FUNCTION SPACE PARTICLE OPTIMIZATION FOR BAYESIAN NEURAL NETWORKS

Anonymous authors
Paper under double-blind review

ABSTRACT

While Bayesian neural networks (BNNs) have drawn increasing attention, their posterior inference remains challenging, due to the high-dimensional and over-parameterized nature. To address this issue, several highly flexible and scalable variational inference procedures based on the idea of *particle optimization* have been proposed. These methods directly optimize a set of particles to approximate the target posterior. However, their application to BNNs often yields sub-optimal performance, as they have a particular failure mode on over-parameterized models. In this paper, we propose to solve this issue by performing particle optimization directly in the space of regression functions. We demonstrate through extensive experiments that our method successfully overcomes this issue, and outperforms strong baselines in a variety of tasks including prediction, defense against adversarial examples, and reinforcement learning.

1 INTRODUCTION

Bayesian nerual networks (BNNs) provide a principled approach to reasoning about the *epistemic uncertainty*—uncertainty in model prediction due to the lack of knowledge. Recent work has demonstrated the potential of BNNs in safety-critical applications like medicine and autonomous driving, deep reinforcement learning, and defense against adversarial samples (see e.g. Ghosh et al., 2018; Zhang et al., 2018b; Feng et al., 2018; Smith & Gal, 2018).

Modeling with BNNs involves placing priors on neural network weights, and performing posterior inference with the observed data. However, posterior inference for BNNs is challenging, due to the multi-modal and high dimensional nature of the posterior. Variational inference (VI) is a commonly used technique for practical approximate inference. Traditional VI methods approximate the true posterior with oversimplified distribution families like factorized Gaussians, which severely limits the approximation quality and induces pathologies such as over-pruning (Trippe & Turner, 2018). These limitations motivate the development of implicit VI methods (Li & Turner, 2018; Shi et al., 2018b), which allow the use of flexible approximate distributions without a tractable density. However, most of the implicit inference methods require to learn a "generator network", which maps a simple distribution to approximate the target posterior. Inclusion of such a generator network can introduce extra complexity, and may become infeasible when the number of parameters is very large, as in the case for BNNs.

Compared with those generator-based methods, particle-optimization-based variational inference (POVI) methods constitute a simpler but more efficient class of implicit VI methods. In an algorithmic perspective, POVI methods iteratively update a set of *particles*, so that the corresponding empirical probability measure approximates the target posterior well. Formally, these methods consider the space of probabilistic measures equipped with different metrics, and simulate a gradient flow that converges to the target distribution. Examples of POVI methods include Stein variational gradient descent (SVGD; Liu & Wang, 2016), gradient flows in the 2-Wasserstein space (Chen et al., 2018), and accelerated first-order methods in the 2-Wasserstein space (Liu et al., 2018).

While POVI methods have shown success in a variety of challenging inference problems, their performance in BNNs is still far from ideal, as with a limited number of particles, it is hard to characterize the highly complex weight-space posterior. The first problem is the curse of dimensionality: Zhuo et al. (2018) and Wang et al. (2018) show that for SVGD with a RBF kernel, particles collapse to the maximum a posteriori (MAP) estimate as the dimension of parameter increases. One may

hope that such a problem is alleviated by switching to other POVI methods, which could be more suitable for BNN inference; however, this is not the case: BNN is *over-parameterized*, and there exist a large number of local modes in the weight-space posterior that are distant from each other, yet correspond to the same regression function. Thus a possible particle approximation is to place each particle in a different mode. In prediction, such an approximate posterior will not perform better than a single point estimate. In other words, good approximations for the weight-space posterior do not necessarily perform well in prediction.

To address the above issue, we propose to perform POVI directly for the posterior of regression functions, i.e. the *function-space posterior*, instead for the weight-space posterior. In our algorithm, particles correspond to regression functions, and will be approximated by weight-space parameters.We derive efficient rules to approximate the infinite-dimensional function-space updates. Our method Extensive experiments show that our proposal avoids the degenerate behavior of weight-space POVI methods, and leads to significant improvements on several tasks, including prediction, model robustness, and exploration in reinforcement learning.

The whole paper is organized as follows. We will first briefly review BNN and POVI in Section 2, then in Section 3 we present our algorithm for function-space POVI. We compare our method with existing work in Section 4, and finally demonstrate our method's effectiveness in Section 5.

2 BACKGROUND

Bayesian Neural Networks (BNNs) Consider a supervised learning problem, with $\mathbf{X} = \{x_1, \ldots, x_N\}$ being the training inputs and $\mathbf{Y} = \{y_1, \ldots, y_N\}$ being the corresponding outputs. Let $\mathcal{X}$ and $\mathcal{Y}$ be their ranges, respectively. A BNN model defines a prior $p(\theta)$ over neural network weights θ. The neural network then parameterizes a function $f(\cdot;\theta) : \mathcal{X} \to \mathcal{Y}$, which could have generated the observed data. The conditional distribution $p(y|x,\theta)$ specified the relation between $f(x;\theta)$ and y; for example, for real-valued regression, the conditional distribution could be $p(y|x,\theta) = \mathcal{N}(y|f(x;\theta),\sigma^2)$, where σ^2 is the variance of observation noise. Given $\mathbf{X}, \mathbf{Y}$, one then infers the posterior distribution $p(\theta|\mathbf{X},\mathbf{Y}) \propto p(\theta)p(\mathbf{Y}|\theta,\mathbf{X})$, and for a test data point x_{N+1}, y is predicted to have the distribution $p(y_{N+1}|x_{N+1},\mathbf{X},\mathbf{Y}) = \int p(y_{N+1}|x_{N+1},\theta)p(d\theta|\mathbf{X},\mathbf{Y})$.

Posterior inference for BNNs is generally difficult due to the high-dimensionality of θ. The *over-parameterized* nature of BNNs further exacerbates the problem: for over-parameterized models, there exist multiple θ that correspond to the same likelihood function $p(y|x,\theta)$. One could easily obtain an exponential number of such θ, by reordering the weights in the network. Each of the θ can be a mode of the posterior, which makes approximate inference particularly challenging for BNN.

Particle-Optimization based Variational Inference (POVI) Approximate inference aims to find an approximation of the true posterior. POVI methods view the approximate inference problem as minimizing some energy functionals over probability measures, which obtain their minimum at the true posterior. The optimization problem is then solved by simulating the corresponding gradient flow in certain metric spaces, i.e. to simulate a PDE of the form

$$\partial_t q_t = -\nabla(\mathbf{v}_t \cdot q_t),$$

where q_t is the approximate posterior at time t, and $\mathbf{v}$ is the gradient flow depending on the choice of metric and energy functional. As q_t can be arbitrarily flexible, it cannot be maintained exactly in simulation. Instead, we approximate it with a set of particles $\{\theta^{(i)}\}_{i=1}^n$, i.e. $q_t(\theta) \approx \frac{1}{n}\sum_{i=1}^n \delta(\theta - \theta_t^{(i)})$, and simulate the gradient flow with a discretized version of the ODE $\mathrm{d}\theta_t^{(i)}/\mathrm{d}t = -\mathbf{v}_t(\theta_t^{(i)})$. In other words, in each iteration, we update the particles with

$$\theta_{\ell+1}^i \leftarrow \theta_\ell^i - \epsilon_\ell \mathbf{v}(\theta_\ell^i), \tag{1}$$

where ϵ_ℓ is the step-size at the ℓ-th iteration, and $\mathbf{v}$ is the gradient flow. Table 1 summarizes the common choices of $\mathbf{v}$.

While the flexibility of POVI methods is unlimited in theory, the use of finite particles makes them less robust in practice, especially when applied to high-dimensional, over-parameterized models. The problem of high dimensionality is investigated in Zhuo et al. (2018) and Wang et al. (2018); we

Table 1: Common choices of $\mathbf{v}$ in POVI procedures. k denotes a kernel, $\mathbf{K}_{ij} := k(\theta^{(i)}, \theta^{(j)})$, and $K \in \mathbb{R}^{n \times n}$ satisfies $K_{ij} = \mathbf{K}_{ij}$. We omit the subscript t for brevity.

