Deep Gaussian Processes for Regression using Approximate Expectation Propagation

Thang D. Bui[1] TDB40@CAM.AC.UK
José Miguel Hernández-Lobato[2] JMH@SEAS.HARVARD.EDU
Daniel Hernández-Lobato[3] DANIEL.HERNANDEZ@UAM.ES
Yingzhen Li[1] YL494@CAM.AC.UK
Richard E. Turner[1] RET26@CAM.AC.UK
[1]University of Cambridge, [2]Harvard University, [3]Universidad Autónoma de Madrid

Abstract

Deep Gaussian processes (DGPs) are multi-layer hierarchical generalisations of Gaussian processes (GPs) and are formally equivalent to neural networks with multiple, infinitely wide hidden layers. DGPs are nonparametric probabilistic models and as such are arguably more flexible, have a greater capacity to generalise, and provide better calibrated uncertainty estimates than alternative deep models. This paper develops a new approximate Bayesian learning scheme that enables DGPs to be applied to a range of medium to large scale regression problems for the first time. The new method uses an approximate Expectation Propagation procedure and a novel and efficient extension of the probabilistic backpropagation algorithm for learning. We evaluate the new method for non-linear regression on eleven real-world datasets, showing that it always outperforms GP regression and is almost always better than state-of-the-art deterministic and sampling-based approximate inference methods for Bayesian neural networks. As a by-product, this work provides a comprehensive analysis of six approximate Bayesian methods for training neural networks.

1 Introduction

Gaussian Processes (GPs) are powerful nonparametric distributions over continuous functions that can be used for both supervised and unsupervised learning problems (Rasmussen & Williams, 2005). In this article, we study a multi-layer hierarchical generalisation of GPs or deep Gaussian Processes (DGPs) (Damianou & Lawrence, 2013) for supervised learning tasks. A GP is equivalent to an infinitely wide neural network with single hidden layer and similarly a DGP is a multi-layer neural network with multiple infinitely wide hidden layers (Neal, 1995). The mapping between layers in this type of network is parameterised by a GP, and, as a result, DGPs retain useful properties of GPs such as nonparametric modelling power and well-calibrated predictive uncertainty estimates. In addition, DGPs employ a hierarchical structure of GP mappings and therefore are arguably more flexible, have a greater capacity to generalise, and are able to provide better predictive performance (Damianou, 2015). This family of models is attractive as it can also potentially discover layers of increasingly abstract data representations, in much the same way as their deep parametric counterparts, but it can also handle and propagate uncertainty in the hierarchy.

The addition of non-linear hidden layers can also potentially overcome practical limitations of *shallow* GPs. First, modelling real-world complex datasets often requires rich, hand-designed covariance functions. DGPs can perform input warping or dimensionality compression or expansion, and automatically learn to construct a kernel that works well for the data at hand. As a result, learning in this model provides a flexible form of Bayesian kernel design. Second, the functional mapping from inputs to outputs specified by a DGP is non-Gaussian which is a more general and flexible modelling choice. Third, DGPs can repair damage done by sparse approximations to the representational power of each GP layer. For example, pseudo datapoint based approximation methods for DGPs trade model complexity for a lower computational complexity of $\mathcal{O}(NLM^2)$ where N is the number of datapoints, L is the number of layers, and M is the number of pseudo datapoints. This complexity scales quadratically in M whereas the dependence on the number of layers L is only linear. Therefore, it can be

Proceedings of the 33^{rd} International Conference on Machine Learning, New York, NY, USA, 2016. JMLR: W&CP volume 48.

cheaper to increase the representation power of the model by adding extra layers rather than by adding more pseudo datapoints.

The focus of this paper is Bayesian learning of DGPs, which involves inferring the posterior over the layer mappings and hyperparameter optimisation via the marginal likelihood. Unfortunately, exact Bayesian learning in this model is analytically intractable and as such approximate inference is needed. Current proposals in the literature do not scale well and have not been compared to alternative deep Bayesian models. We will first review the model and past work in Section 2, and then make the following contributions: (i) we propose a new approximate inference scheme for DGPs for regression, using a sparse GP approximation, a novel approximate Expectation Propagation scheme and the probabilistic backpropagation algorithm, resulting in a computationally efficient, scalable and easy to implement algorithm (Sections 3, 4 and 5), (ii) we evaluate the performance of the new method in supervised learning tasks on various medium to large scale datasets and show that the proposed method is always better than GP regression and is almost always better than state-of-the-art approximate inference techniques for multi-layer neural networks (Section 7).

2 Deep Gaussian processes

We first review DGPs and existing literature on approximate inference and learning for DGPs. Suppose we have a training set comprising of N D-dimensional input and observation pairs $(\mathbf{x}_n, y_n)$. For ease of presentation, the outputs are assumed to be real-valued scalars, but other types of data can be easily accommodated[1]. The probabilistic representation of a DGP comprising L layers can be written as follows,

$$p(f_l|\Theta_l) = \mathcal{GP}(f_l; \mathbf{0}, \mathbf{K}_l), \quad l = 1, \cdots, L$$
$$p(\mathbf{h}_l|f_l, \mathbf{h}_{l-1}, \sigma_l^2) = \prod_n \mathcal{N}(h_{l,n}; f_l(h_{l-1,n}), \sigma_l^2), \quad h_{1,n} = \mathbf{x}_n$$
$$p(\mathbf{y}|f_L, \mathbf{h}_{L-1}, \sigma_L^2) = \prod_n \mathcal{N}(y_n; f_L(h_{L-1,n}), \sigma_L^2)$$

where hidden layers[2] are denoted $h_{l,n}$ and the functions in each layer, f_l. More formally, we place a zero mean GP prior over the mapping f_l, that is, given the inputs to f_l any finite set of function values are distributed under the prior according to a multivariate Gaussian $p(\mathbf{f}_l) = \mathcal{N}(\mathbf{f}; \mathbf{0}, \mathbf{K}_{\mathbf{ff}})$. Note that these function values and consequently the hidden variables are not marginally normally distributed, as the inputs are random variables. When $L = 1$, the model described above collapses back to GP regression. When the inputs $\{\mathbf{x}_n\}$ are unknown and random, the model becomes a DGP latent variable model, which has been studied in (Damianou & Lawrence, 2013).

An example of DGPs when $L = 2$ and $\dim(h_1) = 2$ is shown in Figure 1. We use this network with the proposed approximation and training algorithm to fit a value function of the mountain car problem (Sutton & Barto, 1998) from a small number of noisy evaluations. This function is particularly difficult for models such as GP regression with a standard exponentiated quadratic kernel due to a *steep value function cliff*, but is well handled by a DGP with only two GP layers. Interestingly the functions in the first layer are fairly simple and learn to cover or explain different parts of the input space.

