

Variational Fourier Features for Gaussian Processes

James Hensman*
Nicolas Durrande†
PROWLER.io
66-68 Hills Road
Cambridge, CB2 1LA, UK

JAMES@PROWLER.IO
NICOLAS@PROWLER.IO

Arno Solin
Aalto University
FI-00076 AALTO
Espoo, Finland

ARNO.SOLIN@AALTO.FI

Editor: Manfred Opper

Abstract

This work brings together two powerful concepts in Gaussian processes: the variational approach to sparse approximation and the spectral representation of Gaussian processes. This gives rise to an approximation that inherits the benefits of the variational approach but with the representational power and computational scalability of spectral representations. The work hinges on a key result that there exist spectral features related to a finite domain of the Gaussian process which exhibit almost-independent covariances. We derive these expressions for Matérn kernels in one dimension, and generalize to more dimensions using kernels with specific structures. Under the assumption of additive Gaussian noise, our method requires only a single pass through the data set, making for very fast and accurate computation. We fit a model to 4 million training points in just a few minutes on a standard laptop. With non-conjugate likelihoods, our MCMC scheme reduces the cost of computation from $\mathcal{O}(NM^2)$ (for a sparse Gaussian process) to $\mathcal{O}(NM)$ per iteration, where N is the number of data and M is the number of features.

Keywords: Gaussian processes, Fourier features, variational inference

1. Introduction

Efficient computation in Gaussian process (GP) models is of broad interest across machine learning and statistics, with applications in the fields of spatial epidemiology (Diggle, 2013; Banerjee et al., 2014), robotics and control (Deisenroth et al., 2015), signal processing (e.g. Särkkä et al., 2013), Bayesian optimization and probabilistic numerics, (e.g. Osborne, 2010; Briol et al., 2016; Hennig et al., 2015), and many others.

Computation is challenging for several reasons. First, the computational complexity usually scales cubically with the number of data N , as one must decompose a $N \times N$ dense covariance matrix. Second, the posterior is both high-dimensional and non-Gaussian if the

*. This work was partially completed whilst JH was affiliated to the Centre for Health Information, Computation, and Statistics (CHICAS), Faculty of Health and Medicine, Lancaster University, UK.

†. This work was partially completed whilst ND was affiliated to Institut Fayol-LIMOS, Mines Saint-Étienne, 158 cours Fauriel, Saint-Étienne, France.

data likelihood is not conjugate. Third, GP models have a hierarchical structure, with the latent values of the function being conditioned on parameters of the covariance function.

In this work, we adopt a variational Bayesian framework for dealing with these issues. By this we mean that we minimize the Kullback-Leibler (KL) divergence $\text{KL}[q||p]$, where q denotes an approximation to the posterior, and p is the posterior itself. Many authors have applied this framework to GP problems before: perhaps the earliest example is Csató et al. (2002). Seeger (2003) also made use of this KL criterion in the context of Gaussian processes. Particularly influential works include Titsias (2009), who developed the first sparse GP using the KL divergence, and Opper and Archambeau (2009), who made a case for this approximation in the non-conjugate (but not sparse) case.

The variational framework for Gaussian processes has several advantages. First, the three challenges mentioned above can be tackled by a single unified objective. Second, the accuracy of the approximation can easily be shown to increase monotonically with increasing complexity of the proposed approximating family (i.e. more inducing points is always better). Third, the framework provides an evidence lower bound (or ELBO), which can be evaluated to compare different approximations.

In this work, we combine the variational approach with Fourier features; we refer to our method as Variational Fourier Features (VFF). Whilst most sparse Gaussian process approximations rely on *inducing* or *pseudo* inputs, which lie in the same domain as the inputs to the GP, Fourier features lie in the spectral domain of the process. As inducing-point approximations build up the posterior through kernel functions, Fourier features build up the approximate posterior through sinusoids. Our approach differs from existing works which use random frequencies (Rahimi and Recht, 2008, 2009) or optimized frequencies (Lázaro-Gredilla et al., 2010), in that we use frequencies placed on a regular grid (cf. Solin and Särkkä, 2014), which is key in our derivation of a computationally convenient matrix structure.

Naive combination of Fourier features with sparse GP approaches is not feasible because the standard Fourier transform of a stochastic process does not converge, meaning that these features are associated with random variables of infinite variance. Matthews et al. (2016) have shown that valid inducing features for the variational approach require finite, deterministic projections of the process. To combat this, Figueiras-Vidal and Lázaro-Gredilla (2009) used a windowing technique, which made the integrals converge to valid inducing features (for the squared-exponential kernel), but lost perhaps the most important aspect of a Fourier approach: that the features should be independent. In this work, we combine a finite window with a Reproducing Kernel Hilbert Space (RKHS) projection to construct features that are *almost* independent, having covariance structures that are easily decomposed. This gives rise to a complexity of $\mathcal{O}(NM)$, comparing favourably with the $\mathcal{O}(NM^2)$ complexity of other variational approaches, where M denotes the number of features.

The contributions of this paper are the following:

- We introduce a novel approximation scheme for representing a GP using a Fourier decomposition. We limit our interest to the widely used Matérn family of stationary kernels, for which our proposed decomposition exhibits an almost-independent structure.

- We combine our approximation scheme with the variational Bayesian framework for Gaussian processes, and present how the methodology generalizes to GP inference problems with general likelihoods.

Consequently, this paper aims to provide an approximate Gaussian process inference scheme which is theoretically sound, has strong representative power, is extremely fast, and—most importantly—works well in practice.

The rest of this paper is structured as follows. Section 2 reviews the relevant background related to Gaussian process inference and previous work in sparse and spectral approximations to GP models. In Section 3 the core methodology for Variational Fourier Features is presented for the one-dimensional case. In Section 4 this is extended to multidimensional cases with additive and product structures. Implementation details and computational complexity of the method are covered in Section 5. Section 6 is dedicated to a set of illustrative toy examples and a number of empirical experiments, where practical aspects of Variational Fourier Feature inference are demonstrated. Finally, we conclude the paper with a discussion in Section 7.

2. Background

In this section, we first set up the family of Gaussian process models that we will consider, which allows us to establish some notation. Next, we review some basic results regarding the spectrum of stationary Gaussian processes, and recall how Fourier features approximate the kernel. We relate sparse Gaussian process approximations and Fourier approximations by explicating them as alternative models (Quiñonero-Candela and Rasmussen, 2005). We then recap the variational approximation to Gaussian processes, including expressions for sparse approximations and approximations for non-conjugate likelihoods. The two final subsections in this section discuss decomposition of Gaussian processes: we first link decomposition and conditioning and then discuss inter-domain decomposition.

2.1 Gaussian process models

Gaussian processes are distributions over functions, defined by a mean function and a covariance function (see Williams and Rasmussen, 2006, for an introduction). In this section we consider GPs over the real line, $x \in \mathbb{R}$. Making the standard assumption that the mean function is zero, we write

$$f(x) \sim \mathcal{GP}(0, k(x, x')) . \tag{1}$$

The data $\mathbf{y} = [y_n]_{n=1}^N$ at locations $\mathbf{X} = [x_n]_{n=1}^N$ are conditioned on the function evaluations $\mathbf{f} = [f(x_n)]_{n=1}^N$ through some factorising likelihood

$$p(\mathbf{y} | f(x)) = p(\mathbf{y} | \mathbf{f}) = \prod_n p(y_n | f(x_n)) , \tag{2}$$

which we do not in general assume to be Gaussian. A defining property of Gaussian processes is that any finite number of function evaluations follow a multivariate normal distribution, so

$$p(\mathbf{f}) = \mathcal{N}(\mathbf{0}, \mathbf{K}_{\mathbf{ff}}) , \tag{3}$$

and the process conditioned on these values is given by a conditional Gaussian process:

$$f(x) | \mathbf{f} \sim \mathcal{GP} \left(\mathbf{k}_{\mathbf{f}}(x)^\top \mathbf{K}_{\mathbf{ff}}^{-1} \mathbf{f}, k(x, x') - \mathbf{k}_{\mathbf{f}}(x)^\top \mathbf{K}_{\mathbf{ff}}^{-1} \mathbf{k}_{\mathbf{f}}(x') \right). \quad (4)$$

To compute the posterior over functions given the data $f(x) | \mathbf{y}$, it is possible to compute the posterior only for the finite set $p(\mathbf{f} | \mathbf{y})$ using standard methodologies like Markov Chain Monte Carlo (MCMC) (e.g. Murray et al., 2010; Filippone et al., 2013) or variational Bayes (e.g. Opper and Archambeau, 2009; Khan et al., 2012, 2013). Evaluating the function at a test point means averaging equation (4) over this posterior $p(\mathbf{f} | \mathbf{y})$.

In our notation, the matrix $\mathbf{K}_{\mathbf{ff}}$ is given by evaluating the covariance function at all pairs of data points $\mathbf{K}_{\mathbf{ff}}[i, j] = k(x_i, x_j)$. The vector valued function $\mathbf{k}_{\mathbf{f}}(x)$ is given similarly, $\mathbf{k}_{\mathbf{f}}(x) = [k(x_1, x), k(x_2, x), \dots, k(x_n, x)]^\top$. We omit dependence on covariance function parameters for clarity.

2.2 Spectral representations of stationary covariances

Stationary Gaussian processes are those with a covariance function that can be written as a function of the distance between observations, $k(x, x') = k(|x - x'|) = k(r)$. Of particular interest in this work will be the Matérn family with half-integer orders, the first three of which are

$$\begin{aligned} k_{\frac{1}{2}}(r) &= \sigma^2 \exp(-r/\ell), \\ k_{\frac{3}{2}}(r) &= \sigma^2 (1 + \sqrt{3}r/\ell) \exp(-\sqrt{3}r/\ell), \\ k_{\frac{5}{2}}(r) &= \sigma^2 (1 + \sqrt{5}r/\ell + \frac{5}{3}r^2/\ell^2) \exp(-\sqrt{5}r/\ell), \end{aligned} \quad (5)$$

where σ^2 and ℓ are variance and lengthscale parameters respectively.

Bochner's theorem (Akhiezer and Glazman, 1993) tells us that any continuous positive definite function, such as a covariance function, can be represented as the Fourier transform of a positive measure. If the measure has a density, it is known as the spectral density $s(\omega)$ of the covariance function. This gives rise to the Fourier duality of spectral densities and covariance functions, known as the Wiener-Khintchin theorem. It gives the following relations:

$$s(\omega) = \mathcal{F}\{k(r)\} = \int_{-\infty}^{\infty} k(r) e^{-i\omega r} dr, \quad (6)$$

$$k(r) = \mathcal{F}^{-1}\{s(\omega)\} = \frac{1}{2\pi} \int_{-\infty}^{\infty} s(\omega) e^{i\omega r} d\omega. \quad (7)$$

Since kernels are symmetric real functions, the spectrum of the process is also a symmetric real function. The spectra corresponding to the Matérn covariances are

$$s_{\frac{1}{2}}(\omega) = 2\sigma^2 \lambda (\lambda^2 + \omega^2)^{-1}, \quad \lambda = \ell^{-1}, \quad (8)$$

$$s_{\frac{3}{2}}(\omega) = 4\sigma^2 \lambda^3 (\lambda^2 + \omega^2)^{-2}, \quad \lambda = \sqrt{3}\ell^{-1}, \quad (9)$$

$$s_{\frac{5}{2}}(\omega) = \frac{16}{3}\sigma^2 \lambda^5 (\lambda^2 + \omega^2)^{-3}, \quad \lambda = \sqrt{5}\ell^{-1}. \quad (10)$$

2.3 Model approximations

In order to overcome the cubic computational scaling of Gaussian process methods, many approximations have been proposed. Here we briefly review two types of approximations, those based on the spectrum of the process and others referred to as ‘sparse’ approximations.

2.3.1 RANDOM FOURIER FEATURES

Random Fourier Features (RFF) is a method for approximating kernels. The essential element of the RFF approach (Rahimi and Recht, 2008, 2009) is the realization that the Wiener-Khintchin integral (7) can be approximated by a Monte Carlo sum

$$k(r) \approx \tilde{k}(r) = \frac{\sigma^2}{M} \sum_{m=1}^M \cos(\omega_m r), \quad (11)$$

where the frequencies ω_m are drawn randomly from the distribution proportional to $s(\omega)$. Note that the imaginary part is zero because $s(\omega r)$ is even and $i \sin(\omega r)$ is odd.

This approximate kernel has a finite basis function expansion as

$$\phi(x) = [\cos(\omega_1 x), \cos(\omega_2 x), \dots, \cos(\omega_M x), \sin(\omega_1 x), \sin(\omega_2 x), \dots, \sin(\omega_M x)]^\top, \quad (12)$$

and we can write the approximate Gaussian process as

$$f(x) \sim \mathcal{GP} \left(0, \tilde{k}(x - x') \right) \quad \text{or} \quad f(x) \sim \mathcal{GP} \left(0, \frac{\sigma^2}{M} \phi(x)^\top \phi(x') \right) \quad (13)$$

or equivalently as a parametric model, in the form

$$f(x) = \phi(x)^\top \mathbf{w}, \quad (14)$$

$$\mathbf{w} \sim \mathcal{N} \left(\mathbf{0}, \frac{\sigma^2}{M} \mathbf{I} \right). \quad (15)$$

2.3.2 OPTIMIZED FOURIER FEATURES

In order to have a representative sample of the spectrum, the RFF methodology typically requires the number of spectral sample points to be large. In Lázaro-Gredilla et al. (2010), this was addressed by the attractive choice of optimizing the spectral locations along with the hyperparameters (the ‘Sparse Spectrum GP’). However, as argued by those authors, this option does not converge to the full GP and can suffer from overfitting to the training data. We confirm this empirically in a simple experiment in Section 6.

Noting that this method was prone to overfitting, Gal and Turner (2015) sought to improve the model by integrating out, rather than optimizing the frequencies. Gal and Turner derived a variational approximation that made use of a tractable integral over the frequency space. The result is an algorithm that suffers less overfitting than the Sparse Spectrum GP, yet remains flexible. We emphasize the difference to the approach in this work: Gal and Turner proposed variational inference in a sparse spectrum model that is derived from a GP model; our work aims to directly approximate the posterior of the true models using a variational representation.

2.3.3 REGULAR FOURIER FEATURES

Another way to approximate the integral in (7) is with a regularly spaced set of frequencies ω . We refer to this method as *regular* Fourier features. While the Random Fourier Features approach could be interpreted as a Monte Carlo approximation to the Wiener-Khinchin integral, the regular Fourier features approach can be seen as a quadrature approximation of the same integral. Therefore the regular methods are not *sparse* in spectral sense, but *dense* and deterministic given the domain of interest (Solin and Särkkä, 2014).

The approximation to the original covariance function takes the form of a finite sum

$$k(r) \approx \tilde{k}(r) = \sum_m s(\omega_m) \cos(\omega_m r) \quad \omega_m = m\Delta, \quad (16)$$

where Δ is a small number which defines the grid spacing. Similarly to Random Fourier features, the resulting approximate Gaussian process can be written as a parametric model:

$$f(x) = \boldsymbol{\phi}(x)^\top \mathbf{w}, \quad (17)$$

$$\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{S}) \quad \mathbf{S} = \text{diag}(s(\omega_1), \dots, s(\omega_M), s(\omega_1), \dots, s(\omega_M)). \quad (18)$$

Like in the Random Fourier Features approximation (15), regular Fourier features are independent (the matrix \mathbf{S} is diagonal), and if the inputs are spaced on a regular grid, fast Fourier transform (FFT) methods can speed up the computations to $\mathcal{O}(N \log N)$ (see e.g. Paciorek, 2007; Fritz et al., 2009).

2.3.4 SPARSE GAUSSIAN PROCESS MODELS

In the above we reviewed how spectral approximations result in parametric models. In their review paper, Quiñero-Candela and Rasmussen (2005) also presented a view of sparse Gaussian processes where the result is an approximate, parametric model.

The terminology ‘sparse’ originally referred to methods which restricted computation to a subset of the data points (e.g. Csató and Opper, 2002; Williams and Seeger, 2001), until Snelson and Ghahramani (2005) relaxed this assumption with the idea of pseudo inputs. The pseudo-inputs (or ‘inducing points’) $\mathbf{Z} = [z_m]_{m=1}^M$ lie in the same domain as \mathbf{X} , but with $M < N$. The values of the process associated with these points are denoted $\mathbf{u} = [f(z_m)]_{m=1}^M$. The simplest form of sparse Gaussian process model using these variables is the Deterministic Training Conditional (DTC) model, written

$$f(x) = \mathbf{k}_u(x)^\top \mathbf{K}_{uu}^{-1} \mathbf{u}, \quad (19)$$

$$\mathbf{u} \sim \mathcal{N}(\mathbf{0}, \mathbf{K}_{uu}). \quad (20)$$

This kind of ‘projected process’ approximation has also been discussed by e.g. Banerjee et al. (2008). The literature contains many parametric models that approximate Gaussian process behaviours; for example Bui and Turner (2014) included tree-structures in the approximation for extra scalability, and Moore and Russell (2015) combined local Gaussian processes with Gaussian random fields.

2.4 Variational Gaussian process approximations

We have seen how Fourier approximations and sparse Gaussian processes can be written as approximate parametric models. The variational approximation to Gaussian processes provides a more elegant, flexible and extensible solution in that the *posterior distribution* of the original model is approximated, rather than the model itself. In existing work, the variational approach has been used alongside ideas from the sparse GP literature: the remit of this work is to combine the variational methodology with Fourier based approximations. We provide a short review of the variational methodology here, for a more detailed discussion see Matthews (2016).

In variational Bayes (see e.g. Blei et al., 2016, for a contemporary review) the idea is to approximate the posterior of a model by selecting the optimal distribution from a fixed family. Optimality is usually defined through the Kullback-Leibler divergence

$$\text{KL}[q(f(x)) \parallel p(f(x) \mid \mathbf{y})] = \mathbb{E}_{q(f(x))} [\log q(f(x)) - \log p(f(x) \mid \mathbf{y})] . \quad (21)$$

This equation is a slight abuse of notation, since it is not legitimate to write $p(f(x))$. Nonetheless, the intuition is good, and our overview has the same result as given by a more technical derivation (Matthews et al., 2016).

2.4.1 THE APPROXIMATING STOCHASTIC PROCESS

We are tasked with defining a family of variational distributions q . Similarly to the sparse GP models, we introduce a set of pseudo inputs (or ‘inducing points’) $\mathbf{Z} = [z_m]_{m=1}^M$, which lie in the same domain as \mathbf{X} , but with $M < N$. Note that \mathbf{Z} do not form part of the model, but only the variational approximating distribution $q(f(x))$: we may choose them freely, and we assume here that \mathbf{X} and \mathbf{Z} do not intersect. We collect the values of the function at \mathbf{Z} into a vector $\mathbf{u} = [f(z_m)]_{m=1}^M$, and write the process conditioned on these values as

$$f(x) \mid \mathbf{u} \sim \mathcal{GP} \left(\mathbf{k}_u(x)^\top \mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1} \mathbf{u}, k(x, x') - \mathbf{k}_u(x)^\top \mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1} \mathbf{k}_u(x') \right) , \quad (22)$$

echoing the true posterior of the process, equation (4). We emphasize that we have *selected* this particular form of distribution for $q(f(x) \mid \mathbf{u})$, and the joint approximation for the posterior is $q(\mathbf{u})q(f(x) \mid \mathbf{u})$. Whilst most of the flexibility of the approximation comes from $q(\mathbf{u})$, using this form of conditional GP for $q(f(x) \mid \mathbf{u})$ means that the approximation cannot be written as a parametric model, and does not suffer the degenerate behaviour associated with the DTC approximation (20).

We must also choose some approximate posterior distribution $q(\mathbf{u})$, whose exact form will depend on the likelihood $p(\mathbf{y} \mid f(x))$ as we shall discuss momentarily.