Method	$-\mathbf{v}(\theta^{(i)})$
SVGD (Liu & Wang, 2016)	$\frac{1}{n}\sum_{j=1}^{n} \mathbf{K}_{ij} \nabla_{\theta^{(j)}} \log p(\theta^{(j)}\vert\mathbf{x}) + \nabla_{\theta^{(j)}} \mathbf{K}_{ij}$
w-SGLD-B (Chen et al., 2018)	$\nabla_{\theta^{(i)}} \log p(\theta^{(i)}\vert\mathbf{x}) + \sum_{j=1}^{n} \nabla_{\theta^{(j)}} \mathbf{K}_{ji} / \sum_{k=1}^{n} \mathbf{K}_{jk} + \sum_{j=1}^{n} \nabla_{\theta^{(j)}} \mathbf{K}_{ji} / \sum_{k=1}^{n} \mathbf{K}_{ik}$
pi-SGLD (Chen et al., 2018)	sum of $-\mathbf{v}$ in SVGD and w-SGLD-B
GFSF (Liu et al., 2018)	$\nabla_{\theta^{(i)}} \log p(\theta^{(i)}\vert\mathbf{x}) + \sum_{j=1}^{n} (K^{-1})_{ij} \nabla_{\theta^{(j)}} \mathbf{K}_{ij}$

give an intuitive explanation of the over-parameterization problem here[1]: in an over-parameterized model like BNN, the target posterior has a large number of modes that are sufficiently distant from each other, yet corresponding to the same regression function. A possible convergence point for POVI with finite particles is thus to occupy all these modes, as such a configuration has a small 2-Wasserstein distance to the true posterior. In this case, prediction using these particles will not improve over using the MAP estimate. Such a degeneracy is actually observed in practice; see Section 5.1 and Appendix A.1.

3 Function Space Particle Optimization

Function Space Inference When we model with BNNs, there exists a map from the network weights θ to a corresponding regression function f, $\theta \mapsto f(\cdot;\theta)$. Therefore, the prior on θ implicitly defines a prior measure on the space of f, denoted as $p(f)$. Furthermore, the conditional distribution $p(y|x,\theta)$ also corresponds to a conditional distribution of $p(y|x,f)$. Posterior inference for network weights can thus be viewed as posterior inference for the latent regression function f.

A nice property of the function space inference problem is that it does not suffer from the over-parameterization problem. However, it is hard to implement, as functions are infinite-dimensional, and the corresponding prior measure is implicitly defined. In this section, we present a heuristic approach to function-space inference with POVI.

Notations In the following subsections, we will use $\mathbf{X}$ and $\mathbf{Y}$ to denote the observed training data, boldfaced symbols (e.g. $\mathbf{x}, \mathbf{y}$) to denote a subset of $\mathcal{X}$ or $\mathcal{Y}$, and regular symbols (e.g. x, y) to denote single data points. We denote the approximate posterior measure as $q(f)$. For any finite subset $\mathbf{x} \subset \mathcal{X}$, we use $f(\mathbf{x})$ to denote the vector-valued evaluations of the latent regression function on $\mathbf{x}$, and define $p(f(\mathbf{x})), q(f(\mathbf{x}))$ accordingly.

3.1 A Motivating Setup

To motivate and justify our method, we first consider a simplified setting, and assume for the moment that $\mathcal{X}$ is a finite set, and the gradient for the (log) *function-space prior*, $\nabla_{f(\mathbf{x})} \log p(f(\mathbf{x}))$ is available for any $\mathbf{x} \subseteq \mathcal{X}$. These assumptions will be removed in the subsequent section.

In this case, we can treat the function values $f(\mathcal{X})$ as the parameter to be inferred, and apply POVI in this space. This algorithm is sound when $f(\mathcal{X})$ is finite-dimensional, as theories on the consistency of POVI (e.g. Liu, 2017) directly apply. However, it becomes inefficient as the size of $\mathcal{X}$ increases. We address this issue by approximating function-space particles in the parametric model, and replacing the update rule in POVI with a version based on mini-batches.

3.1.1 Parametric Approximation to Particle Functions

Clearly, instead of maintaining $f(x)$ for all $x \in \mathcal{X}$, we can store the information of f in the neural network parameters θ, which we denote as $f(\cdot;\theta)$. In other words, we will maintain n weight-space

[1] People familiar with SVGD could argue that this issue can be mitigated by choosing a reasonable kernel function in $\mathbf{v}$, e.g. a kernel defined on $f(\cdot;\theta)$. We remark that similar kernels do not work in practice. We provide detailed experiments and discussion in Appendix B.

particles $\theta^1, \ldots, \theta^n$. The initial value of the function-space particles will be $f(\cdot;\theta^1), \ldots, f(\cdot;\theta^n)$, and each step of the function-space POVI update will be approximated as follows:

1. For each $i \in \{1, \ldots, n\}$, compute the particle optimization update $\mathbf{v}(f_\ell^i)$, where $f_\ell^i(\cdot)$ is approximated by $f(\cdot;\theta_\ell^i)$;
2. For each $i \in \{1, \ldots, n\}$, optimize θ_ℓ^i so that $f(\mathcal{X};\theta_\ell^i)$ is close to $f(\mathcal{X};\theta_\ell^i) + \epsilon_\ell \mathbf{v}(f_\ell^i)(\mathcal{X})$. We approximate this with a single step of SGD, i.e. to set

$$\theta_{\ell+1}^i \leftarrow \theta_\ell^i - \epsilon_\ell \left(\frac{\partial f(\mathcal{X};\theta_\ell^i)}{\partial \theta_\ell^i} \right)^\top \mathbf{v}(f_\ell^i)(\mathcal{X}). \tag{2}$$

3.1.2 Batched Version of the Particle Update

While we no longer need to keep track of $f_\ell^i(\mathcal{X})$ for all $\mathcal{X}$, we still need to iterate over $\mathcal{X}$ when calculating $\mathbf{v}$ and the update rule (2). We solve this issue by turning the procedure into a batched version, i.e., in each iteration, we draw a mini-batch with B elements, $\mathbf{x} \sim \mu$, for an arbitrary distribution μ supported on $\mathcal{X}^B$; we then replace the update rule (2) with evaluations on $\mathbf{x}$, i.e.

$$\theta_{\ell+1}^i \leftarrow \theta_\ell^i - \epsilon_\ell \left(\frac{\partial f(\mathbf{x};\theta_\ell^i)}{\partial \theta_\ell^i} \right)^\top \mathbf{v}(f_\ell^i)(\mathbf{x}), \tag{3}$$

in which f_ℓ^i and subsequently $\mathbf{v}(f_\ell^i)$ are *defined* on $\mathbf{x}$.

Theoretical Justification To understand this approach, note that for any fixed, finite $\mathbf{x} \subset \mathcal{X}$, the update rule (3) approximates a gradient flow which minimizes $\mathcal{E}[q(f(\mathbf{x}))]$, whose unique minimizer is $p(f(\mathbf{x})|\mathbf{X}, \mathbf{Y})$. Denote this gradient flow as $\partial_t q = -\nabla(\mathbf{v}_{\mathbf{x},t} \cdot q_t)$. Then the expectation of (3) approximates the "averaged gradient flow", $\partial_t q = -\nabla \mathbb{E}_{\mathbf{x}\sim\mu}(\mathbf{v}_{\mathbf{x},t} \cdot q_t)$, which minimizes the averaged energy $\mathbb{E}_{\mathbf{x}\sim\mu}\mathcal{E}[q(f(\mathbf{x}))]$ (Ambrosio et al., 2008). If the posterior process can be uniquely determined by all B-dimensional marginals $\{p(f(\mathbf{x})|\mathbf{X}, \mathbf{Y}) : \mathbf{x} \in \mathcal{X}^B\}$, it will clearly become the unique minimizer of the averaged energy, and simulation of the averaged gradient flow yields the true posterior. As a concrete example, if the posterior is a Gaussian process, it suffices to set $B \geq 2$.

Computation of (3) The computation of $\mathbf{v}(f_\ell^i)$ in the update rule (3) is mostly straightforward, as it is a single step of POVI update in B-dimensional space. The only exception that requires elaboration is the likelihood term, $p(\mathbf{Y}|\mathbf{X}, f_\ell^i(\mathbf{x}))$. To approximate this term, we note that $\mathbf{Y}$ and $f_\ell^i(\mathbf{x})$ is conditionally independent given the training set $\mathbf{X}$. Therefore, if $\mathbf{x}$ consists of a mini-batch of $\mathbf{X}$, denoted as $\mathbf{x}_b$, we can approximate $\log p(\mathbf{Y}|\mathbf{X}, f_\ell^i(\mathbf{x}))$ with (a scaled version of) $\log p(\mathbf{y}_b|\mathbf{x}_b, f_\ell^i(\mathbf{x}_b))$, where $\mathbf{y}_b$ is the set of outputs corresponding to $\mathbf{x}_b$. For convenience, we choose μ in such a way that samples from μ always consist of a mini-batch from $\mathbf{X}$. A natural choice of μ is thus the union of a mini-batch from $\mathbf{X}$, and samples from the kernel density estimation of $\mathbf{X}$.

Summing up, we have obtained an approximation to the function-space POVI procedure, summarized in Algorithm 1. While the approximation is a bit heuristic, due to the use of single step SGD for function-space POVI update, it can be implemented easily, and as we will show in Section 5, performs well in practice.

3.2 General Scenarios

We now remove the assumptions made in Section 3.1, to make our setting more practical in real world applications.