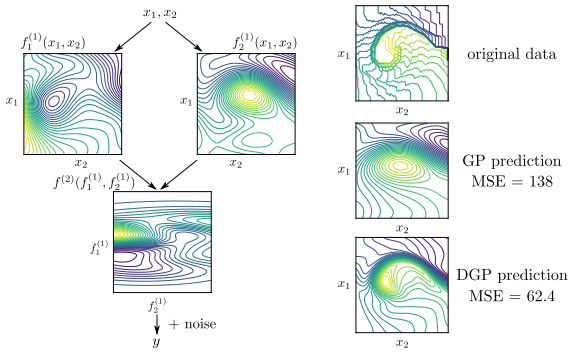

Figure 1. A deep GP example that has two GP layers and one 2-D hidden layer. The training output is the state values of the mountain car problem. The left graphs show latent functions in each layer, two functions in the first layer and one in the second layer, learnt by using the proposed approach. The right graph shows the training data [top] and the predictions of the overall function mapping from inputs to outputs made by a GP [middle] and the DGP on the left [bottom].

We are interested in inferring the posterior distribution over the latent function mappings and the intermediate hidden variables, as well as obtaining a marginal likelihood estimate for hyperparameter tuning and model comparison. Due to the nonlinearity in the hierarchy, these quantities are analytically intractable. As such, approximate inference is required. The simplest approach is to obtain the *maximum a posteriori* estimate of the hidden variables (Lawrence & Moore, 2007). However, this procedure is prone to over-fitting and does not provide uncertainty estimates. An alternative existing approach is based on a variational-free-energy method proposed by Damianou and Lawrence (2013), extending the seminal work on variational sparse GPs by Titsias (2009). In this scheme, a variational approximation over both latent functions and hidden variables is chosen such that a free energy is both computationally and analytically tractable. Critically, as

[1] We also discuss how to handle non-Gaussian likelihoods in the supplementary material.

[2] Hidden variables in the intermediate layers can and will generally have multiple dimensions but we have omitted this here to lighten the notation.

a variational distribution over the hidden variables is used in this approach, in addition to one over the inducing outputs, the number of variational parameters increases linearly with the number of training datapoints which hinders the use of this method for large scale datasets. Furthermore, initialisation for this scheme is a known issue, even for a modest number of datapoints (Turner & Sahani, 2011). An extension of (Damianou & Lawrence, 2013) that has skip links from the inputs to every hidden layer in the network was proposed by Dai, Damianou, González, and Lawrence (2016), based on suggestions provided in (Duvenaud et al., 2014). Recent work by Hensman and Lawrence (2014) introduces a nested variational scheme that only requires a variational distribution over the inducing outputs, removing the parameter scaling problem of (Damianou & Lawrence, 2013). However, both approaches of (Dai et al., 2016) and (Hensman & Lawrence, 2014) have not been fully evaluated on benchmark supervised learning tasks or on medium to large scale datasets, nor compared to alternative deep models.

A special case of DGPs when $L = 2$ and the sole hidden layer h_1 is only one dimensional is warped GPs (Snelson et al., 2004; Lázaro-Gredilla, 2012). In (Lázaro-Gredilla, 2012) a variational approach, in a similar spirit to (Titsias, 2009) and (Damianou & Lawrence, 2013), was used to jointly learn the latent functions. In contrast, the latent function in the second layer is assumed to be deterministic and parameterised by a small set of parameters by Snelson, Rasmussen, and Ghahramani (2004), which can be learnt by maximising the analytically tractable marginal likelihood. However, the performance of warped GPs is often similar to a standard GP, most likely due to the narrow bottleneck in the hidden layer.

This paper differs substantially from previous work developing three seemingly aggressive approximations that significantly enhance scalability of DGPs, and which perhaps surprisingly lead to excellent practical performance. First, in order to sidestep the cubic computational cost of GPs we leverage a well-known pseudo point sparse approximation (Snelson & Ghahramani, 2006; Quiñonero-Candela & Rasmussen, 2005). Second, an approximation to the Expectation Propagation (EP) energy function (Seeger, 2007), a marginal likelihood estimate, is optimised directly to find an approximate posterior over the inducing outputs. Third, the optimisation demands analytically intractable moments that are approximated by nesting Assumed Density Filtering (Hernández-Lobato & Adams, 2015). The proposed algorithm is not restricted to the warped GP case and is applicable to non-Gaussian observation models.

The complexity of our method is similar to that of the variational approach proposed by Damianou and Lawrence (2013), $\mathcal{O}(NLM^2)$, but is much less memory intensive, $\mathcal{O}(LM^2)$ vs. $\mathcal{O}(NL+LM^2)$. These costs are competitive to those of the nested variational approach in (Hensman & Lawrence, 2014).

3 The Fully Independent Training Conditional approximation

The computational complexity of full GP models scales cubically with the number of training instances, making it intractable in practice. Sparse approximation techniques are therefore often necessary. They can be coarsely put into two classes: ones that explicitly sparsify and create a parametric representation that approximates the original model, and ones that retain the original nonparametric properties and perform sparse approximation to the exact posterior. The method used here, Fully Independent Training Conditional (FITC), can be interpreted as a member of the first category (Bui et al., 2016). The FITC approximation is formed by considering a small set of M function values $\mathbf{u}$ in the infinite dimensional vector f and assuming conditional independence between the remaining values given the set $\mathbf{u}$ (Snelson & Ghahramani, 2006; Quiñonero-Candela & Rasmussen, 2005). This set is often called inducing outputs or pseudo targets and their input locations $\mathbf{z}$ can be chosen by optimising the approximate marginal likelihood. The resulting model can be written as follows,

$$p(\mathbf{u}_l|\Theta_l) = \mathcal{N}(\mathbf{u}_l; \mathbf{0}, \mathbf{K}_{\mathbf{u}_{l-1},\mathbf{u}_{l-1}}), \quad l = 1, \cdots, L$$
$$p(\mathbf{h}_l|\mathbf{u}_l, \mathbf{h}_{l-1}, \sigma_l^2) = \prod_n \mathcal{N}(h_{l,n}; \mathbf{C}_{n,l}\mathbf{u}_l, \mathbf{R}_{n,l}),$$
$$p(\mathbf{y}|\mathbf{u}_L, \mathbf{H}_{L-1}, \sigma_L^2) = \prod_n \mathcal{N}(y_n; \mathbf{C}_{n,L}\mathbf{u}_L, \mathbf{R}_{n,L}).$$

