
Function-space Inference with Sparse Implicit Processes

Simón Rodríguez Santana¹ Bryan Zaldivar² Daniel Hernández-Lobato³

Abstract

Implicit Processes (IPs) represent a flexible framework that can be used to describe a wide variety of models, from Bayesian neural networks, neural samplers and data generators to many others. IPs also allow for approximate inference in function-space. This change of formulation solves intrinsic degenerate problems of parameter-space approximate inference concerning the high number of parameters and their strong dependencies in large models. For this, previous works in the literature have attempted to employ IPs both to set up the prior and to approximate the resulting posterior. However, this has proven to be a challenging task. Existing methods that can tune the prior IP result in a Gaussian predictive distribution, which fails to capture important data patterns. By contrast, methods producing flexible predictive distributions by using another IP to approximate the posterior process cannot tune the prior IP to the observed data. We propose here the first method that can accomplish both goals. For this, we rely on an inducing-point representation of the prior IP, as often done in the context of sparse Gaussian processes. The result is a scalable method for approximate inference with IPs that can tune the prior IP parameters to the data, and that provides accurate non-Gaussian predictive distributions.

1. Introduction

Approximate inference has recently attracted an enormous interest from the neural networks (NNs) community due to the high performance and popularity of these models (Gal & Ghahramani, 2016; Sun et al., 2019). Specifically, approximate inference is seen as a fast and approachable way for

¹Institute of Mathematical Sciences (ICMAT-CSIC), Madrid, Spain. ²Institute of Corpuscular Physics, University of Valencia and CSIC, Spain. ³Escuela Politécnica Superior, Universidad Autónoma de Madrid, Spain.. Correspondence to: Simón Rodríguez Santana <simon.rodriguez@icmat.es>.

obtaining more complete information about the confidence on the predictions made by NNs (Blundell et al., 2015b; Gal, 2016). The resulting methods are more informative, which can be useful in real-world scenarios. This is specially important in sensitive tasks that may require taking decisions based on the obtained predictions and their respective uncertainties, *e.g.*, improving the safety of autonomous vehicles (McAllister et al., 2017). Employing the Bayesian framework, in which one uses probability distributions to assign degrees of belief on the weight values of these models, can be successful in this task (Blundell et al., 2015a; Graves, 2011). However, most of these techniques come with intrinsic difficulties such as the choice of a meaningful prior distribution or the intractability of some of the expressions needed, which must be approximated. Moreover, the increasing complexity of models can also lead to difficulties, since performing inference in the space of parameters becomes more challenging in larger models. In such a case, the parameter space presents numerous symmetries and strong dependencies, which makes approximating the target parameter posterior distribution a complicated task. This makes the models harder to train, requiring more data to do so, and compromises the prediction performance (Sun et al., 2019).

There has been an increasing interest recently in *implicit stochastic processes* (IPs) as a way to express prior distributions on the function space to face the problems described (Ma et al., 2019; Sun et al., 2019). Although it requires a slightly more complex formulation, doing approximate inference in the function space presents a way to simplify the problems that arise in the aforementioned parameter-space. Nevertheless, recent work has only been able to show partial success in this matter: the proposed methods so far have only been able to update the prior model parameters according to the data by sacrificing the flexibility of the predictive distribution, which is fixed to be Gaussian. Alternatively, they can only generate a more complex and flexible predictive distribution at the expense of not being capable of adjusting the prior parameters. Thus, a general method that can do both things is missing in the literature. Namely, (i) Update the prior to learn important features of the data. (ii) Provide non-Gaussian and flexible predictive distributions.

Our contribution is a new method for approximate inference using IPs called *Sparse Implicit Process* (SIP). SIP fulfills both objectives (i) and (ii). We have evaluated SIP in the

context of Bayesian NNs and neural samplers (Ma et al., 2019), as two illustrative examples to showcase its flexibility. Importantly, however, IPs and our method, SIP, are very general models and include, as particular cases, other commonly used priors over functions. Some examples include *normalizing flows* (Rezende & Mohamed, 2015), deep GPs (Damianou & Lawrence, 2013) and warped GPs (Snelson et al., 2004). To obtain a scalable method that can address very large datasets we consider an *inducing-point* approach in SIP in which the number of latent variables on which to perform inference is reduced from N to $M \ll N$, with N the training set size (Snelson & Ghahramani, 2005; Titsias, 2009). We have evaluated SIP in extensive regression experiments, including synthetic and real-world problems. When compared against other methods from the literature, SIP often leads to better generalization properties and it captures complex patterns in the predictive distribution. It can also adjust the prior parameters to explain the training data. Finally, in very large datasets, it also remains scalable, reaching good results in less time than other methods.

2. Background

We describe first current methods for approximate inference in parameter-space using implicit distributions. After that, we explain in detail how to carry out approximate inference in the functional space using IPs.

2.1. Parameter-space Approximate Inference using Implicit Distributions

Consider some data $\mathcal{D} = \{\mathbf{x}_i, y_i\}_{i=1}^N$, a prior $p(\mathbf{w})$ over the model parameters \mathbf{w} and a likelihood $p(\mathbf{y}|\mathbf{w}, \mathbf{X})$. The goal is to find a distribution $q_\phi(\mathbf{w})$ to approximate the exact posterior $p(\mathbf{w}|\mathbf{y}, \mathbf{X})$. In variational inference, q maximizes the evidence lower bound (ELBO) (Jordan et al., 1999):

$$\mathcal{L}(\phi) = \mathbb{E}_{q_\phi(\mathbf{w})}[\log p(\mathbf{y}|\mathbf{w}, \mathbf{X})] - \text{KL}(q_\phi(\mathbf{w})|p(\mathbf{w})) \quad (1)$$

where $\text{KL}(\cdot, \cdot)$ is the Kullback-Leibler divergence between distributions. Maximizing (1) is equivalent to minimizing $\text{KL}(q_\phi(\mathbf{w})|p(\mathbf{w}|\mathbf{y}, \mathbf{X}))$. Typical approaches for Bayesian NN (BNN) models use a parametric distribution q that assumes independence among the components of \mathbf{w} (Blundell et al., 2015a; Graves, 2011). By contrast, in Mescheder et al. (2017); Santana & Hernández-Lobato (2020), a more expressive approximation is used, making q implicit, *i.e.*, $q_\phi(\mathbf{w}) = \int q_\phi(\mathbf{w}|\epsilon)p(\epsilon)d\epsilon$, with ϵ some random noise. Therefore, if ϵ is high-dimensional and $q_\phi(\mathbf{w}|\epsilon)$ is complicated enough, the result is a very flexible approximate distribution $q_\phi(\mathbf{w})$. Note, however, that although the first term in (1) can be estimated via Monte Carlo sampling, the KL contribution is intractable since q lacks a closed-form

density. To solve this, the KL term can be re-written as:

$$\begin{aligned} \text{KL}(q_\phi(\mathbf{w})|p(\mathbf{w})) &= \mathbb{E}_{q_\phi(\mathbf{w})} [\log q_\phi(\mathbf{w}) - \log p(\mathbf{w})] \\ &= \mathbb{E}_{q_\phi(\mathbf{w})} [T(\mathbf{w})], \end{aligned} \quad (2)$$

where $T(\mathbf{w})$ is the log-ratio between $q_\phi(\mathbf{w})$ and the prior for \mathbf{w} . Let $\sigma(\cdot)$ be the sigmoid function. $T(\mathbf{w})$ can be estimated using a classifier $T_\omega(\mathbf{w})$, with parameters ω obtained by maximizing the objective:

$$\mathbb{E}_{q_\phi} [\log \sigma(T_\omega(\mathbf{w}))] + \mathbb{E}_{p(\mathbf{w})} [\log(1 - \sigma(T_\omega(\mathbf{w})))] . \quad (3)$$

If the discriminator is flexible enough, its optimal value, T_{ω^*} , is exactly the log ratio between $q_\phi(\mathbf{w})$ and $p(\mathbf{w})$, *i.e.*, $T_{\omega^*}(\mathbf{w}) = \log q_\phi(\mathbf{w}) - \log p(\mathbf{w})$ (Mescheder et al., 2017). Using this, (1) becomes

$$\mathcal{L}(\phi) = \sum_{i=1}^N \mathbb{E}_{q_\phi} [\log p(y_i|\mathbf{w}, \mathbf{x}_i)] - \mathbb{E}_{q_\phi} [T_{\omega^*}(\mathbf{w})], \quad (4)$$

where the KL term is replaced by the discriminator, which is trained simultaneously. This enables the use of an implicit model for $q_\phi(\mathbf{w})$. Finally, instead of minimizing the regular KL-divergence between q and the posterior, other approaches suggest to minimize a more general α -divergence, which includes the KL-divergence as a particular case (Hernández-Lobato et al., 2016). This has shown to overall improve the final results (Santana & Hernández-Lobato, 2020; Ma et al., 2019).

2.2. Function-Space Approximate Inference

Implicit processes (IPs) are defined as a collection of random variables $f(\cdot)$ such that any finite set of evaluations $(f(\mathbf{x}_1), \dots, f(\mathbf{x}_N))^T$ has joint distribution determined by the generative process (Ma et al., 2019):

$$\mathbf{z} \sim p(\mathbf{z}), \quad f(\mathbf{x}_i) = g_\theta(\mathbf{x}_i, \mathbf{z}), \quad \forall \mathbf{x}_i \in \mathcal{X}, \quad (5)$$

where \mathbf{z} is some random variable that summarizes the randomness, θ represents the parameters of the process and \mathcal{X} is the input space. We use the notation $f(\cdot) \sim \mathcal{IP}(g_\theta(\cdot, \cdot), p_{\mathbf{z}})$ to indicate that f is sampled from the corresponding IP with parameters θ , using samples from $p(\mathbf{z})$ (denoted as $p_{\mathbf{z}}$). This definition of IPs results in a framework that is general enough to include many different models. For example, Bayesian NNs can be described using IPs if the randomness is given by the prior $p(\mathbf{w})$ over the NN weights \mathbf{w} . We would then sample a function parameterized by $\mathbf{w} \sim p(\mathbf{w})$, which specifies the output of the NN as $f(\mathbf{x}) = g_\theta(\mathbf{x}, \mathbf{w})$ for every $\mathbf{x} \in \mathcal{X}$. θ are here the parameters of $p(\mathbf{w})$. If $p(\mathbf{w})$ is a factorizing Gaussian, θ will be the corresponding means and variances. Other important models that can be described as IPs include, *e.g.*, neural samplers (NS), warped GPs, and deep GPs (Ma et al., 2019; Snelson et al., 2004;

Damianou & Lawrence, 2013). Previous works using IPs for inference are the variational implicit processes (VIP) and the functional variational Bayesian NN (fBNN) (Ma et al., 2019; Sun et al., 2019). We introduce them next.

