

How Good are Low-Rank Approximations in Gaussian Process Regression?

Constantinos Daskalakis¹, Petros Dellaportas^{2, 4, 5}, Aristeidis Panos³

¹CSAIL, Massachusetts Institute of Technology, USA

²University College London, UK

³University of Warwick, UK

⁴Athens University of Economics and Business, Greece

⁵The Alan Turing Institute, UK

costis@csail.mit.edu, p.dellaportas@ucl.ac.uk, ares.panos@warwick.ac.uk

Abstract

We provide guarantees for approximate Gaussian Process (GP) regression resulting from two common low-rank kernel approximations: based on random Fourier features, and based on truncating the kernel’s Mercer expansion. In particular, we bound the Kullback–Leibler divergence between an exact GP and one resulting from one of the afore-described low-rank approximations to its kernel, as well as between their corresponding predictive densities, and we also bound the error between predictive mean vectors and between predictive covariance matrices computed using the exact versus using the approximate GP. We provide experiments on both simulated data and standard benchmarks to evaluate the effectiveness of our theoretical bounds.

1 Introduction

Gaussian processes (GPs) have long been studied in probability and statistics; see e.g. Rasmussen and Williams (2006). In Bayesian inference, they provide a canonical way to define a probability distribution over functions, which can be used as a prior to build probabilistic frameworks for quantifying uncertainty in prediction. Among many applications, they have been a method of choice for hyperparameter tuning in deep learning (Snoek, Larochelle, and Adams 2012).

In the simplest setting of GP regression, which is the focus of this paper, a measure over functions $f : \mathbf{x} \mapsto y$ is defined such that, for any collection $X = (\mathbf{x}_1, \dots, \mathbf{x}_N)$ of feature vectors, their corresponding responses $\mathbf{y} = (y_1, \dots, y_N)$ are jointly Gaussian, with zero mean and covariance matrix $K(k_\theta, X) := (k_\theta(\mathbf{x}_i, \mathbf{x}_j))_{ij}$, where $k_\theta(\cdot, \cdot)$ is a positive semidefinite kernel indexed by some parameter vector θ . A common inferential practice is to assume that we do not observe the Gaussian sample directly but additional noise drawn from a zero-mean isotropic Gaussian is added to it prior to our observation. Bayesian inference then proceeds by estimating θ and the noise variance as well as computing predictive distributions of unobserved responses \mathbf{y}^* corresponding to a collection of new feature vectors X^* of interest. These inference tasks require computing the inverse and determinant of the covariance matrix $K(k_\theta, X)$, which naively costs $O(N^3)$ operations (or more precisely matrix multiplication time),

making the inferential framework hard to scale computationally beyond a few thousand observations.

The computational burden of GP inference has motivated a large body of work on faster, approximate inference frameworks, as surveyed in Liu et al. (2020). Many rely on the Nyström method, identifying for this purpose a set of “inducing inputs” on the input (i.e. feature vector) domain (Quiñero-Candela and Rasmussen 2005; Snelson and Ghahramani 2006; Titsias 2009; Williams and Seeger 2001; Hensman, Fusi, and Lawrence 2013), or the spectral domain (Lázaro-Gredilla et al. 2010; Gal and Turner 2015; Hensman, Durrande, and Solin 2017). Other approaches are based on approximating the kernel by truncating its Mercer expansion (Ferrari-Trecate, Williams, and Opper 1999; Solin and Särkkä 2020), or using random features (Cutajar et al. 2017). For more discussion see Section 2.

The motivating question for this work is that, while there is substantial work providing approximation guarantees for various low-rank kernel approximations with respect to different metrics, the impact of such approximation guarantees to the quality of approximate GP inference is not sufficiently understood. E.g. many works provide *entry-wise* approximation guarantees between a given kernel and an approximate one constructed via the Nyström method, random features, Mercer expansion truncation, or other approximation technique; see e.g. Rahimi and Recht (2008); Cortes, Mohri, and Talwalkar (2010); Yang et al. (2012). However, it is unclear how to translate such entry-wise guarantees to meaningful approximation guarantees relating GP inference using an exact kernel to inference using an approximate kernel.

Recent work by Burt, Rasmussen, and van der Wilk (2020) pursued an investigation similar to ours for the sparse variational GP regression framework (Titsias 2009; Hensman, Matthews, and Ghahramani 2015). They provide bounds for the Kullback–Leibler (KL) divergence between the true posterior distribution and one obtained using inducing inputs in the above framework.

The goal of our work is to provide bounds for the impact to GP inference of two other prominent low-rank kernel approximation methods, based on random features (Thm 2) and on truncating the kernel’s Mercer expansion (Thm 3). In particular, we provide bounds on the KL divergence between the marginal likelihood of an idealized GP with covariance matrix $K(k_\theta, X)$ and the marginal likelihood of a GP with a

low-rank covariance matrix Σ obtained from k_θ and X using random features, or truncating k_θ 's Mercer expansion. We quantify the KL divergence in terms of the rank of Σ . We show that moderate values of the rank suffice to bring the KL divergence below any desired threshold εN , where $\varepsilon > 0$. We obtain similar bounds for the KL divergence between the predictive densities of the exact and approximate GP, and we also bound the error between the predictive mean vectors and between the predictive covariance matrices computed using the exact vs using the approximate GP. In the balance, our work provides theoretical grounding for the use of two common low-rank kernel approximations in GP regression, quantifying the inferential loss suffered in exchange for the computational benefit of working with a low-rank kernel, as discussed in Sec 3.2.

In Sec 5.1, we provide experiments investigating the effectiveness of our theoretical bounds in capturing the dependence of the KL approximation on the dimension of the input features and the rank of Σ . In particular, by comparing the blue and green curves of Fig 1 we validate our theoretical results suggesting that the Gaussian kernels require lower rank approximations to achieve a desired threshold εN when compared with Matérn kernels with the same feature vector dimensions. Similarly, by comparing the solid and dotted curves of Fig 1 we validate our theoretical results suggesting that the Mercer approximations require lower rank kernels to achieve a desired threshold εN when compared with a random feature approximations with the same feature vector dimensions. Moreover, our theoretical bounds suggest that, for a fixed rank of Σ , approximating the Gaussian kernel using random features results in worst KL approximation compared to approximating it by truncating its Mercer expansion. This is indeed reflected in our experiments on simulated data, when comparing the blue curves of the two panels of Fig 2. Similarly, our theoretical bounds suggest that truncating the Mercer expansion of the Gaussian kernel provides better approximation compared to truncating the expansion of the Matérn kernel, and this is indeed reflected when comparing the blue and green curves of the right panel of Fig 2.

In a series of real data experiments in Sec 5.2, we illustrate how low-rank approximations perform with different kernels and different ranks. The results indicate that Mercer approximations outperform random Fourier features and they perform similarly with the sparse Gaussian process regression (SGPR) of Titsias (2009), analyzed theoretically by Burt, Rasmussen, and van der Wilk (2020). The better performance of Mercer compared to Fourier is consistent with our theoretical bounds. The similar performance of Mercer and SGPR is also consistent with theory, as per our comparison to Burt, Rasmussen, and van der Wilk (2020) in Sec 2.

