
Parametric Gaussian Process Regressors

Martin Jankowiak¹ Geoff Pleiss² Jacob R. Gardner³

Abstract

The combination of inducing point methods with stochastic variational inference has enabled approximate Gaussian Process (GP) inference on large datasets. Unfortunately, the resulting predictive distributions often exhibit substantially underestimated uncertainties. Notably, in the regression case the predictive variance is typically dominated by observation noise, yielding uncertainty estimates that make little use of the input-dependent function uncertainty that makes GP priors attractive. In this work we propose two simple methods for scalable GP regression that address this issue and thus yield substantially improved predictive uncertainties. The first applies variational inference to FITC (Fully Independent Training Conditional; Snelson et. al. 2006). The second bypasses posterior approximations and instead directly targets the posterior predictive distribution. In an extensive empirical comparison with a number of alternative methods for scalable GP regression, we find that the resulting predictive distributions exhibit significantly better calibrated uncertainties and higher log likelihoods—often by as much as half a nat per datapoint.

1. Introduction

Machine learning is finding increasing use in applications where autonomous decisions are driven by predictive models. For example, machine learning can be used to steer dynamic load balancing in critical electrical systems, and autonomous vehicles use machine learning algorithms to detect and classify objects in unpredictable weather conditions and decide whether to brake. As machine learning models increasingly become deployed as components in

¹The Broad Institute, Cambridge, MA, USA ²Dept. of Computer Science, Cornell University, Ithaca, NY, USA ³Dept. of Computer and Information Science, University of Pennsylvania, Philadelphia, PA, USA . Correspondence to: Martin Jankowiak <jankowiak@gmail.com>.

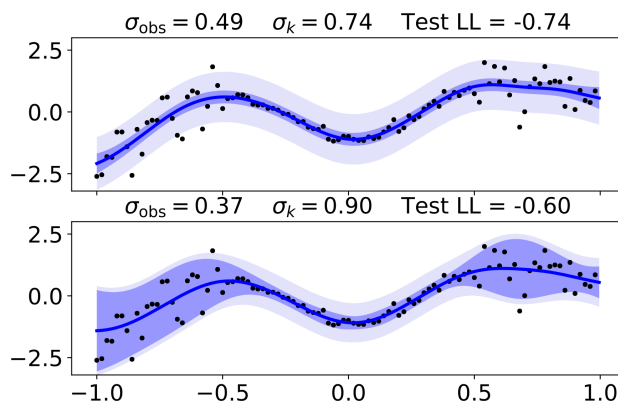


Figure 1. We depict GP regressors fit to a heteroscedastic dataset using two different inference algorithms. Solid lines depict mean predictions and $2\text{-}\sigma$ uncertainty bands are in blue. In the lower panel, fit with the PPGPR approach described in Sec. 3.2, significant use is made of input-dependent function uncertainty (dark blue), while in the upper panel, fit with variational inference (see Sec. 2.3.1), the predictive uncertainty is dominated by the observation noise σ_{obs}^2 (light blue) and the kernel scale σ_k is smaller.

larger decision making pipelines, it is essential that models be able to reason about uncertainty and risk. Techniques drawn from probabilistic machine learning offer the ability to deal with these challenges by offering predictive models with simple and interpretable probabilistic outputs.

Recent years have seen extensive use of variational inference (Jordan et al., 1999) as a workhorse inference algorithm for a variety of probabilistic models. The popularity of variational inference has been driven by a number of different factors, including: i) its amenability to data subsampling (Hoffman et al., 2013); ii) its applicability to black-box non-conjugate models (Kingma and Welling, 2013; Rezende et al., 2014); and iii) its suitability for GPU acceleration.

The many practical successes of variational inference notwithstanding, it has long been recognized that variational inference often results in overconfident uncertainty estimates.¹ This problem can be especially acute for Gaussian Process (GP) models (Bauer et al., 2016). In particular

¹For example see Turner and Sahani (2011) for a discussion of this point in the context of time series models.

GP regressors fit with variational inference tend to apportion most of the predictive variance to the input-independent observation noise, making little use of the input-dependent function uncertainty that makes GP priors attractive in the first place (see Fig. 4 in Sec. 5 for empirical evidence). As we explain in Sec. 3.3, this tendency can be understood as resulting from the asymmetry with which variational inference—through its reliance on Jensen’s inequality—treats the various contributions to the uncertainty in output space, in particular its asymmetric treatment of the observation noise.

In this work we propose two simple solutions that correct this undesirable behavior. In the first we apply variational inference to the well known FITC (Fully Independent Training Conditional; Snelson and Ghahramani (2006)) method for sparse GP regression. In the second we directly target the predictive distribution—bypassing posterior approximations entirely—to formulate an objective function that treats the various contributions to predictive variance on an equal footing. As we show empirically in Sec. 5, the predictive distributions resulting from these two parametric approaches to GP regression exhibit better calibrated uncertainties and higher log likelihoods than those obtained with existing methods for scalable GP regression.

2. Background

This section is organized as follows. In Sec. 2.1-2.2 we review the basics of Gaussian Processes and inducing point methods. In Sec. 2.3 we review various approaches to scalable GP inference that will serve as the baselines in our experiments. In Sec. 2.4 we describe the predictive distributions that result from these methods. Finally in Sec. 2.5 we review FITC (Snelson and Ghahramani, 2006)—an approach to sparse GP regression that achieves scalability by suitably modifying the regression model—as it will serve as one of the ingredients to the approach introduced in Sec. 3.1. We also use this section to establish our notation.

2.1. Gaussian Process Regression

In probabilistic modeling Gaussian Processes offer powerful non-parametric function priors that are useful in a variety of regression and classification tasks (Rasmussen, 2003). For a given input space \mathbb{R}^d GPs are entirely specified by a covariance function or kernel $k : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ and a mean function $\mu : \mathbb{R}^d \rightarrow \mathbb{R}$. Different choices of μ and k allow the modeler to encode prior information about the generative process. In the prototypical case of univariate regression the joint density takes the form²

$$p(\mathbf{y}, \mathbf{f}|\mathbf{X}) = p(\mathbf{y}|\mathbf{f}, \sigma_{\text{obs}}^2)p(\mathbf{f}|\mathbf{X}, k, \mu) \quad (1)$$

²In the following we will suppress dependence on the kernel k and the mean function μ .

where \mathbf{y} are the real-valued targets, \mathbf{f} are the latent function values, $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^N$ are the N inputs with $\mathbf{x}_i \in \mathbb{R}^d$, and σ_{obs}^2 is the variance of the Normal likelihood $p(\mathbf{y}|\cdot)$. The marginal likelihood takes the form

$$p(\mathbf{y}|\mathbf{X}) = \int d\mathbf{f} p(\mathbf{y}|\mathbf{f}, \sigma_{\text{obs}}^2)p(\mathbf{f}|\mathbf{X}) \quad (2)$$

Eqn. 2 can be computed analytically, but doing so is computationally prohibitive for large datasets. This is because the cost scales as $\mathcal{O}(N^3)$ from the terms involving \mathbf{K}_{NN}^{-1} and $\log\det \mathbf{K}_{NN}$ in Eqn. 2, where $\mathbf{K}_{NN} = k(\mathbf{X}, \mathbf{X})$. This necessitates approximate methods when N is large.