VIPs approximate the marginal likelihood of the prior IP by the marginal likelihood of a GP. For this, an empirical covariance function is estimated by sampling from the prior IP. Namely, $f_s^\theta(\cdot) \sim \mathcal{IP}(g_\theta(\cdot, \cdot), p_{\mathbf{z}})$. Then, the prior mean and covariances of the GP are:

$$\begin{aligned} m^*(\mathbf{x}) &= \frac{1}{S} \sum_{s=1}^S f_s^\theta(\mathbf{x}), \\ \mathcal{K}^*(\mathbf{x}_1, \mathbf{x}_2) &= \frac{1}{S} \sum_{s=1}^S \Delta_s^\theta(\mathbf{x}_1) \Delta_s^\theta(\mathbf{x}_2), \end{aligned} \quad (6)$$

where $\Delta_s(\mathbf{x}) = f_s^\theta(\mathbf{x}) - m^*(\mathbf{x})$. They then approximate $p(\mathbf{y}|\mathbf{X}, \theta)$ by $q_{\text{GP}}(\mathbf{y}|\mathbf{X}, \theta)$, *i.e.*, the marginal likelihood of a GP with the estimated means and covariances, and maximize the latter, expecting that it will increase $p(\mathbf{y}|\mathbf{X}, \theta)$ as well. To guarantee scalability and avoid the cubic cost of the GP they further approximate the GP using a linear model with the same mean and covariances. The linear model is efficiently tuned by optimizing the α -energy function, performing gradient descent w.r.t. θ via the method described in Hernández-Lobato et al. (2016). A limitation of VIP, however, is that the final predictive distribution is Gaussian (that of a GP) which may lack enough flexibility.

By contrast, fBNNs rely on VI and, instead of using a GP, they use another IP to approximate the posterior of the prior IP (Sun et al., 2019). In the ELBO in (1), the first term can be estimated by Monte Carlo when using an IP as the approximate posterior. However, the KL term becomes the KL-divergence between stochastic process, which is intractable. To address this, fBNN evaluates the KL-divergence between distributions at a finite set of points $\tilde{\mathbf{X}}$ drawn at random from the input space, leaving the ELBO as:

$$\mathcal{L}(q) = \mathbb{E}_q[\log p(\mathbf{y}|\mathbf{f}^{\tilde{\mathbf{X}}})] - \mathbb{E}_{\tilde{\mathbf{X}}}[\text{KL}(q(\mathbf{f}^{\tilde{\mathbf{X}}})|p(\mathbf{f}^{\tilde{\mathbf{X}}}))], \quad (7)$$

where $\mathbf{f}^{\mathbf{X}}$ and $\mathbf{f}^{\tilde{\mathbf{X}}}$ are the IP values at \mathbf{X} and $\tilde{\mathbf{X}}$, respectively. This objective function is then maximized w.r.t the parameters of the posterior approximate IP q . Critically, $\tilde{\mathbf{X}}$ must cover training and testing regions of the input space. Therefore, fBNN may suffer in large dimensional datasets. Moreover, fBNN is unable to fit the prior IP model by itself due to estimating the gradients of the KL-divergence term using a spectral gradient estimator (Sun et al., 2019).

3. Sparse Implicit Processes

We introduce *Sparse Implicit Processes* (SIP), a new method for approximate inference when using IPs. As in fBNN, in

SIP we consider another IP to approximate the posterior process resulting from combining the prior IP with the data using Bayes' rule. However, unlike in fBNN we perform inference on finite sets of variables, as it is usually done with GPs. This avoids the problem of computing the KL-divergence between stochastic processes. SIP also allows to easily adjust the parameters of the prior IP. However, we will need to address two issues: (i) avoid the number of latent variables increasing with the number of training points N ; and (ii) deal with the intractability of the computations.

The first problem can be sorted out by considering an approximation based on inducing points, as in sparse GPs (Snelson & Ghahramani, 2005; Titsias, 2009). Instead of making inference about $\mathbf{f} = (f(\mathbf{x}_1), \dots, f(\mathbf{x}_N))^T$ we perform inference about the process values at $M \ll N$ inducing points. We denote the set of inducing points $\tilde{\mathbf{X}}$, and denote the IP values at these input locations $\mathbf{u} = (f(\tilde{\mathbf{x}}_1), \dots, f(\tilde{\mathbf{x}}_M))^T$. Next, we focus on approximating $p(\mathbf{f}, \mathbf{u}|\mathcal{D})$, which only depends on finite sets of variables. For this, we consider the approximate posterior distribution:

$$q(\mathbf{f}, \mathbf{u}) = p_\theta(\mathbf{f}|\mathbf{u})q_\phi(\mathbf{u}), \quad (8)$$

where $q_\phi(\mathbf{u})$ is an implicit distribution with parameters ϕ , and θ are the prior IP parameters. Critically, $p_\theta(\mathbf{f}|\mathbf{u})$ is fixed and $q_\phi(\mathbf{u})$ is tunable, as in the variational sparse GP approximation (Titsias, 2009). We will use a Monte Carlo GP approximation of $p_\theta(\mathbf{f}|\mathbf{u})$, following what is done in VIP to approximate the prior IP using a GP. However, unlike VIP, here we only use the Gaussian approximation in part of the posterior. We will explicitly define $p_\theta(\mathbf{f}|\mathbf{u})$ later. For now, using this decomposition, the obtained functional-ELBO is:

$$\begin{aligned} \mathcal{L}(\phi, \theta) &= \mathbb{E}_{q_{\phi, \theta}} \left[\log \frac{p(\mathbf{y}|\mathbf{f})p_\theta(\mathbf{f}|\mathbf{u})p_\theta(\mathbf{u})}{p_\theta(\mathbf{f}|\mathbf{u})q_\phi(\mathbf{u})} \right] \\ &= \mathbb{E}_{q_\phi}[\log p(\mathbf{y}|\mathbf{f})] - \text{KL}(q_\phi(\mathbf{u})|p_\theta(\mathbf{u})). \end{aligned} \quad (9)$$

The first term can be estimated by Monte Carlo sampling, as in standard approaches for VI. However, the second term lacks any closed-form solution, since it is the KL-divergence between two implicit distributions. To estimate it we rely on the method described in Section 2.1, where a classifier is used to estimate the log-ratio (Mescheder et al., 2017; Santana & Hernández-Lobato, 2020). Namely,

$$\text{KL}(q_\phi(\mathbf{u})|p_\theta(\mathbf{u})) = \mathbb{E}_q [T_{\omega^*}(\mathbf{u})], \quad (10)$$

where $T_{\omega^*}(\mathbf{u})$ is approximated by a NN that discriminates between samples from $q_\phi(\mathbf{u})$ and $p_\theta(\mathbf{u})$. Note, however, that $T_{\omega^*}(\mathbf{u})$ depends on ϕ and θ . Nevertheless, as argued in Mescheder et al. (2017), $\mathbb{E}_q(\nabla_\phi [T_{\omega^*}(\mathbf{u})]) = 0$ if $T_{\omega^*}(\mathbf{u})$ is optimal. During the training procedure, we must ensure that the discriminator is trained enough so that it provides a fitting approximation to the original KL term in (9) (see Appendix B for further discussion on the matter).

Regarding $\nabla_{\theta} T_{\omega^*}(\mathbf{u})$, note that it is expected to be small when compared to $\nabla_{\theta} \mathbb{E}_{q_{\phi}}[\log p(\mathbf{y}|\mathbf{f})]$ (see Appendix A). However, although small, these gradients play a crucial role to adjust the prior distribution to the observed data. This can be critical in different cases, *e.g.* when specifying a prior is complicated, or when making predictions far from the observed data. When a classifier is used to approximate the KL-divergence, these gradients will be ignored, which results in not properly adjusting the prior parameters θ (see Appendix A). To partially correct for this we substitute the KL-divergence by the symmetrized KL-divergence:

$$\text{KL}(q_{\phi}|p_{\theta}) \approx \frac{1}{2}(\text{KL}(q_{\phi}|p_{\theta}) + \text{KL}(p_{\theta}|q_{\phi})), \quad (11)$$

where we have omitted the dependence on \mathbf{u} of q_{ϕ} and p_{θ} . The KL-divergence in VI is just a regularizer enforcing q to look similar to the prior. (11) is expected to play a similar role. Moreover, modifying this regularizer in VI is a common approach that often gives better results. See, *e.g.*, Wenzel et al. (2020). Importantly, the reversed KL-divergence, $\text{KL}(p_{\theta}|q_{\phi})$, involves also the log-ratio between the prior and the posterior approximation q . Therefore, it can also be estimated using the same classifier $T_{\omega^*}(\mathbf{u})$. Critically, however, it involves an expectation with respect to $p_{\theta}(\mathbf{u})$. This introduces some easy to compute gradients with respect to the parameters of the prior θ . As shown in our experiments, we have found those dependencies to be enough to provide some prior adaptation to the data. This choice is also supported by good empirical results obtained and because the prior adaptation is not observed when only the first KL-divergence term is considered (see Appendix A). When changing the KL-term the objective is no longer a lower bound on the log-marginal likelihood. However, this is also the case for VIP, which uses a GP approximation to the log-marginal likelihood (Ma et al., 2019). Moreover, as indicated by Knoblauch et al. (2019), the distributions obtained, as alternative approximations to VI, often perform better in practice. Note that VI is optimal only relative to a particular objective, whose origins are assumptions (*e.g.*, formulation of the likelihood and the prior) that could be misaligned with reality.

An important remark is that the data-dependent term in (9) only considers the squared error (in the case of a Gaussian likelihood). We follow Santana & Hernández-Lobato (2020) and resort to the energy function employed in BB- α and VIP (Hernández-Lobato et al., 2016; Ma et al., 2019). Replacing the data-dependent term in (9) results in the approximate optimization of α -divergences. This generalizes the original data-dependent term, making the objective more sensitive to the actual distribution of the data for higher values of α . In particular, for $\alpha \rightarrow 0$, the data-dependent term in (9) is obtained, while for $\alpha \rightarrow 1$, it resembles the log-likelihood of the training data. With this re-formulation, the objective

depends on α and is:

$$\begin{aligned} \mathcal{L}_{\alpha}^*(\phi, \theta) &= \frac{1}{\alpha} \sum_{i=1}^N \log \mathbb{E}_{q_{\phi, \theta}} [p(y_i|f_i)^{\alpha}] \\ &\quad - \frac{1}{2} [\text{KL}(q_{\phi}|p_{\theta}) + \text{KL}(p_{\theta}|q_{\phi})], \end{aligned} \quad (12)$$

where the first term can be estimated via Monte Carlo, using a small number of samples as in Hernández-Lobato et al. (2016). Moreover, α can be chosen to target the data log-likelihood, *i.e.*, when $\alpha = 1$. When $\alpha \rightarrow 0$ the original data-dependent term in (9) is obtained. The bias introduced by the $\log(\cdot)$ function in (12) becomes negligible by using a small number of samples (Hernández-Lobato et al., 2016).