Paper Roadmap. Sec 2 discusses further related work. Sec 3 presents the basic GP regression setting, and well-known facts about the computational benefits of using low-rank kernel approximations. In Sec 4, we provide our theoretical results for the inferential impact to GP regression of using low-rank kernel approximations, based on random features (in Sec 4.1) and based on truncating the Mercer expansion of the kernel (in Sec 4.2). In both cases, we provide

bounds on the KL divergence between a GP and one obtained by a low-rank approximation to its kernel. In Sec 4.3 we state that these approximate guarantees are extended for the corresponding predictive densities. In Sec 5, we provide experiments whose goal is two-fold: to illustrate our theoretical guarantees in simulated data scenarios and to investigate the practical performance of our studied kernel approximations, as suggested by our theoretical bounds, on a broad collection of standard benchmarks.

2 Related Work

The challenge of scaling up GP inference is well-recognized and well-explored. We have already provided several references on approximate GP inference using inducing inputs and kernel approximations. Theoretical guarantees for GP approximations with finite models, have been provided in (Zhu et al. 1998; Ferrari-Trecate, Williams, and Opper 1999), where the notion of Mercer truncation is utilized to provide similar results over the choice of the approximating rank r . In both cases, the quality of their approximation is expressed in terms of expected mean squared error. However, note that their bounds crucially depend on a “large N ” assumption, as several sums are approximated by integrals in their development. Thus their bounds on mean squared error are only approximate and they do not quantify what is the loss resulting from their large N assumption.

The approximation error resulting from low-rank approximations based on random Fourier features has been recently studied by Hoang et al. (2020). In comparison to their results, our bound of Thm 2 is much more general as their guarantees require that the input feature vectors are sampled from a Gaussian mixture and also that the mixture components are (i) well-separated and (ii) they contribute exponentially decaying proportions of the points. In contrast, Thm 2 makes no distributional assumption about the input points.

Finally, a similar to ours theoretical investigation has been pursued by Burt, Rasmussen, and van der Wilk (2020) for the different method of sparse variational GP regression of Titsias (2009); Hensman, Matthews, and Ghahramani (2015). They provide bounds on the number of inducing inputs necessary to bring the KL divergence between the true GP posterior and the variational distribution obtained by the use of inducing inputs below a desired threshold. For the Gaussian kernel, the required number of inducing inputs scales logarithmically in the number N of training inputs, while for the Matérn kernel it scales polynomially. While their paper and ours bound different quantities, our bounds from Thm 3 are quantitatively similar to their bounds in Cor 22 for the Gaussian kernel, and our bounds have a better dependence on N compared to their bounds for the Matérn kernel in Cor 25. (To compare set $\varepsilon = \gamma/N$ in our bounds or $\gamma = \varepsilon N$ in their bounds.)

3 Preliminaries

3.1 GP Regression

In GP regression, we assume that response variables $\mathbf{y} = (y_i)_{i=1}^N \in \mathbb{R}^N$ corresponding to a collection of D -dimensional feature vectors $X = (\mathbf{x}_i)_{i=1}^N \in \mathbb{R}^{N \times D}$ are noisy evaluations of some random function $f(\cdot)$, i.e. y_i is a

noisy observation of $f(\mathbf{x}_i)$. We take the noise, $y_i - f(\mathbf{x}_i)$, for each data entry i to be independent Gaussian with mean 0 and variance σ^2 . Moreover, we place a GP prior over $f(\cdot)$, with zero mean and kernel $k_\theta(\cdot, \cdot)$, so that the collection of function values $f(X) := (f(\mathbf{x}_i))_{i=1}^N$ has a joint Gaussian distribution with zero mean and covariance matrix $K(k_\theta, X) := (k_\theta(\mathbf{x}_i, \mathbf{x}_j))_{ij}$.

The afore-described assumptions on the data generation process can be used in a regression setting in order to make predictions as follows. Suppose that we are given a collection of train feature vectors $X_{\text{train}} \in \mathbb{R}^{N_{\text{train}} \times D}$ and corresponding responses $\mathbf{y}_{\text{train}}$ and our goal is to use this training data to predict the responses \mathbf{y}_{test} on a collection of test feature vectors $X_{\text{test}} \in \mathbb{R}^{N_{\text{test}} \times D}$. Setting $A_{\text{train}} = K(k_\theta, X_{\text{train}}) + \sigma^2 I_{N_{\text{train}}}$, the log-marginal likelihood of the training data becomes $\log p(\mathbf{y}_{\text{train}} | X_{\text{train}}) = -\frac{1}{2} \mathbf{y}_{\text{train}}^\top A_{\text{train}}^{-1} \mathbf{y}_{\text{train}} - \frac{1}{2} \log |A_{\text{train}}| - \frac{N_{\text{train}}}{2} \log(2\pi)$. Conditioning on the training data, the distribution of \mathbf{y}_{test} is normal with mean and variance given by the following:

- $\mathbb{E}(\mathbf{y}_{\text{test}} | \mathbf{y}_{\text{train}}) = K(k_\theta, X_{\text{test}}, X_{\text{train}}) A_{\text{train}}^{-1} \mathbf{y}_{\text{train}}$
- $\text{Var}(\mathbf{y}_{\text{test}} | \mathbf{y}_{\text{train}}) = K(k_\theta, X_{\text{test}}) + \sigma^2 I_{N_{\text{test}}} - K(k_\theta, X_{\text{test}}, X_{\text{train}}) A_{\text{train}}^{-1} K(k_\theta, X_{\text{test}}, X_{\text{train}})^\top$

where $K(k_\theta, X_{\text{test}}, X_{\text{train}}) := (k_\theta(\mathbf{x}_{\text{test}, i}, \mathbf{x}_{\text{train}, j}))_{ij}$.

3.2 Low-rank GP Regression

Consider the setting of Sec 3.1, and additionally suppose that the kernel function $k_\theta(\cdot, \cdot)$ is low-rank in the sense that there exists a feature map $\phi : \mathbb{R}^D \rightarrow \mathbb{R}^r$ such that for all $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^D$: $k_\theta(\mathbf{x}, \mathbf{x}') = \langle \phi(\mathbf{x}), \phi(\mathbf{x}') \rangle$, where $\langle \cdot, \cdot \rangle$ is the Euclidean inner product. It follows that the kernel matrix $K(k_\theta, X)$ computed on a collection of feature vectors $X = (\mathbf{x}_i)_{i=1}^N$ can be written as $K(k_\theta, X) = \Xi \Xi^\top$, where Ξ is an $N \times r$ matrix whose rows are the vectors $\phi(\mathbf{x}_i)$, for $i = 1, \dots, N$. As such, we get that the covariance matrix of the training data $\mathbf{y}_{\text{train}}$ is $A_{\text{train}} = \Xi_{\text{train}} \Xi_{\text{train}}^\top + \sigma^2 I_{N_{\text{train}}}$. We can then use the Woodbury matrix inversion lemma and the Sylvester determinant theorem to obtain explicit forms for the inverse of A_{train} and its determinant: $A_{\text{train}}^{-1} = \sigma^{-2} I_{N_{\text{train}}} - \sigma^{-2} \Xi_{\text{train}} (\sigma^2 I_r + \Xi_{\text{train}}^\top \Xi_{\text{train}})^{-1} \Xi_{\text{train}}^\top$, and $|A_{\text{train}}| = \sigma^{2(N_{\text{train}} - r)} |\sigma^2 I_r + \Xi_{\text{train}}^\top \Xi_{\text{train}}|$. Since these identities involve inversion or determinant calculations of $r \times r$ matrices, by plugging them into the expressions for the log-marginal likelihood of observations \mathbf{y} and the mean and variance of the predictive density of future observations \mathbf{y}^* , we can, with the right ordering of operations, compute the log-likelihood and the predictive density in $O(r^3 + r^2 N_{\text{train}})$ time, i.e. linear in N_{train} , when r is a constant.