2.2. Sparse Gaussian Processes

Over the past two decades significant progress has been made in scaling Gaussian Process inference to large datasets. The key technical innovation was the development of inducing point methods (Snelson and Ghahramani, 2006; Titsias, 2009; Hensman et al., 2013), which we now review. By introducing inducing variables \mathbf{u} that depend on variational parameters $\{\mathbf{z}_m\}_{m=1}^M$, where $M = \dim(\mathbf{u}) \ll N$ and with each $\mathbf{z}_m \in \mathbb{R}^d$, we augment the GP prior as follows:

$$p(\mathbf{f}|\mathbf{X}) \rightarrow p(\mathbf{f}|\mathbf{u}, \mathbf{X}, \mathbf{Z})p(\mathbf{u}|\mathbf{Z})$$

We then appeal to Jensen’s inequality and lower bound the log joint density over the targets and inducing variables:

$$\begin{aligned} \log p(\mathbf{y}, \mathbf{u}|\mathbf{X}, \mathbf{Z}) &= \log \int d\mathbf{f} p(\mathbf{y}|\mathbf{f})p(\mathbf{f}|\mathbf{u})p(\mathbf{u}) \\ &\geq \mathbb{E}_{p(\mathbf{f}|\mathbf{u})} [\log p(\mathbf{y}|\mathbf{f}) + \log p(\mathbf{u})] \\ &= \sum_{i=1}^N \log \mathcal{N}(y_i | \mathbf{k}_i^T \mathbf{K}_{MM}^{-1} \mathbf{u}, \sigma_{\text{obs}}^2) \\ &\quad - \frac{1}{2\sigma_{\text{obs}}^2} \text{Tr} \tilde{\mathbf{K}}_{NN} + \log p(\mathbf{u}) \end{aligned} \quad (3)$$

where $\mathbf{k}_i = k(\mathbf{x}_i, \mathbf{Z})$, $\mathbf{K}_{MM} = k(\mathbf{Z}, \mathbf{Z})$ and $\tilde{\mathbf{K}}_{NN}$ is given by

$$\tilde{\mathbf{K}}_{NN} = \mathbf{K}_{NN} - \mathbf{K}_{NM} \mathbf{K}_{MM}^{-1} \mathbf{K}_{MN} \quad (4)$$

with $\mathbf{K}_{NM} = \mathbf{K}_{MN}^T = k(\mathbf{X}, \mathbf{Z})$. The essential characteristics of Eqn. 3 are that: i) it replaces expensive computations involving \mathbf{K}_{NN} with cheaper computations like \mathbf{K}_{MM}^{-1} that scale as $\mathcal{O}(M^3)$; and ii) it is amenable to data subsampling, since the log likelihood and trace terms factorize as sums over datapoints (y_i, \mathbf{x}_i) .

2.3. Approximate GP Inference

We now describe how Eqn. 3 can be used to construct a variety of algorithms for scalable GP inference. We limit ourselves to algorithms that satisfy two desiderata: i) support for data subsampling,³ and ii) the result of inference

³For this reason we do not consider the method in (Titsias, 2009). Note that all the inference algorithms that we describe

is a compact artifact that enables fast predictions at test time.⁴ The rest of this section is organized as follows. In Sec. 2.3.1 we describe SVGP (Hensman et al., 2013)—currently the most popular method for scalable GP inference. In Sec. 2.3.2 we describe how Eqn. 3 can be leveraged in the context of MAP (maximum a posteriori) inference. Then in Sec. 2.3.3 we briefly review how ideas from robust variational inference can be applied to the GP setting.

2.3.1. SVGP

SVGP proceeds by introducing a multivariate Normal variational distribution $q(\mathbf{u}) = \mathcal{N}(\mathbf{m}, \mathbf{S})$. The parameters \mathbf{m} and \mathbf{S} are optimized using the ELBO (evidence lower bound), which is the expectation of Eqn. 3 w.r.t. $q(\mathbf{u})$ plus an entropy term $H[q(\mathbf{u})]$:

$$\begin{aligned} \mathcal{L}_{\text{svgp}} &= \mathbb{E}_{q(\mathbf{u})} [\log p(\mathbf{y}, \mathbf{u} | \mathbf{X}, \mathbf{Z})] + H[q(\mathbf{u})] \quad (5) \\ &= \sum_{i=1}^N \log \mathcal{N}(y_i | \mathbf{k}_i^T \mathbf{K}_{MM}^{-1} \mathbf{m}, \sigma_{\text{obs}}^2) - \frac{1}{2\sigma_{\text{obs}}^2} \text{Tr} \tilde{\mathbf{K}}_{NN} \\ &\quad - \frac{1}{2\sigma_{\text{obs}}^2} \sum_{i=1}^N \mathbf{k}_i^T \mathbf{K}_{MM}^{-1} \mathbf{S} \mathbf{K}_{MM}^{-1} \mathbf{k}_i - \text{KL}(q(\mathbf{u}) | p(\mathbf{u})) \end{aligned}$$

where KL denotes the Kullback-Leibler divergence. Eqn. 5 can be rewritten more compactly as

$$\mathcal{L}_{\text{svgp}} = \sum_{i=1}^N \left\{ \log \mathcal{N}(y_i | \mu_{\mathbf{f}}(\mathbf{x}_i), \sigma_{\text{obs}}^2) - \frac{\sigma_{\mathbf{f}}(\mathbf{x}_i)^2}{2\sigma_{\text{obs}}^2} \right\} - \text{KL}(q(\mathbf{u}) | p(\mathbf{u})) \quad (6)$$

where $\mu_{\mathbf{f}}(\mathbf{x}_i)$ is the predictive mean function given by

$$\mu_{\mathbf{f}}(\mathbf{x}_i) = \mathbf{k}_i^T \mathbf{K}_{MM}^{-1} \mathbf{m} \quad (7)$$

and where $\sigma_{\mathbf{f}}(\mathbf{x}_i)^2 \equiv \text{Var}[f_i | \mathbf{x}_i]$ denotes the latent function variance

$$\sigma_{\mathbf{f}}(\mathbf{x}_i)^2 = \tilde{\mathbf{K}}_{ii} + \mathbf{k}_i^T \mathbf{K}_{MM}^{-1} \mathbf{S} \mathbf{K}_{MM}^{-1} \mathbf{k}_i \quad (8)$$

$\mathcal{L}_{\text{svgp}}$, which depends on \mathbf{m} , \mathbf{S} , \mathbf{Z} , σ_{obs} and the various kernel hyperparameters θ_{ker} , can then be maximized with gradient methods. Below we refer to the resulting GP regression method as **SVGP**. Note that throughout this work we consider a variant of SVGP in which the KL divergence term in Eqn. 6 is scaled by a positive constant $\beta_{\text{reg}} > 0$.

