
Scalable Gaussian Processes with Grid-Structured Eigenfunctions (GP-GRIEF)

Trefor W. Evans¹ Prasanth B. Nair¹

Abstract

We introduce a kernel approximation strategy that enables computation of the Gaussian process log marginal likelihood and all hyperparameter derivatives in $\mathcal{O}(p)$ time. Our GRIEF kernel consists of p eigenfunctions found using a Nyström approximation from a dense Cartesian product grid of inducing points. By exploiting algebraic properties of Kronecker and Khatri-Rao tensor products, computational complexity of the training procedure can be practically *independent* of the number of inducing points. This allows us to use arbitrarily many inducing points to achieve a globally accurate kernel approximation, even in high-dimensional problems. The fast likelihood evaluation enables type-I or II Bayesian inference on large-scale datasets. We benchmark our algorithms on real-world problems with up to two-million training points and 10^{33} inducing points.

1. Introduction

Gaussian process (GP) modelling is a powerful Bayesian approach for classification and regression, however, it is restricted to modestly sized datasets since training and inference require $\mathcal{O}(n^3)$ time and $\mathcal{O}(n^2)$ storage, where n is the number of training points (Rasmussen & Williams, 2006). This has motivated the development of approximate GP methods that use a set of m ($\ll n$) inducing points to reduce time and memory requirements to $\mathcal{O}(m^2n + m^3)$ and $\mathcal{O}(mn)$, respectively (Smola & Bartlett, 2000; Snelson & Ghahramani, 2006; Titsias, 2009; Peng & Qi, 2015). However, such techniques perform poorly if too few inducing points are used, and computational savings are lost on complex datasets that require m to be large.

Wilson & Nickisch (2015) exploited the structure of inducing points placed on a Cartesian product grid, allowing for

¹University of Toronto, Canada. Correspondence to: Trefor W. Evans <trefor.evans@mail.utoronto.ca>, Prasanth B. Nair <pbn@utias.utoronto.ca>.

$m > n$ while dramatically reducing computational demands over an exact GP. This inducing point structure enables significant performance gains in low-dimensions, however, time and storage complexities scale *exponentially* with the dataset dimensionality, rendering the technique intractable for general learning problems unless a dimensionality reduction procedure is applied. In the present work, a Cartesian product grid of inducing points is also considered, however, we show that these computational bottlenecks can be eliminated by identifying and exploiting further structure of the resulting matrices. The proposed approach leads to a highly scalable algorithm which we call GP-GRIEF (Gaussian Processes with Grid-Structured Eigenfunctions). After an initial setup cost of $\mathcal{O}(np^2 + dnp + dm^{3/d})$, GP-GRIEF requires only $\mathcal{O}(p)$ time and $\mathcal{O}(p)$ memory per log marginal likelihood evaluation, where d denotes the dataset dimensionality, and p is the number of eigenfunctions that we will describe next. We emphasize that our complexity is practically *independent* of m , which can generally be set arbitrarily high.

GP-GRIEF approximates an exact kernel as a finite sum of eigenfunctions which we accurately compute using the Nyström approximation conditioned on a huge number of inducing points. In other words, our model is sparse in the kernel eigenfunctions rather than the number of inducing points, which can greatly exceed the size of the training set due to the structure we introduce. This is attractive since it is well-known that eigenfunctions produce the most compact representation among orthogonal basis functions. Although the eigenfunctions used are approximate, we demonstrate convergence in the limit of large m . Additionally, our ability to fill out the input space with inducing points enables accurate global approximations of the eigenfunctions, even at test locations far from the training data. These basis functions also live in a reproducing kernel Hilbert space, unlike some other sparse GPs whose bases have a pre-specified form (e.g. Lázaro-Gredilla et al. (2010)). We summarize our main contributions below

- We break the *curse of dimensionality* incurred by placing inducing points on a full Cartesian product grid. Typically, a grid of inducing points results in a computational complexity that scales exponentially in d , however, we reduce this complexity to *linear* in d by exploiting algebraic properties of Kronecker and Khatri-Rao products.

- We practically eliminate dependence of the inducing point quantity, m , on computational complexity. This allows us to choose $m \gg n$ to provide a highly accurate kernel approximation, even at locations far from the training data.
- We show that the Nyström eigenfunction approximation becomes exact for large m , which is achievable thanks to the structure and algebra we introduce.
- Applications of the developed algebra are discussed to enable the extension of structured kernel interpolation methods for high-dimensional problems. We also develop an efficient preconditioner for general kernel matrices.
- We discuss a flexible parametrization of the GRIEF kernel through a re-weighting of the kernel eigenfunctions. This admits computation of the log marginal likelihood, along with all $p + 1$ hyperparameter derivatives in $\mathcal{O}(p)$.
- Finally, we demonstrate type-I Bayesian inference on real-world datasets with up to 2 million training points and $m = 10^{33}$ inducing points.

We begin with a review of GPs in section 2, and we outline an eigenfunction kernel approximation in section 3. Section 4 demonstrates why we should use many inducing points and subsequently develops the algebra necessary to make $m \gg n$ efficient and stable. Section 5 outlines a kernel reparameterization that enables efficient type-I Bayesian inference, even for large datasets. We finish with numerical studies in section 6, demonstrating the performance of GP-GRIEF on real-world datasets.

2. Background on Gaussian Processes

We will employ Gaussian processes (GPs) as non-parametric prior distributions over the latent function which generated the training dataset. It is assumed that the dataset is corrupted by independent Gaussian noise with variance $\sigma^2 \geq 0$ and that the latent function is drawn from a Gaussian process with zero mean and covariance determined by the kernel $k : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$. Considering a regression problem, the log marginal likelihood (LML) of the training targets, $\mathbf{y} \in \mathbb{R}^n$, can be written as

$$\log \mathcal{P}(\mathbf{y} | \boldsymbol{\theta}, \sigma^2, \mathbf{X}) = -\frac{1}{2} \log |\mathbf{K}_{\mathbf{X},\mathbf{X}} + \sigma^2 \mathbf{I}_n| - \frac{1}{2} \mathbf{y}^T (\mathbf{K}_{\mathbf{X},\mathbf{X}} + \sigma^2 \mathbf{I}_n)^{-1} \mathbf{y} - \frac{n}{2} \log(2\pi), \quad (1)$$

where $\mathbf{X} = \{\mathbf{x}_i \in \mathbb{R}^d\}_{i=1}^n$ is the set of n training point input locations, $[\mathbf{K}_{\mathbf{A},\mathbf{B}}]_{i,j} = k(\mathbf{a}_i, \mathbf{b}_j)$ such that $\mathbf{K}_{\mathbf{X},\mathbf{X}} \in \mathbb{R}^{n \times n}$ is the kernel covariance matrix evaluated on the training dataset, and we assume the kernel is parametrized by the hyperparameters, $\boldsymbol{\theta}$. If we consider type II Bayesian inference, we would like to select the hyperparameters $\{\sigma^2, \boldsymbol{\theta}\}$ that maximize the LML. After hyperparameter estimation,

inference can be carried out at an untried point, $\mathbf{x}_* \in \mathbb{R}^d$, giving the posterior distribution of the prediction $y_* \in \mathbb{R}$

$$\begin{aligned} y_* | \boldsymbol{\theta}, \sigma^2, \mathbf{X}, \mathbf{x}_* &\sim \mathcal{N}(\mathbb{E}[y_*], \mathbb{V}[y_*]), \\ \mathbb{E}[y_*] &= \mathbf{K}_{\mathbf{x}_*,\mathbf{X}} (\mathbf{K}_{\mathbf{X},\mathbf{X}} + \sigma^2 \mathbf{I}_n)^{-1} \mathbf{y}, \\ \mathbb{V}[y_*] &= \mathbf{K}_{\mathbf{x}_*,\mathbf{x}_*} - \mathbf{K}_{\mathbf{x}_*,\mathbf{X}} (\mathbf{K}_{\mathbf{X},\mathbf{X}} + \sigma^2 \mathbf{I}_n)^{-1} \mathbf{K}_{\mathbf{X},\mathbf{x}_*}. \end{aligned} \quad (2)$$

If we take a fully Bayesian (type I) approach, then we integrate out the hyperparameters by considering the hyperparameter posterior. This generally results in analytically intractable integrals which require techniques such as Markov-Chain Monte Carlo (MCMC) sampling, or one of its variants (Neal, 1997).

