

---

# Generative Particle Variational Inference via Estimation of Functional Gradients

---

Neale Ratzlaff<sup>\*1</sup> Qinxun Bai<sup>\*2</sup> Li Fuxin<sup>1</sup> Wei Xu<sup>2</sup>

## Abstract

Recently, particle-based variational inference (ParVI) methods have gained interest because they can avoid arbitrary parametric assumptions that are common in variational inference. However, many ParVI approaches do not allow arbitrary sampling from the posterior, and the few that do allow such sampling suffer from suboptimality. This work proposes a new method for learning to approximately sample from the posterior distribution. We construct a neural sampler that is trained with the functional gradient of the KL-divergence between the empirical sampling distribution and the target distribution, assuming the gradient resides within a reproducing kernel Hilbert space. Our generative ParVI (GPVI) approach maintains the asymptotic performance of ParVI methods while offering the flexibility of a generative sampler. Through carefully constructed experiments, we show that GPVI outperforms previous generative ParVI methods such as amortized SVGD, and is competitive with ParVI as well as gold-standard approaches like Hamiltonian Monte Carlo for fitting both exactly known and intractable target distributions.

## 1. Introduction

Bayesian inference provides a powerful framework for reasoning and prediction under uncertainty. However, computing the posterior is tractable with only a few parametric distributions, making wider applications of Bayesian inference difficult. Traditionally, MCMC and variational inference methods are utilized to provide tractable approximate inference, but these approaches face their own difficulties if the dimensionality of the space is extremely high. For example, a recent case of interest is Bayesian neural networks (BNNs), which applies Bayesian inference to deep neural

network training in order to provide a principled way to assess model uncertainty. The goal in this regime is to model the posterior of every parameter in all the weight tensors from every layer of a deep network. However, developing efficient computational techniques for approximating this intractable posterior with extremely high dimensionality remains challenging.

Recently, particle-based variational inference (ParVI) methods (Liu & Wang, 2016; Liu et al., 2019; Liu, 2017) have been proposed to represent the variational distribution by a set of particles and update them through a deterministic optimization process to approximate the posterior. While achieving both asymptotic accuracy and computational efficiency, ParVI methods are restricted by the fixed number of particles and lack the ability of drawing new samples beyond the initial set of particles. To address this issue, amortized ParVI methods (Wang & Liu, 2016) have been proposed to amortize the ParVI gradients in training a neural sampler. While being flexible in drawing samples, in Sec. 4 we show that amortized ParVI methods cannot match the convergence behavior of ParVI methods.

In this work, we propose a generative particle variational inference (GPVI) approach that addresses those issues. GPVI trains a neural sampler network by directly estimating the functional gradient w.r.t. the KL-divergence between the distribution of generated particles and the target distribution, and pulls it back to update the neural sampler. As such it allows the neural sampler to directly generate particles that match the posterior distribution, hence achieving the asymptotic accuracy and computational efficiency of ParVI methods. In figure 1, we show that the predictive distribution of 1D regression functions sampled from GPVI nearly matches that from the ParVI solution, while amortized ParVI fails.

The main computational challenge lies in a reliable estimate of the functional gradient that involves the inverse of the input-output Jacobian of the neural sampler. Instead of directly computing this term and paying a high computational cost, we introduce a helper network to estimate the inverse Jacobian vector product and train the helper network via gradient descent. By alternating between this gradient step and the gradient update of the sampler network, the computational cost is distributed over the whole train-

---

<sup>\*</sup>Equal contribution <sup>1</sup>Department of Electrical Engineering and Computer Science, Oregon State University, Corvallis, Oregon <sup>2</sup>Horizon Robotics, Cupertino, California. Correspondence to: Neale Ratzlaff <ratzlafn@oregonstate.edu>.

ing procedure.

In experiments, our proposed approach achieves comparable convergence performance as ParVI methods. It is considerably superior than that of amortized ParVI methods, while still allowing efficient sampling from the posterior. By directly applying our approach as a hypernetwork to generate BNNs, we achieve competitive performance regarding uncertainty estimation.

In summary, our contributions are three-fold,

- We propose GPVI, a new variational inference approach that trains a neural sampler to generate particles from any posterior distribution. GPVI estimates the functional gradient and uses it to update the neural sampler. It enjoys the asymptotic accuracy and computational efficiency of ParVI methods. Comparing with existing amortized ParVI methods, our approach enjoys the same efficiency and flexibility while showing considerable advantage in convergence behavior.
- We design careful techniques for efficient gradient estimates that address the challenges in approximating the product between the inverse of the Jacobian and a vector.
- We apply our approach to BNNs and achieve competitive uncertainty estimation quality for deep neural networks.

## 2. Related Work

ParVI is a recent class of non-parametric variational inference methods. (Liu & Wang, 2016; Liu, 2017) proposed stein variational gradient descent (SVGD), which deterministically updates an empirical distribution of particles toward the target distribution via a series of sequentially constructed smooth maps. Each map defines a perturbation of the particles in the direction of steepest descent towards the target distribution under the KL-divergence metric. Liu et al. (2019) cast ParVI as a gradient flow on the space of probability measures  $\mathcal{P}_2$  equipped with the  $W_2$  Wasserstein metric, and proposed ParVI methods GFSD and GFSF by smoothing the density and the function respectively. A potential limitation of SVGD and GFSF is the restriction of the transport maps to an RKHS. The fisher neural sampler (Hu et al., 2018; Grathwohl et al., 2020) lifts the function space of transportation maps to the space of  $\mathcal{L}_2$  functions, and performs a minimax optimization procedure to find the optimal perturbation. As a non-parametric approach, ParVI does not depend on the specific parameterization of the approximate posterior. ParVI is also more efficient than MCMC methods due to its deterministic updates and gradient-based optimization.

The limitation of ParVI is the inability to draw additional

samples beyond the initial set. To address this limitation, our approach allows for arbitrary draws of new samples once the generator is trained. The most notable generative ParVI approach prior to our work was amortized SVGD (Wang & Liu, 2016), which proposed using a generative model to approximately sample from the SVGD particle distribution. Amortized SVGD performs a one-step update that backpropagates the SVGD gradient back to the generator parameters, which is often not optimal. A more detailed discussion of the differences between our method and GPVI is presented in Sec. 3.5.

Performing direct Bayesian inference in the weight-space of neural networks (Neal, 2012; Buntine & Weigend, 1991; MacKay, 1992) is known to be intractable in the context of deep learning. Thus, approximations have been formulated using Langevin Dynamics (Welling & Teh, 2011; Kottakara Balan et al., 2015), Monte Carlo dropout (Gal & Ghahramani, 2016; Kingma et al., 2015), and variational inference with mean-field Gaussian priors (Graves, 2011). The variational approximation has been iterated upon extensively. Bayes by Backprop (Blundell et al., 2015) proposed an unbiased gradient estimator that allowed for training of deep Bayesian neural networks using the local reparameterization trick from Kingma et al. (2015). However, fully factorized Gaussian approximations are insufficient to capture the complex structure of a high dimensional posterior (Yao et al., 2019). As a result, structured approximate posteriors have been proposed using matrix Gaussians (Louizos & Welling, 2016), normalizing flows (Louizos & Welling, 2017; Krueger et al., 2017), and hypernetworks (Pawlowski et al., 2017).