2.3.2. MAP

In contrast to SVGP, which maintains a distribution over the inducing variables \mathbf{u} , MAP is a particle method in which that make use of Eqn. 3 automatically inherit its support for data subsampling.

⁴Consequently we do not consider MCMC algorithms like the Stochastic gradient HMC algorithm explored in the context of deep gaussian processes in (Havasi et al., 2018), and which also utilizes Eqn. 3.

we directly optimize a single point $\mathbf{u} \in \mathbb{R}^M$ rather than a distribution over \mathbf{u} . In particular we simply maximize Eqn. 3 evaluated at \mathbf{u} . Note that the term $\log p(\mathbf{u})$ serves as a regularizer. In the following we refer to this inference procedure as **MAP**. To the best of our knowledge, it has not been considered before in the sparse GP literature.

2.3.3. ROBUST GAUSSIAN PROCESSES

Knoblauch et al. (2019); Knoblauch (2019) consider modifications to the typical variational objective (e.g. Eqn. 5), which consists of an expected log likelihood and a KL divergence term. In particular, they replace the expected log likelihood loss with an alternative divergence like the gamma divergence. This divergence raises the likelihood to a power

$$\log p(\mathbf{y} | \mathbf{f}) \rightarrow p(\mathbf{y} | \mathbf{f})^{\gamma-1} \quad (9)$$

where typically $\gamma \in (1.0, 1.1)$.⁵ Empirically Knoblauch (2019) demonstrates that this modification can yield better performance than SVGP on regression tasks. We refer to this inference procedure as **γ -Robust**.

2.4. Predictive Distributions

All of the methods in Sec. 2.3 yield predictive distributions of the same form. In particular, conditioned on the inducing variable \mathbf{u} the predictive distribution at input \mathbf{x}^* is given by

$$p(\mathbf{y}^* | \mathbf{x}^*, \mathbf{u}) = \mathcal{N}(\mathbf{y}^* | \mathbf{k}_*^T \mathbf{K}_{MM}^{-1} \mathbf{u}, \tilde{\mathbf{K}}_{**} + \sigma_{\text{obs}}^2) \quad (10)$$

Integrating out $\mathbf{u} \sim \mathcal{N}(\mathbf{m}, \mathbf{S})$ then yields

$$p(\mathbf{y}^* | \mathbf{x}^*) = \mathcal{N}(\mathbf{y}^* | \mu_{\mathbf{f}}(\mathbf{x}^*), \sigma_{\mathbf{f}}(\mathbf{x}^*)^2 + \sigma_{\text{obs}}^2) \quad (11)$$

where $\mu_{\mathbf{f}}(\mathbf{x}^*)$ is the predictive mean function in Eqn. 7 and $\sigma_{\mathbf{f}}(\mathbf{x}^*)^2$ is the latent function variance in Eqn. 8. Note that since MAP can be viewed as a degenerate limit of SVGP, the predictive distribution for MAP can be obtained by taking the limit $\mathbf{S} \rightarrow 0$ in $\sigma_{\mathbf{f}}(\mathbf{x}^*)^2$.

2.5. FITC

FITC (Fully Independent Training Conditional) (Snelson and Ghahramani, 2006) is a method for sparse GP regression that is formulated using a joint probability

$$p(\mathbf{y}, \mathbf{u}) = \prod_{i=1}^N p(y_i | \mathbf{k}_i^T \mathbf{K}_{MM}^{-1} \mathbf{u}, \tilde{\mathbf{K}}_{ii} + \sigma_{\text{obs}}^2) p(\mathbf{u}) \quad (12)$$

with a modified likelihood corresponding to the conditional predictive distribution in Eqn. 10 (for additional interpretations see Bauer et al. (2016)). As noted by Snelson and Ghahramani (2006) this can be viewed as a standard regression model with a particular form of parametric mean

⁵See (Cichocki and Amari, 2010) for a detailed discussion of this and other divergences.

Algorithm 1: Scalable GP Regression. All of the inference algorithms we consider follow the same basic pattern and only differ in the form of the objective function, e.g. $\mathcal{L}_{\text{svgp}}$ (Eqn. 6), $\mathcal{L}_{\text{vfitc}}$ (Eqn. 15) or $\mathcal{L}_{\text{ppgpr}}$ (Eqn. 18). Similarly for all methods the predictive distribution is given by Eqn. 11. See Sec. B in the supplementary materials for a discussion of the time and space complexity of each method.

Input: \mathcal{L} – objective function
 $\text{optim}()$ – gradient-based optimizer
 $\mathcal{D} = \{\mathbf{x}_i, y_i\}_{i=1}^N$ – training data
 $\mathbf{m}, \mathbf{S}, \mathbf{Z}, \theta_{\text{ker}}$ – initial parameters

Output: $\mathbf{m}, \mathbf{S}, \mathbf{Z}, \theta_{\text{ker}}$.

while not converged do
 Choose a random mini-batch $\mathcal{D}_{\text{mb}} \subset \mathcal{D}$
 Form an unbiased estimate $\hat{\mathcal{L}}(\mathcal{D}_{\text{mb}})$
 Gradient step: $\mathbf{m}, \mathbf{S}, \mathbf{Z}, \theta_{\text{ker}} \leftarrow \text{optim}(\hat{\mathcal{L}}(\mathcal{D}_{\text{mb}}))$

end

function and input-dependent observation noise. Integrating out the latent function values \mathbf{u} results in a marginal likelihood

$$p(\mathbf{y}) = \mathcal{N}(\mathbf{y} | \mathbf{0}, \mathbf{K}_{NM} \mathbf{K}_{MM}^{-1} \mathbf{K}_{MN} + \text{diag}(\tilde{\mathbf{K}}_{NN}) + \sigma_{\text{obs}}^2 \mathbf{1}) \quad (13)$$

that can be used for training with $\mathcal{O}(NM^2 + M^3)$ computational complexity per gradient step. Note that the structure of the covariance matrix in Eqn. 13 prevents mini-batch training, limiting FITC to moderately large datasets.