where $\mathbf{C}_{n,l} = \mathbf{K}_{\mathbf{h}_{l,n},\mathbf{u}_l}\mathbf{K}^{-1}_{\mathbf{u}_l,\mathbf{u}_l}$ and $\mathbf{R}_{n,l} = \mathbf{K}_{\mathbf{h}_{l,n},\mathbf{h}_{l,n}} - \mathbf{K}_{\mathbf{h}_{l,n},\mathbf{u}_l}\mathbf{K}^{-1}_{\mathbf{u}_l,\mathbf{u}_l}\mathbf{K}_{\mathbf{u}_l,\mathbf{h}_{l,n}} + \sigma_l^2\mathrm{I}$. Note that the function outputs index the covariance matrices, for example $\mathbf{K}_{\mathbf{h}_{l,n},\mathbf{u}_l}$ denotes the covariance between $\mathbf{h}_{l,n}$ and $\mathbf{u}_l$, and takes $\mathbf{h}_{l-1,n}$ and $\mathbf{z}_l$ as inputs respectively. This is important when calculating how to propagate uncertainty through the network. The FITC approximation creates a parametric model, but one which is cleverly structured so that the induced non-stationary noise captures the uncertainty introduced from the sparsification. The computational complexity of inference and hyperparameter tuning in this approximate model is $\mathcal{O}(NM^2)$ which means M needs to be smaller than N to provide any computational gain (i.e. the approximation should be sparse). The quality of the approximation largely depends on the number of inducing outputs M and the complexity of the underlying function, i.e. if the function's characteristic lengthscale is small, M needs to be large and vice versa. As M tends to N and $\mathbf{z} = \mathbf{X}$, i.e. the inducing inputs and training inputs are shared, the approximate model reverts back to the original GP model. The graphical model is shown in Figure 2 [left].

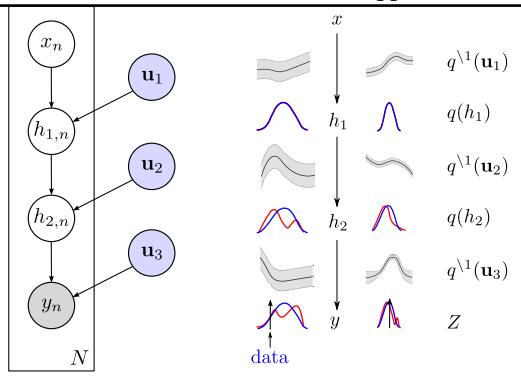

Figure 2. Left: The graphical model of our DGP-FITC model where the inducing outputs $\{\mathbf{u}_l\}$ play a role of global parameters. Right: A Gaussian moment-matching procedure to compute $\log \mathcal{Z}$. The bottom arrows denote the value of the observation and the left and right graphs [before and after an update respectively] show how the algorithm makes the final propagated Gaussian fit to the data, i.e. the model is trained so that training points are more probable after each update. The red curves show the distribution over hidden variables before being approximated by a Gaussian in blue. Best viewed in colour.

4 Approximate Bayesian inference via EP

Having specified a probabilistic model for data using a deep sparse Gaussian processes we now consider inference for the inducing outputs $\mathbf{u} = \{\mathbf{u}_l\}_{l=1}^L$ and learning of the model parameters $\alpha = \{\mathbf{z}_l, \Theta_l\}_{l=1}^L$. The posterior distribution over the inducing outputs can be written as $p(\mathbf{u}|\mathbf{X}, \mathbf{y}) \propto p(\mathbf{u}) \prod_n p(y_n|\mathbf{u}, \mathbf{X}_n)$. This quantity can then be used for output prediction given a test input, $p(y^*|\mathbf{x}^*, \mathbf{X}, \mathbf{y}) = \int \mathrm{d}\mathbf{u}\, p(\mathbf{u}|\mathbf{X}, \mathbf{y})\, p(y^*|\mathbf{u}, \mathbf{x}^*)$. The model hyperparameters can be tuned by maximising the marginal likelihood $p(\mathbf{y}|\alpha) = \int \mathrm{d}\mathbf{u}\, \mathrm{d}\mathbf{h}\, p(\mathbf{u}, \mathbf{h})\, p(\mathbf{y}|\mathbf{u}, \mathbf{h}, \alpha)$. However, both the posterior of $\mathbf{u}$ and the marginal likelihood are not analytically tractable when there is more than one GP layer in the model. As such, approximate inference is needed; here we make use of the EP energy function with a tied factor constraint similar to that proposed in the Stochastic Expectation Propagation (SEP) algorithm (Li et al., 2015) to produce a scalable, convergent approximate inference method.

4.1 EP, Stochastic EP and the EP approximate energy

In EP (Minka, 2001), the approximate posterior is assumed to be $q(\mathbf{u}) \propto p(\mathbf{u}) \prod_n \tilde{t}_n(\mathbf{u})$ where $\{\tilde{t}_n(\mathbf{u})\}_{n=1}^N$ are the approximate data factors. Each factor approximately captures the contribution of datapoint n makes to the posterior and, in this work, they take an unnormalised Gaussian form. The factors can be found by running an iterative procedure which often requires several passes through the training set for convergence[3]. The EP algorithm also provides an approximation to the marginal likelihood,

$$\log p(\mathbf{y}|\alpha) \approx \mathcal{F}(\alpha) = \phi(\theta) - \phi(\theta_{\text{prior}}) + \sum_{n=1}^N \log \tilde{\mathcal{Z}}_n$$
$$\text{where} \quad \log \tilde{\mathcal{Z}}_n = \log \mathcal{Z}_n + \phi(\theta^{\setminus n}) - \phi(\theta),$$

where $\theta, \theta^{\setminus n}$ and θ_{prior} are the natural parameters of $q(\mathbf{u})$, the cavity $q^{\setminus n}(\mathbf{u})$ $[q^{\setminus n}(\mathbf{u}) \propto q(\mathbf{u})/\tilde{t}_n(\mathbf{u})]$ and $p(\mathbf{u})$ respectively, $\phi(\theta)$ is the log normaliser of a Gaussian distribution with natural parameters θ, and $\log \mathcal{Z}_n = \log \int \mathrm{d}\mathbf{u}\, q^{\setminus n}(\mathbf{u})\, p(y_n|\mathbf{u}, \mathbf{X}_n)$ (Seeger, 2007). Unfortunately, EP is not guaranteed to converge, but if it does, the fixed points lie at the stationary points of the EP energy, which is given by $-\mathcal{F}(\alpha)$. Furthermore, EP requires the approximate factors to be stored in memory, which has a cost of $\mathcal{O}(NLM^2)$ in this application as we need to store the mean and the covariance matrix for each factor.

