
Preconditioning Kernel Matrices

Kurt Cutajar

EURECOM, Department of Data Science

KURT.CUTAJAR@EURECOM.FR

Michael A. Osborne

University of Oxford, Department of Engineering Science

MOSB@ROBOTS.OX.AC.UK

John P. Cunningham

Columbia University, Department of Statistics

JPC2181@COLUMBIA.EDU

Maurizio Filippone

EURECOM, Department of Data Science

MAURIZIO.FILIPPONE@EURECOM.FR

Abstract

The computational and storage complexity of kernel machines presents the primary barrier to their scaling to large, modern, datasets. A common way to tackle the scalability issue is to use the conjugate gradient algorithm, which relieves the constraints on both storage (the kernel matrix need not be stored) and computation (both stochastic gradients and parallelization can be used). Even so, conjugate gradient is not without its own issues: the conditioning of kernel matrices is often such that conjugate gradients will have poor convergence in practice. Preconditioning is a common approach to alleviating this issue. Here we propose preconditioned conjugate gradients for kernel machines, and develop a broad range of preconditioners particularly useful for kernel matrices. We describe a scalable approach to both solving kernel machines and learning their hyperparameters. We show this approach is exact in the limit of iterations and outperforms state-of-the-art approximations for a given computational budget.

1. Introduction

Kernel machines, in enabling flexible feature space representations of data, comprise a broad and important class of tools throughout machine learning and statistics; prominent examples include support vector machines (Schölkopf

& Smola, 2001) and Gaussian processes (GPs) (Rasmussen & Williams, 2006). At the core of most kernel machines is the need to solve linear systems involving the Gram matrix $K = \{k(\mathbf{x}_i, \mathbf{x}_j | \boldsymbol{\theta})\}_{i,j=1,\dots,n}$, where the kernel function k , parameterized by $\boldsymbol{\theta}$, implicitly specifies the feature space representation of data points \mathbf{x}_i . Because K grows with the number of data points n , a fundamental computational bottleneck exists: storing K is $\mathcal{O}(n^2)$, and solving a linear system with K is $\mathcal{O}(n^3)$. As the need for large-scale kernel machines grows, much work has been directed towards this scaling issue.

Standard approaches to kernel machines involve a factorization (typically Cholesky) of K , which is efficient and exact but maintains the quadratic storage and cubic runtime costs. This cost is particularly acute when adapting (or learning) hyperparameters $\boldsymbol{\theta}$ of the kernel function, as K must then be factorized afresh for each $\boldsymbol{\theta}$. To alleviate this burden, numerous works have turned to approximate methods (Candela & Rasmussen, 2005; Snelson & Ghahramani, 2007; Rahimi & Recht, 2008) or methods that exploit structure in the kernel (Gilboa et al., 2015). Approximate methods can achieve attractive scaling, often through the use of low-rank approximations to K , but they can incur a potentially severe loss of accuracy. An alternative to factorization is found in the conjugate gradient method (CG), which is used to directly solve linear systems via a sequence of matrix-vector products. Any kernel structure can then be exploited to enable fast multiplications, driving similarly attractive runtime improvements, and eliminating the storage burden (neither K nor its factors need be represented in memory). Unfortunately, in the absence of special structure that accelerates multiplications, CG performs no better than $\mathcal{O}(n^3)$ in the worst case, and in practice finite numerical precision often results in a degradation of run-

time performance compared to factorization approaches.

Throughout optimization, the typical approach to the slow convergence of CG is to apply preconditioners to improve the geometry of the linear system being solved (Golub & Van Loan, 1996). While preconditioning has been explored in domains such as spatial statistics (Chen, 2005; Stein et al., 2012; Ashby & Falgout, 1996), the application of preconditioning to kernel matrices in machine learning has received little attention. Here we design and study preconditioned conjugate gradient methods (PCG) for use in kernel machines, and provide a full exploration of the use of approximations of K as preconditioners.

Our contributions are as follows. (i) Extending the work in (Davies, 2014), we apply a broad range of kernel matrix approximations as preconditioners. Interestingly, this step allows us to exploit the important developments of *approximate* kernel machines to accelerate the *exact* computation that PCG offers. (ii) As a motivating example used throughout the paper, we analyze and provide a general framework to both learn kernel parameters and make predictions in GPs. (iii) We extend stochastic gradient learning for GPs (Filippone & Engler, 2015; Anitescu et al., 2012) to allow any likelihood that factorizes over the data points by developing an unbiased estimate of the gradient of the approximate log-marginal likelihood. We demonstrate this contribution in making the first use of PCG for GP classification. (iv) We evaluate datasets over a range of problem size and dimensionality. Because PCG is exact in the limit of iterations (unlike approximate techniques), we demonstrate a tradeoff between accuracy and computational effort that improves beyond state-of-the-art approximation and factorization approaches.

In all, we show that PCG, with a thoughtful choice of preconditioner, is a competitive strategy which is possibly even superior than existing approximation and CG-based techniques for solving general kernel machines¹.

