NEURAL OPERATOR VARIATIONAL INFERENCE BASED ON REGULARIZED STEIN DISCREPANCY FOR DEEP GAUSSIAN PROCESSES

Anonymous authors
Paper under double-blind review

ABSTRACT

A Deep Gaussian Process (DGP) model is a hierarchical composition of GP models that provides a deep Bayesian nonparametric approach to infer the posterior. Exact Bayesian inference is usually intractable for DGPs, motivating the use of various approximations. We theoretically demonstrate that the traditional alternative of mean-field Gaussian assumptions across the hierarchy leads to lack of expressiveness and efficacy of DGP models, whilst stochastic approximation often incurs a significant computational cost. To address this issue, we propose **N**eural **O**perator **V**ariational **I**nference (NOVI) for Deep Gaussian Processes, where a sampler is obtained from a neural generator through minimizing Regularized Stein Discrepancy in $\mathcal{L}_2$ space between the approximate distribution and true posterior. Wherein a minimax problem is obtained and solved by Monte Carlo estimation and subsampling stochastic optimization. We experimentally demonstrate the effectiveness and efficiency of the proposed model, by applying it to a more flexible and wider class of posterior approximations on data ranging in size from hundreds to tens of thousands. By comparison, NOVI is superior to previous methods in both classification and regression.

1 INTRODUCTION

Gaussian processes (GPs) Rasmussen & Williams (2006) have proven to be extraordinarily effective as a tool for statistical inference and machine learning, for example when combined with thresholding to perform classification tasks via probit models Rasmussen & Williams (2006); Neal (1997) or to find interfaces in Bayesian inversion Iglesias et al. (2016). However, the joint Gaussian assumption of the latent function values can be restrictive in a number of circumstances. This is due to at least two factors: first, not all prior information is purely expressible in terms of mean and covariance, and second, Gaussian marginals are insufficient for many applications such as in the sparse data scenario, where the constructed probability distribution is far from posterior contraction. Thus, Deep Gaussian processes (DGPs) Damianou & Lawrence (2013) have been proposed to circumvent both of these constraints.

A DGP model is a hierarchical composition of GP models that provides a deep probabilistic nonparametric approach with sound uncertity quantification Ober & Aitchison (2021). The non-Gaussian distribution over composition functions yields both expressive capacity and intractable inference Dunlop et al. (2018). Previous work on DGP models utilized variational inference with a combination of sparse Gaussian processes Snelson & Gharahmani (2005); Quiñonero-Candela & Rasmussen (2005) and mean-field Gaussian assumptions for approximate posterior adjoint with stochastic optimization to scale up DGPs to large datasets like DSVI Salimbeni & Deisenroth (2017). These strategies often incorporate a collection of inducing points ($M \ll N$) whose position is learned alongside the other model hyperparameters, reduicng the training cost to $\mathcal{O}\left(NM^2\right)$.

While the mean-field Gaussian assumptions of the approximate posterior simplifies the computation, these assumptions impose overly stringent constraints, potentially limiting the expressiveness and effectiveness of such deterministic approximation approaches for DGP models Havasi et al. (2018); Yu et al. (2019); Ustyuzhaninov et al. (2020); Lindinger et al. (2020). To solve the aforementioned problems, SGHMC Havasi et al. (2018) draws unbiased samples from the posterior belief using the

stochastic approximation approach. However, due to its sequential sampling method, generating such samples is computationally expensive for both training and prediction, and its convergence is more challenging to evaluate in finite time Gao et al. (2021). In order to solve this problem without introducing a higher time complexity, we address the issue by operator variational inference Ranganath et al. (2016), a black-box algorithm that uses operators for optimizing any operator objective with data subsampling, where a minimax problem is obtained and solved by Monte Carlo estimation.

The main contributions are as follows:

- We propose NOVI for DGPs, a novel variational framwork based on Stein Discrepancy and operator variational inference with a neural generator. It minimizes Regularized Stein Discrepancy in $\mathcal{L}_2$ space between the approximate distribution and true posterior to construct a more flexible and wider class of posterior approximations than previous methods based on mean-field Gaussian posterior assumptions.
- We theoretically demonstrate that our training schedule is equivalent to optimizing the Fisher divergence between the approximation and the true posterior while the bias raised by our method can be bounded by Fisher divergence in Section 5.
- We experimentally demonstrate the effectiveness and efficiency of the proposed model on 8 UCI regression datasets and image classification datasets including MNIST, Fashion-MNIST and CIFAR-10, which outperforms state-of-the-art approximation methods. Additional ablation study demonstrates that our method is superior in preventing overfitting.

Our code is publicly available at https://github.com/studying910/NOVI-DGP.

2 Background and Related Work

In this section, we first present necessary notations and settings on single-layer Gaussian Processes (GPs) and Deep Gaussian Processes (DGPs), then we point out the flaws of current model and introduce our motivation.

2.1 Preliminary on Gaussian Processes

Let a random function $f : \mathbb{R}^D \to \mathbb{R}$ map N training inputs $\boldsymbol{X} = \{\boldsymbol{x}_n\}_{n=1}^N$ to a collection of noisy observed outputs $\boldsymbol{y} = \{y_n\}_{n=1}^N$. In general, a zero mean GP prior is imposed on the function f, i.e., $f \sim \mathcal{GP}(0, k)$ where k represents a covariance function $k : \mathbb{R}^D \times \mathbb{R}^D \to \mathbb{R}$. Let $\mathbf{f} = \{f(\boldsymbol{x}_n)\}_{n=1}^N$ represent the latent function values at the inputs $\boldsymbol{X}$. This assumption yields a multivariate Gaussian prior over the function values $p(\mathbf{f}) = \mathcal{N}(\mathbf{f}|0, \boldsymbol{K}_{\boldsymbol{XX}})$ where the i-th row and j-th column element of the covariance matrix $[\boldsymbol{K}_{\boldsymbol{XX}}]_{ij} = k(\boldsymbol{x}_i, \boldsymbol{x}_j)$. In this work, we suppose $\boldsymbol{y}$ is contaminated by an i.i.d noise, thus $p(\boldsymbol{y}|\mathbf{f}) = \mathcal{N}(\boldsymbol{y}|\mathbf{f}, \sigma^2\boldsymbol{I})$ where σ^2 is the noise variance. Figure 1(a) is a graphical illustration of a typical GP . The GP posterior of the latent output $p(\mathbf{f}|\boldsymbol{y})$ has a closed-form solution Rasmussen & Williams (2006) but suffers from $\mathcal{O}(N^3)$ computational cost and $\mathcal{O}(N^2)$ storage requirement, thus limiting its scalability to big data.

Advanced sparse methods have been developed to set so-called *inducing points* $\boldsymbol{Z} = \{\boldsymbol{z}_m\}_{m=1}^M$ ($M \ll N$) from the input space and the associated inducing outputs known as *inducing variables*: $\mathbf{u} = \{\mathrm{u}_m = f(\boldsymbol{z}_m)\}_{m=1}^M$ Titsias (2009a); Snelson & Gharahmani (2005); Quiñonero-Candela & Rasmussen (2005), with a time complexity of $\mathcal{O}(NM^2)$. In this *Sparse GPs* (SGPs) paradigm, as shown in Figure 1(b), *inducing variables* $\mathbf{u}$ share a joint multivariate Gaussian distribution with $\mathbf{f}$: $p(\mathbf{f}, \mathbf{u}) = p(\mathbf{f}|\mathbf{u})p(\mathbf{u})$ where the condition is specified as:

$$p(\mathbf{f}|\mathbf{u}) = \mathcal{N}(\boldsymbol{K}_{\boldsymbol{XZ}}\boldsymbol{K}_{\boldsymbol{ZZ}}^{-1}\mathbf{u}, \boldsymbol{K}_{\boldsymbol{XX}} - \boldsymbol{K}_{\boldsymbol{XZ}}\boldsymbol{K}_{\boldsymbol{ZZ}}^{-1}\boldsymbol{K}_{\boldsymbol{ZX}}) \tag{1}$$

and $p(\mathbf{u}) = \mathcal{N}(\mathbf{u}|0, \boldsymbol{K}_{\boldsymbol{ZZ}})$ is the prior over the inducing outputs.

