
Tractable Function-Space Variational Inference in Bayesian Neural Networks

Tim G. J. Rudner* **Zonghao Chen** **Yee Whye Teh** **Yarin Gal**
University of Oxford University College London University of Oxford University of Oxford

Abstract

Reliable predictive uncertainty estimation plays an important role in enabling the deployment of neural networks to safety-critical settings. A popular approach for estimating the predictive uncertainty of neural networks is to define a prior distribution over the network parameters, infer an approximate posterior distribution, and use it to make stochastic predictions. However, explicit inference over neural network parameters makes it difficult to incorporate meaningful prior information about the data-generating process into the model. In this paper, we pursue an alternative approach. Recognizing that the primary object of interest in most settings is the distribution over functions induced by the posterior distribution over neural network parameters, we frame Bayesian inference in neural networks explicitly as inferring a posterior distribution over functions and propose a scalable function-space variational inference method that allows incorporating prior information and results in reliable predictive uncertainty estimates. We show that the proposed method leads to state-of-the-art uncertainty estimation and predictive performance on a range of prediction tasks and demonstrate that it performs well on a challenging safety-critical medical diagnosis task in which reliable uncertainty estimation is essential.

1 Introduction

Machine learning models succeed at an increasingly wide range of narrowly defined tasks [Krizhevsky et al., 2012, Mnih et al., 2013, Silver et al., 2016, Jumper et al., 2021] but may fail without warning when used on inputs that are meaningfully different from the data they were trained on [Amodei et al., 2016, Hendrycks et al., 2021, Rudner and Toner, 2021a,b]. To deploy machine learning models in safety-critical environments where failures are costly or may endanger human lives, machine learning methods must be reliable and have the ability to ‘fail gracefully.’ A promising tool for incorporating fail-safe mechanisms into machine learning systems, predictive uncertainty quantification allows machine learning models to express their confidence in the correctness of their predictions.

In this paper, we develop a method for obtaining reliable uncertainty estimates in Bayesian neural networks (BNNs, Neal [1996]). While BNNs have promised to combine the advantages of deep learning and Bayesian inference, existing approaches for approximate inference in BNNs fall short of this promise and have been demonstrated to result in approximate posterior predictive distributions that underperform ‘non-Bayesian’ methods both in terms of predictive accuracy and uncertainty quantification—making them of limited use in practice [Ovadia et al., 2019, Foong et al., 2019, Farquhar et al., 2020a, Band et al., 2021]. A potential reason for this shortcoming is that commonly used parameter-space inference methods make it difficult to define meaningful priors that effectively incorporate information about the data-generating process into inference.

*Corresponding author. Email: <tim.rudner@cs.ox.ac.uk>.

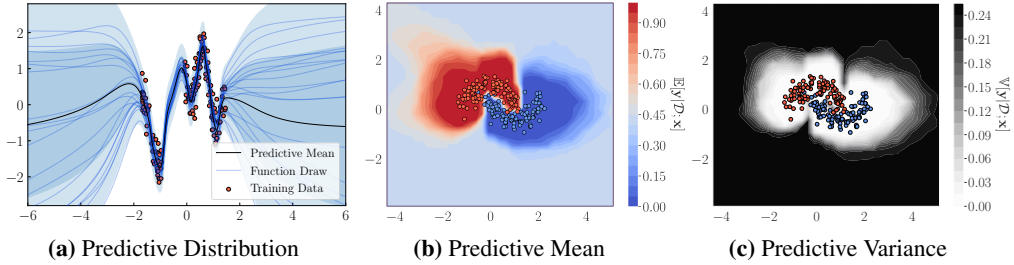


Figure 1: 1D regression on the *Snelson* dataset and binary classification on the *Two Moons* dataset. The plots show the predictive distributions of a BNNs, obtained via function-space variational inference (FSVI). For further illustrative examples and comparisons to deep ensembles and BNNs learned via parameter-space variational inference, see Appendix B.

To avoid this limitation, we follow Sun et al. [2019] and consider a variational objective defined explicitly in terms of distributions over *functions* induced by distributions over parameters. In contrast to prior works that rely on approximation techniques that prevent such function-space variational objectives to be used with high-dimensional inputs and highly-overparameterized neural networks, we propose a simple estimator of the Kullback-Leibler divergence between distributions over functions that enables us to perform stochastic variational inference. The proposed estimation procedure allows defining priors that explicitly encourage high predictive uncertainty away from the training data as well as priors that reflect relevant information about the task at hand.

We demonstrate that this approach leads to posterior approximations that exhibit significantly improved predictive uncertainty estimates compared to a wide array of state-of-the-art Bayesian and non-Bayesian methods. Figure 1 shows examples of predictive distributions obtained via function-space variational inference on low-dimensional, easy-to-visualize datasets. As can be seen in the figures, the predictive distributions fit the training data well while also exhibiting a high degree of predictive uncertainty in parts of the input space far away from the training data, as desired.

Contributions. We propose a simple estimation procedure for performing function-space variational inference in BNNs. The variational method allows for the incorporation of meaningful prior information about the data-generating process into the inference and produces reliable predictive uncertainty estimates. We perform a thorough empirical evaluation in which we compare the proposed approach to a wide array of competing methods and show that it consistently results in high predictive performance and reliable predictive uncertainty estimates, outperforming other methods in terms of predictive accuracy, robustness to distribution shifts, and uncertainty-based detection of distributionally-shifted data samples. We evaluate the proposed method on standard benchmarking datasets as well as on a safety-critical medical diagnosis task in which reliable uncertainty estimation is essential.²

2 Preliminaries

We consider supervised learning tasks on data $\mathcal{D} \doteq \{(\mathbf{x}_n, \mathbf{y}_n)\}_{n=1}^N = (\mathbf{X}_{\mathcal{D}}, \mathbf{y}_{\mathcal{D}})$ with inputs $\mathbf{x}_n \in \mathcal{X} \subseteq \mathbb{R}^D$ and targets $\mathbf{y}_n \in \mathcal{Y}$, where $\mathcal{Y} \subseteq \mathbb{R}^Q$ for regression and $\mathcal{Y} \subseteq \{0, 1\}^Q$ for classification tasks. Bayesian neural networks (BNNs) are stochastic neural networks trained using (approximate) Bayesian inference. Denoting the parameters of such a stochastic neural network by the multivariate random variable $\Theta \in \mathbb{R}^P$ and letting the function mapping defined by a neural network architecture be given by $f: \mathcal{X} \times \mathbb{R}^P \rightarrow \mathbb{R}^Q$, we obtain a random function $f(\cdot; \Theta)$. For a parameter realization θ , we obtain a corresponding function realization, $f(\cdot; \theta)$. When evaluated at a finite collection of points $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^m$, $f(\mathbf{X}; \Theta)$ is a multivariate random variable and $f(\mathbf{X}; \theta)$ is a vector.