Evaluating the performance of BNNs is difficult, as there is often no access to the ground truth posterior. Therefore, prior work often considers a “gold-standard” method such as HMC (Neal et al., 2011), or exact Gaussian processes (Rasmussen, 2003) as a proxy for the true posterior. While these methods set a high bar for BNNs, they are in general not scalable to large-scale problems. Nonetheless there were attempts to bring HMC (Strathmann et al., 2015; Chen et al., 2014) and GPs (Wilson et al., 2016; Cheng & Boots, 2017; Wang et al., 2019; Hensman et al., 2013; 2015) in the realm of tractability for high dimensional problems.

Variational inference via smooth transport maps is also used in normalizing flows (Rezende & Mohamed, 2015). But the transport maps used in normalizing flows and in Marzouk et al. (2016) are bijections which makes them unsuitable for learning distributions over neural network functions.

## 3. Generative Particle Variational Inference

We want to learn a parametric generator:  $f_{\theta} : \mathcal{Z} \rightarrow \mathbb{R}^d$ , parameterized by  $\theta$ , where  $\mathcal{Z} \subset \mathbb{R}^d$  is the convex space of




Figure 1. Predictive uncertainty of methods for a 1-D regression task. (a) HMC predictive posterior matches the uncertainty in the data; (b) SVGD performs comparably to HMC; (c) Our proposed GPVI performs similarly to SVGD, with the additional capability of sampling new particles during inference; (d) Amortized SVGD overestimates the uncertainty when the data is sparse.

input noise and  $\mathbf{x} = \mathbf{f}_\theta(\mathbf{z})$  generates a sample  $\mathbf{x}$  from an input noise  $\mathbf{z}$ . Let  $q(\mathbf{x})$  represents the implicit distribution of samples generated by  $\mathbf{f}_\theta(\mathbf{z})$ , where  $\mathbf{z} \sim N(\mathbf{0}, I_d)$ . Let  $p(\mathbf{x})$  be the target distribution, we want to solve for  $\mathbf{f}_\theta$  that minimizes the objective  $\text{KL}(q(\mathbf{x})\|p(\mathbf{x}))$ .

Our approach treats this problem from a functional optimization perspective, by first computing the functional gradient of the objective, and then pulling it back to parameter space through the function parameterization. In the rest of this section we introduce our algorithm in detail and compare it with amortized ParVI approaches.

### 3.1. Functional gradient and its pullback

Let  $\mathcal{J}(\mathbf{f}) = \text{KL}(q(\mathbf{x})\|p(\mathbf{x}))$  be the objective, where  $\mathbf{x} = \mathbf{f}(\mathbf{z})$ . If  $\mathbf{f}$  is injective, by change of variables for probability measure,

$$q(\mathbf{x}) = \frac{p_{\mathbf{z}}(\mathbf{z})}{\left| \det \left( \frac{\partial \mathbf{f}}{\partial \mathbf{z}} \right) \right|}, \quad (1)$$

where  $p_{\mathbf{z}}(\mathbf{z})$  is the distribution from which  $\mathbf{z}$  is sampled. The minimization objective becomes,

$$\mathcal{J}(\mathbf{f}) = \mathbf{E}_{\mathbf{z}} \left[ -\log p(\mathbf{f}(\mathbf{z})) + \log \frac{p_{\mathbf{z}}(\mathbf{z})}{\left| \det \left( \frac{\partial \mathbf{f}}{\partial \mathbf{z}} \right) \right|} \right]. \quad (2)$$

Consider some function approximation of  $\mathbf{f}$ , say  $\mathbf{f} = \mathbf{f}_\theta$ , then the minimization objective becomes,

$$\mathcal{J}(\theta) = \mathbf{E}_{\mathbf{z}} \left[ -\log p(\mathbf{f}_\theta(\mathbf{z})) + \log \frac{p_{\mathbf{z}}(\mathbf{z})}{\left| \det \left( \frac{\partial \mathbf{f}_\theta}{\partial \mathbf{z}} \right) \right|} \right]. \quad (3)$$

Directly computing the gradient of  $\mathcal{J}(\theta)$  w.r.t.  $\theta$  involves not only an inverse of the Jacobian  $\left( \frac{\partial \mathbf{f}}{\partial \mathbf{z}} \right)^{-1}$ , but also second derivatives of  $\mathbf{f}_\theta$  (see Appendix A for details), which is overly expensive to compute in practice.

We propose, instead, to first compute the functional gradient of (2) w.r.t.  $\mathbf{f}$ , i.e.,  $\nabla_{\mathbf{f}} \mathcal{J}(\mathbf{f})$ , and then back-propagate

it through the generator to get the gradient w.r.t.  $\theta$ , i.e.,

$$\nabla_{\theta} \mathcal{J} = \mathbf{E}_{\mathbf{z}} \left[ \frac{\partial \mathbf{f}(\mathbf{z})}{\partial \theta} \nabla_{\mathbf{f}} \mathcal{J}(\mathbf{f})(\mathbf{z}) \right]. \quad (4)$$

The following theorem gives an explicit formula for computing the functional gradient  $\nabla_{\mathbf{f}} \mathcal{J}(\mathbf{f})$  when  $\mathbf{f}$  is chosen from a Reproducing kernel Hilbert space (RKHS). The proof is provided in Appendix A.

**Theorem 3.1.** *Let  $\mathbf{x} = \mathbf{f}(\mathbf{z})$ , where  $\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})$ , vector function  $\mathbf{f} = (f^1, \dots, f^d) \in \mathcal{H}^d$  with  $f^i \in \mathcal{H}$ , where  $\mathcal{H}$  is the RKHS with kernel  $k(\cdot, \cdot)$ ,  $\mathcal{H}^d$  is equipped with inner product  $\langle \mathbf{f}, \mathbf{g} \rangle_{\mathcal{H}^d} = \sum_{i=1}^d \langle f^i, g^i \rangle_{\mathcal{H}}$ . For  $\mathcal{J}(\mathbf{f})$  well-defined by (2), we have*

$$\nabla_{\mathbf{f}} \mathcal{J}(\mathbf{f})(\mathbf{z}) = \mathbf{E}_{\mathbf{z}'} \left[ -\nabla_{\mathbf{x}} \log p(\mathbf{x}) \Big|_{\mathbf{x}=\mathbf{f}(\mathbf{z}')} k(\mathbf{z}', \mathbf{z}) - \left( \frac{\partial \mathbf{f}}{\partial \mathbf{z}'} \right)^{-1} \nabla_{\mathbf{z}'} k(\mathbf{z}', \mathbf{z}) \right]. \quad (5)$$