In SGPs the posterior distribution of the latent function $p(\mathbf{f}|\boldsymbol{y})$ is intractable since $\mathbf{u}$ need be marginalized, resulting in the following integral:

$$p(\mathbf{f}|\boldsymbol{y}) = \int p(\mathbf{f}|\mathbf{u})p(\mathbf{u}|\boldsymbol{y})d\mathbf{u} \tag{2}$$

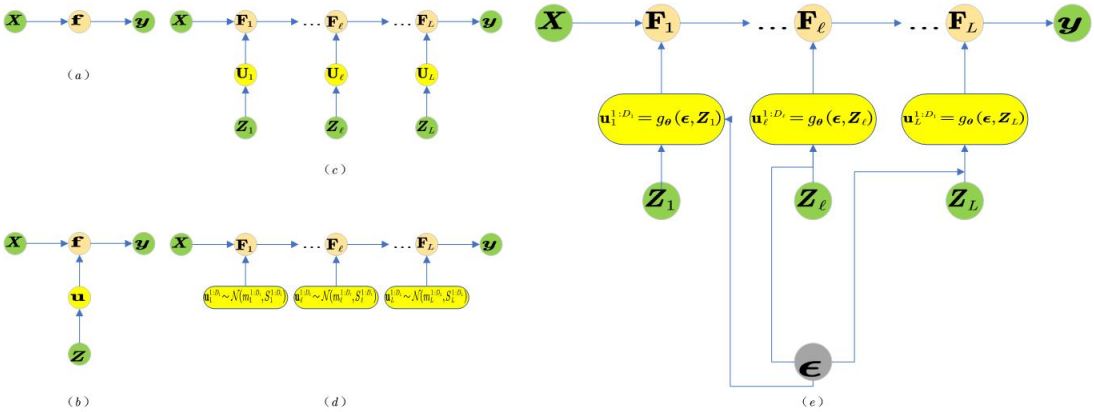

Figure 1: Graphical illustrations of (a) GP, (b) SGP, (c) DGP, (d) DSVI and (e) our model.

which is difficult to solve since the posterior distribution of inducing variables $p(\mathbf{u}|\mathbf{y})$ is intractable.

As an elegant approach, *Sparse variational GPs* (SVGPs) Titsias (2009a); Hensman et al. (2015) reformulate the posterior inference problem as variational inference (VI) and confine the variational distribution to be $q(\mathbf{f},\mathbf{u}) = p(\mathbf{f}|\mathbf{u})q(\mathbf{u})$ Hensman et al. (2013); Titsias (2009a); Gal et al. (2014); Salimbeni & Deisenroth (2017). Then the evidence lower bound (ELBO) Hoffman et al. (2013) for minimizing Kullback-Leibler (KL) divergence between the variational posterior distribution $q(\mathbf{f},\mathbf{u})$ and the joint posterior distribution $p(\mathbf{f},\mathbf{u}|\mathbf{y})$ is simplified as:

$$\text{ELBO} = \mathbb{E}_{q(\mathbf{f},\mathbf{u})}[\log p(\boldsymbol{y},\mathbf{f},\mathbf{u}) - \log q(\mathbf{f},\mathbf{u})] = \mathbb{E}_{q(\mathbf{f})}[\log p(\boldsymbol{y}|\mathbf{f})] - \text{KL}[q(\mathbf{u})\,\|p(\mathbf{u})]\,, \tag{3}$$

where $q(\mathbf{f}) = \int p(\mathbf{f}|\mathbf{u})\, q(\mathbf{u})\, \mathrm{d}\mathbf{u}$. The current method for SVGP is to approximate $q(\mathbf{u}) = \mathcal{N}(\boldsymbol{m},\boldsymbol{S})$ Hensman et al. (2015); Deisenroth & Ng (2015); Gal et al. (2014); Hensman et al. (2013); Hoang et al. (2015; 2016); Titsias (2009b) to produce a Gaussian marginal[1].

2.2 Deep Gaussian Processes

A multi-layer DGP model is a hierarchical composition of GP models constructed by stacking the muti-output SGPs together Damianou & Lawrence (2013), as shown in Figure 1(c). Consider a model with L layers and D_l independent random functions in layer $\ell = 1,\ldots,L$ such that output of the ℓ-1-th layer $\mathbf{F}_{\ell-1}$ is used as an input to the ℓ-th layer, i.e., $\mathbf{F}_\ell = \{\mathbf{f}_{\ell,1} = f_{\ell,1}(\mathbf{F}_{\ell-1}),\cdots,\mathbf{f}_{\ell,D_\ell} = f_{\ell,D_\ell}(\mathbf{F}_{\ell-1})\}$, where $f_{\ell,d} \sim \mathcal{GP}(0,k_\ell)$ for $d = 1,\ldots,D_\ell$ and $\mathbf{F}_0 \triangleq \boldsymbol{X}$. The inducing points and corresponding inducing variables for DGP layers are denoted by $\mathbf{Z} = \{\boldsymbol{Z}_\ell\}_{\ell=1}^L$ and $\mathbf{U} = \{\mathbf{U}_\ell\}_{\ell=1}^L$ respectively where $\mathbf{U}_\ell = \{\mathbf{u}_{\ell,1} = f_{\ell,1}(\boldsymbol{Z}_\ell),\cdots,\mathbf{u}_{\ell,D_\ell} = f_{\ell,D_\ell}(\boldsymbol{Z}_\ell)\}$. Let $\mathbf{F} = \{\mathbf{F}_\ell\}_{\ell=1}^L$, the DGP model design yields the following joint model density:

$$p(\boldsymbol{y},\mathbf{F},\mathbf{U}) = p(\boldsymbol{y}|\mathbf{F}_L)\prod_{\ell=1}^{L} p(\mathbf{F}_\ell|\mathbf{F}_{\ell-1},\mathbf{U}_\ell)p(\mathbf{U})\,. \tag{4}$$

Here we place independent GP priors within and across layers on $\mathbf{U}$: $p(\mathbf{U}) = \prod_{l=1}^L p(\mathbf{U}_l) = \prod_{l=1}^L\prod_{d=1}^{D_\ell}\mathcal{N}(\mathbf{u}_{\ell,d}|0,\boldsymbol{K}_{\boldsymbol{Z}_\ell\boldsymbol{Z}_\ell})$ and in the same way as Equation (1), the condition is defined as:

$$p(\mathbf{F}_\ell|\mathbf{F}_{\ell-1},\mathbf{U}_\ell) = \prod_{d=1}^{D_\ell}\mathcal{N}(\mathbf{f}_{\ell,d}|\boldsymbol{K}_{\mathbf{F}_{\ell-1}\boldsymbol{Z}_\ell}\boldsymbol{K}_{\boldsymbol{Z}_\ell\boldsymbol{Z}_\ell}^{-1}\mathbf{u}_{\ell,d},\boldsymbol{K}_{\mathbf{F}_{\ell-1}\mathbf{F}_{\ell-1}} - \boldsymbol{K}_{\mathbf{F}_{\ell-1}\boldsymbol{Z}_\ell}\boldsymbol{K}_{\boldsymbol{Z}_\ell\boldsymbol{Z}_\ell}^{-1}\boldsymbol{K}_{\boldsymbol{Z}_\ell\mathbf{F}_{\ell-1}}). \tag{5}$$