Letting $p_{\mathbf{y}|f(\mathbf{x}, \Theta)}$ be a likelihood function and $p_{\mathbf{y}_{\mathcal{D}}|f(\mathbf{X}_{\mathcal{D}}, \Theta)}$ be the likelihood of observing the targets $\mathbf{y}_{\mathcal{D}}$ under the stochastic function $f(\cdot; \Theta)$ evaluated at inputs $\mathbf{X}_{\mathcal{D}}$ and letting p_{Θ} be a prior distribution over the stochastic network parameters Θ , we can use Bayes' Theorem to find the posterior distribution, $p_{\Theta|\mathcal{D}}$ [MacKay, 1992, Neal, 1996]. However, since the mapping f

²Our code can be accessed at <https://github.com/timrudner/FSVI>.

is a nonlinear function of the stochastic parameters Θ , exact inference is analytically intractable. Variational inference is an approach that seeks to sidestep this intractability by framing posterior inference as a variational optimization problem, where the goal is to find a distribution q_{Θ} in a variational family $\mathcal{Q}_{q_{\Theta}}$ that solves the variational problem $\min_{q_{\Theta} \in \mathcal{Q}_{q_{\Theta}}} \mathbb{D}_{\text{KL}}(q_{\Theta} \parallel p_{\Theta|\mathcal{D}})$ [Wainwright and Jordan, 2008]. If $\mathcal{Q}_{q_{\Theta}}$ is the family of mean-field Gaussian distributions and the prior distribution over parameters p_{Θ} given by a diagonal Gaussian distribution, the resulting variational objective is amenable to stochastic variational inference and can be optimized using gradient-based methods [Hinton and van Camp, 1993, Graves, 2011, Hoffman et al., 2013, Blundell et al., 2015].

2.1 A Function-Space Perspective on Variational Inference in Bayesian Neural Networks

Instead of seeking to infer an approximate posterior distribution over parameters, we frame variational inference in stochastic neural networks as inferring an approximation to the posterior distribution over *functions* $p_{f(\cdot; \Theta)|\mathcal{D}}$ induced by the posterior distribution over parameters $p_{\Theta|\mathcal{D}}$, that is,

$$p_{f(\cdot; \Theta)|\mathcal{D}}(f(\cdot; \theta) | \mathcal{D}) = \int_{\mathbb{R}^P} p_{\Theta|\mathcal{D}}(\theta' | \mathcal{D}) \delta(f(\cdot; \theta) - f(\cdot; \theta')) d\theta', \quad (1)$$

where $\delta(\cdot)$ is the Dirac delta function [Wolpert, 1993]. Considering the prior distribution over functions $p_{f(\cdot; \Theta)}$ induced by a prior distribution over parameters p_{Θ} ,

$$p_{f(\cdot; \Theta)}(f(\cdot; \theta)) = \int_{\mathbb{R}^P} p_{\Theta}(\theta') \delta(f(\cdot; \theta) - f(\cdot; \theta')) d\theta', \quad (2)$$

and the variational distribution over functions $q_{f(\cdot; \Theta)}$ induced by a variational distribution over parameters q_{Θ} ,

$$q_{f(\cdot; \Theta)}(f(\cdot; \theta)) = \int_{\mathbb{R}^P} q_{\Theta}(\theta') \delta(f(\cdot; \theta) - f(\cdot; \theta')) d\theta', \quad (3)$$

we can express the problem of finding a posterior distribution over functions variationally as

$$\min_{q_{\Theta} \in \mathcal{Q}_{q_{\Theta}}} \mathbb{D}_{\text{KL}}(q_{f(\cdot; \Theta)} \parallel p_{f(\cdot; \Theta)|\mathcal{D}}), \quad (4)$$

which allows us to effectively incorporate meaningful prior information about the underlying data-generating process into training. As discussed by Burt et al. [2021], this variational objective is guaranteed to be well-defined for suitably chosen prior distributions over functions. Specifically, the KL divergence between two distributions over functions generated from different distributions over parameters applied to the same mapping (e.g., the same neural network architecture) is well-defined (i.e., finite) if the KL divergence between the distributions over parameters is finite, since, by the strong data processing inequality [Polyanskiy and Wu, 2017],

$$\mathbb{D}_{\text{KL}}(q_{f(\cdot; \Theta)} \parallel p_{f(\cdot; \Theta)}) \leq \mathbb{D}_{\text{KL}}(q_{\Theta} \parallel p_{\Theta}). \quad (5)$$

As a result, if $\mathbb{D}_{\text{KL}}(q_{\Theta} \parallel p_{\Theta}) < \infty$, which is the case for finite-dimensional parameter vectors Θ and q_{Θ} absolutely continuous with respect to p_{Θ} , then the function-space KL divergence is finite and thus well-defined as a variational objective.

Hence, for a likelihood function defined on a finite set of training targets $\mathbf{y}_{\mathcal{D}}$ and a suitably defined prior distribution over functions, we can express the variational problem above equivalently as the well-defined maximization problem $\max_{q_{\Theta} \in \mathcal{Q}_{\Theta}} \mathcal{F}(q_{\Theta})$ with

$$\mathcal{F}(q_{\Theta}) \doteq \mathbb{E}_{q_{f(\mathbf{x}_{\mathcal{D}}; \Theta)}} [\log p_{\mathbf{y}} | f(\mathbf{x}; \Theta)(\mathbf{y}_{\mathcal{D}} | f(\mathbf{X}_{\mathcal{D}}; \theta))] - \mathbb{D}_{\text{KL}}(q_{f(\cdot; \Theta)} \parallel p_{f(\cdot; \Theta)}), \quad (6)$$

where $\mathbb{D}_{\text{KL}}(q_{f(\cdot; \Theta)} \parallel p_{f(\cdot; \Theta)})$ is also a KL divergence between distributions over functions.

Unfortunately, evaluating the KL divergence in Equation (6) is in general intractable for arbitrary mappings f . To obtain a tractable objective, Sun et al. [2019] showed that $\mathbb{D}_{\text{KL}}(q_{f(\cdot; \Theta)} \parallel p_{f(\cdot; \Theta)})$ can be expressed as the supremum of the KL divergence from $q_{f(\cdot; \Theta)}$ to $p_{f(\cdot; \Theta)}$ over all *finite* sets of evaluation points, resulting in the objective function

$$\mathcal{F}(q_{\Theta}) = \mathbb{E}_{q_{f(\mathbf{x}_{\mathcal{D}}; \Theta)}} [\log p_{\mathbf{y}} | f(\mathbf{x}; \Theta)(\mathbf{y}_{\mathcal{D}} | f(\mathbf{X}_{\mathcal{D}}; \theta))] - \sup_{\mathbf{X} \in \mathcal{X}_{\mathbb{N}}} \mathbb{D}_{\text{KL}}(q_{f(\mathbf{x}; \Theta)} \parallel p_{f(\mathbf{x}; \Theta)}), \quad (7)$$

where $\mathcal{X}_{\mathbb{N}} \doteq \bigcup_{n \in \mathbb{N}} \{\mathbf{X} \in \mathcal{X}_n \mid \mathcal{X}_n \subseteq \mathbb{R}^{n \times D}\}$ is the collection of all finite sets of evaluation points. However, this objective function is still challenging to optimize in practice: The supremum cannot be obtained analytically and the KL divergence term itself is analytically intractable and difficult to estimate in high dimensions—even for a single evaluation point.

In the next section, we will describe an approximation and estimation procedure that allows scaling function-space variational inference to large neural networks and high-dimensional input data.

3 Deriving a Tractable Function-Space Variational Objective

The primary obstacle to computing the objective in Equation (6) is the KL divergence from $q_{f(\cdot; \Theta)}$ to $p_{f(\cdot; \Theta)}$. There are two reasons why the KL divergence in Equation (7) is intractable: First, for BNNs or other non-linear models, we do not have access to the probability density functions of the multivariate distributions $q_{f(\mathbf{X}; \Theta)}$ and $p_{f(\mathbf{X}; \Theta)}$; second, for all but extremely simple input spaces, we are unable to compute the supremum over all possible finite sets of evaluation points. In the remainder of this section, we outline an approach for obtaining an estimator of a locally accurate approximation to the KL divergence that allows for scalable gradient-based optimization of Equation (7).