2.4.2 THE ELBO

Variational Bayes proceeds by maximizing the Evidence Lower Bound (ELBO), which indirectly *minimizes* the Kullback-Leibler objective. By expanding the true posterior using Bayes’ rule, and substituting this into equation (21), we have

$$\begin{aligned} \text{KL} [q(f(x)) \parallel p(f(x) \mid \mathbf{y})] &= -\mathbb{E}_{q(f(x))} \left[\log \frac{p(\mathbf{y} \mid f(x)) p(f(x))}{q(f(x))} \right] + \log p(\mathbf{y}) \\ &\triangleq -\text{ELBO} + \log p(\mathbf{y}) . \end{aligned} \quad (23)$$

To obtain a tractable ELBO for Gaussian processes, we factor both prior and approximate posterior processes into conditional GPs, conditioning on: the inducing input points and values (\mathbf{Z}, \mathbf{u}) ; the data pairs (\mathbf{X}, \mathbf{f}) ; the remainder of the process $f(x)$. The prior distribution on the process $p(f(x))$ can be written $p(\mathbf{u})p(\mathbf{f} | \mathbf{u})p(f | \mathbf{f}, \mathbf{u})$, with

$$\begin{aligned} p(\mathbf{u}) &= \mathcal{N}(\mathbf{0}, \mathbf{K}_{\mathbf{uu}}) \\ p(\mathbf{f} | \mathbf{u}) &= \mathcal{N}\left(\mathbf{K}_{\mathbf{fu}}\mathbf{K}_{\mathbf{uu}}^{-1}\mathbf{u}, \mathbf{K}_{\mathbf{ff}} - \mathbf{K}_{\mathbf{fu}}\mathbf{K}_{\mathbf{uu}}^{-1}\mathbf{K}_{\mathbf{fu}}^\top\right) \\ p(f(x) | \mathbf{f}, \mathbf{u}) &= \mathcal{GP}(m^*(x), k^*(x, x')) . \end{aligned} \tag{24}$$

where $m^*(x)$ and $k^*(x, x')$ are the usual Gaussian process conditional mean and variance, conditioned on both \mathbf{f} and \mathbf{u} . The approximate posterior process can be factored similarly:

$$\begin{aligned} q(\mathbf{f} | \mathbf{u}) &= \mathcal{N}\left(\mathbf{K}_{\mathbf{fu}}\mathbf{K}_{\mathbf{uu}}^{-1}\mathbf{u}, \mathbf{K}_{\mathbf{ff}} - \mathbf{K}_{\mathbf{fu}}\mathbf{K}_{\mathbf{uu}}^{-1}\mathbf{K}_{\mathbf{fu}}^\top\right) , \\ q(f(x) | \mathbf{f}, \mathbf{u}) &= \mathcal{GP}(m^*(x), k^*(x, x')) . \end{aligned} \tag{25}$$

Noting that the two processes are the same, aside for $p(\mathbf{u})$ and $q(\mathbf{u})$, substituting into equation (23) significantly simplifies the ELBO:

$$\text{ELBO} = \mathbb{E}_{q(\mathbf{u})q(\mathbf{f} | \mathbf{u})} [\log p(\mathbf{y} | \mathbf{f})] - \mathbb{E}_{q(\mathbf{u})} \left[\log \frac{q(\mathbf{u})}{p(\mathbf{u})} \right] . \tag{26}$$

2.4.3 THE OPTIMAL APPROXIMATION

It is possible to show that the distribution $\hat{q}(\mathbf{u})$ that maximizes the ELBO is given by

$$\log \hat{q}(\mathbf{u}) = \mathbb{E}_{q(\mathbf{f} | \mathbf{u})} [\log p(\mathbf{y} | \mathbf{f})] + \log p(\mathbf{u}) + \text{const}. \tag{27}$$

This distribution is intractable for general likelihoods, but can be well approximated using MCMC methods. Parameters of the covariance function can be incorporated into the MCMC scheme in a straightforward manner, see Hensman et al. (2015a) for details.

2.4.4 GAUSSIAN APPROXIMATIONS

A computationally cheaper method than sampling $\hat{q}(\mathbf{u})$ is to approximate it using a Gaussian distribution, and to optimize the ELBO with respect to the mean and covariance of the approximating distribution, in order to minimize the KL divergence (Hensman et al., 2015b). If the approximating Gaussian distribution $q(\mathbf{u})$ has mean \mathbf{m} and covariance $\mathbf{\Sigma}$, then the entire approximating process is a GP, with

$$\begin{aligned} q(f(x)) &= \int q(\mathbf{u})q(f(x) | \mathbf{u}) d\mathbf{u} \\ &= \mathcal{GP}\left(\mathbf{k}_{\mathbf{u}}(x)^\top \mathbf{K}_{\mathbf{uu}}^{-1}\mathbf{m}, k(x, x') + \mathbf{k}_{\mathbf{u}}(x)^\top (\mathbf{K}_{\mathbf{uu}}^{-1}\mathbf{\Sigma}\mathbf{K}_{\mathbf{uu}}^{-1} - \mathbf{K}_{\mathbf{uu}}^{-1})\mathbf{k}_{\mathbf{u}}(x')\right) . \end{aligned} \tag{28}$$

The covariance function parameters can be estimated by optimization of the ELBO alongside the variational parameters, which leads to approximate maximum likelihood estimation. This may lead to some bias in the estimation of the parameters (Turner and Sahani, 2011), but the method is reported to work well in practice (Chai, 2012; Hensman et al., 2015b; Lloyd et al., 2015; Dezfouli and Bonilla, 2015).

2.4.5 GAUSSIAN LIKELIHOOD

For the special case where the data-likelihood $p(y_n | f(x_n))$ is Gaussian with noise-variance σ_n^2 , the optimal distribution for q is given by

$$\hat{q}(\mathbf{u}) = \mathcal{N}(\hat{\mathbf{m}}, \hat{\Sigma}), \quad \text{with} \quad \hat{\Sigma} = [\mathbf{K}_{\mathbf{uu}}^{-1} + \sigma_n^{-2} \mathbf{K}_{\mathbf{uu}}^{-1} \mathbf{K}_{\mathbf{uf}} \mathbf{K}_{\mathbf{uf}}^\top \mathbf{K}_{\mathbf{uu}}^{-1}]^{-1}, \quad (29)$$

$$\hat{\mathbf{m}} = \sigma_n^{-2} \hat{\Sigma} \mathbf{K}_{\mathbf{uf}} \mathbf{y},$$

and the ELBO at this optimal point is

$$\text{ELBO}(\hat{q}) = \log \mathcal{N}(\mathbf{y} | 0, \mathbf{K}_{\mathbf{fu}} \mathbf{K}_{\mathbf{uu}}^{-1} \mathbf{K}_{\mathbf{uf}} + \sigma_n^2 \mathbf{I}) - \frac{1}{2} \sigma_n^{-2} \text{tr}(\mathbf{K}_{\mathbf{ff}} - \mathbf{K}_{\mathbf{fu}} \mathbf{K}_{\mathbf{uu}}^{-1} \mathbf{K}_{\mathbf{uf}}). \quad (30)$$

This expression might give the misleading impression that this approximation resembles the DTC method (described by Quiñonero-Candela and Rasmussen, 2005) if one interprets the matrix $\mathbf{K}_{\mathbf{fu}} \mathbf{K}_{\mathbf{uu}}^{-1} \mathbf{K}_{\mathbf{uf}}$ as an approximate prior. However, prediction in the variational approximation does not suffer the same degenerate behaviour as DTC (Titsias, 2009; Hensman et al., 2013).

2.4.6 PREDICTION

Prediction in the variational approximation to Gaussian processes differs from the view given in Quiñonero-Candela and Rasmussen (2005), and it is here that the elegance of the method is apparent. It is not necessary to make any additional approximations at predict-time, since the whole process has already been approximated. For Gaussian posteriors (and Gaussian approximate posteriors), one simply evaluates the GP (equation (28)). For non Gaussian approximations to $q(\mathbf{u})$, one must average the conditional equation (22) under $q(\mathbf{u})$.

2.5 Linking conditioning and decomposition

The sparse-variational methodology that we have presented involved factoring the Gaussian process as $p(f(x)) = p(f(x) | \mathbf{u}) p(\mathbf{u})$. A related and useful concept is decomposition of the process: we will make use of such a decomposition in subsequent sections to derive Variational Fourier Features.

Let $f(x) \sim \mathcal{GP}(0, k(x, x'))$. We can decompose f as the sum of two independent Gaussian processes:

$$g(x) \sim \mathcal{GP}(0, \mathbf{k}_{\mathbf{u}}(x)^\top \mathbf{K}_{\mathbf{uu}}^{-1} \mathbf{k}_{\mathbf{u}}(x')), \quad (31)$$

$$h(x) \sim \mathcal{GP}(0, k(x, x') - \mathbf{k}_{\mathbf{u}}(x)^\top \mathbf{K}_{\mathbf{uu}}^{-1} \mathbf{k}_{\mathbf{u}}(x')), \quad (32)$$

$$f(x) = g(x) + h(x). \quad (33)$$

It is clear from the additivity of the covariance functions that this sum recovers the correct covariance for $f(x)$. Note that the covariance of $h(x)$ is the same as for $p(f(x) | \mathbf{u})$. We say that $g(x)$ is the projection of $f(x)$ onto \mathbf{u} , and $h(x)$ is the orthogonal complement of $g(x)$.

Figure 1 shows the decomposition of a Gaussian process using inducing points and also with the Fourier features that we shall derive in this document.

In the variational approximation, the posterior is controlled through the variables \mathbf{u} , and the conditional $p(f(x) | \mathbf{u})$ remains the same in the approximate posterior as in the

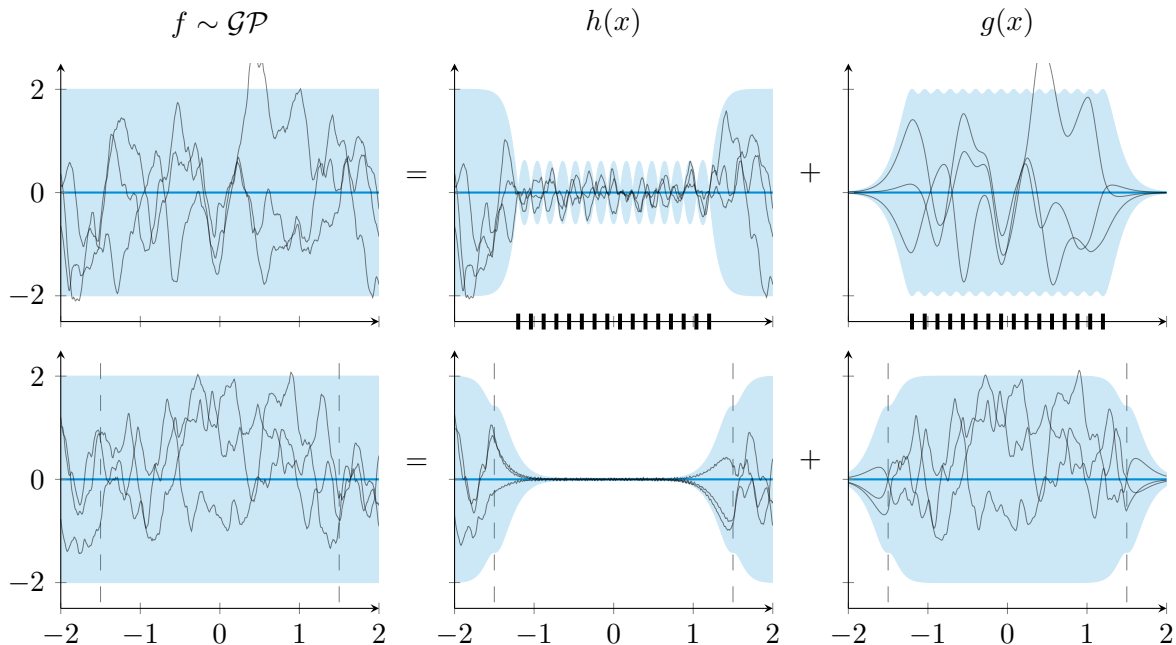


Figure 1: Decomposition of a Gaussian process with a Matérn- $\frac{3}{2}$ covariance function ($\sigma = 1, \ell = 0.2$) using inducing points (top) and Variational Fourier Features (bottom). The shaded region represents two standard deviations, and samples from the processes are shown as black lines. In the top row, the positions of the 16 inducing points are marked in black; on the bottom row 31 frequencies (hence 63 inducing variables) are used and vertical lines denote the boundary of the projection $[a, b]$. Note the variance of $h(x)$ is zero wherever there is an inducing point (top). In the bottom row, the variance is close to zero for a continuous stretch (where the data should live); in this region $h(x)$ contains only high-frequency functions that we have failed to capture with the Fourier features. The variance increases towards the boundaries: since all Fourier basis functions have the same value at a and b , they cannot fit functions with different values in a and b and this is left as a residual. Instead of adding more basis functions, such as a linear one, to account for this, we advocate for choosing an interval $[a, b]$ larger than the data.

prior. In terms of the projection, we see that the process $g(x)$ is completely controlled by the variables \mathbf{u} and the process $h(x)$ is completely independent of \mathbf{u} . This is borne out by the independence relations of the processes, formally:

$$g(x) \mid \mathbf{u} = \mathcal{GP} \left(\mathbf{k}_{\mathbf{u}}(x)^\top \mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1} \mathbf{u}, 0 \right), \quad (34)$$

$$h(x) \mid \mathbf{u} = h(x). \quad (35)$$

The assumption of the sparse variational framework is that the posterior process can be captured by the random variables \mathbf{u} , or equivalently by the process $g(x)$. The power of the framework lies in the fact that the ELBO can be used to evaluate whether the projection is sufficient, and that the orthogonal part $h(x)$ is not discarded but taken into account in the ELBO computation and incorporated fully at predict time. The ELBO encapsulates the degree to which the posterior can be captured by $g(x)$ (or equivalently \mathbf{u}).

2.6 Inter-domain approaches

The most common method for efficient approximation of Gaussian processes is based on the idea of inducing variables, denoted $u_m = f(z_m)$ in the above. A powerful idea is to generalize this, allowing a different decomposition of the process by constructing linear combinations of the process values as projections, instead of simple evaluations.

Previous work that has considered inter-domain approaches has suggested that the corresponding features have better ability to represent complicated functions (Figueiras-Vidal and Lázaro-Gredilla, 2009) than the inducing points approach, and have promise for applications in multiple-output settings (Alvarez and Lawrence, 2009).

In inter-domain Gaussian process approximations, the idea is to change

$$u_m = f(z_m) \quad (36)$$

to a projection

$$u_m = \int f(x) d\mu_m(x), \quad (37)$$

so that the variable u_m is more informative about the process, and in turn the projection onto \mathbf{u} is able to capture more information about the posterior.

The utility of the variable u_m depends mostly on its covariance with the remainder of the process, which is encapsulated in the vector-valued function $\mathbf{k}_{\mathbf{u}}(x)$. This vector plays a similar role in the variational method to the feature vector $\phi(x)$ in the Fourier-based approximate models of Sections 2.3.1 and 2.3.2. The remit of this work is to construct inducing variables \mathbf{u} such that $\mathbf{k}_{\mathbf{u}}(x)$ contains sinusoidal functions as is $\phi(x)$, which we will do by using Fourier projections of the process.

The central challenge is that the Fourier transform of the process does not make a valid inducing variable because its variance would be infinite. To make valid and useful inducing variables, Matthews et al. (2016) have shown that the inducing variables u_m must be “deterministic, conditioned on the whole latent function”. In the following section, we construct valid inducing variables and examine their properties.

3. Variational Fourier Features

Our goal is to combine the variational sparse GP idea with Fourier features, so that the approximating process in (22) contains a mean function which is built from sinusoids, like the $\phi(x)$ features in the Fourier-features approximate model (15).

This is more tricky than it first appears. An initial idea would be to define u_m as the Fourier transform of the process, so that $\mathbf{k}_u(x) = \text{cov}(\mathbf{u}, x) \propto \phi(x)$. The question is, what would the \mathbf{K}_{uu} matrix be? One might hope for a diagonal matrix related to the spectral content of the kernel $s(\omega)$, but as we discuss in Section 3.1, the Fourier transform of a stationary GP has diverging variance, i.e. this matrix is not defined.

To proceed, we need to make two alterations to the initial idea, which we outline in Sections 3.2 and 3.3. First we must window the domain, in order to obtain variables with finite variance (Section 3.2). Second, since these variables do not have an elegant form, we switch from an L^2 integral to one under the RKHS norm (Section 3.3). Subsequently, we show that our approach has covariance structures that make for efficient computations (Section 3.4), and conclude this section by explaining the behaviour of our approximation outside the windowing box (Section 3.5).

3.1 The trouble with Fourier features

In the variational sparse GP (Section 2.4) we made use of the random variables $u_m = f(z_m)$. To obtain a Fourier-based approximation, we might consider the inter-domain variables $u_m = \int_{-\infty}^{\infty} f(t)e^{-i\omega_m t} dt$. It is straightforward to show that such variables have zero mean and infinite variance, and so will not be useful in the variational framework; nonetheless the properties of the variables are interesting as we now demonstrate.

In order to use u_m as an inducing variable, we would require both $\text{cov}(u_m, f(x))$ and $\text{cov}(u_m, u_{m'})$ which make up the entries in $\mathbf{k}_u(x)$ and \mathbf{K}_{uu} respectively. If f is drawn from a GP with zero mean and covariance $k(x, x')$, then formally we may write

$$\text{cov}(u_m, f(x)) = \mathbb{E}[u_m f(x)] = \int_{-\infty}^{\infty} \mathbb{E}[f(t)f(x)]e^{-i\omega_m t} dt = \int_{-\infty}^{\infty} k(t, x)e^{-i\omega_m t} dt. \quad (38)$$

To solve this integral, we can first plug in the definition of the kernel function in terms of the spectrum (7), and then change the order of integration, giving

$$\text{cov}(u_m, f(x)) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} s(\omega)e^{i\omega(t-x)} d\omega e^{-i\omega_m t} dt = s(\omega_m)e^{-i\omega_m x}. \quad (39)$$

This initially appears promising: the covariance function is a sinusoid rescaled by the relevant spectrum and the elements of $\mathbf{k}_u(x)$ have the desired form. The corresponding elements of \mathbf{K}_{uu} are given by the (co)variance of u_m . It is here that problems arise. Denoting the complex conjugate of u_m as \bar{u}_m ,

$$\text{cov}(u_m, u_{m'}) = \mathbb{E}[u_m \bar{u}_{m'}] = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} k(t, t')e^{-i\omega_m t} dt e^{i\omega_{m'} t'} dt' = s(\omega_m)\delta(\omega_m - \omega_{m'}), \quad (40)$$

where $\delta(\cdot)$ is Dirac's delta. This implies that the matrix \mathbf{K}_{uu} is diagonal, but with undefined entries on the diagonal. This result shows that it is simply not possible to use the Fourier transform of the whole GP as a valid inducing variable.

One approach to address this issue might be to take a power-averaged Fourier transform of the form

$$u_m = \lim_{a \rightarrow \infty} \frac{1}{a} \int_{-a/2}^{a/2} f(t) e^{-i\omega_m t} dt. \quad (41)$$

Unfortunately, this still does not yield useful features, because the variables \mathbf{u}_m are now independent of the function at any particular value. Writing formally:

$$\text{cov}(u_m, f(x)) = \mathbb{E}[u_m f(x)] = \lim_{a \rightarrow \infty} \frac{1}{a} \int_{-a/2}^{a/2} k(x, t) e^{-i\omega_m t} dt = 0. \quad (42)$$

As recognised by Figueiras-Vidal and Lázaro-Gredilla (2009), the only way to obtain valid features is to specify an input density (cf. Williams and Seeger, 2000). We proceed by considering a uniform input density on the interval $[a, b]$.

3.2 L_2 Fourier features on $[a, b]$

The reason that the Fourier transform of a GP behaves so strangely is that it must explain the behaviour over the whole real line. One might interpret the Fourier transform at a particular frequency as a sum of all the random variables in the GP (i.e. integrating over the real line) multiplied by a sinusoidal weight. The result is a Gaussian random variable, but unless the weight decays to zero at infinity, the result has infinite variance.

Our solution to this diverging integral is to window the integral. Figueiras-Vidal and Lázaro-Gredilla (2009) also use a window to ensure convergence of the integral, but their choice of a Gaussian windowing function means that the result is tractable only for Gaussian (squared exponential) covariance functions, and does not admit a diagonal form for $\mathbf{K}_{\mathbf{u}\mathbf{u}}$.