3. Eigenfunction Kernel Approximation

We consider a compact representation of the GP prior using a truncated Mercer expansion of the kernel k . We use the first p eigenfunctions which we approximate numerically using a Nyström approximation (Peng & Qi, 2015)

$$\begin{aligned} \tilde{k}(\mathbf{x}, \mathbf{z}) &= \sum_{i=1}^p \underbrace{(\lambda_i^{-\frac{1}{2}} \mathbf{K}_{\mathbf{x},\mathbf{U}} \mathbf{q}_i)}_{\phi_i(\mathbf{x})} \underbrace{(\lambda_i^{-\frac{1}{2}} \mathbf{K}_{\mathbf{z},\mathbf{U}} \mathbf{q}_i)}_{\phi_i(\mathbf{z})} \\ &= \mathbf{K}_{\mathbf{x},\mathbf{U}} \mathbf{Q} \mathbf{S}_p^T \boldsymbol{\Lambda}_p^{-1} \mathbf{S}_p \mathbf{Q}^T \mathbf{K}_{\mathbf{U},\mathbf{z}} \approx k(\mathbf{x}, \mathbf{z}), \end{aligned} \quad (3)$$

where $\mathbf{U} = \{\mathbf{u}_i \in \mathbb{R}^d\}_{i=1}^m$ refers to the set of m inducing point locations; $\boldsymbol{\Lambda}, \mathbf{Q} \in \mathbb{R}^{m \times m}$ are diagonal and unitary matrices containing the eigenvalues and eigenvectors of $\mathbf{K}_{\mathbf{U},\mathbf{U}}$, respectively; λ_i and \mathbf{q}_i denote the i th largest eigenvalue and corresponding eigenvector of $\mathbf{K}_{\mathbf{U},\mathbf{U}}$, respectively; $\mathbf{S}_p \in \mathbb{R}^{p \times m}$ is a sparse selection matrix where $\mathbf{S}_p(i, :)$ contains one value set to unity in the column corresponding to the index of the i th largest value on the diagonal of $\boldsymbol{\Lambda}$; and we use the shorthand notation $\boldsymbol{\Lambda}_p = \mathbf{S}_p \boldsymbol{\Lambda} \mathbf{S}_p^T = \text{diag}(\boldsymbol{\lambda}_p) \in \mathbb{R}^{p \times p}$ to denote a diagonal matrix containing the p largest eigenvalues of $\mathbf{K}_{\mathbf{U},\mathbf{U}}$, sorted in descending order. $\phi_i(\mathbf{x})$ is the numerical approximation of the i th eigenfunction evaluated at the input \mathbf{x} , scaled by the root of the i th eigenvalue. We only explicitly compute this scaled eigenfunction for numerical stability, as we will discuss later. Using the kernel \tilde{k} , the prior covariance matrix on the training set becomes

$$\tilde{\mathbf{K}}_{\mathbf{X},\mathbf{X}} = \underbrace{\mathbf{K}_{\mathbf{X},\mathbf{U}} \mathbf{Q} \mathbf{S}_p^T \boldsymbol{\Lambda}_p^{-\frac{1}{2}}}_{\boldsymbol{\Phi}} \underbrace{\boldsymbol{\Lambda}_p^{-\frac{1}{2}} \mathbf{S}_p \mathbf{Q}^T \mathbf{K}_{\mathbf{U},\mathbf{X}}}_{\boldsymbol{\Phi}^T}, \quad (4)$$

where the columns of $\boldsymbol{\Phi} \in \mathbb{R}^{n \times p}$ contain the p scaled eigenfunctions of our kernel evaluated on the training set. Observe that if \mathbf{U} is randomly sampled from \mathbf{X} , then $\tilde{\mathbf{K}}_{\mathbf{X},\mathbf{X}}$ is the same covariance matrix from the ‘‘Nyström method’’ of Williams & Seeger (2001), however, since we have replaced the kernel and not just the covariance matrix, we recover a valid probabilistic model (Peng & Qi, 2015).

While $\tilde{\mathbf{K}}_{X,X}$ has a rank of at most p , Peng & Qi (2015) show how a correction can be added to \tilde{k} (eq. (3)) to give a full rank covariance matrix (provided k does also). The resulting GP will be non-degenerate. We can write this correction as $\delta(\mathbf{x} - \mathbf{z})(k(\mathbf{x}, \mathbf{z}) - \tilde{k}(\mathbf{x}, \mathbf{z}))$, where $\delta(a) = 1$ if $a = 0$, else 0. This correction term does not affect the computational complexity of GP training, however, we find it does not generally improve performance over the unmodified \tilde{k} . We do not consider this correction in further discussion.

4. Grid-Structured Eigenfunctions (GRIEF)

Previous work employing Nyström approximations in kernel methods require m to be small (often $\ll n$) to yield computational benefits. As a result, the choice of inducing point locations, U , has a great influence on the approximation accuracy, and many techniques have been proposed to choose U effectively (Smola & Schölkopf, 2000; Drineas & Mahoney, 2005; Zhang et al., 2008; Belabbas & Wolfe, 2009; Kumar et al., 2012; Wang & Zhang, 2013; Gittens & Mahoney, 2013; Li et al., 2016; Musco & Musco, 2017). In this work, we would instead like to use *so many* inducing points that carefully optimizing the distribution of U is unnecessary. We will even consider $m \gg n$. The following result shows how an eigenfunction approximation can be improved by using many inducing points.

Theorem 1. *If the i th eigenvalue of k is simple and non-zero and $U \supset X$, a Nyström approximation of the i th kernel eigenfunction converges in the limit of large m ,*

$$\mathbf{q}_i^{(n)} = \lim_{m \rightarrow \infty} \sqrt{\frac{m}{n}} \frac{1}{\lambda_i^{(m)}} \mathbf{K}_{X,U} \mathbf{q}_i^{(m)}, \quad (5)$$

where $\lambda_i^{(m)} \in \mathbb{R}$ and $\mathbf{q}_i^{(m)} \in \mathbb{R}^m$ are the i th largest eigenvalue and corresponding eigenvector of $\mathbf{K}_{U,U}$, respectively. $\mathbf{q}_i^{(n)}$ is the kernel eigenfunction corresponding to the i th largest eigenvalue, evaluated on the set X .

Proof. We begin by constructing a Nyström approximation of the eigenfunction evaluated on U , using X as inducing points. From theorem 3.5 of (Baker, 1977), as $m \rightarrow \infty$,

$$\mathbf{q}_i^{(m)} = \sqrt{\frac{n}{m}} \frac{1}{\lambda_i^{(n)}} \mathbf{K}_{U,X} \mathbf{q}_i^{(n)}, \quad (6)$$

where we assume that the i th eigenvalue of k is simple and non-zero. Multiplying both sides by $\mathbf{K}_{X,U} \mathbf{K}_{U,U}^{-1}$,

$$\mathbf{K}_{X,U} \mathbf{K}_{U,U}^{-1} \mathbf{q}_i^{(m)} = \sqrt{\frac{n}{m}} \frac{1}{\lambda_i^{(n)}} \mathbf{K}_{X,U} \mathbf{K}_{U,U}^{-1} \mathbf{K}_{U,X} \mathbf{q}_i^{(n)}. \quad (7)$$

Since $U \supset X$, $\mathbf{K}_{X,U} = \mathbf{S}_n \mathbf{K}_{U,U}$ is a subset of the rows of $\mathbf{K}_{U,U}$, where $\mathbf{S}_n \in \mathbb{R}^{n \times m}$ is a selection matrix. We can write $\mathbf{K}_{X,U} \mathbf{K}_{U,U}^{-1} \mathbf{K}_{U,X} = \mathbf{S}_n \mathbf{K}_{U,U} \mathbf{K}_{U,U}^{-1} \mathbf{K}_{U,U} \mathbf{S}_n^T =$

$\mathbf{S}_n \mathbf{K}_{U,U} \mathbf{S}_n^T = \mathbf{K}_{X,X}$. Additionally, since the eigenvector $\mathbf{q}_i^{(m)}$ satisfies, $\mathbf{K}_{U,U}^{-1} \mathbf{q}_i^{(m)} = \frac{1}{\lambda_i^{(m)}} \mathbf{q}_i^{(m)}$, we get

$$\frac{1}{\lambda_i^{(m)}} \mathbf{K}_{X,U} \mathbf{q}_i^{(m)} = \sqrt{\frac{n}{m}} \frac{1}{\lambda_i^{(n)}} \mathbf{K}_{X,X} \mathbf{q}_i^{(n)}. \quad (8)$$

Noting that $\mathbf{K}_{X,X} \mathbf{q}_i^{(n)} = \lambda_i^{(n)} \mathbf{q}_i^{(n)}$ completes the proof. \square

For multiple eigenvalues, it can similarly be shown that the i th approximated eigenfunction converges to lie within the linear space of eigenfunctions corresponding to the i th eigenvalue of k as $m \rightarrow \infty$.