We first approach the problem of computing the KL divergence between two BNNs evaluated at a finite set of points. To do so, we first derive tractable approximations to the distributions over functions $q_{f(\mathbf{X}; \Theta)}$ and $p_{f(\mathbf{X}; \Theta)}$. Next, we show that under these approximations, we are able to obtain a closed-form approximation to the KL divergence and describe a simple Monte Carlo estimator of the supremum in the function-space KL divergence.

3.1 Approximating Distributions over Functions via Local Linearization

To obtain an approximation to the probability distributions of $q_{f(\mathbf{X}; \Theta)}$ and $p_{f(\mathbf{X}; \Theta)}$, we use a first-order Taylor expansion of the *mapping* f about the mean parameters of q_{Θ} and p_{Θ} , respectively, and derive the induced distributions under the linearized mapping.

For a stochastic function $f(\cdot; \Theta)$ defined in terms of stochastic parameters Θ distributed according to distribution g_{Θ} with $\mathbf{m} \doteq \mathbb{E}_{g_{\Theta}}[\Theta]$ and $\mathbf{S} \doteq \text{Cov}_{g_{\Theta}}[\Theta]$, we denote the linearization of the stochastic function $f(\cdot; \Theta)$ about \mathbf{m} by

$$f(\cdot; \Theta) \approx \tilde{f}(\cdot; \mathbf{m}, \Theta) \doteq f(\cdot; \mathbf{m}) + \mathcal{J}(\cdot; \mathbf{m})(\Theta - \mathbf{m}), \quad (8)$$

where $\mathcal{J}(\cdot; \mathbf{m}) \doteq (\partial f(\cdot; \Theta) / \partial \Theta)|_{\Theta=\mathbf{m}}$ is the Jacobian of $f(\cdot; \Theta)$ evaluated at $\Theta = \mathbf{m}$, and the mean and covariance of the distribution over the linearized mapping \tilde{f} at $\mathbf{X}, \mathbf{X}' \in \mathcal{X}$ are given by

$$\mathbb{E}[\tilde{f}(\mathbf{X}; \Theta)] = f(\mathbf{X}; \mathbf{m}) \quad (9)$$

$$\text{Cov}[\tilde{f}(\mathbf{X}; \Theta), \tilde{f}(\mathbf{X}'; \Theta)] = \mathcal{J}(\mathbf{X}; \mathbf{m}) \mathbf{S} \mathcal{J}(\mathbf{X}'; \mathbf{m})^{\top}. \quad (10)$$

For a derivation of this result, see Appendix A. Since Gaussianity is preserved under affine transformations, if g_{Θ} is a multivariate Gaussian distribution with mean \mathbf{m} and diagonal co-variance \mathbf{S} , then the distribution \tilde{g} over $\tilde{f}(\mathbf{X}; \Theta)$ is given by

$$\tilde{g}_{\tilde{f}(\mathbf{X}; \mathbf{m}, \Theta)} = \mathcal{N}(f(\mathbf{X}; \mathbf{m}), \mathcal{J}(\mathbf{X}; \mathbf{m}) \mathbf{S} \mathcal{J}(\mathbf{X}; \mathbf{m})^{\top}). \quad (11)$$

For stochastic functions parameterized by many millions of parameters, obtaining the covariance of $\tilde{g}_{\tilde{f}(\mathbf{X}; \Theta)}$ —which requires computing an inner product of two Jacobian matrices—can be computationally expensive. Instead of computing the distribution over the linearized mapping exactly, we can construct a suitable Monte Carlo estimator. To do so, we consider a partition of the set of parameters into sets α and β (with $|\beta| \ll |\alpha|$) and note that the linearized mapping can then be expressed as

$$\tilde{f}(\cdot; \mathbf{m}, \Theta) = f(\cdot; \mathbf{m}) + \tilde{f}_{\alpha}(\cdot; \mathbf{m}, \Theta_{\alpha}) + \mathcal{J}_{\beta}(\cdot; \mathbf{m})(\Theta_{\beta} - \mathbf{m}_{\beta}), \quad (12)$$

with

$$\tilde{f}_{\alpha}(\cdot; \mathbf{m}, \Theta_{\alpha}) \doteq \mathcal{J}_{\alpha}(\cdot; \mathbf{m})(\Theta_{\alpha} - \mathbf{m}_{\alpha}), \quad (13)$$

where $\mathcal{J}_{\alpha}(\cdot; \mathbf{m})$ and $\mathcal{J}_{\beta}(\cdot; \mathbf{m})$ are the columns of the Jacobian matrix corresponding to the sets of parameters α and β , respectively, and Θ_{α} and Θ_{β} are the corresponding random parameter vectors. Noting that Equation (12) expresses \tilde{f} as a sum of (affine transformations of) random variables, we can use the fact that for independent Gaussian random variables \mathbf{X} and \mathbf{Y} , the distribution $h_{\mathbf{Z}}$ of $\mathbf{Z} = \mathbf{X} + \mathbf{Y}$ is equal to the convolution of the distributions $h_{\mathbf{X}}$ and $h_{\mathbf{Y}}$ to obtain an approximation to \tilde{f} . In particular, we can show that if g_{Θ} is a multivariate Gaussian distribution with $\Theta_{\alpha} \perp \Theta_{\beta}$, the distribution $\tilde{g}_{\tilde{f}(\mathbf{X}; \Theta)}$ can be approximated by the Monte Carlo estimator

$$\hat{g}_{\tilde{f}(\mathbf{X}; \mathbf{m}, \Theta)} = \frac{1}{R} \sum_{j=1}^R \mathcal{N}\left(f(\mathbf{X}; \mathbf{m}) + \tilde{f}_{\alpha}(\mathbf{X}; \mathbf{m}, \Theta_{\alpha}^{(j)}), \mathcal{J}_{\beta}(\mathbf{X}; \mathbf{m}) \mathbf{S}_{\beta} \mathcal{J}_{\beta}(\mathbf{X}; \mathbf{m})^{\top}\right), \quad (14)$$

where $g_{\Theta_{\beta}} = \mathcal{N}(\mathbf{m}_{\beta}, \mathbf{S}_{\beta})$ and samples $\tilde{f}_{\alpha}(\mathbf{X}; \mathbf{m}, \Theta_{\alpha}^{(j)})$ are obtained by sampling parameters from the distribution $g_{\Theta_{\alpha}} = \mathcal{N}(\mathbf{m}_{\alpha}, \mathbf{S}_{\alpha})$. For a derivation of this result, see Appendix A. This estimator is biased for finite R but converges to $\tilde{g}_{\tilde{f}(\mathbf{X}; \mathbf{m}, \Theta)}$ as $R \rightarrow \infty$. Similarly, for finite R , the smaller $[\mathbf{S}_{\alpha}]_{ii}$, the more accurate and less biased the estimator will be. In our empirical evaluation, we use a single Monte Carlo sample, $R = 1$, to preserve Gaussianity and choose α to be the set of parameters in neural network layers $1 : L - 1$ and β to be the set of parameters in the final neural network layer.