4.2 Direct EP energy minimisation with a tied factor constraint

In order to reduce the expensive memory footprint of EP, the data-factors are tied. That is the posterior $p(\mathbf{u}|\mathbf{X}, \mathbf{y})$ is approximated by $q(\mathbf{u}) \propto p(\mathbf{u}) g(\mathbf{u})^N$, where the factor $g(\mathbf{u})$ could be thought of as an *average* data factor that captures the average effect of a likelihood term on the posterior. Approximations of this form were recently used in the SEP algorithm (Li et al., 2015) and although seemingly limited, in practice were found to perform almost as well as full EP while significantly reducing EP's memory requirement, from $\mathcal{O}(NLM^2)$ to $\mathcal{O}(LM^2)$ in our case.

The original SEP work devised modified versions of the EP updates appropriate for the new form of the approximate posterior. Originally we applied this method to DGPs (details of this approach including hyperparameter optimisation are included in the supplementary material). However, an alternative approach was found to have superior performance, which is to optimise the EP energy function directly to refine the approximating factors. The benefit is that the approximate EP energy can be (jointly) optimised for the hyperparameters (including the inducing inputs) at the same time. Normally, optimisation of the EP energy requires a double-loop algorithm, which is computationally inefficient, however the use of tied factors simplifies the approximate marginal likelihood and allows direct optimisation. The energy becomes,

$$\mathcal{F}(\alpha) = \phi(\theta) - \phi(\theta_{\text{prior}}) + \sum_{n=1}^N \Big[\log \mathcal{Z}_n + \phi(\theta^{\setminus 1}) - \phi(\theta)\Big]$$
$$= (1-N)\phi(\theta) + N\phi(\theta^{\setminus 1}) - \phi(\theta_{\text{prior}}) + \sum_{n=1}^N \log \mathcal{Z}_n$$

since the cavity distribution $q^{\setminus n}(\mathbf{u}) \propto q(\mathbf{u})/\tilde{t}_n(\mathbf{u}) = q(\mathbf{u})/g(\mathbf{u}) = q^{\setminus 1}(\mathbf{u})$ is the same for all training points.

[3] We summarise the EP steps in the supplementary material.

This elegantly removes the need for a double-loop algorithm, since we can posit a form for the approximate posterior and optimise the above approximate marginal likelihood directly. However, it is important to note that, in general, optimising this objective will not give the same solution as optimising the full EP energy. The new energy produces an approximation formed by averaging the moments of $q^{\backslash 1}(\mathbf{u})\, p(y_n|\mathbf{u}, \mathbf{x}_n)$ over datapoints, whereas EP averages natural parameters, which is arguably more sensible but less tractable.

In detail, we assume the tied factor takes a Gaussian form with natural parameters θ_1. As a result, the approximate posterior and the cavity are also Gaussian with natural parameters $\theta = \theta_{\text{prior}} + N\theta_1$ and $\theta^{\backslash 1} = \theta_{\text{prior}} + (N-1)\theta_1$ respectively. This means that we can compute the first three terms in the energy function exactly. However, it remains to compute $\log \mathcal{Z}_n = \log \int \mathrm{d}\mathbf{u}\, q^{\backslash 1}(\mathbf{u})\, p(y_n|\mathbf{u}, \mathbf{x}_n)$ which we will discuss next.

5 Probabilistic backpropagation for deep Gaussian processes

Computing $\log \mathcal{Z}_n$ in the objective function above is analytically intractable for $L \geq 1$ since the likelihood given the inducing outputs $\mathbf{u}$ is nonlinear and the propagation of the Gaussian cavity through each layer results in a complex distribution. However, for certain choices of covariance functions $\{\mathbf{K}_l\}_{l=1}^{L}$, it is possible to use an efficient and accurate approximation which propagates a Gaussian through the first layer of the network and projects this non-Gaussian distribution back to a moment matched Gaussian before propagating through the next layer and repeating the same steps. This scheme is algorithmically identical to Assumed Density Filtering (Kushner & Budhiraja, 2000) and a central part of the probabilistic backpropagation algorithm that has been applied to standard neural networks (Hernández-Lobato & Adams, 2015).

The aim is to compute $\log \mathcal{Z}$ and its gradients with respect to the parameters such as θ_1 or the hyperparameters of the model[4]. By reintroducing the hidden variables in the middle layers, we perform a Gaussian approximation to $\mathcal{Z}$ in a sequential fashion, as illustrated in Figure 2 [right]. We take a two layer case as a running example:

$$\begin{aligned}\mathcal{Z} &= \int \mathrm{d}\mathbf{u}\, p(y|\mathbf{x}, \mathbf{u})\, q^{\backslash 1}(\mathbf{u}) \\ &= \int \mathrm{d}h_1\, \mathrm{d}\mathbf{u}_2\, p(y|h_1, \mathbf{u}_2)\, q^{\backslash 1}(\mathbf{u}_2) \int \mathrm{d}\mathbf{u}_1\, p(h_1|\mathbf{x}, \mathbf{u}_1)\, q^{\backslash 1}(\mathbf{u}_1)\end{aligned}$$

One key difference between our approach and the variational free energy method in (Damianou & Lawrence, 2013) is that our algorithm does not retain an explicit approximate distribution over the hidden variables. Instead, we approximately integrate them out when computing $\log \mathcal{Z}$ as follows.

First, we can exactly marginalise out the inducing outputs for each GP layer, leading to $\mathcal{Z} = \int \mathrm{d}h_1\, q(y|h_1)\, q(h_1)$ where $q(h_1) = \mathcal{N}(h_1; m_1, v_1)$, $q(y|h_1) = \mathcal{N}(y|h_1; m_{2|h_1}, v_{2|h_1})$ and $m_1 = \mathbf{K}_{h_1,\mathbf{u}_1}\mathbf{K}_{\mathbf{u}_1,\mathbf{u}_1}^{-1}\mathbf{m}_1^{\backslash 1}$, $v_1 = \sigma_1^2 + K_{h_1,h_1} - \mathbf{K}_{h_1,\mathbf{u}_1}\mathbf{K}_{\mathbf{u}_1,\mathbf{u}_1}^{-1}\mathbf{K}_{\mathbf{u}_1,h_1} + \mathbf{K}_{h_1,\mathbf{u}_1}\mathbf{K}_{\mathbf{u}_1,\mathbf{u}_1}^{-1}\mathbf{V}_1^{\backslash 1}\mathbf{K}_{\mathbf{u}_1,\mathbf{u}_1}^{-1}\mathbf{K}_{\mathbf{u}_1,h_1}$, $m_{2|h_1} = \mathbf{K}_{h_2,\mathbf{u}_2}\mathbf{K}_{\mathbf{u}_2,\mathbf{u}_2}^{-1}\mathbf{m}_2^{\backslash 1}$, $v_{2|h_1} = \sigma_2^2 + K_{h_2,h_2} - \mathbf{K}_{h_2,\mathbf{u}_2}\mathbf{K}_{\mathbf{u}_2,\mathbf{u}_2}^{-1}\mathbf{K}_{\mathbf{u}_2,h_2} + \mathbf{K}_{h_2,\mathbf{u}_2}\mathbf{K}_{\mathbf{u}_2,\mathbf{u}_2}^{-1}\mathbf{V}_1^{\backslash 1}\mathbf{K}_{\mathbf{u}_2,\mathbf{u}_2}^{-1}\mathbf{K}_{\mathbf{u}_2,h_2}$.