3. Parametric Gaussian Process Regressors

Before introducing the two scalable methods for GP regression we consider in this work, we examine one of the salient characteristics of the SVGP objective Eqn. 6 referred to in the introduction. As discussed in Sec. 2.4, in SVGP the predictive variance $\text{Var}[y^* | \mathbf{x}^*]$ at an input \mathbf{x}^* has two contributions, one that is input-dependent, i.e. $\sigma_{\mathbf{f}}(\mathbf{x}^*)^2$, and one that is input-independent, i.e. σ_{obs}^2 :

$$\begin{aligned} \text{Var}[y^* | \mathbf{x}^*] &= \sigma_{\text{obs}}^2 + \sigma_{\mathbf{f}}(\mathbf{x}^*)^2 \\ &= \sigma_{\text{obs}}^2 + \tilde{\mathbf{K}}_{**} + \mathbf{k}_*^T \mathbf{K}_{MM}^{-1} \mathbf{S} \mathbf{K}_{MM}^{-1} \mathbf{k}_* \end{aligned} \quad (14)$$

These two contributions appear *asymmetrically* in Eqn. 6; in particular $\sigma_{\mathbf{f}}(\mathbf{x}_i)^2$ does not appear in the data fit term that results from the expected log likelihood $\mathbb{E}_{q(\mathbf{u})} [\log p(\mathbf{y}_i | \mathbf{x}_i, \mathbf{u})]$. We expect this behavior to be undesirable, since it leads to a mismatch between the training objective and the predictive distribution used at test time.

In the following we introduce two approaches that address this asymmetry. Crucially, unlike the inference strategies outlined in Sec. 2.3, neither approach makes use of the lower-bound energy surface in Eqn. 3. As we will see, the

first approach, Variational FITC, introduced in Sec. 3.1, only partially removes the asymmetry, while the second Parametric Predictive GP approach, introduced in Sec. 3.2, restores full symmetry between the training objective and the test time predictive distribution. For more discussion of this point see Sec. 3.3.

3.1. Variational FITC

Variational FITC proceeds by applying variational inference to the FITC model defined in Eqn. 12. That is, we introduce a multivariate Normal variational distribution $q(\mathbf{u}) = \mathcal{N}(\mathbf{m}, \mathbf{S})$ and compute the ELBO:

$$\begin{aligned} \mathcal{L}_{\text{vfitc}} &= \mathbb{E}_{q(\mathbf{u})} [\log p(\mathbf{y}, \mathbf{u} | \mathbf{X}, \mathbf{Z})] + H[q(\mathbf{u})] \\ &= \sum_{i=1}^N \log \mathcal{N}(y_i | \mu_{\mathbf{f}}(\mathbf{x}_i), \tilde{\mathbf{K}}_{ii} + \sigma_{\text{obs}}^2) \\ &\quad - \frac{1}{2} \sum_{i=1}^N \frac{\mathbf{k}_i^T \mathbf{K}_{MM}^{-1} \mathbf{S} \mathbf{K}_{MM}^{-1} \mathbf{k}_i}{\tilde{\mathbf{K}}_{ii} + \sigma_{\text{obs}}^2} - \text{KL}(q(\mathbf{u}) | p(\mathbf{u})) \end{aligned} \quad (15)$$

The parameters $\mathbf{m}, \mathbf{S}, \mathbf{Z}$, as well as the observation noise σ_{obs} and kernel hyperparameters θ_{ker} can then be optimized by maximizing Eqn. 15 using gradient methods (see Algorithm 1). Note that since the inducing point locations \mathbf{Z} appear in the model, this is properly understood as a parametric model.

We note that, unlike FITC, the objective in Eqn. 15 readily supports mini-batch training, and is thus suitable for very large datasets. Below we refer to this method as **VFITC**.

3.2. Parametric Predictive GP Regression

As we discuss further in the next section, using the VFITC objective only partially addresses the asymmetric treatment of σ_{obs}^2 and $\sigma_{\mathbf{f}}(\mathbf{x}^*)^2$ in Eqn. 6. We now describe our second approach to scalable GP regression, which is motivated by the goal of restoring full symmetry between the training objective and the test time predictive distribution. To accomplish this, we introduce a parametric GP regression model formulated directly in terms of the family of predictive distributions given in Eqn. 11. We then introduce an objective function based on the KL divergence between $p(y | \mathbf{x})$ and $p_{\text{data}}(y | \mathbf{x})$

$$\begin{aligned} \mathcal{L}'_{\text{ppgpr}} &= -\mathbb{E}_{p_{\text{data}}(\mathbf{x})} \text{KL}(p_{\text{data}}(y | \mathbf{x}) || p(y | \mathbf{x})) \\ &\quad \rightarrow \mathbb{E}_{p_{\text{data}}(y, \mathbf{x})} [\log p(y | \mathbf{x})] \end{aligned} \quad (16)$$

where $p_{\text{data}}(y, \mathbf{x})$ is the empirical distribution over training data. In the second line we have dropped the entropy term $-\mathbb{E}_{p_{\text{data}}(y | \mathbf{x})} [\log p_{\text{data}}(y | \mathbf{x})]$, since it is a constant w.r.t. to the maximization problem. We obtain our final objective function by adding an optional regularization term modu-

lated by a positive constant $\beta_{\text{reg}} > 0$:

$$\mathcal{L}_{\text{ppgpr}} = \mathbb{E}_{p_{\text{data}}(y, \mathbf{x})} [\log p(y|\mathbf{x})] - \beta_{\text{reg}} \text{KL}(q(\mathbf{u})||p(\mathbf{u})) \quad (17)$$

Note that apart from the regularization term, maximizing this objective function corresponds to doing maximum likelihood estimation of a parametric model defined by $p(y|\mathbf{x})$. The objective in Eqn. 17 can be expanded as

$$\mathcal{L}_{\text{ppgpr}} = \sum_{i=1}^N \log \mathcal{N}(y_i | \mu_{\mathbf{f}}(\mathbf{x}_i), \sigma_{\text{obs}}^2 + \sigma_{\mathbf{f}}(\mathbf{x}_i)^2) - \beta_{\text{reg}} \text{KL}(q(\mathbf{u})||p(\mathbf{u})) \quad (18)$$

where $\sigma_{\mathbf{f}}(\mathbf{x}_i)^2 = \text{Var}[f_i|\mathbf{x}_i]$ is the function variance defined in Eqn. 8. The parameters \mathbf{m} , \mathbf{S} , \mathbf{Z} , as well as the observation noise σ_{obs} and kernel hyperparameters θ_{ker} can then be optimized by maximizing Eqn. 17 using gradient methods (see Algorithm 1). We refer to this model class as **PPGPR**.