2. Motivating example – Gaussian Processes

Gaussian processes (GPs) are the fundamental building block of many probabilistic kernel machines that can be applied in a large variety of modeling scenarios (Rasmussen & Williams, 2006). Throughout the paper, we will denote by $X = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ a set of n input vectors and use $\mathbf{y} = (y_1, \dots, y_n)^\top$ for the corresponding labels. GPs are formally defined as collections of random variables characterized by the property that any finite number of them is jointly Gaussian distributed. The specification of a kernel function determines the covariance structure of such ran-

dom variables

$$\text{cov}(f(\mathbf{x}), f(\mathbf{x}')) = k(\mathbf{x}, \mathbf{x}' | \boldsymbol{\theta}).$$

In this work we focus in particular on the popular Radial Basis Function (RBF) kernel

$$k(\mathbf{x}_i, \mathbf{x}_j | \boldsymbol{\theta}) = \sigma^2 \exp \left[-\frac{1}{2} \sum_{r=1}^d \frac{(\mathbf{x}_i - \mathbf{x}_j)_r^2}{l_r^2} \right], \quad (1)$$

where $\boldsymbol{\theta}$ represents the collection of the kernel parameters σ^2 and l_r^2 . Defining $f_i = f(\mathbf{x}_i)$ and $\mathbf{f} = (f_1, \dots, f_n)^\top$, and assuming a zero mean GP, we have

$$\mathbf{f} \sim \mathcal{N}(\mathbf{f} | \mathbf{0}, K),$$

where K is the $n \times n$ Gram matrix with elements $K_{ij} = k(\mathbf{x}_i, \mathbf{x}_j | \boldsymbol{\theta})$. Note that the kernel above and many popular kernels in machine learning give rise to dense kernel matrices. Observations are then modeled through a transformation h of a set of GP-distributed latent variables, specifying the model

$$y_i \sim p(y_i | h(f_i)), \quad \mathbf{f} \sim \mathcal{N}(\mathbf{f} | \mathbf{0}, K).$$

2.1. The need for preconditioning

The success of nonparametric models based on kernels hinges on the adaptation of kernel parameters $\boldsymbol{\theta}$. The motivation for preconditioning begins with an inspection of the log-marginal likelihood of GP models with prior $\mathcal{N}(\mathbf{f} | \mathbf{0}, K)$. In Gaussian processes with a Gaussian likelihood $y_i \sim \mathcal{N}(y_i | f_i, \lambda)$, we have analytic forms for

$$\log[p(\mathbf{y} | \boldsymbol{\theta}, X)] = -\frac{1}{2} \log(|K_{\mathbf{y}}|) - \frac{1}{2} \mathbf{y}^\top K_{\mathbf{y}}^{-1} \mathbf{y} + \text{const},$$

and its derivatives with respect to kernel parameters θ_i ,

$$g_i = -\frac{1}{2} \text{Tr} \left(K_{\mathbf{y}}^{-1} \frac{\partial K_{\mathbf{y}}}{\partial \theta_i} \right) + \frac{1}{2} \mathbf{y}^\top K_{\mathbf{y}}^{-1} \frac{\partial K_{\mathbf{y}}}{\partial \theta_i} K_{\mathbf{y}}^{-1} \mathbf{y}. \quad (2)$$

where $K_{\mathbf{y}} = K + \lambda I$. The traditional approach involves factorizing the kernel matrix $K_{\mathbf{y}} = LL^\top$ using the Cholesky algorithm (Golub & Van Loan, 1996) which costs $\mathcal{O}(n^3)$ operations. After that, all other operations cost $\mathcal{O}(n^2)$ except for the trace term in the calculation of g_i which once again requires $\mathcal{O}(n^3)$ operations. Similar computations are required for computing mean and variance predictions for test data (Rasmussen & Williams, 2006). Note that the solution of a linear system is required for computing the variance at every test point.

This approach is not viable for large n and, consequently, many approaches have been proposed to approximate these computations, thus leading to approximate optimal values for $\boldsymbol{\theta}$ and approximate predictions. Here we investigate the

¹Code to replicate all results in this paper is available at http://github.com/mauriziofilippone/preconditioned_GPs

possibility of avoiding approximations altogether, by arguing that for parameter optimization it is sufficient to obtain an unbiased estimate of the gradient g_i . In particular, when such an estimate is available, it is possible to employ stochastic gradient optimization that has strong theoretical guarantees (Robbins & Monro, 1951). In the case of GPs, the problematic terms in eq. 2 are the solution of the linear system $K_{\mathbf{y}}^{-1}\mathbf{y}$ and the trace term. In this work we make use of a stochastic linear algebra result that allows for an approximation of the trace term,

$$\text{Tr} \left(K_{\mathbf{y}}^{-1} \frac{\partial K_{\mathbf{y}}}{\partial \theta_i} \right) \approx \frac{1}{N_{\mathbf{r}}} \sum_{i=1}^{N_{\mathbf{r}}} \mathbf{r}^{(i)\top} K_{\mathbf{y}}^{-1} \frac{\partial K_{\mathbf{y}}}{\partial \theta_i} \mathbf{r}^{(i)},$$

where the $N_{\mathbf{r}}$ vectors $\mathbf{r}^{(i)}$ have components drawn from $\{-1, 1\}$ with probability 1/2. Verifying that this is an unbiased estimate of the trace term is straightforward considering that $\text{E} \left(\mathbf{r}^{(i)} \mathbf{r}^{(i)\top} \right) = I$ (Gibbs, 1997).

