
Structured Variational Inference for Coupled Gaussian Processes

Adam Vincent

Gatsby Computational Neuroscience Unit
University College London
London, W1T 4JG
vincenta@gatsby.ucl.ac.uk

Abstract

Sparse variational approximations allow for principled and scalable inference in Gaussian Process (GP) models. In settings where several GPs are part of the generative model, these GPs are a posteriori coupled. For many applications such as regression where predictive accuracy is the quantity of interest, this coupling is not crucial. However if one is interested in posterior uncertainty, it cannot be ignored. A key element of variational inference schemes is the choice of the approximate posterior parameterization. When the number of latent variables is large, mean field (MF) methods provide fast and accurate posterior means while more structured posterior lead to inference algorithm of greater computational complexity. Here, we extend previous sparse GP approximations and propose a novel parameterization of variational posteriors in the multi-GPs setting allowing for fast and scalable inference capturing posterior dependencies.

1 Introduction

Gaussian Processes (GPs) are powerful nonparametric distributions over continuous functions which can be used for both supervised and unsupervised learning problems [1]. Here, we propose an extension of the variational pseudo-point GP (VFE) approximation [2, 3] to the case of regression models composed of multiple GPs. Such models include Generalized Additive Models (GAM) which are a commonly used non-linear extension of generalized linear models when interpretability of models is required [4]. In these models, inference leads to posterior dependencies across GPs. For this reason we will refer to this setting as Coupled GPs (CGPs).

The VFE approximation provides state-of-the art performance for GP regression and provides approximations to the posterior distributions in the form of a GP. This approach has been successfully extended to the CGPs setting using a factorized (mean-field) approximation of the posterior across GPs [5, 6]. However, it suffers from the known variance underestimation of mean-field approximations and therefore can lead to poor predictions or can bias learning [7]. Furthermore the mean-field approaches provide no information about how the uncertainties across GPs interact.

Several general approaches have been proposed to alleviate this problem of mean field approximations [8, 9, 10], however, they do not readily apply to the CGPs setting or cannot leverage the sparse approximation.

This paper introduces a new posterior approximation tailored to the CGPs setting that explicitly captures posterior dependencies and allows for scalable inference. We now briefly review the VFE approach to GP regression and develop the proposed Variationally Coupled GPs (VCGPs) approach.

2 Regression with multiple GPs

We are interested in the broad class of regression models with predictors of the form $\rho(x) = \Phi(f_1(x), \dots, f_C(x))$ where f_c are functions from $\mathcal{X}_c \rightarrow \mathbb{R}$. We consider models with factorized likelihood of the form $p(y|X) = \prod_i p(y_i|\rho(x_i))$. The specific form of the likelihood is arbitrary. Φ is an arbitrary function from $\mathbb{R}^C \rightarrow \mathbb{R}$ and is assumed to be known. This model class generalizes the classical GP regression setting as well as GLMs and GAMs.

We are interested in inferring the functions $f_{1..C}$ from data. Without further assumptions, this is an ill-posed problem. Here, we are interested in the case where independent GPs are used as priors over the different latent functions and we wish to perform Bayesian inference. For most choices of kernels, likelihood and function Φ , latents are coupled a posteriori.

We denote $\mathcal{F} = f_1, \dots, f_C$ such that $p(\mathcal{F})$ constitutes the joint distribution over the processes. $\mathcal{F}(x) = [f_1(x), \dots, f_C(x)]$ is the vector of function evaluations at x . We are interested in computing the joint posterior $p(\mathcal{F}|X, Y)$.

3 Pseudo-point approximations in multiple GPs regression

3.1 Variational lower bound to the log evidence

The classical variational lower bound is

$$\log p(y|X) \geq E_{q(\mathcal{F})} \log \frac{p(y|\mathcal{F}, X) p(\mathcal{F})}{q(\mathcal{F})} = \mathcal{L}(q) \quad (1)$$

where q is the variational posterior approximation.