When $\beta_{\text{reg}} = 1$ the form of the objective in Eqn. 17 can be motivated by a connection to Expectation Propagation (Minka, 2004); see Sec. G in the supplementary materials and (Li and Gal, 2017) for further discussion.⁶

3.3. Discussion

What are the implications of employing the scalable GP regressors described in Sec. 3.1 and Sec. 3.2?⁷ It is helpful to compare the objective functions in Eqn. 15 and Eqn. 17 to the SVGP objective in Eqn. 6. In Eqn. 15 we obtain a data fit term

$$\mathcal{L}_{\text{vfitc}} \supset -\frac{1}{2} \frac{1}{\sigma_{\text{obs}}^2 + \tilde{\mathbf{K}}_{ii}} |y_i - \mu_{\mathbf{f}}(\mathbf{x}_i)|^2 \quad (19)$$

while in Eqn. 18 we obtain a data fit term

$$\mathcal{L}_{\text{ppgpr}} \supset -\frac{1}{2} \frac{1}{\sigma_{\text{obs}}^2 + \sigma_{\mathbf{f}}(\mathbf{x}_i)^2} |y_i - \mu_{\mathbf{f}}(\mathbf{x}_i)|^2 \quad (20)$$

In contrast in SVGP the corresponding data fit term is

$$\mathcal{L}_{\text{svgp}} \supset -\frac{1}{2} \frac{1}{\sigma_{\text{obs}}^2} |y_i - \mu_{\mathbf{f}}(\mathbf{x}_i)|^2 \quad (21)$$

with σ_{obs}^2 in the denominator. We reiterate that in all three approaches the predictive variance $\text{Var}[y^*|\mathbf{x}^*]$ at an input \mathbf{x}^* is given by the formula

$$\text{Var}[y^*|\mathbf{x}^*] = \sigma_{\text{obs}}^2 + \sigma_{\mathbf{f}}(\mathbf{x}^*)^2 \quad (22)$$

where $\sigma_{\mathbf{f}}(\mathbf{x}^*)^2$ is the latent function variance at \mathbf{x}^* (see Eqn. 8). Thus SVGP—and, more generally, variational inference for any regression model with a Normal likelihood—

makes an arbitrary⁸ distinction between the observation variance σ_{obs}^2 and the latent function variance $\sigma_{\mathbf{f}}(\mathbf{x}^*)^2$, even though both terms contribute symmetrically to the total predictive variance in Eqn. 22. Moreover, this asymmetry will be inherited by any method that makes use of Eqn. 3, e.g. the MAP procedure described in Sec. 2.3.2.

When the priority is predictive performance—the typical case in machine learning—this asymmetric treatment of the two contributions to the predictive variance is troubling. As we will see in experiments (Sec. 5), the difference between Eqn. 20 and Eqn. 21 has dramatic consequences. In particular the data fit term in SVGP does nothing to encourage function variance $\sigma_{\mathbf{f}}(\mathbf{x}^*)$. As a consequence for many datasets $\sigma_{\mathbf{f}}(\mathbf{x}^*) \ll \sigma_{\text{obs}}$; i.e. *most of the predictive variance is explained by the input-independent observation noise*.

In contrast, the data fit term in Eqn. 20 directly encourages large $\sigma_{\mathbf{f}}(\mathbf{x}^*)$, typically resulting in behavior opposite to that of SVGP, i.e. $\sigma_{\mathbf{f}}(\mathbf{x}^*) \gg \sigma_{\text{obs}}$. This is gratifying because—after having gone to the effort to introduce an input-dependent kernel and learn an appropriate geometry on the input space—we end up with predictive variances that make substantial use of the input-dependent kernel.

The case of Variational FITC (see Eqn. 19) is intermediate in that the objective function directly incentivizes non-negligible latent function variance through the term $\tilde{\mathbf{K}}_{ii}(\mathbf{x}_i)$, while the \mathbf{S} -dependent contribution to $\sigma_{\mathbf{f}}(\mathbf{x}_i)$ is still treated asymmetrically (since it does not appear in the data fit term).

As argued by (Bauer et al., 2016), methods based on FITC—and by extension our Parametric Predictive GP approach—tend to *underestimate* the observation noise σ_{obs}^2 , while variational methods like SVGP or the closely related method introduced by (Titsias, 2009) tend to *overestimate* σ_{obs}^2 . While the possibility of underestimating σ_{obs}^2 is indeed a concern for our approach, our empirical results in Sec. 5 suggest that, this tendency notwithstanding, our methods yield excellent predictive performance.

3.4. Additional Variants

A number of variants to the Parametric Predictive GP approach outlined in Sec. 3.2 immediately suggest themselves. One possibility is to take the formal limit $\mathbf{S} \rightarrow 0$ in $q(\mathbf{u})$. In this limit $q(\mathbf{u})$ is a Dirac delta distribution, the function variance $\sigma_{\mathbf{f}}(\mathbf{x}_i)^2 \rightarrow \tilde{\mathbf{K}}_{ii}$, and the number of parameters is now linear in M instead of quadratic.⁹ Below we refer to this variant as **PPGPR- δ** . Another possibility is to restrict

⁸Arbitrary from the point of view of output space. Depending on the particular application and structure of the model, distinctions between different contributions to the predictive variance may be of interest.

⁹Additionally in the regularizer we make the replacement $-\beta_{\text{reg}} \text{KL}(q(\mathbf{u})||p(\mathbf{u})) \rightarrow \beta_{\text{reg}} \log p(\mathbf{u})$.

⁶We thank Thang Bui for pointing out this connection.

⁷See Sec D in the supplementary materials for a comparison to exact GPs.

Parametric Gaussian Process Regressors




Figure 2. We depict test negative log likelihoods (NLL) for twelve univariate regression datasets (lower is better). Results are averaged over ten random train/test/validation splits. See Sec. 5.1 for details. Here and elsewhere horizontal error bars depict standard errors.




Figure 3. We depict test root mean squared errors (RMSE) for twelve univariate regression datasets (lower is better). Results are averaged over ten random train/test/validation splits. See Sec. 5.1 for details.

the covariance matrix \mathbf{S} in $q(\mathbf{u})$ to be diagonal; we refer to this ‘mean field’ variant as **PPGPR-MF**. Yet another possibility is to decouple the parametric mean function $\mu_f(\mathbf{x})$ and variance function $\sigma_f(\mathbf{x})^2$ used to define $p(y|\mathbf{x})$ in Eqn. 17, i.e. use separate inducing point locations \mathbf{Z}_μ and \mathbf{Z}_σ for additional flexibility. This is conceptually similar to the decoupled approach introduced in Cheng and Boots (2017), with the difference that our parametric modeling approach is not constrained by the need to construct a well-defined variational inference problem in an infinite-dimensional RKHS. Below we refer to this decoupled approach together with a diagonal covariance \mathbf{S} as **PPGPR-MFD**. We refer to the variant of PPGPR that is closest to SVGP (because it utilizes a full-rank covariance matrix $\mathbf{S} = \mathbf{L}\mathbf{L}^T$) as **PPGPR-Chol**.