As an extension of Variational Inference with DGPs, DSVI Salimbeni & Deisenroth (2017) approximates the posterior by requiring the distribution across the inducing outputs to be a-posteriori Gaussian and independent amongst distinct GPs (known as the mean-field assumption Opper & Saad

[1]The solution is given in App. A

(2001); Hoffman et al. (2013), see also in Figure 1(d)), $q\left(\mathbf{u}_{\ell,1:D_\ell}\right)=\mathcal{N}\left(\boldsymbol{m}_{\ell,1:D_\ell}, \boldsymbol{S}_{\ell,1:D_\ell}\right)$, where $\boldsymbol{m}_{\ell,1:D_\ell}$ and $\boldsymbol{S}_{\ell,1:D_\ell}$ are variational parameters and the ELBO can be simplified to:

$$\text{ELBO}=\mathbb{E}_{q(\mathbf{F}_L)}\left[\log p\left(\boldsymbol{y}|\mathbf{F}_L\right)\right]-\text{KL}\left[q\left(\mathbf{U}\right)\parallel p\left(\mathbf{U}\right)\right]. \tag{6}$$

Since the ELBO decomposes across the data points, the samples acquired from $q\left(\mathbf{F}_L\right)$ can be used to evaluate Equation (6) through minibatches. By iteratively sampling the layer outputs and utilizing the reparameterisation trick Kingma & Welling (2013), DSVI enables scalability to big datasets.

As mentioned in Section 1, while the mean-field Gaussian assumptions of the variational posterior $q(\mathcal{U})$ makes it simple to analytically marginalise out the inducing outputs, these assumptions impose overly stringent constraints, potentially limiting the expressiveness and effectiveness of such deterministic approximation approaches for DGP models. In particular, according to Bayesian formula, the true posterior distrbution can be written as:

$$p\left(\mathbf{U}|\boldsymbol{y}\right)=\frac{p\left(\mathbf{U}\right)p\left(\boldsymbol{y}|\mathbf{U}\right)}{p\left(\boldsymbol{y}\right)}=\frac{\int p\left(\boldsymbol{y},\mathbf{F},\mathbf{U}\right)d\mathbf{F}}{p\left(\boldsymbol{y}\right)} \tag{7}$$

Due to the fact that the latent functions $\mathbf{F}_1,\cdots,\mathbf{F}_{L-1}$ are inputs to the non-linear kernel function, the likelihood term $p\left(\boldsymbol{y}|\mathbf{U}\right)$ in Equation (7) is intractable and $p\left(\mathbf{U}|\boldsymbol{y}\right)$ is often non-Gaussian in reality. Moreover, the true posterior belief often demonstrates a substantial degree of correlation across the DGP layers Havasi et al. (2018).

To address this issue, we present a new variational family that provides both efficient computation and expressiveness based on *Operator Variational Inference* (OVI) Ranganath et al. (2016), while simultaneously learning preservable transformations and generating unbiased posterior samples constructed by neural networks, as detailed in Section 3 and Section 4.

3 OVI AND STEIN DISCREPANCY

Before using OVI and Stein Discrepancy to develop a unique inference strategy for DGP model, we provide a quick introduction to these concepts that forms the foundation of our method.

Definition 1. *Let $p(\boldsymbol{x})$ be a probability density supported on $\mathcal{X}\subseteq\mathbb{R}^d$ and $\boldsymbol{\phi}:\mathcal{X}\to\mathbb{R}^d$ be a differentiable function, we define Langevin-Stein Operator (LSO) Ranganath et al. (2016) as:*

$$\mathcal{A}_p\boldsymbol{\phi}(\boldsymbol{x})\triangleq\nabla_{\boldsymbol{x}}\log p(\boldsymbol{x})^T\boldsymbol{\phi}(\boldsymbol{x})+\text{Tr}(\nabla_{\boldsymbol{x}}\boldsymbol{\phi}(\boldsymbol{x})). \tag{8}$$

Definition 2. *(Stein's Discrepancy) Hu et al. (2018); Grathwohl et al. (2020); di Langosco et al. (2021) Let $p(\boldsymbol{x})$, $q(\boldsymbol{x})$ be probability densities supported on $\mathcal{X}\subseteq\mathbb{R}^d$. Stein discrepancy is defined by considering the maximum violation of Stein identity for $\boldsymbol{\phi}$ in some proper function set $\mathcal{F}$*

$$\mathcal{S}\left(q,p\right)\triangleq\sup_{\boldsymbol{\phi}\in\mathcal{F}}\mathbb{E}_{\boldsymbol{x}\sim q}[\mathcal{A}_p\boldsymbol{\phi}(\boldsymbol{x})]. \tag{9}$$

For example, if $\mathcal{F}$ is taken to be a ball in a *Reproducing Kernel Hilbert Space* (RKHS) related to kernel $k\left(\cdot,\cdot\right)$,Liu et al. (2016) showed that the supremum in Equation (9) has a closed form solution [2].

$$\text{KSD}\left(q,p\right)\triangleq\max_{\boldsymbol{\phi}\in\mathcal{H}^d}\left\{\mathbb{E}_{x\sim q}\left[\mathcal{A}_p\boldsymbol{\phi}\left(\boldsymbol{x}\right)\right]\right\}\qquad s.t.\qquad\|\boldsymbol{\phi}\|_{\mathcal{H}^d}\leqslant 1. \tag{10}$$

Kernelized Stein Discrepancy (KSD) Hu et al. (2018); Liu et al. (2016); Gorham & Mackey (2017); Jitkrittum et al. (2017) has been employed to learn implicit samplers for unnormalized densities. Nevertheless, when dimension increases, the ineffectiveness of this method becomes evidentRamdas et al. (2015). Instead of an unit-ball in RKHS, previous methods extend the function space $\mathcal{F}$ in Stein discrepancy (9) to be the $\mathcal{L}_2$ space Hu et al. (2018); Grathwohl et al. (2020) and parameterize $\boldsymbol{\phi}$ with a neural network $\boldsymbol{\phi}_{\boldsymbol{\eta}}$ as a discriminator

$$\text{LSD}\left(q,p;\boldsymbol{\eta}\right)\triangleq\max_{\boldsymbol{\eta}}\{\mathbb{E}_{\boldsymbol{x}\sim q}[\nabla_{\boldsymbol{x}}\log p\left(\boldsymbol{x}\right)^T\boldsymbol{\phi}_{\boldsymbol{\eta}}\left(\boldsymbol{x}\right)+\text{Tr}\left(\nabla_{\boldsymbol{x}}\boldsymbol{\phi}_{\boldsymbol{\eta}}\left(\boldsymbol{x}\right)\right)]\}. \tag{11}$$

which is referred to the Learned Stein Discrepancy (LSD)Grathwohl et al. (2020). Neural networks as functions are not by definition square integrable, as they do not by default disappear at infinity.