4 Approximation Guarantees

We consider the setting of Sec 3.1. In particular, we suppose that $f(\cdot)$ is sampled from a GP with mean zero and kernel function $k_\theta : \mathbb{R}^D \times \mathbb{R}^D \rightarrow \mathbb{R}$, and suppose that a collection $X = (\mathbf{x}_i)_{i=1}^N$ of feature vectors maps to a collection of responses $\mathbf{y} = (y_i)_{i=1}^N$ sampled as follows

$$\mathbf{y} \sim \mathcal{N}(0, K + \sigma^2 I_N), \quad (1)$$

where $K := K(k_\theta, X) \equiv (k_\theta(\mathbf{x}_i, \mathbf{x}_j))_{ij}$.

A well-studied topic in mathematics, statistics, and machine learning is approximating kernels with low-rank kernels. Given a kernel function k_θ , a long line of research has aimed to identify feature maps $\phi_{\theta, \varepsilon} : \mathbb{R}^D \rightarrow \mathbb{R}^r$ satisfying that, for a collection of features vectors $X = (\mathbf{x}_i)_{i=1}^N$,

$$K(k_\theta, X) \approx_\varepsilon \Sigma(\phi_{\theta, \varepsilon}, X), \quad (2)$$

where $\Sigma(\phi_{\theta, \varepsilon}, X) = (\phi_{\theta, \varepsilon}(\mathbf{x}_i)^\top \phi_{\theta, \varepsilon}(\mathbf{x}_j))_{ij}$. In (2), we have left the notion of approximation “ \approx_ε ” intentionally vague, for now, as there are many notions of approximation that have been pursued in the literature. We will soon discuss some instantiations. The parameter ε is a tunable parameter controlling the quality of the approximation.

The goal of this paper is to quantify the loss of using the approximate kernel $\Sigma(\phi_{\theta, \varepsilon}, X)$ in place of the original kernel $K(k_\theta, X)$ for the purposes of GP regression. In particular, we want to compare, in some precise sense, doing inference using the “idealized model” (1) versus an approximate model, which samples responses $\mathbf{y} = (y_i)_{i=1}^N$ for a collection of feature vectors $X = (\mathbf{x}_i)_{i=1}^N$ as follows:

$$\mathbf{y} \sim \mathcal{N}(0, \Sigma + \sigma^2 I_N), \quad (3)$$

where $\Sigma := \Sigma(\phi_{\theta, \varepsilon}, X)$. Notice that Σ can be written as $\Sigma = \Xi \Xi^\top$, where Ξ is a $N \times r$ matrix whose rows are $\phi_{\theta, \varepsilon}(\mathbf{x}_i)$ for $i = 1, \dots, N$. Thus, Σ is a rank- r matrix and, as discussed in Sec 3.2, GP regression under (3) is computationally cheap when r is small. Our goal is to quantify the inferential loss suffered in exchange for the computational benefit of working with a low-rank kernel.

The sense in which we aim to quantify the inferential loss is by bounding the KL divergence between the marginal likelihood under (3) and under (1). In Secs 4.1 and 4.2, we provide such bounds for two common low-rank kernel approximation methods, based on random features and truncated Mercer expansion respectively. In Sec 4.3, we show that our bounds also bound the KL divergence between the predictive densities, as well as the prediction error.

We start with a generic result, providing bounds on the KL divergence between Gaussians whose covariance matrices have special structure.

Proposition 1 (Proof in the supplementary material). *Suppose that Σ_1 and Σ_2 are $N \times N$ positive definite (symmetric) matrices, such that $(1 + \gamma)\Sigma_1 - \Sigma_2$ is positive semi-definite for some $\gamma \geq 0$. Then*

$$\begin{aligned} & \text{KL}(\mathcal{N}(0, \Sigma_1) \parallel \mathcal{N}(0, \Sigma_2)) \leq \\ & \frac{1}{2} \text{Tr}(\Sigma_2^{-1/2} (\Sigma_1 - (1 + \gamma)\Sigma_2) \Sigma_2^{-1/2}). \end{aligned} \quad (4)$$

If additionally $\Sigma_2 \succeq (1 + \gamma)^{-1} \Sigma_1$, then we obtain

$$\text{KL}(\mathcal{N}(0, \Sigma_1) \parallel \mathcal{N}(0, \Sigma_2)) \leq \gamma N. \quad (5)$$

If $\Sigma_1 = \sigma^2 I_N + K_1$ and $\Sigma_2 = \sigma^2 I_N + K_2$, where K_1 and K_2 are positive semi-definite, $\sigma^2 > 0$, and $(1 + \gamma)\Sigma_1 - \Sigma_2$ is positive semi-definite, then

$$\begin{aligned} & \text{KL}(\mathcal{N}(0, \Sigma_1) \parallel \mathcal{N}(0, \Sigma_2)) \leq \\ & \frac{1}{2\sigma^2} \text{Tr}(K_1 - (1 + \gamma)K_2 + \gamma\sigma^2 I_N). \end{aligned} \quad (6)$$

In the next sections we will instantiate Prop 1 by taking $K_1 = K(k_\theta, X)$ and $K_2 = \Sigma(\phi_{\theta, \varepsilon}, X)$, resulting respectively in the idealized data generation process of (1) and the approximate one of (3). Our theorem states that the KL divergence between these two processes is controlled by (4)–(6), which as we will see next can become smaller than any desired εN for relatively modest values of the rank r , namely poly-logarithmic in N (Thm 2), or even an absolute constant (Thm 3), whenever the dimension D is an absolute constant.

4.1 Guarantees for Random Feature Approximation

A common framework for low-rank kernel approximations defines a parametric family of functions $e_\eta : \mathbb{R}^D \rightarrow \mathbb{R}$ and a distribution $p(\eta)$ over η , picking a random feature map $\phi(\mathbf{x}) = (e_{\eta_1}(\mathbf{x}), \dots, e_{\eta_r}(\mathbf{x}))$ by sampling $\eta_1, \dots, \eta_r \sim p(\eta)$. The goal is that the resulting feature map ϕ results in a good approximation of some target kernel matrix $K(k, X)$ by $\Sigma(\phi, X)$, as discussed earlier in this section.

For example, the celebrated work by (Rahimi and Recht 2008) exploits Bochner’s theorem for shift invariant kernels to derive from it a kernel-specific density $p(\eta)$ that it uses in conjunction with the family of cosine functions $e_\eta(\cdot)$ with frequency and phase determined by η . Specifically $e_\eta(\cdot)$ is derived from a random Fourier feature with spectral frequency η ; see also (Cutajar et al. 2017).