A number of other variants are also possible. For example we might replace the KL regularizer in Eqn. 17 with another divergence, for example a Rényi divergence (Knoblauch et al., 2019). Alternatively we could use another divergence measure in Eqn. 16—e.g. the gamma divergence used in Sec. 2.3.3—to control the qualitative features of $p_{\text{data}}(y|\mathbf{x})$ that we would like to capture in $p(y|\mathbf{x})$. We leave the exploration of these and other variants to future work.

4. Related Work

The use of pseudo-inputs and inducing point methods to scale-up Gaussian Process inference has spawned a large literature, especially in the context of variational inference (Csató and Opper, 2002; Seeger et al., 2003; Quiñero-

Candela and Rasmussen, 2005; Snelson and Ghahramani, 2006; Titsias, 2009; Hensman et al., 2013; 2015a; Cheng and Boots, 2017). While variational inference remains the most popular inference algorithm in the scalable GP setting, researchers have also explored different variants of Expectation Propagation (Hernández-Lobato and Hernández-Lobato, 2016; Bui et al., 2017) as well as Stochastic gradient Hamiltonian Monte Carlo (Havasi et al., 2018) and other MCMC algorithms (Hensman et al., 2015b). For a recent review of scalable methods for GP inference we refer the reader to (Liu et al., 2018).

Our approach bears some resemblance to that described in (Raissi et al., 2019) in that, like them, we consider GP regression models that are parametric in nature. There are important differences, however. First because Raissi et al. (2019) consider a two-step training procedure that does not benefit from a single coherent objective, they are unable to learn inducing point locations \mathbf{Z} . Second, inconsistent treatment of latent function uncertainty between the two training steps degrades the quality of the predictive uncertainty. For these reasons we find that the approach in (Raissi et al., 2019) significantly underperforms¹⁰ all the other methods we consider and hence we do not consider it further.

Several of our objective functions can also be seen as instances of Direct Loss Minimization, which emerges from a

¹⁰Both in terms of RMSE and log likelihood (especially the latter, with the predictive uncertainty drastically underestimated in some cases). See Sec. D for details.

view of approximate inference as regularized loss minimization (Sheth and Khardon, 2016; 2017; 2019).¹¹ We refer the reader to Sec. F in the supplementary materials for further discussion of this important connection.

Our focus on the predictive distribution also recalls (Snelson and Ghahramani, 2005), in which the authors construct parsimonious approximations to Bayesian predictive distributions. Their approach differs from the approach adopted here, since the posterior distribution is still computed (or approximated) as an intermediate step, whereas in PPGPR we completely bypass the posterior. Similarly PPGPR recalls (Gordon et al., 2018), where the authors consider training objectives that explicitly target posterior predictive distributions in the context of meta-learning.

5. Experiments

In this section we compare the empirical performance of the approaches to scalable GP regression introduced in Sec. 3 to the baseline inference strategies described in Sec. 2.3. All our models use a prior mean of zero and a Matérn kernel with independent length scales for each input dimension.¹²

5.1. Univariate regression

We consider a mix of univariate regression datasets from the UCI repository (Dua and Graff, 2017), with the number of datapoints ranging from $N \sim 10^4$ to $N \sim 10^6$ and the number of input dimensions in the range $\dim(\mathbf{x}) \in [3, 380]$. Among the methods introduced in Sec. 3, we focus on Variational FITC (Sec. 3.1) and PPGPR-MFD (Sec. 3.4). In addition to the baselines reviewed in Sec. 2.3, we also compare to the orthogonal parameterization of the basis decoupling method described in Cheng and Boots (2017) and Salimbeni et al. (2018), which we refer to as **OD-SVGP**. For all but the two largest datasets we also compare to **Exact** GP inference, leveraging the conjugate gradient approach described in (Wang et al., 2019).

We summarize the results in Fig. 2-4 and Table 1 (see the supplementary materials for additional results). Both our approaches yield consistently lower negative log likelihoods (NLLs) than the baseline approaches, with PPGPR performing particularly well. Averaged across all twelve datasets, PPGPR outperforms the strongest baseline, OD-SVGP, by ~ 0.35 nats. Interestingly on most datasets PPGPR outperforms Exact GP inference. We hypothesize that this is at least partially due to the ability of our parametric models to encode heteroscedasticity, a “bonus” feature that was al-

¹¹Note, however, that this interpretation is not applicable to our best performing model class, PPGPR-MFD, which benefits from mean and variance functions that are decoupled.

¹²For an implementation of our method in GPyTorch see <https://git.io/JJy9b>.

Table 1. Average ranking of methods (lower is better). CRPS is the Continuous Ranking Probability Score, a popular calibration metric for regression (Gneiting and Raftery, 2007). Rankings are averages across datasets and splits. See Sec. 5.1 for details.

	MAP	γ -Robust	SVGP	OD-SVGP	VFITC	PPGPR
NLL	5.47	4.76	4.33	3.08	2.20	1.17
RMSE	4.53	4.59	2.92	2.49	4.35	2.12
CRPS	5.55	4.12	4.21	3.20	2.90	1.02

ready noted by Snelson and Ghahramani (2006). Perhaps surprisingly, we note that on most datasets MAP yields comparable NLLs to SVGP. Indeed averaged across all twelve datasets, SVGP only outperforms MAP by ~ 0.05 nats.

The results for root mean squared errors (RMSE) exhibit somewhat less variability (see Fig. 3), with PPGPR performing the best and OD-SVGP performing second best among the scalable methods. In particular, in aggregate PPGPR attains the lowest RMSE (see Table 1), though it is outperformed by other methods on 2/12 datasets. We hypothesize that the RMSE performance of VFITC could be substantially improved if the variational distribution took the structured form that is implicitly used in OD-SVGP. Not surprisingly Exact GP inference yields the lowest RMSE for most datasets.

Strikingly, both VFITC and PPGPR yield predictive variances that are in a qualitatively different regime than those resulting from the scalable inference baselines. Fig. 4 depicts the fraction of the total predictive variance $\text{Var}[y^* | \mathbf{x}]$ due to the observation noise σ_{obs}^2 . The predictive variances from the baseline methods make relatively little use of input-dependent function uncertainty, instead relying primarily on the observation noise. By contrast the variances of our parametric GP regressors are dominated by function uncertainty. This substantiates the discussion in Sec. 3.3. Additionally, this observation explains the similar log likelihoods exhibited by SVGP and MAP: since neither method makes much use of function uncertainty, the uncertainty encoded in $q(\mathbf{u})$ is of secondary importance.¹³

Finally we note that for most datasets PPGPR prefers small values of β_{reg} . This suggests that choosing $M \ll N$ is sufficient for ensuring well-regularized models and that overfitting is not much of a concern in practice.