We can use a large m by distributing inducing points on a Cartesian tensor product grid¹. Saatçi (2011) demonstrated efficient GP inference when training points are distributed in this way by exploiting Kronecker matrix algebra. We will assume this grid structure for our inducing points, i.e. U will form a grid. If the covariance kernel satisfies the product correlation rule (as many popular multidimensional kernels do), i.e. $k(\mathbf{x}, \mathbf{z}) = \prod_{i=1}^d k_i(x_i, z_i)$, then $\mathbf{K}_{U,U} \in \mathbb{R}^{m \times m}$ inherits the Kronecker product form $\mathbf{K}_{U,U} = \bigotimes_{i=1}^d \mathbf{K}_{U,U}^{(i)}$, where \bigotimes is the Kronecker product (Van Loan, 2000). $\mathbf{K}_{U,U}^{(i)} \in \mathbb{R}^{\bar{m} \times \bar{m}}$ are one-dimensional kernel covariance matrices for a slice of the input space grid along the i th dimension, and $\bar{m} = \sqrt[d]{m} \approx \mathcal{O}(10)$ is the number of inducing points we choose along each dimension of the full grid. It is evident that the Kronecker product leads to a large, expanded matrix from smaller ones, therefore, it is very advantageous to manipulate and store these small matrices without “unpacking” them, or explicitly computing the Kronecker product. Exploiting this structure decreases the storage of $\mathbf{K}_{U,U}$ from $\mathcal{O}(m^2) \rightarrow \mathcal{O}(dm^{2/d}) = \mathcal{O}(d\bar{m}^2)$, and the cost of matrix-vector products with $\mathbf{K}_{U,U}$ from $\mathcal{O}(m^2) \rightarrow \mathcal{O}(dm^{(d+1)/d}) = \mathcal{O}(d\bar{m}^{d+1})$. Additionally, the cost of the eigen-decomposition of $\mathbf{K}_{U,U}$ decreases from $\mathcal{O}(m^3) \rightarrow \mathcal{O}(dm^{3/d}) = \mathcal{O}(d\bar{m}^3)$, and the eigenvector matrix $\mathbf{Q} = \bigotimes_{i=1}^d \mathbf{Q}^{(i)}$ and eigenvalue matrix $\mathbf{\Lambda} = \text{diag}(\bigotimes_{i=1}^d \lambda^{(i)})$ both inherit a Kronecker product structure, enabling matrix-vector products with $\tilde{\mathbf{K}}_{X,X}$ in $\mathcal{O}(d\bar{m}^{d+1})$ operations (Van Loan, 2000; Saatçi, 2011).

In low-dimensions, exploiting the Kronecker product structure of $\mathbf{K}_{U,U} = \bigotimes_{i=1}^d \mathbf{K}_{U,U}^{(i)}$ can be greatly advantageous, however, we can immediately see from the above complexities that the cost of matrix-vector products² with $\tilde{\mathbf{K}}_{X,X}$

¹ We want U to be sampled from the same distribution as the training data. Approximating the data distribution by placing U on a grid is easy to do by various means as a quick preprocessing step.

² We assume that a conjugate gradient method would be employed for GP training requiring matrix-vector products. Alternative formulations would require columns of $\mathbf{Q} = \bigotimes_{i=1}^d \mathbf{Q}^{(i)}$ to be expanded which similarly scales exponentially ($\mathcal{O}(\bar{m}^d)$).

increases *exponentially* in d . The storage requirements will similarly increase exponentially since a vector of length $m=\bar{m}^d$ needs to be stored when a matrix-vector product is made with $\mathbf{Q} = \bigotimes_{i=1}^d \mathbf{Q}^{(i)}$, and $\mathbf{K}_{X,U}$ requires $\mathcal{O}(\bar{m}^d n)$ storage. This poor scaling poses a serious impediment to the application of this approach to high-dimensional datasets.

We now show how to massively decrease time and storage requirements from exponential to *linear* in d by identifying further matrix structure in our problem.

We begin by identifying structure in the exact cross-covariance between train (or test) points and inducing points. These matrices, e.g. $\mathbf{K}_{X,U}$, admit a row-partitioned Khatri-Rao product structure as follows (Nickson et al., 2015)

$$\begin{aligned} \mathbf{K}_{X,U} &= \underset{i=1}{*} \mathbf{K}_{X,U}^{(i)} \\ &= \begin{pmatrix} \mathbf{K}_{X,U}^{(1)}(1,:) \otimes \mathbf{K}_{X,U}^{(2)}(1,:) \otimes \cdots \otimes \mathbf{K}_{X,U}^{(d)}(1,:) \\ \mathbf{K}_{X,U}^{(1)}(2,:) \otimes \mathbf{K}_{X,U}^{(2)}(2,:) \otimes \cdots \otimes \mathbf{K}_{X,U}^{(d)}(2,:) \\ \vdots \\ \mathbf{K}_{X,U}^{(1)}(n,:) \otimes \mathbf{K}_{X,U}^{(2)}(n,:) \otimes \cdots \otimes \mathbf{K}_{X,U}^{(d)}(n,:) \end{pmatrix}, \end{aligned} \quad (9)$$

where $*$ is the Khatri-Rao product whose computation gives a block Kronecker product matrix (Liu & Trenkler, 2008). We will always mention how Khatri-Rao product blocks are partitioned. Since $\mathbf{K}_{X,U}^{(i)}$ are only of size $n \times \bar{m}$, the storage of $\mathbf{K}_{X,U}$ has decreased from exponential to linear in d : $\mathcal{O}(\bar{m}^d n) \rightarrow \mathcal{O}(dn\bar{m}) \approx \mathcal{O}(dn)$. We also observe that the selection matrix $\mathbf{S}_p = \underset{i=1}{*} \mathbf{S}_p^{(i)}$ can be written as a row-partitioned Khatri-Rao product matrix where each sub-matrix contains one non-zero per row. Further, by exploiting both Kronecker and Khatri-Rao matrix algebra, our main result below shows that $\mathbf{K}_{X,U} \mathbf{Q} \mathbf{S}_p^T$ can be computed in $\mathcal{O}(dnp)$ time. This is a substantial reduction over the naive cost of $\mathcal{O}(\bar{m}^d np)$ time.

Theorem 2. *The product of a row-partitioned Khatri-Rao matrix $\mathbf{K}_{X,U} = \underset{i=1}{*} \mathbf{K}_{X,U}^{(i)} \in \mathbb{R}^{n \times \bar{m}^d}$, a Kronecker product matrix $\mathbf{Q} = \bigotimes_{i=1}^d \mathbf{Q}^{(i)} \in \mathbb{R}^{\bar{m}^d \times \bar{m}^d}$, and a column-partitioned Khatri-Rao matrix $\mathbf{S}_p^T = \underset{i=1}{*} (\mathbf{S}_p^{(i)})^T \in \mathbb{R}^{\bar{m}^d \times p}$ can be computed as follows*

$$\mathbf{K}_{X,U} \mathbf{Q} \mathbf{S}_p^T = \bigodot_{i=1}^d \mathbf{K}_{X,U}^{(i)} \mathbf{Q}^{(i)} (\mathbf{S}_p^{(i)})^T, \quad (10)$$

where \bigodot is the (element-wise) Hadamard product. This computation only requires products of the smaller matrices $\mathbf{K}_{X,U}^{(i)} \in \mathbb{R}^{n \times \bar{m}}$, $\mathbf{Q}^{(i)} \in \mathbb{R}^{\bar{m} \times \bar{m}}$ and $\mathbf{S}_p^{(i)} \in \mathbb{R}^{p \times \bar{m}}$.

Proof. First, observe that $\mathbf{K}_{X,U} \mathbf{Q} = \underset{i=1}{*} \mathbf{K}_{X,U}^{(i)} \mathbf{Q}^{(i)} = \underset{i=1}{*} \mathbf{R}^{(i)}$ is a row-partitioned Khatri-Rao product matrix using theorem 2 of (Liu & Trenkler, 2008). Now we must compute a matrix product of row- and column-partitioned

Khatri-Rao matrices. We observe that each element of this matrix product is an inner product between two Kronecker product vectors, i.e. $[\mathbf{K}_{X,U} \mathbf{Q} \mathbf{S}_p^T]_{ij} = (\bigotimes_{l=1}^d \mathbf{R}^{(l)}(i, :)) (\bigotimes_{l=1}^d \mathbf{S}_p^{(l)}(:, j)) = \prod_{l=1}^d \mathbf{R}^{(l)}(i, :) \mathbf{S}_p^{(l)}(:, j)$. Writing this in matrix form completes the proof. \square

If all the sub-matrices were dense, $\mathcal{O}(dn\bar{m}p)$ time would be required to compute $\mathbf{K}_{X,U} \mathbf{Q} \mathbf{S}_p^T$, however, since \mathbf{S}_p is a sparse selection matrix with one non-zero per row, computation requires just $\mathcal{O}(dn \max(p, \bar{m}c)) \approx \mathcal{O}(dnp)$ time. We achieve this time by computing only the necessary columns of $\mathbf{K}_{X,U}^{(i)} \mathbf{Q}^{(i)}$ and avoiding redundant computations. The constant c is the average number of non-zeros columns in each of $\{\mathbf{S}_p^{(i)}\}_{i=1}^d$ which is typically $\mathcal{O}(1)$, however, may be \bar{m} in the worst case. Evidently the time complexity is effectively *independent* of the number of inducing points. Even in the rare worst case where $c = \bar{m}$ and $\bar{m}^2 > p$, the scaling is extremely weak; $\mathcal{O}(dnm^{2/d})$.