Following (Girard et al., 2003; Barber & Schottky, 1998; Deisenroth & Mohamed, 2012), we can use the law of iterated conditionals to approximate the difficult integral in the equation above by a Gaussian $\mathcal{Z} \approx \mathcal{N}(y|m_2, v_2)$ where the mean and variance take the following form,

$$\begin{aligned} m_2 &= \mathrm{E}_{q(h_1)}[m_{2|h_1}] \\ v_2 &= \mathrm{E}_{q(h_1)}[v_{2|h_1}] + \mathrm{var}_{q(h_1)}[m_{2|h_1}] \end{aligned}$$

which results in $m_2 = \mathrm{E}_{q(h_1)}[\mathbf{K}_{h_2,\mathbf{u}_2}]\mathbf{A}$ and $v_2 = \sigma_2^2 + \mathrm{E}_{q(h_1)}[K_{h_2,h_2}] + \mathrm{tr}\left(\mathbf{B}\mathrm{E}_{q(h_1)}[\mathbf{K}_{\mathbf{u}_2,h_2}\mathbf{K}_{h_2,\mathbf{u}_2}]\right) - m_2^2$ where $\mathbf{A} = \mathbf{K}_{\mathbf{u}_2,\mathbf{u}_2}^{-1}\mathbf{m}_2^{\backslash 1}$ and $\mathbf{B} = \mathbf{K}_{\mathbf{u}_2,\mathbf{u}_2}^{-1}(\mathbf{V}_2^{\backslash 1} + \mathbf{m}_2^{\backslash 1}\mathbf{m}_2^{\backslash 1,\mathrm{T}})\mathbf{K}_{\mathbf{u}_2,\mathbf{u}_2}^{-1} - \mathbf{K}_{\mathbf{u}_2,\mathbf{u}_2}^{-1}$. The equations above require the expectations of the kernel matrix under a Gaussian distribution over the inputs, which are analytically tractable for widely used kernels such as exponentiated quadratic, linear or a more general class of spectral mixture kernels (Titsias & Lawrence, 2010; Wilson & Adams, 2013). In addition, the approximation above can be improved for networks that have multidimensional intermediate variables, by using a Gaussian with a non-diagonal covariance matrix. We discuss this in the supplementary material.

As the mean and variance of the Gaussian approximation in each intermediate layer can be computed analytically, their gradients with respect to the mean and variance of the input distribution, as well as the parameters of the current layers are also available. Since we require the gradients of the approximation to $\log \mathcal{Z}$, we need to store these results in the forward propagation step, compute the approximate $\log \mathcal{Z}$ and its gradients at the output layer and use the chain rule in the backward step to differentiate through the ADF procedure. This procedure is reminiscent of the backpropagation algorithm in standard parametric neural networks, hence the name *probabilistic backpropagation* (Hernández-Lobato & Adams, 2015).

6 Stochastic optimisation for scalable training

The propagation and moment-matching as described above costs $\mathcal{O}(LM^2)$ and needs to be repeated for all datapoints

[4] We ignore the data index here to lighten the notation

in the training set in batch mode, resulting in an overall complexity of $\mathcal{O}(NLM^2)$. Fortunately, the last term of the objective in Section 4.2 is a sum of independent terms, i.e. its computation can be distributed, resulting in a substantial decrease in computational cost. Furthermore, the objective is suitable for stochastic optimisation. In particular, an unbiased noisy estimate of the objective and its gradients can be obtained using a minibatch of training datapoints,

$$\mathcal{F} \approx -(N-1)\phi(\theta) + N\phi(\theta^{\backslash 1}) - \phi(\theta_{\text{prior}}) + \frac{N}{|B|}\sum_{b=1}^{|B|} \log \mathcal{Z}_b,$$

where $|B|$ denotes the minibatch size.

7 Experiments

We implement and compare the proposed approximation scheme to state-of-the-art methods for Bayesian neural networks. We first detail our implementation in Section 7.1 and then discuss the experimental results in Sections 7.2 and 7.3. We released our Theano and Python implementations on https://github.com/thangbui/deepGP_approxEP.

7.1 Experimental details

In all the experiments reported here, we use Adam with the default learning rate (Kingma & Ba, 2015) for optimising our objective function. We use an exponentiated quadratic kernel with ARD lengthscales for each layer. The hyperparameters and pseudo point locations are different between functions in each layer. The lengthscales and inducing inputs of the first GP layer are sensibly initialised based on the median distance between datapoints in the input space and the k-means cluster centers respectively. We use long lengthscales and initial inducing inputs between $[-1, 1]$ for the higher layers to force them to start with an identity mapping. We parameterise the natural parameters of the average factor and initialise them with small random values. We evaluate the predictive performance on the test set using two popular metrics: root mean squared error (RMSE) and mean log likelihood (MLL).

7.2 Regression on UCI datasets

We validate the proposed approach for training DGPs in regression experiments using several datasets from the UCI repository. In particular, we use the ten datasets and train/test splits used by Hernández-Lobato and Adams (2015) and Gal and Ghahramani (2016): 1 split for the *year* dataset [$N \approx 500000, D = 90$], 5 splits for the *protein* dataset [$N \approx 46000, D = 9$], and 20 for the others.