We will apply a square window, effectively changing the integration limits to a and b :

$$u_m = \int_a^b f(t) e^{-i\omega_m(t-a)} dt. \quad (43)$$

In addition, we will assume that the frequency ω_m is harmonic on the interval $[a, b]$; that is

$$\omega_m = \frac{2\pi m}{b-a}. \quad (44)$$

For $x \in [a, b]$, the covariance between such inducing variables and the GP at x is

$$\text{cov}(u_m, f(x)) = \mathbb{E}[u_m f(x)] = \int_a^b \mathbb{E}[f(t) f(x)] e^{-i\omega_m(t-a)} dt = \int_a^b k(t, x) e^{-i\omega_m(t-a)} dt. \quad (45)$$

These integrals are tractable for Matérn kernels. As detailed in Appendix A, we obtain the following results for Matérn- $\frac{1}{2}$ kernels:

$$\text{cov}(u_m, f(x)) = s_{1/2}(\omega_m) e^{i\omega_m(x-a)} + s_{1/2}(\omega_m) \frac{1}{2\lambda} \left(\lambda[-e^{a-x} - e^{x-b}] + i\omega_m[e^{a-x} - e^{x-b}] \right). \quad (46)$$

Similarly, the covariance between two inducing variables also has a closed form expression. A complete derivation is given in Appendix A for the Matérn- $\frac{1}{2}$ case. These inter-domain inducing variables have two desirable properties: they have finite variance and their associated covariance matrix $\mathbf{K}_{\mathbf{uu}}$ can be written as a diagonal matrix plus some rank one matrices (two in the Matérn- $\frac{1}{2}$ case).

Figure 2 illustrates entries of the feature vector $\mathbf{k}_{\mathbf{u}}(x) = \text{cov}(u_m, f(x))$, and compares them to the RKHS features that we shall derive in the next section. We see from the Figure that the covariance function $\text{cov}(u_m, f(x))$ is almost sinusoidal in x for a region sufficiently far from the boundaries a and b . The ‘edge effects’ depend on the involved frequency as well as the lengthscale parameter of the kernel.

We have constructed inducing variables that are valid (in the sense that they have finite variance), but the result is somewhat inelegant: the expressions for $\text{cov}(u_m, f(x))$ and $\text{cov}(u_m, u_{m'})$ are long for the Matérn- $\frac{1}{2}$ kernel and become trickier (but still tractable) for higher orders. We abandon them at this stage in preference of the RKHS-based inducing variables that we derive in the next section. These have preferable edge effects and an elegant expressions for the required covariances.

3.3 RKHS Fourier features on $[a, b]$

In the previous section we saw that choosing inducing variables of the form $u_i = \langle f, \cos(\omega(xa)) \rangle_{L^2}$ or $u_i = \langle f, \sin(\omega(xa)) \rangle_{L^2}$ resulted in features that are very distinct from cosine and sine functions. In this section, we replace the L^2 inner product by the RKHS inner product $u_i = \langle f, \cos(\omega(xa)) \rangle_{\mathcal{H}}$ in order to guarantee that the features are exactly the sine and cosine functions. As we will discuss later, a direct asset of this approach is that the simple expression of the features makes the computation of the covariances between inducing variables much easier.

We start with a truncated Fourier basis defined similarly to (12)

$$\phi(x) = [1, \cos(\omega_1(x-a)), \dots, \cos(\omega_M(x-a)), \sin(\omega_1(x-a)), \dots, \sin(\omega_M(x-a))]^\top, \quad (47)$$

where we include the constant basis function, $\phi_0(x) = 1$, accounting for $\omega_m = 0$, and define $\omega_m = \frac{2\pi m}{b-a}$ as previously.

A key RKHS result (see Berlinet and Thomas-Agnan, 2004, Theorem 11), is that if $\mathcal{F} = \text{span}(\phi)$ is a subspace of a RKHS \mathcal{H} , then \mathcal{F} has the kernel

$$k_{\mathcal{F}}(x, x') = \phi(x)^\top \mathbf{K}_{\phi\phi}^{-1} \phi(x'), \quad (48)$$

where $\mathbf{K}_{\phi\phi}[m, m'] = \langle \phi_m, \phi_{m'} \rangle_{\mathcal{H}}$ is the Gram matrix of ϕ in \mathcal{H} . By this, we mean that for any function $f \in \mathcal{F}$, $\langle f, k_{\mathcal{F}}(\cdot, x) \rangle_{\mathcal{H}} = f(x)$. Furthermore, the coordinate of the projection of a function $h \in \mathcal{H}$ onto ϕ_m is defined as

$$\mathcal{P}_{\phi_m}(h) = \langle h, \phi_m \rangle_{\mathcal{H}}. \quad (49)$$

Since $\mathcal{F} \subset \mathcal{H}$, $k(x, x') - k_{\mathcal{F}}(x, x')$ is a positive definite function corresponding to the kernel of \mathcal{F}^\perp , and we can decompose the GP in a similar way to (33) which gives

$$g(x) \sim \mathcal{GP}(0, \phi(x)^\top \mathbf{K}_{\phi\phi}^{-1} \phi(x')), \quad (50)$$

$$h(x) \sim \mathcal{GP}(0, k(x, x') - \phi(x)^\top \mathbf{K}_{\phi\phi}^{-1} \phi(x')), \quad (51)$$

$$f(x) = g(x) + h(x). \quad (52)$$

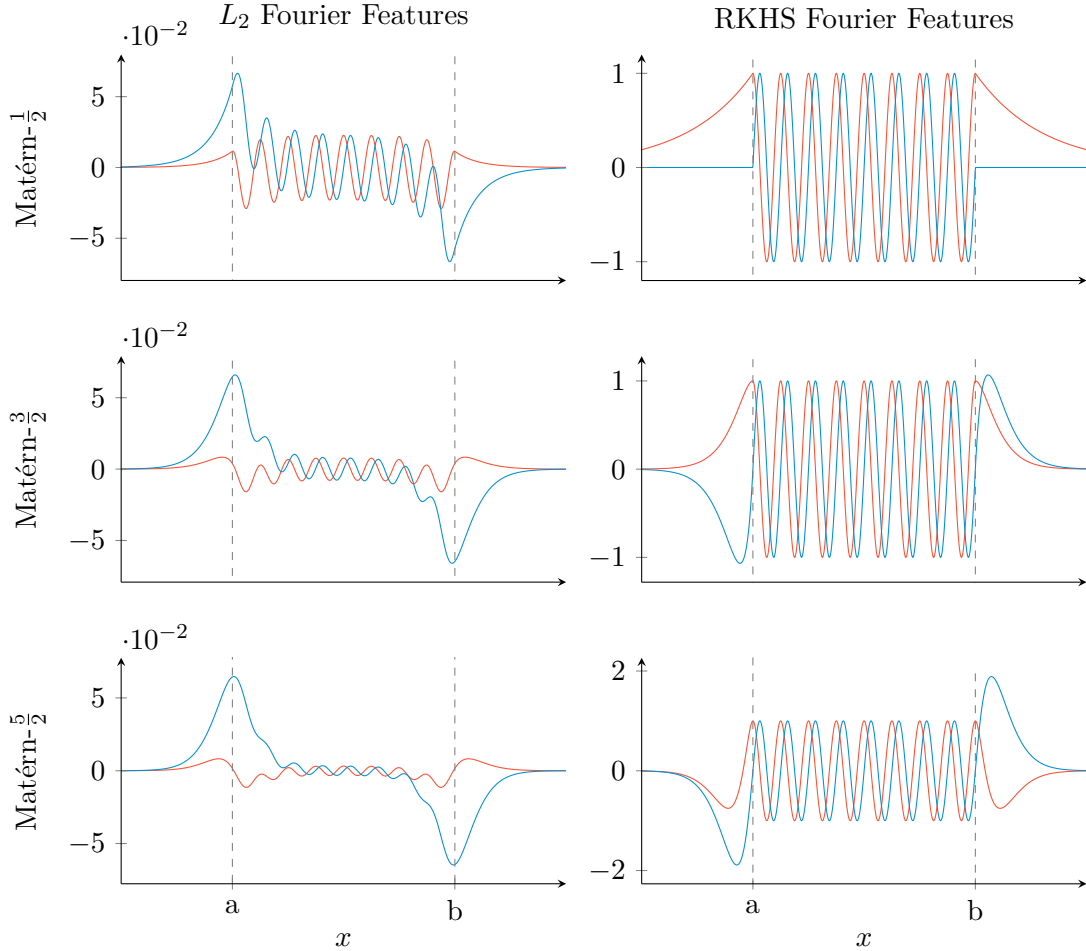


Figure 2: The covariance $\text{cov}(u_m, f(x))$ as a function of x for the Matérn- $\frac{1}{2}$ (top), Matérn- $\frac{3}{2}$ (middle) and Matérn- $\frac{5}{2}$ (bottom) kernels, using L_2 Fourier features (left column) and RKHS Fourier features (right column). The red and blue lines represent the real and imaginary parts of the covariance. The boundaries of the projection $[a, b]$ are marked by vertical dashed lines. The L_2 features are wiggly with overwhelming edge effects, whereas the RKHS features are exactly sinusoids for $x \in [a, b]$. Beyond the boundaries, the covariance reverts to zero in a way which depends on the smoothness of the kernel and the lengthscale parameter. We have used $\omega_m = 16\pi/(b-a)$, and a lengthscale of $\ell = (b-a)/10$.

For this to be valid we need two things: \mathcal{F} must be included in \mathcal{H} and the matrix $\mathbf{K}_{\phi\phi}$ must be invertible. The first is not true in general (Matérn RKHS over \mathbb{R} or the RBF and Brownian RKHS on $[0, 1]$ are counter examples) but has been proved to be true for Matérn kernels over $[a, b]$ (Durrande et al., 2016). The second is a direct conclusion of the linear independence between the elements of ϕ . Furthermore, Durrande et al. (2016) also detail the closed-form expressions of the inner products for the Matérn- $\frac{1}{2}$, Matérn- $\frac{3}{2}$ and Matérn- $\frac{5}{2}$ RKHSs. These expressions, which are repeated in Appendix B, are of particular interest here since they allow us to compute the entries of the $\mathbf{K}_{\phi\phi}$ matrix. Although we refer to the above reference for the constructive proof of these expressions, it is easy to check that the reproducing property is satisfied. For example, the expression given in the appendix for a Matérn- $\frac{1}{2}$ RKHS over $[a, b]$ is

$$\langle g, h \rangle_{\mathcal{H}_{1/2}} = \frac{1}{2\lambda\sigma^2} \int_a^b (\lambda g(t) + g'(t)) (\lambda h(t) + h'(t)) dt + \frac{1}{\sigma^2} g(a) h(a), \quad (53)$$

so we can write for any function g that is continuous and (weakly) differentiable

$$\left\langle g, k_{\frac{1}{2}}(x, \cdot) \right\rangle_{\mathcal{H}_{1/2}} = \frac{1}{2\lambda\sigma^2} \left(\int_a^x (\lambda g(t) + g'(t)) 2\lambda\sigma^2 e^{\lambda(t-x)} dt + \int_x^b 0 dt \right) + g(a) e^{\lambda(a-x)} \quad (54)$$

$$= \int_a^x \lambda g(t) e^{\lambda(t-x)} dt + \left[g(t) e^{\lambda(t-x)} \right]_a^x - \int_a^x g(t) \lambda e^{\lambda(t-x)} dt + g(a) e^{\lambda(a-x)} \quad (55)$$

$$= g(x) - g(a) e^{\lambda(a-x)} + g(a) e^{\lambda(a-x)} \quad (56)$$

$$= g(x). \quad (57)$$

Another property of this inner product is that the Gram matrix $\mathbf{K}_{\phi\phi}$ has a particular structure which allows to reduce drastically the computational burden of computing and inverting $\mathbf{K}_{\phi\phi}$. This will be addressed in Section 3.4.

3.3.1 THE CORRESPONDING RANDOM VARIABLES

To interpret these results from a Gaussian process point of view, we need to find the inducing variables \mathbf{u} that would lead to the exact same decomposition.

We would like to use the RKHS inner product between the sinusoids and the GP sample path in place of the L_2 inner product (43), but since GP samples do not belong the RKHS (Driscoll, 1973), it is a priori not possible to apply \mathcal{P}_{ϕ_m} to f and define $u_m = \langle \phi_m, f \rangle_{\mathcal{H}}$. For example, the Matérn- $\frac{1}{2}$ inner product involves derivatives whereas the Matérn- $\frac{1}{2}$ samples are not differentiable anywhere. However, using the fact that the functions from ϕ are very regular, it is possible to extend the operators $\mathcal{P}_{\phi_m} : h \mapsto \langle \phi_m, h \rangle_{\mathcal{H}}$ to square integrable functions using integration by parts. For the Matérn- $\frac{1}{2}$ kernel, integrating (53) results in

$$\mathcal{P}_{\phi_m}(h) = \frac{1}{2\lambda\sigma^2} \left(\int_a^b h(\lambda^2 \phi_m - \phi_m'') dt + h(b) (\lambda \phi_m(b) + \phi_m'(b)) + h(a) (\lambda \phi_m(a) - \phi_m'(a)) \right). \quad (58)$$

Note that the function h does not need to be differentiated, and so it is possible to apply this functional to a sample from the GP. Similar results apply for the Matérn- $\frac{3}{2}$ and Matérn- $\frac{5}{2}$ kernels. It is now possible to apply these operators to the Gaussian process in order to construct the inducing variables,

$$u_m = \mathcal{P}_{\phi_m}(f). \quad (59)$$

The covariance of the inducing variables with the function values is given by

$$\text{cov}(u_m, f(x)) = \mathbb{E}[u_m f(x)] = \mathbb{E}[\mathcal{P}_{\phi_m}(f)f(x)] = \mathcal{P}_{\phi_m}(k(x, \cdot)) = \phi_m(x) \quad (60)$$

which is valid for $a \leq x \leq b$. This means that $\mathbf{k}_{\mathbf{u}}(x) = \phi(x)$ when $x \in [a, b]$. Similarly the covariance of the features is given by

$$\text{cov}(u_m, u_{m'}) = \mathbb{E}[u_m u_{m'}] = \mathbb{E}[\mathcal{P}_{\phi_m}(f)\mathcal{P}_{\phi_{m'}}(f)] = \mathcal{P}_{\phi_m}(\phi_{m'}) = \langle \phi_m, \phi_{m'} \rangle_{\mathcal{H}}, \quad (61)$$

and so we have $\mathbf{K}_{\mathbf{uu}} = \mathbf{K}_{\phi\phi}$.

3.4 Covariance structures for RKHS Fourier features

To compute the covariance matrix $\mathbf{K}_{\mathbf{uu}}$, we return again to the closed-form expressions for the inner product provided by Durrande et al. (2016), substituting the basis functions ϕ_m and $\phi_{m'}$ appropriately. The solutions are tractable, and we detail in the appendix the expressions of $\mathbf{K}_{\mathbf{uu}}$ for the first three half-integer Matérn kernels. One striking property of the resulting $\mathbf{K}_{\mathbf{uu}}$ is that they can be written as a diagonal matrix plus a few rank one matrices. For example in the Matérn- $\frac{1}{2}$ case we find that $\mathbf{K}_{\mathbf{uu}}$ is equal to a diagonal matrix plus a rank one matrix,

$$\mathbf{K}_{\mathbf{uu}} = \text{diag}(\boldsymbol{\alpha}) + \boldsymbol{\beta}\boldsymbol{\beta}^\top, \quad (62)$$

with

$$\boldsymbol{\alpha} = \frac{b-a}{2} [2s(0)^{-1}, s(\omega_1)^{-1}, \dots, s(\omega_M)^{-1}, s(\omega_1)^{-1}, \dots, s(\omega_M)^{-1}]^\top, \quad (63)$$

$$\boldsymbol{\beta} = [\sigma^{-1}, \sigma^{-1}, \dots, \sigma^{-1}, 0, \dots, 0]^\top. \quad (64)$$

As shown in the appendix, $\mathbf{K}_{\mathbf{uu}}$ still has a similar structure for higher order Matérn kernels: in the Matérn- $\frac{3}{2}$ case, it is the sum of a diagonal matrix and two rank one matrices and in the Matérn- $\frac{5}{2}$ case, it is the sum of a diagonal matrix and three rank one matrices.

Since $\mathbf{K}_{\mathbf{uu}}$ has a low-rank plus diagonal form, one of the usual computational bottlenecks in the variational framework, solving $\mathbf{K}_{\mathbf{uu}}^{-1}\mathbf{K}_{\mathbf{uf}}$, can be done in $\mathcal{O}(NM)$ operations, rather than the standard $\mathcal{O}(NM^2)$, by a straightforward application of the Woodbury matrix identity.

This low-rank plus diagonal structure was not published previously in Durrande et al. (2016); we believe that we are the first to make use of this result for computational efficiency and we expect it to provide similar superior computational benefits as the pure diagonal structure in Solin and Särkkä (2014).

$\phi_m(x), x \in [a, b]$	Matérn- $\frac{1}{2}$	Matérn- $\frac{3}{2}$	Matérn- $\frac{5}{2}$
$\cos(\omega_m(x - a))$	$e^{-\lambda r}$	$(1 + \lambda r)e^{-\lambda r}$	$(1 + \lambda r + \frac{1}{2}(\lambda^2 - \omega_m^2)r^2)e^{-\lambda r}$
$\sin(\omega_m(x - a))$	0	$s r \omega_m e^{-\lambda r}$	$s r \omega_m (1 + \lambda r)e^{-\lambda r}$

Table 1: The covariance $\text{cov}(u_m, f(x))$ for x outside the interval $[a, b]$. Here, we define r as the absolute distance to the closest edge (a or b), and s to be 1 for $x < a$ and -1 for $x > b$.

3.5 Predicting outside the interval $[a, b]$

In the preceding sections, we have defined our inducing variables and how they covary with the function; we have assumed throughout that we wish to examine a point on the latent function in the pre-defined interval $[a, b]$. This does not pose a problem at training time, since we can easily define $[a, b]$ to contain all the data. For prediction we may need to predict outside this interval. For a point on the function $f(x)$ outside the interval $[a, b]$, the covariance with the inducing variables is still well defined, though the form is slightly more complicated than the simple sinusoids inside the interval.

To obtain the expressions for the covariance beyond $[a, b]$, we apply the closed-form expression for the covariance (58), this time for $\text{cov}(u_m, f(x)) = \text{cov}(\mathcal{P}_{\phi_m}(f), f(x))$ with $x > b$ and $x < a$. After some fairly elementary simplifications, the results are shown in Table 1, and illustrated in Figure 2. The result is that the covariance function $\text{cov}(u_m, f(x))$ returns to zero beyond the boundary, with smoothness that depends on the order of the kernel.

4. Extending the applicable kernel family

In the above we have focused on the Matérn family of kernels in one dimension. In this section we will expand the approach to higher-dimensional inputs using sums and products of kernels.

4.1 Additive kernels

A straightforward way to produce a Gaussian process prior on multiple inputs (say D) is to use a sum of independent Gaussian processes, one for each input:

$$f(\mathbf{x}) = \sum_{d=1}^D f_d(x_d), \quad f_d \sim \mathcal{GP}(0, k_d(x_d, x'_d)), \quad (65)$$

where x_d is the d th element in \mathbf{x} and $k_d(\cdot, \cdot)$ is a kernel defined on a scalar input space. This construction leads to a Gaussian process with an additive kernel, that is we can write

$$f(\mathbf{x}) \sim \mathcal{GP}\left(0, \sum_{d=1}^D k_d(x_d, x'_d)\right). \quad (66)$$

Additive kernel structures have been explored by Durrande et al. (2012) and Duvenaud et al. (2011), who have shown that these kernels are well suited to high-dimensional problems. To use our Variational Fourier Features approach with an additive kernel, we assume that each function $f_d(x_d)$ has a Matérn covariance function, and decompose each of these GPs. We construct a matrix of features, with elements

$$u_{m,d} = \mathcal{P}_{\phi_m}(f_d). \quad (67)$$

The result is that we have DM features. By construction, we see that features from different GPs are independent, $\text{cov}(u_{m,d}, u_{m,d'}) = 0$, and that the covariance between features for the same dimension follows the construction for a single dimension. It is straightforward to exploit these independence structures for computational scalability during inference.

4.2 Separable kernels

A kernel is said to be separable if it can be written as a product of kernels with no shared inputs. This means that for any D -dimensional input \mathbf{x} , we can write

$$k(\mathbf{x}, \mathbf{x}') = \prod_{d=1}^D k_d(x_d, x'_d). \quad (68)$$

We can construct Variational Fourier Features for such a product kernel if each sub-kernel is in the Matérn family.

This formulation of separability has been extensively used in speeding up GP inference (see, e.g., Saatçi, 2012; Stegle et al., 2011) by writing the covariance matrix as a Kronecker product of the individual covariance matrices. These speed-ups usually require the data to be aligned with a regular or rectilinear grids (see, e.g., Solin et al., 2016), though some approaches exist to extend to observations beyond the grid (Wilson et al., 2014; Nickson et al., 2015), requiring additional approximations.