Following [5, 6] we introduce for each GP indexed c some 'inducing points' $Z_c = [z_c^{(1)}, \dots, z_c^{(m)}] \in \mathcal{X}_c^m$. We note \mathbf{U}_c the vector of associated function evaluations $\mathbf{U}_c = [u_c^{(1)}, \dots, u_c^{(m)}] = [f_c(z_c^{(1)}), \dots, f_c(z_c^{(m)})]$. We also note $\mathbf{U} = [\mathbf{U}_1, \dots, \mathbf{U}_c]$ the stacked vector.

3.2 Parameterization of the coupled variational posterior

We then choose the following form for our posterior approximation $q(\mathcal{F}) = q(\mathbf{U}) \prod_c p(f_{c-\mathbf{U}_c}|\mathbf{U}_c)$. This intuitively means that the posterior is compressed into the lower dimensional $q(\mathbf{U})$ since the full posterior $q(\mathcal{F})$ is built for each latent by interpolating from $q(\mathbf{U})$ using the prior only.

This choice leads to a simplification of the lower bound (1) as

$$\mathcal{L}(q) = E_{q(\mathcal{F})} \log \frac{p(y|\mathcal{F}, X) p(\mathbf{U}) \prod_c p(f_{c-\mathbf{U}_c}|\mathbf{U}_c)}{q(\mathbf{U}) \prod_c p(f_{c-\mathbf{U}_c}|\mathbf{U}_c)} \quad (2)$$

$$= E_{q(\mathcal{F})} \log p(y|\mathcal{F}, X) - KL[q(\mathbf{U}) || p(\mathbf{U})] \quad (3)$$

where KL denotes the Kullback-Leibler divergence between distributions.

We also restrict ourselves to the case of factorizing likelihoods which simplifies the expectations terms in the bound

$$\mathcal{L}(q) = \sum_i E_{q(\mathcal{F}^{(i)})} \log p(y^{(i)}|\mathcal{F}^{(i)}, x^{(i)}) - KL[q(\mathbf{U}) || p(\mathbf{U})] \quad (4)$$

Saul et al [5] considered the mean field case $q(\mathbf{U}) = \prod_c q(\mathbf{U}_c)$ with each factor parameterized as a multivariate Gaussian distribution $N(\mu_{\mathbf{U}_c}, \Sigma_{\mathbf{U}_c})$. This approach does not capture posterior dependencies across GPs. Instead, we here directly parameterize $q(\mathbf{U})$ a multivariate Normal distribution $N(\mu_{\mathbf{U}}, \Sigma_{\mathbf{U}})$ and show how it only induces small additional complexity and how methods to train VFE for single GPs readily apply to our setting [11].

4 Optimizing the variational lower bound

The Kullback divergence term in equation (4) is between two multivariate Gaussian distributions. Its expression and derivatives are available in closed form. The expectation terms in equation (4) are intractable in most cases and need to be approximated.

Table 1: Results - Conjugate Additive Regression

Model	RMSE	$\sqrt{\mathbb{E}_{\mathcal{D}}[\sigma_{post}^2(\sum f)]}$	$\mathbb{E}_{\mathcal{D}}[\rho_{post}(f_1, f_2)]$	$\sqrt{\mathbb{E}_{\mathcal{D}}[\sigma_{post}^2(\Delta f)]}$	$-\mathcal{L}(q)$
Exact	0.4675	0.0967	-0.9306	0.509	381.924
VCGP[10]	0.4698	0.1237	-0.892	0.5109	388.667
VCGP[30]	0.4678	0.1036	-0.9265	0.5286	383.714
MF[10]	0.469	0.1331	0	0.1331	398.376
MF[30]	0.468	0.125	0	0.125	409.01

4.1 Stochastic optimization

We approximate the expectation terms in (4) with Monte Carlo samples from the variational posterior for each data point. For each expectation $E_{q(v)}f(v)$ under $v \sim \mathcal{N}(\mu, \Sigma = LL^T)$, we reparameterize the Gaussian distribution as $v = L\epsilon + \mu$ with $\epsilon \sim \mathcal{N}(0, I)$. Using a sampling procedure in terms of isotropic Gaussians, we can compute unbiased gradients of the bound with respect to the variational parameters [12]. Since the bound is expressed as a sum of terms depending on each data-point separately, we can scale the methods to large datasets by sub-sampling the data.