The other critical point of SIP is how to compute $p_{\theta}(\mathbf{f}|\mathbf{u})$ in (8). We approximate this conditional distribution by the conditional of a GP with the same covariance and mean function as the prior IP, as it is done in VIP (Ma et al., 2019). More precisely, given samples of \mathbf{f} and \mathbf{u} , we employ (6) to estimate the means and covariances needed via Monte Carlo. We then resort to the GP predictive equations described in Rasmussen & Williams (2006). Namely, $p_{\theta}(\mathbf{f}|\mathbf{u})$ is approximated as Gaussian with mean and covariance

$$\begin{aligned} \mathbb{E}[\mathbf{f}] &= m(\mathbf{x}) + \mathbf{K}_{\mathbf{f}, \mathbf{u}}(\mathbf{K}_{\mathbf{u}, \mathbf{u}} + \mathbf{I}\sigma^2)^{-1}(\mathbf{u} - m(\bar{\mathbf{X}})), \\ \text{Cov}(\mathbf{f}) &= \mathbf{K}_{\mathbf{f}, \mathbf{f}} - \mathbf{K}_{\mathbf{f}, \mathbf{u}}(\mathbf{K}_{\mathbf{u}, \mathbf{u}} + \mathbf{I}\sigma^2)^{-1}\mathbf{K}_{\mathbf{u}, \mathbf{f}}, \end{aligned} \quad (13)$$

with σ^2 a small noise variance (set to 10^{-5}). Moreover, $m(\cdot)$ and each entry in $\mathbf{K}_{\mathbf{f}, \mathbf{u}}$ and $\mathbf{K}_{\mathbf{u}, \mathbf{u}}$ is estimated empirically using (6), as in VIP (Ma et al., 2019).

Finally, predictions at a new point \mathbf{x}^* are estimated via Monte Carlo with S samples. Let $\mathbf{u}_s \sim q_{\phi}(\mathbf{u})$. Then,

$$p(f(\mathbf{x}_*)|\mathbf{y}, \mathbf{X}) \approx \frac{1}{S} \sum_{s=1}^S p_{\theta}(f(\mathbf{x}_*)|\mathbf{u}_s). \quad (14)$$

Thus, the predictive distribution is a mixture of Gaussians, where each Gaussian is determined by one sample extracted from the approximate posterior IP at $\bar{\mathbf{X}}$. This means that SIP can produce flexible predictive distributions that need not be Gaussian, unlike VIP.

To conclude, we illustrate the SIP method with its graphical model in Figure 1. Here θ are the prior IP parameters. The inducing points' locations are explicitly depicted as \mathbf{z} , with their respective sampled functions values locations \mathbf{u} . The function values on the input data are \mathbf{f}_n , with N training data points $\{\mathbf{x}_n, n = 1, \dots, N\}$, while f^* is used for the test data (\mathbf{x}^*). Finally, each noise contribution is sampled from $\mathcal{N}(0, \sigma_{noise})$ (*i.e.* ϵ_n for training and ϵ^* for testing). We then estimate the output, which can correspond to an observed (y_n , in training) or unobserved variable (\tilde{y}^* , in testing).

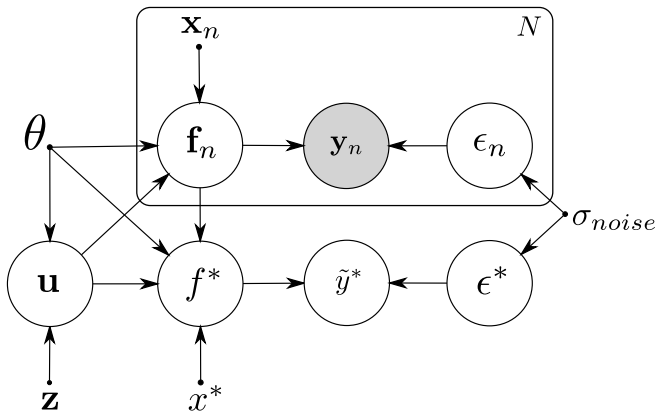


Figure 1. Graph model for the SIP method. Point-like vertices denote deterministic variables and circles represent random variables, which can either be observed (grey) or unobserved (white).

4. Related Work

Traditionally, approximate inference has been conducted in parameter space. Bayesian Neural Networks (BNNs) are a notable example. They enable predictive distributions accounting for prediction uncertainty in the context of NN. A widely explored technique for this is Bayes by back-propagation (BBB) (Blundell et al., 2015a; Graves, 2011; Jordan et al., 1999). BBB performs variational inference (VI) in the space of weights using a parametric distribution (MacKay, 1992; Graves, 2011). Novel approaches are leading to new interpretations and generalizations based on VI, from which the resulting methods can have appealing theoretical properties (Knoblauch et al., 2019). However, constraining the approximate solution to a certain parametric family which in most cases assumes independence may be too restrictive. More precisely, in complex NNs this approach leads to pathological behaviors that may result in a poor generalization (Sun et al., 2019). An alternative to BBB is Probabilistic Back-propagation (PBP) (Hernández-Lobato & Adams, 2015). This method propagates probabilities through the NN. PBP has proven to be efficient and scalable, although it also has limited expressiveness in the posterior approximation, which is a factorizing Gaussian. Therefore, as BBB, PBP suffers from problematic approximation bias (Graves, 2011; Blundell et al., 2015a).

Recent works have analyzed the properties of simple variational approximation methods. In Foong et al. (2020) they are shown to underestimate the prediction uncertainty. Wide BNNs are also subject of study in Coker et al. (2021) under the mean-field assumption. Pathological behaviors arise for deeper BNN models in the approximate posterior, which strongly differs from the exact one. Therefore, simple VI approximations based on, *e.g.*, mean-field should be avoided. More flexible approximations can be obtained using nor-

malizing flows (NF) (Rezende & Mohamed, 2015). NFs perform a series of non-linear invertible transformations on the variables of a tractable parametric distribution, obtaining a more complex distribution with closed-form density. For this, the transformations must be invertible, which may limit the flexibility of the approximate solution.

There has been a growing interest in increasing the flexibility of the approximate distribution (Liu & Wang, 2016; Salimans et al., 2015; Tran et al., 2017). An implicit model can be useful for this (Li & Liu, 2016; Mescheder et al., 2017; Santana & Hernández-Lobato, 2020). There, the approximate distribution lacks a closed-form density, but one can easily sample from it. Since q lacks a density expression, it becomes difficult to evaluate the KL-divergence term of the VI objective. Adversarial Variational Bayes (AVB) solves this problem (Mescheder et al., 2017) using an auxiliary discriminator. See Section 2.1 for further details. AVB has been extended by Adversarial α -divergence minimization (AADM), which combines the BB- α objective (Hernández-Lobato et al., 2016) with an implicit model for q to locally minimize an α -divergence (Santana & Hernández-Lobato, 2020). In AADM $\alpha \in (0, 1]$ is an adjustable parameter which allows to interpolate between targeting the direct and the reversed KL-divergence. Furthermore, AADM can model complex predictive distributions, unlike AVB.

Other works consider using inducing points and sparse models in the context of BNN (Immer et al., 2021; Ritter et al., 2021). However, in contrast to SIP, the locations of the inducing points are fixed, and also, unlike SIP, approximate inference is carried out in the parameter space. SIP is also more general since it is not restricted to work with BNNs.

The methods described so far suffer from the problems of working in the space of parameters, which is high-dimensional and includes strong dependencies. Recent works have shown better results by performing approximate inference in the space of functions (Ma et al., 2019; Sun et al., 2019). Characterizing function-space inference, as well as constructing effective frameworks for it focuses many research efforts nowadays (Burt et al., 2021). In some of the most successful approaches, the underlying model is constructed using an implicit stochastic process. Two successful methods for approximate inference in the context of IPs are VIP (Ma et al., 2019) and fBNN (Sun et al., 2019), described in detail in Section 2.2. VIP is limited to having a Gaussian predictive distribution, which may lack flexibility. By contrast, in fBNN is difficult to infer the parameters of the prior IP. Moreover, fBNN also relies on uniformly covering the input space to guarantee that the posterior IP looks similar to the prior IP in regions with no data. This is challenging in high dimensions. SIP does not have these limitations and can produce flexible predictive distributions and adjust the prior parameters to the data.

Functional Variational Inference (FVI) is a recent method proposed for approximate inference in the context of IPs (Ma & Hernández-Lobato, 2021). FVI approximately minimizes the KL-divergence between stochastic processes. This is done efficiently in a two step approach. First, the prior IP is approximated using a flexible model called stochastic process generator (SPG). Then, another SPG is used to efficiently approximate the posterior of the prior SPG. Both SPGs share key aspects that facilitate this task. In any case, FVI suffers from the same limitation as fBNN. It cannot adjust the prior parameters to the data, unlike in SIP.

In the context of BNNs, Rudner et al. (2020) approximate the KL-divergence between stochastic processes by considering a NN linearized around its mean parameters. After using an inducing point approximation, the result is the KL-divergence between two Gaussians. This approach is, however, constrained to (i) using BNNs with (ii) Gaussian weights and a (iii) mean-field approximation, unlike SIP.

Approximate inference in the functional space is not a new concept. Methods based on GPs have been used extensively (Rasmussen & Williams, 2006). In GPs, the posterior can be computed analytically. However, they have a big training cost. Inducing points approximations, similar to those of SIP, can be combined with stochastic VI in GPs for better scalability (Hensman et al., 2013). However, these methods are limited to using GP priors and only produce Gaussian predictions. IP based methods, such as SIP, can use a wider range of priors and can produce non-Gaussian predictions.

5. Experiments

We evaluate the proposed method, SIP, using different NN models for the prior: a Bayesian NN (BNN) and a neural sampler (NS), both with 2 hidden layers of 50 units. Such NN size is similar to that used in, e.g., Ma et al. (2019); Sun et al. (2019); Santana & Hernández-Lobato (2020). In the NS, input noise is standard Gaussian with 10 dimensions. The posterior implicit model in SIP uses another NS with 100 noise dimensions. The tensor-flow code for SIP is included in the [github repository](#). Each method is trained until convergence using a mini-batch size of 10. For training, in SIP we use 100 samples to estimate (12) and its gradients. In test, 500 samples are used in (14). The exact number of inducing points used is specified in each experiment, although through preliminary experiments we saw little improvements using more than 50 – 100 inducing points. Adding more is expected to improve the accuracy of the method, but also raises the computational cost (Titsias, 2009). Finally, we use $\alpha = 1$ in SIP in the synthetic experiments (this targets the data log-likelihood for visualization of the predictive distribution) and $\alpha = 0.5$ in the real-world ones. $\alpha = 0.5$ gives good results for the metrics used and is recommended in Santana & Hernández-Lobato (2020); Ma

et al. (2019). We compare SIP with BBB, VIP, fBNN, and AADM (see Section 4). The noise variance is chosen by maximizing the estimate of the log-marginal likelihood of each method. We follow Ma et al. (2019); Sun et al. (2019) and focus our experiments on similar regression problems.