5.2. PPGPR Ablation Study

Encouraged by the predictive performance of PPGPR-MFD in Sec. 5.1, we perform a detailed comparison of the different PPGPR variants introduced in Sec. 3.4. Our results are summarized in Table 2 (see supplementary materials for additional figures). We find that as we increase the capacity of

¹³Note that MAP can be viewed as a degenerate limit of SVGP in which $q(\mathbf{u})$ is a Dirac delta function.

Parametric Gaussian Process Regressors




Figure 4. We depict the mean fraction of the predictive variance that is due to the observation noise (as measured on the test set). Results are averaged over ten random train/test/validation splits. See Sec. 5.1 for details.




Figure 5. We visualize the calibration of three probabilistic deep learning models fit to trip data in three cities. For each method we depict the empirical CDF of the z-scores $z = (y^* - \mu_f(\mathbf{x}^*)) / \sigma(\mathbf{x}^*)$ computed on the test set $\{(\mathbf{x}_k^*, y_k^*)\}$. The “Ideal CDF” is the Normal CDF, which corresponds to the best possible calibration for a model with a Normal predictive distribution. See Sec. 5.3 for details.

the variance function $\sigma_f(\mathbf{x})^2$ (i.e. PPGPR- $\delta \Rightarrow$ PPGPR-MF \Rightarrow PPGPR-Chol) the test NLL tends to decrease, while the test RMSE tends to increase. This makes sense since PPGPR prefers large predictive uncertainty in regions of input space where good data fit is hard to achieve. Consequently, there is less incentive to move the predictive mean function away from the prior in those regions, which can then result in higher RMSEs; this effect becomes more pronounced as the variance function becomes more flexible. In contrast, since the mean and variance functions are entirely decoupled in the case of PPGPR-MFD (apart from shared kernel hyperparameters), this model class obtains low NLLs without sacrificing performance on RMSE.

Note, finally, that PPGPR- δ and PPGPR-Chol utilize the same family of predictive distributions as MAP and SVGP, respectively, and only differ in the objective function used during training. If we compare these pairs of methods in Table 2, we find that PPGPR- δ and PPGPR-Chol yield the best log likelihoods, with PPGPR-Chol exhibiting degraded RMSE performance for the reason described in the previous paragraph. Indeed these observations motivated the introduction of PPGPR-MFD.

Table 2. Average rank of PPGPR variants (lower is better).

	MAP	SVGP	PPGPR δ	PPGPR MF	PPGPR Chol	PPGPR MFD
NLL	5.90	5.08	3.96	2.66	1.59	1.81
RMSE	3.64	2.36	3.65	4.16	5.38	1.82
CRPS	5.72	4.78	3.87	2.77	2.47	1.40

5.3. Calibration in DKL Regression

In this section we demonstrate that PPGPR (introduced in Sec. 3.2) offers an effective mechanism for calibrating deep neural networks for regression. To evaluate the potential of this approach, we utilize a real-world dataset of vehicular trip durations in three large cities. In this setting, uncertainty estimation is critical for managing risk when estimating transportation costs.

We compare three methods: i) deep kernel learning (Wilson et al., 2016) using SVGP (SVGP+DKL); ii) deep kernel learning using PPGPR-MFD (PPGPR+DKL); and iii) MCDropout (Gal and Ghahramani, 2016), a popular method for calibrating neural networks that does not rely on Gaussian

Processes.

In Fig. 5 we visualize how well each of the three methods is calibrated as compared to the best possible calibration for a model with Normal predictive distributions. Overall PPGPR+DKL performs the best, with MCDropout outperforming or matching SVGP+DKL. Using PPGPR has a number of additional advantages over MCDropout. In particular, because the predictive variances can be computed analytically for Gaussian Process models, the PPGPR+DKL model is significantly faster at test time than MCDropout, which requires forwarding data points through many sampled models (here 50). This is impractically slow for many applied settings, especially for large neural networks.

6. Conclusions

Gaussian Process regression with a Normal likelihood represents a peculiar case in that: i) we can give an analytic formula for the exact posterior predictive distribution; but ii) it is impractical to compute for large datasets. In this work we have argued that if our goal is high quality predictive distributions, it is sensible to bypass posterior approximations and directly target the quantity of interest. While this may be a bad strategy for an arbitrary probabilistic model, in the case of GP regression inducing point methods provide a natural family of parametric predictive distributions whose capacity can be controlled to prevent overfitting. As we have shown empirically, the resulting predictive distributions exhibit significantly better calibrated uncertainties and higher log likelihoods than those obtained with other methods, which tend to yield overconfident uncertainty estimates that make little use of the kernel.

More broadly, our empirical results suggest that the good predictive performance of approximate GP methods like SVGP may have more to do with the power of kernel methods than adherence to Bayes. The approach we have adopted here can be viewed as maximum likelihood estimation for a model with a high-powered likelihood. This likelihood leverages flexible mean and variance functions whose parametric form is motivated by inducing point methods. We suspect that a good ansatz for the variance function is of particular importance for good predictive uncertainty. We explore one possible alternative—and much more flexible—parameterization in Jankowiak et al. (2020). Here we suggest that one simple and natural variant of PPGPR would replace the mean function with a flexible neural network so that inducing points are only used to define the variance function.

ACKNOWLEDGEMENTS

MJ would like to thank Jeremias Knoblauch, Jack Jewson, and Theodoros Damouls for engaging conversations about

robust variational inference. MJ would also like to thank Felipe Petroski Such for help with Uber compute infrastructure. This work was completed while MJ and JG were at Uber AI.