3.2 Approximating the Function-Space Kullback-Leibler Divergence

From Section 3.1, we know that if q_{Θ} and p_{Θ} are both Gaussian distributions, then the induced distributions under the linearized mapping \tilde{f} evaluated at a finite set of evaluation points will be Gaussian as well. This means that for Gaussian variational and prior distributions over Θ , we can obtain locally accurate approximations to the induced distributions $q_{f(\cdot; \Theta)}$ to $p_{f(\cdot; \Theta)}$ and use them to approximate the KL divergence in the variational objective by $\mathbb{D}_{\text{KL}}(\tilde{q}_{\tilde{f}(\mathbf{x}; \Theta)} \parallel \tilde{p}_{\tilde{f}(\mathbf{x}; \Theta)})$. Moreover, for an isotropic Gaussian prior and a mean-field Gaussian variational distribution, $\mathbb{D}_{\text{KL}}(\tilde{q}_{\tilde{f}(\mathbf{x}; \Theta)} \parallel \tilde{p}_{\tilde{f}(\mathbf{x}; \Theta)})$ is a KL divergence between two multivariate Gaussians and can be obtained analytically.

Using this approximation, we obtain an estimator of the variational objective given by

$$\tilde{\mathcal{F}}(q_{\Theta}) \doteq \mathbb{E}_{q_{f(\mathbf{x}_{\mathcal{D}}; \Theta)}} [\log p_{\mathbf{y}|f(\mathbf{x}; \Theta)}(\mathbf{y}_{\mathcal{D}} | f(\mathbf{X}_{\mathcal{D}}; \theta))] - \sup_{\mathbf{X} \in \mathcal{X}_{\mathbb{N}}} \mathbb{D}_{\text{KL}}(\tilde{q}_{\tilde{f}(\mathbf{x}; \Theta)} \parallel \tilde{p}_{\tilde{f}(\mathbf{x}; \Theta)}), \quad (15)$$

where the arguments of the KL divergence have been replaced by the (locally accurate) approximations to the variational and prior distributions over functions evaluated at \mathbf{X} , respectively. Since the stochastic functions $\tilde{f}(\cdot; \Theta)$ induced by q_{Θ} and p_{Θ} under the linearized mapping will be closer to the stochastic function under f the smaller the variance of q_{Θ} and p_{Θ} , respectively, the approximation to the KL divergence will be more accurate the smaller the variance of q_{Θ} and p_{Θ} .

Next, we turn to computing the supremum. Unlike Sun et al. [2019], who consider the supremum as a separate optimization problem, we do not seek to compute the supremum by searching over points $\mathbf{X} \in \mathcal{X}_{\mathbb{N}}$ but instead propose to estimate the supremum at every gradient step via a simple finite-sample estimator. Specifically, letting $I(\mathbf{X}) \doteq \mathbb{D}_{\text{KL}}(\tilde{q}_{\tilde{f}(\mathbf{x}; \Theta)} \parallel \tilde{p}_{\tilde{f}(\mathbf{x}; \Theta)})$, we estimate $G = \sup_{\mathbf{X} \in \mathcal{X}_{\mathbb{N}}} I(\mathbf{X})$ using the Monte Carlo estimator

$$\hat{G}(\mathcal{X}_{\mathcal{C}}^S) = \max_{\mathbf{X} \in \mathcal{X}_{\mathcal{C}}^S} I(\mathbf{X}), \quad (16)$$

where $\mathcal{X}_{\mathcal{C}}^S \doteq \{\mathbf{X}_{\mathcal{C}}^{(i)}\}_{i=1}^S$ is a collection of S sets of *context points* $\mathbf{X}_{\mathcal{C}}^{(i)} \doteq \{\mathbf{x}^{(j)}\}_{j=1}^K$ jointly sampled from a context distribution $p_{\mathcal{X}_{\mathcal{C}}}$. Each context set $\mathbf{X}_{\mathcal{C}}^{(i)}$ can be viewed as a single Monte Carlo sample from the input space so that the estimator $\hat{G}(\mathcal{X}_{\mathcal{C}}^S)$ provides an S -sample Monte Carlo estimate of the supremum. While this estimator is crude and only provides a rough approximation to the true supremum, it encourages the variational distribution over functions to match the prior distribution over functions on the sets of context points. The choice of the context distribution $p_{\mathcal{X}_{\mathcal{C}}}$ can be informed by knowledge about the prediction task and should be viewed as a problem-specific modeling choice. Similarly, the numbers of samples S and K are hyperparameters to be optimized with a validation set. For details on how $p_{\mathcal{X}_{\mathcal{C}}}$ is chosen for the empirical evaluation in Section 5, see Appendix D.

3.3 Stochastic Estimation of the Approximate Function-Space Variational Objective

Let q_{Θ} be a Gaussian mean-field variational distribution, let p_{Θ} be an isotropic Gaussian prior, let $(\mathbf{X}_{\mathcal{B}}, \mathbf{y}_{\mathcal{B}})$ be a mini-batch of the training data, and reparameterize Θ as $\hat{\Theta}(\boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\epsilon}^{(j)}) \doteq \boldsymbol{\mu} + \boldsymbol{\Sigma} \odot \boldsymbol{\epsilon}^{(j)}$. Using the estimator $\hat{G}(\mathcal{X}_{\mathcal{C}}^S)$ defined above and estimating the expected log-likelihood via Monte Carlo sampling, we obtain a Monte Carlo estimator for the function-space variational objective:

$$\bar{\mathcal{F}}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) = \frac{1}{M} \sum_{j=1}^M \log p_{\mathbf{y}|f(\mathbf{x}; \Theta)}(\mathbf{y}_{\mathcal{B}} | f(\mathbf{X}_{\mathcal{B}}; \hat{\Theta}(\boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\epsilon}^{(j)}))) - \max_{\mathbf{X} \in \mathcal{X}_{\mathcal{C}}^S} \mathbb{D}_{\text{KL}}(\tilde{q}_{\tilde{f}(\mathbf{x}; \hat{\Theta})} \parallel \tilde{p}_{\tilde{f}(\mathbf{x}; \hat{\Theta})}) \quad (17)$$

with $\boldsymbol{\epsilon}^{(j)} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_P)$ and $\mathcal{X}_{\mathcal{C}}^S$ as defined above. This Monte Carlo estimator is biased due to the linearization and context-set approximations but allows for scalable gradient-based stochastic optimization.

Selection of Prior. For all experiments that involve uncertainty quantification, we chose a prior distribution over parameters that induces a prior distribution over functions $p_{f(\cdot; \Theta)}$ and a prior predictive distribution that exhibit a high degree of predictive uncertainty at evaluation points from regions in input space where $p_{\mathcal{X}_{\mathcal{C}}}$ has non-zero support and, under smoothness constraints, on evaluation points in nearby regions. For settings where prior information is encoded in data—for example, in the form of expert demonstrations of robotic manipulation tasks [Rudner et al., 2021] or in the form of pre-trained networks in continual or transfer learning [Rudner et al., 2022]—an empirical prior that reflects this information can be specified. For further details, see Appendix D.

Selection of Context Distribution. The distribution $p_{\mathcal{X}_C}$ allows us to incorporate information about the data-generating process into training and encourage the variational distribution to match the prior over functions in relevant parts of the input space. By taking advantage of the abundance of data available in real-world settings, context distributions can be constructed from large datasets like ImageNet [Krizhevsky et al., 2012], from small but diverse datasets like CIFAR-100, or by using any set of task-related unlabeled data. In our experiments, we choose two types of context distributions. One of the context distributions is constructed from the training data and only contains randomly sampled monochrome images, and one is constructed from a real-world dataset generated from a data distribution related to that of the training data. For example, when training on FashionMNIST, we use KMNIST as the context distribution, and when training on CIFAR-10, we use CIFAR-100 as the context distribution. For further details, see Appendix D.