### 3.2. Reparameterization of $\mathbf{f}$

There are two considerations for reparameterizing  $\mathbf{f}$ . Firstly, in order for the inverse  $\left( \frac{\partial \mathbf{f}}{\partial \mathbf{z}'} \right)^{-1}$  in (5) to be well-defined, the Jacobian  $\frac{\partial \mathbf{f}}{\partial \mathbf{z}'}$  should be a square matrix, i.e., input noise  $\mathbf{z}$  should have the same dimension as the output  $\mathbf{x} = \mathbf{f}(\mathbf{z})$ . In practice, especially those applications involving BNNs, each  $\mathbf{x}$  represents parameters of a sampled neural network and therefore can be extremely high dimensional. As a result, a high dimensional  $\mathbf{f}$  can be computationally prohibitive. Secondly, in order for the change of variables density formula (1) to hold,  $\mathbf{f}$  needs to be injective, which is in general not guaranteed for an arbitrary neural network function.

To overcome the above two concerns, we consider the following parameterization,

$$\mathbf{f}_\theta(\mathbf{z}) = \mathbf{g}_\theta(\mathbf{z}^{(:k)}) + \lambda \mathbf{z}, \quad \forall \mathbf{z} \in \mathbb{R}^d, \quad (6)$$

where  $\mathbf{z}^{(:k)} \in \mathbb{R}^k$  denotes the vector consisting of the first  $k$  components of  $\mathbf{z}$ , and  $\mathbf{g}_\theta : \mathbb{R}^k \rightarrow \mathbb{R}^d$  with parameters

$\theta$  is a much slimmer neural network. In our experiments,  $g_\theta$  is designed with an input dimension  $k$  less than 30% the size of  $d$ . For high dimensional open-category experiments where  $d > 60,000$ , we use a  $k$  of less than 2% of  $d$ . For  $f_\theta$  defined by (6), the Jacobian is,

$$\left[ \frac{\partial f_\theta}{\partial z} \right]_{d \times d} = \left[ \left[ \frac{\partial g_\theta}{\partial z^{(:k)}} \right]_{d \times k} \middle| \mathbf{0}_{d \times (d-k)} \right]_{d \times d} + \lambda \mathbf{I}_d, \quad (7)$$

where  $\lambda$  is a hyper-parameter. Note that for sufficiently large  $\lambda$ , the Jacobian defined by (7) is positive definite. Since the domain  $\mathcal{Z}$  of  $z$  is convex, it is straightforward to show that  $f_\theta$  defined by (6) is injective. We include the proof in Appendix A.

In practice we set  $\lambda$  to be 1.0 and find it sufficient throughout our experiments.

### 3.3. Estimating the Jacobian inverse

The main computational challenge of (5) lies in computing the term

$$(J_f(z'))^{-1} \nabla_{z'} k(z', z), \quad (8)$$

where  $J_f(z') = \frac{\partial f}{\partial z'}$ , especially considering that we need an efficient implementation for batched  $z$  and  $z'$ .

Directly evaluating and storing the full Jacobian  $J_f(z')$  for each  $z'$  of the sampled batch is not acceptable from the standpoint of either time or memory consumption. There exist iterative methods for solving the linear equation system  $y = (J_f(z'))^{-1} \nabla_{z'} k(z', z)$  (Young, 1954; Fletcher, 1976), which involves computing the vector-Jacobian product  $J_f(z') \nabla_{z'} k(z', z)$  at each iteration. By alternating between this iterative solver and the gradient update (4), it is possible to get a computationally amenable algorithm. However, as shown in the Appendix C, such an algorithm does not converge to the target distribution  $p$  even for a simple Bayesian linear regression task. This is due to the fact that batches of both  $z$  and  $z'$  for evaluating (5) need to be re-sampled for each gradient update (4) to avoid the cumulative sampling error. Therefore, the above alternating procedure of iterative solver for  $y = (J_f(z'))^{-1} \nabla_{z'} k(z', z)$  ends up shooting a moving target for different batches of  $z$  and  $z'$  at each iterate, which is difficult for convergence.

To overcome this computational challenge, we propose a helper network, denoted by  $h_\eta(z', \nabla_{z'} k)$ , and parameterized by  $\eta$ , that consumes both  $z'$  and  $\nabla_{z'} k(z', z)$  and predicts  $(J_f(z'))^{-1} \nabla_{z'} k(z', z)$ . With the helper network, the functional gradient (5) can be computed by,

$$\nabla_f \mathcal{J}(f)(z) = \mathbf{E}_{z'} \left[ -\nabla_x \log p(x) \middle|_{x=f(z')} k(z', z) - h_\eta(z', \nabla_{z'} k(z', z)) \right]. \quad (9)$$

We use the following loss to train the helper network,

$$\mathcal{L}(\eta) = \|J_f(z') h_\eta(z', \nabla_{z'} k(z', z)) - \nabla_{z'} k(z', z)\|^2, \quad (10)$$

where  $J_f(z') h_\eta$  can be computed by the following formula given the reparameterization (6) of  $f_\theta$ ,

$$h_\eta^T \left( \frac{\partial f_\theta}{\partial z'} \right) = h_\eta^T \left( \frac{\partial g_\theta}{\partial z^{(:k)}} \right) + \lambda h_\eta^T, \quad (11)$$

where the Vector-Jacobian Product (VJP)  $h_\eta^T \left( \frac{\partial g_\theta}{\partial z^{(:k)}} \right)$  can be computed by one backward pass of the function  $g_\theta$ .

### 3.4. Summary of the algorithm

Both  $g_\theta$  and  $h_\eta$  can be trained with stochastic gradient descent (SGD), and our algorithm alternates between the SGD updates of  $g_\theta$  and  $h_\eta$ . Note that the helper network  $h_\eta$  only has to chase the update of  $g_\theta$ , but no more extra moving targets due to re-sampling of  $z$  and  $z'$ . Our experiments in Sec. 4 and the appendix show that the helper network is able to efficiently approximate (8) which enables the convergence of the gradient update (4). Our overall Generative Particle VI (GPVI) algorithm is summarized in Algorithm 1.