[2]The optimal solution of Equation (10) has been shown in Oates et al. (2017); Liu et al. (2016); Chwialkowski et al. (2016); Liu & Wang (2016) to $\boldsymbol{\phi}=\boldsymbol{\phi}^*_{q,p}/\left\|\boldsymbol{\phi}^*_{q,p}\right\|_{\mathcal{H}^d}$, where $\boldsymbol{\phi}^*_{q,p}\left(\cdot\right)=\mathbb{E}_{\boldsymbol{x}\sim q}\left[\mathcal{A}_pk\left(x,\cdot\right)\right]$

To enforce the $\mathcal{L}_2$ constraint, an $\mathcal{L}_2$ regularizer with strength $\lambda \in \mathbb{R}^+$ is applied to LSD to gain a Regularized Stein Discrepancy (RSD)

$$\mathrm{RSD}\left(q, p; \boldsymbol{\eta}\right) \triangleq \max_{\boldsymbol{\eta}} \{\mathbb{E}_{\boldsymbol{x}\sim q}[\nabla_{\boldsymbol{x}} \log p\left(\boldsymbol{x}\right)^T \boldsymbol{\phi}_{\boldsymbol{\eta}}\left(\boldsymbol{x}\right) + \mathrm{Tr}\left(\nabla_{\boldsymbol{x}} \boldsymbol{\phi}_{\boldsymbol{\eta}}\left(\boldsymbol{x}\right)\right)] - \lambda \mathbb{E}_{\boldsymbol{x}\sim q}[\boldsymbol{\phi}_{\boldsymbol{\eta}}(\boldsymbol{x})^T \boldsymbol{\phi}_{\boldsymbol{\eta}}\left(\boldsymbol{x}\right)]\}. \tag{12}$$

In Bayesian posterior inference, we take p and $q_{\boldsymbol{\theta}}$ as the true posterior and approximate posterior, respectively, where $\boldsymbol{\theta} \in \Theta$ and Θ is a set of variational parameters. Stein divergence in Equation (9) is usually used as an objictive of OVIRanganath et al. (2016), which is a black-box algorithm that uses operators for optimizing any operator objective with data subsampling and a wider class of posterior approximations that does not require a tractable density. Given parameterizations of the variational family Θ and the discriminator $\boldsymbol{\phi}_{\boldsymbol{\eta}}$, OVI seeks to solve a minimax problem

$$\theta^\star = \arg\inf_{\boldsymbol{\theta}\in\Theta} \sup_{\boldsymbol{\eta}} \mathbb{E}_{\boldsymbol{x}\sim q_\theta}\left[\mathcal{A}_p \boldsymbol{\phi}_{\boldsymbol{\eta}}\left(\boldsymbol{x}\right)\right]. \tag{13}$$

4 Deep Gaussian Processes with Neural Operator Variational Inference

Now we will discuss the algorithm design for the Bayesian inference problem of sampling the posterior $p\left(\mathbf{U}|\mathcal{D}\right)$ for DGPs. For consistency, we continue to use notation in Section 2.2. Let $\mathcal{D} = \{\boldsymbol{x}_n, y_n\}_{n=1}^N$ represent the training dataset, $\mathbf{U} \triangleq \{\mathbf{U}_\ell\}_{\ell=1}^L$ represent inducing variables and $\boldsymbol{\nu}$ reprensent the DGP model hyperparameters including inducing points locations, kernel hyperparameters and noise variance.

4.1 Neural network as Generator

Let $q_0(\boldsymbol{\epsilon})$ be the reference distribution that generates noise $\boldsymbol{\epsilon} \in \mathbb{R}^{d_0}$. Let $g_{\boldsymbol{\theta}}$ represent our sampler, which is a black-box generator parameterized by a multi-layer neural network. Let $q_{\boldsymbol{\theta}}\left(\mathbf{U}\right)$ be the underlying density of the generated samples $\mathbf{U} = g_{\boldsymbol{\theta}}(\boldsymbol{\epsilon})$. In summary, our setup is as follows:

$$\boldsymbol{\epsilon} \sim q_0\left(\boldsymbol{\epsilon}\right), \qquad g_{\boldsymbol{\theta}}\left(\boldsymbol{\epsilon}\right) = \mathbf{U} \sim q_{\boldsymbol{\theta}}\left(\mathbf{U}\right)$$

Neural networks as a generator have a high capacity and can well approximate almost any distribution by transforming simple ones such as Gaussian or uniform distribution with many applications in deep generative models Huszár (2017); Mescheder et al. (2017); Titsias & Ruiz (2019); Cybenko (1989); Lu & Lu (2020); Perekrestenko et al. (2020); Yang et al. (2022). Since the generative distribution $q_{\boldsymbol{\theta}}\left(\mathbf{U}\right)$ is implict, KL divergence is not applicable as a measure between $q_{\boldsymbol{\theta}}\left(\mathbf{U}\right)$ and the true posterior $p(\mathbf{U}|\mathcal{D})$ in this case. Therefore, it is reasonable to use OVI and RSD to construct a better objective.

4.2 Training Schedule

In Section 3 we have developed OVI, a method using Langevin-Stein operator and allowing for a more flexible representation of the posterior gemoetry beyond the commonly used Gaussian distribution used in vanilla VI. We extend it to applications in inducing points posterior inference for DGP model by learning the parameters of neural network generator to best fit the data. Since our discriminator $\boldsymbol{\phi}_{\boldsymbol{\eta}}$ is sufficiently expressive, we produce an objective whose expectation[3] is 0 if and only if the true posterior $p(\mathbf{U}|\mathcal{D})$ and the approximate distribution $q(\mathbf{U})$ are equivalent. During training, we will minimize

$$\mathcal{L}(\boldsymbol{\theta}, \boldsymbol{\nu}) = \mathrm{RSD}\left(q_{\boldsymbol{\theta}}(\mathbf{U}), p(\mathbf{U}|\mathcal{D}, \boldsymbol{\nu}); \boldsymbol{\phi}_{\boldsymbol{\eta}}\right) \tag{14}$$

with respect to $\boldsymbol{\theta}$ and jointly optimise the model hyperparameters $\boldsymbol{\nu}$ by maximizing log-likelihood via Monte Carlo sampling. However, this procedure is difficult due to the supremum on r.h.s. of Equation (12). In order to obtain the optimized network parameters $\boldsymbol{\theta}$, we iteratively update the generator $g_{\boldsymbol{\theta}}$ and the discriminator $\boldsymbol{\phi}_{\boldsymbol{\eta}}$ in an alternating manner where the discriminator is trained to more accurately estimate the Stein Discrepancy and the generator is trained to minimize the estimation of the discrepancy. The proposed training algorithm is summarized in Algorithm 1 which we refer to it as *Neural Operator Variational Inference* (NOVI) for DGP.

[3]The expectation doesn't include the regularization term.