References

- Matthias Bauer, Mark van der Wilk, and Carl Edward Rasmussen. Understanding probabilistic sparse gaussian process approximations. In *Advances in neural information processing systems*, pages 1533–1541, 2016.
- Eli Bingham, Jonathan P Chen, Martin Jankowiak, Fritz Obermeyer, Neeraj Pradhan, Theofanis Karaletsos, Rohit Singh, Paul Szerlip, Paul Horsfall, and Noah D Goodman. Pyro: Deep universal probabilistic programming. *The Journal of Machine Learning Research*, 20(1):973–978, 2019.
- Thang D Bui, Josiah Yan, and Richard E Turner. A unifying framework for gaussian process pseudo-point approximations using power expectation propagation. *The Journal of Machine Learning Research*, 18(1):3649–3720, 2017.
- Ching-An Cheng and Byron Boots. Variational inference for gaussian process models with linear complexity. In *Advances in Neural Information Processing Systems*, pages 5184–5194, 2017.
- Andrzej Cichocki and Shun-ichi Amari. Families of alpha-beta-and gamma-divergences: Flexible and robust measures of similarities. *Entropy*, 12(6):1532–1568, 2010.
- Lehel Csató and Manfred Opper. Sparse on-line gaussian processes. *Neural computation*, 14(3):641–668, 2002.
- Dheeru Dua and Casey Graff. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059, 2016.
- Jacob Gardner, Geoff Pleiss, Kilian Q Weinberger, David Bindel, and Andrew G Wilson. Gpytorch: Blackbox matrix-matrix gaussian process inference with gpu acceleration. In *Advances in Neural Information Processing Systems*, pages 7576–7586, 2018.
- Tilmann Gneiting and Adrian E Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.
- Jonathan Gordon, John Bronskill, Matthias Bauer, Sebastian Nowozin, and Richard E Turner. Meta-learning probabilistic inference for prediction. *arXiv preprint arXiv:1805.09921*, 2018.

- Marton Havasi, José Miguel Hernández-Lobato, and Juan José Murillo-Fuentes. Inference in deep gaussian processes using stochastic gradient hamiltonian monte carlo. In *Advances in Neural Information Processing Systems*, pages 7506–7516, 2018.
- James Hensman, Nicolo Fusi, and Neil D Lawrence. Gaussian processes for big data. *arXiv preprint arXiv:1309.6835*, 2013.
- James Hensman, Alexander Matthews, and Zoubin Ghahramani. Scalable variational gaussian process classification. 2015a.
- James Hensman, Alexander G Matthews, Maurizio Filippone, and Zoubin Ghahramani. Mcmc for variationally sparse gaussian processes. In *Advances in Neural Information Processing Systems*, pages 1648–1656, 2015b.
- Daniel Hernández-Lobato and José Miguel Hernández-Lobato. Scalable gaussian process classification via expectation propagation. In *Artificial Intelligence and Statistics*, pages 168–176, 2016.
- José Miguel Hernández-Lobato, Yingzhen Li, Mark Rowland, Daniel Hernández-Lobato, Thang Bui, and Richard Turner. Black-box α -divergence minimization. 2016.
- Matthew D Hoffman, David M Blei, Chong Wang, and John Paisley. Stochastic variational inference. *The Journal of Machine Learning Research*, 14(1):1303–1347, 2013.
- Martin Jankowiak, Geoff Pleiss, and Jacob R Gardner. Deep sigma point processes. In *Proceedings of the 36th conference on Uncertainty in artificial intelligence*, 2020.
- Michael I Jordan, Zoubin Ghahramani, Tommi S Jaakkola, and Lawrence K Saul. An introduction to variational methods for graphical models. *Machine learning*, 37(2): 183–233, 1999.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Jeremias Knoblauch. Robust deep gaussian processes. *arXiv preprint arXiv:1904.02303*, 2019.
- Jeremias Knoblauch, Jack Jewson, and Theodoros Damoulas. Generalized variational inference. *arXiv preprint arXiv:1904.02063*, 2019.
- Yingzhen Li and Yarin Gal. Dropout inference in bayesian neural networks with alpha-divergences. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 2052–2061. JMLR. org, 2017.
- Yingzhen Li, José Miguel Hernández-Lobato, and Richard E Turner. Stochastic expectation propagation. In *Advances in neural information processing systems*, pages 2323–2331, 2015.
- Haitao Liu, Yew-Soon Ong, Xiaobo Shen, and Jianfei Cai. When gaussian process meets big data: A review of scalable gps. *arXiv preprint arXiv:1807.01065*, 2018.
- Alexander Graeme de Garis Matthews. *Scalable Gaussian process inference using variational methods*. PhD thesis, University of Cambridge, 2017.
- Thomas Minka. Power ep. *Dep. Statistics, Carnegie Mellon University, Pittsburgh, PA, Tech. Rep*, 2004.
- Joaquin Quiñonero-Candela and Carl Edward Rasmussen. A unifying view of sparse approximate gaussian process regression. *Journal of Machine Learning Research*, 6 (Dec):1939–1959, 2005.
- Maziar Raissi, Hessam Babaei, and George Em Karniadakis. Parametric gaussian process regression for big data. *Computational Mechanics*, 64(2):409–416, 2019.
- Carl Edward Rasmussen. Gaussian processes in machine learning. In *Summer School on Machine Learning*, pages 63–71. Springer, 2003.
- Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. *arXiv preprint arXiv:1401.4082*, 2014.
- Hugh Salimbeni, Ching-An Cheng, Byron Boots, and Marc Deisenroth. Orthogonally decoupled variational gaussian processes. In *Advances in neural information processing systems*, pages 8711–8720, 2018.
- Matthias Seeger, Christopher Williams, and Neil Lawrence. Fast forward selection to speed up sparse gaussian process regression. Technical report, 2003.
- Rishit Sheth and Roni Khardon. Monte carlo structured svi for two-level non-conjugate models. *arXiv preprint arXiv:1612.03957*, 2016.
- Rishit Sheth and Roni Khardon. Excess risk bounds for the bayes risk using variational inference in latent gaussian models. In *Advances in Neural Information Processing Systems*, pages 5151–5161, 2017.
- Rishit Sheth and Roni Khardon. Pseudo-bayesian learning via direct loss minimization with applications to sparse gaussian process models. 2019.
- Edward Snelson and Zoubin Ghahramani. Compact approximations to bayesian predictive distributions. In *Proceedings of the 22nd international conference on Machine learning*, pages 840–847. ACM, 2005.

Edward Snelson and Zoubin Ghahramani. Sparse gaussian processes using pseudo-inputs. In *Advances in neural information processing systems*, pages 1257–1264, 2006.

Michalis Titsias. Variational learning of inducing variables in sparse gaussian processes. In *Artificial Intelligence and Statistics*, pages 567–574, 2009.

R. E. Turner and M. Sahani. Two problems with variational expectation maximisation for time-series models. In D. Barber, T. Cemgil, and S. Chiappa, editors, *Bayesian Time series models*, chapter 5, pages 109–130. Cambridge University Press, 2011.

Ke Alexander Wang, Geoff Pleiss, Jacob R Gardner, Stephen Tyree, Kilian Q Weinberger, and Andrew Gordon Wilson. Exact gaussian processes on a million data points. *arXiv preprint arXiv:1903.08114*, 2019.

Andrew Gordon Wilson, Zhiting Hu, Ruslan Salakhutdinov, and Eric P Xing. Deep kernel learning. In *Artificial Intelligence and Statistics*, pages 370–378, 2016.