5.1. Synthetic Experiments

These experiments compare the quality of the predictive distribution given by SIP with that of VIP and fBNN. We train each method using the same BNN prior for a fair assessment. Here, in VIP we have disabled the regularization term to favor a better prior fit. If included, VIP’s heterocedastic behavior is lost (see Appendix D.1). Moreover, Appendix D.2 has extra results for another dataset. The results obtained are similar to these ones. Here we also evaluate the performance of Hamilton Monte Carlo (HMC) applied on the same model. The prior for HMC is the BNN prior learned by SIP. The number of inducing points in SIP is set to 50.

We first study a bimodal dataset we generated taking 1000 samples of x from a uniform distribution $\mathcal{U}(-4, 4)$. Then, we generate one of two possible values for y :

$$y_1 = 10 \cos(x - 0.5) + \epsilon, \quad y_2 = 10 \sin(x - 0.5) + \epsilon$$

with $\epsilon \sim \mathcal{N}(0, 1)$. We select either y_1 or y_2 randomly, producing bimodal data. For comparison, we plot the samples from the learned prior and the predictive distribution of each method after 2000 training epochs.

Figure 2 show samples from the learned prior distribution and the predictive distribution for y of each method. The original training data is shown in the top left corner. We observe that in the case of VIP and SIP, the prior model captures the mean value of the training data, while for fBNN seems to have not learned any pattern at all. However, VIP’s predictive distribution, which is Gaussian, is unable to represent the bimodality of the data. fBNN’s predictive distribution, although more flexible than the one of VIP by construction, cannot capture the bimodality either. The reason behind this is that the data-dependent term of fBNNs focuses on minimizing the squared error, and hence it simply outputs the average prediction between the two modes, as illustrated in Santana & Hernández-Lobato (2020). In summary, SIP is the only method that learns a sensible prior distribution and whose predictions capture the bimodality of the data (see Appendix D for extra experiments).

An unexpected result in Figure 2 is that HMC, which is expected to be the most accurate method for approximate inference, cannot capture the bimodal predictions. This is simply because the assumed model (a BNN) is wrong and the exact posterior need not be optimal. In particular, if one randomly generates functions from the NN prior to then contaminate them with additive Gaussian noise, the bimodal predictive distribution is never observed, in practice. A uni-

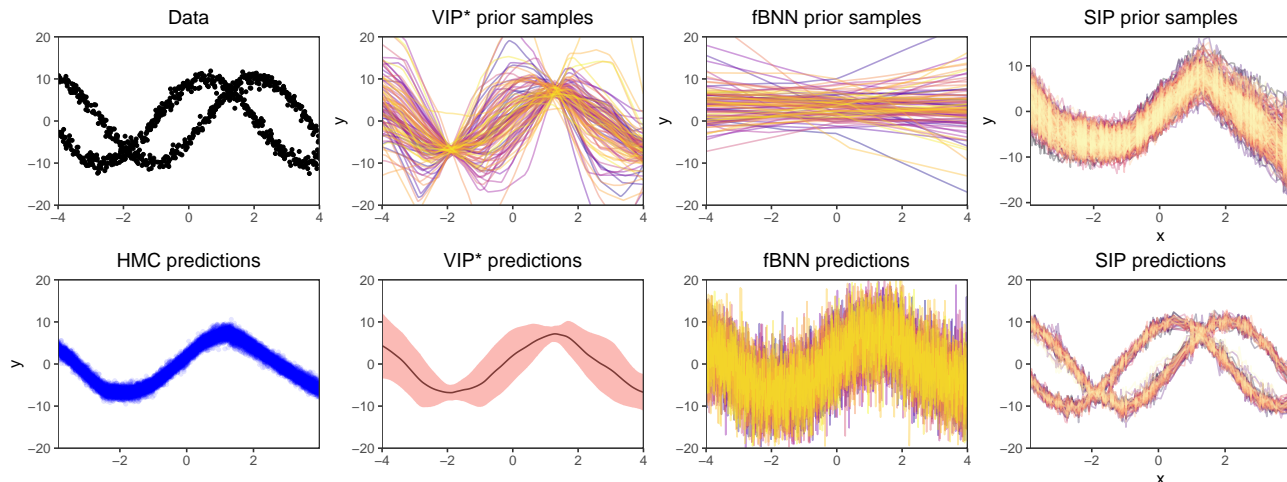


Figure 2. Samples from the prior and the predictive distribution of each method. First column contains the data (top row, in black) and the HMC predictions (bottom row, in blue). For the rest of the methods, the first row shows samples from the learned prior distribution. The second row shows the samples from the predictive distribution. Best seen in color.

modal predictive distribution is obtained. This is the one captured by HMC. This makes sense since a BNN prior will converge to a GP when number of hidden units is large (Neal, 1994). A GP will always output a Gaussian predictive distribution and not a bi-modal predictive distribution. By contrast, SIP is more flexible and thanks to the approximate inference mechanism is able to bypass the wrong model specification and produces a more accurate predictive distribution. Therefore, the approximate inference method used by SIP is robust to model mis-specification, as a consequence of the data-dependent term, which enforces to produce accurate predictions for the training instances.

5.1.1. LOCATION OF THE INDUCING POINTS

Using inducing points is critical in SIP, and hence it is important to assess that it is able of learning their locations, $\bar{\mathbf{X}}$, correctly. These points specify the potential values that the posterior IP can take in different regions of the input space. More precisely, consider some level of smoothness implied by the prior IP. Then, in regions of the input space close to $\bar{\mathbf{X}}$ one should expect a similar process value as the one specified by the implicit distribution $q_\phi(\mathbf{u})$. To test the ability to learn the location of the inducing points we consider the synthetic dataset in (Snelson & Ghahramani, 2005) and 15 inducing points in SIP. We initialize them adversarially concentrated in one input location and train the model. In this case we do not fit the parameters of q_ϕ so that the model focuses on the inducing points locations and the prior for prediction, preventing $q_\phi(\mathbf{u})$ to compensate for locations with not enough inducing points nearby.

Figure 3 shows the predictive distribution obtained by SIP,

as well as the changes in position of the inducing points across epochs (scaled by $1e3$). We see that the inducing points locate themselves covering the whole range of the training data. After epoch 2000 they move very little. This shows that the method is able to successfully locate the inducing points to the most convenient position (see Appendix D.3). Moreover, the predictive distribution seems good even though we do not fit the parameters of $q_\phi(\mathbf{u})$ here.




Figure 3. Predictive distribution (top) and evolution of the location of the inducing points (bottom) for the dataset in (Snelson & Ghahramani, 2005). The crosses at the top and bottom represent the starting and finishing positions of the inducing points, respectively. Training epochs are scaled by 10^3 . Best seen in color.

5.2. Regression on UCI Datasets

We compare each method and SIP with each prior (*i.e.*, a BNN and a NS) on multivariate regression problems from the public UCI dataset repository (Dua & Graff, 2017).

Table 1. Ranking analysis between methods across every multivariate regression problem (lower is better). The method with the lowest average rank across data sets and splits is highlighted in bold.

Method	BBB	AVB	AADM	fBNN _{BNN}	VIP	SIP _{NS}	SIP _{BNN}
RMSE	6.15±0.04	4.14±0.08	3.97±0.09	3.82±0.10	3.29±0.09	3.10±0.09	3.53±0.08
NLL	5.39±0.05	3.91±0.06	2.82±0.06	5.61±0.07	3.70±0.12	3.81±0.07	2.76±0.08
CRPS	5.84±0.05	4.47±0.08	4.03±0.08	3.97±0.07	2.57±0.08	4.67±0.08	2.45±0.08

We refer to these methods as SIP_{BNN} and SIP_{NS}, respectively. Following Ma et al. (2019), we select 8 datasets: Boston Housing, Concrete, Energy Efficiency, Kin8nm, Naval, Combined Cycle Power Plant, Wine and Yatch. We split the data 20 times into train and test with 80% and 20% of the instances, respectively. The performance metrics employed are the RMSE, the test log-likelihood (LL) and the Continuous Ranked Probability Score (CRPS). CRPS is a proper scoring rule that can be used as an alternative metric of the accuracy of the predictive distribution (Gneiting & Raftery, 2007). In the case of fBNN we report results for a BNN prior since it performs better than a GP prior. In SIP we use 100 inducing points. In SIP_{NS} we implemented dropout with a rate of 0.05. SIP, fBNN and AADM use a *warm-up* period on which the KL in 12 is multiplied by a factor β that linearly increases from 0 to 1 for the first 20% of the total number of epochs. Finally, each method is trained until convergence on each dataset. More details on the experimental setup, as well as detailed results for each dataset, are provided in Appendix F.

To get an overall comparison among each method we have ranked them from best to worst for each data split: the best method gets rank 1, the second best rank 2, etc. These ranks are then averaged for each metric. Final results are displayed in Table 1, where the average rank across all datasets is shown for each method and performance metric (lower is better). We observe that for RMSE, SIP_{NS} is the best method. Similarly, SIP_{BNN} has the smallest average rank for test log-likelihood (followed by AADM) and CRPS (followed by VIP). VIP also performs very well in terms of RMSE. AADM is also performs well and is a flexible method that can also produce accurate predictive distributions (Santana & Hernández-Lobato, 2020). However, it performs inference in parameter space instead of function space. This may explain the advantage of SIP over AADM. In terms of CRPS, VIP is the second best method, probably because this metric is more influenced by the good result in the prediction error than the test log-likelihood. In summary, SIP gives competitive results with state-of-the-art methods for approximate inference in the context of NNs.

5.3. Experiments on Large Datasets

We also compared each method’s performance as a function of the computational training time. We considered two

datasets: (i) *Year Prediction MSD*, with 515,345 instances and 90 attributes (Dua & Graff, 2017); (ii) *Airplanes Delay* with 2,127,068 instances and 8 attributes (Santana & Hernández-Lobato, 2020). A random split of 10,000 instances is used as the test on which each performance metric is computed for each method over training. In SIP we employ 100 inducing points, with dropout rate 0.05 for SIP_{NS}. In VIP and fBNN_{GP}, a pre-training procedure on a subset of the data is carried out to estimate the hyper-parameters. The time for this is added to the total training time.