Infinite Set $\mathcal{X}$ While we assumed $\mathcal{X}$ is a finite set previously to make Section 3.1.1 is more easily understood, our algorithm works no matter $\mathcal{X}$ is finite or not: as our algorithm works with mini-batches, when $\mathcal{X}$ is infinite, we can also sample $\mathbf{x}$ from $\mathcal{X}$ and apply the whole procedure.

Function-Space Prior Gradient The function space prior for BNN is implicitly defined, and we do not have access to its exact gradient. While we could in principle utilize gradient estimators for implicit models (Li & Turner, 2018; Shi et al., 2018a), we opt to use a more scalable work-around in implementation, which is to approximate the prior measure with a Gaussian process (GP).

Algorithm 1 Function Space POVI for Bayesian Neural Network

Input: Function-space prior $p(f(\mathbf{x}))$ for finite $\mathbf{x} \subset \mathbf{X}$ (or an approximation), training set $(\mathbf{X}, \mathbf{Y})$, and a set of initial particles $\{\theta_0^i\}_{i=1}^n$.
Output: A set of particles $\{\theta^i\}_{i=1}^n$, such that $f(\cdot;\theta^i)$ approximates the target distribution.
for iteration ℓ **do**
Sample a mini-batch $\mathbf{x}, \mathbf{y}$ from the training set, and $\tilde{x}_{1...B} \sim^{i.i.d.} \mu$.
For each $i \in \{1, \ldots, n\}$, calculate the function space POVI update $\mathbf{v}(f_\ell^i)$, where

$$f_\ell^i := (f(\mathbf{x};\theta_\ell^i), f(\tilde{x}_{1,\ldots,B};\theta_\ell^i)),$$

and $\mathbf{v}$ is selected from Table 1.
For each $i \in \{1...K\}$, set

$$\theta_{\ell+1}^i \leftarrow \theta_\ell^i - \epsilon_\ell \left(\frac{\partial f_\ell^i}{\partial \theta_\ell^i}\right)^\top \mathbf{v}(f_\ell^i).$$

end for

More specifically, given input $\mathbf{x}$, we draw samples $\tilde{\theta}^1, \ldots, \tilde{\theta}^k$ from $p(\theta)$ and construct a multivariate normal distribution that matches the first two moments of $\tilde{p}(f(\mathbf{x})) = \frac{1}{k}\sum_{j=1}^k \delta(f(\mathbf{x}) - f(\mathbf{x};\tilde{\theta}^j))$.

A small batch size is needed to reduce the sample size k. While our procedure works for fairly small B (e.g. $B \geq 2$ for a GP posterior), we choose to use separate batches of samples to estimate the gradient of the log prior and log likelihood. In this way, a much larger batch size could be used for the log likelihood estimate.

We expect this approximation to be accurate for BNNs, because under assumptions like Gaussian prior on weights, as each layer becomes infinitely wide, the prior measure determined by BNNs will converge to a GP with a composite kernel (de G. Matthews et al., 2018; Garriga-Alonso et al., 2018).

4 Related Work

It is known that treating Bayesian regression as inference in the Euclidean space of model parameters is inefficient, and introducing alternative metrics in weight space improves the performance. A classical choice is the Fisher information metric, which has been utilized to improve Markov Chain Monte Carlo methods (Girolami & Calderhead, 2011), variational inference (Zhang et al., 2018a) and gradient descent (Amari, 1997). While such methods explore locally non-identifiable parameter regions more rapidly than their weight-space counterparts (Amari, 2016), they still suffer from global non-identifiability[2], which frequently occurs in models like Bayesian neural networks. Our work takes one step further, and deal with local and global non-identifiability simultaneously.

Closely related to our work is the variational implicit process (VIP; Ma et al., 2018), which shares the idea of function space inference. The main limitation of VIP is the lack of flexibility in inference: in the inference phase, the algorithm draws S prior functions from $p(f)$, and fits a Bayesian linear regression model using these functions as features. As S is limited by the computational budget, such an approximation family will have problem scaling to more complex models. As will be shown in Appendix A.2.3, our method compares favorably to VIP on BNN inference; and we provide evaluations on more complex network architectures like ResNet.

Approximate inference for BNN is a rich field. Under the VI framework, apart from the implicit VI methods mentioned in Section 1, Louizos & Welling (2017) proposed a hierarchical variational model, which approximates $p(\theta|\mathbf{x},\mathbf{y})$ with $q(\theta) = \int q(\theta|\mathbf{z})q(\mathbf{z})d\mathbf{z}$, where $\mathbf{z}$ represents layer-wise multiplicative noise, parameterized by normalizing flows. While this approach improves upon plain single-level variational models, its flexibility is limited by a oversimplified choice of $q(\theta|\mathbf{z})$. Such a trade-off between approximation quality and computational efficiency is inevitable for weight-space

[2]Local non-identifiability means for any neighborhood U of parameter θ, there exists $\theta' \in U$ s.t. θ' and θ represent the same model, i.e., they correspond to the same likelihood function. Global non-identifiability means there exists such θ', but not necessarily in all neighborhoods.

VI procedures. Another line of work use stochastic gradient Markov Chain Monte Carlo (SG-MCMC) for approximate inference (Li et al., 2016; Chen et al., 2014). While SG-MCMC converges to the true posterior asymptotically, within finite time it produces correlated samples, and has been shown to be less particle-efficient than the deterministic POVI procedures (Liu & Wang, 2016; Chen et al., 2018). Finally, there are other computationally efficient approaches to uncertainty estimation, e.g. Monte-Carlo dropout (Gal & Ghahramani, 2016), batch normalization (Hron et al., 2018), and efficient implementations of factorized Gaussian approximation (e.g., Blundell et al., 2015; Zhang et al., 2018a; Khan et al., 2018).

5 Evaluation

In this section, we evaluate our method on a variety of tasks. First, we present a qualitative evaluation on a synthetic regression dataset. We then evaluate the predictive performance on the standard regression and classification datasets. Finally, we assess the uncertainty quality of our proposed method on two tasks: defense against adversarial attacks, and contextual bandits.

We compare with strong baselines. For our method, we only present results implemented with SVGD for brevity (abbreviated as "f-SVGD"); results using other POVI methods are similar, and can be found in Appendix A.1 and A.2.2. Unless otherwise stated, baseline results are directly taken from the original papers.

5.1 Synthetic Data and the Over-parameterization Problem

To evaluate the approximation quality of our method qualitatively, and to demonstrate the curse-of-dimensionality problem encountered by weight space POVI methods, we first experiment on a simulated dataset. We follow the simulation setup in Sun et al. (2017): for input, we randomly generate 12 data points from Uniform$(0, 0.6)$ and 8 from Uniform$(0.8, 1)$. The output y_n for input x_n is modeled as $y_n = x_n + \epsilon_n + \sin(4(x_n + \epsilon_n)) + \sin(13(x_n + \epsilon_n))$, where $\epsilon_n \sim \mathcal{N}(0, 0.0009)$. The model is a feed-forward network with 2 hidden layers and ReLU activation; each hidden layer has 50 units. We use 50 particles for weight space SVGD and our method, and use Hamiltonian Monte Carlo (HMC) to approximate the ground truth posterior. We plot 95% credible intervals for prediction and mean estimate, representing epistemic and aleatoric[3] uncertainties respectively.

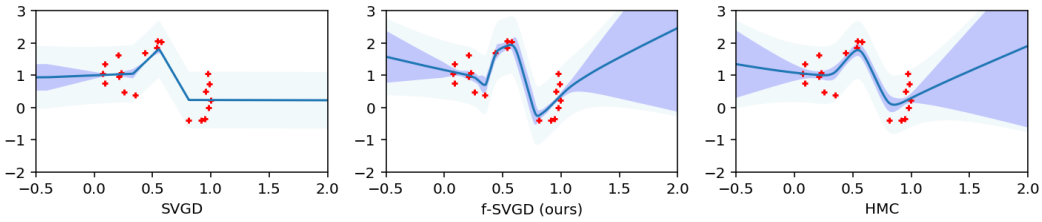

Figure 1: Approximate posterior obtained by different methods. Dots indicate observations, solid line indicates predicted mean, light shaded area corresponds to the predictive credible interval, and dark shaded area corresponds to the credible interval for mean estimate.

The result is shown in Figure 1. We can see our method provides a reasonable approximation for epistemic uncertainty, roughly consistent with HMC; on the other hand, weight-space POVI methods severely underestimate uncertainty. Furthermore, we found that such pathology exists in all weight-space POVI methods, and amplifies as model complexity increases; eventually, all weight-space methods yield degenerated posteriors concentrating on a single function. We thus conjecture it is caused by the over-parameterization problem in weight space. See Appendix A.1 for related experiments.

[3] predictive uncertainty due to the noise in the data generating process.