What results is a kernel composed of basis eigenfunctions that are accurately approximated on a grid of inducing points using a Nyström approximation. Although m increases exponentially in d given this inducing point structure, the cost of GP training and inference is not affected. We call the resulting model GP-GRIEF (GP with GRId-structured EigenFunctions). Completing the computations required for GP training and inference require straightforward application of the matrix-determinant and matrix-inversion lemmas which we demonstrate later in eq. (13).

Eigenvalue Search What has not been addressed is how to form $\mathbf{S}_p = \underset{i=1}{*} \mathbf{S}_p^{(i)}$ and compute λ_p efficiently. This requires finding the index locations and values of the largest p eigenvalues in a vector of length $m = \bar{m}^d$. In high-dimensions, this task is daunting considering m can easily exceed the number of atoms in the observable universe. Fortunately, the resulting vector of eigenvalues, $\text{diag}(\boldsymbol{\Lambda}) = \bigotimes_{i=1}^d \boldsymbol{\lambda}^{(i)}$, has a Kronecker product structure which we can exploit to develop a fast search algorithm that requires only $\mathcal{O}(d\bar{m}p)$ time. To do this, we compute a truncated Kronecker product expansion by keeping only the p largest values after each sequential Kronecker product such that only Kronecker products between length p and length \bar{m} vectors are computed. Algorithm 1 outlines a more numerically stable version of this search strategy that computes the log of the eigenvalues and also demonstrates how \mathbf{S}_p is computed.

Computation in High Dimensions Direct use of theorem 2 may lead to finite-precision rounding inaccuracies and overflow errors in high dimensions because of the Hadamard product over d matrices. We can write a more numerically stable version of this algorithm by taking the log of eq. (10), allowing us to write the computation as a sum of d matrices,

Algorithm 1 Computes \mathbf{S}_p , and $\log \lambda_p$ (the log of the p largest eigenvalues of $\mathbf{K}_{\text{U,U}}$). We use zero-based array indexing, $\text{mod}(a, b)$ computes $a \bmod b$, $|\mathbf{a}|$ computes the length of \mathbf{a} , $\text{sort}_b(\mathbf{a})$ returns the $\min(|\mathbf{a}|, b)$ largest elements of \mathbf{a} in descending order, as well as the indices of these elements in \mathbf{a} , and $\lfloor \mathbf{a} \rfloor$ computes the floor of the elements in \mathbf{a} .

Input: $\{\boldsymbol{\lambda}^{(i)} \in \mathbb{R}^{\bar{m}}\}_{i=1}^d$
Output: $\{\mathbf{S}_p^{(i)} \in \mathbb{R}^{p \times \bar{m}}\}_{i=1}^d$ & $\log \lambda_p \in \mathbb{R}^p$
 $\log \lambda_p$, $\text{idxs} = \text{sort}_p(\log(\boldsymbol{\lambda}^{(1)}))$
for $i = 2$ **to** d **do**
 $\log \lambda_p$, $\text{ord} = \text{sort}_p(\log \lambda_p \otimes \mathbf{1}_{\bar{m}} + \mathbf{1}_{\log \lambda_p} \otimes \log(\boldsymbol{\lambda}^{(i)}))$
 $\text{idxs} = \left[\text{idxs}(\lfloor \text{ord} / \bar{m} \rfloor, :), \text{mod}(\text{ord}, \bar{m}) \right]$
end for
 $\{\mathbf{S}_p^{(i)} = \mathbf{I}_{\bar{m}}(\text{idxs}(:, i-1), :)\}_{i=1}^d$

rather than a product

$$\mathbf{K}_{\text{X,U}} \mathbf{Q} \mathbf{S}_p^T = \bigodot_{i=1}^d \text{sign}(\mathbf{B}^{(i)}) \odot \exp \left(\sum_{i=1}^d \log(\text{abs } \mathbf{B}^{(i)}) \right),$$

where $\mathbf{B}^{(i)} = \mathbf{K}_{\text{X,U}}^{(i)} \mathbf{Q}^{(i)} (\mathbf{S}_p^{(i)})^T$, and \exp , \log are computed element-wise. While the sign matrix is the Hadamard product of d matrices, it contains only $\{-1, 0, 1\}$ so it is not susceptible to numerical issues. Also, when the sign of an element is zero, we do not compute the log. Unfortunately, the \exp computation can still lead to numerical issues, however, Φ suffers less because of the rescaling provided by the eigenvalues (i.e. elements of Φ are the quotient of possibly very large or small values). Since all eigenvalues are positive, we can stably compute Φ as follows

$$\Phi = (\mathbf{K}_{\text{X,U}} \mathbf{Q} \mathbf{S}_p^T) \Lambda_p^{-\frac{1}{2}} = \bigodot_{i=1}^d \text{sign}(\mathbf{B}^{(i)}) \odot \exp \left(\sum_{i=1}^d \log(\text{abs } \mathbf{B}^{(i)}) - \frac{1}{2} \mathbf{1}_n \log \lambda_p^T \right),$$

where $\log \lambda_p$ is computed by algorithm 1.

4.1. Preconditioning Applications

As an aside remark, we discuss the application of $(\tilde{\mathbf{K}}_{\text{X,X}} + \sigma^2 \mathbf{I}_n)^{-1}$ as a preconditioner for the exact kernel matrix $\mathbf{K}_{\text{X,X}} + \sigma^2 \mathbf{I}_n$ in moderately sized problems where $\mathcal{O}(n^2)$ storage is not prohibitive. The use of $\tilde{\mathbf{K}}_{\text{X,X}}$ for matrix preconditioning was explored with notable empirical success by Cutajar et al. (2016) where a sub-set of training data was used as inducing points giving $\text{U} \subset \text{X}$ and $m < n$. By theorem 1, we know that the Nyström approximation converges for large m , and we have shown that we can accommodate $m \gg n$ to provide an accurate low-rank kernel matrix approximation.

4.2. SKI Applications

As a further aside, we discuss how the developed algebra can be applied in a general kernel interpolation setting. Wilson & Nickisch (2015) introduced a kernel interpolation perspective to unify inducing point methods wherein the kernel is interpolated from a covariance matrix on the inducing points. For instance, the subset of regressors (SoR) method (Silverman, 1985; Quiñonero-Candela & Rasmussen, 2005) can be viewed as a zero-mean GP interpolant of the true kernel while Wilson & Nickisch (2015) proposed a sparse approximate interpolant. We can denote the interpolated covariance matrix as $\mathbf{E} \mathbf{K}_{\text{U,U}} \mathbf{E}^T$, where $\mathbf{E} \in \mathbb{R}^{n \times m}$ is the interpolation matrix.

In structured kernel interpolation (SKI) the inducing points form a grid such that $\mathbf{K}_{\text{U,U}}$ inherits a Kronecker product form. This can provide dramatic computational advantages, however, SKI suffers from the exponential scaling discussed earlier and so is recommended only for very low-dimensional problems, $d \leq 5$ (Wilson et al., 2016). We observe that in the case of the GP interpolant (e.g. SoR), as well as the sparse interpolant suggested by Wilson & Nickisch (2015), the interpolation matrix \mathbf{E} inherits a row-partitioned Khatri-Rao structure. This enables direct use of theorem 2 to reduce the exponential scaling in d to a linear scaling, and allows SKI to scale to high-dimensional problems. However, time complexity would scale quadratically in n , unlike the proposed GRIEF methods.