Posterior Predictive Distribution. After optimizing the variational objective with respect to the parameters of the variational distribution q_{Θ} , we use the fact that we can obtain function draws by sampling from the distribution over parameters to obtain an approximate posterior predictive distribution

$$\begin{aligned} q(y_* | \mathbf{x}_*) &= \int p(y_* | f(\mathbf{x}_*; \theta)) q_{f(\mathbf{x}_*; \Theta)} d f(\mathbf{x}_*; \theta) \\ &\approx \frac{1}{M_*} \sum_{j=1}^{M_*} p(y_* | f(\mathbf{x}_*; \Theta^{(j)})) \quad \text{with } \Theta^{(j)} \sim q_{\Theta}, \end{aligned} \tag{18}$$

where M_* is the number of Monte Carlo samples used to estimate the predictive distribution.

4 Related Work

There is a growing body of work on function-space approaches to inference in BNNs, deep learning, and applications such as continual learning [Benjamin et al., 2019, Sun et al., 2019, Titsias et al., 2020, Burt et al., 2021, Pan et al., 2020, Ma and Hernández-Lobato, 2021, Rudner et al., 2022].

Function-Space Inference in Bayesian Neural Networks. Previously proposed methods for FSVI in BNNs are based on approximate gradient estimators and either replace the supremum in Equation (7) with an expectation [Sun et al., 2019] or do not define an explicit variational objective [Wang et al., 2019]. Sun et al. [2019] and Carvalho et al. [2020] use Gaussian process priors over functions for which the function-space variational inference problem is not well-defined (see Section 2.1 and Burt et al. [2021]). More recent work has attempted to circumvent the intractability of the variational objective in Equation (6) by proposing alternative objectives for function-space inference in BNNs [Ma et al., 2019, Ober and Aitchison, 2020, Ma and Hernández-Lobato, 2021]. Rudner et al. [2022] extend the approach presented in Section 3 to sequential inference problems and apply it to continual learning.

Linear Models. Immer et al. [2020] and Khan et al. [2019] show that approximate BNN posterior distribution via the Laplace and Generalized-Gauss-Newton approximation corresponds to exact posteriors under linearizations of different models. Unlike in our approach, they use a Laplace approximation and do not perform variational inference and do not optimize the variance parameters. Furthermore, Immer et al. [2020] and Khan et al. [2019] use a neural network model to obtain a parameter maximum a posteriori estimate, but then use a linearization of the neural network model to compute a posterior predictive distribution. In contrast, our work only uses the linearization to obtain an estimator of the variational objective but uses the unlinearized model to construct a posterior predictive distribution.

Pathologies of Variational Inference in Bayesian Neural Networks. Burt et al. [2021] consider the function-space variational objective in Equation (6) and show that the KL divergence between BNNs with different networks architectures are not well-defined. A parallel line of research showed that posterior predictive distributions of shallow BNNs with mean-field variational distributions have a limited ability to represent complex covariance structures in function space [Foong et al., 2019, 2020] but that deep BNNs do not suffer from this limitation [Farquhar et al., 2020b]. Our results are consistent with the findings of Farquhar et al. [2020b] that mean-field variational distributions are able to represent complex covariance structures in function space.

Table 1: Comparison of in- and out-of-distribution performance metrics on FashionMNIST (mean \pm standard error over ten random seeds). The last two columns show the AUROC for binary in- vs. out-of-distribution detection on MNIST (M) and NotMNIST (NM). MNIST and NotMNIST are used as out-of-distribution datasets. Best overall results for single and ensemble models are printed in boldface with gray shading. Results within a 95% confidence interval of the best overall result are printed in boldface only. All methods use the same four-layer CNN architecture. For further details about model architectures and training and evaluation protocols, see Appendix D.

Method	Accuracy \uparrow	ECE \downarrow	AUROC M \uparrow	AUROC NM \uparrow
MAP	91.73 \pm 0.08	0.037 \pm 0.001	87.00 \pm 0.30	74.85 \pm 1.31
MFVI [Blundell et al., 2015]	91.03 \pm 0.04	0.038 \pm 0.001	93.10 \pm 0.34	88.88 \pm 0.74
MFVI (tempered)	91.38 \pm 0.05	0.058 \pm 0.001	86.30 \pm 0.29	80.78 \pm 0.68
MFVI (radial) [Farquhar et al., 2020a]	90.31 \pm 0.11	0.035 \pm 0.001	84.40 \pm 0.68	82.11 \pm 1.15
MC DROPOUT [Gal and Ghahramani, 2016]	90.55 \pm 0.04	0.012 \pm 0.001	88.46 \pm 0.57	80.02 \pm 1.04
SWAG [Maddox et al., 2019]	92.56 \pm 0.05	0.043 \pm 0.001	85.18 \pm 0.35	80.31 \pm 0.30
DUQ [van Amersfoort et al., 2020]	92.40 \pm 0.20	—	95.50 \pm 0.70	94.60 \pm 1.80
BNN-LAPLACE [Immer et al., 2020]	92.25 \pm 0.10	0.012 \pm 0.003	95.55 \pm 0.60	—
SPG [Ma and Hernández-Lobato, 2021]	91.60 \pm 0.14	—	95.60 \pm 6.00	—
FSVI ($p_{\mathbf{x}_c}$ = random monochrome)	93.13 \pm 0.13	0.012 \pm 0.002	96.23 \pm 0.46	95.02 \pm 0.69
FSVI ($p_{\mathbf{x}_c}$ = KMNIST)	93.48 \pm 0.12	0.010 \pm 0.001	99.80 \pm 0.20	97.26 \pm 0.23
Deep Ensemble	92.49 \pm 0.01	0.019 \pm 0.000	89.22 \pm 0.09	83.17 \pm 0.91
FSVI Ensemble ($p_{\mathbf{x}_c}$ = random monochrome)	94.44 \pm 0.07	0.020 \pm 0.001	97.85 \pm 0.15	96.95 \pm 0.20

5 Empirical Evaluation

In this section, we evaluate FSVI on high-dimensional classification tasks that were out of reach for function-space variational inference methods proposed in prior works and compare FSVI to several well-established and state-of-the-art Bayesian deep learning and deterministic uncertainty quantification methods. We show that FSVI (sometimes *significantly*) outperforms existing Bayesian and non-Bayesian methods in terms of their in-distribution uncertainty calibration and out-of-distribution predictive uncertainty estimation. For a details on models, training and validation procedures, and datasets used, see Appendix D. For a comparison to Sun et al. [2019] on small-scale regression tasks, see Appendix B.2.

5.1 Predictive Performance, Uncertainty Estimation, and Distribution Shift Detection

In this set of experiments, we assess the reliability of the uncertainty estimates generated by FSVI. If a BNN trained via FSVI is able to perform reliable uncertainty estimation, its predictive uncertainty will be significantly higher on input points that were generated according to a different data-generating distribution than the training data. For models trained on the FashionMNIST dataset, we use the MNIST and NotMNIST datasets as out-of-distribution evaluation points, while for models trained on the CIFAR-10 dataset, we use the SVHN dataset as out-of-distribution evaluation points.