Algorithm 1: NOVI for DGP

Input: training data $\mathcal{D} = \{\boldsymbol{x}_n, y_n\}_{n=1}^N$, penalty parameter λ, n_c number of iterations for training the critic, learning rate α, β, γ, M batch size, sample number K
Initialize discriminator $\boldsymbol{\eta}$, generator $\boldsymbol{\theta}$, DGP hyperparameters $\boldsymbol{\nu}$
repeat
  for $j = 1$ **to** n_c **do**
    Sample a minibatch $\{\boldsymbol{x}_i, y_i\}_{i=1}^M \sim \mathcal{D}$
    Generate i.i.d. noise inputs $\boldsymbol{\epsilon}_1 \dots \boldsymbol{\epsilon}_K$ from q_0
    Obtain fake sample $g_{\boldsymbol{\theta}}(\boldsymbol{\epsilon}_1) \dots g_{\boldsymbol{\theta}}(\boldsymbol{\epsilon}_K)$
    Compute empirical loss $\widehat{\mathrm{RSD}}(q_{\boldsymbol{\theta}}, p; \boldsymbol{\phi}_{\boldsymbol{\eta}})$
    $\boldsymbol{\eta} \leftarrow \boldsymbol{\eta} - \alpha \nabla_{\eta} \widehat{\mathrm{RSD}}(q_{\boldsymbol{\theta}}, p; \boldsymbol{\phi}_{\boldsymbol{\eta}})$
  end for
  Compute empirical loss $\widehat{\mathcal{L}}(\boldsymbol{\theta}, \boldsymbol{\nu})$
  $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} - \beta \ \nabla_{\boldsymbol{\theta}} \widehat{\mathcal{L}}(\boldsymbol{\theta}, \boldsymbol{\nu})$
  $\boldsymbol{\nu} \leftarrow \boldsymbol{\nu} - \gamma \ \frac{1}{K} \sum_{k=1}^K \nabla_{\boldsymbol{\nu}} \log p(\boldsymbol{y}, \mathbf{U}^k | \boldsymbol{\nu})$
until $\boldsymbol{\theta}, \boldsymbol{\nu}$ converge

In our implementation, we utilize Monte Carlo method to estimate the objective (14) and RSD(12):

$$\begin{aligned}\widehat{\mathrm{RSD}}(q_{\boldsymbol{\theta}}, p; \boldsymbol{\phi}_{\boldsymbol{\eta}}) = \frac{1}{K} \sum\nolimits_{k=1}^K & \left(\nabla_{\mathbf{U}} \log p(\mathbf{U}|\mathcal{D}, \boldsymbol{\nu})^T|_{\mathbf{U}=\mathbf{U}^k} \boldsymbol{\phi}_{\boldsymbol{\eta}}(\mathbf{U}^k) + \mathbb{E}_{\boldsymbol{\omega} \sim \mathcal{N}(0, \boldsymbol{I})}(\boldsymbol{\omega}^T \nabla_{\mathbf{U}} \boldsymbol{\phi}_{\boldsymbol{\eta}}(\mathbf{U})|_{\mathbf{U}=\mathbf{U}^k} \boldsymbol{\omega})\right) \\ & - \lambda \frac{1}{K} \sum\nolimits_{k=1}^K \left(\boldsymbol{\phi}_{\boldsymbol{\eta}}(\mathbf{U}^k)^T \boldsymbol{\phi}_{\boldsymbol{\eta}}(\mathbf{U}^k)\right) \\ \widehat{\mathcal{L}}(\boldsymbol{\theta}, \boldsymbol{\nu}) = & \ \widehat{\mathrm{RSD}}(q_{\boldsymbol{\theta}}, p; \boldsymbol{\phi}_{\boldsymbol{\eta}^\star}),\end{aligned} \tag{15}$$

where $\boldsymbol{\phi}_{\boldsymbol{\eta}^\star}$ is the supremum of RSD estimate and the gradient with $\boldsymbol{\theta}$ and $\boldsymbol{\nu}$ is computed via automatic differentiation. We use Hutchinson estimator Hutchinson (1989) to compute the expensive divergence of $\boldsymbol{\phi}_{\boldsymbol{\eta}}$ in Equation (15), which is a simple yet effective way to obtain a stochastic estimate of the trace of a matrix. It can reduce the time complexity from $\mathcal{O}\left(D^2\right)$ to $\mathcal{O}(D)$ where D is the dimensionality of the matrix. In Theorem 1, we prove that the score function $\nabla_{\mathbf{U}} \log p(\mathbf{U}|\mathcal{D}, \boldsymbol{\nu})$ can be evaluated by Monte Carlo method, which shows that RSD can be utilized as a reasonable objective to update the parameters of the generator network.

Theorem 1. *The score function $\nabla_{\mathbf{U}} \log p(\mathbf{U}|\mathcal{D}, \boldsymbol{\nu})$ in Equation (15) can be evaluated by Monte Carlo sampling (detailed proof can be seen in App. B):*

$$\nabla_{\mathbf{U}} \log p(\mathbf{U}|\mathcal{D}, \boldsymbol{\nu}) = -\left(\boldsymbol{\Delta}_1, \dots, \boldsymbol{\Delta}_\ell, \dots, \boldsymbol{\Delta}_L\right) + \nabla_{\mathbf{U}} \log \sum\nolimits_{s=1}^S p(\boldsymbol{y}|\widehat{\mathbf{F}}_L^{(s)}) \tag{16}$$

where $\boldsymbol{\Delta}_\ell = (\boldsymbol{K}_{\boldsymbol{Z}_\ell \boldsymbol{Z}_\ell}^{-1} \mathbf{u}_{\ell,1}, ..., \boldsymbol{K}_{\boldsymbol{Z}_\ell \boldsymbol{Z}_\ell}^{-1} \mathbf{u}_{\ell,d}, ..., \boldsymbol{K}_{\boldsymbol{Z}_\ell \boldsymbol{Z}_\ell}^{-1} \mathbf{u}_{\ell,D_\ell})$ *and* $\widehat{\mathbf{f}}_{\ell,d}^{(s)} \sim \mathcal{N}(\boldsymbol{K}_{\widehat{\mathbf{F}}_{\ell-1} \boldsymbol{Z}_\ell} \boldsymbol{K}_{\boldsymbol{Z}_\ell \boldsymbol{Z}_\ell}^{-1} \mathbf{u}_{\ell,d}, \boldsymbol{K}_{\widehat{\mathbf{F}}_{\ell-1} \widehat{\mathbf{F}}_{\ell-1}} - \boldsymbol{K}_{\widehat{\mathbf{F}}_{\ell-1} \boldsymbol{Z}_\ell} \boldsymbol{K}_{\boldsymbol{Z}_\ell \boldsymbol{Z}_\ell}^{-1} \boldsymbol{K}_{\boldsymbol{Z}_\ell \widehat{\mathbf{F}}_{\ell-1}})$ *for* $\ell = 1, \dots, L$, S *is the number of samples involved in estimation.*

The derivation of a more concise approximation for regression task is also detailed in App. B.

4.3 Prediction

Let $\mathcal{D}^\star = \{\boldsymbol{x}_n^\star, y_n^\star\}_{n=1}^T$ be the test data, to predict its value, we sample from the optimized generator and convert the input locations $\boldsymbol{x}$ to the test location $\boldsymbol{x}^\star$ in formula. We denote the function values at the test location as $\mathbf{F}_\ell^\star$. To obtain the final layer density we use

$$q(\mathbf{F}_L^\star) = \int \prod\nolimits_{\ell=1}^L \prod\nolimits_{d=1}^{D_\ell} p(\mathbf{f}_{\ell,d}^\star | \mathbf{F}_{\ell-1}^\star, \mathbf{u}_{\ell,d}) q_{\boldsymbol{\theta}^\star}(\mathbf{u}_{\ell,d}) \, d\mathbf{F}_{\ell-1}^\star d\mathbf{u}_{\ell,d} \tag{17}$$

where $\boldsymbol{\theta}^\star$ is the optimal of the generator and the first term of the integral $p(\mathbf{f}_{\ell,d}^\star | \mathbf{F}_{\ell-1}^\star, \mathbf{u}_{\ell,d})$ is conditional Gaussian. We leverage this consequence to draw samples from $q(\mathbf{F}_L^\star)$, and further perform the sampling using re-parameterization trick Salimbeni & Deisenroth (2017); Rezende et al.