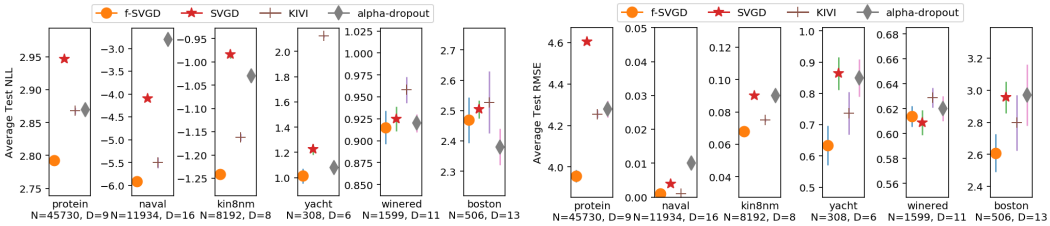

Figure 2: Average test RMSE and predictive negative log-likelihood, on UCI regression datasets. Smaller (lower) is better. Best viewed in color.

5.2 Predictive Performance

Following previous work on Bayesian neural networks (e.g. Hernández-Lobato & Adams, 2015), we evaluate the predictive performance of our method on two sets of real-world datasets: a number of UCI datasets for real-valued regression, and the MNIST dataset for classification.

5.2.1 UCI Regression Dataset

On the UCI datasets, our experiment setup is close to Hernández-Lobato & Adams (2015). The model is a single-layer neural network with ReLU activation and 50 hidden units, except for a larger dataset, Protein, in which we use 100 units. The only difference to Hernández-Lobato & Adams (2015) is that we impose an inverse-Gamma prior on the observation noise, which is also used in e.g. Shi et al. (2018b) and Liu & Wang (2016). Detailed experiment setup are included in Appendix A.2.1, and full data for our method in Appendix A.2.2.

We compare with the original weight-space Stein variational gradient descent (SVGD), and two strong baselines in BNN inference: kernel implicit variational inference (KIVI, Shi et al., 2018b), and variational dropout with α-divergences (Li & Gal, 2017).The results are summarized in Figure 2. We can see that our method has superior performance in almost all datasets.

In addition, we compare with another two state-of-the-art methods for BNN inference: multiplicative normalizing flows and Monte-Carlo batch normalization. As the experiment setup is slightly different following (Azizpour et al., 2018), we report the results in Appendix A.2.3. In most cases, our method also compares favorably to these baselines.

5.2.2 MNIST Classification Dataset

Following previous work such as Blundell et al. (2015), we report results on the MNIST handwriting digit dataset.We use a feed-forward network with two hidden layers, 400 units in each layer, and ReLU activation, and place a standard normal prior on the network weights. We choose this setting so the results are comparable with previous work.

We compare our results with vanilla SGD, Bayes-by-Backprop (Blundell et al. (2015)), and KIVI. For our method, we use a mini-batch size of 100, learning rate of 2×10^{-4} and train for 1,000 epochs. We hold out the last 10,000 examples in training set for model selection. The results are summarized in Table 2. We can see that our method outperform all baselines.

Table 2: Test error on the MNIST dataset. **Boldface** indicates the best result.

Method	BBB (Gaussian Prior)	BBB (Scale Mixture Prior)	KIVI	f-SVGD
Test Error	1.82%	1.36%	1.29%	**1.21%**

5.3 Robustness against Adversarial Examples

It is hypothesized that Bayesian models are more robust against adversarial examples due to their ability to handle epistemic uncertainty (Rawat et al., 2017; Smith & Gal, 2018). This hypothesis is supported by Li & Gal (2017), in a relatively easier setup with feed-forward networks on MNIST; to our knowledge, no results are reported using more flexible approximate inference techniques. In this section, we evaluate the robustness of our method on a setting compatible to previous work, as

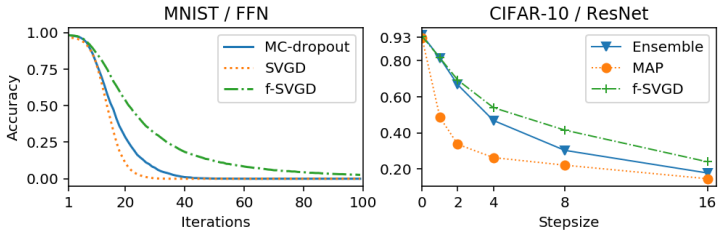

Figure 3: Accuracy on adversarial examples.

well as a more realistic setting with ResNet-32 on the CIFAR-10 dataset. We briefly introduce the experiment setup here; detailed settings are included in Appendix A.3.

On the MNIST dataset, we follow the setup in Li & Gal (2017), and experiment with a feed-forward network. We use the iterative fast gradient sign method (I-FGSM) to construct targeted white-box attack samples. In each iteration, we limit the ℓ^∞ norm of the perturbation to 0.01 (pixel values are normalized to the range of $[0, 1]$). We compare our method with vanilla SVGD, and MC Dropout.

On the CIFAR-10 dataset, we use the ResNet-32 architecture (He et al., 2016a). As dropout requires modification of the model architecture, we only compare with the single MAP estimate, and an ensemble model. We use 8 particles for our method and the ensemble baseline. We use the FGSM method to construct white-box untargeted attack samples.

Results for both experiments are presented in Figure 3. We can see our method improves robustness significantly, both when compared to previous approximate BNN models, and baselines in the more realistic setting.

5.4 Improved Exploration in Contextual Bandit

Finally, we evaluate the approximation quality of our method on several contextual bandit problems, Contextual bandit is a standard reinforcement learning problem. It is an arguably harder task than supervised learning for BNN approximation methods, as it requires the agent to balance between exploitation and exploration, and decisions based on poorly estimated uncertainty will lead to catastrophic performance through a feedback loop (Riquelme et al., 2018). Problem background and experiment details are presented in Appendix A.4.

We consider the Thompson sampling algorithm with Bayesian neural networks. We use a feed-forward network with 2 hidden layers and 100 ReLU units in each layer. Baselines include other approximate inference methods including Bayes-by-Backprop and vanilla SVGD, as well as other uncertainty estimation procedures including Gaussian process and frequentist bootstrap. We use the mushroom and wheel bandits from Riquelme et al. (2018).

The cumulative regret is summarized in Table 3. We can see that our method provides competitive performance compared to the baselines, and outperforming all baselines by a large margin in the wheel bandit, in which high-quality uncertainty estimate is especially needed.

Table 3: Cumulative regret in different bandits. Results are averaged over 10 trials.

	BBB	GP	Bootstrap	f-SVGD
Mushroom	19.15 ± 5.98	16.75 ± 1.63	$\mathbf{2.71 \pm 0.22}$	4.39 ± 0.39
Wheel	55.77 ± 8.29	60.80 ± 4.40	42.16 ± 7.80	$\mathbf{7.54 \pm 0.41}$

6 Conclusion

We presented a flexible approximate inference method for Bayesian regression models, building upon particle-optimization based variational inference procedures. The newly proposed method performs POVI on function spaces, which is scalable and easy to implement and overcomes the degeneracy problem in direct applications of POVI procedures. Extensive experiments demonstrated the effectiveness of our proposal.

REFERENCES

Shun-ichi Amari. Neural learning in structured parameter spaces-natural riemannian gradient. In *Advances in neural information processing systems*, pp. 127–133, 1997.

Shun-ichi Amari. *Information geometry and its applications*. Springer, 2016.

Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient flows: in metric spaces and in the space of probability measures*. Springer Science & Business Media, 2008.

Hossein Azizpour, Mattias Teye, and Kevin Smith. Bayesian uncertainty estimation for batch normalized deep networks. In *International Conference on Machine Learning (ICML)*, 2018.

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

Olivier Chapelle and Lihong Li. An empirical evaluation of thompson sampling. In *Advances in neural information processing systems*, pp. 2249–2257, 2011.

Changyou Chen, Ruiyi Zhang, Wenlin Wang, Bai Li, and Liqun Chen. A unified particle-optimization framework for scalable bayesian sampling. *arXiv preprint arXiv:1805.11659*, 2018.

Tianqi Chen, Emily Fox, and Carlos Guestrin. Stochastic gradient hamiltonian monte carlo. In *International Conference on Machine Learning*, pp. 1683–1691, 2014.

Alexander G. de G. Matthews, Jiri Hron, Mark Rowland, Richard E. Turner, and Zoubin Ghahramani. Gaussian process behaviour in wide deep neural networks. In *International Conference on Learning Representations*, 2018. URL `https://openreview.net/forum?id=H1-nGgWC-`.

Di Feng, Lars Rosenbaum, and Klaus Dietmayer. Towards safe autonomous driving: Capture uncertainty in the deep neural network for lidar 3d vehicle detection. *arXiv preprint arXiv:1804.05132*, 2018.