In contrast, our approach of decomposing the kernel naturally leads to a Kronecker structure, even for irregularly spaced data. Let $f(x) \sim \mathcal{GP}\left(0, \prod_d k_d(x_d, x'_d)\right)$, and define a vector of features as the Kronecker product of features over each dimension,

$$\phi(\mathbf{x}) = \bigotimes_{d=1}^D [\phi_1(x_d), \dots, \phi_M(x_d)]^\top, \quad (69)$$

so that each of the M^D elements of $\phi(\mathbf{x})$ is a product of one-dimensional functions $\prod_d \phi_i^{(d)}(x_d)$. We define a hyper-rectangular boundary given by $\prod_d [a_d, b_d]$, and define the inducing variables \mathbf{u} using a projection similar to the one-dimensional case

$$u_m = \mathcal{P}_{\phi_m}(f), \quad (70)$$

where we use the norm of the Hilbert space associated with the product kernel. The covariance between an inducing point u_m and the function is then given by

$$\text{cov}(u_m, f(\mathbf{x})) = \mathbb{E}[\mathcal{P}_{\phi_m}(f)f(x)] = \mathcal{P}_{\phi_m}(k(\mathbf{x}, \cdot)) = \prod_{d=1}^D \phi_m^{(d)}(x_d) = \phi_m(\mathbf{x}). \quad (71)$$

Extending the covariance function beyond the boundary follows similarly to the one-dimensional case above. The covariance between inducing variables is given by

$$\text{cov}(u_m, u'_m) = \prod_d \left\langle \phi_m^{(d)}, \phi_{m'}^{(d)} \right\rangle_{\mathcal{H}_d}. \quad (72)$$

This means that the covariance matrix $\mathbf{K}_{\mathbf{u}\mathbf{u}}$ has a Kronecker structure, $\mathbf{K}_{\mathbf{u}\mathbf{u}} = \bigotimes_{d=1}^D \mathbf{K}_{\mathbf{u}\mathbf{u}}^{(d)}$, where each sub-matrix $\mathbf{K}_{\mathbf{u}\mathbf{u}}^{(d)}$ has the same structure as for the one-dimensional case (62).

5. Implementation details and computational complexity

In this section we give some details of our implementation of the methods proposed, and consider the theoretical computational complexity of each. To make use of the RKHS Fourier features that we have derived, the expressions given for $\mathbf{K}_{\mathbf{u}\mathbf{u}}$ and $\mathbf{k}_{\mathbf{u}}(x)$ can simply be plugged in to the variational framework described in Section 2.4.

5.1 Implementation

All of the methods in this paper and all of the code to replicate the experiments in Section 6 are available at <http://github.com/jameshensman/VFF>. We made use of TensorFlow (Abadi et al., 2015) and GPflow (Matthews et al., 2017) to construct model classes in Python. Some simple matrix classes in Python assisted with computing efficient solutions to low-rank matrix problems (e.g. solving $\mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1}\mathbf{m}$). Using TensorFlow’s automatic differentiation made application of gradient based optimization straightforward.

5.2 The one-dimensional case

Consider first problems with only one input dimension. The sparse variational framework discussed in Section 2.4 gives three methods: if the likelihood is Gaussian, then it is possible to compute the optimal posterior in closed form (as per Titsias, 2009), if the likelihood is not Gaussian we can approximate the posterior with a Gaussian or use MCMC on the optimal distribution.

For the Gaussian case, we see that the computation of the optimal posterior (equation (29)) is dominated by the matrix multiplication $\mathbf{K}_{\mathbf{u}\mathbf{f}}\mathbf{K}_{\mathbf{f}\mathbf{u}}$, which costs $\mathcal{O}(NM^2)$ operations. In our case, since $\mathbf{K}_{\mathbf{f}\mathbf{u}}$ does not depend on the kernel parameters, this multiplication only needs to be computed once before optimizing (or sampling) the covariance function parameters. To compute the posterior mean and covariance, we must invert an $M \times M$ matrix, so the cost per iteration is $\mathcal{O}(M^3)$, after the initial cost of $\mathcal{O}(NM^2)$.

For the non-Gaussian case with a Gaussian approximation we must compute the marginals of $q(\mathbf{f})$, in order to evaluate the ELBO (equation (26)) for optimization. We parameterize the Gaussian covariance using a lower-triangular matrix \mathbf{L} , so $\mathbf{\Sigma} = \mathbf{L}\mathbf{L}^\top$, and the computational cost is dominated by the matrix multiplication $\mathbf{K}_{\mathbf{f}\mathbf{u}}\mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1}\mathbf{L}$. Even though the matrix $\mathbf{K}_{\mathbf{u}\mathbf{u}}$ has computationally-convenient structure, we cannot avoid a dense matrix multiplication in this case. It may be possible that by assuming some structure for the approximate covariance, computation could be reduced, but we leave investigation of that line for future research.

We note however that the Gaussian approximation does lend itself to stochastic optimization, where the data are randomly sub-sampled (Hensman et al., 2013). In this case the computational complexity is $\mathcal{O}(\tilde{N}M^2)$, where \tilde{N} is the size of the minibatch.

For the MCMC approach, evaluating the unnormalised distribution (27) and its derivative costs $\mathcal{O}(NM)$ per iteration, since again the cost is dominated by the computation of the marginals of $q(\mathbf{f} | \mathbf{u})$, which requires only the product $\mathbf{K}_{\mathbf{f}\mathbf{u}}\mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1}$, and the low-rank structure of $\mathbf{K}_{\mathbf{u}\mathbf{u}}$ enables this to be computed efficiently.

5.3 The additive case

If we consider the kernel to be a sum of kernels over D input dimensions (Section 4.1), then we again have access to the three methods provided by the sparse GP framework. In this case, the number of inducing variables is $2MD + 1$, since we have M frequencies per dimension, with sine and cosine terms.

For the method with a Gaussian likelihood, the cost is $\mathcal{O}(D^3M^3)$ per iteration, with an initial cost of $\mathcal{O}(NM^2D^2)$, following the arguments above. When the likelihood is non-Gaussian, with a Gaussian approximation to the posterior, the cost is $\mathcal{O}(NM^2D^2)$. However, if we assume that the posterior factorizes over the dimensions, then the cost can be reduced to $\mathcal{O}(NM^2D)$ (Adam et al., 2016).

For an MCMC approach, the cost is again dominated by the computation of $\mathbf{K}_{\mathbf{f}\mathbf{u}}\mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1}$, which must be done once for each dimension and so the cost is $\mathcal{O}(NMD)$. We note that these operations can be effectively distributed on multiple machines (in our implementation this can be done by TensorFlow), since $\mathbf{K}_{\mathbf{u}\mathbf{u}}$ is a block-diagonal matrix.

5.4 The Kronecker case

If we use a separable kernel made from a product of kernels over the D dimensions as discussed in Section 4.2, then the number of inducing variables grows exponentially with the input dimension to $(2M)^D$. This gives the potential for a very detailed representation, with many basis functions, and the Kronecker structure can be exploited for efficiency in some cases. However, this approximation will be unsuitable for large D .

For the case with Gaussian noise, it is not straightforward to avoid computing the inverse of the $M^D \times M^D$ matrix, and so the cost is $\mathcal{O}(M^{3D})$. It may be possible, however, that some efficient methods for solving linear systems might be applicable (e.g. Filippone and Engler, 2015), though we leave that avenue for future work.

For the case with a Gaussian approximation to $q(\mathbf{u})$, we can exploit the Kronecker structure of the model. We note the relationship to the work of Nickson et al. (2015), which used the inducing-point approach, with the points confined to a regular lattice. In their work, the covariance structure of the approximation was constrained to also have a Kronecker structure, which appeared to work well in the cases considered. However, our experiments suggest (Section 6.3) that this approximation may not be sufficient in all cases. Our proposal is to use a sum of two Kronecker structured matrices:

$$q(\mathbf{u}) = \mathcal{N}(\mathbf{m}, \mathbf{S}), \quad \text{where } \mathbf{S} = \bigotimes_{d=1}^D \mathbf{L}_d \mathbf{L}_d^\top + \bigotimes_{d=1}^D \mathbf{J}_d \mathbf{J}_d^\top. \quad (73)$$

Noting that we can compute the determinant of such a matrix structure in $\mathcal{O}(DM^3)$ (see e.g. Rakitsch et al., 2013), the computational cost of the method is then dominated by the cost of computing the mean and marginal variances of $q(\mathbf{f})$. For the mean, we first compute $\mathbf{K}_{\mathbf{uu}}^{-1}\mathbf{m}$ which costs $\mathcal{O}(M^D)$ (making use of the low-rank structure of $\mathbf{K}_{\mathbf{uu}}$), and then $\mathbf{K}_{\mathbf{fu}}(\mathbf{K}_{\mathbf{uu}}^{-1}\mathbf{m})$, which costs $\mathcal{O}(NM^D)$. In computing the marginal variances of $q(\mathbf{f})$, the additive structure of \mathbf{S} poses no problem, and the Kronecker structures mean that $\mathbf{K}_{\mathbf{fu}}\mathbf{K}_{\mathbf{uu}}^{-1}\mathbf{S}$ can be computed in $\mathcal{O}(NM^2D)$. Again, it is possible to make use of stochastic optimization to reduce the cost per iteration.

In the MCMC case, the cost is once again dominated by computing the marginals of $q(\mathbf{f} | \mathbf{u})$, which costs $\mathcal{O}(NM^D)$.

5.5 Centring of variables for MCMC

For an effective MCMC scheme, it is necessary to re-centre the variables \mathbf{u} using a square-root of $\mathbf{K}_{\mathbf{uu}}$ (Christensen et al., 2006; Vanhatalo and Vehtari, 2007; Murray and Adams, 2010; Filippone et al., 2013; Hensman et al., 2015a) such that

$$\mathbf{v} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \tag{74}$$

$$\mathbf{u} = \mathbf{K}_{\mathbf{uu}}^{1/2}\mathbf{v}, \tag{75}$$

which gives $\mathbf{u} \sim \mathcal{N}(\mathbf{0}, \mathbf{K}_{\mathbf{uu}})$. For dense $\mathbf{K}_{\mathbf{uu}}$ matrices, the square root $\mathbf{K}_{\mathbf{uu}}^{1/2}$ is often chosen to be the Cholesky factor, but the structured $\mathbf{K}_{\mathbf{uu}}$ described above has a more convenient square root, given by concatenating the two parts into

$$\mathbf{R} = \left[\text{diag}(\boldsymbol{\alpha})^{\frac{1}{2}}, \boldsymbol{\beta} \right] \tag{76}$$

so that $\mathbf{R}\mathbf{R}^\top = \text{diag}(\boldsymbol{\alpha}) + \boldsymbol{\beta}\boldsymbol{\beta}^\top = \mathbf{K}_{\mathbf{uu}}$. Since \mathbf{R} has an extra column (or two or three extra columns for Matérn- $\frac{3}{2}$ and Matérn- $\frac{5}{2}$ respectively) we require an extra variable(s) in the vector \mathbf{v} .

6. Experiments

This Section provides empirical evidence and examples of how the Variational Fourier Features method works in practice. We first cover some illustrative toy examples, which underline the accuracy of the approximation and the influence of the approximation parameters. After this we address three standard test examples for GP regression (for high- and low-dimensional inputs) and classification, where we show the method provides accurate predictions, and does so extremely fast. Finally, we examine how the approximation approaches the posterior in the MCMC case by examining a simple log Gaussian Cox process data set.

6.1 Illustrative examples

Our Variational Fourier Features (VFF) method has two tuning parameters which affect the approximation: the number of Fourier features M , and the choice of the bounds a and b . In this Section the effects related to tuning these parameters are demonstrated.

6.1.1 REGRESSION

To perform GP regression with Variational Fourier Features, one simply takes the expressions given for standard GP approximations (equations (28) and (29)) and substitutes $\mathbf{K}_{\mathbf{uu}}$, $\mathbf{K}_{\mathbf{uf}}$ and $\mathbf{k}_{\mathbf{u}}(x)$ appropriately. Figure 3 compares the VFF approximation with Random Fourier Features on a trivial but illuminating regression experiment.

We started by fitting a ‘full’ GP model (i.e. with inversion of the dense covariance matrix), maximising the likelihood with respect to the kernel’s variance parameter, lengthscale parameter and the noise variance. We then used these fitted values to compare variational and Random Fourier Features.

The Random Fourier Features method has some attractive convergence properties (Rahimi and Recht, 2008), as well as being simple to implement. Here we show empirically that the VFF method approaches the true posterior much faster, in terms of the number of basis functions used.

This experiment uses a Matérn- $\frac{1}{2}$ kernel. The corresponding spectral density is proportional to a Cauchy distribution, which has heavy tails. The frequencies in Random Fourier Features are drawn from this heavy-tailed distribution. We see from Figure 3 that this leads to a poor Monte Carlo estimate of equation (7). With 20 random features, the fit of the model is poor, whilst with 20 variational features, the approximation is qualitatively better. With 100 features, the random features approach still cannot replicate the full GP, whilst the variational method is already exact since the ELBO is tight to the likelihood to within numerical precision. Even with 500 features, the RFF method cannot approximate the model well.

This demonstration is exacerbated by the choice of kernel function. For other kernels in the Matérn series, with spectra that behave more like a Gaussian, the problem is less severe. Nonetheless, this demonstrates the power of Variational Fourier Features for these models.

6.1.2 SETTING a , b , AND M

To use Variational Fourier Features, we must select the projection interval $[a, b]$, as well as the frequencies used. In all our experiments, we use the first M frequencies that are harmonic on the interval, that is $\omega = [\frac{2\pi m}{(b-a)}]_{m=1}^M$.

Figure 4 compares the effect of changing the interval size and the number of inducing frequencies, for a simple Gaussian regression problem. We drew data from a Gaussian process with a Matérn- $\frac{3}{2}$ kernel ($\sigma^2 = 1$, $\ell = 0.2$) and added Gaussian noise ($\sigma_n^2 = 0.05$). We used these known parameters for the series of regression fits shown, varying $[a, b]$ and M . The true marginal likelihood was computed to be -15.99 .

The top row has the interval set too narrow, within the region of the training data (this is valid but not recommended!). Accordingly, the ELBO (marked in the upper left corner of each plot) is low, no matter how many frequencies are used. In the next row, the boundary is still too close to the data: some edge effects are visible. The low value of the ELBO reflects the low quality of the approximation.

In the third row, the boundary $[a, b]$ is sufficiently far from the data. No edge effects are visible, but we notice extra wiggles in the approximation when insufficient frequencies




Figure 3: A comparison of our Variational Fourier Features (VFF) with Random Fourier Features (RFF). Top row: Random Fourier Features with 20, 100, and 500 frequencies. Bottom row: Variational Fourier Features with 20 and 100 frequencies, and the full-GP fit. The kernel is a Matérn- $\frac{1}{2}$, with variance, lengthscale and noise variance set by maximum likelihood according to the full GP. The bounds of the problem are set to $[-1, 2]$ for the VFF.

are used. For sufficient frequencies, and this sensible setting of the interval, the ELBO is close to the marginal likelihood.

The fourth row of Figure 4 shows the effect of choosing an interval that is too large. Because we have set the frequencies to be dependent on $(b - a)$, making this quantity large makes the frequencies too low. Sinusoidal behaviour is visible with $M = 8$ or $M = 16$, although a sensible solution is recovered once $M = 32$. Again, the ELBO reflects the quality of the approximation.

The purpose of this experiment was to illustrate how the properties of the approximation parameters (a, b , and M) affect the approximation. We have emphasized how the ELBO can be examined to ensure that two key considerations are met. First, that a and b are sufficiently far from the edges of the data. Second, that a sufficient number of frequencies are used. We emphasize again the monotonic nature of the approximation inherited from the variational framework: adding more features (increasing M) necessarily improves the approximation in the KL sense.

6.1.3 COMPARISON WITH SPARSE GP FOR VARIOUS INPUT DIMENSIONS

The computational complexity of the VFF approach increases with the dimension of the input space. For additive kernels, we have shown that the complexity increases linearly with input dimension, and for models with a product kernel, the computation increases

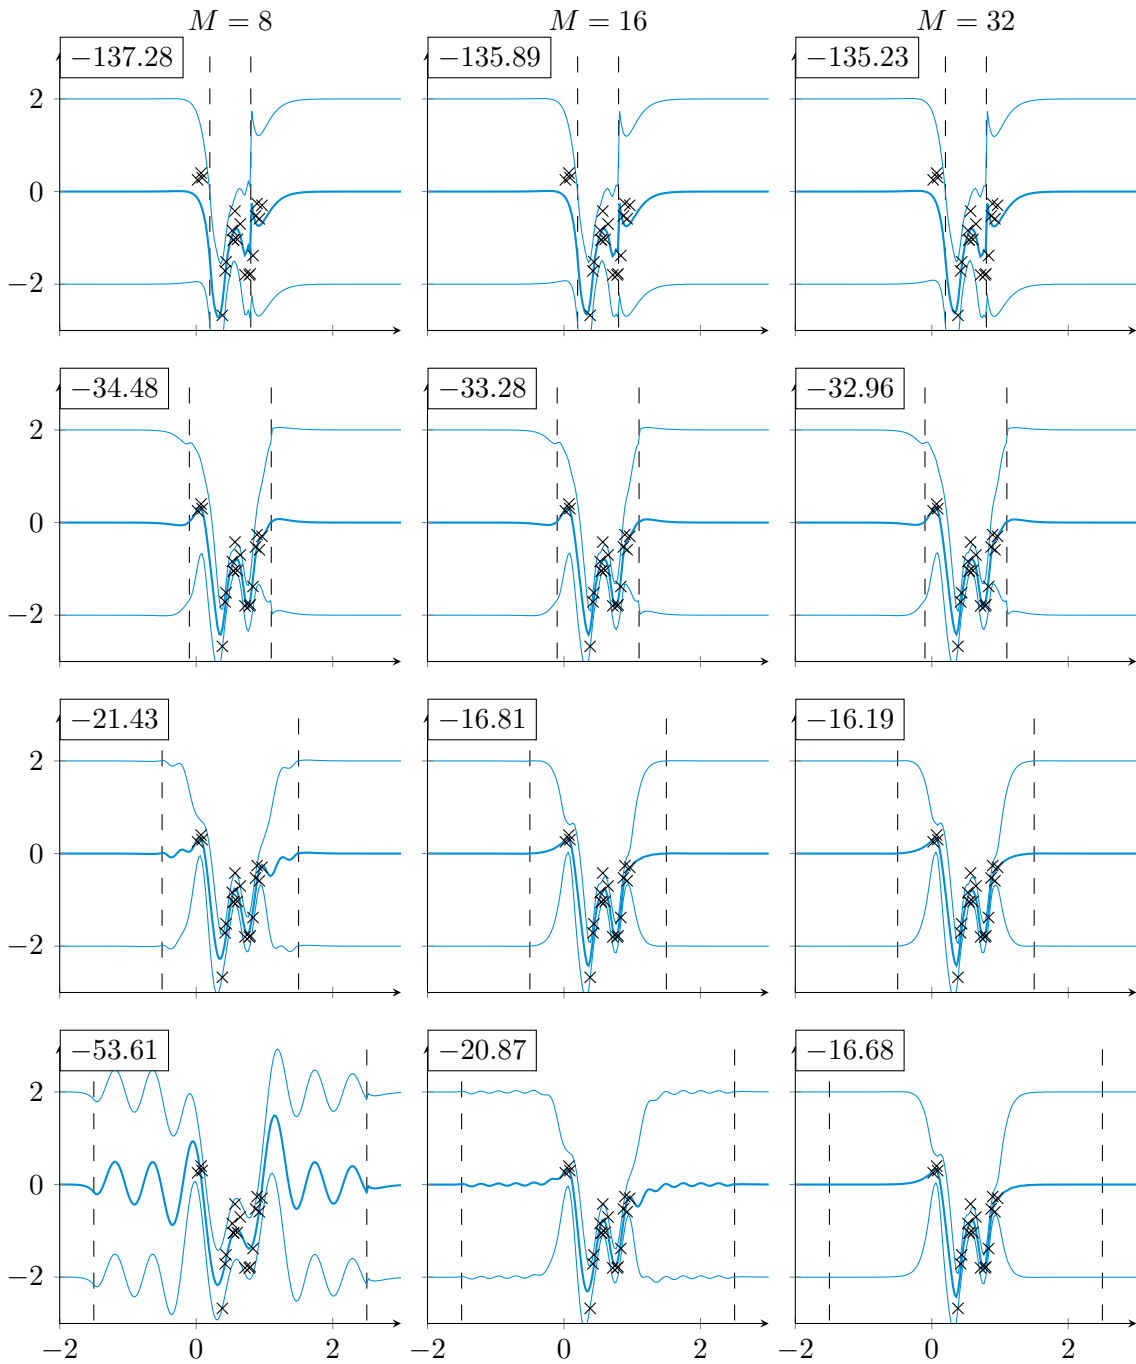


Figure 4: Demonstration of the effect of changing the interval size $[a, b]$ and the number of frequencies used. Each column has an increasing number of frequencies ($M = 8, M = 16, M = 32$), and each row has an increasing interval size, marked by vertical dashed lines. The data are drawn from a Matérn- $\frac{3}{2}$ kernel with Gaussian noise, and inference is performed with the kernel parameters set to the known values. The ELBO is shown in the top left of each plot. For reference, the true marginal likelihood in this problem is -15.99 .

exponentially. In this section we explore the quality of our approximation in the product-kernel case, and compare with the sparse (inducing input) approach. We measure the time/accuracy tradeoff as a function of the input dimension and the number of inducing variables.