(2014); Kingma et al. (2015). Specifically, we first sample $\boldsymbol{\epsilon}^{\ell} \sim \mathcal{N}(0, \boldsymbol{I}_{D^{\ell}})$ and $\mathbf{U} \sim q_{\boldsymbol{\theta}^{\star}}(\mathbf{U})$, then recursively draw the sampled variables $\widehat{\mathbf{f}}^{\star}_{\ell,d} \sim p(\mathbf{f}^{\star}_{\ell,d}|\widehat{\mathbf{F}}^{\star}_{\ell-1}, \mathbf{u}_{\ell,d})$ for $\ell = 1, \ldots, L$ as:

$$\widehat{\mathbf{f}}^{\star}_{\ell,d} = \boldsymbol{K}_{\widehat{\mathbf{F}}^{\star}_{\ell-1}\boldsymbol{Z}_{\ell}}\boldsymbol{K}^{-1}_{\boldsymbol{Z}_{\ell}\boldsymbol{Z}_{\ell}}\mathbf{u}_{\ell,d} + \boldsymbol{\epsilon}_{\ell} \odot \sqrt{\operatorname{diag}\left(\boldsymbol{K}_{\widehat{\mathbf{F}}^{\star}_{\ell-1}\widehat{\mathbf{F}}^{\star}_{\ell-1}} - \boldsymbol{K}_{\widehat{\mathbf{F}}^{\star}_{\ell-1}\boldsymbol{Z}_{\ell}}\boldsymbol{K}^{-1}_{\boldsymbol{Z}_{\ell}\boldsymbol{Z}_{\ell}}\boldsymbol{K}_{\boldsymbol{Z}_{\ell}\widehat{\mathbf{F}}^{\star}_{\ell-1}}\right)}, \tag{18}$$

where the square root is element-wise. We define $\mathbf{F}^{\star}_{0} \triangleq \boldsymbol{X}^{\star}$ for the first layer and use $\operatorname{diag}(\cdot)$ to denote the vector of diagonal elements of a matrix. The diagonal approximation in Equation (18) holds since in DGP model, the i-th marginal of approximate posterior $q(\mathbf{f}_{(\ell,d)[i]})$ depends only on the corresponding inputs $\boldsymbol{x}_i$ Quiñonero-Candela & Rasmussen (2005). As illustrated in Figure 1(e), in our experiment, we concatenate $\boldsymbol{Z}_{\ell}$ and $\boldsymbol{\epsilon}$ to generate $\mathbf{U}$ to avoid overfitting Yu et al. (2019).

5 Convergence Guarantees

In this section, we provide convergence guarantees for NOVI DGP, detailed proof can be seen in App. C.

Definition 3. *The Fisher divergence Sriperumbudur et al. (2017) between two suitably smooth density functions is defined as*

$$F(q, p) = \int_{\mathbb{R}^d} \|\nabla \log q(\boldsymbol{x}) - \nabla \log p(\boldsymbol{x})\|_2^2 \, q(\boldsymbol{x}) \, d\boldsymbol{x}.$$

Theorem 2. *Supposed that the discriminator and the generator network has enough capacity. Training the generator with the optimal discriminator corresponds to minimizing the fisher divergence between p_{θ} and q. The corresponding optimal loss is*

$$\mathcal{L}(\boldsymbol{\theta}, \boldsymbol{\nu}) = \frac{1}{4\lambda} F\left(q_{\boldsymbol{\theta}}(\mathbf{U}), p(\mathbf{U}|\mathcal{D}, \boldsymbol{\nu})\right)$$

Theorem 3. *The bias of the estimation for prediction $\widehat{\mathbf{F}}^{\star}_{L}$ in Equation (18) from the DGPs exact evaluation can be bounded by the square root of the Fisher divergence between $q_{\boldsymbol{\theta}}(\mathbf{U})$ and $p(\mathbf{U}|\mathcal{D}, \boldsymbol{\nu})$ up to multiplying a constant.*

6 Experiments

We empirically evaluate and compare the performance of our method with doubly stochastic VI (DSVI) (Salimbeni & Deisenroth, 2017) for DGPs, which is implemented as our baseline model, and state-of-the-art SGHMC model Havasi et al. (2018) using real-world datasets in regression and classification tasks both in small and large data regimes. All our experiments were run with exactly the same hyper-parameters and initializations. Detailed training information can be seen in App. E.

6.1 UCI Regression Benchmark

Our experiments are conducted on 8 UCI regression datasets with size ranging from 308 to 45730. The performance metric is average RMSE of the test data. The results are shown in Figure 2 (tabular version can be seen in App. D.3). On four of the eight datasets, simply using 2-layer NOVI model achieves the best result and a huge performance gap against other two methods. On larger datasets, like 'Power', 'Concrete', 'Qsar' and 'Protein', the deepest NOVI model outperforms other methods. We attribute this phenomenon to the overfitting of the deep model on small data sets. Additional results for real-world regression datasets can be seen in App. D.5.

6.2 Image Classification

We apply our method to MNIST LeCun et al. (1998), Fashion-MNIST Xiao et al. (2017) and CIFAR-10 Krizhevsky et al. (2009) multiclass classification problem. Both MNIST and Fashion-MNIST datasets are grey-scale images of 28×28 pixels. The CIFAR-10 dataset consists of colored images of 32×32 pixels. Results are shown in Table 1 [4]. For all three datasets, NOVI outperforms other two methods with significant less training time and iterations. We also perform experiments using three UCI classification datasets and present results in App. D.1.

[4] For 3-layer DSVI and SGHMC models, since they have not yet released the corresponding code to reproduce it, we only test the training time and report its iterations according to the original paper.

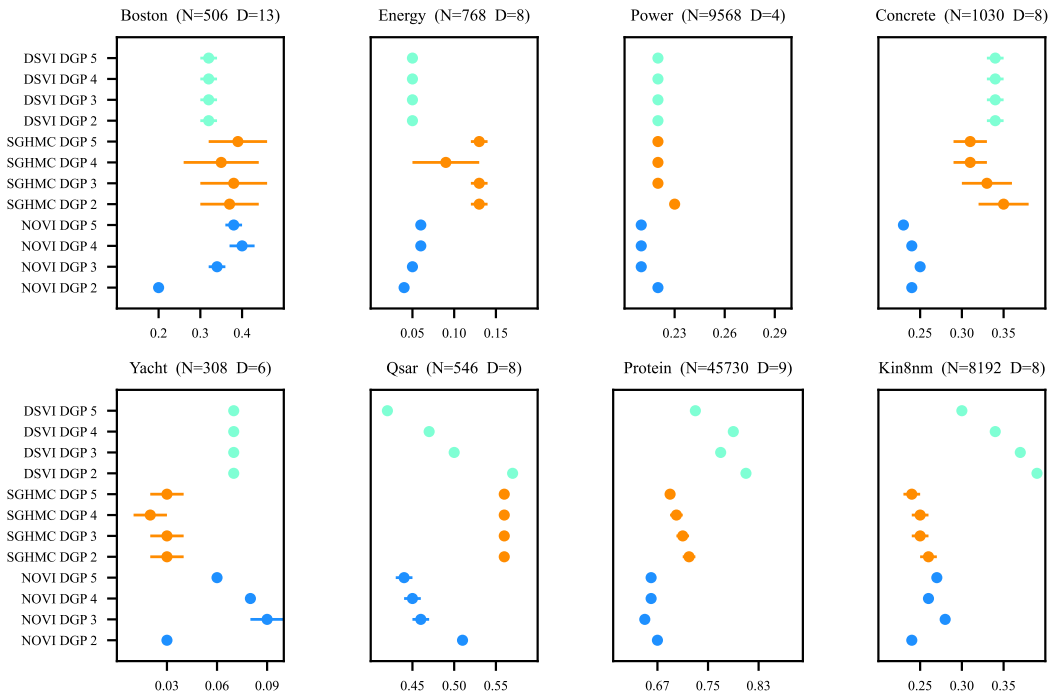

Figure 2: Regression mean test RMSE results by our NOVI method (blue), SGHMC (orange) and DSVI (cyan) for DGPs on UCI benchmark datasets. Lower is better. The mean is shown with error bars of one standard error.