For models trained on either FashionMNIST or CIFAR-10, we evaluate their in-distribution performance in terms of test accuracy, test log-likelihood, and test calibration. To evaluate the quality of different models’ uncertainty estimates, we compute uncertainty estimates for the pairs FashionMNIST/MNIST, FashionMNIST/NotMNIST, and CIFAR-10/SVHN to and measure for a range of thresholds how well the datasets in each pair can be separated solely based on the uncertainty estimates. This experiment setup follows prior work by van Amersfoort et al. [2020] and Immer et al. [2020]. We report the area under the receiver operating characteristic (ROC) curve in Tables 1 and 2.

Predictive Performance and Calibration. To assess in-distribution predictive performance and calibration, we report the test accuracy, negative log-likelihood (NLL), and expected calibration error (ECE) for models trained on FashionMNIST and CIFAR-10 in Tables 1 and 2. On both FashionMNIST and CIFAR-10, FSVI achieves the lowest NLL and either the best or second-best predictive accuracy and ECE, respectively, across all methods. Notably, FSVI significantly outperforms SPG [Ma and Hernández-Lobato, 2021], an alternative function-space variational inference method.

Predictive Uncertainty under Distribution Shift. In Tables 1 and 2, we report evaluation metrics that elucidate the reliability of different methods’ predictive uncertainty under distribution shift. FSVI exhibits reliable predictive uncertainty estimates that allow distinguishing between in- and out-of-distribution inputs with high accuracy. As would be expected, we observe that using context

Table 2: Comparison of in- and out-of-distribution performance metrics on CIFAR-10 (mean \pm standard error over ten random seeds). SVHN and corrupted CIFAR-10 (C-CIFAR) are used as an out-of-distribution datasets. The penultimate column shows the AUROC for binary in- vs. out-of-distribution detection on SVHN. Best overall results for single and ensemble models are printed in boldface with gray shading. Results within a 95% confidence interval of the best overall result are printed in boldface only. All methods use a ResNet-18 architecture. For further details about model architectures and training and evaluation protocols, see [Appendix D](#).

Method	Accuracy \uparrow	ECE \downarrow	OOD-AUROC \uparrow	C-CIFAR Acc \uparrow
MAP	93.19 \pm 0.11	0.043 \pm 0.001	94.65 \pm 0.27	78.87 \pm 1.39
MFVI [Blundell et al., 2015]	89.98 \pm 0.09	0.040 \pm 0.001	92.14 \pm 0.34	79.36 \pm 1.35
MFVI (tempered)	90.87 \pm 0.11	0.048 \pm 0.001	91.82 \pm 0.90	79.86 \pm 1.32
MC DROPOUT [Gal and Ghahramani, 2016]	93.55 \pm 0.07	0.040 \pm 0.001	92.44 \pm 0.57	80.13 \pm 1.37
SWAG [Maddox et al., 2019]	93.13 \pm 0.14	0.067 \pm 0.002	89.79 \pm 0.50	76.12 \pm 0.51
VOGN [Osawa et al., 2019]	84.27 \pm 0.20	0.040 \pm 0.002	87.60 \pm 0.20	—
DUQ [van Amersfoort et al., 2020]	94.10 \pm 0.20	—	92.70 \pm 1.30	—
SPG [Ma and Hernández-Lobato, 2021]	77.69 \pm 0.64	—	88.30 \pm 4.00	—
FSVI ($p_{\mathbf{x}_c}$ = random monochrome)	93.35 \pm 0.04	0.034 \pm 0.001	94.76 \pm 0.24	80.81 \pm 0.43
FSVI ($p_{\mathbf{x}_c}$ = CIFAR-100)	93.57 \pm 0.04	0.026 \pm 0.001	98.07 \pm 0.10	81.20 \pm 0.42
Deep Ensemble	95.13 \pm 0.06	0.019 \pm 0.001	98.04 \pm 0.07	81.22 \pm 0.37
FSVI Ensemble ($p_{\mathbf{x}_c}$ = random monochrome)	95.19 \pm 0.03	0.013 \pm 0.001	99.19 \pm 0.41	81.35 \pm 0.48

distributions that reflect our knowledge about the data-generating process can significantly improve uncertainty quantification under FSVI. For the FashionMNIST experiment, we used the KMNIST dataset, which contains grayscale images of Kuzushiji letters, and for the CIFAR-10 experiment, we used the CIFAR-100 dataset, which contains RGB images of 100 classes. Both KMNIST and CIFAR-100 differ from the OOD datasets (MNIST and NotMNIST and SVHN, respectively) used to compute OOD-AUROC metrics in Tables 1 and 2, but using them as context distributions significantly increased the ability of BNNs trained via FSVI to identify distributionally shifted samples. Since the variational objective encourages matching the prior (which we chose to have high variance) on samples from the context distribution can improve uncertainty estimation in regions of the input space far from the training data.

5.2 Generalization and Reliability of Predictive Uncertainty under Distribution Shift

To assess the reliability of predictive models in deep learning, [Ovadia et al. \[2019\]](#) propose the following desiderata: In order for a model to be considered reliable, it ought to (i) exhibit low predictive uncertainty on training data and high predictive uncertainty on out-of-distribution inputs, (ii) generate predictive uncertainty estimates that allow distinguishing in- from out-of-distribution inputs, and (iii) if possible, maintain high predictive accuracy even under distribution shift. Models that satisfy these desiderata are less likely to make poor, high-confidence predictions and more amenable for use in safety-critical downstream tasks.

To illustrate these desiderata, we follow [Ovadia et al. \[2019\]](#) and consider the rotated MNIST task, where a model is trained on MNIST and evaluated on rotated MNIST digits. The goal is to maintain a high level of predictive accuracy (measured in terms of Brier scores) while exhibiting an increasing level of predictive uncertainty on distribution shifts of increasing magnitude. [Figure 2](#) shows Brier scores (lower is better) and predictive entropy estimates (higher means more uncertain) of four different models. As rotating the MNIST digits gradually shifts the data distributions, we would expect Brier scores to increase (corresponding to worse predictive accuracy) as the rotation angle increases. A model with reliable predictive entropy estimates would only experience a small decrease under distribution shift while exhibiting a large increase in predictive uncertainty. As can be seen in the plot, the Brier scores of FSVI decreases the least, while FSVI’s uncertainty is significantly higher than other models’. To assess the




Figure 2: Predictive uncertainty and accuracy on rotated MNIST. Models with reliable uncertainty estimates would exhibit higher predictive uncertainty the more the digits are rotated. Ideally, such models would maintain high predictive accuracy (low Brier score).