However, the kernel approximation guarantees obtained by (Rahimi and Recht 2008) (as well as by much work in this literature) only bound the *element-wise* distance between the kernel matrices $K(k, X)$ and $\Sigma(\phi, X)$. To bound the KL divergence between (1) and (3) such entry-wise bounds are insufficient. Rather, we need a spectral approximation of $K(k, X) + \sigma^2 I$ by $\Sigma(\phi, X) + \sigma^2 I$, as per Prop 1. Making use of spectral approximations by (Avron et al. 2018) for modified Fourier features, we show that the KL divergence between (1) and (3) can indeed be controlled for the Gaussian kernel. We provide our statement for the Gaussian kernel with the same fixed scaling in every direction for notational simplicity. It extends to the general Gaussian kernel with different scaling per direction in an obvious way (namely by rescaling coordinates).

Theorem 2 (Proof in the supplementary material). *Consider the D -dimensional Gaussian kernel $k(\mathbf{x}, \mathbf{x}') = \exp(-2\pi^2 \|\mathbf{x} - \mathbf{x}'\|_2^2)$, and the kernel matrix $K = K(k, X) = (k(\mathbf{x}_i, \mathbf{x}_j))_{i,j}$, where $X = (\mathbf{x}_1, \dots, \mathbf{x}_N)$ is a collection of points in \mathbb{R}^D such that, for some $R > 0$, $\|\mathbf{x}_i - \mathbf{x}_j\|_\infty \leq R, \forall i, j$. Suppose $D \leq 5 \log(N/\sigma^2) + 1$ and $\varepsilon \in (0, 0.5)$. There exists (a samplable in $O(D)$ -time) distribution $p(\eta)$ and a parameterized family $e_\eta(\cdot)$ of modified Fourier Features such that, if we take $r \geq \Omega\left(\frac{R^D}{\varepsilon^2} (\log \frac{N}{\sigma^2})^{2D} \log(\frac{N}{\delta})\right)$ random $\eta_1, \dots, \eta_r \sim p(\eta)$ and define the rank- r matrix $\Sigma = \Sigma(\phi, X)$ using the feature map $\phi(\mathbf{x}) = (e_{\eta_1}(\mathbf{x}), \dots, e_{\eta_r}(\mathbf{x}))$, then with probability at least $1 - \delta$, the KL divergence from distribution (3) to distribution (1) is at most εN .*

4.2 Guarantees for Truncated Mercer Approximation

In this section, we discuss an alternative approach for obtaining low-rank kernel approximations, based on truncating the Mercer expansion of the kernel (Mercer 1909), and our associated approximation guarantees when this low-rank kernel approximation is used in GP regression.

Suppose that k_θ is a Mercer kernel on some probability space $\mathcal{X} \subseteq \mathbb{R}^D$ with probability measure μ , which means that $k_\theta(\cdot, \cdot)$ can be written as:

$$k_\theta(\mathbf{x}, \mathbf{x}') = \sum_{t=1}^{\infty} \lambda_t e_t(\mathbf{x}) e_t(\mathbf{x}'), \quad (7)$$

where $(\lambda_t)_{t \in \mathbb{N}}$ is a sequence of summable non-negative, non-increasing numbers, i.e. *eigenvalues*, and $(e_t)_{t \in \mathbb{N}}$ is a family of mutually orthogonal unit-norm functions with respect to the inner product $\langle f, g \rangle = \int_{\mathcal{X}} f(\mathbf{x}) g(\mathbf{x}) d\mu(\mathbf{x})$, defined by μ , i.e. *eigenfunctions*. Now suppose that $X = (\mathbf{x}_i)_{i=1}^N$ is a collection of vectors $\mathbf{x}_i \in \mathcal{X}$. It follows from Eq. (7) that the kernel matrix $K(k_\theta, X)$ can be written as:

$$K(k_\theta, X) \equiv \sum_{t=1}^{\infty} \lambda_t \omega_t \omega_t^\top, \quad (8)$$

where $\omega_t = (e_t(\mathbf{x}_1), e_t(\mathbf{x}_2), \dots, e_t(\mathbf{x}_N))$, for all $t \in \mathbb{N}$. Recall that the sequence $(\lambda_t)_t$ is summable so $\lambda_t \rightarrow 0$ as $t \rightarrow \infty$. The rate of decay is very fast for many kernels. For example, the decay is exponentially fast for the Gaussian kernel, and polynomially fast for the Matérn kernel when the input distribution is compact or concentrated. These are standard facts (see e.g. (Rasmussen and Williams 2006)), but for completeness we illustrate how to derive the eigendecomposition of the high-dimensional Gaussian kernel under a Gaussian input density in Sec Bof the supplement.

The fast decay of the eigenvalues motivates approximating $K(k_\theta, X)$ by keeping the first few terms of (8). In our theorem, we quantify the impact of that truncation to GP regression in terms of the KL divergence between the data likelihood of the GP process with kernel $K(k_\theta, X)$ and the GP process with the truncated kernel.

Theorem 3 (Proof in the supplementary material). *For some $B > 0$, let $k(\cdot, \cdot)$ be a Mercer kernel on probability space (\mathcal{X}, μ) with $k(\mathbf{x}, \mathbf{x}) \leq B$, for all $\mathbf{x} \in \mathcal{X}$. Let $X = (\mathbf{x}_1, \dots, \mathbf{x}_N)$ comprise N samples from μ , let $K = K(k, X)$ (which satisfies (8)), and let $\Sigma = \sum_{t=1}^r \lambda_t \omega_t \omega_t^\top$, for some $r \in \mathbb{N}$. (So the rank of Σ is r .) With probability at least $1 - \delta$ (with respect to the samples X), the KL divergence from distribution (3) to distribution (1) is at most*

$$\frac{N}{2\sigma^2} \cdot \left(\Lambda_{>r} + \sqrt{\frac{B \Lambda_{>r}}{N \delta}} \right), \quad (9)$$

where $\Lambda_{>r} = \sum_{t>r} \lambda_t$. Two example instantiations of the bound are as follows:

- Suppose $k(\mathbf{x}, \mathbf{x}') = \exp(-2\pi^2 \|\mathbf{x} - \mathbf{x}'\|_2^2)$ is the multivariate Gaussian kernel over \mathbb{R}^D , endowed with a Gaussian density $\mu(\mathbf{x}) = (2\pi/R^2)^{\frac{D}{2}} \exp(-2\pi^2 \|\mathbf{x}\|_2^2/R^2)$,

where $R > 0$. For any absolute constant $0 < c < 1$, choosing rank

$$r = \begin{cases} \left(\Omega(RD \log(RD \vee e) + R \log \frac{1}{\varepsilon\sigma\delta}) \right)^D, & \text{if } R \geq c \\ \left(\Omega \left(\frac{D}{\log \frac{1}{R}} \left(\log \frac{D}{\log \frac{1}{R}} \vee \log \frac{2}{R^2} \right) + \frac{\log \frac{1}{\varepsilon\sigma\delta}}{\log \frac{1}{R}} \right) \right)^D, & \text{otherwise} \end{cases}$$

makes (9) at most εN . In both bounds the constant hidden by the $\Omega(\cdot)$ notation depends on c and no other parameter. Moreover, the bound easily extends to Gaussian kernels with different length scales per dimension and other product Gaussian input measures μ .