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

Adrià Garriga-Alonso, Laurence Aitchison, and Carl Edward Rasmussen. Deep convolutional networks as shallow gaussian processes. *arXiv preprint arXiv:1808.05587*, 2018.

Soumya Ghosh, Jiayu Yao, and Finale Doshi-Velez. Structured variational learning of Bayesian neural networks with horseshoe priors. In Jennifer Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 1744–1753, Stockholmsmssan, Stockholm Sweden, 10–15 Jul 2018. PMLR. URL `http://proceedings.mlr.press/v80/ghosh18a.html`.

Mark Girolami and Ben Calderhead. Riemann manifold langevin and hamiltonian monte carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2): 123–214, 2011.

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

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European Conference on Computer Vision*, pp. 630–645. Springer, 2016b.

José Miguel Hernández-Lobato and Ryan Adams. Probabilistic backpropagation for scalable learning of bayesian neural networks. In *International Conference on Machine Learning*, pp. 1861–1869, 2015.

Jiri Hron, Alex Matthews, and Zoubin Ghahramani. Variational Bayesian dropout: pitfalls and fixes. In Jennifer Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 2019–2028, Stockholmsmssan, Stockholm Sweden, 10–15 Jul 2018. PMLR. URL http://proceedings.mlr.press/v80/hron18a.html.

Mohammad Emtiyaz Khan, Didrik Nielsen, Voot Tangkaratt, Wu Lin, Yarin Gal, and Akash Srivastava. Fast and scalable bayesian deep learning by weight-perturbation in adam. *arXiv preprint arXiv:1806.04854*, 2018.

Chunyuan Li, Changyou Chen, David E Carlson, and Lawrence Carin. Preconditioned stochastic gradient langevin dynamics for deep neural networks. In *AAAI*, volume 2, pp. 4, 2016.

Yingzhen Li and Yarin Gal. Dropout inference in bayesian neural networks with alpha-divergences. *arXiv preprint arXiv:1703.02914*, 2017.

Yingzhen Li and Richard E. Turner. Gradient estimators for implicit models. In *International Conference on Learning Representations*, 2018. URL https://openreview.net/forum?id=SJi9WOeRb.

Chang Liu, Jingwei Zhuo, Pengyu Cheng, Ruiyi Zhang, Jun Zhu, and Lawrence Carin. Accelerated first-order methods on the wasserstein space for bayesian inference. *arXiv preprint arXiv:1807.01750*, 2018.

Qiang Liu. Stein variational gradient descent as gradient flow. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 30*, pp. 3115–3123. Curran Associates, Inc., 2017. URL http://papers.nips.cc/paper/6904-stein-variational-gradient-descent-as-gradient-flow.pdf.

Qiang Liu and Dilin Wang. Stein variational gradient descent: A general purpose bayesian inference algorithm. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 29*, pp. 2378–2386. Curran Associates, Inc., 2016. URL http://papers.nips.cc/paper/6338-stein-variational-gradient-descent-a-general-purpose-bayesian-inference-algor pdf.

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

Chao Ma, Yingzhen Li, and José Miguel Hernández-Lobato. Variational implicit processes. *arXiv preprint arXiv:1806.02390*, 2018.

A. Rawat, M. Wistuba, and M.-I. Nicolae. Adversarial Phenomenon in the Eyes of Bayesian Deep Learning. *ArXiv e-prints*, November 2017.

Carlos Riquelme, George Tucker, and Jasper Snoek. Deep bayesian bandits showdown: An empirical comparison of bayesian deep networks for thompson sampling. In *International Conference on Learning Representations*, 2018. URL https://openreview.net/forum?id=SyYe6k-CW.

Jeff Schlimmer. Mushroom records drawn from the audubon society field guide to north american mushrooms. *GH Lincoff (Pres), New York*, 1981.

Jiaxin Shi, Shengyang Sun, and Jun Zhu. A spectral approach to gradient estimation for implicit distributions. In Jennifer Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 4644–4653, Stockholmsmssan, Stockholm Sweden, 10–15 Jul 2018a. PMLR. URL http://proceedings.mlr.press/v80/shi18a.html.

Jiaxin Shi, Shengyang Sun, and Jun Zhu. Kernel implicit variational inference. In *International Conference on Learning Representations*, 2018b. URL https://openreview.net/forum?id=r1l4eQW0Z.

Lewis Smith and Yarin Gal. Understanding measures of uncertainty for adversarial example detection. *arXiv preprint arXiv:1803.08533*, 2018.

Shengyang Sun, Changyou Chen, and Lawrence Carin. Learning structured weight uncertainty in bayesian neural networks. In *Artificial Intelligence and Statistics*, pp. 1283–1292, 2017.

William R. Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.

Brian Trippe and Richard Turner. Overpruning in variational bayesian neural networks. *arXiv preprint arXiv:1801.06230*, 2018.

Dilin Wang, Zhe Zeng, and Qiang Liu. Stein variational message passing for continuous graphical models. In Jennifer Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 5219–5227, Stockholmsmssan, Stockholm Sweden, 10–15 Jul 2018. PMLR. URL http://proceedings.mlr.press/v80/wang18l.html.

Guodong Zhang, Shengyang Sun, David Duvenaud, and Roger Grosse. Noisy natural gradient as variational inference. In Jennifer Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 5852–5861, Stockholmsmssan, Stockholm Sweden, 10–15 Jul 2018a. PMLR. URL http://proceedings.mlr.press/v80/zhang18l.html.

Ruiyi Zhang, Chunyuan Li, Changyou Chen, and Lawrence Carin. Learning structural weight uncertainty for sequential decision-making. In *International Conference on Artificial Intelligence and Statistics*, pp. 1137–1146, 2018b.

Jingwei Zhuo, Chang Liu, Jiaxin Shi, Jun Zhu, Ning Chen, and Bo Zhang. Message passing stein variational gradient descent. In Jennifer Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 6018–6027, Stockholmsmssan, Stockholm Sweden, 10–15 Jul 2018. PMLR. URL http://proceedings.mlr.press/v80/zhuo18a.html.

A EXPERIMENT DETAILS AND ADDITIONAL RESULTS

A.1 SYNTHETIC DATA

Comparison with Other POVI Methods We present posterior approximations obtained by weight-space and function-space versions of other POVI methods. The simulation setup is the same as in Section 5.1. As shown in Figure 4, function-space methods provide improvement in all cases, avoid the degenerate behavior of weight-space methods.

Experiments with Increasing Model Complexity To obtain a better understanding of the degenerate behavior of weight-space POVI methods, we repeat the experiment with increasingly complex models. Specifically, we repeat the experiment while varying the number of hidden units in each layer from 5 to 100. Other settings are the same as Section 5.1. The posteriors are shown in Figure 5. We can see that weight space methods provide accurate posterior estimates when the number of weights is small, and degenerate gradually as model complexity increases, eventually all particles collapse into a single function. On the contrary, function space methods produce stable approximations all the time.

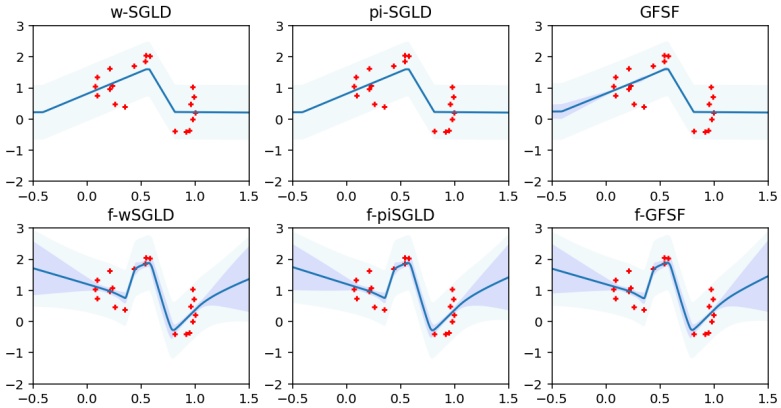

Figure 4: Posterior approximations obtained by weight space (up) and function space (down) variants of other POVI methods.