5. Re-weighted Eigenfunction Kernel

We can approximately recover a wide class of kernels by modifying the weights associated with the kernel eigenfunctions (Buhmann, 2003). Here we consider this flexible kernel parametrization for the GRIEF kernel. Extending eq. (3), we can write the re-weighted GRIEF kernel as

$$\tilde{k}(\mathbf{x}, \mathbf{z}) = \sum_{i=1}^p w_i \phi_i(\mathbf{x}) \phi_i(\mathbf{z}), \quad (11)$$

where $\mathbf{W} \in \mathbb{R}^{p \times p} = \text{diag}(\mathbf{w})$, $\mathbf{w} = \{w_i > 0\}_{i=1}^p$ are the eigenfunction weights. The covariance matrix then becomes

$$\tilde{\mathbf{K}}_{\text{X,X}} = \Phi \mathbf{W} \Phi^T, \quad (12)$$

and the full set of hyperparameters is $\{\sigma^2, \mathbf{w}, \boldsymbol{\theta}\}$, however, we choose to fix $\boldsymbol{\theta}$ so that Φ remains constant. We can then compute the log marginal likelihood (LML) and all $p+1$ derivatives with respect to $\{\sigma^2, \mathbf{w}\}$ in $\mathcal{O}(p)$ which is independent of n . To do this, we first assume that $\mathbf{y}^T \mathbf{y} \in \mathbb{R}$, $\Phi^T \mathbf{y} = \mathbf{r} \in \mathbb{R}^p$, and $\Phi^T \Phi = \mathbf{A} \in \mathbb{R}^{p \times p}$ are precomputed, which requires $\mathcal{O}(np^2 + dnp + d\bar{m}^3)$ time, however, this step only needs to be done once before LML iterations begin. Then, to compute the LML (eq. (1)), we use the

matrix inversion and determinant lemmas to give

$$\mathbf{y}^T (\tilde{\mathbf{K}}_{X,X} + \sigma^2 \mathbf{I}_n)^{-1} \mathbf{y} = \sigma^{-2} (\mathbf{y}^T \mathbf{y} - \mathbf{r}^T \mathbf{P}^{-1} \mathbf{r}), \quad (13)$$

$$\log |\tilde{\mathbf{K}}_{X,X} + \sigma^2 \mathbf{I}_n| = \log |\mathbf{P}| + \sum_{i=1}^p \log w_i + (n-p) \log \sigma^2,$$

where $\mathbf{P} \in \mathbb{R}^{p \times p} = \sigma^2 \mathbf{W}^{-1} + \mathbf{A}$. Using these relations, the LML can be computed within $\mathcal{O}(p^3)$ time. The LML derivatives with respect to all hyperparameters can also be computed in $\mathcal{O}(p^3)$ as shown in the following expression which is derived in appendix A of the supplement

$$\frac{\partial \text{LML}}{\partial \mathbf{w}} = \frac{(\mathbf{r} - \mathbf{A} \mathbf{P}^{-1} \mathbf{r})^2}{2\sigma^4} - \frac{\text{diag}(\mathbf{A}) - (\mathbf{A} \odot \mathbf{P}^{-1} \mathbf{A})^T \mathbf{1}_p}{2\sigma^2},$$

$$\frac{\partial \text{LML}}{\partial \sigma^2} = \frac{\mathbf{y}^T \mathbf{y} - 2\mathbf{r}^T \mathbf{P}^{-1} \mathbf{r} + \mathbf{r}^T \mathbf{P}^{-1} \mathbf{A} \mathbf{P}^{-1} \mathbf{r}}{2\sigma^4} - \frac{n - \text{Tr}(\mathbf{P}^{-1} \mathbf{A})}{2\sigma^2}.$$

To further reduce the per-iteration computational complexity from $\mathcal{O}(p^3) \rightarrow \mathcal{O}(p)$ we apply a linear transformation to the basis functions to make them mutually orthogonal when evaluated on the training data. We can write the i th transformed basis function as $\tilde{\phi}_i(\mathbf{x}) = \sum_{j=1}^p \tilde{v}_{ji} \tilde{\Sigma}_{ii}^{-1} \phi_j(\mathbf{x})$, where $\tilde{\Sigma} \in \mathbb{R}^{\tilde{p} \times \tilde{p}}$ is a diagonal matrix containing the non-zero singular values of Φ , $\tilde{\mathbf{V}} \in \mathbb{R}^{p \times \tilde{p}}$ contains the corresponding right-singular vectors of Φ , and $\tilde{p} \leq \min(p, n)$. Using these transformed basis functions, both $\mathbf{A} = \mathbf{I}_{\tilde{p}}$ and $\mathbf{P} = \sigma^2 \mathbf{W}^{-1} + \mathbf{A}$ become diagonal matrices, enabling evaluation of the LML and all derivatives in $\mathcal{O}(p)$ using eq. (13) and the derivative expressions above. The transformation requires the singular-value decomposition of Φ before LML iterations, however, this precomputation is no more expensive than those discussed previously at $\mathcal{O}(n\tilde{p}^2)$.

Since the kernel is now heavily parametrized, maximizing the LML for type-II Bayesian inference is susceptible to overfitting. Instead, we may choose to take a fully Bayesian type-I approach and integrate out the hyperparameters using hybrid MCMC sampling. This type-I approach requires far more LML evaluations than type-II (typically $\mathcal{O}(10^5)$), however, the fast $\mathcal{O}(p)$ (or $\mathcal{O}(p^3)$) evaluations make this tractable even for very large problems since the cost per iteration is *independent* of the number of training points. In appendix B of the supplement we discuss further extensions to this type-I inference procedure.

6. Experimental Studies

Two-Dimensional Visualization Figure 1 shows a comparison between GP-GRIEF and the Variational Free Energy (VFE) inducing point approximation (Titsias, 2009) on a two-dimensional test problem with $n = 10$ training points generated by the function $f(\mathbf{x}) = \sin(x_1) \sin(x_2)$ and corrupted with $\mathcal{N}(0, 0.1)$ noise. For both models, we use a squared-exponential base kernel, and we estimate the kernel lengthscale and noise variance, σ^2 , by maximizing the

log marginal likelihood. VFE can also select its inducing point locations. In this study, we do not consider optimizing the GRIEF weights, \mathbf{w} . VFE with $m = 4$ inducing points achieves a root-mean squared error (RMSE) of 0.47 on the test set, whereas GP-GRIEF with the same number of basis functions³, $p = 4$, achieves an RMSE of 0.34, identical to the reconstruction provided by a full GP using the exact kernel. While GP-GRIEF uses a dense grid of $m = 25$ inducing points, it has a computational complexity equivalent to VFE. This demonstrates the reconstruction power of GP-GRIEF, even when very few eigenfunctions are considered.

Kernel Reconstruction Accuracy We compare the kernel covariance reconstruction error of the GRIEF Nyström method to competing techniques in fig. 2. We sample 5000 points from $\mathcal{U}(-\sqrt{3}, \sqrt{3})$ in $d = 100$ dimensions, randomly taking half for training and half for testing, and we consider a squared-exponential kernel. Given only the training set, we attempt to reconstruct the exact prior covariance matrices between the training set (fig. 2a), and the joint train/test set (2b). This allows us to study the kernel reconstruction accuracy between points both within, and beyond the training set. In both studies, the proposed GRIEF Nyström method greatly outperforms a random Fourier features reconstruction (Rahimi & Recht, 2007), and a randomized Nyström reconstruction where $m = p$ inducing points are uniformly sampled from the training set. We emphasize that both randomized Nyström and GRIEF Nyström have the same computational complexity for a given number of basis functions, p , even though the GRIEF Nyström method uses $m = 10^{200}$ inducing points ($\bar{m} = 100$).

In the joint train/test study of fig. 2b, we observe a larger gap between GRIEF Nyström and randomized Nyström than in fig. 2a. This is not surprising since the goal of the randomized Nyström method (and indeed all existing extensions to this technique) is to improve the accuracy of eigenfunctions evaluated on the training set which does not guarantee performance of the eigenfunctions evaluated on points outside this set. For example, if a test point is placed far from the training set then we expect a poor approximation from existing Nyström methods. However, our GRIEF Nyström approach attempts to fill out the input space with inducing points everywhere, not just near training points. This guarantees an accurate approximation at test locations even if they are far from training data. Comparatively, the random Fourier features technique samples from a distribution that is independent of the training data, so it is also expected to perform no worse on the joint set than the training set. However, we observe that it provides an equally poor reconstruction on both sets.

The black curves in fig. 2 show the exact eigen-

³Technically, VFE gives an infinite basis function expansion through a correction term, however, we will assume $p = m = 4$.




Figure 1: Reconstruction using GP-GRIEF outperforms VFE. Both techniques use an equal number of basis functions and have the same computational complexity. Crosses denote training point positions and dots denote inducing point locations.

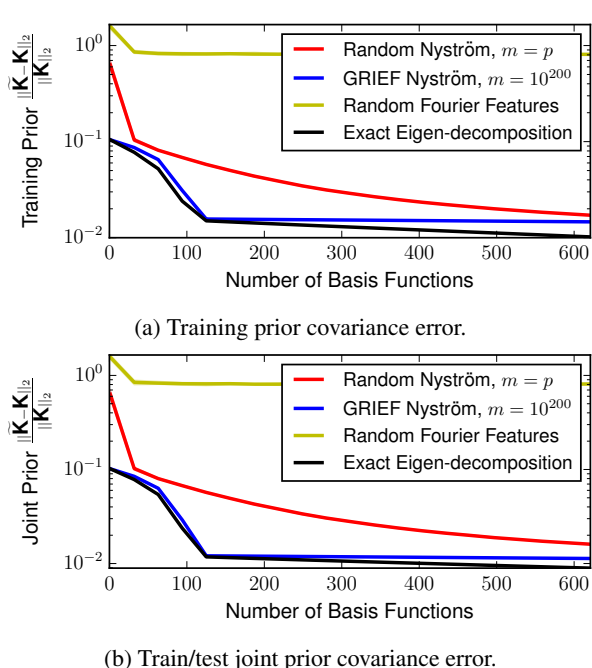


Figure 2: Covariance matrix reconstruction error of GP-GRIEF beats randomized Nyström with uniform sampling averaged over 10 samples. GP-GRIEF approaches the optimal reconstruction accuracy of the black curve.

decomposition of the covariance matrices which demonstrates the optimal reconstruction accuracy for a kernel approximation with a given number of basis functions. We observe that GP-GRIEF approaches this optimal accuracy in both studies, even though the test point distribution is not known at training time.