Data Set	Model	Time (L=3)	Iter(L=3)	Acc (L=3)	Time (L=4)	Iter (L=4)	Acc (L=4)
MNIST	DSVI	0.34s/iter	20K	-	0.54s/iter	20K	97.41
	SGHMC	1.14s/iter	20K	-	1.22s/iter	20K	97.55
	NOVI (ours)	0.38s/iter	10K	**97.94**	0.50s/iter	10K	**98.01**
Fashion-MNIST	DSVI	0.34s/iter	20K	-	0.50s/iter	20K	87.99
	SGHMC	1.21s/iter	20K	-	1.25s/iter	20K	87.08
	NOVI (ours)	0.40s/iter	10K	**88.96**	0.55s/iter	10K	**89.15**
CIFAR-10	DSVI	0.43s/iter	20K	-	0.66s/iter	20K	51.79
	SGHMC	8.04s/iter	20K	-	8.61s/iter	20K	52.81
	NOVI (ours)	0.43s/iter	10K	**53.32**	0.52s/iter	10K	**53.42**

Table 1: Mean test accuracy (%) and training details achieved by DSVI, SGHMC and NOVI (ours) DGP model for three image classification datasets. Batch size is set to 256 for all methods. L denotes the number of hidden layers. Our proposed method can also be combined with convolution kernels Kumar et al. (2018) to obtain a better result, for a fair comparison, we have not implemented here.

Type	DSVI 2	DSVI 3	DSVI 4	DSVI 5
Time (s)	0.835	0.903	0.965	1.339
Iteration	20K	20K	20K	20K
Type	SGHMC 2	SGHMC 3	SGHMC 4	SGHMC 5
Time (s)	0.630	1.000	1.490	1.870
Iteration	20K	20K	20K	20K
Type	NOVI 2	NOVI 3	NOVI 4	NOVI 5
Time (s)	0.391	0.613	0.863	1.123
Iteration	500	500	500	500

Table 2: Comparison of training time (s) of a single iteration and total training iterations on Energy dataset. Batch size is set to 1000 for all three methods.

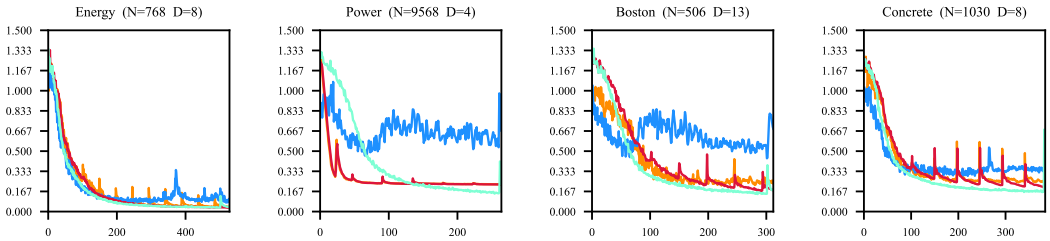

Figure 3: The mean RMSE comparison of NOVI (test: orange, train: red) with Monte Carlo log-likelihood maximization method (test: blue, train: cyan) using 2-layer DGP model on four UCI regression datasets.

6.3 Computational complexity

We have compared training efficiency with other two methods on a single GPU [5] using Energy dataset. Results are shown in Table 2. It can be seen that when our model takes less time per iteration than the other two methods, we only need less than one-tenth of the number of iterations to converge. Also, as shown in Table 1, for high-dimensional image datasets, NOVI also requires significant less training time and iterations to converge, which shows that the proposed method is scalable to larger datasets. Comparison about numbers of inducing points can be seen in App. D.4.

6.4 Ablation study

To demonstrate the effectiveness of NOVI, we directly maximize log-likelihood with random initialized $\mathbf{U}$ and hyperparameters $\boldsymbol{\nu}$ and compare it with our method using 2-layer DGP model. Results are shown in Figure 3. For all datasets, it can be observed that NOVI yields lower test RMSE and higher train RMSE, hence indicating that our optimization method reduces overfitting. Although the loss fluctuation occurs during the training of our method, it is caused by the unique adversarial training and converges to a stable value after only several hundred iterations. Additional results for ablation study on classification datasets can be seen in App. D.2.

7 Conclusion

This paper presented a novel NOVI framework to incorporate Stein Discrepency with DGPs that can effectively model a non-Gaussian and hierarchy-related posterior, thus further enhancing the flexibility of DGP models. To achieve this, we generate inducing variables from a neural generator and optimize it jointly with variational parameters through adversarial training. Furthermore, we theoretically demonstrate that the bias raised by our method can be bounded by Fisher divergence, which provides a clear and concise tool to optimize the neural generator. Empirical evaluation shows that NOVI outperforms the state-of-art approximation methods both in regression and classification. The proposed method also requires significant less training time and iterations to converge, which shows that NOVI is more scalable to larger datasets. Due to the nature of adversarial training, NOVI will inevitably encounter fluctuations in loss during training, causing certain difficulties in optimization, but experimental results show that the fluctuations are greatly alleviated at convergence point. Future work includes implementing convolution structure to better extract features from images and utilizing Neural Architecture Search (NAS) method to obtain a more suitable network architecture for practical applications.

References

Kacper Chwialkowski, Heiko Strathmann, and Arthur Gretton. A kernel test of goodness of fit. In *International conference on machine learning*, pp. 2606–2615. PMLR, 2016.

George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989.

[5] We use Tesla V100 for all computations.

Andreas Damianou and Neil Lawrence. Deep Gaussian processes. In *Proc. AISTATS*, pp. 207–215, 2013.

M. P. Deisenroth and J. W. Ng. Distributed Gaussian processes. In *Proc. ICML*, pp. 1481–1490, 2015.

Lauro Langosco di Langosco, Vincent Fortuin, and Heiko Strathmann. Neural variational gradient descent. *arXiv preprint arXiv:2107.10731*, 2021.

Matthew M Dunlop, Mark A Girolami, Andrew M Stuart, and Aretha L Teckentrup. How deep are deep gaussian processes? *Journal of Machine Learning Research*, 19(54):1–46, 2018.

Y. Gal, M. van der Wilk, and C. E. Rasmussen. Distributed variational inference in sparse Gaussian process regression and latent variable models. In *Proc. NeurIPS*, pp. 3257–3265, 2014.

Xuefeng Gao, Mert Gürbüzbalaban, and Lingjiong Zhu. Global convergence of stochastic gradient hamiltonian monte carlo for nonconvex stochastic optimization: Nonasymptotic performance bounds and momentum-based acceleration. *Operations Research*, 2021.

Jackson Gorham and Lester Mackey. Measuring sample quality with kernels. In *International Conference on Machine Learning*, pp. 1292–1301. PMLR, 2017.

