.
- Hwang, J. G. and Ding, A. A. (1997). Prediction intervals for artificial neural networks. *Journal of the American Statistical Association*, 92(438):748–757.
- Ioffe, S. and Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. PMLR.
- Izmailov, P., Podoprikin, D., Garipov, T., Vetrov, D., and Wilson, A. G. (2018). Averaging weights leads to wider optima and better generalization. *arXiv preprint arXiv:1803.05407*.
- Jiang, X., Osl, M., Kim, J., and Ohno-Machado, L. (2012). Calibrating predictive model estimates to support personalized medicine. *Journal of the American Medical Informatics Association*, 19(2):263–274.
- Johnson, A. A. and Jones, G. L. (2015). Geometric ergodicity of random scan gibbs samplers for hierarchical one-way random effects models. *Journal of Multivariate Analysis*, 140:325–342.
- Jospin, L. V., Laga, H., Boussaid, F., Buntine, W., and Bennamoun, M. (2022). Hands-on bayesian neural networks—a tutorial for deep learning users. *IEEE Computational Intelligence Magazine*, 17(2):29–48.


- Kendall, A. and Gal, Y. (2017). What uncertainties do we need in bayesian deep learning for computer vision? *arXiv preprint arXiv:1703.04977*.
- Kilroy, E., Liu, C. Y., Yan, L., Kim, Y. C., Dapretto, M., Mendez, M. F., and Wang, D. J. (2011). Relationships between cerebral blood flow and iq in typically developing children and adolescents. *Journal of cognitive science*, 12(2):151.
- Kingma, D. P. and Welling, M. (2013). Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Lakshminarayanan, B., Pritzel, A., and Blundell, C. (2016). Simple and scalable predictive uncertainty estimation using deep ensembles. *arXiv preprint arXiv:1612.01474*.
- Lee, J., Humt, M., Feng, J., and Triebel, R. (2020). Estimating model uncertainty of neural networks in sparse information form. In *International Conference on Machine Learning*, pages 5702–5713. PMLR.
- Lee, J., Shridhar, K., Hayashi, H., Iwana, B. K., Kang, S., and Uchida, S. (2019). Probact: A probabilistic activation function for deep neural networks. *arXiv preprint arXiv:1905.10761*, 5:13.
- Lei, J. and Wasserman, L. (2014). Distribution-free prediction bands for non-parametric regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(1):71–96.
- Leibig, C., Allken, V., Ayhan, M. S., Berens, P., and Wahl, S. (2017). Leveraging uncertainty information from deep neural networks for disease detection. *Scientific reports*, 7(1):1–14.
- Li, R., Reich, B. J., and Bondell, H. D. (2021). Deep distribution regression. *Computational Statistics & Data Analysis*, 159:107203.
- Liang, F., Kim, J., and Song, Q. (2016). A bootstrap metropolis–hastings algorithm for bayesian analysis of big data. *Technometrics*, 58(3):304–318.
- Louizos, C. and Welling, M. (2016). Structured and efficient variational deep learning with matrix gaussian posteriors. In *International conference on machine learning*, pages 1708–1716. PMLR.
- Louizos, C. and Welling, M. (2017). Multiplicative normalizing flows for variational bayesian neural networks. In *International Conference on Machine Learning*, pages 2218–2227. PMLR.
- Lu, Y. and Lu, J. (2020). A universal approximation theorem of deep neural networks for expressing probability distributions. *Advances in Neural Information Processing Systems*, 33.
- Maas, A. L., Hannun, A. Y., Ng, A. Y., et al. (2013). Rectifier nonlinearities improve neural network acoustic models. In *Proc. icml*, volume 30, page 3. Citeseer.
- MacKay, D. J. (1995). Probable networks and plausible predictions-a review of practical bayesian methods for supervised neural networks. *Network: computation in neural systems*, 6(3):469.
- Mandt, S., Hoffman, M. D., and Blei, D. M. (2017). Stochastic gradient descent as approximate bayesian inference. *Journal of Machine Learning Research*, 18:1–35.
- Molchanov, D., Ashukha, A., and Vetrov, D. (2017). Variational dropout sparsifies deep neural networks. In *International Conference on Machine Learning*, pages 2498–2507. PMLR.