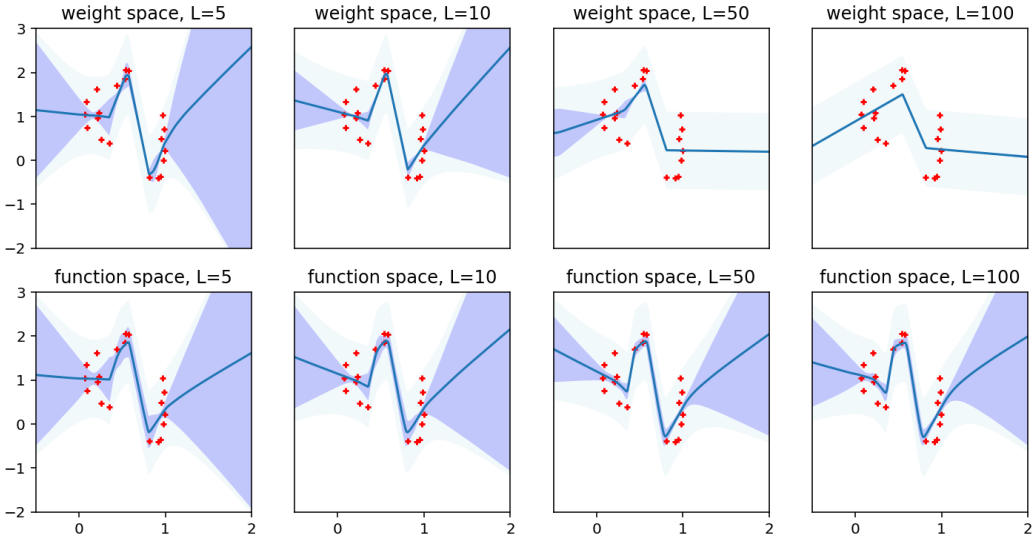

Figure 5: Posterior approximations on increasingly complex models. L denotes the number of hidden units in each layer. We can see a clear trend of degeneration for weight-space method.

A.2 UCI DATASETS

A.2.1 EXPERIMENT DETAILS IN 5.2.1

For our method in all datasets, we use the AdaM optimizer with a learning rate of 0.004. For datasets with fewer than 1000 samples, we use a batch size of 100 and train for 500 epochs. For the larger datasets, we set the batch size to 1000, and train for 1000 epochs. We use a 90-10 random train-test split, repeated for 20 times, except for Protein in which we use 5 replicas. For our method and weight-space POVI methods, we use 20 particles. We use the RBF kernel, with the bandwidth chosen by the median trick (Liu & Wang, 2016).

For our method, we approximate the function-space prior with GP, and separate the mini-batch used for prior gradient and other parts in $\mathbf{v}$, as discussed in the main text. We construct the multivariate normal prior approximation with 40 draws and a batch size of 4.

A.2.2 FULL RESULTS WITH DIFFERENT POVI PROCEDURES

In this section, we provide full results on UCI datasets, using weight-space and function-space versions of different POVI procedures. The experiment setup is the same as Section 5.2.1. The predictive RMSE and NLL are presented in Table 4 and 5, respectively. We can see that for all POVI methods, function space variants provide substantial improvement over their weight-space counterparts.

Table 4: Average test RMSE on UCI datasets. Bold indicates statistically significant best results ($p < 0.05$ with t-test).

Dataset	Weight Space			Function Space (Ours)		
	SVGD	w-SGLD	pi-SGLD	SVGD	w-SGLD	pi-SGLD
Boston	2.96 ± 0.10	$\mathbf{2.84 \pm 0.15}$	$\mathbf{2.84 \pm 0.15}$	$\mathbf{2.61 \pm 0.12}$	$\mathbf{2.62 \pm 0.12}$	$\mathbf{2.62 \pm 0.12}$
Concrete	5.32 ± 0.10	5.51 ± 0.10	5.49 ± 0.10	$\mathbf{4.73 \pm 0.13}$	$\mathbf{4.78 \pm 0.14}$	$\mathbf{4.77 \pm 0.13}$
Kin8nm	0.09 ± 0.00	0.07 ± 0.00	0.07 ± 0.00	$\mathbf{0.07 \pm 0.00}$	0.07 ± 0.00	0.07 ± 0.00
Naval	0.00 ± 0.00	$\mathbf{0.00 \pm 0.00}$	$\mathbf{0.00 \pm 0.00}$	0.00 ± 0.00	$\mathbf{0.00 \pm 0.00}$	$\mathbf{0.00 \pm 0.00}$
Power	3.94 ± 0.03	3.97 ± 0.03	3.97 ± 0.03	$\mathbf{3.80 \pm 0.03}$	$\mathbf{3.83 \pm 0.03}$	$\mathbf{3.83 \pm 0.03}$
Protein	4.61 ± 0.01	4.40 ± 0.02	4.40 ± 0.02	$\mathbf{3.96 \pm 0.03}$	$\mathbf{4.02 \pm 0.03}$	$\mathbf{4.02 \pm 0.03}$
Winered	$\mathbf{0.61 \pm 0.01}$	$\mathbf{0.63 \pm 0.01}$	$\mathbf{0.63 \pm 0.01}$	$\mathbf{0.61 \pm 0.01}$	$\mathbf{0.61 \pm 0.01}$	$\mathbf{0.61 \pm 0.01}$
Yacht	0.86 ± 0.05	0.95 ± 0.07	0.93 ± 0.07	$\mathbf{0.63 \pm 0.06}$	$\mathbf{0.63 \pm 0.06}$	$\mathbf{0.63 \pm 0.06}$

Table 5: Average test NLL on UCI datasets. Bold indicates best results.

Dataset	Weight Space			Function Space (Ours)		
	SVGD	w-SGLD	pi-SGLD	SVGD	w-SGLD	pi-SGLD
Boston	$\mathbf{2.50 \pm 0.03}$	$\mathbf{2.52 \pm 0.07}$	$\mathbf{2.52 \pm 0.07}$	$\mathbf{2.47 \pm 0.08}$	$\mathbf{2.52 \pm 0.09}$	$\mathbf{2.51 \pm 0.08}$
Concrete	3.08 ± 0.02	3.15 ± 0.03	3.15 ± 0.03	$\mathbf{2.96 \pm 0.04}$	$\mathbf{3.02 \pm 0.04}$	$\mathbf{3.01 \pm 0.04}$
Kin8nm	-0.98 ± 0.01	-1.20 ± 0.01	-1.20 ± 0.01	$\mathbf{-1.24 \pm 0.00}$	$\mathbf{-1.25 \pm 0.00}$	$\mathbf{-1.25 \pm 0.00}$
Naval	-4.09 ± 0.01	-6.33 ± 0.02	$\mathbf{-6.39 \pm 0.02}$	-5.92 ± 0.03	$\mathbf{-6.43 \pm 0.04}$	$\mathbf{-6.45 \pm 0.05}$
Power	2.79 ± 0.01	2.80 ± 0.01	2.80 ± 0.01	$\mathbf{2.76 \pm 0.01}$	$\mathbf{2.76 \pm 0.01}$	$\mathbf{2.76 \pm 0.01}$
Protein	2.95 ± 0.00	2.90 ± 0.00	2.90 ± 0.00	$\mathbf{2.79 \pm 0.01}$	$\mathbf{2.81 \pm 0.01}$	$\mathbf{2.81 \pm 0.01}$
Winered	$\mathbf{0.93 \pm 0.01}$	0.97 ± 0.01	0.97 ± 0.01	$\mathbf{0.92 \pm 0.02}$	$\mathbf{0.93 \pm 0.02}$	$\mathbf{0.94 \pm 0.02}$
Yacht	1.23 ± 0.04	1.39 ± 0.07	1.37 ± 0.07	$\mathbf{1.01 \pm 0.06}$	$\mathbf{0.99 \pm 0.06}$	$\mathbf{1.00 \pm 0.06}$

A.2.3 COMPARISON WITH OTHER METHODS

In this section, we provide comparison to a few other strong baselines in BNN inference.

MNF and MCBN We first present comparisons with multiplicative normalizing flow (MNF; Louizos & Welling, 2017) and Monte-Carlo batch normalization (MCBN; Azizpour et al., 2018). We make the experiment setups consistent with Azizpour et al. (2018), and cite the baseline results from their work. The model is a BNN with 2 hidden layer, each with 50 units for the smaller datasets, and 100 for the protein dataset. We use 10% data for test, and hold out an additional fraction of 10% data to determine the optimal number of epochs. The hyperparameters and optimization scheme for our model is the same as in Section 5.2.1. The results are presented in Table 6. We can see that in most cases, our algorithm performs better than the baselines.

VIP We present comparison to the variational implicit process (VIP; Ma et al., 2018). We follow the experiment setup in their work, and cite results from it. The model is a BNN with 2 hidden layer, each with 10 units. For our method, we follow the setup in Section 5.2.1. We note that this is not entirely a fair comparison, as the method in Ma et al. (2018) also enables model hyper-parameter learning, while for our method we keep the hyper-parameters fixed. However, as shown in Table 7, our method still compares even or favorably to them, outperforming them in NLL in 5 out of 9 datasets. This supports our hypothesis in Section 4 that the linear combination posterior used in Ma et al. (2018) limits their flexibility.

Table 6: Average test RMSE and NLL following the setup of Azizpour et al. (2018). Bold indicates best results.