---

#### Algorithm 1 Generative Particle VI (GPVI)

---

**Initialize** generator  $g_\theta$ , helper  $h_\eta$ , and learning rate  $\epsilon$

**while** Not converge **do**

1. sample two batches  $\{z_i\}, \{z'_i\} \sim N(\mathbf{0}, I_d)$
2. compute  $k(z', z)$  and  $\nabla_{z'} k(z', z)$
3. forward  $g_\theta$  to compute  $f_\theta(z)$  and  $f_\theta(z')$  by (6)
4. forward  $h_\eta$  to compute  $h_\eta(z', \nabla_{z'} k(z', z))$
5. backward  $g_\theta$  to compute the VJP  $h_\eta^T \left( \frac{\partial g_\theta}{\partial z^{(:k)}} \right)$  and then construct  $J_f(z') h_\eta$  by (11),
6. update  $h_\eta$  by  $\eta \leftarrow \eta - \epsilon \nabla_\eta \mathcal{L}$ , where  $\nabla_\eta \mathcal{L}$  is computed by back-propagating (10)
7. compute the functional gradient by (9)
8. update  $\theta$  by  $\theta \leftarrow \theta - \epsilon \nabla_\theta \mathcal{J}$ , where  $\nabla_\theta \mathcal{J}$  is computed by (4)

**end**

---

### 3.5. Comparison with Amortized SVGD

Stein variational gradient descent (SVGD) represents  $q(x)$  by a set of particles  $\{x_i\}_{i=1}^n$ , which are updated iteratively by,

$$x_i \leftarrow x_i + \epsilon \phi^*(x_i), \quad (12)$$

where  $\epsilon$  is a step size and  $\phi^* : \mathbb{R}^d \rightarrow \mathbb{R}^d$  is a vector field (perturbation) on the space of particles that corresponds to the optimal direction to perturb particles, i.e.,

$$\phi^* = \arg \min_{\phi \in \mathcal{F}} \left\{ \frac{d}{d\epsilon} \text{KL}(q_{[\epsilon\phi]}(x) \| p(x)) \middle|_{\epsilon=0} \right\},$$

where  $q_{[\epsilon\phi]}(\mathbf{x})$  denotes the density of particles updated by (12) using the perturbation  $\phi$ , where the density of original particles is  $q(\mathbf{x})$ . When  $\mathcal{F}$  is chosen to be the unit ball of some RKHS  $\mathcal{H}$  with kernel function  $k(\cdot, \cdot)$ , SVGD gives the following closed form solution for  $\phi^*$ ,

$$\phi^*(\mathbf{x}) = \mathbf{E}_{\mathbf{x}' \sim q} [\nabla_{\mathbf{x}'} \log p(\mathbf{x}') k(\mathbf{x}', \mathbf{x}) + \nabla_{\mathbf{x}'} k(\mathbf{x}', \mathbf{x})]. \quad (13)$$

To turn SVGD into a neural sampler, Amortized SVGD (Wang & Liu, 2016) first samples particles from a generator and then back-propagates the particle gradients (13) through the generator to update the generator parameters. Let  $\mathbf{x} = \mathbf{f}_\theta(\mathbf{z})$ ,  $\mathbf{z} \sim N(\mathbf{0}, I_d)$  be the particle generating process, where  $\mathbf{f}_\theta$  is the generator parameterized by  $\theta$ , amortized SVGD updates  $\theta$  by,

$$\theta \leftarrow \theta + \epsilon \sum_{i=1}^m \frac{\partial \mathbf{f}_\theta(\mathbf{z}_i)}{\partial \theta} \phi^*(\mathbf{f}_\theta(\mathbf{z}_i)), \quad (14)$$

where  $\phi^*(\mathbf{f}_\theta(\mathbf{z}_i))$  is computed by (13).

The following Lemma gives an explicit view regarding what functional gradient amortized SVGD back-propagates through the generator. The proof is given in the Appendix A.

**Lemma 3.2.** *If particles are generated by  $\mathbf{x} = \mathbf{f}(\mathbf{z})$ ,  $\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})$ , Eqn. (13) is the functional gradient of the KL objective w.r.t. the perturbation function  $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^d$  applied on the output space of  $\mathbf{f}$ , i.e.,*

$$\nabla_\phi \Big|_{\phi=id} KL(q(\mathbf{x}) \| p(\mathbf{x})) = -\phi^*, \quad (15)$$

where  $\mathbf{x} = \phi(\mathbf{f}(\mathbf{z}))$  and  $\phi = (\phi^1, \dots, \phi^d) \in \mathcal{H}^d$  with  $\phi^i \in \mathcal{H}$ , where  $\mathcal{H}$  is the RKHS with kernel  $k(\cdot, \cdot)$ ,  $\mathcal{H}^d$  is equipped with inner product  $\langle \phi, \xi \rangle_{\mathcal{H}^d} = \sum_{i=1}^d \langle \phi^i, \xi^i \rangle_{\mathcal{H}}$ .

To see the difference between our GPVI and amortized SVGD, both methods learn a generator  $\mathbf{f}$  that generates particles  $\mathbf{x}$  from input noise  $\mathbf{z}$ . To update  $\mathbf{f}$ , gradients of the KL objective  $\mathcal{J}(\mathbf{f})$  have to be backpropagated to  $\mathbf{f}$  in both cases. Here, the best is to use the steepest descent direction for  $\mathbf{f}$ , which GPVI does (proved in Thm 3.1). As Lemma 3.2 shows, the SVGD gradient  $\phi^*$  given by (13) is optimal for a transformation  $\phi$  applied to the set of current particles. In both amortized SVGD and GPVI, such a set is not kept and  $\phi$  does not exist since they only maintain and update  $\mathbf{f}$ . But amortized SVGD directly uses  $\phi^*(\mathbf{f})$  in (14) as the direction to update  $\mathbf{f}$ , which is an unproven use – it only coincides with the steepest descent direction for  $\mathbf{f}$  used by GPVI in special cases such as  $\mathbf{f} = \text{id}$ , where (5) and (13) are equivalent because  $\mathbf{z} = \mathbf{x}$  and the Jacobian is identity. Regarding the definition of RKHS, GPVI’s RKHS approximates the tangent space of  $\mathbf{f}$ , it is therefore naturally defined on  $\mathbf{z}$ . (Amortized) SVGD uses

RKHS to approximate the tangent space of  $\phi$ , which is naturally defined on  $\mathbf{x}$ .

The amortizing step applied in both methods, i.e., back-propagating some functional gradient to update the generator parameters, is in general not guaranteed to keep the original descent direction of the functional gradient. However, GPVI amortizes the steepest descent direction which is optimal in first order sense, while amortized SVGD amortizes a non-steepest descent direction, which is more likely to result in a non-descending direction after amortizing. In practice, as shown in the next section, our approach consistently outperforms amortized SVGD in approximating the target distribution and capturing model uncertainty.

## 4. Experiments

To demonstrate the effectiveness of our approach for approximate Bayesian inference, we evaluated GPVI in two different settings: density estimation and BNNs. In our density estimation experiment, we trained our generator to draw samples from a target distribution, showing that GPVI can accurately fit the posterior from the data. For our BNN experiments, we evaluated on regression, classification, and high dimensional open-category tasks. Our experiments show that among all methods compared, GPVI is the only method to excel in both sampling efficiency and asymptotic performance. We compare GPVI with ParVI methods: SVGD (Liu & Wang, 2016), GFSF (Liu et al., 2019), and KSD (Hu et al., 2018; Grathwohl et al., 2020), as well as their corresponding amortized versions. We also compared with a mean-field VI approach Bayes by Back-prop (MF-VI) (Blundell et al., 2015), as well as deep ensembles (Lakshminarayanan et al., 2017), and HMC (Neal et al., 2011). We emphasize that our aim is not to only maximize the likelihood of generated samples, rather we want to closely approximate the posterior of parameterized functions given the data. Thus, predictions w.r.t data unseen during training should reflect the epistemic uncertainty of the model.