- Murray, A. L., Eisner, M., and Ribeaud, D. (2016). The development of the general factor of psychopathology ‘p factor’ through childhood and adolescence. *Journal of abnormal child psychology*, 44(8):1573–1586.
- Neal, R. M. (2012). *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media.
- Neal, R. M. et al. (2011). MCMC using hamiltonian dynamics. *Handbook of markov chain monte carlo*, 2(11):2.
- O’Shea, A., Cohen, R., Porges, E. C., Nissim, N. R., and Woods, A. J. (2016). Cognitive aging and the hippocampus in older adults. *Frontiers in aging neuroscience*, 8:298.


- Papaspiliopoulos, O. and Roberts, G. (2008). Stability of the Gibbs sampler for Bayesian hierarchical models.
- Pearce, T., Brintrup, A., Zaki, M., and Neely, A. (2018). High-quality prediction intervals for deep learning: A distribution-free, ensembled approach. In *International conference on machine learning*, pages 4075–4084. PMLR.
- Plataniotis, K. N. and Hatzinakos, D. (2017). Gaussian mixtures and their applications to signal processing. In *Advanced signal processing handbook*, pages 89–124. CRC Press.
- Pouyanfar, S., Sadiq, S., Yan, Y., Tian, H., Tao, Y., Reyes, M. P., Shyu, M.-L., Chen, S.-C., and Iyengar, S. S. (2018). A survey on deep learning: Algorithms, techniques, and applications. *ACM Computing Surveys (CSUR)*, 51(5):1–36.
- Quinonero-Candela, J., Rasmussen, C. E., Sinz, F., Bousquet, O., and Schölkopf, B. (2005). Evaluating predictive uncertainty challenge. In *Machine Learning Challenges Workshop*, pages 1–27. Springer.
- Rezende, D. and Mohamed, S. (2015). Variational inference with normalizing flows. In *International conference on machine learning*, pages 1530–1538. PMLR.
- Ritter, H. and Karaletsos, T. (2022). Tyxe: Pyro-based bayesian neural nets for pytorch. *Proceedings of Machine Learning and Systems*, 4.
- Roberts, G. O. and Smith, A. F. (1994). Simple conditions for the convergence of the gibbs sampler and metropolis-hastings algorithms. *Stochastic processes and their applications*, 49(2):207–216.
- Román, J. C. and Hobert, J. P. (2012). Convergence analysis of the gibbs sampler for bayesian general linear mixed models with improper priors.
- Romano, Y., Patterson, E., and Candes, E. (2019). Conformalized quantile regression. *Advances in neural information processing systems*, 32.
- Sasabayashi, D., Takayanagi, Y., Takahashi, T., Nishiyama, S., Mizukami, Y., Katagiri, N., Tsujino, N., Nemoto, T., Sakuma, A., Katsura, M., et al. (2021). Reduced cortical thickness of the paracentral lobule in at-risk mental state individuals with poor 1-year functional outcomes. *Translational psychiatry*, 11(1):1–9.
- Scarselli, F. and Tsoi, A. C. (1998). Universal approximation using feedforward neural networks: A survey of some existing methods, and some new results. *Neural networks*, 11(1):15–37.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. (2014). Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958.
- Ståhl, N., Falkman, G., Karlsson, A., and Mathiason, G. (2020). Evaluation of uncertainty quantification in deep learning. In *International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems*, pages 556–568. Springer.
- Sun, Y. and Liang, F. (2022). A kernel-expanded stochastic neural network. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, n/a(n/a).
- Sun, Y., Xiong, W., and Liang, F. (2021). Sparse deep learning: A new framework immune to local traps and miscalibration. *Advances in Neural Information Processing Systems*, 34.
- Tagasovska, N. and Lopez-Paz, D. (2019). Single-model uncertainties for deep learning. *Advances in Neural Information Processing Systems*, 32.
- Teye, M., Azizpour, H., and Smith, K. (2018). Bayesian uncertainty estimation for batch normalized deep networks. In *International Conference on Machine Learning*, pages 4907–4916. PMLR.