We consider toy data sets with input dimension $d = 1 \dots 4$, with input locations drawn uniformly over $[0, 1]^d$. The response data are given by GP samples with products of Matérn- $\frac{3}{2}$ kernels, with additive Gaussian observation noise:

$$\begin{aligned}
 x_{i,j} &\sim \mathcal{U}(0, 1) && \text{for } 1 \leq i \leq 10^4 \text{ and } 1 \leq j \leq d, \\
 \mathbf{f} &\sim \mathcal{N}\left(\mathbf{0}, \prod_{j=1}^d k_{3/2}(\mathbf{X}_j, \mathbf{X}_j)\right), \\
 \mathbf{y} &\sim \mathcal{N}(\mathbf{f}, 0.1 \times \mathbf{I}).
 \end{aligned} \tag{77}$$

Given these data, we compare two approximations: one sparse GP with $M = m^d$ inducing inputs located on a regular grid with minimal and maximal values equal to 0 and 1 for each coordinate and a VFF approximation based on a Kronecker structure with $M = (m - 1)/2$ frequencies for each input dimension which also leads to m^d inducing variables. For both approximations, the parameters of the covariance function were fixed to the known values used to generate the data (lengthscales $\ell = 0.2$, variance $\sigma^2 = 1$). The domain boundaries $[a, b]$ that are specific to the VFF model were fixed to $[-0.3, 1.3]$, which seemed to be a reasonable trade-off for the various dimensions after testing a few other values. We then compute the KL divergence between each approximation and the truth, by comparing the ELBO to the marginal likelihood from a ‘full’ GP regression model. We recorded the time required to compute the ELBO, and repeated the whole experiment five times.

The results are reported in Figure 5. In one dimension, both methods have a similar accuracy in term of KL for a given number of inducing variables but VFF appears to be faster than sparse GPs. However the low complexity of this one dimensional model does not allow us to measure the time accurately enough to distinguish the influence of the number of inducing variables on the execution time. More interesting conclusions can be drawn when the input space is greater than one: for the same number of inducing variables, VFF is significantly faster than the sparse GP (one fewer Cholesky decomposition is required) but the quality of the approximation is not as good as the sparse GP for the same number of basis functions.

For a given computational time, both methods have an equivalent accuracy but with a different number of inducing variables. The VFF basis functions appear to have less representational power, but their structure allows faster computations. Our understanding is that this smaller representational power of VFF can be explained by the fact that the inducing functions need to account for variations on $[a, b]^d$ whereas the sparse GP basis functions only need account for that on $[0, 1]^d$: a much smaller volume, especially in high dimensions. Note that the accuracy of the sparse GP can be improved by optimization of the location of the inducing inputs but that would result in computation times many times larger.

We have shown empirically that in the worst case, VFF is comparable to the sparse GP in terms of the available compute-time/accuracy tradeoff. For other models including

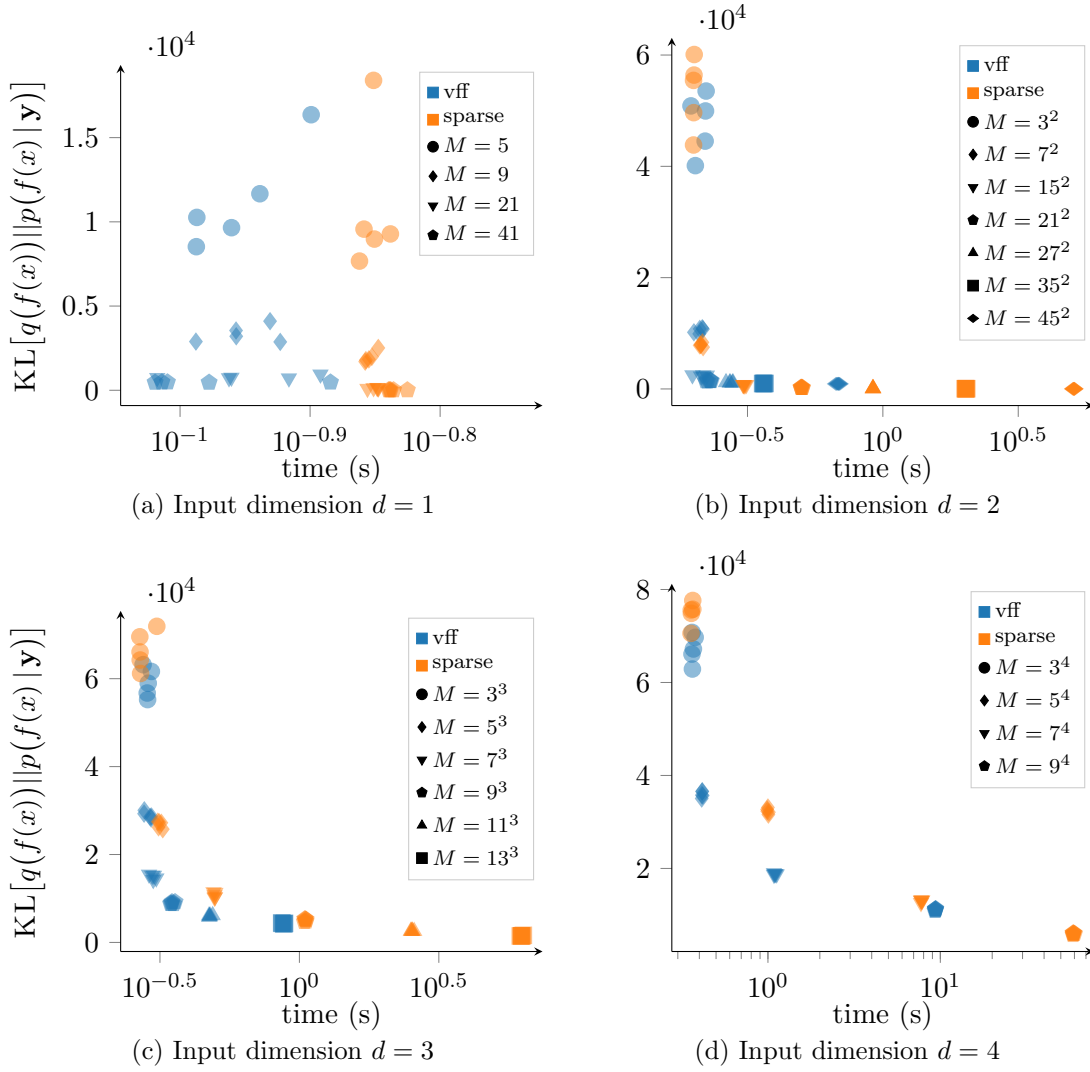


Figure 5: KL divergence versus wall clock time for sparse GPs and VFF with various number M of inducing inputs. The data is obtained by sampling a GP with Matérn- $\frac{3}{2}$ kernel ($\sigma^2 = 1$, $\ell = 0.2$) at 10^4 uniformly distributed input locations and $\mathcal{N}(0, 0.1)$ observation noise. The KL is computed with respect to the full GP regression model with known parameters for all models. The experiments are replicated five times with different draws of \mathbf{X} and \mathbf{y} .

additive models, we expect much stronger performance from the VFF approximation, as the following sections demonstrate.

6.2 Additive modelling of airline delays

We demonstrate the feasibility of the variational spectral approximation in a high-dimensional GP regression example for predicting airline delays. The US flight delay prediction example (see Hensman et al., 2013, for the original example) has reached the status of a standard test data set in Gaussian process regression, partly because of its large-scale non-stationary nature and partly because of the massive size of the data set with nearly 6 million records.

This example has recently been used by Deisenroth and Ng (2015), where it was solved using distributed Gaussian processes. Samo and Roberts (2016) use this example for demonstrating the computational efficiency of string Gaussian processes. Furthermore, Adam et al. (2016) bring this problem forward as an example of a data set, where the model can be formed by the addition of multiple underlying components.

The data set consists of flight arrival and departure times for every commercial flight in the USA for the year 2008. Each record is complemented with details on the flight and the aircraft. The eight covariates \mathbf{x} are the same as in Hensman et al. (2013), namely the age of the aircraft (number of years since deployment), route distance, airtime, departure time, arrival time, day of the week, day of the month, and month. We predict the delay of the aircraft at landing (in minutes), y .

Following the proposed solution by Adam et al. (2016) for this estimation problem, we use a Gaussian process regression model with a prior covariance structure given as a sum of Matérn- $\frac{3}{2}$ covariance functions for each input dimension, and assume the observations to be corrupted by independent Gaussian noise, $\varepsilon_i \sim \mathcal{N}(0, \sigma_n^2)$. The model is

$$f(\mathbf{x}) \sim \mathcal{GP} \left(0, \sum_{d=1}^8 k(x_d, x'_d) \right), \quad (78)$$

$$y_i = f(\mathbf{x}_i) + \varepsilon_i, \quad (79)$$

for $i = 1, 2, \dots, N$. For the Variational Fourier Feature method (VFF), we used $M = 30$ frequencies per input dimension.

We consider the entire data set of 5.93 million records. To provide comparable predictions, we follow the test setup in the paper by Samo and Roberts (2016), where results were calculated using String GPs (their method), the Bayesian committee machine (BCM, Tresp, 2000), the robust Bayesian committee machine (rBCM, Deisenroth and Ng, 2015), and stochastic variational GP inference (SVIGP, Hensman et al., 2013). The SVIGP used a squared-exponential kernel with one lengthscale per dimension. We repeated the SVIGP experiments of Samo and Roberts (2016) to complete the table for larger data sets and predictive densities.

Predictions are made for several subset sizes of the data, each selected uniformly at random: 10,000, 100,000, 1,000,000, and 5,929,413 (all data). In each case, two thirds of the data is used for training and one third for testing. For each subset size, the training is repeated ten times, and for each run, the outputs are normalized by subtracting the training sample mean from the outputs and dividing the result by the sample standard deviation. In the VFF method, the inputs were normalized to $[0, 1]$ and the domain size was set to $[-2, 3]$, i.e. the boundary is at a distance of two times the range of the data for each dimension.

N	10,000		100,000		1,000,000		5,929,413	
	MSE	NLPD	MSE	NLPD	MSE	NLPD	MSE	NLPD
VFF	0.89 ± 0.15	1.362 ± 0.091	0.82 ± 0.05	1.319 ± 0.030	0.83 ± 0.01	1.326 ± 0.008	0.827 ± 0.004	1.324 ± 0.003
Full-RBF	0.89 ± 0.16	1.349 ± 0.098	N/A	N/A	N/A	N/A	N/A	N/A
Full-additive	0.89 ± 0.16	1.362 ± 0.096	N/A	N/A	N/A	N/A	N/A	N/A
SVIGP	0.89 ± 0.16	1.354 ± 0.096	0.79 ± 0.05	1.299 ± 0.033	0.79 ± 0.01	1.301 ± 0.009	0.791 ± 0.005	1.300 ± 0.003
String GP [†]	1.03 ± 0.10	N/A	0.93 ± 0.03	N/A	0.93 ± 0.01	N/A	0.90 ± 0.01	N/A
rBCM [†]	1.06 ± 0.10	N/A	1.04 ± 0.04	N/A	N/A	N/A	N/A	N/A

Table 2: Predictive mean squared errors (MSEs) and negative log predictive densities (NLPDs) with one standard deviation on the airline arrival delays experiment. [†]MSE results as reported by Samo and Roberts (2016), who also list results for the BCM and rBCM methods.

Table 2 shows the predictive mean squared errors (MSEs) and the negative log predictive densities (NLPDs) with one standard deviation on the airline arrival delays experiment. In this form the MSEs over the normalized data can be interpreted as a fraction of the sample variance of airline arrival delays. Thus a MSE of 1.00 is an accuracy equivalent to using the training mean as predictor. The results for the String GP method are included for reference only and they are given as reported by Samo and Roberts (2016). They also list results for the rBCM method, which performed worse than the String GP.

Running the VFF experiment with all 5.93 million data using our Python implementation took 265 ± 6 seconds (626 ± 11 s CPU time) on a two-core MacBook Pro laptop (with all calculation done on the CPU). The results for 10,000 data were calculated in 21 ± 2 seconds (27 ± 4 s CPU time). For comparison, according to Samo and Roberts (2016), running the String GP experiment took 91.0 hours total CPU time (or 15 h of wall-clock time on an 8-core machine). Even the SVIGP method required 5.1 ± 0.1 hours of computing (27.0 ± 0.8 h CPU time) on a cluster.

For comparison, we have also included the results for the naive full GP solution for 10,000 data. These results were computed on a computing cluster and the computation time was several hours per repetition. We show results both for a squared-exponential kernel with one lengthscale per dimension, and also with an additive kernel similar to the one defined for the VFF method. As one might expect, the RBF model had a slightly lower average MSE (0.8898) than the additive model (0.89274). The difference is small: only 3 seconds of RMSE in the original units of flight-delay. The VFF result (0.8934) for the same data is close to the full GP solution (the difference being only about 0.66 s in terms of non-normalized RMSE). In terms of NLPD the results agree as well; the full RBF model is able to capture the phenomenon slightly better than the full additive model, and the VFF model approximates well the full additive model. The variability in the data is large, and including more data in the GP training gives the model more power for prediction. The SVIGP model (with a squared-exponential kernel) is the most flexible one and reaches the lowest RMSE and NLPD in the experiments.

Figure 6 shows the component-wise predictions for each of the covariates, thus summarizing the underlying structure of the delay prediction model. As one might expect, the effect of the month, day of month, and day of week are small. The only interpretable effects are seen in slightly higher delays in the summer months and around the holiday season

towards the end of the year, as well as some effects for the weekend traffic on Friday and Sunday. The age of the aircraft has a barely noticeable effect. The four remaining inputs (e–h) explain most of the delay. The effects of airtime and travel distance are strongly correlated. Strong correlation can also be seen in departure and arrival time effects, both of which show peaks during the night. The model also catches the periodic nature of the departure and arrival time effects. These results are in line with the analysis of Adam et al. (2016).

6.3 Classification with a Gaussian approximation

To perform variational classification using Variational Fourier Features, we can use a Gaussian approximation to the posterior $q(\mathbf{u})$ and optimize the ELBO with respect to the mean and variance of the approximation, as described in Section 2.4.4. To illustrate, we replicate the experiment on the banana data set from Hensman et al. (2015b). In that work, the authors showed that a variational inducing point method performed better than competing alternatives because the positions of the inducing points could be effectively optimized using the ELBO as the objective. In this work, we have selected the frequencies of the Variational Fourier Features to be placed on a regular grid: the only aspect of the Fourier features to be tuned is M , controlling the number of features used.

Since the variational framework provides a guarantee that more inducing variables must be monotonically better (Titsias, 2009), we expect that increasing the number of Variational Fourier Features will provide improved approximations. However, if we restrict the form of the Gaussian approximation by requiring a structured covariance, this may no longer be true. In practice, we have found that a freely-structured covariance matrix works well in one dimension and for additive models, but for Kronecker-structured models, optimization over an exponentially large matrix is intolerably slow. A suggestion for a related method provided by Nickson et al. (2015) is to force the Gaussian approximation to have a Kronecker-structured covariance matrix. This reduces the number of variables over which we have to optimize, but removes the guarantee that more inducing variables is always better. We find in practice that this structure does not work well: Figure 7 shows the ELBO as a function of the number of inducing variables M , for both a trained covariance and a Kronecker-structured one. We see that for the Kronecker structure, the ELBO decreases as M increases, implying that more inducing variables are actually making the approximation worse in this case.

Our suggestion is to use a sum of two Kronecker-structured matrices. The right-most plot of Figure 7 shows that the ELBO increases monotonically with M when we use this structure. We provide no guarantee that the sum of two Kronecker matrices is optimal, and suggest that future work might consider generalizing the form of this approximation.

To compare our method with the inducing-point method, we refer to Figure 8, which shows the results of fitting a Gaussian process classifier to the banana data set using increasing numbers of inducing points (IP) and increasing numbers of Fourier features (VFF). We observe the effect noted in Hensman et al. (2015b) that the inducing points move toward the decision boundary, though some slight differences from that experiment exist because of our choice of the product of Matérn- $\frac{5}{2}$ kernels, instead of the squared exponential kernel. The Figure also shows that the number of frequencies required to obtain a good approximation

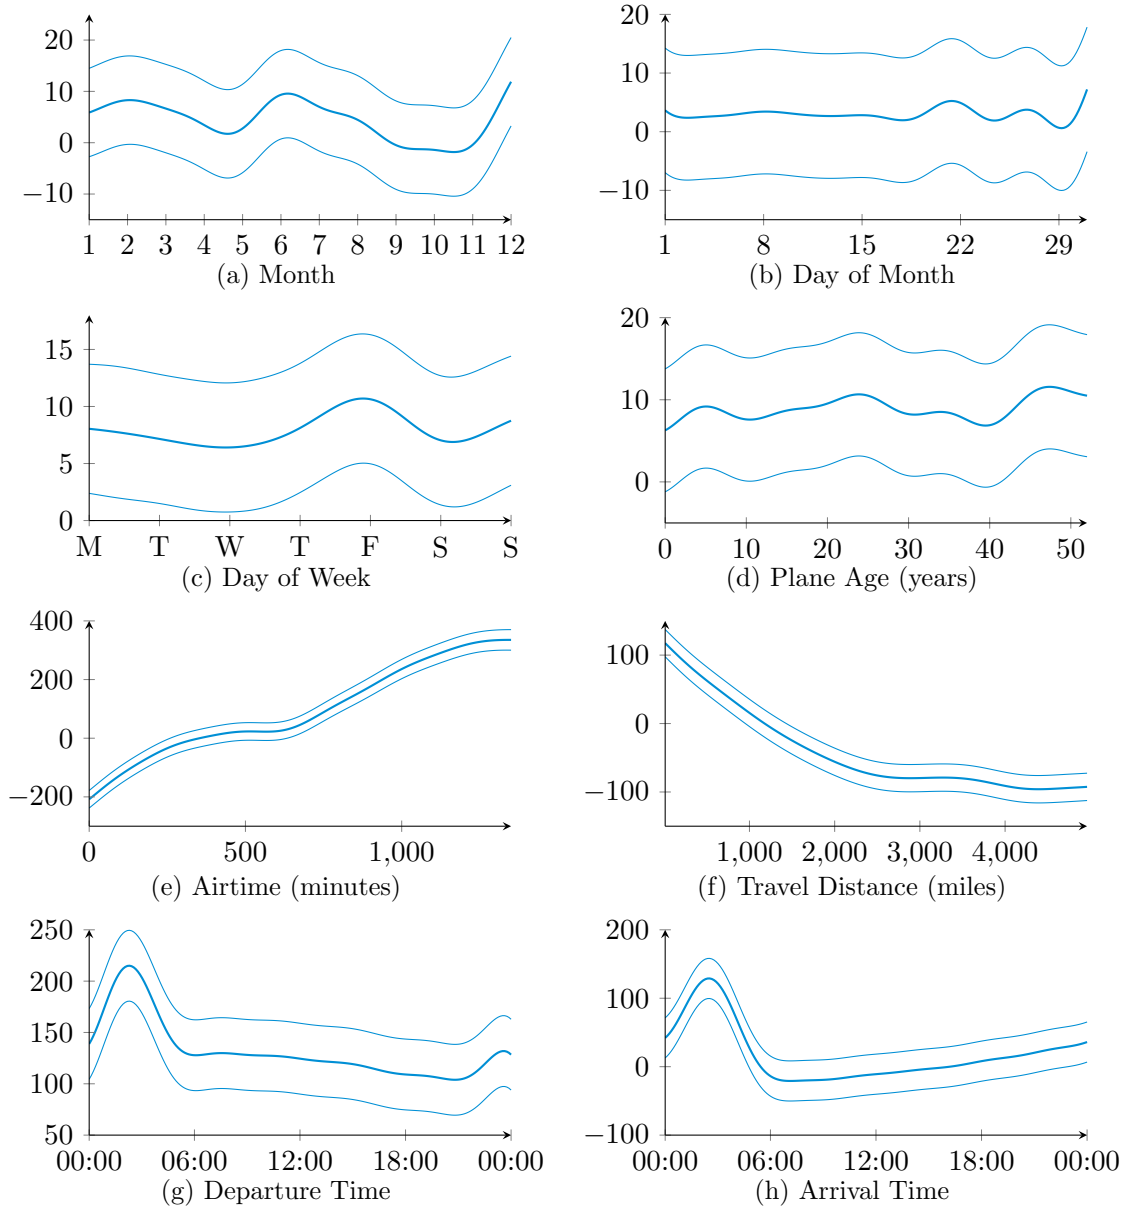


Figure 6: The posterior mean and two posterior standard deviations in the airline delay prediction experiment estimated from 5.93 million data. Each panel show the posterior over the effect from one covariate under an additive model. The vertical axis represents delay in minutes in each case.