The details of our experimental setup are as follows. In our BNN experiments, we parameterized samples from the target distribution as neural networks with a fixed architecture. In the low dimensional regression and classification settings, we drew 100 samples from the approximate posterior for both training and evaluation, allowing us to evaluate the predictive mean and variance. For methods without a sampler e.g. ParVI and deep ensembles, we initialized a 100 member ensemble. In the high dimensional open-category tasks, we instead used 10 samples due to the larger computation cost. For methods that utilize a hypernetwork such as GPVI and amortized ParVI, we used Gaussian input noise  $\mathbf{z} \sim \mathcal{N}(0, I)$  and varied the hypernetwork architecture depending on the task. For ParVI and deep

Method	$\Sigma$ error (2d) ↓	$\Sigma$ error (5d) ↓
Amortized SVGD	0.10 ± .09	0.37 ± .32
Amortized GFSF	0.18 ± .04	0.21 ± .13
Amortized KSD	0.28 ± .32	1.68 ± .52
GPVI	0.14 ± .08	0.14 ± .04

Table 1. Comparison of generative Particle VI approaches for density estimation of 2d and 5d Gaussian distributions.

ensembles we randomly initialized each member of the ensemble. In all methods except HMC we used the Adam optimizer (Kingma & Ba, 2014). A detailed description of network architectures, chosen hyperparameters, and experimental settings is given in the appendix.

#### 4.1. Density Estimation

We first evaluate the ability of generative approaches to fit a target distribution from data. We used 2D and 5D zero mean Gaussian distributions with non-diagonal covariances as our target distributions. We consider sampler networks with one hidden layer of width 2 and 5, respectively, and input noise with the same dimensionality as the output. The error is defined as the difference between the true variance and the variance of the sampler output distribution  $WW^T - \Sigma$ , where  $W$  is the weights of the sampler. As shown in Table 1, while the estimation problem becomes harder with increased dimension, our approach performs consistently well while amortized ParVI methods suffer more or less from a performance drop and all end up inferior to our approach. We have included more density estimation results in the Appendix.

#### 4.2. Bayesian Linear Regression

We evaluated all methods on Bayesian linear regression to investigate how well each method can fit a unimodal normal distribution over linear function weights. The target function is a linear regressor with parameter vector  $\beta \in \mathbb{R}^d$ , i.e.,  $\mathbf{y} = \mathbf{X}\beta + \epsilon$ , where  $\mathbf{X} \sim \mathcal{N}(0, I_d)$ ,  $\epsilon \sim \mathcal{N}(0, I)$ , and  $\beta^i \sim U(0, 1) + 5$ . We chose such linear Gaussian settings where we can explicitly compute the target posterior  $p(\beta|\mathbf{X}, \mathbf{y}) \sim \mathcal{N}((\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{y}, \mathbf{X}^T \mathbf{X})$  given observations  $(\mathbf{X}, \mathbf{y})$ , allowing us to numerically evaluate how well each method fits the target distribution. Each method was trained to regress  $\mathbf{y}$  from  $\mathbf{X}$ . For GPVI and amortized ParVI methods that make use of a hypernetwork, we used a linear generator with input noise  $\mathbf{z} \in \mathbb{R}^d$ . In this setting, the output of a generator with bias  $\mathbf{b}$  and weights  $W$  follows the distribution  $\mathcal{N}(\mathbf{b}, WW^T)$ , which should match the posterior  $p(\beta|\mathbf{X}, \mathbf{y})$ . For ParVI methods and deep ensembles, we initialized an ensemble of linear regressors with  $d$  parameters. We computed the mean and covariance of the learned parameters to measure the quality of posterior fit. For MF-VI we chose a standard normal prior on the weights, and likewise computed the mean and covari-

Method	$\mu$ error ↓	$\Sigma$ error ↓
SVGD	0.006 ± .0024	<b>0.125 ± .03</b>
GFSF	0.003 ± .0014	0.139 ± .07
KSD	0.009 ± .0006	0.373 ± .11
Amortized SVGD	<b>0.002 ± .0009</b>	0.158 ± .04
Amortized GFSF	<b>0.002 ± .0003</b>	0.209 ± .05
Amortized KSD	0.004 ± .0007	0.430 ± .10
MF-VI	0.004 ± .0035	0.303 ± .04
Deep Ensemble	0.004 ± .0002	1.0 ± 0
HMC	0.009 ± .0003	0.181 ± .05
GPVI	<b>0.002 ± .0007</b>	<b>0.128 ± .04</b>
GPVI Exact Jac	<b>0.002 ± .0004</b>	<b>0.106 ± .07</b>

Table 2. Bayesian linear regression. Reported error is the  $L_2$  norm of the difference between the learned mean and covariance parameters, and the ground truth after 50000 iterations.

ance of weight samples to measure posterior fit. We also test a variant of GPVI where we exactly compute the Jacobian inverse instead of using our helper network. We set  $d = 3$  in our experiments. While simple, this quantitative sanity check is crucial before evaluating on more complicated domains where the true posterior is not available. It will be evident in later experiments that the ability to closely approximate the true posterior under this simple setup is indicative of performance on higher dimensional tasks.

Results are shown in table 2. Our GPVI outperforms amortized ParVI, MF-VI, and surprisingly even HMC. GPVI is also very competitive with the best ParVI methods, SVGD and GFSF. We note that GPVI is also competitive with the “Exact-Jac” variant, where we invert the Jacobian explicitly, instead of using the helper network. This “Exact-Jac” version of GPVI is not scalable, but it serves as a benchmark to show that our helper network is able to closely approximate the exact solution. Finally, as deep ensembles lacks any mechanism for correctly estimating the covariance, it performs poorly as expected.

#### 4.3. Multimodal Classification

The previous experiment possesses an analytically known target distribution. However, most problems of interest involve distributions without closed-form expressions. Therefore, we further tested on a 2-dimensional, 4-class classification problem, where each class consists of samples from one component of a mixture distribution. The mixture distribution is defined as  $p(x) = \sum_{i=1}^4 \mathcal{N}(\mu_i, 0.3)$ , with means  $\mu_i \in \{(-2, -2), (-2, 2), (2, -2), (2, 2)\}$ . We assigned labels  $y_i \in \{1, 2, 3, 4\}$  according to the index of the mixture component the samples were drawn from. For this task, samples are weight parameters  $\theta$  of two-layer neural networks with 10 hidden units in each layer and ReLU activations, denoted by

$$f_{\theta} : \mathbb{R}^2 \rightarrow \mathbb{R}^4.$$

To train each method, we drew a total of 100 training points, and 200 testing points from the target distribution. To evaluate the posterior predictive distribution, we drew points from a grid spaced from  $\{-10, 10\}$ , then plot the predictive distribution as measured by the standard deviation in predictions among model samples in figure 2.