UCI Regression Studies We next assess performance on real-world regression datasets from the UCI repository. Using the authors’ code⁴, we report the mean and standard deviation

⁴https://github.com/treforevans/gp_grief

of the RMSE from 10-fold cross validation⁵. Also presented is the mean training time per fold on a machine with two E5-2680 v3 processors. We use a squared-exponential kernel with automatic relevance determination (SE-ARD) and we compare our test errors to those reported by Yang et al. (2015) using type-II inference on the same train-test splits. Yang et al. (2015) used an exact GP with an SE-ARD kernel for datasets with $n < 2000$, and Fastfood expansions were used to approximate the SE-ARD kernel for the larger datasets ($n > 2000$).

Before training the GP-GRIEF models, we initialize the base kernel hyperparameters, θ , by maximizing the marginal likelihood of an exact GP constructed on $\min(n, 1000)$ points randomly selected from the dataset. We then train GP-GRIEF-II or GP-GRIEF-I which differ by not only type-II or type-I inference but also the kernel parametrizations used. GP-GRIEF-II uses the kernel from eq. (3) which is parametrized by the base kernel hyperparameters: $\{\theta, \sigma^2\}$. The presented training time includes log marginal likelihood maximization to estimate the hyperparameters, beginning with the initialized values. GP-GRIEF-I uses the kernel from eq. (11) which is parametrized by the basis function weights: $\{\mathbf{w}, \sigma^2\}$; the base kernel hyperparameters, θ , are fixed to the initialized values. We integrate out $\{\mathbf{w}, \sigma^2\}$ using Metropolis adjusted Langevin dynamics MCMC which uses gradient information (Girolami & Calderhead, 2011). The training time includes MCMC sampling, which we run for 10000 iterations. We use log-normal priors with {mode, variance} of {1, 100} for \mathbf{w} , and $\{\sigma_0^2, 0.04\}$ for σ^2 , where σ_0^2 is the initialized value. We begin sampling at the prior mode, burning the first 1000 samples and thinning every 50 thereafter. For datasets with $n > 10^6$ we use the $\mathcal{O}(p)$ LML computations described in section 5. For all studies, we fix $\bar{m} = 10$, and we fix $p = 1000$ for GP-GRIEF-I. For

⁵90% train, 10% test per fold. We use folds from <https://people.orie.cornell.edu/andrew/code>

Scalable Gaussian Processes with Grid-Structured Eigenfunctions

Dataset	n	d	$m=\bar{m}^d$	p	Time (hrs)	GP-GRIEF-II	p	Time (hrs)	GP-GRIEF-I	Yang et al. (2015)
challenger	23	4	10^4	10	0	0.554 ± 0.277	1000	0.178	0.519 ± 0.261	0.63 ± 0.26
fertility	100	9	10^9	100	0.001	0.172 ± 0.055	1000	0.66	0.166 ± 0.051	0.21 ± 0.05
slump	103	7	10^7	100	0	3.972 ± 1.891	1000	0.566	3.470 ± 1.712	4.72 ± 2.42
automobile	159	25	10^{25}	100	0.007	0.145 ± 0.057	1000	0.604	0.111 ± 0.036	0.18 ± 0.07
servo	167	4	10^4	100	0	0.280 ± 0.085	1000	0.265	0.268 ± 0.075	0.28 ± 0.09
cancer	194	33	10^{33}	100	0.007	27.843 ± 3.910	1000	0.667	30.568 ± 3.340	35 ± 4
hardware	209	7	10^7	100	0	0.408 ± 0.046	1000	0.637	0.402 ± 0.045	0.43 ± 0.04
yacht	308	6	10^6	100	0.001	0.170 ± 0.083	1000	0.595	0.120 ± 0.070	0.16 ± 0.11
autompg	392	7	10^7	100	0.001	2.607 ± 0.356	1000	0.594	2.563 ± 0.369	2.63 ± 0.38
housing	506	13	10^{13}	100	0.004	3.212 ± 0.864	1000	0.62	2.887 ± 0.489	2.91 ± 0.54
forest	517	12	10^{12}	100	0.001	1.386 ± 0.14	1000	0.621	1.384 ± 0.139	1.39 ± 0.16
stock	536	11	10^{11}	100	0.001	0.005 ± 0.000	1000	0.567	0.005 ± 0.000	0.005 ± 0.001
energy	768	8	10^8	100	0.002	0.49 ± 0.057	1000	0.622	0.461 ± 0.064	0.46 ± 0.07
concrete	1030	8	10^8	1000	0.008	5.232 ± 0.723	1000	0.57	5.156 ± 0.766	4.95 ± 0.77
solar	1066	10	10^{10}	1000	0.003	0.786 ± 0.198	1000	0.628	0.809 ± 0.193	0.83 ± 0.20
wine	1599	11	10^{11}	1000	0.012	0.483 ± 0.052	1000	0.583	0.477 ± 0.047	0.47 ± 0.08
skillcraft	3338	19	10^{19}	1000	0.011	0.248 ± 0.016	1000	0.573	0.248 ± 0.016	0.25 ± 0.02
pumadyn	8192	32	10^{32}	1000	0.156	0.20 ± 0.00	1000	0.645	0.212 ± 0.004	0.20 ± 0.00
elevators	16599	18	10^{18}	1000	0.283	0.091 ± 0.002	1000	0.664	0.097 ± 0.001	0.090 ± 0.001
kin40k	40000	8	10^8	1000	0.38	0.206 ± 0.004	1000	0.649	0.206 ± 0.004	0.28 ± 0.01
keggu	63608	27	10^{27}	1000	3.642	0.118 ± 0.003	1000	0.704	0.134 ± 0.005	0.12 ± 0.00
3droad	434874	3	10^3	1000	0.869	11.648 ± 0.281	1000	0.221	12.966 ± 0.077	10.91 ± 0.05
electric	2049280	11	10^{11}	1000	8.019	0.064 ± 0.002	1000	0.418	0.058 ± 0.006	0.12 ± 0.12

Table 1: Mean and standard deviation of test error and average training time (including hyperparameter estimation or MCMC sampling) from 10-fold cross validation (90% train, 10% test per fold) on UCI regression datasets.

GP-GRIEF-II, we make p proportional to n by rounding n down to the nearest power of ten, or take 1000 if it is lesser, i.e. $p = \min(1000, 10^{\lfloor \log_{10} n \rfloor})$.

It is firstly evident that both GP-GRIEF-I and GP-GRIEF-II outperform the exact GP presented by Yang et al. (2015) on nearly every small dataset ($n < 2000$). In particular, GP-GRIEF-I performs extremely well on these small datasets as we would expect since it uses a very flexible kernel and is robust to over-fitting as a result of the principled Bayesian approach employed. On larger datasets, where we expect the hyperparameter posterior to be more peaked, we see that the type-II techniques begin to be competitive. On these larger datasets, both GP-GRIEF techniques show comparable test error to Yang et al. (2015) on all datasets but perform considerably better on kin40k and the electric dataset with two-million training points. With respect to time, we note that the GP-GRIEF-I model trained extremely rapidly considering a fully Bayesian approach was taken; only 25 minutes were required for the two-million point electric dataset even though this size is prohibitive for most GP models taking a type-II empirical Bayes approach.

The independence of computational complexity on m allows enormous numbers of inducing points to be used. We use $m=10^{33}$ inducing points for the cancer dataset which demonstrates the efficiency of the matrix algebra employed

since storing a double-precision vector of this length requires 8 quadrillion exabytes; far more than all combined hard-disk space on earth.

7. Conclusion

Our new technique, GP-GRIEF, has been outlined along with promising initial results on large real-world datasets where we demonstrated GP training and inference in $\mathcal{O}(p)$ time with $\mathcal{O}(p)$ storage. This fast training enables type-I Bayesian inference to remain computationally attractive even for very large datasets as we had shown in our studies. We showed that our complexities are independent of m , allowing us to break the *curse of dimensionality* inherent to methods that manipulate distributions of points on a full Cartesian product grid. Asymptotic results were also presented to show why a choice of large m is important to provide an accurate global kernel approximation. Lastly, we considered the use of up to 10^{33} inducing points in our regression studies, demonstrating the efficiency of the matrix algebra employed. We discussed how the developed algebra can be used in areas beyond the focus of the numerical studies, such as in a general kernel interpolation framework, or in general kernel matrix preconditioning applications. However, it will be interesting to explore what other applications could exploit the developed matrix algebra techniques.