- Suppose $k(\mathbf{x}, \mathbf{x}')$ is a Matérn kernel with parameter $\nu > 0$ and length scale $\alpha > 0$ over \mathbb{R}^D ,¹ endowed with a bounded measure $\mu(\mathbf{x})$ over a bounded set. Then choosing rank $r \geq A \left(\frac{1}{\varepsilon\sigma\delta} \right)^{\Omega(D/\nu)}$ makes (9) at most εN , for some constant A that depends on ν, α, D and the bounds on μ and its support, but does not depend on $\varepsilon, \sigma, \delta, N$.

4.3 Approximation Guarantees for GP Regression

There are two ways to use Thms 2 and 3 to obtain bounds on the approximation error resulting from using an approximate GP model based on either random Fourier features or truncating the Mercer expansion of the kernel. Indeed, if we apply those theorems using $X = X_{\text{train}}$ and $\mathbf{y} = \mathbf{y}_{\text{train}}$ we immediately get bounds on $\text{KL}(P(\mathbf{y}_{\text{train}}) \parallel Q(\mathbf{y}_{\text{train}}))$ where $P(\mathbf{y}_{\text{train}})$ and $Q(\mathbf{y}_{\text{train}})$ are respectively the densities of the training data under the exact GP and the approximate one. Indeed, in the settings of Thms 2 and 3 and the choice of rank made in these theorems for $N = N_{\text{train}}$, we get that with probability at least $1 - \delta$ (with respect to the randomness in the sampling of modified Fourier Features in the setting of Thm 2 and the sampling of X_{train} in the setting of Thm 3):

$$0 \leq \text{KL}(P(\mathbf{y}_{\text{train}}) \parallel Q(\mathbf{y}_{\text{train}})) \leq \varepsilon N_{\text{train}}.$$

By the definition of KL divergence, this bound can be interpreted as a bound on the difference of the marginal likelihoods under the true and the approximate GP, in expectation over data sampled from the true GP. Indeed, we equivalently get that, with probability $\geq 1 - \delta$:

$$0 \leq \mathbb{E}_{\mathbf{y}_{\text{train}} \sim P} [\log P(\mathbf{y}_{\text{train}}) - \log Q(\mathbf{y}_{\text{train}})] \leq \varepsilon N_{\text{train}}.$$

Moreover, it is straightforward to use Thms 2 and 3 to obtain bounds on the Kullback–Leibler divergence between the predictive densities of unobserved responses \mathbf{y}_{test} on new features X_{test} corresponding to the exact GP and that obtained by either random Fourier features or by truncating the Mercer expansion of the kernel. Indeed, if we apply those theorems using $X = \begin{bmatrix} X_{\text{train}} \\ X_{\text{test}} \end{bmatrix}$ and $\mathbf{y} = \begin{bmatrix} \mathbf{y}_{\text{train}} \\ \mathbf{y}_{\text{test}} \end{bmatrix}$, we immediately get

¹Specifically the kernel takes the form

$$k(\mathbf{x}, \mathbf{x}') = \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{\|\mathbf{x} - \mathbf{x}'\|_2}{\alpha} \right)^\nu K_\nu \left(\frac{\|\mathbf{x} - \mathbf{x}'\|_2}{\alpha} \right),$$

where K_ν is a modified Bessel function and $\alpha = \ell/\sqrt{2\nu}$, $\ell > 0$; see Chapter 4 of Rasmussen and Williams (2006).

bounds on $\text{KL}(P(\mathbf{y}_{\text{train}}, \mathbf{y}_{\text{test}}) \parallel Q(\mathbf{y}_{\text{train}}, \mathbf{y}_{\text{test}}))$ where $P(\mathbf{y}_{\text{train}}, \mathbf{y}_{\text{test}})$ and $Q(\mathbf{y}_{\text{train}}, \mathbf{y}_{\text{test}})$ are the joint densities of the combined vector of observed and unobserved responses, under respectively the exact GP and the approximate GP. Indeed, in the settings of Thms 2 and 3 and the choice of rank made in these theorems for $N = N_{\text{train}} + N_{\text{test}} = N_{\text{total}}$, we get that with probability at least $1 - \delta$ (w.r.t. the randomness in the sampling of modified Fourier Features in the setting of Thm 2 and the sampling of X_{train} in the setting of Thm 3):

$$\text{KL}(P(\mathbf{y}_{\text{train}}, \mathbf{y}_{\text{test}}) \parallel Q(\mathbf{y}_{\text{train}}, \mathbf{y}_{\text{test}})) \leq \varepsilon N_{\text{total}}. \quad (10)$$

By the chain rule of KL divergence, the LHS of (10) equals

$$\text{KL}(P(\mathbf{y}_{\text{train}}) \parallel Q(\mathbf{y}_{\text{train}})) + \text{KL}(P(\mathbf{y}_{\text{test}}|\mathbf{y}_{\text{train}}) \parallel Q(\mathbf{y}_{\text{test}}|\mathbf{y}_{\text{train}})).$$

Because of the non-negativity of Kullback–Leibler divergence, Eq (10) implies that with probability at least $1 - \delta$:

$$\text{KL}(P(\mathbf{y}_{\text{test}}|\mathbf{y}_{\text{train}}) \parallel Q(\mathbf{y}_{\text{test}}|\mathbf{y}_{\text{train}})) \leq \varepsilon N_{\text{total}}.$$

By the definition of conditional KL divergence, we equivalently get that with probability $\geq 1 - \delta$ (w.r.t. the randomness in the sampling of modified Fourier Features in the setting of Thm 2 and the sampling of X_{train} in the setting of Thm 3) the expected (w.r.t. $\mathbf{y}_{\text{train}} \sim P$) KL divergence between the predictive densities of the true and the approximate GPs are close:

$$\mathbb{E}_{\mathbf{y}_{\text{train}} \sim P} [\text{KL}(P(\mathbf{y}_{\text{test}}|\mathbf{y}_{\text{train}}) \parallel Q(\mathbf{y}_{\text{test}}|\mathbf{y}_{\text{train}}))] \leq \varepsilon N_{\text{total}}. \quad (11)$$

In turn, by using Markov's inequality, we get that for η of our choosing, with probability at least $1 - \delta - \eta$ (w.r.t. the sampling of both $\mathbf{y}_{\text{train}} \sim P$ and the randomness in the sampling of modified Fourier Features in the setting of Thm 2 and the sampling of X_{train} in the setting of Thm 3):

$$\text{KL}(P(\mathbf{y}_{\text{test}}|\mathbf{y}_{\text{train}}) \parallel Q(\mathbf{y}_{\text{test}}|\mathbf{y}_{\text{train}})) \leq \frac{\varepsilon}{\eta} N_{\text{total}}. \quad (12)$$