We compare our method (FITC-DGP) against sparse GP regression using FITC (FITC-GP) and Bayesian neural network (BNN) regression using several state-of-the-art deterministic and sampling-based approximate inference techniques. As baselines, we include the results for BNNs reported in (Hernández-Lobato & Adams, 2015), BNN-VI(G)-1 and BNN-PBP-1, and in (Gal & Ghahramani, 2016), BNN-Dropout-1. The results reported for these methods are for networks with one hidden layer of 50 units (100 units for *protein* and *year*). Specifically, BNN-VI(G) uses a mean-field Gaussian approximation for the weights in the network, and obtains the stochastic estimates of the bound and its gradient using a Monte Carlo approach (Graves, 2011). BNN-PBP employs Assumed Density Filtering and the probabilistic backpropagation algorithm to obtain a Gaussian approximation for the weights (Hernández-Lobato & Adams, 2015). BNN-Dropout is a recently proposed technique that employs *dropout* during training as well as at prediction time, that is to average over several predictions, each made by the entire network with a random proportion of the weights set to zero (Gal & Ghahramani, 2016). We implement other methods as follows,

- DGP: we evaluate three different architectures of DGPs, each with two GP layers and one hidden layer of one, two and three dimensions respectively (DGP-1, DGP-2 and DGP-3). We include the results for two settings of the number of inducing outputs, $M = 50$ and $M = 100$ respectively. Note that for the bigger datasets *protein* and *year*, we use $M = 100$ and $M = 200$ but do not annotate this in Figure 3. We choose these settings to ensure the run time for our method is smaller or comparable to that of other methods for BNNs.
- GP: we use the same number of pseudo datapoints as in DGP (GP 50 and GP 100).
- BNN-VI(KW): this method, similar to Graves (2011), employs a mean-field Gaussian variational approximation but evaluates the variational free energy using the *reparameterisation trick* proposed by Kingma and Welling (2014). We use a diagonal Gaussian prior for the weights and fix the prior variance to 1. The noise variance of the Gaussian noise model is optimised together with the means and variances of the variational approximation using the variational free energy. We test two different network architectures with the rectified linear activation function, and one and two hidden layers, each of 50 units (100 for the two big datasets), denoted by VI(KW)-1 and VI(KW)-2 respectively.
- BNN-SGLD: we reuse the same networks with one and two hidden layers as with VI(KW) and approximately sample from the posterior over the weights using Stochastic Gradient Langevin Dynamics (SGLD) (Welling & Teh, 2011). We place a diagonal Gaussian prior over the weights, and parameterise the observation noise variance as $\sigma^2 = \log(1 + \exp(\kappa))$, a broad Gaussian prior over κ and sample κ using the same SGLD

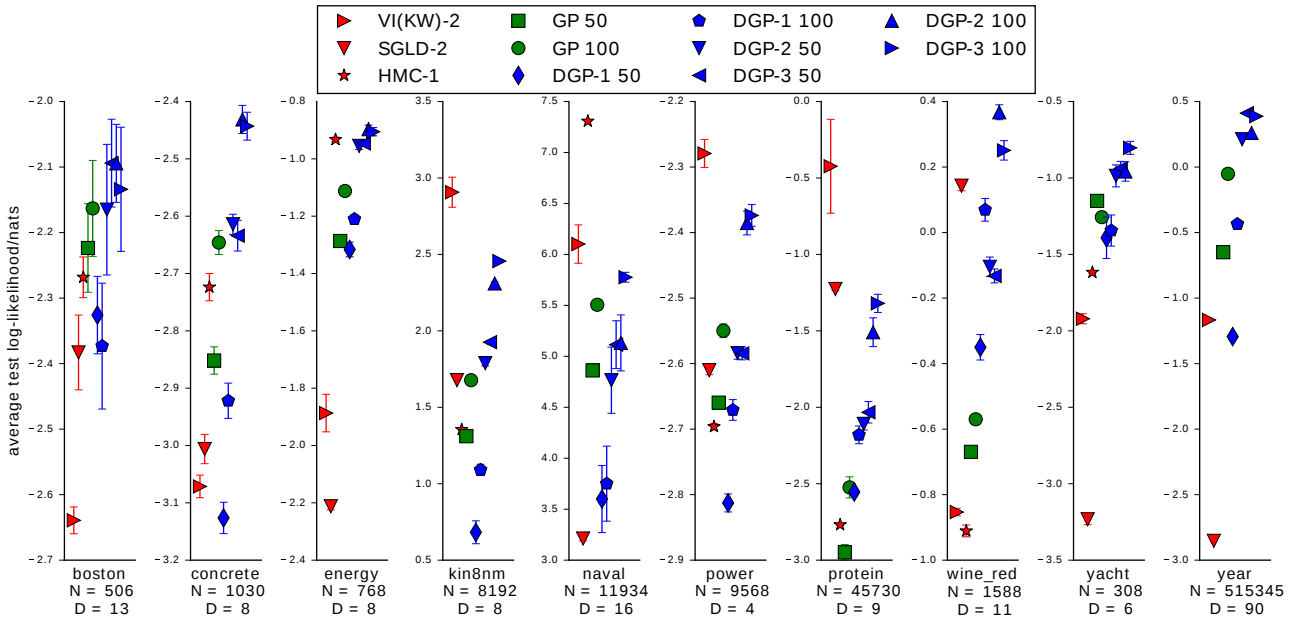

Figure 3. Average predictive log likelihood of existing approaches for BNNs and GPs, and the proposed method for DGPs, across 10 datasets. The higher the better, and best viewed in colour. Full results are included in the supplementary material.

procedure. Two step sizes, one for the weights and one for κ, were manually tuned for each dataset. We use Autograd for the implementation of BNN-SGLD and BNN-VI(KW) (github.com/HIPS/autograd).

- BNN-HMC: We run Hybrid Monte Carlo (HMC) (Neal, 1993) using the MCMCstuff toolbox (Vanhatalo & Vehtari, 2006) for networks with one hidden layer. We place a Gaussian prior over the network weights and a broad inverse Gamma hyper-prior for the prior variance. We also assume an inverse Gamma prior over the observation noise variance. The number of leapfrog steps and step size are first tuned using Bayesian optimisation using the pybo package (github.com/mwhoffman/pybo). Note that this procedure takes a long time (e.g. 3 days for protein) and the *year* dataset is too large to be handled in this way.

Figure 3 shows the average test log likelihood (MLL) for a subset of methods with their standard errors. We exclude methods that perform consistently poorly to improve readability. Full results and many more comparisons are included in the supplementary material. We also evaluate the average rank of the MLL performance of all methods across the datasets and their train/test splits and include the results in Figure 4. This figure is generated using the comparison scheme provided by Demšar (2006), and shows statistical differences in the performance of the methods. More precisely, if the gap between the average ranks of any two methods is above the critical distance (shown on the top right), the two methods' performances are statistically significantly different. Methods that are not significantly different from each other are linked by a solid line.