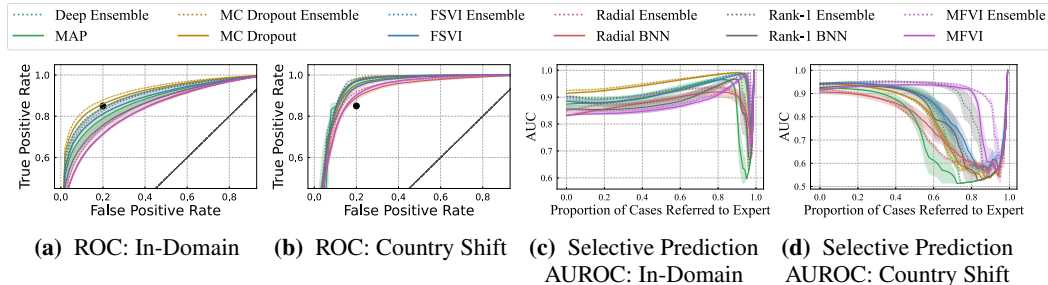


Figure 4: We jointly assess model predictive performance and uncertainty quantification on both in-domain and distributionally shifted data. **Left:** The *receiver operating characteristic curve* (ROC) for in-population diagnosis on the (a) EyePACS [2015] test set and for (b) changing medical equipment and patient populations on the APTOS [2019] test set. The dot in **black** denotes the NHS-recommended 85% sensitivity and 80% specificity ratios [Widdowson, 2016]. **Right:** *Selective prediction* on AUROC in (c) EyePACS [2015] and (d) APTOS [2019] settings. Shading denotes standard error over six random seeds. See Appendix B.1 for tabular results.

reliability of different uncertainty quantification methods on a more challenging distribution-shift task, we consider corrupted CIFAR-10 inputs under the second-mildest corruption level used in [Ovadia et al., 2019] and report our results in Table 2. Consistent with the rotated MNIST results, FSVI achieves the highest accuracy on the corrupted data.

5.3 Safety-Critical Uncertainty-Aware Selective Prediction: Diabetic Retinopathy Diagnosis

To evaluate the reliability of the predictive uncertainty of FSVI in a real-world safety-critical setting, we consider the task of diagnosing diabetic retinopathy (DR), a medical condition that can lead to impaired vision, from retina scans [Leibig et al., 2017, Filos et al., 2019, Band et al., 2021]. We use two publicly available datasets, EyePACS [2015] and APTOS [2019], each containing RGB images of a human retina graded by a medical expert on the following scale: 0 (no DR), 1 (mild DR), 2 (moderate DR), 3 (severe DR), and 4 (proliferative DR). The Kaggle dataset was collected from patients in the United States, while the APTOS dataset was collected from patients in India using cheaper but more modern scanning devices. We follow Leibig et al. [2017], Filos et al. [2019], and Band et al. [2021] and binarize all examples from both the EyePACS and APTOS datasets by dividing the classes up into sight-threatening diabetic retinopathy—defined as moderate diabetic retinopathy or worse (classes {2, 3, 4})—and non-sight-threatening diabetic retinopathy—defined as no or mild diabetic retinopathy (classes {0, 1}). This results in a binary prediction task.




Figure 3: Retina scan examples. **Top:** healthy. **Bottom:** unhealthy.

To assess the reliability of predictive models when medical training and test data are obtained from different patient populations or collected with the same medical equipment, we follow Band et al. [2021] and use the Kaggle dataset for training and the distributionally shifted APTOS dataset for evaluation. The results are shown in Figure 4, which plot the ROC curves for the binary prediction problems as well as the area under the ROC curve for an uncertainty aware selective prediction task. For further details about the uncertainty-aware selective prediction evaluation protocol, see Appendix D.4. Figure 4 shows that FSVI performs well on all four tasks and is only outperformed by MC DROPOUT. For full tabular results, see Appendix B.1.

6 Conclusion

The paper proposed a scalable and effective approach to function-space variational inference in BNNs. We demonstrated that the proposed estimator of the function-space variational objective can be scaled up to high-dimensional data and large neural network architectures and that FSVI exhibits consistently reliable in- and out-of-distribution predictive performance on a wide range of datasets when compared to well-established and state-of-the-art uncertainty quantification methods. We hope that this work will lead to further research into function-space variational inference and the development of more sophisticated data-driven prior distributions over functions.

Acknowledgements

We thank Bryn Elesedy, Bobby He, and Andrew Jesson for feedback on an early draft of this paper. We thank Joost van Amersfoort for helpful discussions about experiment design and implementations. Tim G. J. Rudner is funded by the Rhodes Trust and the Engineering and Physical Sciences Research Council (EPSRC). We gratefully acknowledge donations of computing resources by the Alan Turing Institute.

References

- Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. Concrete problems in ai safety, 2016.
- APTOS. APTOS 2019 Blindness Detection Dataset, 2019.
- Neil Band, Tim G. J. Rudner, Qixuan Feng, Angelos Filos, Zachary Nado, Michael W. Dusenberry, Ghassen Jerfel, Dustin Tran, and Yarin Gal. Benchmarking Bayesian Deep Learning on Diabetic Retinopathy Detection Tasks. 2021.
- Ari Benjamin, David Rolnick, and Konrad Kording. Measuring and regularizing networks in function space. In *International Conference on Learning Representations*, 2019.
- Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural networks. volume 37 of *Proceedings of Machine Learning Research*, pages 1613–1622, Lille, France, 07–09 Jul 2015. PMLR.
- David R. Burt, Sebastian W. Ober, Adrià Garriga-Alonso, and Mark van der Wilk. Understanding variational inference in function-space. In *Third Symposium on Advances in Approximate Bayesian Inference*, 2021.
- Eduardo D. C. Carvalho, Ronald Clark, Andrea Nicastro, and Paul H. J. Kelly. Scalable uncertainty for computer vision with functional variational inference. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- EyePACS. Diabetic Retinopathy Detection Dataset, 2015.
- Sebastian Farquhar, Michael A. Osborne, and Yarin Gal. Radial Bayesian neural networks: Beyond discrete support in large-scale Bayesian deep learning. In Silvia Chiappa and Roberto Calandra, editors, *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 1352–1362. PMLR, 26–28 Aug 2020a.
- Sebastian Farquhar, Lewis Smith, and Yarin Gal. Liberty or depth: Deep Bayesian neural nets do not need complex weight posterior approximations. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020b.
- Angelos Filos, Sebastian Farquhar, Aidan N. Gomez, Tim G. J. 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, 2019.
- Andrew Y. K. Foong, Yingzhen Li, José Miguel Hernández-Lobato, and Richard E. Turner. ‘in-between’ uncertainty in Bayesian neural networks, 2019.
- Andrew Y. K. Foong, David R. Burt, Yingzhen Li, and Richard E. Turner. On the expressiveness of approximate inference in Bayesian neural networks. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- Yarin Gal and Zoubin Ghahramani. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48, ICML 2016*, pages 1050–1059, 2016.
- Alex Graves. Practical variational inference for neural networks. In *Proceedings of the 24th International Conference on Neural Information Processing Systems, NIPS’11*, page 2348–2356, Red Hook, NY, USA, 2011. Curran Associates Inc. ISBN 9781618395993.
- Dan Hendrycks, Nicholas Carlini, John Schulman, and Jacob Steinhardt. Unsolved problems in ml safety, 2021.