Dataset	Test RMSE			Test NLL		
	MCBN	MNF	f-SVGD	MCBN	MNF	f-SVGD
Boston	$\mathbf{2.75 \pm 0.05}$	2.98 ± 0.06	2.95 ± 0.08	$\mathbf{2.38 \pm 0.02}$	2.51 ± 0.06	2.57 ± 0.07
Concrete	$\mathbf{4.78 \pm 0.09}$	6.57 ± 0.04	$\mathbf{5.03 \pm 0.10}$	3.45 ± 0.11	3.35 ± 0.04	$\mathbf{3.16 \pm 0.05}$
Kin8nm	0.07 ± 0.00	0.09 ± 0.00	$\mathbf{0.06 \pm 0.00}$	-1.21 ± 0.01	-1.04 ± 0.00	$\mathbf{-1.33 \pm 0.00}$
Power	3.74 ± 0.01	4.19 ± 0.01	$\mathbf{3.50 \pm 0.01}$	2.75 ± 0.00	2.86 ± 0.01	$\mathbf{2.67 \pm 0.00}$
Protein	3.66 ± 0.01	4.10 ± 0.01	$\mathbf{3.36 \pm 0.03}$	2.73 ± 0.00	2.83 ± 0.01	$\mathbf{2.56 \pm 0.01}$
Winered	0.62 ± 0.00	$\mathbf{0.61 \pm 0.00}$	0.63 ± 0.01	0.95 ± 0.01	$\mathbf{0.93 \pm 0.00}$	0.98 ± 0.01
Yacht	1.23 ± 0.05	2.13 ± 0.05	$\mathbf{0.85 \pm 0.06}$	1.39 ± 0.03	1.96 ± 0.05	$\mathbf{1.03 \pm 0.03}$

Table 7: Average test RMSE and NLL on UCI datasets, following the setup in Ma et al. (2018). Bold indicates best results.

	NLL		RMSE	
	VIP-BNN	f-SVGD	VIP-BNN	f-SVGD
Boston	2.45 ± 0.04	$\mathbf{2.33 \pm 0.05}$	2.88 ± 0.14	$\mathbf{2.58 \pm 0.12}$
Concrete	3.02 ± 0.02	$\mathbf{2.93 \pm 0.02}$	4.81 ± 0.13	$\mathbf{4.63 \pm 0.12}$
Energy	$\mathbf{0.60 \pm 0.03}$	0.69 ± 0.03	0.45 ± 0.01	0.43 ± 0.01
Kin8nm	-1.12 ± 0.01	-1.11 ± 0.01	0.07 ± 0.00	0.07 ± 0.00
Power	2.92 ± 0.00	$\mathbf{2.76 \pm 0.00}$	4.11 ± 0.05	$\mathbf{3.78 \pm 0.02}$
Protein	2.87 ± 0.00	$\mathbf{2.81 \pm 0.01}$	4.25 ± 0.07	$\mathbf{4.04 \pm 0.03}$
Winered	0.97 ± 0.02	$\mathbf{0.89 \pm 0.01}$	0.64 ± 0.01	$\mathbf{0.61 \pm 0.01}$
Yacht	$\mathbf{-0.02 \pm 0.07}$	1.17 ± 0.01	0.32 ± 0.06	0.67 ± 0.07
Naval	$\mathbf{-5.62 \pm 0.04}$	-5.41 ± 0.10	0.00 ± 0.00	0.00 ± 0.00

A.3 Adversarial Examples

MNIST Experiment Details We follow the setup in Li & Gal (2017). We use a feed-forward network with ReLU activation and 3 hidden layers, each with 1000 units. For both SVGD and f-SVGD, we use the AdaM optimizer with learning rate 5×10^{-4}; we use a batch size of 1000 and train for 1000 epochs. The attack method is the targeted iterated FGSM with ℓ^∞ norm constraint, i.e. for $t = 1, \ldots, T$, set

$$x_{adv}^{(t)} := x_{adv}^{(t-1)} + \epsilon \cdot \mathrm{sgn}(\nabla \log \mathbb{E}_{q(f)}[p(y = 0|x_{adv}^{(t-1)}, f)]),$$

where $\epsilon = 0.01$, and $x_{adv}^{(0)}$ denotes the clean input. We remove clean images labeled with "0".

For our method and weight-space SVGD, we use 20 particles, and scale the output logits to have a prior variance of 10. We use an additional 40 particles to generate 8-dimensional prior samples.

A few notes on applying our method to classification problems: it is important to use the *normalized logits* as f, so that the model is fully identified; also as for classification problems, $f(x)$ is already high-dimensional for a single data point x, one should down-sample $f(x)$ within each x. We choose to always include the true label in the down-sampled version, for labeled samples, and share negative labels across different x.

CIFAR Experiment Details We use the ResNet-32 architecture, defined in He et al. (2016b), and uses the same training scheme. We use 8 particles for our method and the ensemble method. For our method, we use 32 particles to generate 6-dimensional prior samples.

The ResNet architecture consists of batch normalization layers, which needs to be computed with the full batch of input. We approximate it with a down-sampled batch, so that more prior samples can be produced more efficiently. In our implementation, we down-sample the input batch to $1/4$ of its original size.

A.4 CONTEXTUAL BANDIT

Contextual bandit is a classical online learning problem. The problem setup is as follows: for each time $t = 1, 2, \cdots, N$, a context $s_t \in \mathcal{S}$ is provided to the online learner, where $\mathcal{S}$ denotes the given context set. The online learner need to choose one of the K available actions $I_t \in \{1, 2, \cdots, K\}$ based on context s_t, and get a (stochastic) reward $\ell_{I_t,t}$. The goal of the online learner is to minimize the pseudo-regret

$$\overline{R}_n^{\mathcal{S}} = \max_{g:\mathcal{S}\to\{1,2,\cdots,K\}} \mathbb{E}\left[\sum_{t=1}^{n} \ell_{g(s_t),t} - \sum_{t=1}^{n} \ell_{I_t,t}\right]. \tag{4}$$

where g denotes the mapping from context set $\mathcal{S}$ to available actions $\{1, 2, \cdots, K\}$. Pseudo-regret measures the regret of not following the best g, thus pseudo-regret is non-negative, and minimize the pseudo-regret is equal to find such the best g.

For contextual bandits with non-adversarial rewards, Thompson sampling (a.k.a. posterior sampling; Thompson, 1933) is a classical algorithm that achieves state-of-the-art performance in practice (Chapelle & Li, 2011). Denote the underlying ground-truth reward distribution of context s and action I_t as ν_{s,I_t}. In Thompson sampling, we place a prior $\mu_{s,i,0}$ on reward for context s and action i, and maintain $\mu_{s,i,t}$, the corresponding posterior distribution at time t. For each time $t = 1, 2, \ldots, N$, Thompson sampling selects action by

$$I_t \in \operatorname*{arg\,max}_{i=\{1,2,\cdots,K\}} \hat{\ell}_{i,t}, \quad \hat{\ell}_{i,t} \sim \mu_{s,i,t}.$$

The corresponding posterior is then updated with the observed reward $\ell_{I_t,t}$. The whole procedure of Thompson sampling is shown in Algorithm 2.

Algorithm 2 Thompson Sampling

Input: Prior distribution $\mu_{s,i,0}$, time horizon N
for time $t = 1, 2, \cdots, N$ **do**
  Observe context $s_t \in \mathcal{S}$
  Sample $\hat{\ell_{i,t}} \sim \mu_{s,i,t}$.
  Select $I_t \in \arg\max_{i=\{1,2,\cdots,K\}} \hat{\ell_{i,t}}$ and get $\ell_{I_t,t}$
  Update the posterior of $\mu_{s_t,I_t,t+1}$ with $\ell_{I_t,t}$.
end for

Notice that contextual bandit problems always face the *exploration-exploitation dilemma*. Exploration should be appropriate, otherwise, we can either exploit too much sub-optimal actions or explore too much meaningless actions. Thompson sampling addresses this issue by each time selecting the actions greedy with the sampled rewards, which is equal to selecting action i with the probability that i can get the highest reward under context s_t. This procedure need an accurate posterior uncertainty. Either over-estimate or under-estimate the posterior uncertainty can lead to the failure of balancing exploration and exploitation, which further lead to the failure of Thompson sampling.

Here, we focus on two benchmark contextual bandit problems, called *mushroom* and *wheel*.

Mushroom Bandit In mushroom, we use the data from mushroom dataset (Schlimmer, 1981), which contains 22 attributes per mushroom and two classes: poisonous and safe. Eating a safe mushroom provides reward +5. Eating a poisonous mushroom delivers reward +5 with probability 1/2 and reward -35 otherwise. If the agent does not eat a mushroom, then the reward is 0. We run 50000 rounds in this problem.

Wheel Bandit Wheel bandit is a synthetic problem that highlights the need for exploration. Let $\delta \in (0, 1)$ be an "exploration parameter". Context X is sampled uniformly at random in the unit circle in $\mathbb{R}^2$. There are $k = 5$ possible actions: the first action results in a constant reward $\ell_1 \sim \mathcal{N}(\mu_1, \sigma^2)$; the reward corresponding to other actions is determined by X:

- For contexts inside the blue circle in Figure 6, i.e. for X s.t. $\|X\| \leq \delta$, the other four actions all result in a suboptimal reward $\mathcal{N}(\mu_2, \sigma^2)$ for $\mu_2 < \mu_1$.