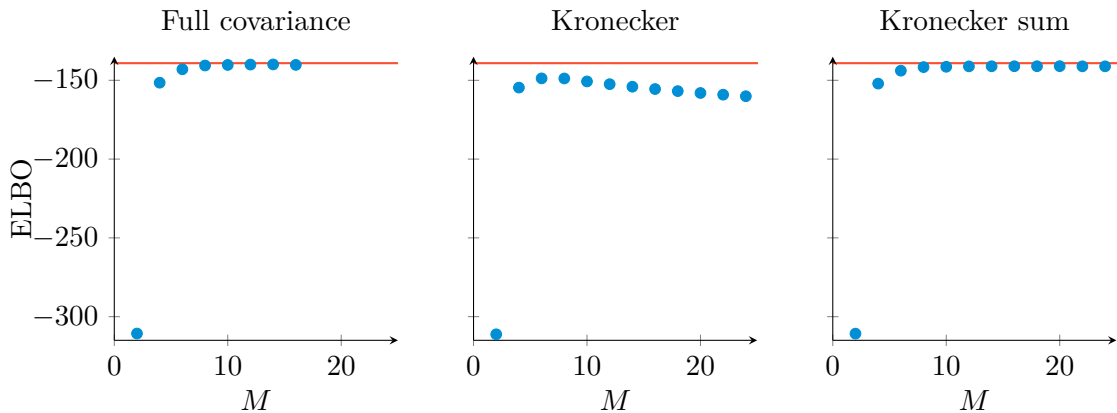


Figure 7: The increasing ELBO for the banana data set with increasing numbers of frequencies M , and a Gaussian approximation to the posterior. Left: $q(\mathbf{u})$ has an unconstrained covariance matrix. Middle: $q(\mathbf{u})$ has a Kronecker-structured covariance matrix. Right: $q(\mathbf{u})$ has a covariance matrix with a sum of two Kronecker-structured matrices. In each case the ELBO achieved using a full matrix decomposition is shown as a horizontal red line.

is reasonable, though it is not straightforward to compare directly between the VFF and IP methods. The total number of basis functions used in the Fourier features is twice the number of frequencies (M) listed, squared (since we have real and imaginary components, and take the Kronecker product). Still, the VFF method requires fewer computations per frequency than the IIP method does per inducing point, since the $\mathbf{K}_{\mathbf{u}\mathbf{u}}$ matrix is easily decomposed.

6.4 Solar irradiance

An alternative, but related idea to our proposed Variational Fourier Features approach is presented by Gal and Turner (2015). In that work, the authors used a Random Fourier Features model (as in equation (15)), but additionally performed a variational-Bayesian treatment of the frequencies. We repeat an experiment from that work on solar irradiance data (Lean, 2004), shown in Figure 9. In this experiment, the time series was normalized and segments of the data were removed; in Figure 9, removed data are shown in red, and vertical dashed lines denote the removed sections. The remaining data are shown in black, and we show the predicted mean (heavy blue line) and 95% confidence intervals (light blue lines). We fitted five models (or model approximations) to the remaining data, using 50 inducing points (or frequencies) except in the case of the Random Fourier Features, where we used 500 (following Gal and Turner, 2015). In each case, we used a Matérn- $\frac{5}{2}$ kernel initialized with a lengthscale of 10, and optimized the parameters with respect to the marginal likelihood (or ELBO, accordingly).

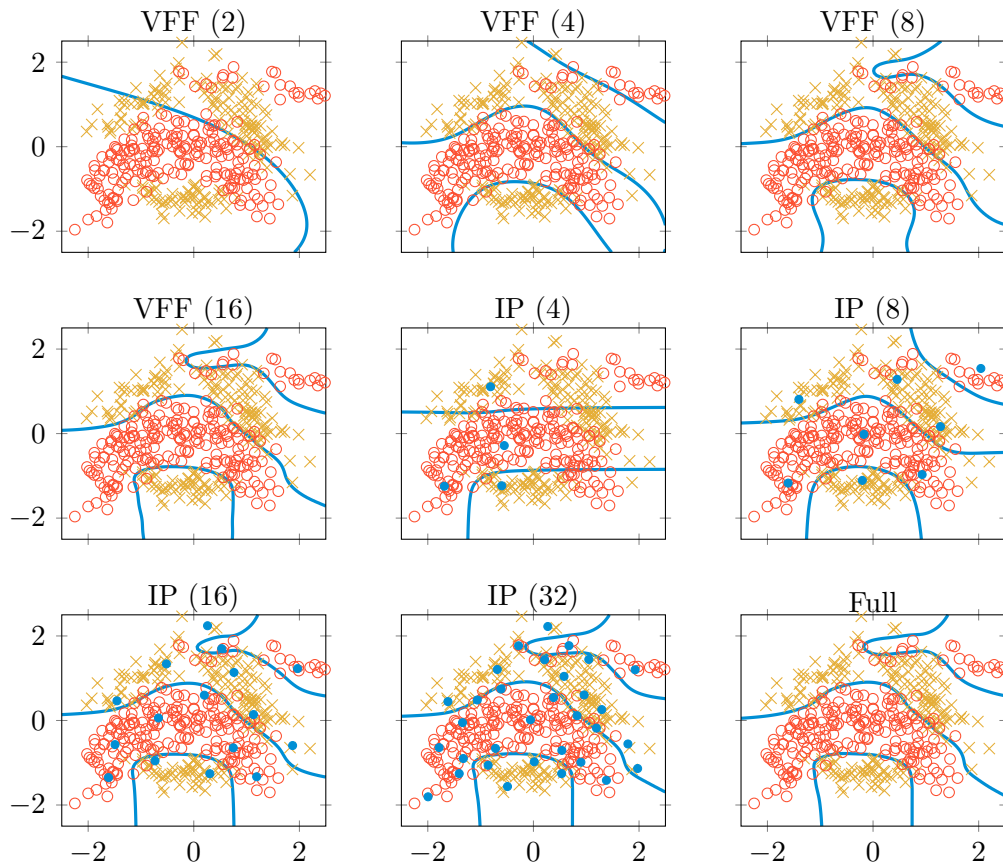


Figure 8: Classification of the banana data set with increasing number of Variational Fourier Features (VFF) and Inducing Points (IP). The two classes of data are shown as red circles and yellow crosses, and the decision boundary is shown as a solid blue line. An approximation with a full inversion of the covariance matrix is shown in the last plot. The blue dots show optimized positions of the inducing points. The Variational Fourier Features approach the true posterior more rapidly than the inducing point method.

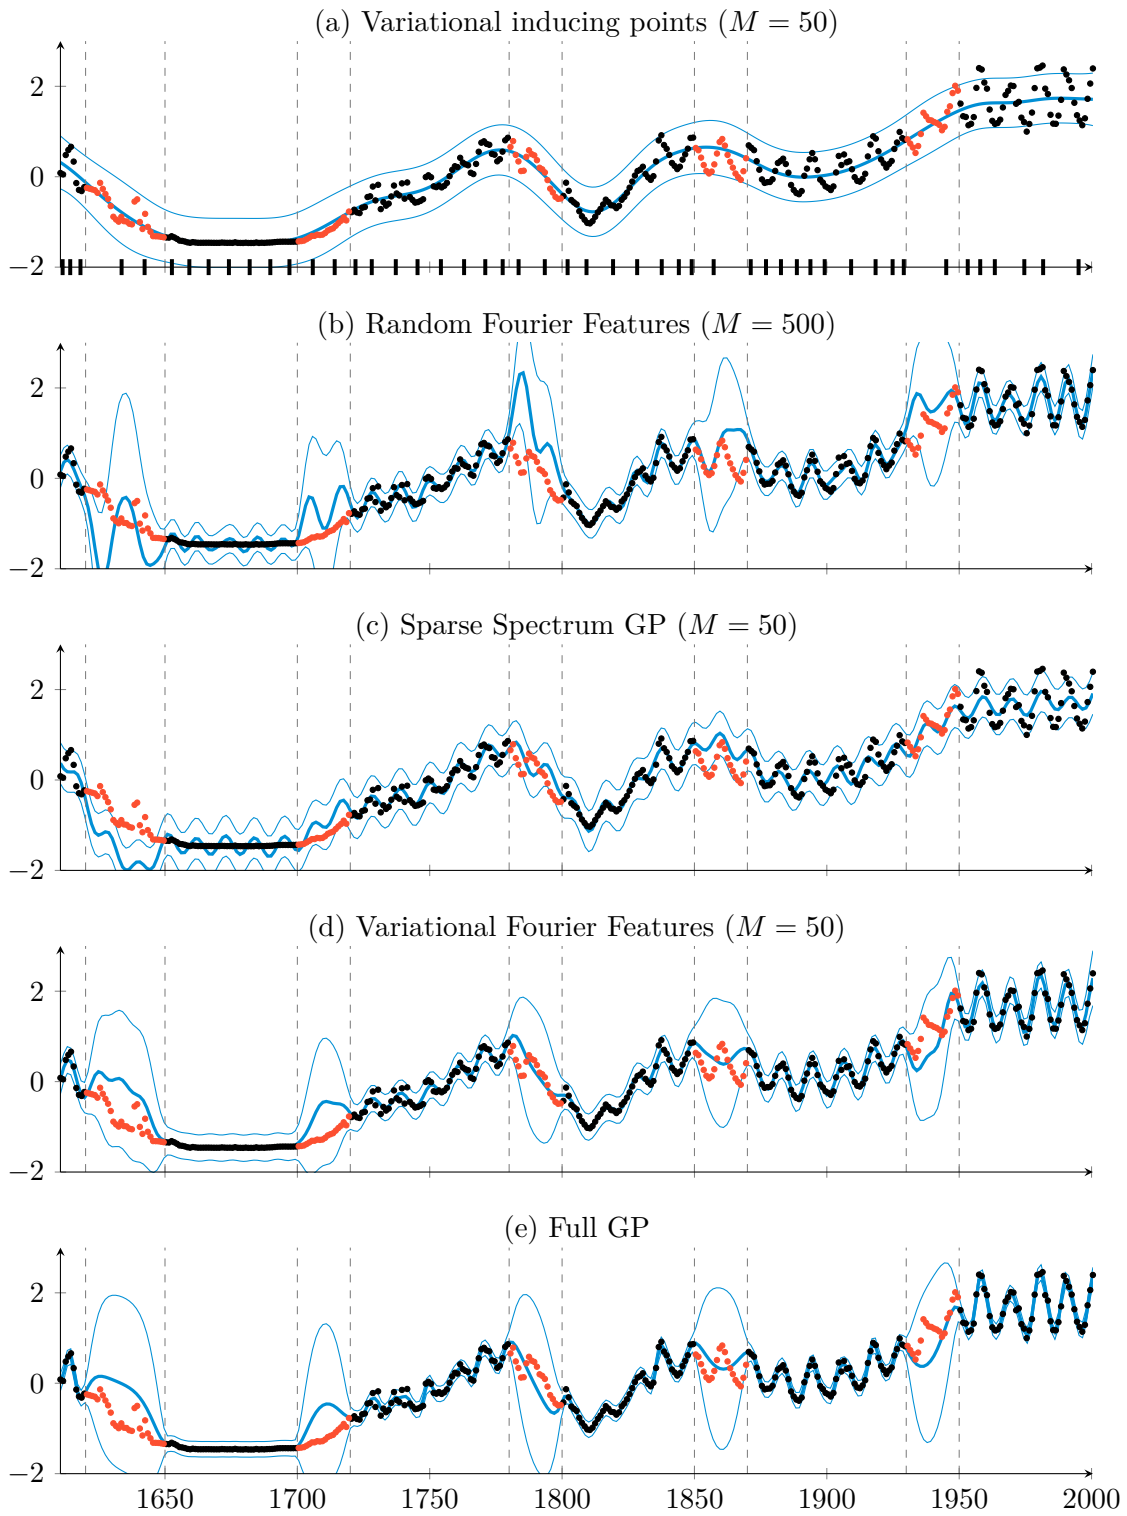


Figure 9: Replicating the experiment on solar irradiance data from Gal and Turner (2015).

Figure 9(a) shows the results of an inducing point method (still within the variational framework, based on Titsias, 2009). The positions of the inducing points z are shown as black marks on the horizontal axis. In this case, there are insufficient points to capture the complexity of the data, and the model reverts to a long lengthscale. This is an interesting case of underfitting, where the variational approximation to the function values gives rise to a bias in the estimate of the lengthscale parameter. This bias occurs because when maximising the ELBO with respect to the parameters, the KL divergence is smaller for longer lengthscales, biasing the parameter estimation in that direction. Turner and Sahani (2011) give an excellent discussion of this issue.

Figure 9(b) shows the results of the Random Fourier Features method. This method is given 500 frequencies; we concur with the findings of Gal and Turner (2015) that using only 50 frequencies gives very poor performance. Where the data are missing, the error bars provided by the model are large, similarly to the full GP (Figure 9(e)). The approximation is unable to recover the correct mean function, however, and is somewhat subject to overfitting: the model emphasises the high-frequency content in the data. In some sense, this is a desirable property as the model is more flexible than the original GP. For example, Yang et al. (2015) exploited this property to fit complex GP models. However, we advocate the variational method precisely because of the asymptotic convergence (in M) to the true posterior distribution of the original model.

Figure 9(c) shows the result of fitting the Sparse Spectrum GP approximation (Lázaro-Gredilla et al., 2010). This model is similar to the Random Fourier Features model, except that the frequencies are optimized. Here the overfitting is more severe: the model focusses entirely on the high-frequency content of the data. Although this gives good predictions in some cases (the segment around 1930 for example), the predictions are wild in others (e.g. around 1630, where none of the held-out data fall within the 95% intervals).

Figure 9(d) and (e) show the Variational Fourier Features method and the full GP method. The VFF approximation strives to approach the true posterior, and it is clear that the method is superior to the inducing points method in (a) in that the lengthscale is well estimated and the confidence intervals match the behaviour of the full GP. Some small discrepancies exist, but we note that these disappear completely at 100 inducing frequencies.

In summary, this experiment emphasises again how the variational approach attempts to approximate the posterior of the model, whereas previous approximations have *changed* the model, giving different behaviours. Within the variational methods, the Variational Fourier Features outperform the inducing point approach in this case, since our orthogonal basis functions are capable of capturing the posterior more effectively.

6.5 MCMC approximations

Above we have used the Gaussian approximation to the posterior process. Hensman et al. (2015a) propose a methodology which combines the variational approach with MCMC by sampling from the optimal approximating distribution, as we described in equation (27). The advantages of this approach are that the posterior need not be assumed Gaussian, and that the covariance parameters can also be treated in a Bayesian fashion. The VFF framework that we have described fits well into this scheme, and as we have already described,

the computational cost per iteration is linear in the number of data and in the number of required frequencies.

Here we replicate an experiment from Hensman et al. (2015a) (therein based on inducing point representations) in estimating the rate of a two-dimensional log Gaussian Cox process. In this experiment, we aim to emphasise the convergent properties of Variational Fourier Features, and show how they might be used in practice within the MCMC framework.

The data consist of a series of 127 points collected in a square area, denoting the location of pine saplings. This simple illustration was also used by Møller et al. (e.g. 1998). The model is

$$f(\mathbf{s}) \sim \mathcal{GP}(0, k(\mathbf{s}, \mathbf{s}')), \quad (80)$$

$$\log \lambda(\mathbf{s}) = f(\mathbf{s}) + c, \quad (81)$$

$$\mathbf{y} \sim \mathcal{PP}(\lambda), \quad (82)$$

where \mathcal{PP} denotes an inhomogeneous Poisson process with intensity λ , \mathbf{y} are spatial locations representing a draw from that Poisson process, k is a product of Matérn- $\frac{3}{2}$ covariance functions of the two-dimensional space indexed by \mathbf{s} , and c is a parameter to be inferred alongside the covariance function parameters. The model contains a doubly-intractable likelihood, which can be approximated by gridding the space and assuming that the intensity is constant across the bin size. The resulting approximation to the likelihood is

$$p(\mathbf{y} | \mathbf{f}) = \prod_{i=1}^G \mathcal{P}(N_i | \lambda(\mathbf{s}_i)\Delta), \quad (83)$$

where G is the number of grid cells, \mathbf{s}_i represent the centres of those cells, N_i is the number of data in the i^{th} cell, Δ is the area of a grid cell and \mathcal{P} represents the Poisson distribution. To perform inference in this model, we decompose the GP with Variational Fourier features, use the rotation described in Section 5.5 to centre the \mathbf{u} variables and run Hamiltonian Monte Carlo (HMC) to jointly sample the centred variables alongside the covariance function parameters and c .

To illustrate how the VFF method approaches the posterior in the MCMC case, we ran the HMC algorithm using $M = 14, 16, \dots, 30$ frequencies. We selected the boundaries of the decomposition $[a, b]$ to be $[-1, 2]^2$, whilst the data were normalized to the unit square $[0, 1]^2$. The posterior mean intensity of the process on the unit square, along with the data, are shown in Figure 10. The figure shows that as M increases, the behaviour of the method stabilizes: there is little to distinguish $M = 30$ from $M = 28$. If M is too low, then there is a bias towards longer lengthscales, but this is alleviated as M increases. Since more inducing variables must always move the approximation toward the posterior in the KL sense (Titsias, 2009), once the optimal posterior (for each fixed M) remains unchanging with an increase in M , we must be approximating the posterior closely. A practical approach then, is to use some reasonable number for M (say, 100) and then ensure that the sampled distribution is unmoving with increased M . Since our experiments suggest that the lengthscale parameter is particularly sensitive to bias due to M being too low, we suggest examining the convergent behaviour of the lengthscale posterior. Figure 11 shows how the posterior distribution of the lengthscales converges as M increases.




Figure 10: The posterior mean of the process intensity with increasing M (denoted above each figure). The spatial locations of the data are shown in the last pane. As M increases, the behaviour of the method stabilizes: we conclude that at $M = 30$ there is no benefit from increasing M further. If M is too low, the method is biased towards a longer lengthscale (see also Figure 11).




Figure 11: The posterior distribution of the two lengthscale parameters for the pines problem (see also Figure 10), estimated using kernel density smoothing on the MCMC trace. As more frequencies are used, the posterior distribution converges.

We note that the regular-lattice nature of this problem means that it can also be efficiently solved using fast Fourier transform (FFT) methods (e.g. Taylor and Diggle, 2014). However, our approach offers several advantages. First, in our method the number of frequencies used is a variational parameter: we have discussed how to select M above. Since the dimension of the vector which is subject to MCMC depends on the number of frequencies, it is desirable to decouple it from the grid density: fine grids are desired to resolve the spatial nature of the process, whilst the length of the vector \mathbf{u} is desired to be short for computational reasons. Second, our approach does not require the embedding of the kernel matrix into a circulant matrix which can occasionally suffer from non-positive-definiteness. Third and most importantly, our method is not restricted to problems on a regular grid.

7. Discussion and future directions

In this paper we have presented a method which we refer to as the Variational Fourier Features (VFF) approach to Gaussian process modelling. It combines the variational approach to sparse approximations in Gaussian processes with a representation based on the spectrum of the process. This approach inherits its appealing approximative construction from the variational approach, but with the representative power and computational scalability of spectral representations.

We generalized the method to several dimensions in cases where the kernel exhibits special structure in the form of additive or separable (product) covariance functions. This choice made it possible to apply the scheme to higher-dimensional test problems in the experiments, for which we demonstrate good performance both in terms of computational speed and predictive power.

Our example application of predicting airline delays suggested that the additive structure worked well for that problem: for data subsets where the full model could be fitted, there was little performance difference between the additive and isotropic case, and indeed the VFF method was competitive with both, suggesting that the posterior was very well approximated. It seems implausible that this structure could work well for all problems, and clearly it will not work well where interesting interactions occur between covariates.