- Geoffrey E. Hinton and Drew van Camp. Keeping the neural networks simple by minimizing the description length of the weights. In *Proceedings of the Sixth Annual Conference on Computational Learning Theory*, COLT '93, page 5–13, New York, NY, USA, 1993. Association for Computing Machinery. ISBN 0897916115.
- Matthew D. Hoffman, David M. Blei, Chong Wang, and John Paisley. Stochastic variational inference. *Journal of Machine Learning Research*, 14(1):1303–1347, May 2013. ISSN 1532-4435.
- Alexander Immer, Maciej Korzepa, and Matthias Bauer. Improving predictions of Bayesian neural networks via local linearization, 2020.
- Pavel Izmailov, Wesley J. Maddox, Polina Kirichenko, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. Subspace inference for Bayesian deep learning. In Ryan P. Adams and Vibhav Gogate, editors, *Proceedings of The 35th Uncertainty in Artificial Intelligence Conference*, volume 115 of *Proceedings of Machine Learning Research*, pages 1169–1179. PMLR, 22–25 Jul 2020.
- John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Zidek, Anna Potapenko, Alex Bridgland, Clemens Meyer, Simon A A Kohl, Andrew J Ballard, Andrew Cowie, Bernardino Romera-Paredes, Stanislav Nikolov, Rishub Jain, Jonas Adler, Trevor Back, Stig Petersen, David Reiman, Ellen Clancy, Michal Zielinski, Martin Steinegger, Michalina Pacholska, Tamas Berghammer, Sebastian Bodenstein, David Silver, Oriol Vinyals, Andrew W Senior, Koray Kavukcuoglu, Pushmeet Kohli, and Demis Hassabis. Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873):583–589, 2021. doi: 10.1038/s41586-021-03819-2.
- Mohammad Emtiyaz E Khan, Alexander Immer, Ehsan Abedi, and Maciej Korzepa. Approximate inference turns deep networks into Gaussian processes. In *Advances in Neural Information Processing Systems 32*, pages 3094–3104. Curran Associates, Inc., 2019.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. ImageNet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems 25.*, pages 1106–1114, 2012.
- Christian Leibig, Vaneeda Allken, Murat Seçkin Ayhan, Philipp Berens, and Siegfried Wahl. Leveraging Uncertainty Information From Deep Neural Networks for Disease Detection. *Nature Scientific Reports*, 7(1):17816, 2017.
- Chao Ma and José Miguel Hernández-Lobato. Functional variational inference based on stochastic process generators. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021.
- Chao Ma, Yingzhen Li, and Jose Miguel Hernandez-Lobato. Variational implicit processes. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 4222–4233. PMLR, 09–15 Jun 2019.
- David J. C. MacKay. A practical Bayesian framework for backpropagation networks. *Neural Comput.*, 4(3):448–472, May 1992. ISSN 0899-7667. doi: 10.1162/neco.1992.4.3.448.
- Wesley J Maddox, Pavel Izmailov, Timur Garipov, Dmitry P Vetrov, and Andrew Gordon Wilson. A simple baseline for Bayesian uncertainty in deep learning. In *Advances in Neural Information Processing Systems*, pages 13153–13164, 2019.
- Alexander G. de G. Matthews, James Hensman, Richard Turner, and Zoubin Ghahramani. On sparse variational methods and the Kullback-Leibler divergence between stochastic processes. volume 51 of *Proceedings of Machine Learning Research*, pages 231–239, Cadiz, Spain, 09–11 May 2016. PMLR.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. In *NIPS Deep Learning Workshop*. 2013.
- Radford M Neal. Bayesian Learning for Neural Networks. 1996.
- Sebastian W. Ober and Laurence Aitchison. Global inducing point variational posteriors for Bayesian neural networks and deep Gaussian processes, 2020.

- Kazuki Osawa, Siddharth Swaroop, Mohammad Emtiyaz E Khan, Anirudh Jain, Runa Eschenhagen, Richard E Turner, and Rio Yokota. Practical deep learning with Bayesian principles. In *Advances in Neural Information Processing Systems*, volume 32, pages 4287–4299. Curran Associates, Inc., 2019.
- Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, D. Sculley, Sebastian Nowozin, Joshua Dillon, Balaji Lakshminarayanan, and Jasper Snoek. Can you trust your model’s uncertainty? Evaluating predictive uncertainty under dataset shift. In *Advances in Neural Information Processing Systems* 32. 2019.
- Pingbo Pan, Siddharth Swaroop, Alexander Immer, Runa Eschenhagen, Richard E. Turner, and Mohammad Emtiyaz Khan. Continual deep learning by functional regularisation of memorable past. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- Yury Polyanskiy and Yihong Wu. Strong data-processing inequalities for channels and Bayesian networks. In Eric Carlen, Mokshay Madiman, and Elisabeth M. Werner, editors, *Convexity and Concentration*, pages 211–249, New York, NY, 2017. Springer New York. ISBN 978-1-4939-7005-6.
- Tim G. J. Rudner and Helen Toner. Key Concepts in AI Safety: An Overview. In *CSET Issue Briefs*, 2021a.
- Tim G. J. Rudner and Helen Toner. Key Concepts in AI Safety: Robustness and Adversarial Examples. In *CSET Issue Briefs*, 2021b.
- Tim G. J. Rudner, Cong Lu, Michael A. Osborne, Yarín Gal, and Yee Whye Teh. On Pathologies in KL-Regularized Reinforcement Learning from Expert Demonstrations. 2021.
- Tim G. J. Rudner, Freddie Bickford Smith, Qixuan Feng, Yee Whye Teh, and Yarín Gal. Continual Learning via Sequential Function-Space Variational Inference. In *Proceedings of the 38th International Conference on Machine Learning*, Proceedings of Machine Learning Research. PMLR, 2022.
- M. J. Schervish. *Theory of Statistics*. Springer-Verlag, New York, NY, 1995.
- David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the game of Go with deep neural networks and tree search. 529, 2016.
- Edward Snelson and Zoubin Ghahramani. Sparse Gaussian processes using pseudo-inputs. In Y. Weiss, B. Schölkopf, and J. C. Platt, editors, *Advances in Neural Information Processing Systems 18*, pages 1257–1264. MIT Press, 2006.
- Shengyang Sun, Guodong Zhang, Jiaxin Shi, and Roger B. Grosse. Functional variational Bayesian neural networks. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019.
- Michalis K. Titsias, Jonathan Schwarz, Alexander G. de G. Matthews, Razvan Pascanu, and Yee Whye Teh. Functional regularisation for continual learning with Gaussian processes. In *International Conference on Learning Representations*, 2020.
- Joost van Amersfoort, Lewis Smith, Yee Whye Teh, and Yarín Gal. Uncertainty estimation using a single deep deterministic neural network. In *International Conference on Machine Learning*, 2020.
- Joost van Amersfoort, Lewis Smith, Andrew Jesson, Oscar Key, and Yarín Gal. Variational deterministic uncertainty quantification, 2021.
- Martin J Wainwright and Michael I Jordan. *Graphical Models, Exponential Families, and Variational Inference*. Now Publishers Inc., Hanover, MA, USA, 2008. ISBN 1601981848.
- Ziyu Wang, Tongzheng Ren, Jun Zhu, and Bo Zhang. Function space particle optimization for Bayesian neural networks. In *International Conference on Learning Representations*, 2019.
- D. T. S. Widdowson. The Management of Grading Quality: Good Practice in the Quality Assurance of Grading. *Tech. Rep.*, 2016.
- David H. Wolpert. Bayesian backpropagation over i-o functions rather than weights. In J. Cowan, G. Tesauro, and J. Alspector, editors, *Advances in Neural Information Processing Systems*, volume 6. Morgan-Kaufmann, 1993.

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] See Section 5.
 - (b) Did you describe the limitations of your work? [Yes] See Section 5.
 - (c) Did you discuss any potential negative societal impacts of your work? [N/A]
 - (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 Appendix A.
 - (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 <https://github.com/timrudner/FSVI>.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Appendix D.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See Table 1, Table 2, and Table 3.
 - (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 D.9.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [Yes]
 - (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? [Yes] See <https://github.com/timrudner/FSVI>.
 - (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]