Finally, since both $P(\mathbf{y}_{\text{test}}|\mathbf{y}_{\text{train}})$ and $Q(\mathbf{y}_{\text{test}}|\mathbf{y}_{\text{train}})$ are Gaussian distributions, our bounds from (11) and (12) directly bound the error between the predictive mean vectors and between the predictive covariance matrices computed using the approximate vs using the true GP, as per the following proposition:

Proposition 4. Consider arbitrary N -dimensional Gaussians $\mathcal{N}(\boldsymbol{\mu}_1, \Sigma_1)$ and $\mathcal{N}(\boldsymbol{\mu}_2, \Sigma_2)$. Suppose that Σ_1 and Σ_2 are non-singular, and suppose that $\text{KL}(\mathcal{N}(\boldsymbol{\mu}_1, \Sigma_1) \parallel \mathcal{N}(\boldsymbol{\mu}_2, \Sigma_2)) \leq \gamma$ for some $\gamma \geq 0$. Then

$$\frac{1}{2}(\boldsymbol{\mu}_2 - \boldsymbol{\mu}_1)^\top \Sigma_2^{-1}(\boldsymbol{\mu}_2 - \boldsymbol{\mu}_1) \leq \gamma, \quad (13)$$

$$b(2\gamma) \cdot \Sigma_2 \preceq \Sigma_1 \preceq t(2\gamma) \cdot \Sigma_2, \quad (14)$$

where $b(2\gamma)$ and $t(2\gamma)$ are respectively the smallest and largest roots of $x - 1 - \ln(x) = 2\gamma$. In particular, the Mahalanobis distance of $\boldsymbol{\mu}_1$ from $(\boldsymbol{\mu}_2, \Sigma_2)$ is bounded by $\sqrt{2\gamma}$ and Σ_1 and Σ_2 are spectrally close.²

²It can be shown that $b(2\gamma) \geq \max(1 - 2\sqrt{\gamma}, \exp(-1 - 2\gamma))$ and $t(2\gamma) \leq 1 + \max(\sqrt{8\gamma}, 8\gamma)$ so these explicit expressions can be plugged in place of $b(2\gamma)$ and $t(2\gamma)$ respectively in (14).

We set $P(y_{\text{test}}|\mathbf{y}_{\text{train}})$ in place of $\mathcal{N}(\boldsymbol{\mu}_1, \Sigma_1)$ and $Q(y_{\text{test}}|\mathbf{y}_{\text{train}})$ in place of $\mathcal{N}(\boldsymbol{\mu}_2, \Sigma_2)$ in Prop 4 and combine it with the bound of (12) to get the following:

- With probability $\geq 1 - \delta - \eta$ (w.r.t. the randomness in the sampling of modified Fourier features in the setting of Thm 2 and the sampling of X_{train} in the setting of Thm 3 as well as the sampling of $\mathbf{y}_{\text{train}} \sim P$) the Mahalanobis distance between the predictive mean vector under the true and the approximate GP is at most $\sqrt{2\varepsilon/\eta N_{\text{total}}}$. Moreover, the predictive covariances of the true and the approximate GPs are sandwiched as follows $b(2\varepsilon/\eta N_{\text{total}}) \cdot \Sigma_2 \preceq \Sigma_1 \preceq t(2\varepsilon/\eta N_{\text{total}}) \cdot \Sigma_2$.

For the above bounds to be most effective, it makes sense to choose ε to scale with N_{total} , perhaps as $1/N_{\text{total}}^\kappa$ for some κ . We note that depending on the choice of κ the rank bound of Thm 2 may or may not be effective. (It is effective if $\kappa < 1/2$). On the other hand, the rank bound of Thm 3 remains effective regardless the choice of κ as the appearance of ε in the rank bound is milder.

5 Experiments

We perform a series of simulated and real-data experiments for studying our theoretical bounds in practise. The low-rank approximation techniques of Secs 4.1 and 4.2 are referred to as *Fourier GP (FGP)* and *Mercer GP (MGP)* respectively.

Inference for FGP is based on guidelines given by (Rahimi and Recht 2008). We sample, for even r , $\frac{r}{2}$ spectral frequencies $\boldsymbol{\eta}_1, \dots, \boldsymbol{\eta}_{\frac{r}{2}}$ from the spectral density $p(\boldsymbol{\eta})$ of the kernel, and compute the feature map $\phi(\mathbf{x}) : \mathbb{R}^D \rightarrow \mathbb{R}^r$, defined by $\sqrt{\frac{2}{r}}[\cos(\boldsymbol{\eta}_1^\top \mathbf{x}), \dots, \cos(\boldsymbol{\eta}_{\frac{r}{2}}^\top \mathbf{x}), \sin(\boldsymbol{\eta}_1^\top \mathbf{x}), \dots, \sin(\boldsymbol{\eta}_{\frac{r}{2}}^\top \mathbf{x})]^\top$. The spectral frequencies are only sampled once, before training, and are kept fixed throughout the optimization of the log-marginal likelihood. The spectral density of the Gaussian kernel (27), which we use in our experiments, is $p(\boldsymbol{\eta}) = \sqrt{2\pi\Delta^{-1}}\sigma_f^{-2} \exp(-2\pi^2\boldsymbol{\eta}^\top \Delta^{-1}\boldsymbol{\eta})$.

Inference for MGP is straightforward when $D = 1$. However, when $D > 1$ the eigenvalues and eigenvectors are constructed as tensor products so their computational complexity is prohibitively large for even small values of D ; see Sec B of the supplement. Therefore, for our real data experiments we devise a simple computational trick to circumvent this problem as follows. We linearly project the D -dimensional features to a lower d -dimensional space via a weight matrix W of dimension $D \times d$ that is readily estimated by maximizing the marginal log-likelihood. The resulting projected feature matrix $Z = (\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_N)$ has dimension $N \times d$. Next, we compute a low-rank Σ by keeping, from the Mercer expansion of the kernel, the top r , under some ordering, tensor products of the eigenfunctions: $\Sigma = \sum_{n=1}^r \lambda_n \xi_n \xi_n^\top$, where $\xi_n = [e_n(\mathbf{z}_1), \dots, e_n(\mathbf{z}_N)]^\top \in \mathbb{R}^N$ and e_r is the eigenfunction of the kernel indexed by r . The choice of the ordering of the tensor product of eigenvectors is first based on the total degree of the corresponding orders in each dimension and second on lexicographical order. Finally, note that the parameters a_j in (28) of Sec B of the supplement have to be prefixed or learnt from the data. We choose to keep them