In this setting, the true posterior  $p(\theta|\mathbf{X}, y)$  is unknown and past work often relies on “gold-standard” approaches like HMC to serve as the ground truth. In the previous Bayesian linear regression tests, however, we saw that HMC was outperformed by GPVI as well as ParVI methods SVGD and GFSF. In the current test, intuitively, a “gold-standard” approach should yield a predictive distribution with high variance (high uncertainty) in regions far from the training data, and low variance (low uncertainty) in regions near each mixture component. As shown in figure 2, GPVI and SVGD both have higher uncertainty in no-data regions than HMC, while remaining confident on the training data. On the other hand, MF-VI and amortized SVGD both underestimate the uncertainty. The performance of MF-VI is as expected, as mean-field VI approaches are known to characteristically underestimate uncertainty. Results of all methods as well as more tests on an easier 2-class variant are included in the Appendix.

#### 4.4. Open Category

To evaluate the stability of our approach we turn to large-scale image classification experiments. In this setting, it is not possible to exactly measure the accuracy of posterior fits. We instead utilize the open-category task to test if our uncertainty estimations can help detect outlier examples. The open category task defines a set of inlier classes that are seen during training, and a set of outlier classes only used for evaluation. While in principle, the content of the outlier classes can be arbitrary, we use semantically similar outlier classes by splitting the training dataset into inlier and outlier classes. The open-category task is in general more difficult than out-of-distribution experiments where different datasets are used as outlier classes, since distributions of categories from the same dataset may be harder to discriminate.

We evaluated on the MNIST and CIFAR-10 image datasets, following Neal et al. (2018) to split each dataset into 6 inlier classes and 4 outlier classes. We performed standard fully supervised training on the 6 inlier classes, and measured uncertainty in the 4 unseen outlier classes. We evaluated the uncertainty using two widely-used statistics: area under the ROC curve (AUC), and the expected calibration error (ECE). The AUC score measures how well a binary classifier discriminates between predictions made on inlier inputs, vs predictions made on outlier inputs. A perfect

AUC score of 1.0 indicates a perfect discrimination, while a score of 0.5 indicates that the two sets of predictions are indistinguishable. ECE partitions predictions into equally sized bins, and computes the  $L_1$  difference in expected accuracy and confidence between bins, which represents the calibration error of the bin. ECE computes a weighted average of the calibration error of each bin. Together AUC and ECE tell us how well a model can detect outlier inputs, as well as how well the model fits the training distribution.

**MNIST** consists of 70,000 grayscale images of handwritten digits at 28x28 resolution, divided into 60,000 training images and 10,000 testing images. We further split the dataset by only using the first six classes for training and testing. The remaining four classes are only used to compute the AUC and ECE statistics. We chose the LeNet-5 classifier architecture for all models, and trained for 100 epochs. Due to the larger computational burden, we only consider 10 samples from each method’s approximate posterior for both training and evaluation. Note that our approach is capable of generating many more, but it would be computationally costly to train much more for ensembles and particle VI approaches. For GPVI and amortized ParVI methods we used a 3 layer MLP hypernetwork with layer widths [256, 512, 1024], ReLU activations, and input noise  $z \in \mathbb{R}^{256}$ . We did not test HMC in this setting, as the computational demand is too high.

Table 3 shows the results. All RKHS-based methods (GPVI, ParVI, amortized ParVI) as well as deep ensembles achieve high supervised (“clean”) accuracy with the LeNet architecture. MF-VI underfits slightly, while KSD struggled to achieve competitive accuracy even after 100 epochs. All methods achieved an AUC over 0.95, but SVGD, GFSF and GPVI have the highest AUC values, respectively. In terms of calibration, GPVI and RKHS-based ParVI methods are the best calibrated. KSD is the worst calibrated model, and deep ensembles/MF-VI have middling performance.

Method	Clean	AUC $\uparrow$	ECE $\downarrow$
SVGD	99.3	<b>.989 <math>\pm</math> .001</b>	<b>.001 <math>\pm</math> .0002</b>
GFSF	99.2	<b>.988 <math>\pm</math> .003</b>	.002 $\pm$ .0003
KSD	97.7	.964 $\pm$ .005	.014 $\pm$ .0007
Amortized SVGD	99.1	.958 $\pm$ .015	<b>.002 <math>\pm</math> .0007</b>
Amortized GFSF	99.2	.978 $\pm$ .005	.004 $\pm$ .0013
Amortized KSD	97.7	.951 $\pm$ .008	.017 $\pm$ .0010
MF-VI	98.6	.951 $\pm$ .008	.014 $\pm$ .0027
Deep Ensemble	99.3	.972 $\pm$ .002	.008 $\pm$ .0060
GPVI	99.3	<b>.988 <math>\pm</math> .001</b>	<b>.001 <math>\pm</math> .0005</b>

Table 3. Results for open-category classification on MNIST. We show the result of standard supervised training (Clean), as well as AUC and ECE statistics computed from training on a subset of classes and testing on the rest of the classes as outliers.