4.2 Complexity

Given C latents, with M inducing points each, this leads to a vector \mathbf{U} of size MC and hence a number of variational parameters that is $O(M^2C^2)$. The evaluation of the lower bound requires the computation of the joint $q(\mathcal{F}(x)) = \int d\mathbf{U} p(\mathcal{F}(x)|\mathbf{U})q(\mathbf{U})$ for each data point which is $O(C^4M^2)$, leading to a cost of $O(NC^4M^2)$ that is linear in N . Finally, computing the KL divergence of the bound requires $O(C^3M^3)$.

5 Experiment: Conjugate Additive Regression

We demonstrate the performance of our VCGP approach in a preliminary experiment. We show that we can achieve tighter bounds to the marginal evidence in the conjugate additive setting where both the marginal evidence and the full coupled posterior are available in closed form. This experiment was implemented using GPflow [13]

5.1 Setting

As a toy example, we consider the case of GP additive regression [6] with Gaussian observation noise and $N = 500$ data points. Ground truth functions are set to $f_1(x) = \sin(x)^3$ and $f_2(x) = \cos(3x)$. Covariates are sampled as $x_1, x_2 \sim \text{Uniform}[-3, 3]$. The observed output is $y^i = f_1(x_1^i) + f_2(x_2^i) + e^i$ with $e^i \sim \mathcal{N}(0, \sigma^2)$ and $\sigma = 0.5$. Priors for both latents are set to be GPs with 0-mean and Gaussian kernels $k_c(x, x') = s_c^2 \exp\left[-\frac{(x-x')^2}{2l_c^2}\right]$. Parameters $\sigma, s_1, s_2, l_1, l_2$ are set to maximize the marginal evidence $p(y|x)$ that is available in closed form, and held fixed throughout the optimization of the MF and VCGP models. Variational inference is then performed in these two models. In both cases, each latent is approximated using 10 and 30 inducing points whose position is optimized.

5.2 Results

All results are summarized in Table 1. Both the exact, MF and SVGP models perform equally well in terms of predictive accuracy (RMSE) with a small tendency to overestimate the posterior predictive uncertainty on the sum of the latents ($\Sigma f = f_1 + f_2$).

However, only the SVGP model can provide accurate marginal variance predictions for each latent as is shown in Figure 1(left)(bottom-right). It is also the only model to accurately predict the posterior correlation across latents and can thus correctly predict the posterior uncertainty on the difference between the latents ($\Delta f = f_1 - f_2$). A histogram of these posterior correlations $\rho(f_1, f_2)$ for the SVGP model is shown in Figure 1(top-right).

Furthermore, our model provides a tighter bound to the log-evidence (written $\mathcal{L}(q)$ for the exact model in Table 1 for simplicity). This means that our coupled variational posterior q_{coupled} is closer to the true posterior p than MF posterior q_{MF} in the sense that $KL(q_{\text{coupled}}|p) < KL(q_{\text{MF}}|p)$.

Similarly to the standard VFE approach to GP regression [3], the performance of our VCGP model improves as we increase the number of inducing points.




Figure 1: Output of inference for exact, VCGP[30] and MF[30] model. *Left*: Dashed line are the true latents, tubes represent the marginal posterior GP for each individual source. For the variational algorithms, vertical bars show the positions of the pseudo inputs associated to the inducing points. *Right, Top*: Histogram of posterior correlation $\rho(f_1, f_2) = \frac{\text{cov}(f_1, f_2)}{\sqrt{\text{var}(f_1)\text{var}(f_2)}}$ across the two latents for the exact and the VCGP model. *Right, Bottom*: Mean and variance of the marginal posterior standard deviations for all models.