Figure 4 shows the results of each method for each performance metric as a function of training time. We have removed the first 3000 seconds of time for the sake of visualization. The fastest method in terms of performance vs. computational time is SIP_{NS}. In the case of the test log-likelihood (LL) and CRPS it obtains the best results after a small training time. In terms of LL, SIP is the clear winner in Airplanes. In Year, however, AADM also gives good results. The same happens in terms of CRPS, where SIP performs best with both priors, *i.e.*, BNN and NS. SIP_{NS} is faster than SIP_{BNN} by the more efficient sampling approach taking place in the prior in this model, which exploits a specific architecture that simplifies this task (see Mescheder et al. (2017)). The rest of methods perform similarly in general here, with fBNN_{GP} obtaining worse results than the other methods and far behind those of fBNN_{BNN}. Finally, VIP performs bad in general. Here, VIP under-estimates the noise variance producing too confident predictions that result in low CRPS and LL (worse than the ones of the other methods, not shown in the figures). Nevertheless, VIP’s RMSE is good in general. SIP and AADM tend to perform better in terms of LL and CRPS than RMSE, which is targeted for medium values of α (Santana & Hernández-Lobato, 2020) (smaller α values target the RMSE). In terms of performance vs. training time, SIP_{NS} is the fastest method. However, given enough training time, SIP_{BNN} outperforms the other methods in terms of LL and CRPS. In terms of RMSE it performs similarly to the other methods in Year. These experiments confirm again that SIP is competitive with state-of-the-art methods.

6. Conclusions

We have proposed SIP for approximate inference. SIP can be used in several models (*e.g.*, Bayesian NNs and neural

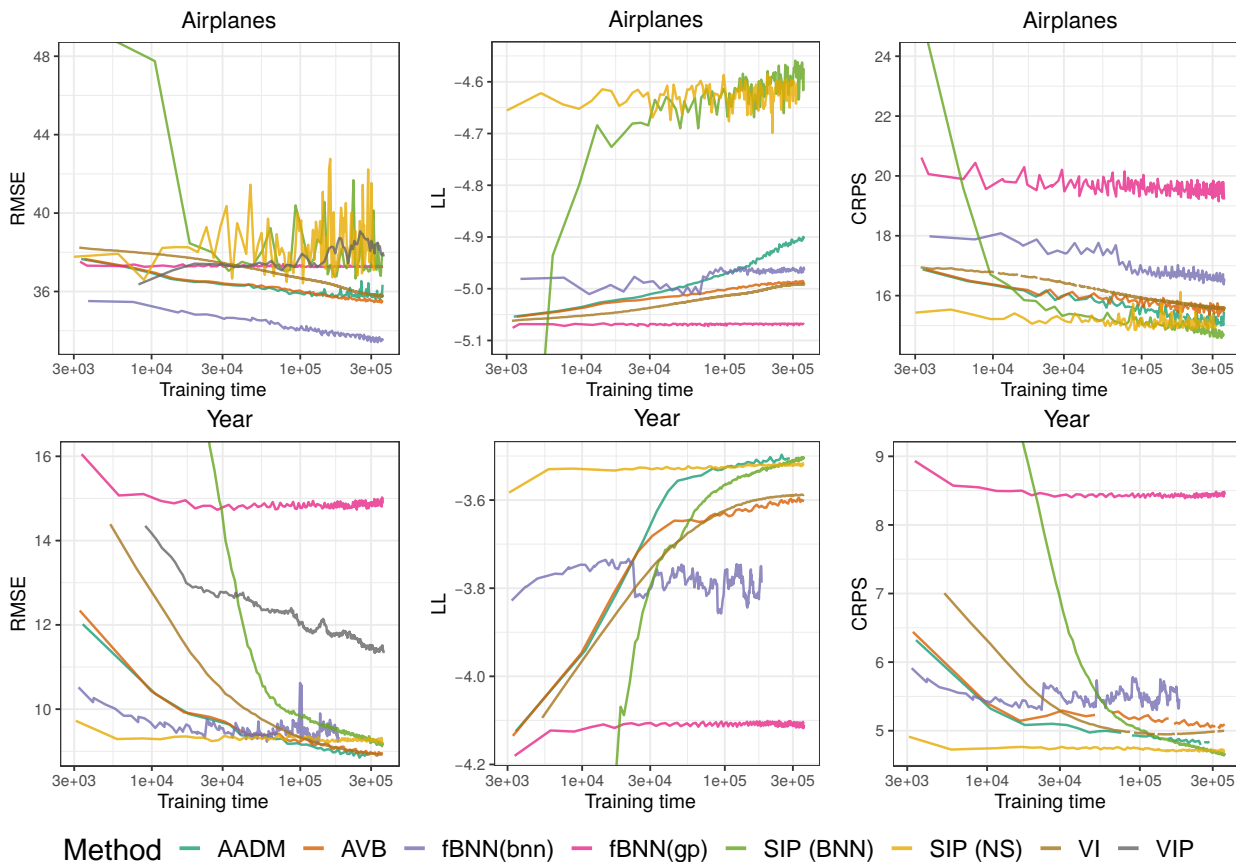


Figure 4. Performance as a function of training time (in seconds) in a log scale for each dataset, method and metric. Airplanes (first row) and Year (second row). The legend of each method is shown at the bottom. Best seen in color.

samplers). Moreover, SIP can adjust the prior parameters to the data. It can also use a flexible implicit process (IP) to approximate the posterior distribution. Current methods cannot perform these two tasks simultaneously. Importantly, SIP can generate flexible predictive distributions that capture complicated patterns in the data such as bimodality or heteroscedasticity (see Appendix D.2). We have evaluated SIP on several tasks for regression. It gives similar and often better results, in terms of several performance metrics, than state-of-the-art methods for approximate inference in the context of Bayesian NNs. SIP is scalable and can be used in datasets with millions of instances. This is achieved by a sparse approximation based on inducing points similar to the one often used in GPs. Our experiments also show that SIP can learn a sensible location for the inducing points. A limitation of SIP is, however, that it requires the evaluation of complex conditional distributions. Nevertheless, they can be approximated by a GP with the same covariances as the prior IP. The covariances can be estimated via Monte Carlo methods. SIP could have an important societal impact, specially when accurate predictive distributions are critical. For example, when the decisions made can have an influence on

people’s life, such as in autonomous vehicles (McAllister et al., 2017) or in medical diagnosis tools (Sajda, 2006).

Acknowledgements

The authors gratefully acknowledge the use of the facilities of Centro de Computacion Cientifica (CCC) at Universidad Autónoma de Madrid. SRS acknowledges the Spanish Ministry of Economy for the FPI SEV-2015-0554-16-4 Ph.D. grant. DHL and SRS also acknowledge financial support from Spanish Plan Nacional I+D+i, PID2019-106827GB-I00/AEI/10.13039/501100011033. BZ has been supported by the Programa Atracción de Talento de la Comunidad de Madrid under grant n. 2017- T2/TIC-5455, from the Comunidad de Madrid/UAM “Proyecto de Jóvenes Investigadores” grant n. SI1/PJI/2019-00294, as well as from Spanish “Proyectos de I+D de Generación de Conocimiento” via grants PGC2018-096646-A-I00 and PGC2018-095161-B-I00. BZ finally acknowledge the support from Generalitat Valenciana through the plan GenT program (CIDE-GENT/2020/055).

References

- Blundell, C., Cornebise, J., Kavukcuoglu, K., and Wierstra, D. Weight uncertainty in neural network. In *International Conference on Machine Learning*, pp. 1613–1622, 2015a.
- Blundell, C., Cornebise, J., Kavukcuoglu, K., and Wierstra, D. Weight uncertainty in neural network. In *International Conference on Machine Learning*, pp. 1613–1622. PMLR, 2015b.
- Burt, D. R., Ober, S. W., Garriga-Alonso, A., and van der Wilk, M. Understanding variational inference in function-space. In *Third Symposium on Advances in Approximate Bayesian Inference*, jan 2021. URL <https://arxiv.org/abs/2011.09421>.
- Coker, B., Pan, W., and Doshi-Velez, F. Wide mean-field variational Bayesian neural networks ignore the data. *arXiv preprint arXiv:2106.07052*, 2021.
- Damianou, A. and Lawrence, N. D. Deep gaussian processes. In *Artificial intelligence and statistics*, pp. 207–215. PMLR, 2013.
- Depeweg, S., Hernández-Lobato, J. M., Doshi-Velez, F., and Udluft, S. Learning and policy search in stochastic dynamical systems with Bayesian neural networks. *arXiv preprint arXiv:1605.07127*, 2016.
- Dua, D. and Graff, C. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- Foong, A., Burt, D., Li, Y., and Turner, R. On the expressiveness of approximate inference in Bayesian neural networks. In *Advances in Neural Information Processing Systems*, pp. 15897–15908, 2020.
- Gal, Y. *Uncertainty in deep learning*. PhD thesis, PhD thesis, University of Cambridge, 2016.
- 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. PMLR, 2016.
- Gneiting, T. and Raftery, A. E. Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association*, 102(477):359–378, 2007.
- Graves, A. Practical variational inference for neural networks. In *Advances in neural information processing systems*, pp. 2348–2356, 2011.
- Hensman, J., Fusi, N., and Lawrence, N. D. Gaussian processes for big data. In *Uncertainty in Artificial Intelligence*, pp. 282–290, 2013.
- Hernández-Lobato, J. M. and Adams, R. P. Probabilistic backpropagation for scalable learning of Bayesian neural networks. In *International Conference on Machine Learning*, pp. 1861–1869, 2015.
- Hernández-Lobato, J. M., Li, Y., Rowland, M., Hernández-Lobato, D., Bui, T. D., and Turner, R. E. Black-box α -divergence minimization. In *International Conference on Machine Learning*, pp. 1511–1520, 2016.
- Immer, A., Korzepa, M., and Bauer, M. Improving predictions of Bayesian neural nets via local linearization. In *International Conference on Artificial Intelligence and Statistics*, pp. 703–711, 2021.
- Jordan, M. I., Ghahramani, Z., Jaakkola, T. S., and Saul, L. K. An introduction to variational methods for graphical models. *Machine learning*, 37:183–233, 1999.
- Knoblauch, J., Jewson, J., and Damoulas, T. Generalized variational inference: Three arguments for deriving new posteriors. *arXiv preprint arXiv:1904.02063*, 2019.
- Li, Y. and Liu, Q. Wild variational approximations. In *NIPS workshop on advances in approximate Bayesian inference*, 2016.
- Liu, Q. and Wang, D. Stein variational gradient descent: A general purpose Bayesian inference algorithm. In *Advances In Neural Information Processing Systems*, pp. 2378–2386, 2016.
- Ma, C. and Hernández-Lobato, J. M. Functional variational inference based on stochastic process generators. In *Advances in Neural Information Processing Systems*, 2021.
- Ma, C., Li, Y., and Hernández-Lobato, J. M. Variational implicit processes. In *International Conference on Machine Learning*, pp. 4222–4233, 2019.
- MacKay, D. J. A practical Bayesian framework for back-propagation networks. *Neural computation*, 4(3):448–472, 1992.
- McAllister, R., Gal, Y., Kendall, A., Van Der Wilk, M., Shah, A., Cipolla, R., and Weller, A. Concrete problems for autonomous vehicle safety: Advantages of Bayesian deep learning. International Joint Conferences on Artificial Intelligence, Inc., 2017.
- Mescheder, L., Nowozin, S., and Geiger, A. Adversarial Variational Bayes: Unifying variational autoencoders and generative adversarial networks. In *International Conference on Machine Learning*, pp. 2391–2400, 2017.
- Neal, R. M. *Bayesian Learning for Neural Networks*. PhD thesis, 1994.