The rank result shows that DGPs with our inference scheme are the best performing methods overall. Specifically, the DGP-3-100 architecture obtains the best performance on 6 out of 10 datasets and are competitive on the remaining four datasets. The performance of other DGP variants follow closely with the exception for DGP-1 which is a standard warped GP, the network with one dimensional hidden layer. DGP-1 performs poorly compared to GP regression, but is still competitive with several methods for BNNs. The results also strongly indicate that the predictive performance is almost always improved by adding extra hidden layers or extra hidden dimensions or extra inducing outputs.

The best non-GP method is BNN-VI(KW)-2 which obtains the best performance on three datasets. However, this method performs poorly on 6 out of 7 remaining datasets, pushing down the corresponding average rank. Despite this, VI(KW) is the best method of all deterministic approximations for BNNs with one or two hidden layers. Overall, the VI approach without the *reparameterisation trick* of Graves, Dropout and PBP perform poorly in comparison and give inaccurate predictive uncertainty.

Sampling based methods such as SGLD and HMC obtain good predictive performance overall, but often require more tuning compared to other methods. In particular, HMC appears superior on one dataset, and competitive with DGPs on three other datasets; however, this method does not scale to large datasets.

The results for the RMSE metric follow the same trends with DGP-2 and DGP-3 performing as well or better compared to other methods. Interestingly, BNN-SGLD, despite

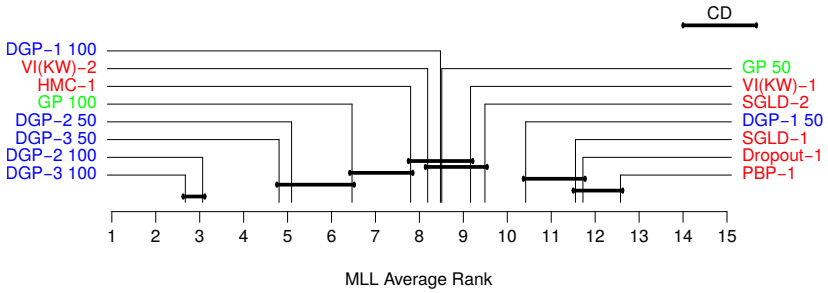

Figure 4. The average rank of all methods across the datasets and their train/test splits, generated based on Demšar (2006). See the text for more details.

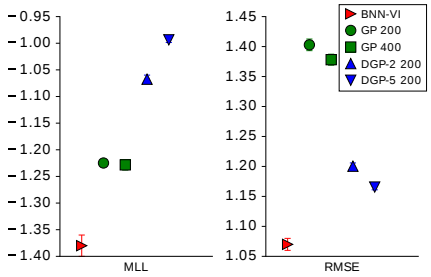

Figure 5. MLL and RMSE results for the photovoltaic molecule regression experiment.

being ranked relatively low according to the MLL metric, often provides good RMSE results perhaps unsurprisingly given the algorithm's similarity to stochastic gradient ascent. In addition, we compare the proposed method against the nested variational approach (Nested VI) by Hensman & Lawrence (2014), demonstrating that Nested VI consistently performs poorly compared to our method. Full results are included in the supplementary material.

7.3 Predicting the efficiency of organic photovoltaic molecules

Having demonstrated the performance of our inference scheme for DGPs, we carry out an additional regression experiment on a challenging dataset. We obtain a subset of 60,000 organic molecules and their power conversion efficiency from the Harvard Clean Energy Project (HCEP) (available at http://www.molecularspace.org) (Hachmann et al., 2011). We use 50,000 molecules for training and 10,000 for testing. The molecules are represented using 512-dimensional binary feature vectors, which were generated using the RDKit package, based on the molecular structures in the canonical SMILES format and a bond radius of 2. The power conversion efficiency of these molecules was estimated using density functional theory, determining whether a molecule could be potentially used as solar cell. The overall aim of the HCEP is to find *organic* molecules that are as efficient as their *silicon* counterparts. Our aim here is to show DGPs are effective predictive models that provide good uncertainty estimates, which can be used for tasks such as Bayesian Optimisation.

We test the method on two DGPs with one hidden layer of 2 and 5 dimensions, denoted by DGP-2 and DGP-5 respectively and each GP is sparsified using 200 inducing outputs. We compare these against two FITC-GP architectures with 200 and 400 pseudo datapoints respectively. We also repeat the experiment using a Bayesian neural network with two hidden layers, each of 400 hidden units. We use the variational approach with the *reparameterisation trick* by Kingma and Welling (2014) to perform inference in this model. The noise variance was fixed to 0.16 based on a suggestion in (Pyzer-Knapp et al., 2015). Figure 5 shows the predictive performance by five architectures. The DGP with a five dimensional hidden layer significantly outperforms others in terms of test MLL, including the shallow structure with considerably more pseudo datapoints. This result demonstrates the efficacy of DGPs in providing good predictive uncertainty estimates, even when the kernel used is a *simple* exponentiated quadratic kernel and the input features are binary. Surprisingly, VI(KW), although performing poorly as measured by the MLL, makes good predictions for the mean.

8 Summary

This paper has introduced a new and powerful deterministic approximation scheme for DGPs based upon an approximate EP algorithm and the FITC approximation to sidestep the computational and analytical intractability. A novel extension of the probabilistic backpropagation algorithm was developed to address a difficult marginalisation problem in the approximate EP algorithm used. The new method was evaluated on eleven datasets and compared against a number of state-of-the-art algorithms for Bayesian neural networks. The results show that the new method for training DGPs is superior on 7 out of 11 datasets considered, and performs comparably on the remainder, demonstrating that DGPs are a competitive alternative to multi-layer Bayesian neural networks for supervised learning tasks.

The proposed method, in principle, can be applied to classification and unsupervised learning. However, initial work on classification using DGPs, as included in the supplementary, does not show a substantial gain over a GP. This issue is potentially related to the diagonal Gaussian approximation currently used for the hidden layers from the second layer onwards. A non-diagonal approximation is feasible but more expensive. This can be easily addressed because the computation of our training method can be distributed on GPUs for example, making it even more scalable. We will investigate both problems in future work.

Acknowledgements

The authors would like to thank Nilesh Tripuraneni, Alex Matthews, Jes Frellsen and Carl Rasmussen for insightful comments and discussion. TDB thanks Google for funding his European Doctoral Fellowship. JMHL acknowledges support from the Rafael del Pino Foundation. DHL and JMHL acknowledge support from Plan Nacional I+D+i, Grant TIN2013-42351-P, and from Comunidad de Madrid, Grant S2013/ICE-2845 CASI-CAM-CM. YL thanks the Schlumberger Foundation for her Faculty for the Future PhD fellowship. RET thanks EPSRC grants EP/G050821/1 and EP/L000776/1.