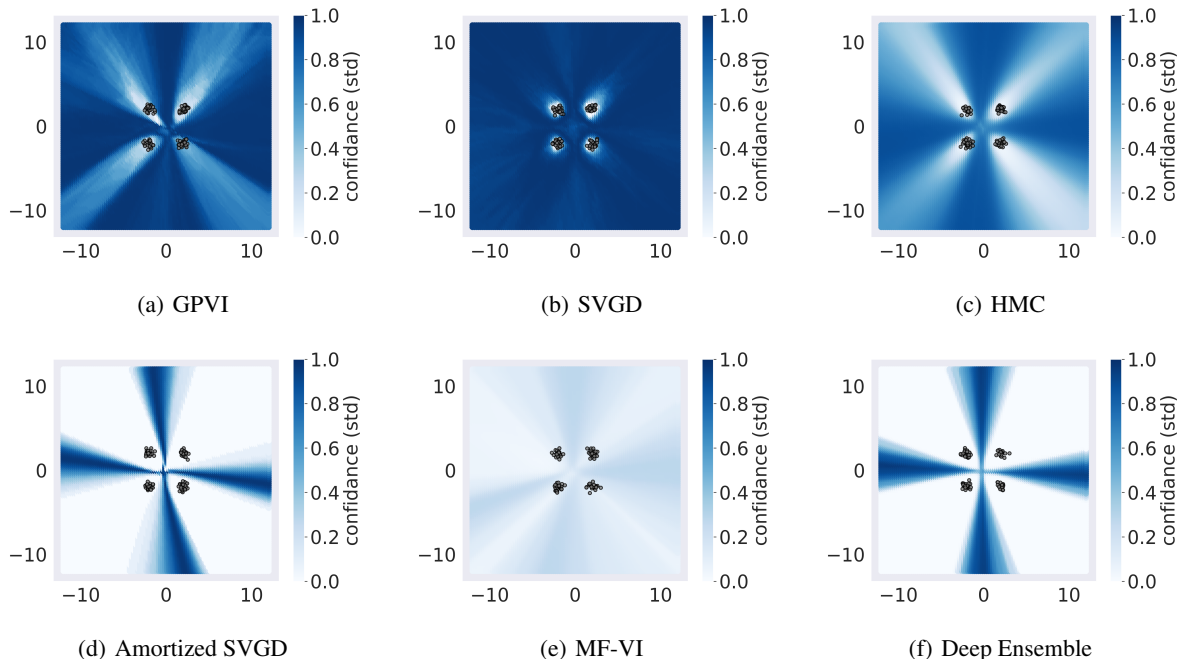


Figure 2. Predictive uncertainty of each method on the 4-class classification task, as measured by the standard deviation between predictions of sampled functions. Regions of high uncertainty are shown as darker, while lighter regions correspond to lower uncertainty. The training data is shown as samples from four unimodal normal distributions. It can be seen that amortized SVGD, MF-VI and deep ensembles significantly underestimate the uncertainty in regions with no training data

**CIFAR-10** consists of 60,000 RGB images depicting 10 object classes at 32x32 resolution, divided into 50,000 training images and 10,000 testing images. We adopted the same 6 inlier / 4 outlier split used in (Neal et al., 2018) for the open-category setting. For this task we used a CNN with 3 convolutional layers and two linear layers, which is much smaller than SOTA classifiers for CIFAR-10. Though the classification accuracy would suffer, it allows us to clearly evaluate our method without considering interactions with architectural components such as BatchNorm or residual connections. We also used 10 samples from each method’s approximate posterior for training and evaluation. For GPVI and amortized ParVI methods we used the same hypernetwork architecture as in the MNIST setup. Table 4 shows the results where it can be seen that our GPVI almost match the performance of SVGD and GFSF in terms of AUC while doing a bit better on ECE.

#### 4.5. Discussion

Much of the focus regarding recent work on Bayesian neural networks concerns their performance on open-category and out-of-distribution tasks with high dimensional image datasets. Instead, we show that our method closely approximates the target posterior, both in tasks where the posterior is explicitly known, as well as when it is intractable. The Bayesian linear regression and density estimation tasks served as sanity checks. Because the poste-

Method	Clean	AUC $\uparrow$	ECE $\downarrow$
SVGD	80.3	<b>.683 <math>\pm</math> .008</b>	.055 $\pm$ .004
GFSF	80.6	<b>.681 <math>\pm</math> .004</b>	.068 $\pm$ .012
Amortized SVGD	71.12	.636 $\pm$ .018	.073 $\pm$ .029
Amortized GFSF	71.09	.583 $\pm$ .007	.042 $\pm$ .029
MF-VI	70.0	.649 $\pm$ .006	<b>.016 <math>\pm</math> .002</b>
Deep Ensemble	73.54	.652 $\pm$ .018	.033 $\pm$ .011
GPVI	76.2	.677 $\pm$ .008	<b>.018 <math>\pm</math> .015</b>

Table 4. Open-category classification on CIFAR-10. We show results of standard supervised training (Clean), as well as AUC and ECE of each method trained in the open-category setting.

rior was known explicitly, we could quantitatively test how well each method fit the posterior. For density estimation, GPVI outperforms amortized approaches at matching the target covariance (Table 1). In Bayesian linear regression, while the approximation was close for all methods, there was a clear hierarchy in terms of which types of methods produced the tightest approximation (Table 2). GPVI and RKHS-based ParVI achieved the best posterior fit overall, and we see in further experiments that quality of fit here is indicative of performance in more difficult tasks.

The four-class classification problem (Fig. 2), while seemingly simple, is particularly difficult for most methods we evaluated since many methods tend to overgeneralize to the corners. MF-VI underestimates the uncertainty as expected



(Yao et al., 2019; Minka, 2001; Bishop, 2006). Notably, amortized SVGD also underestimates the uncertainty with a predictive distribution resembling that of a standard ensemble. We believe this is due to the compounding approximation error explained in Section 3.5, i.e., naively back-propagating the Stein variational gradient to update the generator is more likely to end up with a non-descent direction compared with back-propagating the exact functional gradient as GPVI. As shown in Fig. 2, the posterior approximation of GPVI is tighter, with uncertainty that better matches the data distribution. SVGD performs the best in this task, with high uncertainty everywhere except in regions near observed data. Surprisingly, GPVI and SVGD outperform HMC here, with clearly higher uncertainty near the corners of the sample space. Note that with the functional approximation by neural networks, it is hard to determine if the true posterior exactly matches the intuitively “ideal” uncertainty plot where low variance only shows up around each mixture component. On the other hand, HMC is guaranteed to converge to the true posterior over time (Durmus et al., 2017). Given a fixed computational budget, it is possible that GPVI or SVGD could achieve better performance. GPVI also has the extra benefit of the ability to draw additional samples, which is not possible with ParVI or HMC.

Our final experimental setting, open-category, reveals that GPVI consistently has a higher AUC than all other scalable sampling approaches. On MNIST, GPVI is among the overall top performers, together with two ParVI methods SVGD and GFSE. On CIFAR-10, GPVI is among the top two performers in ECE and performs only slightly behind the two best ParVI methods in AUC. Most importantly, GPVI outperforms with clear margins all amortized ParVI methods on both datasets under all metrics. This is consistent with all other qualitative and quantitative experiments we have conducted, which again shows the advantage of GPVI over existing generative ParVI methods.

We found that KSD completely failed in CIFAR-10, achieving slightly better than random accuracy. We believe this is due to the way that KSD performs the particle update. Where SVGD has a closed form expression for the particle transportation map, KSD parameterizes it with a critic network that has the input-output dimensionality of the particle parameters. Training this critic naturally becomes difficult when applying to neural network functions.

In the open-category setting, there is added time complexity for GPVI relative to amortized methods due to training our helper network. Compared to amortized SVGD, the extra time complexity of training GPVI is a constant factor less than 1. In our MNIST experiments, this additive factor is 0.28; in CIFAR-10 experiments, this factor is 0.78. The difference in MNIST and CIFAR-10 is due to that the out-

put dimension of the helper network scales with the output dimension of the generator. We believe this is acceptable given the performance gain and the possible real (offline training) applications of GPVI.