- Neal, R. M. et al. Mcmc using hamiltonian dynamics. *Handbook of markov chain monte carlo*, 2(11):2, 2011.
- Rasmussen, C. E. and Williams, C. K. I. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2006.
- Rezende, D. and Mohamed, S. Variational inference with normalizing flows. In *International Conference on Machine Learning*, pp. 1530–1538, 2015.
- Ritter, H., Kukla, M., Zhang, C., and Li, Y. Sparse uncertainty representation in deep learning with inducing weights. *arXiv preprint arXiv:2105.14594*, 2021.
- Rudner, T. G. J., Chen, Z., and Gal, Y. Rethinking function-space variational inference in bayesian neural networks. In *NeurIPS Symposium on Advances in Approximate Bayesian Inference*, 2020.
- Sajda, P. Machine learning for detection and diagnosis of disease. *Annu. Rev. Biomed. Eng.*, 8:537–565, 2006.
- Salimans, T., Kingma, D., and Welling, M. Markov chain Monte Carlo and variational inference: Bridging the gap. In *International Conference on Machine Learning*, pp. 1218–1226, 2015.
- Santana, S. R. and Hernández-Lobato, D. Adversarial α -divergence minimization for Bayesian approximate inference. *Neurocomputing*, 2020.
- Snelson, E. and Ghahramani, Z. Sparse Gaussian processes using pseudo-inputs. *Advances in neural information processing systems*, 18:1257–1264, 2005.
- Snelson, E., Rasmussen, C. E., and Ghahramani, Z. Warped Gaussian processes. *Advances in neural information processing systems*, 16:337–344, 2004.
- Sun, S., Zhang, G., Shi, J., and Grosse, R. Functional variational Bayesian neural networks. In *International Conference on Learning Representations*, 2019.
- Titsias, M. Variational learning of inducing variables in sparse Gaussian processes. In *Artificial intelligence and statistics*, pp. 567–574, 2009.
- Tran, D., Ranganath, R., and Blei, D. Hierarchical implicit models and likelihood-free variational inference. In *Advances in Neural Information Processing Systems*, pp. 5523–5533, 2017.
- Wenzel, F., Roth, K., Veeling, B., Swiatkowski, J., Tran, L., Mandt, S., Snoek, J., Salimans, T., Jenatton, R., and Nowozin, S. How good is the Bayes posterior in deep neural networks really? In *International Conference on Machine Learning*, pp. 10248–10259, 2020.

A. Gradients analysis

The definition of the original functional ELBO objective for approximate inference, using implicit distribution and a discriminator to estimate the KL divergence term, is the following:

$$\mathcal{L}(q) = \sum_{i=1}^N \mathbb{E}_{q_{\phi, \theta}} [\log p(y_i | f_i)] - \mathbb{E}_q [T_{\omega^*}(\mathbf{u})] \quad (15)$$

where $T_{\omega^*}(\mathbf{u})$ is the optimal discriminator, approximated by a deep neural network that tells apart samples from $q_{\phi}(\mathbf{u})$ and $p_{\theta}(\mathbf{u})$. However, as mentioned in the main text of the article, training the model employing this expression turns out ineffective in terms of updating the prior, since the gradients w.r.t the prior parameters of the second term are either zero or assumed to be negligible (Sun et al., 2019). To check this argument we performed a few experiments using SIP_{BNN} .

The gradients in the data term of (15) w.r.t. any parameter can easily be estimated by automated derivation methods. However, in order to make the comparisons, we need to estimate the gradients in the second term w.r.t. the same parameters as well, and since we are mainly interested in the parameters regarding the prior, they will be θ and the inducing points' locations $\bar{\mathbf{X}}$. Here, in order to obtain the gradients of the estimated KL term we must first obtain the optimal discriminator value, which results in the most precise estimate of the log-ratio between $q_{\phi}(\mathbf{u})$ and $p_{\theta}(\mathbf{u})$. Then, we perturb slightly the value of the selected parameter and retrain the discriminator, obtaining the new estimate for the log-ratio between the distributions and allowing us to estimate the gradient w.r.t. the parameter using finite differences. In both phases, we train the model until convergence to ensure that we have reached $T_{\omega^*}(\mathbf{u})$.

To compare different contributions to the total gradient of (15) we study separately θ and $\bar{\mathbf{X}}$. In the case of θ , since we use a 2-layered BNN with 50 units in each layer as prior, there are a lot of parameters to choose from. In order to select among them the most relevant ones, we have chosen the 500 parameters that showed the biggest contribution to the gradient in the data term. Afterwards we perform the previous finite differences procedure and compare the gradients obtained in both terms for the same parameters. On the other hand, for the gradients regarding $\bar{\mathbf{X}}$, we estimate the KL gradient for every inducing point, since in this case we are only employing 50 inducing points (as in the synthetic experiments). Therefore, we will be able to compare the gradients of both terms for every inducing point employed.




Figure 5. Comparison between the gradients of the data term (x-axis) and the two sets of prior parameters in the model, *i.e.* θ (left, in blue) and the inducing points' locations $\bar{\mathbf{X}}$ (right, in red). The scale of the x-axis is in both cases much bigger than the y-axis for visualization purposes. Error bars are included in both plots, although may not be visible when compared against the size of the points. Best seen in color.

In Figure 5 we show the comparisons between the gradients of the data term in the f-ELBO and the two possible gradient contributions from the prior parameters, *i.e.* θ (left image, in blue), and the inducing points' locations (right image, in red). We have included error bars, although specially in the second plot most of them are smaller than the size of the points themselves. In both figures, the x-axis represents the gradient in the data term, and the y-value for each point is the estimated gradient value for the KL term in (15) w.r.t. either θ and $\bar{\mathbf{X}}$. As can be seen, in both cases the x-axis has a wider range than the y-axis by a factor bigger than 10. As we mentioned, when comparing the gradient of the data term against the gradient of the KL term w.r.t. θ we selected the 500 parameters in θ that had the largest contributions to the data term, which explains the gap in $x = 0$. Nonetheless, we see here that in every case $|y| < 1$, which means that $\nabla_{\theta} \text{KL} \ll \nabla_{\theta}(\text{Data term})$. Moreover, in the second plot we see the same behavior, meaning that $\nabla_{\bar{\mathbf{X}}} \text{KL} \ll \nabla_{\bar{\mathbf{X}}}(\text{Data term})$.

SIP can be trained using this original formulation of the objective function. Although the resulting predictive distribution produces good results, these negligible contributions to the gradients makes the system unable to train its prior model. The results from this process, using the SIP_{BNN} model, are shown in Figure 6. There, the samples from the prior model can be seen to not follow the behavior of the original data, while the predictive distribution is capable of reproducing the original

bimodal behavior. This means that the prior is not being properly trained, although the posterior model is flexible enough to compensate for this fact and accurately follow the data behavior.




Figure 6. Results of training SIP with the original f-ELBO formulation with only one KL contribution for the bimodal data presented in the main text (*left*). The predictive distribution (*right*) reproduces the bimodal behavior correctly. However, the functions sampled from the final prior model (*center*) do not follow the same behavior of the data.

From these results we conclude that the contribution to the gradient of the original objective function from the prior parameters is much less relevant than the one coming from the data term. Therefore, we need to introduce the new $\mathcal{L}^*(q)$ objective, which allows us to train simultaneously both the prior and posterior parameters.

B. Adaptive contrast for SIP

Adaptive contrast (AC) is introduced by (Mescheder et al., 2017) as an attempt to improve the accuracy of the estimation of the log ratio between distributions in the KL. This estimation will follow from the result of the auxiliary discriminator problem, whose optima is T_{ω^*} . However, since the distributions being compared are usually very different, the discriminator has no problem telling apart samples from each of them, which does not encourage for an optimal fit of its parameters. Therefore, to make this task harder, an extra Gaussian distribution is introduced in (Mescheder et al., 2017). In our case, since the KL distribution is estimated between two IPs, we need to reformulate the original AC to accommodate for this fact. This will imply extending the original AC to use two discriminators instead of one. Let us define two Gaussian distributions \bar{q} and \bar{p} with the same moments as the samples from both q and p (our IP approximating posterior and IP prior)

$$\bar{q} \sim \mathcal{N}(\mu(q), \sigma^2(q)), \quad \bar{p} \sim \mathcal{N}(\mu(p), \sigma^2(p)) \quad (16)$$

with $\mu(\cdot)$ and $\sigma^2(\cdot)$ as the mean and variance across samples from q or p . Using this, we rewrite the two-KL part of the f-ELBO objective as the following:

$$\begin{aligned} \text{KL}(q|p) + \text{KL}(p|q) &= \mathbb{E}_q \left[\log \frac{q\bar{q}\bar{p}}{p\bar{q}\bar{p}} \right] + \mathbb{E}_p \left[\log \frac{p\bar{p}\bar{q}}{q\bar{p}\bar{q}} \right] \\ &= \mathbb{E}_q \left[\log \frac{q}{\bar{q}} \right] + \mathbb{E}_q \left[\log \frac{\bar{p}}{p} \right] + \mathbb{E}_q \left[\log \frac{\bar{q}}{\bar{p}} \right] \\ &\quad + \mathbb{E}_p \left[\log \frac{p}{\bar{p}} \right] + \mathbb{E}_p \left[\log \frac{\bar{q}}{q} \right] + \mathbb{E}_p \left[\log \frac{\bar{p}}{\bar{q}} \right] \\ &= \mathbb{E}_q [T(q, \bar{q}) - T(p, \bar{p})] + \mathbb{E}_p [T(p, \bar{p}) - T(q, \bar{q})] \\ &\quad + \mathbb{E}_q \left[\log \frac{\bar{q}}{\bar{p}} \right] + \mathbb{E}_p \left[\log \frac{\bar{p}}{\bar{q}} \right]. \end{aligned} \quad (17)$$

Now we will have to employ **two** discriminators, one to separate samples from q and \bar{q} ($T(q, \bar{q})$), and another to separate samples from p and \bar{p} ($T(p, \bar{p})$). The two last contributions to this expression are given by the log-ratio of two Gaussian distributions with given mean and variances, and are thus tractable and have a closed-form solution.