Acknowledgements

Research funded by an NSERC Discovery Grant and the Canada Research Chairs program.

References

- Baker, C. T. H. *The numerical treatment of integral equations*. Oxford: Clarendon press, 1977.
- Belabbas, M.-A. and Wolfe, P. J. Spectral methods in machine learning: New strategies for very large datasets. In *National Academy of Sciences of the USA*, number 106, pp. 369–374, 2009.
- Bilionis, I. PredictiveScience lab: py-mcmc. <https://github.com/PredictiveScienceLab/py-mcmc>, December 2014.
- Buhmann, M. D. *Radial basis functions: theory and implementations*. Cambridge university press, 2003.
- Cutajar, K., Osborne, M. A., Cunningham, J. P., and Filippone, M. Preconditioning kernel matrices. In *International Conference on Machine Learning*, pp. 2529–2538, 2016.
- Drineas, P. and Mahoney, M. W. On the Nyström method for approximating a Gram matrix for improved kernel-based learning. *Journal of Machine Learning Research*, (6):2153–2175, 2005.
- Girolami, M. and Calderhead, B. Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2):123–214, 2011.
- Gittens, A. and Mahoney, M. W. Revisiting the Nyström method for improved large-scale machine learning. In *International Conference on Machine Learning*, 2013.
- GPY. GPY: A gaussian process framework in python. <http://github.com/SheffieldML/GPY>, since 2012.
- Kumar, S., Mohri, M., and Talwalkar, A. Sampling methods for the Nyström method. *Journal of Machine Learning Research*, (13):981–1006, 2012.
- Lázaro-Gredilla, M., Quiñero-Candela, J., Rasmussen, C. E., and Figueiras-Vidal, A. Sparse spectrum Gaussian process regression. *Journal of Machine Learning Research*, 11(6):1865–1881, 2010.
- Li, C., Jegelka, S., and Sra, S. Fast DPP sampling for Nyström with application to kernel methods. In *International Conference on Machine Learning*, 2016.
- Liu, S. and Trenkler, G. Hadamard, Khatri-Rao, Kronecker and other matrix products. *International Journal of Information and Systems Sciences*, 4(1):160–177, 2008.
- Musco, C. and Musco, C. Recursive sampling for the Nyström method. In *Advances in Neural Information Processing Systems*, 2017.
- Neal, R. M. Monte Carlo implementation of Gaussian process models for Bayesian regression and classification. Technical Report 9702, University of Toronto, 1997.
- Nickson, T., Gunter, T., Lloyd, C., Osborne, M. A., and Roberts, S. Blitzkriging: Kronecker-structured stochastic Gaussian processes. *arXiv preprint arXiv:1510.07965*, 2015.
- Peng, H. and Qi, Y. EigenGP: Gaussian process models with adaptive eigenfunctions. In *International Joint Conference on Artificial Intelligence*, pp. 3763–3769, 2015.
- Pinheiro, J. C. and Bates, D. M. Unconstrained parametrizations for variance-covariance matrices. *Statistics and computing*, 6(3):289–296, 1996.
- Quiñero-Candela, J. and Rasmussen, C. E. A unifying view of sparse approximate Gaussian process regression. *Journal of Machine Learning Research*, 6(12):1939–1959, 2005.
- Rahimi, A. and Recht, B. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems*, pp. 1177–1184, 2007.
- Rasmussen, C. E. and Williams, C. K. I. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- Saatçi, Y. *Scalable inference for structured Gaussian process models*. PhD thesis, University of Cambridge, 2011.
- Silverman, B. W. Some aspects of the spline smoothing approach to non-parametric regression curve fitting. *Journal of the Royal Statistical Society B*, 47(1):1–52, 1985.
- Smola, A. J. and Bartlett, P. Sparse greedy Gaussian process regression. In *Advances in Neural Information Processing Systems*, pp. 619–625, 2000.
- Smola, A. J. and Schölkopf, B. Sparse greedy matrix approximation for machine learning. In *International Conference on Machine Learning*, pp. 911–918, 2000.
- Snelson, E. and Ghahramani, Z. Sparse Gaussian processes using pseudo-inputs. In *Advances in Neural Information Processing Systems*, pp. 1257–1264, 2006.
- Titsias, M. Variational learning of inducing variables in sparse Gaussian processes. In *Artificial Intelligence and Statistics*, pp. 567–574, 2009.
- Van Loan, C. F. The ubiquitous Kronecker product. *Journal of Computational and Applied Mathematics*, 123(1):85–100, 2000.

- Wang, S. and Zhang, Z. Improving CUR matrix decomposition and the Nyström approximation via adaptive sampling. *Journal of Machine Learning Research*, (14): 2729–2769, 2013.
- Williams, C. K. I. and Seeger, M. Using the Nyström method to speed up kernel machines. In *Advances in Neural Information Processing Systems*, pp. 682–688, 2001.
- Wilson, A. G. and Nickisch, H. Kernel interpolation for scalable structured Gaussian processes (KISS-GP). In *International Conference on Machine Learning*, pp. 1775–1784, 2015.
- Wilson, A. G., Hu, Z., Salakhutdinov, R., and Xing, E. P. Deep kernel learning. In *Artificial Intelligence and Statistics*, pp. 370–378, 2016.
- Yang, Z., Smola, A. J., Song, L., and Wilson, A. G. À la carte – learning fast kernels. In *Artificial Intelligence and Statistics*, pp. 1098–1106, 2015.
- Zhang, K., Tsang, I. W., and Kwok, J. T. Improved Nyström low-rank approximation and error analysis. In *International Conference on Machine Learning*, pp. 1232–1239, 2008.

A. Re-Weighted Eigenfunction Kernel Derivatives

Calculating the $p + 1$ derivatives of the log marginal likelihood (LML) with respect to $\{\mathbf{w}, \sigma^2\}$ using finite difference approximations would require $\mathcal{O}(p^4)$ time. We show how all these derivatives can be analytically computed in $\mathcal{O}(p^3)$ time. We also discuss how the use of transformed basis functions (i.e. replacing Φ with $\tilde{\Phi}$) allows derivative computations in $\mathcal{O}(p)$.

The LML can be written as follows

$$\log \mathcal{P}(\mathbf{y}|\boldsymbol{\theta}, \sigma^2, \mathbf{X}) = -\frac{n}{2} \log(2\pi) - \frac{1}{2} \underbrace{\log |\mathbf{K}_{X,X} + \sigma^2 \mathbf{I}_n|}_{\text{Complexity}} - \frac{1}{2} \underbrace{\mathbf{y}^T \boldsymbol{\alpha}}_{\text{Data Fit}},$$

where $\boldsymbol{\alpha} = (\mathbf{K}_{X,X} + \sigma^2 \mathbf{I}_n)^{-1} \mathbf{y}$. For clarity, we will derive the gradients of the complexity and data-fit terms separately.

Data-Fit Weight Derivatives First we show how the derivative of the data-fit term can be computed within this time

$$\frac{\partial \mathbf{y}^T \boldsymbol{\alpha}}{\partial w_i} = -\boldsymbol{\alpha}^T \frac{\partial \tilde{\mathbf{K}}}{\partial w_i} \boldsymbol{\alpha} = -(\phi_i^T \boldsymbol{\alpha})^2$$

where $\boldsymbol{\alpha} = (\tilde{\mathbf{K}}_{X,X} + \sigma^2 \mathbf{I}_n)^{-1} \mathbf{y} \in \mathbb{R}^n$, and we make the observation that $\frac{\partial \tilde{\mathbf{K}}}{\partial w_i} = \phi_i \phi_i^T$. We can vectorize this to compute all data-fit derivatives

$$\begin{aligned} \frac{\partial \mathbf{y}^T \boldsymbol{\alpha}}{\partial \mathbf{w}} &= -(\Phi^T \boldsymbol{\alpha})^2, \\ &= -\sigma^{-4} (\Phi^T \mathbf{y} - \mathbf{A} \mathbf{P}^{-1} \Phi^T \mathbf{y})^2, \end{aligned}$$

where $\mathbf{A} = \Phi^T \Phi \in \mathbb{R}^{p \times p}$ and $\Phi^T \mathbf{y} \in \mathbb{R}^p$ are both precomputed before LML iterations begin, and $\mathbf{P} = \sigma^2 \mathbf{W}^{-1} + \mathbf{A} \in \mathbb{R}^{p \times p}$ is also required to compute the LML (see eq. (13)) so it is already computed and factorized. Evidently, the data-fit term derivatives can be computed in $\mathcal{O}(p^3)$ at each LML iteration.