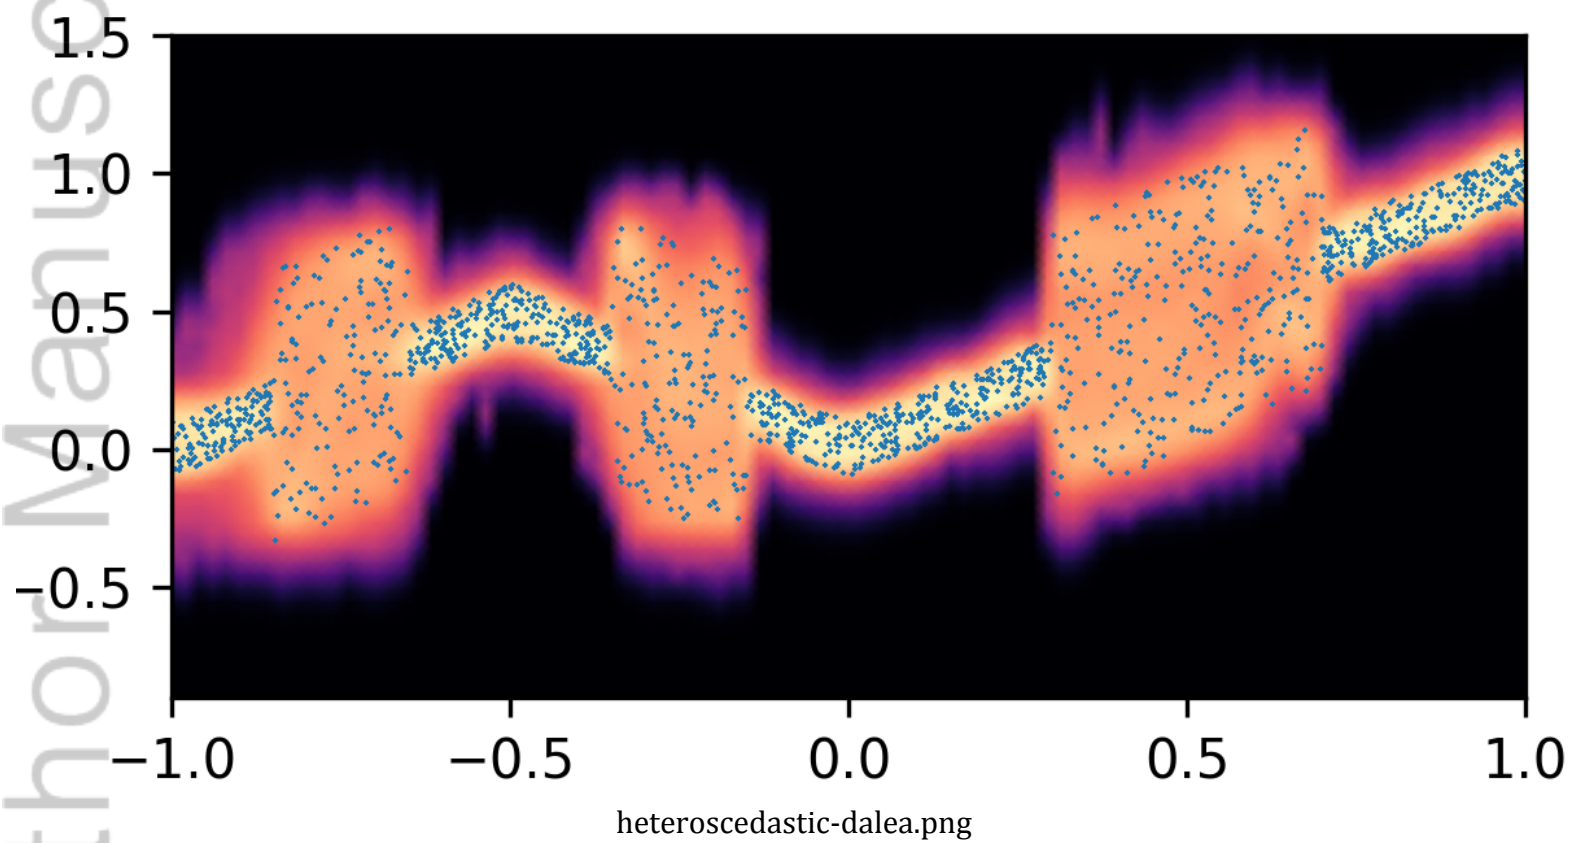
- Thompson, W. K., Barch, D. M., Bjork, J. M., Gonzalez, R., Nagel, B. J., Nixon, S. J., and Luciana, M. (2019). The structure of cognition in 9 and 10 year-old children and associations with problem behaviors: Findings from the abcd study’s baseline neurocognitive battery. *Developmental cognitive neuroscience*, 36:100606.
- Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., Mazoyer, B., and Joliot, M. (2002). Automated anatomical labeling of activations in spm using a macroscopic anatomical parcellation of the mni mri single-subject brain. *Neuroimage*, 15(1):273–289.
- Wang, Y. and Rocková, V. (2020). Uncertainty quantification for sparse deep learning. In *International Conference on Artificial Intelligence and Statistics*, pages 298–308. PMLR.
- Welling, M. and Teh, Y. W. (2011). Bayesian learning via stochastic gradient langevin dynamics. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pages 681–688. Citeseer.
- Wenzel, F., Roth, K., Veeling, B. S., Świątkowski, J., Tran, L., Mandt, S., Snoek, J., Salimans, T., Jenatton, R., and Nowozin, S. (2020). How good is the bayes posterior in deep neural networks really? *arXiv preprint arXiv:2002.02405*.
- Wilson, A. G. and Izmailov, P. (2020). Bayesian deep learning and a probabilistic perspective of generalization. *arXiv preprint arXiv:2002.08791*.
- Wu, T.-Y., Rachel Wang, Y., and Wong, W. H. (2020). Mini-batch metropolis-hastings with reversible sgld proposal. *Journal of the American Statistical Association*, pages 1–9.
- Xue, Y., Cheng, S., Li, Y., and Tian, L. (2019). Reliable deep-learning-based phase imaging with uncertainty quantification. *Optica*, 6(5):618–629.
- Yarotsky, D. (2017). Error bounds for approximations with deep relu networks. *Neural Networks*, 94:103–114.