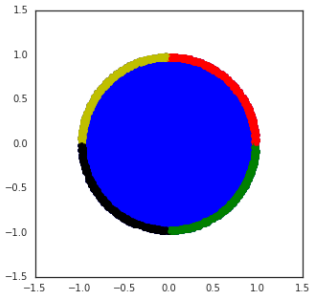

Figure 6: Visualization of the wheel bandit. Best viewed in color.

- Otherwise, one of the four contexts becomes optimal depend on the quarter X is in. The optimal action results in a reward of $\mathcal{N}(\mu_3, \sigma^2)$ for $\mu_3 \gg \mu_1$, and other actions still have the reward $\mathcal{N}(\mu_2, \sigma^2)$.

As the probability that X corresponds to a high reward is $1 - \delta^2$, the need for exploration increases as δ increases, and it is expected that algorithm with poorly calibrated uncertainty will stuck in choosing the suboptimal action a_1 in these regions. Such a hypothesis is confirmed in (Riquelme et al., 2018), making this bandit a particularly challenging problem. In our experiments, we use 50000 contexts, and set $\delta = 0.95$.

Experiment Setup The model for neural networks is a feed-forward network with 2 hidden layers, each with 100 units. The hyper-parameters for all models are tuned on the mushroom bandit, and kept the same for the wheel bandit. For the baselines, we use the hyper-parameters provided in (Riquelme et al., 2018). The main difference from (Riquelme et al., 2018) is that we use 20 replicas for bootstrap. For our method, we also use 20 particles.

B On Alternative Kernels for Weight-Space POVI

Many POVI methods use a kernel to make the gradient flow well-defined for discrete distributions. A natural question is whether we could design the kernel carefully to alleviate the problem of over-parameterization and high dimensionality. This idea is tempting: intuitively, for most POVI methods listed in Table 1, the kernel defines a *repulsive force term*, which push a particle away from its "most similar" neighbors. A better choice of the kernel makes the similarity measure more sensible. In this section, we will list candidates for such a kernel, and show that they do not work in general.

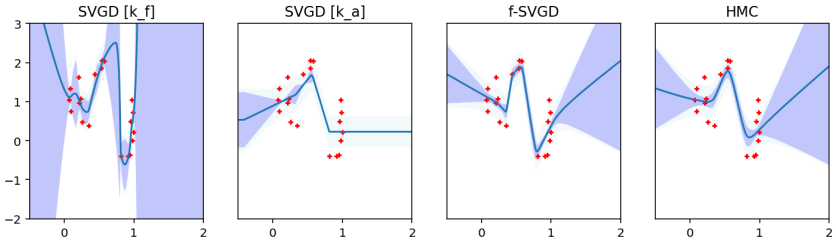

Figure 7: Posterior approximations with weight-space SVGD and alternative kernels. `k_f` corresponds to the function-value kernel, and `k_a` the activation kernel. We include the results for f-SVGD and HMC for reference.

Evaluation Setup For kernels considered in this section, we evaluate their empirical performance on the synthetic dataset and the UCI regression datasets, following the same setup as in the main text. We test on two POVI methods: SVGD and w-SGLD. As results are similar, we only report the results for SVGD for brevity. The results are presented in Figure 7 and Table 8.

Table 8: Test NLL on UCI datasets for SVGD with alternative kernels.

	Weight-Space SVGD			Function-Space SVGD
	RBF	k_a	k_f	
Boston	2.50 ± 0.03	2.50 ± 0.07	2.49 ± 0.06	2.47 ± 0.08
Yacht	1.23 ± 0.04	1.35 ± 0.06	1.20 ± 0.07	$\mathbf{1.01 \pm 0.06}$
Concrete	3.08 ± 0.02	3.12 ± 0.03	3.11 ± 0.03	$\mathbf{2.96 \pm 0.04}$

The "Function Value" Kernel Similar to our proposed method, one could define a *weight-space* kernel on function values, so it directly measures the difference of regression functions; i.e.

$$k_f(\theta^{(1)}, \theta^{(2)}) := \mathbb{E}_{\mathbf{x}\sim\mu}[k(f(\mathbf{x};\theta^{(1)}), f(\mathbf{x};\theta^{(2)}))],$$

where k is an ordinary kernel on $\mathbb{R}^B$ (e.g. the RBF kernel), $\mathbf{x} \in \mathcal{X}^B$, and μ is an arbitrary measure supported on $\mathcal{X}^B$. We first remark that k_f is not positive definite (p.d.) due to over-parameterization, and there is no guarantee that weight-space inference with k_f will converge to the true posterior in the asymptotic limit[4]. Furthermore, it does not improve over the RBF kernel empirically: predictive performance does not improve, and it drastically overestimates the epistemic uncertainty on the synthetic data.

To understand the failure of this kernel, take SVGD as an example: the update rule $\theta_{\ell+1}^{(i)} \leftarrow \theta_\ell^{(i)} - \epsilon_\ell \mathbf{v}(\theta_\ell^{(i)})$ is defined with

$$-\mathbf{v}(\theta^{(i)}) = \frac{1}{n}\sum_{j=1}^{n}\left(\underbrace{\mathbf{K}_{ij}\nabla_{\theta^{(j)}}\log p(\theta^{(j)}|\mathbf{x})}_{\mathrm{grad}_{ij}} + \underbrace{\nabla_{\theta^{(j)}}\mathbf{K}_{ij}}_{\mathrm{rf}_{ij}}\right).$$

Finite-sample SVGD is usually understood as follows (Liu & Wang, 2016): the first term follows a smoothed gradient direction, which pushes the particles towards high-probability region in posterior; the second term acts as a repulsive force, which prevents particles from collapse together. However, for inference with non positive definite kernels such as k_f, it is less clear if these terms still play the same role:

1. When there are particles corresponding to the same regression function, their gradient for log posterior will be averaged in grad. This is clearly undesirable, as these particles can be distant to each other in the weight space, so their gradient contains little learning signal for each other.
2. While for stationary kernels, the repulsive force

   $$\mathrm{rf}_{ij} = \frac{1}{n}\sum_j \nabla_{\theta^{(j)}}\mathbf{K}_{ij} = -\frac{1}{n}\sum_j \nabla_{\theta^{(i)}}\mathbf{K}_{ij}$$

   drives particle i away from other particles, this equality does not hold for k_f which is non-stationary. Furthermore, in such summation over $\nabla_{\theta^{(j)}}\mathbf{K}_{ij}$, as $\nabla_{\theta^{(j)}}\mathbf{K}_{ij} = \nabla_{f(\mathbf{x};\theta^{(j)})}\mathbf{K}_{ij}\nabla_{\theta^{(j)}}f(\mathbf{x};\theta^{(j)})$, assuming $\nabla_{\theta^{(j)}}f(\mathbf{x};\theta^{(j)})$ is of the same scale for j, $\nabla_{\theta^{(j)}}f(\mathbf{x};\theta^{(j)})$ will contribute to particle i's repulsive force most if the function-space repulsive force, $\nabla_{f(\mathbf{x};\theta^{(j)})}\mathbf{K}_{ij}$ is the largest. However, unlike in identifiable models, there is no guarantee that this condition imply $\theta^{(j)}$ is sufficiently close to $\theta^{(i)}$, and mixing their gradient could be detrimental for the learning process.

As similar terms also exist in other POVI methods, such an argument is independent to the choice of POVI methods. Also note that both parts of this argument depends on the over-parameterization property of the model, and the degeneracy (not being p.d.) of the kernel.

The Activation Kernel Another kernel that seems appropriate for BNN inference can be defined using network layer activations. Specifically, let $h(\mathbf{x};\theta)$ be the activations of all layers in the NN parameterized by θ, when fed with a batch of inputs $\mathbf{x}$. The kernel is defined as

$$k_a(\theta^{(1)}, \theta^{(2)}) := \mathbb{E}_{\mathbf{x}\sim\mu}[k(h(\mathbf{x};\theta^{(1)}), h(\mathbf{x};\theta^{(2)}))].$$

[4]E.g. Liu & Wang (2016) requires a p.d. kernel for such guarantees.

This kernel is positive definite if the batch size B is sufficiently large, but it does not work either: predictive performance is worse than weight-space SVGD with RBF kernel, and as shown in Figure 7, it also suffers from the collapsed prediction issue. Intuitively, our argument on over-parameterization in Section 2 should apply to all positive definite kernels.

In conclusion, constructing sensible kernels for BNN in weight-space POVI is non-trivial; on the other hand, our proposed algorithm provides a more natural solution, and works better in practice.