References

Barber, D. and Schottky, B. Radial basis functions: a Bayesian treatment. In *Advances in Neural Information Processing Systems 10*, 1998.

Bui, Thang D., Yan, Josiah, and Turner, Richard E. A unifying framework for sparse Gaussian process approximation using Power Expectation Propagation. *arXiv preprint arXiv:1605.07066*, 2016.

Dai, Zhenwen, Damianou, Andreas, González, Javier, and Lawrence, Neil. Variational auto-encoded deep Gaussian processes. In *4th International Conference on Learning Representations*, 2016.

Damianou, Andreas. *Deep Gaussian processes and variational propagation of uncertainty*. PhD thesis, University of Sheffield, 2015.

Damianou, Andreas C and Lawrence, Neil D. Deep Gaussian processes. In *16th International Conference on Artificial Intelligence and Statistics*, pp. 207–215, 2013.

Deisenroth, Marc and Mohamed, Shakir. Expectation propagation in Gaussian process dynamical systems. In *Advances in Neural Information Processing Systems 25*, pp. 2609–2617, 2012.

Demšar, Janez. Statistical comparisons of classifiers over multiple data sets. *The Journal of Machine Learning Research*, 7:1–30, 2006.

Duvenaud, David, Rippel, Oren, Adams, Ryan P., and Ghahramani, Zoubin. Avoiding pathologies in very deep networks. In *17th International Conference on Artificial Intelligence and Statistics*, 2014.

Gal, Yarin and Ghahramani, Zoubin. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *33rd International Conference on Machine Learning*, 2016.

Girard, Agathe, Rasmussen, Carl Edward, Quiñonero-Candela, Joaquin, and Murray-Smith, Roderick. Gaussian process priors with uncertain inputs — application to multiple-step ahead time series forecasting. In *Advances in Neural Information Processing Systems 15*, pp. 529–536, 2003.

Graves, Alex. Practical variational inference for neural networks. In *Advances in Neural Information Processing Systems 25*, pp. 2348–2356, 2011.

Hachmann, Johannes, Olivares-Amaya, Roberto, Atahan-Evrenk, Sule, Amador-Bedolla, Carlos, Sánchez-Carrera, Roel S, Gold-Parker, Aryeh, Vogt, Leslie, Brockway, Anna M, and Aspuru-Guzik, Alán. The Harvard clean energy project: large-scale computational screening and design of organic photovoltaics on the world community grid. *The Journal of Physical Chemistry Letters*, 2(17):2241–2251, 2011.

Hensman, James and Lawrence, Neil D. Nested variational compression in deep Gaussian processes. *arXiv preprint arXiv:1412.1370, visited on 25/05/2016*, 2014.

Hernández-Lobato, José Miguel and Adams, Ryan P. Probabilistic backpropagation for scalable learning of Bayesian neural networks. In *32nd International Conference on Machine Learning*, 2015.

Kingma, D. P. and Ba, J. Adam: a method for stochastic optimization. In *3rd International Conference on Learning Representations*, 2015.

Kingma, Diederik P. and Welling, Max. Stochastic gradient VB and the variational auto-encoder. In *The International Conference on Learning Representations*, 2014.

Kushner, H. J. and Budhiraja, A. S. A nonlinear filtering algorithm based on an approximation of the conditional distribution. *IEEE Transactions on Automatic Control*, 45(3):580–585, Mar 2000.

Lawrence, Neil D. and Moore, Andrew J. Hierarchical Gaussian process latent variable models. In *24th International Conference on Machine Learning*, pp. 481–488, 2007.

Lázaro-Gredilla, Miguel. Bayesian warped Gaussian processes. In *Advances in Neural Information Processing Systems 25*, pp. 1619–1627, 2012.

Li, Yingzhen, Hernández-Lobato, José Miguel, and Turner, Richard E. Stochastic expectation propagation. In *Advances in Neural Information Processing Systems 29*, 2015.

Minka, Thomas P. *A family of algorithms for approximate Bayesian inference*. PhD thesis, Massachusetts Institute of Technology, 2001.

Neal, Radford M. Bayesian learning via stochastic dynamics. In *Advances in Neural Information Processing Systems 6*, pp. 475–482, 1993.

Neal, Radford M. *Bayesian learning for neural networks*. PhD thesis, University of Toronto, 1995.

Pyzer-Knapp, Edward O, Li, Kewei, and Aspuru-Guzik, Alan. Learning from the Harvard clean energy project: The use of neural networks to accelerate materials discovery. *Advanced Functional Materials*, 25(41):6495–6502, 2015.

Quiñonero-Candela, Joaquin and Rasmussen, Carl Edward. A unifying view of sparse approximate Gaussian process regression. *The Journal of Machine Learning Research*, 6:1939–1959, 2005.

Rasmussen, Carl Edward and Williams, Christopher K. I. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2005.

Seeger, Matthias. Expectation propagation for exponential families. Technical report, Department of EECS, University of California at Berkeley, 2007.

Snelson, Edward and Ghahramani, Zoubin. Sparse Gaussian processes using pseudo-inputs. In *Advances in Neural Information Processing Systems 19*, pp. 1257–1264, 2006.

Snelson, Edward, Rasmussen, Carl Edward, and Ghahramani, Zoubin. Warped Gaussian processes. In *Advances in Neural Information Processing Systems 17*, pp. 337–344, Cambridge, MA, USA, 2004.

Sutton, Richard S. and Barto, Andrew G. *Introduction to Reinforcement Learning*. MIT Press, Cambridge, MA, USA, 1st edition, 1998. ISBN 0262193981.

Titsias, Michalis K. Variational learning of inducing variables in sparse Gaussian processes. In *12th International Conference on Artificial Intelligence and Statistics*, pp. 567–574, 2009.

Titsias, Michalis K and Lawrence, Neil D. Bayesian Gaussian process latent variable model. In *13th International Conference on Artificial Intelligence and Statistics*, pp. 844–851, 2010.

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

Vanhatalo, Jarno and Vehtari, Aki. MCMC methods for MLP-network and Gaussian process and stuff–a documentation for Matlab toolbox MCMCstuff. 2006. Laboratory of computational engineering, Helsinki University of Technology.

Welling, Max and Teh, Yee W. Bayesian learning via stochastic gradient Langevin dynamics. In *28th International Conference on Machine Learning*, pp. 681–688, 2011.

Wilson, Andrew and Adams, Ryan. Gaussian process kernels for pattern discovery and extrapolation. In *30th International Conference on Machine Learning*, pp. 1067–1075, 2013.