- You, Z., Ye, J., Li, K., Xu, Z., and Wang, P. (2019). Adversarial noise layer: Regularize neural network by adding noise. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 909–913. IEEE.
- Zhang, D., Li, L., Sripada, C., and Kang, J. (2020). Image response regression via deep neural networks. *arXiv preprint arXiv:2006.09911*.
- Zhu, Y., Zabararas, N., Koutsourelakis, P.-S., and Perdikaris, P. (2019). Physics-constrained deep learning for high-dimensional surrogate modeling and uncertainty quantification without labeled data. *Journal of Computational Physics*, 394:56–81.























influence-mean.png





influence-mean-colorbar.png

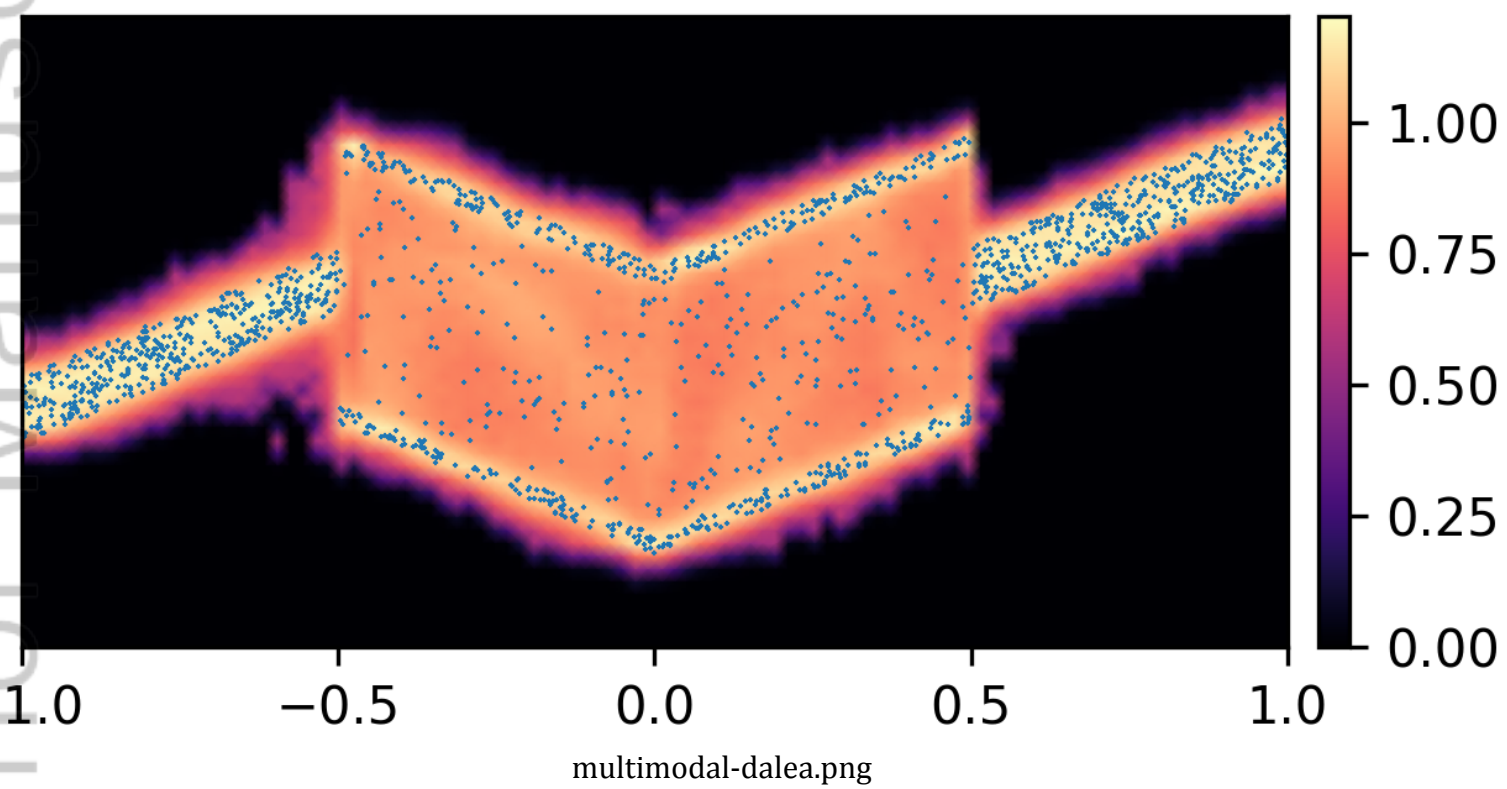






influence-variance-colorbar.png



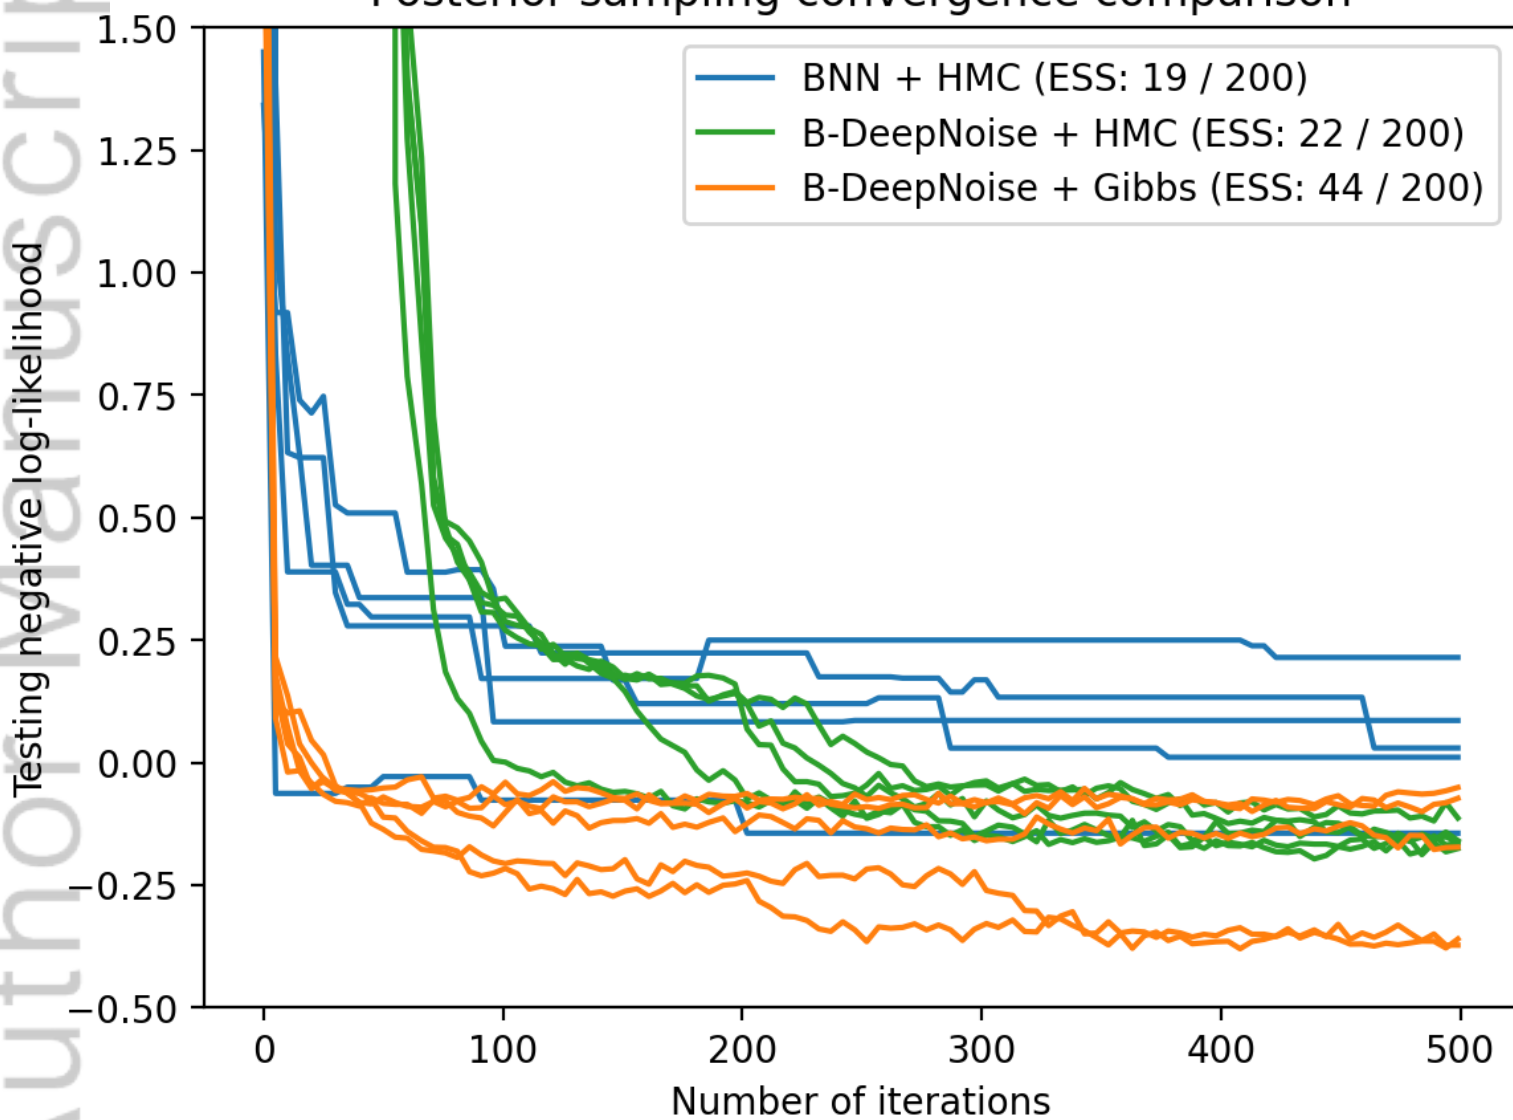








Posterior sampling convergence comparison



nll-trace.png