The method is capable of modelling interactions between covariates using separable kernel structures. As was noted in the text, the computational scalability in this case limits the method to low-dimensional problems, because the number of inducing frequencies scales exponentially with the input dimension. In this sense this approach is not a sparse approximation, but rather based on a ‘dense’ basis function projection. Like related approximation schemes (Solin and Särkkä, 2014; Wilson and Nickisch, 2015; Nickson et al., 2015), VFF scales well as the number of data increases, with the drawback of difficulty in dealing with high-dimensional problems out-of-the-box.

To further improve the computational scaling in high-dimensional problems, there are several approaches which could apply. Wilson et al. (2015) presents a projection approach (based on the ideas in Wilson and Nickisch, 2015). These ideas could be expanded to our approach as well, and we might expect the variational properties of the approximation to be useful in this case. Alternatively, replacing the dense grid spectrum with a set of sparsely well-chosen spectral points is an option that can — at least for models or data with suitable structure — provide computational remedy in high input dimensions. Future work might incorporate some ideas from the sparse spectrum method (Lázaro-Gredilla et al., 2010) where the frequencies are optimized, though we might hope that the variational formulation could prevent over-fitting.

We note that one of the key computations in our method, the multiplication $\mathbf{K}_{\text{uf}}\mathbf{y}$, is precisely a non-uniform FFT (NUFFT, see Dutt and Rokhlin, 1993; Greengard and Lee, 2004; Bernstein et al., 2004), which can be approximated in $\mathcal{O}(N \log N)$ operations. Although we have so far avoided using this technology in our implementations, some initial experiments suggest that further speed-up is possible at the cost of a very small loss in accuracy. In the special case where the data lie on a regular grid, it is possible to compute the product $\mathbf{K}_{\text{uf}}\mathbf{y}$ exactly in $\mathcal{O}(N \log N)$ operations using an FFT. Again we have not exploited this in our presentation so far, and future work may consider the relations between VFF and FFT methods based on the circulant embedding trick (Taylor and Diggle, 2014; Turner, 2010).

A further limitation of our current presentation is that we have only considered Matérn covariance functions with half-integer orders. A productive future direction will be to expand the number of kernels which can be decomposed using Variational Fourier Features. An interesting class of kernels may be the spectral mixture kernel (Wilson and Adams, 2013). Although that work was based on frequency-shifted versions of the RBF covariance, we anticipate that frequency-shifted Matérn kernels would work just as well, and may be more amenable to the VFF framework.

In our derivations, we have considered uniform input densities on some bounded interval $[a, b]$. Instead of considering such compact bounded subsets of \mathbb{R} , it might be possible to extend our methodology to more general input densities (cf. Williams and Seeger, 2000). In special cases this can lead to convenient approximations for the eigen-decomposition of the

kernel. For example, for a RBF covariance function and a Gaussian input density, the eigenfunctions can be given in closed-form in terms of Hermite polynomials (see Williams and Rasmussen, 2006, Ch. 4.3). There is a connection between harmonic Fourier approximations like ours and the covariance operator defined through the covariance function and the input density. The connection goes back to Sturm-Liouville theory and is further discussed in Solin and Särkkä (2014).

Finally, we note that VFF may be particularly well-suited to machine learning methods where Gaussian processes are embedded in some further model structure, such as the Gaussian process latent variable model (GPLVM, Lawrence, 2005) and its variational-Bayesian variant (Titsias and Lawrence, 2010), as well as deep Gaussian processes (Damianou and Lawrence, 2013). As described by Damianou et al. (2016), variational inference in these models involves propagating uncertainty by convolving the $\mathbf{K}_{\mathbf{u}f}$ matrix with Gaussian approximations to the input distribution. Since the VFF method has precisely sinusoidal $\mathbf{K}_{\mathbf{u}f}$ matrices, this convolution will take a simple form, and we envisage that combining the VFF framework with variational uncertainty propagation will lead to substantial improvements in inference for these models.

Acknowledgements

Part of this research was conducted during the invitation of J. Hensman at Mines Saint-Étienne. This visit was funded by the Chair in Applied Mathematics OQUAIDO, which gather partners in technological research (BRGM, CEA, IFPEN, IRSN, Safran, Storengy) and academia (Ecole Centrale de Lyon, Mines Saint-Étienne, University of Grenoble, University of Nice, University of Toulouse) on the topic of advanced methods for computer experiments. J. Hensman gratefully acknowledges a fellowship from the Medical Research Council UK. A. Solin acknowledges the Academy of Finland grant 308640. We acknowledge the computational resources provided by the Aalto Science-IT project. J. Hensman would like to thank T. Smith for insightful discussions. The authors would like to thank T. Bui, A. Vehtari, S.T. John and the anonymous reviewers who helped to improve this manuscript.

Appendix A. L_2 Projection on Fourier Features for Matérn- $\frac{1}{2}$

In this Appendix we derive expressions for Fourier features using the L_2 norm.

A.1 Covariance between inducing variable and GP

$$\text{cov}(u_m, f(x)) = \mathbb{E}[u_m f(x)] = \int_a^b \mathbb{E}[f(t)f(x)]e^{-i\omega_m(t-a)} dt = \int_a^b k(t, x)e^{-i\omega_m(t-a)} dt \quad (84)$$

$$= e^{-i\omega_m a} \sigma^2 \left(\int_a^x e^{\lambda(t-x)+i\omega_m t} dt + \int_x^b e^{\lambda(x-t)+i\omega_m t} dt \right) \quad (85)$$

$$= e^{-i\omega_m a} \sigma^2 \left(\left[\frac{e^{\lambda(t-x)+i\omega_m t}}{i\omega_m + \lambda} \right]_a^x + \left[\frac{e^{\lambda(x-t)+i\omega_m t}}{i\omega_m - \lambda} \right]_x^b \right) \quad (86)$$

$$= e^{-i\omega_m a} \sigma^2 \left(\frac{e^{i\omega_m x} - e^{\lambda(a-x)+i\omega_m a}}{i\omega_m + \lambda} + \frac{e^{\lambda(x-b)+i\omega_m b} - e^{i\omega_m x}}{i\omega_m - \lambda} \right) \quad (87)$$

$$= \sigma^2 \left(\frac{e^{i\omega_m(x-a)} - e^{\lambda(a-x)}}{i\omega_m + \lambda} + \frac{e^{\lambda(x-b)} - e^{i\omega_m(x-a)}}{i\omega_m - \lambda} \right) \quad (88)$$

$$= \frac{2\sigma^2\lambda}{\lambda^2 + \omega_m^2} e^{i\omega_m(x-a)} + \frac{\sigma^2}{\lambda^2 + \omega_m^2} \left(\lambda[-e^{a-x} - e^{x-b}] + i\omega_m[e^{a-x} - e^{x-b}] \right) \quad (89)$$

$$= s_{1/2}(\omega_m) e^{i\omega_m(x-a)} + s_{1/2}(\omega_m) \frac{1}{2\lambda} \left(\lambda[-e^{a-x} - e^{x-b}] + i\omega_m[e^{a-x} - e^{x-b}] \right). \quad (90)$$

A.2 Covariance between inducing variables

For the L_2 Fourier features, the covariance between two features u_m and $u_{m'}$ depends on the basis functions being sines or cosines. We will thus detail the two cases considering either the real or imaginary part of \mathbf{u} . The following integrals, which can be computed using $\cos(\omega x) = \mathcal{R}e(e^{i\omega x})$, will be of particular interest:

$$\int_x^y e^{\lambda s} \cos(\omega s) ds = \frac{1}{\lambda^2 + \omega^2} [e^{\lambda x}(-\lambda \cos(\omega x) - \omega \sin(\omega x)) + e^{\lambda y}(\lambda \cos(\omega y) + \omega \sin(\omega y))], \quad (91)$$

$$\int_x^y e^{\lambda s} \sin(\omega s) ds = \frac{1}{\lambda^2 + \omega^2} [e^{\lambda x}(\omega \cos(\omega x) - \lambda \sin(\omega x)) + e^{\lambda y}(-\omega \cos(\omega y) + \lambda \sin(\omega y))]. \quad (92)$$

case 1: $i, j \leq M$ (cosine block)

$$\text{cov}[u_m, u_{m'}] = \mathbb{E} \left[\int_a^b f(s) \cos(-\omega_m(s-a)) ds \int_a^b f(t) \cos(-\omega_{m'}(t-a)) dt \right] \quad (93)$$

$$= \int_a^b \int_a^b k(s, t) \cos(-\omega_m(s-a)) \cos(-\omega_{m'}(t-a)) ds dt \quad (94)$$

$$= \int_0^{a-b} \int_0^{b-a} k(s, t) \cos(\omega_m s) \cos(\omega_{m'} t) ds dt \quad (95)$$

$$= \frac{1}{\lambda^2 + \omega_m^2} \int_0^{b-a} \left[2\lambda \cos(\omega_m t) - \lambda(e^{-\lambda t} + e^{\lambda(t+a-b)}) \right] \cos(\omega_{m'} t) dt \quad (96)$$

$$= \frac{2\lambda}{\lambda^2 + \omega_m^2} \int_a^b \left[\cos(\omega_m t) - e^{-\lambda t} \right] \cos(\omega_{m'} t) dt. \quad (97)$$

case 1a: $m \neq m'$

$$\text{cov}[u_m, u_{m'}] = \frac{-2\lambda}{\lambda^2 + \omega_m^2} \int_0^{b-a} e^{-\lambda t} \cos(\omega_{m'} t) dt \quad (98)$$

$$= \frac{-2\lambda^2}{(\lambda^2 + \omega_m^2)(\lambda^2 + \omega_{m'}^2)} [1 - e^{\lambda(a-b)}]. \quad (99)$$

case 1b: $m = m' \neq 0$

$$\text{cov}[u_m, u_{m'}] = \frac{-2\lambda^2}{(\lambda^2 + \omega_m^2)(\lambda^2 + \omega_{m'}^2)} [1 - e^{\lambda(a-b)}] + (b-a) \frac{\lambda}{\lambda^2 + \omega_m^2}. \quad (100)$$

case 1c: $m = m' = 0$

$$\text{cov}[u_m, u_{m'}] = \frac{-2\lambda^2}{(\lambda^2 + \omega_m^2)(\lambda^2 + \omega_{m'}^2)} [1 - e^{\lambda(a-b)}] + 2(b-a) \frac{\lambda}{\lambda^2 + \omega_m^2}. \quad (101)$$

case 2: $m, m' > M$ (sine block)

$$\text{cov}[u_m, u_{m'}] = \int_0^{b-a} \int_0^{b-a} k(s, t) \sin(\omega_m s) \sin(\omega_{m'} t) ds dt \quad (102)$$

$$= \frac{2}{\lambda^2 + \omega_m^2} \int_0^{b-a} \left[\lambda \sin(\omega_m t) + \omega_m e^{-\lambda t} \right] \sin(\omega_{m'} t) dt. \quad (103)$$

case 2a: $m \neq m'$

$$\text{cov}[u_m, u_{m'}] = \frac{2\omega_m}{\lambda^2 + \omega_m^2} \int_0^{b-a} e^{-\lambda t} \sin(\omega_{m'} t) dt \quad (104)$$

$$= \frac{2\omega_m \omega_{m'}}{(\lambda^2 + \omega_m^2)(\lambda^2 + \omega_{m'}^2)} [1 - e^{\lambda(a-b)}]. \quad (105)$$

case 2b: $m = m'$

$$\text{cov}[u_m, u_{m'}] = \frac{2\omega_{m'}^2}{(\lambda^2 + \omega_m^2)(\lambda^2 + \omega_{m'}^2)} [1 - e^{\lambda(a-b)}] + (b-a) \frac{\lambda}{\lambda^2 + \omega_m^2}. \quad (106)$$

Appendix B. Matérn inner products

The expressions of inner products for Matérn RKHS on $[a, b]$ can be found in Durrande et al. (2016). We adopt here the following notations to obtain compact expressions: for any function $h \in \mathcal{H}$, $\mathcal{I} : h \rightarrow h$ is the identity operator and $\mathcal{D} : g \rightarrow g'$ is the differentiation operator. As a consequence, $(\lambda\mathcal{I} + \mathcal{D})^2(h)$ is a shorthand for $\lambda^2 h + 2\lambda h' + h''$.

$$\text{Matérn-}\frac{1}{2}: \langle g, h \rangle_{\mathcal{H}_{1/2}} = \frac{1}{2\lambda\sigma^2} \int_a^b (\lambda\mathcal{I} + \mathcal{D})(g)(\lambda\mathcal{I} + \mathcal{D})(h) dt + \frac{1}{\sigma^2} g(a)h(a), \quad (107)$$

$$\begin{aligned} \text{Matérn-}\frac{3}{2}: \langle g, h \rangle_{\mathcal{H}_{3/2}} &= \frac{1}{4\lambda^3\sigma^2} \int_a^b (\lambda\mathcal{I} + \mathcal{D})^2(g)(\lambda\mathcal{I} + \mathcal{D})^2(h) dt \\ &\quad + \frac{1}{\sigma^2} g(a)h(a) + \frac{1}{\lambda^2\sigma^2} g'(a)h'(a), \end{aligned} \quad (108)$$

$$\begin{aligned} \text{Matérn-}\frac{5}{2}: \langle g, h \rangle_{\mathcal{H}_{5/2}} &= \frac{3}{16\lambda^5\sigma^2} \int_a^b (\lambda\mathcal{I} + \mathcal{D})^3(g)(\lambda\mathcal{I} + \mathcal{D})^3(h) dt \\ &\quad + \frac{9}{8\sigma^2} g(a)h(a) + \frac{9}{8\lambda^4\sigma^2} g(a)''h''(a) \\ &\quad + \frac{3}{\lambda^2\sigma^2} \left(g'(a)h'(a) + \frac{1}{8} g''(a)h(a) + \frac{1}{8} g(a)h''(a) \right). \end{aligned} \quad (109)$$

B.1 Gram matrix between Fourier features for the exponential kernel

We detail here the computations of $\mathbf{K}_{\mathbf{uu}}[i, j] = \langle \phi_i, \phi_j \rangle_{\mathcal{H}}$ for the exponential kernel. We recall that $\phi_0 = 1$ and that $\phi_i = \cos(\omega_i(x-a))$ and $\phi_{i+M} = \sin(\omega_i(x-a))$ for $i \in \{1, \dots, M\}$. Furthermore, the frequencies ω_i are harmonic on $[a, b]$.

case 1 : $i, j \leq M$ (cosine block)

$$\mathbf{K}_{\mathbf{uu}}[i, j] = \frac{1}{2\sigma^2\lambda} \int_0^{b-a} (\lambda \cos(\omega_i s) - \omega_i \sin(\omega_i s))(\lambda \cos(\omega_j s) - \omega_j \sin(\omega_j s)) ds + \frac{1}{\sigma^2}. \quad (110)$$

The integral is zero for all non-diagonal terms ($i \neq j$). As a consequence, the block of $\mathbf{K}_{\mathbf{uu}}$ associated with the cosine basis functions are $\text{diag}(\alpha_{\cos}) + \sigma^{-2}$, where

$$\alpha_{\cos}[i] = \frac{1}{2\sigma^2\lambda} \int_0^{b-a} \lambda^2 \cos(\omega_i s)^2 + \omega_i^2 \sin(\omega_i s)^2 ds = \begin{cases} \frac{\lambda(b-a)}{2\sigma^2} & \text{if } i = 0, \\ \frac{b-a}{4\sigma^2\lambda}(\lambda^2 + \omega_i^2) & \text{if } i \neq 0, \end{cases} \quad (111)$$

which leads to $\alpha_{\cos} = \frac{1}{2}(b-a) [2s(0)^{-1}, s(\omega_1)^{-1}, \dots, s(\omega_M)^{-1}]$.

case 2 : $i, j > M$ (sine block)

$$\mathbf{K}_{\mathbf{uu}}[i, j] = \frac{1}{2\sigma^2\lambda} \int_0^{b-a} (\lambda \sin(\omega_i s) - \omega_i \cos(\omega_i s))(\lambda \sin(\omega_j s) - \omega_j \cos(\omega_j s)) ds. \quad (112)$$

The $\mathbf{K}_{\mathbf{uu}}$ block associated to the sine basis functions is exactly a diagonal matrix with:

$$\mathbf{K}_{\mathbf{uu}}[i, i] = \frac{1}{2\sigma^2\lambda} \int_0^{b-a} (\lambda^2 \sin(\omega_i s)^2 + \omega_i^2 \cos(\omega_i s)^2) ds = \frac{b-a}{4\sigma^2\lambda}(\lambda^2 + \omega_i^2). \quad (113)$$

Similarly to the cosine block, we can write the sine block as $\text{diag}(\alpha_{sin})$ with $\alpha_{sin} = \frac{1}{2}(b - a) [s(\omega_1)^{-1}, \dots, s(\omega_M)^{-1}]$.

case 3 : $i \leq 2M + 1 < j$ (off-diagonal block) leads to $\mathbf{K}_{\phi\phi}[i, j] = 0$.

B.2 Gram matrix associated to Fourier features for the Matérn- $\frac{3}{2}$ kernel

case 1 : $i, j \leq M$ (cosine block)

$$\begin{aligned}
 \mathbf{K}_{\phi\phi}[i, j] &= \frac{1}{4\sigma^2\lambda^3} \int_0^{b-a} (\lambda^2 \cos(\omega_i s) - 2\lambda\omega_i \sin(\omega_i s) - \omega_i^2 \cos(\omega_i s)) \times \\
 &\quad (\lambda^2 \cos(\omega_j s) - 2\lambda\omega_j \sin(\omega_j s) - \omega_j^2 \cos(\omega_j s)) ds \\
 &\quad + \frac{1}{\sigma^2} \cos(\omega_i a) \cos(\omega_j a) + \frac{1}{\lambda^2\sigma^2} \omega_i \omega_j \sin(\omega_i a) \sin(\omega_j a) \\
 &= \begin{cases} \frac{1}{\sigma^2} & \text{if } i \neq j, \\ \frac{1}{4\sigma^2\lambda^3} \int_0^{b-a} \lambda^4 \cos^2(\omega_i s) + \omega_i^4 \cos^2(\omega_i s) \\ \quad - 2\lambda^2\omega^2 \cos^2(\omega_i s) + 4\lambda^2\omega_i^2 \sin^2(\omega_i s) ds + \frac{1}{\sigma^2} & \text{if } i = j, \end{cases} \quad (114) \\
 &= \begin{cases} \frac{1}{\sigma^2} & \text{if } i \neq j, \\ \frac{b-a}{8\sigma^2\lambda^3} (\lambda^2 + \omega_i^2)^2 + \frac{1}{\sigma^2} & \text{if } i = j \neq 0, \\ \frac{\lambda(b-a)}{4\sigma^2} + \frac{1}{\sigma^2} & \text{if } i = j = 0. \end{cases}
 \end{aligned}$$

case 2 : $i, j > M$ (sine block)

$$\begin{aligned}
 \mathbf{K}_{\phi\phi}[i, j] &= \frac{1}{4\sigma^2\lambda^3} \int_0^{b-a} (\lambda^2 \sin(\omega_i s) + 2\lambda\omega_i \cos(\omega_i s) - \omega_i^2 \sin(\omega_i s)) \times \\
 &\quad (\lambda^2 \sin(\omega_j s) + 2\lambda\omega_j \cos(\omega_j s) - \omega_j^2 \sin(\omega_j s)) ds \\
 &\quad + \frac{1}{\sigma^2} \sin(\omega_i a) \sin(\omega_j a) + \frac{1}{\lambda^2\sigma^2} \omega_i \omega_j \cos(\omega_i a) \cos(\omega_j a) \\
 &= \begin{cases} \frac{\omega_i \omega_j}{\lambda^2\sigma^2} & \text{if } i \neq j, \\ \frac{1}{4\sigma^2\lambda^3} \int_0^{b-a} \lambda^4 \sin^2(\omega_i s) + \omega_i^4 \sin^2(\omega_i s) \\ \quad - 2\lambda^2\omega^2 \sin^2(\omega_i s) + 4\lambda^2\omega_i^2 \cos^2(\omega_i s) ds + \frac{\omega_i \omega_j}{\lambda^2\sigma^2} & \text{if } i = j, \end{cases} \quad (115) \\
 &= \begin{cases} \frac{\omega_i \omega_j}{\lambda^2\sigma^2} & \text{if } i \neq j, \\ \frac{b-a}{8\sigma^2\lambda^3} (\lambda^2 + \omega_i^2)^2 + \frac{\omega_i \omega_j}{\lambda^2\sigma^2} & \text{if } i = j. \end{cases}
 \end{aligned}$$

case 3: $i \leq M < j$ (off-diagonal block)

$$\begin{aligned}
 \mathbf{K}_{\phi\phi}[i, j] &= \frac{1}{4\sigma^2\lambda^3} \int_0^{b-a} (\lambda^2 \cos(\omega_i s) - 2\lambda\omega_i \sin(\omega_i s) - \omega_i^2 \cos(\omega_i s)) \times \\
 &\quad (\lambda^2 \sin(\omega_j s) + 2\lambda\omega_j \cos(\omega_j s) - \omega_j^2 \sin(\omega_j s)) ds \\
 &\quad + \frac{1}{\sigma^2} \cos(\omega_i a) \sin(\omega_j a) - \frac{1}{\lambda^2\sigma^2} \omega_i \omega_j \sin(\omega_i a) \cos(\omega_j a) \\
 &= \begin{cases} 0 & \text{if } i \neq j, \\ \frac{1}{4\sigma^2\lambda^3} \int_0^{b-a} 2\lambda^3 \omega_i \cos^2(\omega_i s) - 2\lambda\omega_i^3 \cos^2(\omega_i s) \\ \quad - 2\lambda^3 \omega_i \sin^2(\omega_i s) + 2\lambda\omega_i^3 \sin^2(\omega_i s) ds & \text{if } i = j, \end{cases} \quad (116) \\
 &= 0
 \end{aligned}$$