6 Related work

Variational inference for the CGP setting has so far only used the mean-field approximation as described in [5]. When posterior dependencies are a quantity of interest, a natural approach is to increase the complexity of the variational posterior to capture these dependencies. This often results in a prohibitive increase in the complexity of the inference.

Different solutions have been proposed to tackle this problem. A first approach in [8] consists in a two step scheme where MF inference is *assumed* to provide accurate posterior mean estimates. A perturbation analysis is then performed around the MF posterior means to provide second order (covariance) estimates.

A second approach consists in ‘relaxing’ the MF approximation by extending the variational posterior q with additional multiplicative terms capturing dependencies while keeping the computational complexity of the resulting inference scheme low [9, 10].

Our approach fits in this second family of extensions of the the MF parameterization. It is tailored to the VFE approximation to GP models and leverages its sparsity to provide a fast and scalable inference algorithm.

7 Conclusion

We presented an extension of the VFE approximation to regression settings with multiple GPs that takes advantage of its scalability and generality. Our approach goes beyond the classical mean-field approximation usually used and provides richer posteriors capturing dependencies between inferred latent functions. This approach is very general and provides an efficient way to perform inference in a class of models broadly used in statistics such that of GAMs.

Acknowledgments

I would like to thank Lea Duncker, Eszter Vertes and Maneesh Sahani from the Gatsby Unit for helpful discussions.

References

- [1] Carl Edward Rasmussen and Christopher KI Williams. *Gaussian processes for machine learning*, volume 1. MIT press Cambridge, 2006.
- [2] Michalis K Titsias. Variational learning of inducing variables in sparse gaussian processes. In *International Conference on Artificial Intelligence and Statistics*, pages 567–574, 2009.
- [3] 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.
- [4] Trevor Hastie and Robert Tibshirani. *Generalized additive models*. Wiley Online Library, 1990.
- [5] Alan D Saul, James Hensman, Aki Vehtari, and Neil D Lawrence. Chained gaussian processes. In *Artificial Intelligence and Statistics*, pages 1431–1440, 2016.
- [6] Vincent Adam, James Hensman, and Maneesh Sahani. Scalable transformed additive signal decomposition by non-conjugate gaussian process inference. In *Machine Learning for Signal Processing (MLSP), 2016 IEEE 26th International Workshop on*, pages 1–6. IEEE, 2016.
- [7] Richard E Turner and Maneesh Sahani. Two problems with variational expectation maximisation for time-series models. *Bayesian Time series models*, pages 115–138, 2011.
- [8] Ryan J Giordano, Tamara Broderick, and Michael I Jordan. Linear response methods for accurate covariance estimates from mean field variational bayes. In *Advances in Neural Information Processing Systems*, pages 1441–1449, 2015.
- [9] Dustin Tran, David Blei, and Edo M Airolidi. Copula variational inference. In *Advances in Neural Information Processing Systems*, pages 3564–3572, 2015.
- [10] Matthew Hoffman and David Blei. Stochastic structured variational inference. In *Artificial Intelligence and Statistics*, pages 361–369, 2015.
- [11] James Hensman, Alexander G de G Matthews, and Zoubin Ghahramani. Scalable variational Gaussian process classification. In *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, 2015.
- [12] Hugh Salimbeni and Marc Deisenroth. Doubly stochastic variational inference for deep gaussian processes. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 4591–4602. Curran Associates, Inc., 2017.
- [13] Alexander G. de G. Matthews, Mark van der Wilk, Tom Nickson, Keisuke. Fujii, Alexis Boukouvalas, Pablo León-Villagr a, Zoubin Ghahramani, and James Hensman. GPflow: A Gaussian process library using TensorFlow. *Journal of Machine Learning Research*, 18(40):1–6, apr 2017.