In practice, both discriminators needed here are estimated once with the expected value w.r.t. samples from q and from p . To help with this task, one could standardize the samples from q and p in beforehand, following the description in (Mescheder et al., 2017; Santana & Hernández-Lobato, 2020)

$$\text{KL}(q|p) + \text{KL}(p|q) = \text{KL}(q_0|p_q) + \text{KL}(p_0|q_p), \quad (18)$$

where

$$q_0 = \frac{q - \mu(q)}{\sigma(q)}, \quad p_0 = \frac{p - \mu(p)}{\sigma(p)}, \quad q_p = \frac{q - \mu(p)}{\sigma(p)}, \quad p_q = \frac{p - \mu(q)}{\sigma(q)}. \quad (19)$$

This simplifies the involved calculations for the discriminators involved, since now several terms employ standardized Gaussian distributions.

We opt to include the description of AC here for the sake of completeness, although in the experiments conducted in the main text, we found AC produced overfitting in the training of the prior. This would then require regularization techniques such as dropout or the *llh-regularization* used in VIP (Ma et al., 2019). We have tested these models as well, and thus far we have seen they also provide good performance. However, we think the added difficulties of training an extra discriminator and the need for a regularization term surpasses the benefits for this extra additions to the original SIP model.

C. HMC implementation and remarks

Ideally, the ground truth for Bayesian inference is provided by Markov Chain Monte Carlo (MCMC) sampling methods, where a Markov Chain is designed to converge to the true underlying distribution after some -potentially very large- number of steps. In particular, the Hybrid -or Hamilton- Monte Carlo (HMC) method (Neal et al., 2011) is one of the most popular methods nowadays. The HMC method avoids the inefficient random-walk behavior of other MCMC methods by using the gradient information of the target distribution. This is done by solving the Hamilton equations, where a set of latent "momentum" variables are introduced and associated in a one-to-one correspondence to the original variables making up the parameter space of the model. Specific implementation of HMC involves choosing a concrete choice for the "kinetic energy" of the Hamiltonian, which is equivalent to the choice for the distribution of the momentum variables, as well as a particular algorithm for solving the Hamilton equations. A common practice for the former which has been adopted in this work, is to assume a multivariate standard Gaussian. On the other hand, we also follow the standard practice and solve the Hamiltonian equations with the so-called leapfrog algorithm, which solves the system of differential equations along a trajectory consisting of L integration steps of size ϵ . For the datasets below, we have found that $L=25$ and $\epsilon = 5 \cdot 10^{-5}$ give sufficiently good results. Changing the covariance matrix of the momentum variables (e.g. by considering a more general diagonal matrix) gives roughly the same overall performance after properly tuning the above leapfrog parameters. In this work we have found that, for all the problems in consideration, a Markov chain of 10,000 steps was enough to both attain the equilibrium distribution.

In practice, however, HMC may not be able to capture the true underlying distribution. Beyond the obvious limitation of computational resources, which could prevent the Markov chain to converge as desired or to have a sufficiently large number of truly independent samples, there is another reason typically overlooked by the common lore. That is, in order for the HMC to attain the true target distribution the assumed model should be the correct one. For example, in the case of the bimodal dataset analyzed below, if the proposed joint distribution of the parameters and the data is unimodal, then the HMC will never converge to the desired distribution, even in the limit of infinite number of Markov chain steps. An analogous situation is encountered with the heterocedastic dataset shown above. In general, if the data comes from a distribution having features not captured by the proposed model, the MCMC will not deliver the desired outcome. In Figure 1 of the main document we observe the HMC applied to the bimodal dataset. As already commented there, it does not capture the bimodality, simply because the model assumed a Gaussian prior and a Gaussian likelihood. This cannot explain the bimodal data. In a similar way, in Figure 8 we observe that the HMC is unable to capture the heteroscedasticity of the data, because the model assumes homocedasticity instead.

D. Extra synthetic data comparisons

D.1. Alternative setups for bimodal data

The experiments conducted with synthetic data in the main text were all performed using the same prior model, a BNN with 2 layers with 50 units per layer. However, to produce those comparisons we have made changes on both methods that could

affect their final results: the original code for VIP has a regularization term that was removed for the comparisons there, helping both the prior and the predictive distribution to represent heterocedastic behavior (Ma et al., 2019). On the other hand, fBNN can employ a GP prior (or a sparse GP), which would have to be trained in beforehand (Sun et al., 2019). For the sake of completeness, we have run these tests as well with the same bimodal data being used in the main text (Depeweg et al., 2016; Santana & Hernández-Lobato, 2020).

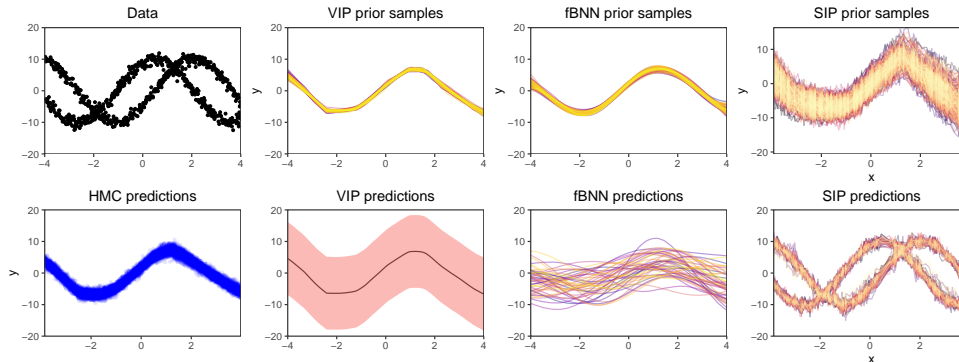


Figure 7. Comparison between methods using bimodal data and alternative setups for VIP and fBNN: VIP includes the regularization term as in the UCI regression experiments, and fBNN uses the GP prior. The first column represents the data distribution (first row, in black) and the HMC predictive distribution (second row, in blue). The remaining figures represent the prior samples (first row) and the predictive distribution (second row) for each method, namely VIP* (with regularizer), fBNN_{GP} and SIP_{BNN}. For the predictive distribution of VIP* the black line represents the predictive mean and the shaded area represents the $\pm 2\sigma$ region. For comparison w.r.t. the original image in the main text, the y-axis range has been increased here so that every plot could be easily seen. Best seen in color.

In Figure 7 we can see the comparisons between methods using these alternative setups, namely VIP with the regularization term (denoted as VIP*) and fBNN with the GP prior (fBNN_{GP}). As in the figure in the main text, the first column represents the training data (first row, in black), and the predictive distribution obtained from applying HMC to the same system on which SIP is implemented, sharing the trained prior parameters (second row, in blue). The other figures represent both the prior samples (first row) or the predictive distribution (second row) for VIP*, fBNN_{GP} and SIP_{BNN} respectively. In the case of fBNN_{GP} we see that the prior is now being trained (although that needs to happen separately from the rest of the model), resulting in samples that fit better the behavior of the data here. The predictive distribution follows the mean of the data as it did in the case of fBNN_{BNN}, and although the noise is slightly smaller in this case, it is not able to reproduce the original bimodality. However, an important detail when comparing these figures against the ones present in the main text is that the y axis range has been increased here w.r.t. the one used in the original one, which may make these figures seem more concentrated towards the mean of the data.

On the other hand, for VIP* we can clearly see that both figures are now not showing any of the heterocedastic behavior they had in the main text. In this case, the prior, as well as the predictive distribution, fit closely the mean of the data, and the regions of $\pm 2\sigma$ in the latter have a fixed width, unlike in the original figure. Thus, removing the regularizer has helped the method to fit a more expressive prior and to obtain better predictive distribution samples, since introducing it hinders the expressiveness of the model overall. However, since this is the original setup the authors in (Ma et al., 2019) propose, this is the one we employed in both the UCI multivariate regression problems as well as in the convergence experiments. Finally, the good performance of the method regarding the RMSE metric can be explained by its focus on adjusting the predictions to the mean of the training data.

D.2. Heterocedastic data

To test the ability of SIP to reproduce other complex features present in the training data, we have employed as well an heterocedastic dataset (Depeweg et al., 2016; Santana & Hernández-Lobato, 2020). We constructed this dataset by uniformly sampling 1000 values for x between $[-4, 4]$ and then obtaining y for each sample as

$$y = 7 \sin(x) + \epsilon \sin(x) + 10, \quad \epsilon \sim \mathcal{N}(0, \sigma = 2). \quad (20)$$

To perform these experiments we employ the same setup as the one in the main text for the bimodal problem.

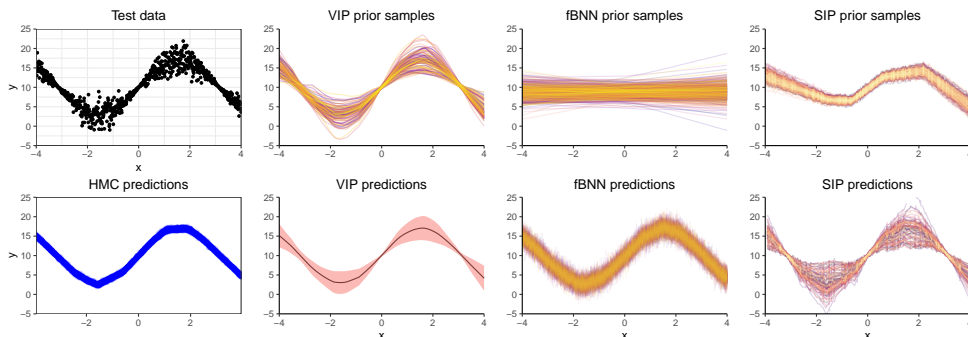


Figure 8. Comparison between methods using heterocedastic data. The first column represents the data distribution (first row, in black) and the HMC predictive distribution (second row, in blue). The remaining figures represent the prior samples (first row) and the predictive distribution (second row) for each method, namely VIP, fBNN_{BNN} and SIP_{BNN}. For the predictive distribution of VIP the line represents the predictive mean and the shaded area represents the $\pm 2\sigma$ region. Best seen in color.

In Figure 8 we present the results of training the same models we used in the main text for the bimodal dataset (HMC, VIP, fBNN_{BNN} and SIP_{BNN}) but with the new heterocedastic synthetic data (represented in the first column, first row, in black). Here, as was the case earlier, HMC (first column, second row, in blue) is unable to reproduce the heterocedastic behavior. Even though it uses the trained prior from SIP, since the selected NN model is not originally capable of reproducing this behavior, HMC is unable to reproduce it either: the prior does not behave in an heterocedastic manner, and neither does the likelihood, resulting in a non-heterocedastic predictive distribution. For the rest of the methods we have the prior samples (first row) and the predictive distribution samples (second row) to compare against them. In the case of VIP (second column), opposite to what is done in Section D.1, the regularization term here is turned off again, thus obtaining heterocedastic behavior both in the prior and predictive distributions. If the regularization term is turned on, however, the heteroscedasticity here is lost, as happened in Section D.1. For fBNN (third column) we see that the BNN prior is not being trained. Its predictive distribution follows the original data mean closely although without any heteroscedasticity. Finally, in SIP (last column) we can see that the prior behaves somewhat similarly to the original data distribution mean, as it did in the bimodal case. In this case, SIP’s predictive distribution also closely follows that of the original data, with heterocedastic behavior as well. In this case, if VIP uses the regularization term, SIP is the only method able to reproduce this behavior, resulting in a better final predictive distribution than the one provided by HMC. This implies what we already mentioned in the main text: when selecting the wrong model, HMC may be unable to reproduce important features from the data that SIP is capable of capturing in its predictive distribution, making it much less susceptible to implicit bias errors caused by choosing the incorrect model for a given problem.