Will Grathwohl, Kuan-Chieh Wang, Jörn-Henrik Jacobsen, David Duvenaud, and Richard Zemel. Learning the stein discrepancy for training and evaluating energy-based models without sampling. In *International Conference on Machine Learning*, pp. 3732–3747. PMLR, 2020.

Marton Havasi, José Miguel Hernández-Lobato, and Juan José Murillo-Fuentes. Inference in deep Gaussian processes using stochastic gradient Hamiltonian Monte Carlo. In *Proc. NeurIPS*, pp. 7517–7527, 2018.

J. Hensman, N. Fusi, and N. Lawrence. Gaussian processes for big data. In *Proc. UAI*, pp. 282–290, 2013.

James Hensman, Alexander Matthews, and Zoubin Ghahramani. Scalable Variational Gaussian Process Classification. In Guy Lebanon and S. V. N. Vishwanathan (eds.), *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, volume 38 of *Proceedings of Machine Learning Research*, pp. 351–360, San Diego, California, USA, 09–12 May 2015. PMLR. URL `https://proceedings.mlr.press/v38/hensman15.html`.

T. N. Hoang, Q. M. Hoang, and K. H. Low. A unifying framework of anytime sparse Gaussian process regression models with stochastic variational inference for big data. In *Proc. ICML*, pp. 569–578, 2015.

T. N. Hoang, Q. M. Hoang, and K. H. Low. A distributed variational inference framework for unifying parallel sparse Gaussian process regression models. In *Proc. ICML*, pp. 382–391, 2016.

M. D. Hoffman, D. M. Blei, C. Wang, and J. Paisley. Stochastic variational inference. *JMLR*, 14(1): 1303–1347, 2013.

Tianyang Hu, Zixiang Chen, Hanxi Sun, Jincheng Bai, Mao Ye, and Guang Cheng. Stein neural sampler. *arXiv preprint arXiv:1810.03545*, 2018.

Ferenc Huszár. Variational inference using implicit distributions. arxiv:1702.08235, 2017.

Michael F Hutchinson. A stochastic estimator of the trace of the influence matrix for laplacian smoothing splines. *Communications in Statistics-Simulation and Computation*, 18(3):1059–1076, 1989.

Marco A Iglesias, Yulong Lu, and Andrew M Stuart. A bayesian level set method for geometric inverse problems. *Interfaces and free boundaries*, 18(2):181–217, 2016.

Wittawat Jitkrittum, Wenkai Xu, Zoltán Szabó, Kenji Fukumizu, and Arthur Gretton. A linear-time kernel goodness-of-fit test. *Advances in Neural Information Processing Systems*, 30, 2017.

Diederik P. Kingma and Max Welling. Auto-encoding variational Bayes. In *Proc. ICLR*, 2013.

Durk P Kingma, Tim Salimans, and Max Welling. Variational dropout and the local reparameterization trick. In *Proc. NeurIPS*, pp. 2575–2583, 2015.

Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images.(2009), 2009.

Vinayak Kumar, Vaibhav Singh, PK Srijith, and Andreas Damianou. Deep gaussian processes with convolutional kernels. *arXiv preprint arXiv:1806.01655*, 2018.

Yann LeCun, Corinna Cortes, and C Burges. Mnist handwritten digit database, 1998. *URL http://www. research. att. com/˜ yann/ocr/mnist*, 1998.

Jakob Lindinger, David Reeb, Christoph Lippert, and Barbara Rakitsch. Beyond the mean-field: Structured deep gaussian processes improve the predictive uncertainties. *Advances in Neural Information Processing Systems*, 33:8498–8509, 2020.

Qiang Liu and Dilin Wang. Stein variational gradient descent: A general purpose bayesian inference algorithm. *Advances in neural information processing systems*, 29, 2016.

Qiang Liu, Jason Lee, and Michael Jordan. A kernelized stein discrepancy for goodness-of-fit tests. In *International conference on machine learning*, pp. 276–284. PMLR, 2016.

Yulong Lu and Jianfeng Lu. A universal approximation theorem of deep neural networks for expressing probability distributions. *Advances in neural information processing systems*, 33: 3094–3105, 2020.

Lars Mescheder, Sebastian Nowozin, and Andreas Geiger. Adversarial variational Bayes: Unifying variational autoencoders and generative adversarial networks. In *Proc. ICML*, pp. 2391–2400, 2017.

Radford M Neal. Monte carlo implementation of gaussian process models for bayesian regression and classification. *arXiv preprint physics/9701026*, 1997.

Chris J Oates, Mark Girolami, and Nicolas Chopin. Control functionals for monte carlo integration. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 79(3):695–718, 2017.

Sebastian W Ober and Laurence Aitchison. Global inducing point variational posteriors for bayesian neural networks and deep gaussian processes. In *International Conference on Machine Learning*, pp. 8248–8259. PMLR, 2021.

Manfred Opper and David Saad. *Advanced mean field methods: Theory and practice*. MIT press, 2001.

Dmytro Perekrestenko, Stephan Müller, and Helmut Bölcskei. Constructive universal high-dimensional distribution generation through deep relu networks. In *International Conference on Machine Learning*, pp. 7610–7619. PMLR, 2020.

J. Quiñonero-Candela and C. E. Rasmussen. A unifying view of sparse approximate Gaussian process regression. *JMLR*, 6:1939–1959, 2005.

Aaditya Ramdas, Sashank Jakkam Reddi, Barnabás Póczos, Aarti Singh, and Larry Wasserman. On the decreasing power of kernel and distance based nonparametric hypothesis tests in high dimensions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29, 2015.

Rajesh Ranganath, Dustin Tran, Jaan Altosaar, and David Blei. Operator variational inference. *Advances in Neural Information Processing Systems*, 29, 2016.

C. E. Rasmussen and C. K. I. Williams. *Gaussian processes for machine learning*. MIT Press, 2006.

Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *Proc. ICML*, pp. 1278–1286, 2014.

Hugh Salimbeni and Marc Deisenroth. Doubly stochastic variational inference for deep Gaussian processes. In *Proc. NeurIPS*, pp. 4588–4599, 2017.

E. L. Snelson and Z. Gharahmani. Sparse Gaussian processes using pseudo-inputs. In *Proc. NeurIPS*, pp. 1257–1264, 2005.

Bharath Sriperumbudur, Kenji Fukumizu, Arthur Gretton, Aapo Hyvärinen, and Revant Kumar. Density estimation in infinite dimensional exponential families. *Journal of Machine Learning Research*, 18, 2017.

M. K. Titsias. Variational learning of inducing variables in sparse Gaussian processes. In *Proc. AISTATS*, pp. 567–574, 2009a.

M. K. Titsias. Variational model selection for sparse Gaussian process regression. Technical report, School of Computer Science, University of Manchester, 2009b.

Michalis K. Titsias and Francisco J. R. Ruiz. Unbiased implicit variational inference. In *Proc. AISTATS*, pp. 167–176, 2019.

Ivan Ustyuzhaninov, Ieva Kazlauskaite, Markus Kaiser, Erik Bodin, Neill Campbell, and Carl Henrik Ek. Compositional uncertainty in deep gaussian processes. In *Conference on Uncertainty in Artificial Intelligence*, pp. 480–489. PMLR, 2020.

Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.

Yunfei Yang, Zhen Li, and Yang Wang. On the capacity of deep generative networks for approximating distributions. *Neural Networks*, 145:144–154, 2022.

Haibin Yu, Yizhou Chen, Bryan Kian Hsiang Low, Patrick Jaillet, and Zhongxiang Dai. Implicit posterior variational inference for deep gaussian processes. *Advances in Neural Information Processing Systems*, 32, 2019.