## 5. Conclusion

We have presented a new method that fuses the best aspects of parametric VI with non-parametric ParVI. GPVI has asymptotic convergence on par with ParVI. Additionally, GPVI can efficiently draw samples from the posterior. We also presented a method for efficiently estimating the product between the inverse of the Jacobian of a deep network, and a gradient vector. Our experiments showed that GPVI performs on par with ParVI, and outperforms amortized ParVI and other competing methods in Bayesian linear regression, a classification task, as well as open-category tasks on MNIST and CIFAR-10. In the future we want to explore the efficacy of our method applied to large scale tasks like image generation.

## Acknowledgements

Neale Ratzlaff and Li Fuxin were partially supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. HR001120C0011 and HR001120C0022. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of DARPA.

## References

- Bishop, C. M. *Pattern recognition and machine learning*. springer, 2006.
- Blundell, C., Cornebise, J., Kavukcuoglu, K., and Wierstra, D. Weight uncertainty in neural networks. *arXiv preprint arXiv:1505.05424*, 2015.
- Buntine, W. L. and Weigend, A. S. Bayesian back-propagation. *Complex systems*, 5(6):603–643, 1991.
- Chen, T., Fox, E., and Guestrin, C. Stochastic gradient hamiltonian monte carlo. In *International conference on machine learning*, pp. 1683–1691, 2014.
- Cheng, C.-A. and Boots, B. Variational inference for gaussian process models with linear complexity. In *Advances in Neural Information Processing Systems*, pp. 5184–5194, 2017.
- Durmus, A., Moulines, E., and Saksman, E. On the convergence of hamiltonian monte carlo. *arXiv preprint arXiv:1705.00166*, 2017.
- Fletcher, R. Conjugate gradient methods for indefinite systems. In *Numerical analysis*, pp. 73–89. Springer, 1976.

- Gal, Y. and Ghahramani, Z. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pp. 1050–1059, 2016.
- Grathwohl, W., Wang, K.-C., Jacobsen, J.-H., Duvenaud, D., and Zemel, R. Learning the stein discrepancy for training and evaluating energy-based models without sampling. In *International Conference on Machine Learning*, pp. 3732–3747. PMLR, 2020.
- Graves, A. Practical variational inference for neural networks. *Advances in neural information processing systems*, 24:2348–2356, 2011.
- Hensman, J., Fusi, N., and Lawrence, N. D. Gaussian processes for big data. *arXiv preprint arXiv:1309.6835*, 2013.
- Hensman, J., Matthews, A., and Ghahramani, Z. Scalable variational gaussian process classification. 2015.
- Hu, T., Chen, Z., Sun, H., Bai, J., Ye, M., and Cheng, G. Stein neural sampler. *arXiv preprint arXiv:1810.03545*, 2018.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kingma, D. P., Salimans, T., and Welling, M. Variational dropout and the local reparameterization trick. *Advances in neural information processing systems*, 28: 2575–2583, 2015.
- Korattikara Balan, A., Rathod, V., Murphy, K. P., and Welling, M. Bayesian dark knowledge. *Advances in Neural Information Processing Systems*, 28:3438–3446, 2015.
- Krueger, D., Huang, C.-W., Islam, R., Turner, R., Lacoste, A., and Courville, A. Bayesian hypernetworks. *arXiv preprint arXiv:1710.04759*, 2017.
- Lakshminarayanan, B., Pritzel, A., and Blundell, C. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in neural information processing systems*, pp. 6402–6413, 2017.
- Liu, C., Zhuo, J., Cheng, P., Zhang, R., and Zhu, J. Understanding and accelerating particle-based variational inference. In *International Conference on Machine Learning*, pp. 4082–4092. PMLR, 2019.
- Liu, Q. Stein variational gradient descent as gradient flow. In *Advances in neural information processing systems*, pp. 3115–3123, 2017.
- Liu, Q. and Wang, D. Stein variational gradient descent: A general purpose bayesian inference algorithm. *Advances in neural information processing systems*, 29: 2378–2386, 2016.
- Louizos, C. and Welling, M. Structured and efficient variational deep learning with matrix gaussian posteriors. In *International Conference on Machine Learning*, pp. 1708–1716, 2016.
- Louizos, C. and Welling, M. Multiplicative normalizing flows for variational bayesian neural networks. *arXiv preprint arXiv:1703.01961*, 2017.
- MacKay, D. J. A practical bayesian framework for back-propagation networks. *Neural computation*, 4(3):448–472, 1992.
- Marzouk, Y., Moselhy, T., Parno, M., and Spantini, A. An introduction to sampling via measure transport. *arXiv preprint arXiv:1602.05023*, 2016.
- Minka, T. P. *A family of algorithms for approximate Bayesian inference*. PhD thesis, Massachusetts Institute of Technology, 2001.
- Neal, L., Olson, M., Fern, X., Wong, W.-K., and Li, F. Open set learning with counterfactual images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 613–628, 2018.
- Neal, R. M. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012.
- Neal, R. M. et al. Mcmc using hamiltonian dynamics. *Handbook of markov chain monte carlo*, 2(11):2, 2011.
- Pawlowski, N., Brock, A., Lee, M. C., Rajchl, M., and Glocker, B. Implicit weight uncertainty in neural networks. *arXiv preprint arXiv:1711.01297*, 2017.
- Rasmussen, C. E. Gaussian processes in machine learning. In *Summer School on Machine Learning*, pp. 63–71. Springer, 2003.
- Rezende, D. and Mohamed, S. Variational inference with normalizing flows. In *International Conference on Machine Learning*, pp. 1530–1538. PMLR, 2015.
- Strathmann, H., Sejdinovic, D., Livingstone, S., Szabo, Z., and Gretton, A. Gradient-free hamiltonian monte carlo with efficient kernel exponential families. *Advances in Neural Information Processing Systems*, 28:955–963, 2015.
- Wang, D. and Liu, Q. Learning to draw samples: With application to amortized mle for generative adversarial learning. *arXiv preprint arXiv:1611.01722*, 2016.

- Wang, K., Pleiss, G., Gardner, J., Tyree, S., Weinberger, K. Q., and Wilson, A. G. Exact gaussian processes on a million data points. In *Advances in Neural Information Processing Systems*, pp. 14648–14659, 2019.
- Welling, M. and Teh, Y. W. Bayesian learning via stochastic gradient langevin dynamics. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pp. 681–688, 2011.
- Wilson, A. G., Hu, Z., Salakhutdinov, R. R., and Xing, E. P. Stochastic variational deep kernel learning. *Advances in Neural Information Processing Systems*, 29:2586–2594, 2016.
- Yao, J., Pan, W., Ghosh, S., and Doshi-Velez, F. Quality of uncertainty quantification for bayesian neural network inference. *arXiv preprint arXiv:1906.09686*, 2019.
- Young, D. Iterative methods for solving partial difference equations of elliptic type. *Transactions of the American Mathematical Society*, 76(1):92–111, 1954.