Complexity Term Weight Derivatives Now we derive the complexity term gradient.

$$\frac{\partial \log |\tilde{\mathbf{K}} + \sigma^2 \mathbf{I}_n|}{\partial w_i} = \text{Tr} \left[(\tilde{\mathbf{K}} + \sigma^2 \mathbf{I}_n)^{-1} \frac{\partial \tilde{\mathbf{K}}}{\partial w_i} \right],$$

using $\frac{\partial \tilde{\mathbf{K}}}{\partial w_i} = \phi_i \phi_i^T$ and the cyclic permutation invariance of the trace operation, we get

$$\frac{\partial \log |\tilde{\mathbf{K}} + \sigma^2 \mathbf{I}_n|}{\partial w_i} = \phi_i^T (\tilde{\mathbf{K}} + \sigma^2 \mathbf{I}_n)^{-1} \phi_i.$$

Using the matrix inversion lemma, the preceding equation becomes

$$\begin{aligned} \frac{\partial \log |\tilde{\mathbf{K}} + \sigma^2 \mathbf{I}_n|}{\partial w_i} &= \sigma^{-2} (\phi_i^T \phi_i - \phi_i^T \Phi \mathbf{P}^{-1} \Phi^T \phi_i), \\ &= \sigma^{-2} (a_{ii} - \mathbf{a}_i^T \mathbf{P}^{-1} \mathbf{a}_i). \end{aligned}$$

Evidently, the complexity term derivatives each require $\mathcal{O}(p^2)$ time so all p derivatives can be computed in $\mathcal{O}(p^3)$. We can write the vectorized computation as

$$\frac{\partial \log |\tilde{\mathbf{K}} + \sigma^2 \mathbf{I}_n|}{\partial \mathbf{w}} = \sigma^{-2} [\text{diag}(\mathbf{A}) - (\mathbf{A} \odot \mathbf{P}^{-1} \mathbf{A})^T \mathbf{1}_p],$$

where it is evident that the dominating expense $\mathbf{P}^{-1} \mathbf{A}$ is also required for the data-fit derivatives.

Noise Variance Derivatives We show here how the derivatives of the LML with respect to σ^2 can be computed in $\mathcal{O}(n + p)$, as follows

$$\begin{aligned} \frac{\partial \mathbf{y}^T \boldsymbol{\alpha}}{\partial \sigma^2} &= -\boldsymbol{\alpha}^T \frac{\partial (\sigma^2 \mathbf{I}_n)}{\partial \sigma^2} \boldsymbol{\alpha} = -\boldsymbol{\alpha}^T \boldsymbol{\alpha}, \\ &= -\sigma^{-4} (\mathbf{y}^T \mathbf{y} - 2\mathbf{y}^T \Phi \mathbf{P}^{-1} \Phi^T \mathbf{y} + \mathbf{y}^T \Phi \mathbf{P}^{-1} \mathbf{A} \mathbf{P}^{-1} \Phi^T \mathbf{y}), \end{aligned}$$

and

$$\begin{aligned} \frac{\partial \log |\tilde{\mathbf{K}} + \sigma^2 \mathbf{I}_n|}{\partial \sigma^2} &= \text{Tr} \left((\tilde{\mathbf{K}} + \sigma^2 \mathbf{I}_n)^{-1} \frac{\partial (\sigma^2 \mathbf{I}_n)}{\partial \sigma^2} \right) \\ &= \text{Tr} \left(\sigma^{-2} [\mathbf{I}_n - \Phi \mathbf{P}^{-1} \Phi^T] \right) \\ &= \sigma^{-2} [n - \text{Tr}(\mathbf{P}^{-1} \mathbf{A})]. \end{aligned}$$

Evidently the first relation can be computed in $\mathcal{O}(p^3)$ if $\mathbf{y}^T \mathbf{y}$, $\Phi^T \mathbf{y}$, and \mathbf{A} are precomputed, and the second relation can be computed in $\mathcal{O}(p)$ since the matrix product $\mathbf{P}^{-1} \mathbf{A}$ has already been explicitly computed to compute the derivatives with respect to \mathbf{w} .

Final Expressions Combining the derived expressions for the derivatives of the LML with respect to the $p + 1$ hyperparameters $\{\mathbf{w}, \sigma^2\}$, it is evident that all computations can be performed in $\mathcal{O}(p^3)$. We can write the final expressions as follows

$$\begin{aligned} \frac{\partial \text{LML}}{\partial \mathbf{w}} &= \frac{(\mathbf{r} - \mathbf{A} \mathbf{P}^{-1} \mathbf{r})^2}{2\sigma^4} - \frac{\text{diag}(\mathbf{A}) - (\mathbf{A} \odot \mathbf{P}^{-1} \mathbf{A})^T \mathbf{1}_p}{2\sigma^2}, \\ \frac{\partial \text{LML}}{\partial \sigma^2} &= \frac{\mathbf{y}^T \mathbf{y} - 2\mathbf{r}^T \mathbf{P}^{-1} \mathbf{r} + \mathbf{r}^T \mathbf{P}^{-1} \mathbf{A} \mathbf{P}^{-1} \mathbf{r}}{2\sigma^4} - \frac{n - \text{Tr}(\mathbf{P}^{-1} \mathbf{A})}{2\sigma^2}, \end{aligned}$$

where $\mathbf{r} = \Phi^T \mathbf{y} \in \mathbb{R}^p$. If we transform the basis functions by replacing Φ with $\tilde{\Phi}$ then it is evident that the derivative computations can be made in $\mathcal{O}(p)$ since both $\mathbf{A} = \tilde{\Phi}^T \tilde{\Phi} = \mathbf{I}_{\tilde{p}}$ and $\mathbf{P} = \sigma^2 \mathbf{W}^{-1} + \mathbf{A}$ will be diagonal.

B. Type-I Inference Extensions

Here we discuss extensions to the type-I inference approach described in section 5 for the re-weighted eigenfunction kernel. Namely, we consider more flexible kernel parameterizations that can be used while still admitting MCMC iterations with a complexity independent of the size of the training set, n . To begin, consider the re-weighted kernel in eq. (11) from a weight space perspective (Rasmussen & Williams, 2006). In this case, we are constructing the generalized linear model $\sum_{i=1}^p \alpha_i \phi_i(\mathbf{x})$, where $\alpha \in \mathbb{R}^p$. We assume the observed responses are corrupted by independent Gaussian noise with variance σ^2 , and we specify the prior $\alpha \sim \mathcal{N}(\mathbf{0}, \mathbf{W})$. This model is identical to the GP specified in section 5, which was introduced from a function space perspective. There we had shown that we can specify a hyper-prior on \mathbf{w} and perform MCMC with a complexity independent of n .

We had previously taken \mathbf{W} to be diagonal in section 5 which assumes no prior correlation between the basis functions. The technique can easily be extended by taking \mathbf{W} to be dense and symmetric positive-definite. Pinheiro & Bates (1996) discuss how such a matrix could be parameterized. In this case, it is easily observed that computation of the log-marginal likelihood can still be computed with a complexity independent of n from eq. (13). With a dense \mathbf{W} , the kernel can be written as

$$\tilde{k}(\mathbf{x}, \mathbf{z}) = \sum_{i=1}^p \sum_{j=1}^p w_{ij} \phi_i(\mathbf{x}) \phi_j(\mathbf{z}). \quad (14)$$

We may also consider specifying a non-zero prior mean for the basis function weights; $\alpha \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{W})$, where $\boldsymbol{\mu} \in \mathbb{R}^p$. This is equivalent to specifying a GP with the kernel in eq. (14) and the prior mean function $\sum_{i=1}^p \mu_i \phi_i(\mathbf{x})$. In this case, it can be observed that computation of the log-marginal likelihood can still be performed with a complexity independent of n through the relations in eq. (13) where \mathbf{r} is replaced by $\mathbf{r} - \mathbf{A}\boldsymbol{\mu}$, and $\mathbf{y}^T \mathbf{y}$ is replaced by $\mathbf{y}^T \mathbf{y} - 2\mathbf{r}^T \boldsymbol{\mu} + \boldsymbol{\mu}^T \mathbf{A} \boldsymbol{\mu}$. Similar to \mathbf{W} , we may also specify a hyper-prior on the elements of $\boldsymbol{\mu}$ and perform MCMC with these variables as well. Both of these extensions increase the flexibility of the Bayesian model while ensuring that the computational complexity remains independent of the training dataset size. This enables type-I inference to be performed on massive datasets.

C. Source Code

Source code that implements the methods discussed in the paper along with several tutorials can be found at https://github.com/treforevans/gp_grief. The code is implemented in Python and depends upon the `py-mcmc` package (Bilionis, 2014) for an implementation

of Metropolis adjusted Langevin dynamics MCMC. The code also depends on `GPY` (GPY, since 2012) for its broad library of kernels.