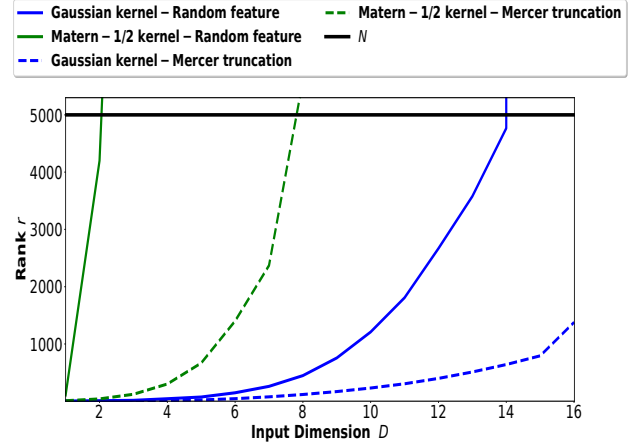


Figure 1: The lowest rank needed for Σ so that $\text{KL}(\mathcal{N}(0, K + \sigma^2 I_N) \parallel \mathcal{N}(0, \Sigma + \sigma^2 I_N)) \leq \varepsilon N$ as the input dimension D increases, in the low-rank approximation settings of Thms 2 (solid lines) and 3 (dashed lines). We sample $N = 5000$ points from $\mathcal{N}(0, \frac{1}{16^2} I_D)$ for the GP with Gaussian kernel, and sample from a D -dimensional $\mathcal{U}(-\frac{\sqrt{3}}{16}, \frac{\sqrt{3}}{16})$ for Matérn-1/2. We set $\sigma^2 = 1$ and $\varepsilon = 0.01$, respectively. For reference, we show the full rank line in black.

fixed with their values set to $1/\sqrt{2}$, which corresponds to a standard d -dimensional Gaussian measure over \mathbf{z} . Thus, we also standardize the projected features \mathbf{z} .

5.1 Experimental Rates of Convergence

In this Section we use simulated data experiments to investigate the effectiveness of our theoretical bounds in capturing the dependence of the KL approximation on the dimension of the input features and the rank of K .

Fig 1 depicts, for fixed εN , the lowest rank r needed for Σ such that $\text{KL}(\mathcal{N}(0, K + \sigma^2 I_N) \parallel \mathcal{N}(0, \Sigma + \sigma^2 I_N)) \leq \varepsilon N$ against the feature dimension D . Our theoretical results are supported by comparing the blue against green and the solid against the dotted curves: the former indicates that the Gaussian kernels require lower rank kernels to achieve the desired threshold when compared with Matérn kernels³ whereas the latter supports the theoretical results stating that the Mercer approximations require lower rank kernels to achieve a desired threshold εN when compared with a random feature approximations.

Fig 2 illustrates the practical behaviour, as a function of the rank, of the KL-divergence between an exact GP and an approximate GP obtained using random features or truncated Mercer expansion, in two example settings covered by Thms 2 and 3. Our theoretical bounds suggest that, for a fixed rank of Σ , approximating the Gaussian kernel using random features results in worst KL approximation compared to approximating it by truncating its Mercer expansion.

³Due to the analytical difficulty of getting closed forms for the eigenvalues/eigenfunctions of Matérn kernels, we resort to approximating their values by eigendecomposing the data kernel matrix.

NEGATIVE LOG-PREDICTIVE DENSITY							
	BIKE	ELEVATORS	SUPER	PROTEIN	SARCOS	KEGGDIR	3DROAD
N_{train}	15641	14939	19136	41157	44039	48071	391386
N_{test}	1738	1660	2127	4573	4894	5342	43488
D	57	18	81	9	21	19	3
SGPR6	1.43(0.01)	0.52(0.01)	0.69(0.00)	1.22(0.00)	0.35(0.01)	0.65(0.01)	1.32(0.01)
FGP6	0.97(0.02)	1.13(0.18)	0.76(0.01)	1.29(0.01)	0.41(0.00)	1.27(0.01)	1.28(0.00)
MGP6	0.99(0.01)	0.56(0.01)	0.97(0.01)	1.25(0.00)	0.33(0.00)	0.98(0.00)	1.31(0.01)
SGPR10	0.90(0.28)	0.50(0.01)	0.64(0.01)	1.19(0.00)	0.16(0.02)	0.58(0.01)	1.06(0.00)
FGP10	0.92(0.01)	1.16(0.24)	0.76(0.00)	1.22(0.01)	0.34(0.00)	1.06(0.01)	1.21(0.00)
MGP10	0.73(0.01)	0.48(0.01)	0.68(0.01)	1.19(0.00)	-0.04(0.00)	0.67(0.00)	1.15(0.01)
SGPR50	0.19(0.01)	0.46(0.01)	0.57(0.01)	1.11(0.01)	-0.21(0.00)	0.35(0.01)	0.92(0.00)
FGP50	0.92(0.01)	0.56(0.01)	0.69(0.01)	1.19(0.00)	-0.11(0.00)	0.71(0.01)	1.06(0.00)
MGP50	0.29(0.01)	0.45(0.01)	0.51(0.01)	1.15(0.00)	-0.26(0.00)	0.46(0.00)	1.02(0.00)
SGPR100	0.05(0.01)	0.44(0.01)	0.53(0.01)	1.07(0.01)	-0.30(0.00)	0.27(0.00)	0.86(0.01)
FGP100	0.87(0.04)	0.53(0.01)	0.64(0.00)	1.16(0.00)	-0.20(0.00)	0.60(0.01)	1.01(0.00)
MGP100	0.04(0.01)	0.43(0.01)	0.51(0.01)	1.14(0.00)	-0.29(0.00)	0.39(0.00)	0.85(0.00)
SGPR200	0.02(0.01)	0.43(0.01)	0.50(0.01)	1.01(0.01)	-0.39(0.00)	0.20(0.01)	0.83(0.01)
FGP200	0.81(0.03)	0.47(0.01)	0.58(0.01)	1.13(0.01)	-0.30(0.00)	0.47(0.01)	0.89(0.00)
MGP200	0.05(0.01)	0.41(0.01)	0.52(0.01)	1.13(0.00)	-0.29(0.00)	0.36(0.00)	0.81(0.00)
SGPR300	0.01(0.01)	0.42(0.01)	0.48 (0.01)	0.97 (0.01)	-0.45 (0.00)	0.16 (0.01)	0.82(0.00)
FGP300	0.72(0.04)	0.46(0.01)	0.55(0.01)	1.10(0.01)	-0.37(0.00)	0.40(0.00)	0.82(0.00)
MGP300	-0.01 (0.00)	0.40 (0.01)	0.50(0.01)	1.02(0.01)	-0.30(0.00)	0.21(0.00)	0.79 (0.00)

Table 1: Negative log-predictive density comparison (standard deviations reported in parentheses) on seven standard benchmark real-world datasets. The lowest negative log-predictive density is in bold.

This is indeed reflected in our experiments on simulated data, when comparing the blue curves of the two panels of Fig 2. Similarly, our theoretical bounds suggest that truncating the Mercer expansion of the Gaussian kernel provides better KL approximation compared to truncating the expansion of the Matérn kernel, and this is indeed reflected when comparing the blue and green curves of the right panel of Fig 2.

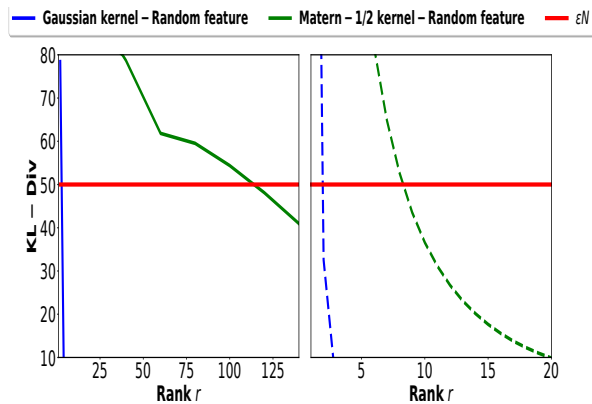


Figure 2: Decay of $KL(\mathcal{N}(0, K + \sigma^2 I_N) \parallel \mathcal{N}(0, \Sigma + \sigma^2 I_N))$ as the rank r of Σ increases, in the same settings as those in Fig 1 where $D = 1$. Results from the random feature-based low-rank approximation (i.e. Thm 2 setting) and the Mercer expansion truncation-based low-rank approximation (i.e. Thm 3 setting) are depicted on the left and right panel, respectively.