D.3. Inducing points evolution

To complement the analysis of the evolution of the location of the inducing points in the main text, we conducted an extra experiment using an alternative dataset. This dataset is generated by sampling uniformly 1000 values for x between $[-4, 5]$, and then obtaining y from the following definition:

$$y = \begin{cases} 10 + \epsilon, & x < 0 \\ 10(1 + \sin(x)) + \epsilon, & x \geq 0, \end{cases} \quad (21)$$

with $\epsilon \sim \mathcal{N}(0, \sigma = 2)$. The piece-wise definition of y is constructed so that y is continuous but with a non-defined derivative in $x = 0$. We train SIP with this dataset until it converges, and register the location of the inducing points for each epoch of training (2000 here). We employ 50 inducing points here for visualization purposes.

In Figure 9 we see the results of training SIP with this alternative dataset, using the same representation we employed in the experiment included in the main text: the first part of the figure contains the original data (black dots), the predictive distribution samples (with the mean as a blue line), and the initial and final location of the inducing points represented as crosses both in the top and bottom of the figure respectively. The second plot represents the location of each inducing point for each epoch in the y-axis (scaled by a factor of 100, ranging from 0 to 2000). Moreover, in this case, we have set the starting positions of the 50 inducing points in random locations throughout the whole range of the training set. In the figure,




Figure 9. Results of the SIP method applied to the piece-wise defined synthetic dataset. The first figure represents the samples from the predictive distribution of the method, with the mean as the blue line and the training points as black dots. The crosses at the top and bottom represent the starting and finishing locations of the inducing points. The second figure represent the locations of the inducing points for every training epoch in the y-axis. The number of epochs is scaled by $1e2$, ranging from 0 to 2000. Best seen in color.

we see that the predictive distribution seems to follow closely that of the original data. However, the most interesting fact here is that the inducing points that started in the region where the data is constant ($x < 0$) seem to have a tendency to concentrate around the region of $x = 0$, moving right to the matching point between both pieces of the function of y . Due to the simple nature of the data before this point, the model focuses the inducing points originally positioned in this region and places them closer to $x = 0$, helping it when modeling the change in behavior between the two parts of y . On the other side, for the $x > 0$ region, the inducing points are distributed more homogeneously to model the sine behavior. Furthermore, the right-most points seem to stray away from the region of the dataset, which is caused by employing far more inducing points than needed in such a simple dataset. This has been done for illustration purposes only, which also causes the overlapping of points across the dataset when there are more than what are needed in the same region. From all of this, we conclude that SIP seems able to distinguish between simpler and more complex regions of the data, and employs its resources effectively to model both at the same time.

E. Further details for the convergence experiments

The convergence experiments have been conducted employing the same computational resources for each method: Each of them employed on its own 2 CPU Intel(R) Xeon(R) Gold 5218 CPU at 2.30GHz (16 cores, 22 Mb L3 cache), [32 cores in total], with 192 GB RAM at 2,4 GHz. In the experiments, only the training time is reported. The same setup is used for both experiments (*Airplanes delay* and *Year Prediction MSD* datasets).

To obtain the results shown in the main text, we have chosen to execute VIP without the regularization term described in the supplementary material for the article (Ma et al., 2019). After conducting several tests, we have concluded that the effect of such regularization term is negligible for these experiments given the size of the datasets.

Finally, the fast rate of convergence shown by SIP_{NS} can be explained by its architecture. The NS follows the system described in (Mescheder et al., 2017) to be able to efficiently sample the moments of the approximate distribution for all data points in each mini-batch, which is a time-consuming task in other cases. This allows SIP to be much faster when employing the NS model for the prior, in comparison to what is shown for SIP_{BNN} , although given enough time the latter usually surpasses the performance of SIP_{NS} .

F. Regression results for UCI datasets

We include here the mean results for the UCI datasets for each method and metric. Each method is trained on 20 different random train/test splits of the dataset, and the values reported here represent the mean performance of each method in each dataset for the given metric in each case. The *Continuous Ranked Probability Score* (CRPS) is estimated using the *proprscoring*¹ package, both in cases where the predictive distribution is implicit (*e.g.* SIP, AVB and AADM), or when it has an explicit form (*e.g.* VI, VIP). For each dataset, the winning method is highlighted in bold while the one in second place is marked in red (in LL, the higher the better; in RMSE and CRPS, the lower the better) This is done as well for the rankings. The final ranking analysis follows the description of the main text, placing SIP as the leading method for each metric, with the BNN prior for LL and CRPS, and for RMSE with the NS prior.

Table 2. LL results in UCI datasets

Dataset	BBB	AVB	AADM	fBNN _{BNN}	VIP	SIP _{BNN}	SIP _{NS}
boston	-2.816±0.183	-2.409±0.153	-2.348±0.151	-3.193±0.687	-2.559±0.409	-2.51±0.321	-2.595±0.521
ccpp	-2.844±0.0449	-2.819±0.0499	-2.816±0.0495	-2.789±0.0805	-4.018±0.965	-2.761±0.0507	-2.775±0.0494
concrete	-3.331±0.0567	-2.976±0.0928	-2.896±0.102	-3.514±0.491	-3.208±1.04	-3.364±0.424	-2.934±0.165
ee	-2.155±0.0733	-1.678±0.0825	-1.239±0.101	-1.284±0.0485	-0.995±0.146	-1.002±0.292	-1.106±0.0888
wine	-0.979±0.0481	-0.955±0.0479	-0.949±0.0476	-19.77±4.32	-0.992±0.0691	-0.972±0.0683	-1.149±0.168
yatch	-2.588±0.0968	-1.742±0.085	-1.559±0.157	-2.114±0.0299	-0.041±0.511	-0.207±0.427	-0.559±0.172
kin8nm	1.114±0.0204	1.283±0.0354	1.295±0.0312	0.999±0.0687	0.979±0.0513	1.178±0.0264	1.223±0.023
naval	6.544±0.0724	6.097±0.628	6.483±0.366	6.157±0.235	5.944±1.8	7.28±0.0598	5.651±0.0445
Avg. ranking	5.39±0.05	3.91±0.06	2.82±0.06	5.61±0.07	3.70±0.12	2.76±0.08	3.81±0.07

Table 3. RMSE results in UCI datasets

Dataset	BBB	AVB	AADM	fBNN _{BNN}	VIP	SIP _{BNN}	SIP _{NS}
boston	3.52 ± 0.804	2.53 ± 0.461	2.55 ± 0.481	5.66 ± 10.2	2.62 ± 0.697	2.88 ± 0.616	2.29 ± 0.452
ccpp	4.15 ± 0.201	4.05 ± 0.206	4.04 ± 0.207	3.77 ± 0.205	3.88 ± 0.224	3.92 ± 0.214	3.87 ± 0.198
concrete	6.64 ± 0.562	4.95 ± 0.569	4.92 ± 0.584	4.78 ± 0.694	4.53 ± 0.548	4.9 ± 0.825	4.45 ± 0.56
ee	2.02 ± 0.231	1.32 ± 0.133	1.77 ± 0.208	0.639 ± 0.116	0.789 ± 0.13	1.47 ± 0.316	0.836 ± 0.106
wine	0.649 ± 0.0388	0.633 ± 0.0342	0.631 ± 0.0339	0.78 ± 0.0645	0.622 ± 0.0437	0.636 ± 0.0367	0.625 ± 0.0452
yatch	2.3 ± 0.98	1.1 ± 0.42	1.3 ± 0.64	0.82 ± 0.31	0.58 ± 0.22	0.36 ± 0.11	0.54 ± 0.27
kin8nm	0.079 ± 0.0022	0.067 ± 0.0024	0.067 ± 0.0023	0.074 ± 0.0022	0.079 ± 0.0025	0.075 ± 0.0021	0.07 ± 0.0021
naval	(3.02 ± 4.73)e⁻⁵	(4.96 ± 0.21) e ⁻⁵	(3.61 ± 0.17) e ⁻⁵	(3.78 ± 2.17) e ⁻⁵	(3.18 ± 1.06) e ⁻⁵	(1.73 ± 0.15) e⁻⁵	(6.37 ± 0.98) e ⁻⁵
Avg. ranking	6.15±0.04	4.14±0.08	3.97±0.09	3.82±0.10	3.29±0.09	3.53±0.08	3.10±0.09

Table 4. CRPS results in UCI datasets

Dataset	BBB	AVB	AADM	fBNN _{BNN}	VIP	SIP _{BNN}	SIP _{NS}
boston	1.859±0.363	1.475±0.281	1.456±0.298	1.861±1.137	1.348±0.241	1.573±0.278	1.45±0.259
ccpp	2.641±0.103	2.615±0.102	2.608±0.104	2.013±0.08	2.229±0.081	2.114±0.082	2.2±0.072
concrete	3.755±0.294	2.811±0.262	2.615±0.214	2.35±0.301	2.273±0.196	2.485±0.238	2.683±0.24
ee	1.121±0.139	0.703±0.07	0.694±0.082	0.41±0.03	0.387±0.061	0.532±0.078	0.591±0.078
wine	0.362±0.021	0.356±0.019	0.363±0.019	0.479±0.049	0.341±0.02	0.353±0.017	0.366±0.028
yatch	1.026±0.185	0.475±0.105	0.519±0.128	0.838±0.056	0.212±0.071	0.154±0.03	0.384±0.094
kin8nm	0.047±0.001	0.042±0.002	0.041±0.001	0.042±0.002	0.044±0.001	0.042±0.001	0.047±0.002
naval	(1.6±0.28) e⁻⁴	(3.46 ± 1.56) e ⁻⁴	(2.21 ± 1.22) e ⁻⁴	(2.56 ± 0.95) e ⁻⁵	(1.81 ± 0.68) e ⁻⁵	(9.4 ± 0.70) e⁻⁵	(6.8 ± 0.28) e ⁻⁵
Avg. ranking	5.84±0.05	4.47±0.08	4.03±0.08	3.97±0.07	2.57±0.08	2.45±0.08	4.67±0.08

¹<https://pypi.org/project/proprscoring/>