B.3 Gram matrix of Fourier features for the Matérn- $\frac{5}{2}$ kernel

$$\langle g, h \rangle_{\mathcal{H}_{5/2}} = \frac{3}{16\lambda^5\sigma^2} \int_0^{b-a} L_t(g)L_t(h) dt + G(a), \quad (117)$$

where $L_t(g) = (\lambda^3 g(t) + 3\lambda^2 g'(t) + 3\lambda g''(t) + g'''(t))$,

$$\begin{aligned}
 G(a) &= \frac{9}{8\sigma^2} g(a)h(a) + \frac{9}{8\lambda^4\sigma^2} g(a)''h''(a) \\
 &\quad + \frac{3}{\lambda^2\sigma^2} \left(g'(a)h'(a) + \frac{1}{8}g''(a)h(a) + \frac{1}{8}g(a)h''(a) \right). \quad (118)
 \end{aligned}$$

case 1: $i, j \leq M$ (cosine block)

$$\begin{aligned}
 L_x(\cos(\omega_i \cdot)) &= \lambda^3 \cos(\omega_i x) - 3\lambda\omega_i^2 \sin(\omega_i x) - 3\lambda\omega_i^2 \cos(\omega_i x) + \omega_i^3 \sin(\omega_i x) \\
 &= (\lambda^3 - 3\lambda\omega_i^2) \cos(\omega_i x) + (\omega_i^3 - 3\lambda\omega_i^2) \sin(\omega_i x) \\
 G(a) &= \frac{9}{8\sigma^2} + \frac{9}{8\lambda^4\sigma^2} \omega_i^2 \omega_j^2 - \frac{3}{8\lambda^2\sigma^2} (\omega_i^2 + \omega_j^2) \\
 &= \frac{1}{\sigma^2} + \frac{1}{8\sigma^2} \left(\frac{3\omega_i^2}{\lambda^2} - 1 \right) \left(\frac{3\omega_j^2}{\lambda^2} - 1 \right). \quad (119)
 \end{aligned}$$

$$\mathbf{K}_{\phi\phi}[i, j] = \begin{cases} G(a) & \text{if } i \neq j, \\ \frac{3(b-a)}{32\lambda^5\sigma^2} (\lambda^6 - 6\lambda^4\omega_i^2 + 9\lambda^2\omega_i^4 + \omega_i^6 - 6\lambda^2\omega_i^4 + 9\omega_i^2\lambda^4) + G(a) & \text{if } i = j \neq 0, \\ \frac{3(b-a)}{16\lambda^5\sigma^2} (\lambda^6 - 6\lambda^4\omega_i^2 + 9\lambda^2\omega_i^4) + G(a) & \text{if } i = j = 0, \end{cases} \quad (120)$$

which boils down to

$$\mathbf{K}_{\phi\phi}[i, j] = \begin{cases} G(a) & \text{if } i \neq j, \\ \frac{3(b-a)}{32\lambda^5\sigma^2}(\lambda^2 + \omega_i^2)^3 + G(a) & \text{if } i = j \neq 0, \\ \frac{3\lambda(b-a)}{16\sigma^2} + G(a) & \text{if } i = j = 0. \end{cases} \quad (121)$$

case 2: $i, j > M$ (sine block)

$$\begin{aligned} L_x(\sin(\omega_i \cdot)) &= \lambda^3 \sin(\omega_i x) + 3\omega_i \lambda^2 \cos(\omega_i x) - 3\lambda \omega_i^2 \sin(\omega_i x) - \omega_i^3 \cos(\omega_i x) \\ &= (\lambda^3 - 3\lambda \omega_i^2) \sin(\omega_i x) - (\omega_i^3 - 3\omega_i \lambda^2) \cos(\omega_i x), \\ G(a) &= \frac{3\omega_i \omega_j}{\lambda^2 \sigma^2}, \end{aligned} \quad (122)$$

$$\mathbf{K}_{\phi\phi}[i, j] = \begin{cases} G(a) & \text{if } i \neq j, \\ \frac{3(b-a)}{32\lambda^5\sigma^2}(\lambda^2 + \omega_i^2)^3 + G(a) & \text{if } i = j \neq 0. \end{cases} \quad (123)$$

case 3: $i \leq M < j$ (off-diagonal block). As previously, calculations give $\mathbf{K}_{\phi\phi}[i, j] = 0$.

References

- Martin Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- Vincent Adam, James Hensman, and Maneesh Sahani. Scalable transformed additive signal decomposition by non-conjugate Gaussian process inference. In IEEE International Workshop on Machine Learning for Signal Processing (MLSP), 2016.
- Naum I Akhiezer and Izrail' M Glazman. Theory of Linear Operators in Hilbert Space. Dover, New York, 1993.
- Mauricio Alvarez and Neil D Lawrence. Sparse convolved Gaussian processes for multi-output regression. In Advances in Neural Information Processing Systems 21, pages 57–64. Curran Associates, Inc., 2009.
- Sudipto Banerjee, Alan E Gelfand, Andrew O Finley, and Huiyan Sang. Gaussian predictive process models for large spatial data sets. Journal of the Royal Statistical Society: Series B (Statistical Methodology), 70(4):825–848, 2008.
- Sudipto Banerjee, Bradley P Carlin, and Alan E Gelfand. Hierarchical Modeling and Analysis for Spatial Data. CRC Press, 2014.
- Alain Berlinet and Christine Thomas-Agnan. Reproducing Kernel Hilbert Spaces in Probability and Statistics. Kluwer Academic Publishers, 2004.
- Matt A Bernstein, Kevin F King, and Xiaohong Joe Zhou. Handbook of MRI Pulse Sequences. Elsevier, 2004.
- David M Blei, Alp Kucukelbir, and Jon D McAuliffe. Variational inference: A review for statisticians. arXiv preprint arXiv:1601.00670, 2016.
- François-Xavier Briol, Chris J Oates, Mark Girolami, Michael A Osborne, and Dino Sejdinovic. Probabilistic integration: A role for statisticians in numerical analysis? arXiv preprint arXiv:1512.00933, 2016.
- Thang D Bui and Richard E Turner. Tree-structured Gaussian process approximations. In Advances in Neural Information Processing Systems 27, pages 2213–2221. Curran Associates, Inc., 2014.
- Kian Ming A Chai. Variational multinomial logit Gaussian process. Journal of Machine Learning Research, 13:1745–1808, 2012.
- Ole F Christensen, Gareth O Roberts, and Martin Sködl. Robust Markov chain Monte Carlo methods for spatial generalized linear mixed models. Journal of Computational and Graphical Statistics, 15(1):1–17, 2006.
- Lehel Csató and Manfred Opper. Sparse on-line Gaussian processes. Neural Computation, 14(3):641–668, 2002.

- Lehel Csató, Manfred Opper, and Ole Winther. TAP Gibbs free energy, belief propagation and sparsity. In Advances in Neural Information Processing Systems, pages 657–663, 2002.
- Andreas Damianou, Michalis K Titsias, and Neil D Lawrence. Variational inference for latent variables and uncertain inputs in Gaussian processes. Journal of Machine Learning Research, 17:1–62, 2016.
- Andreas C Damianou and Neil D Lawrence. Deep Gaussian processes. In International Conference on Artificial Intelligence and Statistics (AISTATS), pages 207–215, 2013.
- Marc P Deisenroth and Jun Wei Ng. Distributed Gaussian processes. In International Conference on Machine Learning (ICML), pages 1481–1490, 2015.
- Marc P Deisenroth, Dieter Fox, and Carl E Rasmussen. Gaussian processes for data-efficient learning in robotics and control. IEEE Transactions on Pattern Analysis and Machine Intelligence, 37(2):408–423, 2015.
- Amir Dezfouli and Edwin V Bonilla. Scalable inference for Gaussian process models with black-box likelihoods. In Advances in Neural Information Processing Systems, pages 1414–1422, 2015.
- Peter J Diggle. Statistical Analysis of Spatial and Spatio-Temporal Point Patterns. CRC Press, 2013.
- Michael F Driscoll. The reproducing kernel Hilbert space structure of the sample paths of a Gaussian process. Probability Theory and Related Fields, 26(4):309–316, 1973.
- Nicolas Durrande, David Ginsbourger, and Olivier Roustant. Additive covariance kernels for high-dimensional Gaussian process modeling. In Annales de la Faculté de Sciences de Toulouse, volume 21, pages p–481, 2012.
- Nicolas Durrande, James Hensman, Magnus Rattray, and Neil D Lawrence. Detecting periodicities with Gaussian processes. PeerJ Computer Science, 2:e50, 2016.
- Alok Dutt and Vladimir Rokhlin. Fast Fourier transforms for nonequispaced data. SIAM Journal on Scientific computing, 14(6):1368–1393, 1993.
- David K Duvenaud, Hannes Nickisch, and Carl E Rasmussen. Additive Gaussian processes. In Advances in Neural Information Processing Systems 24, pages 226–234. Curran Associates, Inc., 2011.
- Anibal Figueiras-Vidal and Miguel Lázaro-Gredilla. Inter-domain Gaussian processes for sparse inference using inducing features. In Advances in Neural Information Processing Systems 22, pages 1087–1095. Curran Associates, Inc., 2009.
- Maurizio Filippone and Raphael Engler. Enabling scalable stochastic gradient-based inference for Gaussian processes by employing the Unbiased Linear System SolvEr (ULISSE). In International Conference on Machine Learning (ICML), pages 6–11, 2015.

- Maurizio Filippone, Mingjun Zhong, and Mark Girolami. A comparative evaluation of stochastic-based inference methods for Gaussian process models. Machine Learning, 93(1):93–114, 2013.
- Jochen Fritz, Insa Neuweiler, and Wolfgang Nowak. Application of FFT-based algorithms for large-scale universal kriging problems. Mathematical Geosciences, 41(5):509–533, 2009.
- Yarin Gal and Richard Turner. Improving the Gaussian process sparse spectrum approximation by representing uncertainty in frequency inputs. In International Conference on Machine Learning (ICML), pages 655–664, 2015.
- Leslie Greengard and June-Yub Lee. Accelerating the nonuniform fast Fourier transform. SIAM Review, 46(3):443–454, 2004.
- Philipp Hennig, Michael A Osborne, and Mark Girolami. Probabilistic numerics and uncertainty in computations. In Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences, volume 471, page 20150142. The Royal Society, 2015.
- James Hensman, Nicolò Fusi, and Neil D Lawrence. Gaussian processes for big data. In Conference on Uncertainty in Artificial Intelligence (UAI), pages 282–290. AUAI Press, 2013.
- James Hensman, Alexander G de G Matthews, Maurizio Filippone, and Zoubin Ghahramani. Mcmc for variationally sparse Gaussian processes. In Advances in Neural Information Processing Systems 28, pages 1639–1647. Curran Associates, Inc., 2015a.
- James Hensman, Alexander G de G Matthews, and Zoubin Ghahramani. Scalable variational Gaussian process classification. In International Conference on Artificial Intelligence and Statistics (AISTATS), pages 351–360, 2015b.
- Emtiyaz Khan, Shakir Mohamed, and Kevin P Murphy. Fast Bayesian inference for non-conjugate Gaussian process regression. In Advances in Neural Information Processing Systems, pages 3140–3148, 2012.
- Emtiyaz Khan, Aleksandr Aravkin, Michael Friedlander, and Matthias Seeger. Fast dual variational inference for non-conjugate latent gaussian models. In International Conference on Machine Learning, pages 951–959, 2013.
- Neil D Lawrence. Probabilistic non-linear principal component analysis with Gaussian process latent variable models. Journal of Machine Learning Research, 6:1783–1816, 2005.
- Miguel Lázaro-Gredilla, Joaquin Quiñonero-Candela, Carl E Rasmussen, and Aníbal R Figueiras-Vidal. Sparse spectrum Gaussian process regression. Journal of Machine Learning Research, 11:1865–1881, 2010.
- Judith Lean. Solar irradiance reconstruction. Data contribution series # 2004-035, IGBP PAGES/World Data Center for Paleoclimatology NOAA/NGDC Paleoclimatology Program, Boulder, CO, USA, 2004.

- Chris Lloyd, Tom Gunter, Michael Osborne, and Stephen Roberts. Variational inference for Gaussian process modulated Poisson processes. In International Conference on Machine Learning (ICML), pages 1814–1822, 2015.
- Alexander G de G Matthews. Scalable Gaussian Process Inference Using Variational Methods. PhD thesis, University of Cambridge, Cambridge, UK, 2016.
- Alexander G de G Matthews, James Hensman, Richard E Turner, and Zoubin Ghahramani. On sparse variational methods and the Kullback-Leibler divergence between stochastic processes. In International Conference on Artificial Intelligence and Statistics (AISTATS), pages 231–238, 2016.
- 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, 2017.
- Jesper M oller, Anne R Syversveen, and Rasmus P Waagepetersen. Log Gaussian cox processes. Scandinavian Journal of Statistics, 25(3):451–482, 1998.
- David Moore and Stuart J Russell. Gaussian process random fields. In Advances in Neural Information Processing Systems 28, pages 3357–3365. Curran Associates, Inc., 2015.
- Iain Murray and Ryan P Adams. Slice sampling covariance hyperparameters of latent Gaussian models. In Advances in Neural Information Processing Systems 23, pages 1732–1740. Curran Associates, Inc., 2010.
- Iain Murray, Ryan P Adams, and David JC MacKay. Elliptical slice sampling. In International Conference on Artificial Intelligence and Statistics (AISTATS), pages 541–548, 2010.
- Thomas Nickson, Tom Gunter, Chris Lloyd, Michael A Osborne, and Stephen Roberts. Blitzkriging: Kronecker-structured stochastic Gaussian processes. arXiv preprint arXiv:1510.07965, 2015.
- Manfred Opper and C edric Archambeau. The variational Gaussian approximation revisited. Neural Computation, 21(3):786–792, 2009.
- Michael A Osborne. Bayesian Gaussian Processes for Sequential Prediction, Optimisation and Quadrature. PhD thesis, University of Oxford, Oxford, UK, 2010.
- Christopher J Paciorek. Bayesian smoothing with Gaussian processes using Fourier basis functions in the spectralGP package. Journal of Statistical Software, 19(2):1–38, 2007.
- Joaquin Qui onero-Candela and Carl E Rasmussen. A unifying view of sparse approximate Gaussian process regression. Journal of Machine Learning Research, 6:1939–1959, 2005.
- Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In Advances in Neural Information Processing Systems 20, pages 1177–1184. Curran Associates, Inc., 2008.

- Ali Rahimi and Benjamin Recht. Weighted sums of random kitchen sinks: Replacing minimization with randomization in learning. In Advances in Neural Information Processing Systems 21, pages 1313–1320. Curran Associates, Inc., 2009.
- Barbara Rakitsch, Christoph Lippert, Karsten Borgwardt, and Oliver Stegle. It is all in the noise: Efficient multi-task Gaussian process inference with structured residuals. In Advances in Neural Information Processing Systems 26, pages 1466–1474. Curran Associates, Inc., 2013.
- Yunus Saatçi. Scalable Inference for Structured Gaussian Process Models. PhD thesis, University of Cambridge, Cambridge, UK, 2012.
- Yves-Laurent Kom Samo and Stephen J Roberts. String and membrane Gaussian processes. Journal of Machine Learning Research, 17:1–87, 2016.
- Simo Särkkä, Arno Solin, and Jouni Hartikainen. Spatiotemporal learning via infinite-dimensional Bayesian filtering and smoothing. IEEE Signal Processing Magazine, 30(4): 51–61, 2013.
- Matthias Seeger. Bayesian Gaussian process models: PAC-Bayesian generalisation error bounds and sparse approximations. PhD thesis, University of Edinburgh, Edinburgh, UK, 2003.
- Edward Snelson and Zoubin Ghahramani. Sparse Gaussian processes using pseudo-inputs. In Advances in Neural Information Processing Systems 18, pages 1257–1264. MIT Press, 2005.
- Arno Solin and Simo Särkkä. Hilbert space methods for reduced-rank Gaussian process regression. arXiv preprint arXiv:1401.5508, 2014.
- Arno Solin, Pasi Jylänki, Jaakko Kauramäki, Tom Heskes, Marcel AJ van Gerven, and Simo Särkkä. Regularizing solutions to the MEG inverse problem using space-time separable covariance functions. arXiv preprint arXiv:1604.04931, 2016.
- Oliver Stegle, Christoph Lippert, Joris M Mooij, Neil D Lawrence, and Karsten M Borgwardt. Efficient inference in matrix-variate Gaussian models with iid observation noise. In Advances in Neural Information Processing Systems 24, pages 630–638. Curran Associates, Inc., 2011.
- Benjamin M Taylor and Peter J Diggle. INLA or MCMC? A tutorial and comparative evaluation for spatial prediction in log-Gaussian Cox processes. Journal of Statistical Computation and Simulation, 84(10):2266–2284, 2014.
- Michalis K Titsias. Variational learning of inducing variables in sparse Gaussian processes. In International Conference on Artificial Intelligence and Statistics (AISTATS), pages 567–574, 2009.
- Michalis K Titsias and Neil D Lawrence. Bayesian Gaussian process latent variable model. In International Conference on Artificial Intelligence and Statistics (AISTATS), pages 844–851, 2010.

- Volker Tresp. A Bayesian committee machine. Neural Computation, 12(11):2719–2741, 2000.
- Richard E Turner. Statistical Models for Natural Sounds. PhD thesis, UCL (University College London), London, UK, 2010.
- Richard E Turner and Maneesh 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.
- Jarno Vanhatalo and Aki Vehtari. Sparse log Gaussian processes via MCMC for spatial epidemiology. In Gaussian Processes in Practice, JMLR: Workshop and Conference Proceedings, pages 73–89, 2007.
- Christopher KI Williams and Carl E Rasmussen. Gaussian Processes for Machine Learning. MIT Press, 2006.
- Christopher KI Williams and Matthias Seeger. The effect of the input density distribution on kernel-based classifiers. In International Conference on Machine Learning (ICML), 2000.
- Christopher KI Williams and Matthias Seeger. Using the Nyström method to speed up kernel machines. In Advances in Neural Information Processing Systems 13, pages 682–688. MIT Press, 2001.
- Andrew G Wilson and Ryan P Adams. Gaussian process kernels for pattern discovery and extrapolation. In International Conference on Machine Learning (ICML), pages 1067–1075, 2013.
- Andrew G Wilson and Hannes Nickisch. Kernel interpolation for scalable structured Gaussian processes (KISS-GP). In International Conference on Machine Learning (ICML), pages 1775–1784, 2015.
- Andrew G Wilson, Elad Gilboa, John P Cunningham, and Arye Nehorai. Fast kernel learning for multidimensional pattern extrapolation. In Advances in Neural Information Processing Systems 27, pages 3626–3634. Curran Associates, Inc., 2014.
- Andrew G Wilson, Christoph Dann, and Hannes Nickisch. Thoughts on massively scalable Gaussian processes. arXiv preprint arXiv:1511.01870, 2015.
- Zichao Yang, Andrew G Wilson, Alex Smola, and Le Song. À la carte — Learning fast kernels. In International Conference on Artificial Intelligence and Statistics (AISTATS), pages 1098–1106, 2015.