5.2 Real Data Experiments

We conduct thorough experiments, testing the quality of FGP and MGP over seven datasets from UCI repository (Dua and Graff 2017). We also compare them to Sparse GP Regression (SGPR) (Titsias 2009) which uses the Nyström approximation to conduct inference.

Both input data and their corresponding responses are standardized for all datasets. For Bike dataset, we use the standard dataset, however, we one-hot encoded some of the categorical features which led to an increased dimensionality. We train all methods for 300 epochs using Adam (Kingma and Ba 2014). All GPs use Gaussian kernels with separate length-scale per dimension. All results have been averaged over five random splits (90% train, 10% test). The last number r in an acronym indicates that a method was trained using rank r . We use $r = 6, 10, 50, 100, 200, 300$ for all three methods. For MGP, the projection dimension d is determined by cross-validation on training data, with its value ranging in $3 \leq d \leq 7$ across all seven datasets. Table 1 presents comparisons of all methods in terms of negative log-predictive density where a similar table with root mean squared error scores is given in appendix. Table 1 indicates that MGP has similar behavior with SGPR. This is consistent with our theoretical predictions as our bounds for MGP and those of Burt, Rasmussen, and van der Wilk (2020) for SGPR are quantitatively similar; see discussion in Sec 2. The slight inferior performance of FGP compared to MGP is also consistent with the predictions of the theory as the bound of Thm 2 scales worse than that of Thm 3.

Acknowledgements

C.D. and P.D. acknowledge partial financial support by the Alan Turing Institute under the EPSRC grant EP/N510129/1. C.D. was supported by NSF Awards CCF-1901292, DMS-2022448 and DMS-2134108, by a Simons Investigator Award, by the Simons Collaboration on the Theory of Algorithmic Fairness, by a DSTA grant, and by the DOE PhILMs project (No. DE-AC05-76RL01830). The authors would like to thank Andrew Ilyas for helping them with setting up the Linux machine used for the experiments.

References

- Avron, H.; Kapralov, M.; Musco, C.; Musco, C.; Velingker, A.; and Zandieh, A. 2018. Random Fourier Features for Kernel Ridge Regression: Approximation Bounds and Statistical Guarantees. *arXiv preprint arXiv:1804.09893*.
- Burt, D. R.; Rasmussen, C. E.; and van der Wilk, M. 2020. Convergence of Sparse Variational Inference in Gaussian Processes Regression. *Journal of Machine Learning Research*, 21: 1–63.
- Cortes, C.; Mohri, M.; and Talwalkar, A. 2010. On the impact of kernel approximation on learning accuracy. In *Artificial Intelligence and Statistics*, 113–120.
- Cutajar, K.; Bonilla, E. V.; Michiardi, P.; and Filippone, M. 2017. Random feature expansions for deep Gaussian processes. In *International Conference on Machine Learning*, 884–893. JMLR. org.
- Dua, D.; and Graff, C. 2017. UCI Machine Learning Repository. <http://archive.ics.uci.edu/ml>. Accessed: 2021-03-27.
- Ferrari-Trecate, G.; Williams, C. K.; and Oppor, M. 1999. Finite-dimensional approximation of Gaussian processes. In *Advances in Neural Information Processing Systems*, 218–224.
- Gal, Y.; and Turner, R. 2015. Improving the Gaussian Process Sparse Spectrum Approximation by Representing Uncertainty in Frequency Inputs. In *International Conference on Machine Learning*, 655–664.
- Hensman, J.; Durrande, N.; and Solin, A. 2017. Variational Fourier features for Gaussian processes. *The Journal of Machine Learning Research*, 18(1): 5537–5588.
- Hensman, J.; Fusi, N.; and Lawrence, N. D. 2013. Gaussian Processes for Big Data. In *Uncertainty in Artificial Intelligence*, 282. Citeseer.
- Hensman, J.; Matthews, A.; and Ghahramani, Z. 2015. Scalable variational Gaussian process classification. In *Artificial Intelligence and Statistics*, 351–360. PMLR.
- Hoang, Q. M.; Hoang, T. N.; Pham, H.; and Woodruff, D. P. 2020. Revisiting the Sample Complexity of Sparse Spectrum Approximation of Gaussian Processes. *arXiv preprint arXiv:2011.08432*.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Lázaro-Gredilla, M.; Quiñero-Candela, J.; Rasmussen, C. E.; and Figueiras-Vidal, A. R. 2010. Sparse spectrum Gaussian process regression. *The Journal of Machine Learning Research*, 11: 1865–1881.
- Liu, H.; Ong, Y.-S.; Shen, X.; and Cai, J. 2020. When Gaussian process meets big data: A review of scalable GPs. *IEEE Transactions on Neural Networks and Learning Systems*.
- Mercer, J. 1909. Functions of positive and negative type and their connection with the theory of integral equations, Philosophical Transaction of the Royal Society of London, Ser.
- Quiñero-Candela, J.; and Rasmussen, C. E. 2005. A unifying view of sparse approximate Gaussian process regression. *The Journal of Machine Learning Research*, 6(Dec): 1939–1959.
- Rahimi, A.; and Recht, B. 2008. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems*, 1177–1184.
- Rasmussen, C. E.; and Williams, C. K. I. 2006. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA.
- Snelson, E.; and Ghahramani, Z. 2006. Sparse Gaussian processes using pseudo-inputs. In *Advances in Neural Information Processing Systems*, 1257–1264.
- Snoek, J.; Larochelle, H.; and Adams, R. P. 2012. Practical Bayesian Optimization of Machine Learning Algorithms. *Advances in Neural Information Processing Systems*.
- Solin, A.; and Särkkä, S. 2020. Hilbert space methods for reduced-rank Gaussian process regression. *Statistics and Computing*, 30(2): 419–446.
- Titsias, M. 2009. Variational learning of inducing variables in sparse Gaussian processes. In *Artificial Intelligence and Statistics*, 567–574.
- Williams, C. K.; and Seeger, M. 2001. Using the Nyström method to speed up kernel machines. In *Advances in Neural Information Processing Systems*, 682–688.
- Yang, T.; Li, Y.-F.; Mahdavi, M.; Jin, R.; and Zhou, Z.-H. 2012. Nyström method vs random fourier features: A theoretical and empirical comparison. *Advances in neural information processing systems*, 25: 476–484.
- Zhu, H.; Williams, C. K.; Rohwer, R.; and Morciniec, M. 1998. Gaussian regression and optimal finite dimensional linear models. *NATO ASI Series. Series F: Computer and System Sciences*, 167–184.