This result shows that all it takes to calculate stochastic gradients is the ability to efficiently solve linear systems. Linear systems can be iteratively solved using *conjugate gradient* (CG) (Golub & Van Loan, 1996). The advantage of this formulation is that we can attempt to optimize kernel parameters using stochastic gradient optimization without having to store $K_{\mathbf{y}}$ and, given that the most expensive operation is now multiplying the kernel matrix by vectors, only $\mathcal{O}(n^2)$ computations are required. However, it is well known that the convergence of the CG algorithm depends on the condition number $\kappa(K_{\mathbf{y}})$ (ratio of largest to smallest eigenvalues), so the suitability of this approach may also be curtailed if $K_{\mathbf{y}}$ is badly conditioned. To this end, a well-known approach for improving the conditioning of a matrix, which in turn accelerates convergence, is *preconditioning*. This necessitates the introduction of a preconditioning matrix, P , which should be chosen in such a way that $P^{-1}K_{\mathbf{y}}$ approximates the identity matrix, I . Intuitively, this can be obtained by setting $P = K_{\mathbf{y}}$; however, given that in Preconditioned CG (PCG) we are required to solve linear systems involving P , this choice would be no easier than solving the original system. Thus we must choose P which approximates $K_{\mathbf{y}}$ as closely as possible, but which can also be easily inverted. The PCG algorithm is shown in Algorithm 1.

2.2. Non-Gaussian Likelihoods

When the likelihood $p(y_i | f_i)$ is not Gaussian, it is no longer possible to analytically integrate out latent variables. Instead, techniques such as Gaussian approximations (see, e.g., (Kuss & Rasmussen, 2005; Nickisch & Rasmussen, 2008)) and methods attempting to characterize the full posterior $p(\mathbf{f}, \boldsymbol{\theta} | \mathbf{y})$ (Murray et al., 2010; Filippone et al., 2013) may be required. Among the various schemes to recover tractability in the case of models with

Algorithm 1 The Preconditioned CG Algorithm, adapted from (Golub & Van Loan, 1996)

Require: data X , vector \mathbf{v} , convergence threshold ϵ , initial vector \mathbf{x}_0 , maximum no. of iterations T
 $\mathbf{r}_0 = \mathbf{v} - K_{\mathbf{y}}\mathbf{x}_0$; $\mathbf{z}_0 = P^{-1}\mathbf{r}_0$; $\mathbf{p}_0 = \mathbf{z}_0$
for $i = 0 : T$ **do**
 $\alpha_i = \frac{\mathbf{r}_i^\top \mathbf{z}_i}{\mathbf{r}_i^\top K_{\mathbf{y}} \mathbf{z}_i}$
 $\mathbf{x}_{i+1} = \mathbf{x}_i + \alpha_i \mathbf{p}_i$
 $\mathbf{r}_{i+1} = \mathbf{r}_i + \alpha_i K_{\mathbf{y}} \mathbf{p}_i$
if $\|\mathbf{r}_{i+1}\| < \epsilon$ **then**
 return $\mathbf{x} = \mathbf{x}_{i+1}$
end if
 $\mathbf{z}_{i+1} = P^{-1}\mathbf{r}_{i+1}$
 $\beta_i = \frac{\mathbf{r}_{i+1}^\top \mathbf{r}_{i+1}}{\mathbf{r}_i^\top \mathbf{r}_i}$
 $\mathbf{p}_{i+1} = \mathbf{p}_{i+1} + \beta_i \mathbf{p}_i$
end for

a non-Gaussian likelihood, we choose the Laplace approximation, as we can formulate it in a way that only requires the solution of linear systems. The GP models we consider assume that the likelihood factorizes across all data points $p(\mathbf{y} | \mathbf{f}) = \prod_{i=1}^n p(y_i | f_i)$. The use of CG for computing the Laplace approximation has been proposed elsewhere (Flaxman et al., 2015), but we make the first use of preconditioning and stochastic gradient estimation within the Laplace approximation to compute stochastic gradients for non-conjugate models.

Defining $W = -\nabla_{\mathbf{f}} \nabla_{\mathbf{f}} \log[p(\mathbf{y} | \mathbf{f})]$ (a diagonal matrix), carrying out the Laplace approximation algorithm, computing its derivatives wrt $\boldsymbol{\theta}$, and making predictions, all possess the same computational bottleneck: the solution of linear systems involving the matrix $B = I + W^{\frac{1}{2}} K W^{\frac{1}{2}}$ (Rasmussen & Williams, 2006). For a given $\boldsymbol{\theta}$, each iteration of the Laplace approximation algorithm requires solving one linear system involving B and two matrix-vector multiplications involving K ; the linear system involving B can be solved using CG or PCG. The Laplace approximation yields the mode $\hat{\mathbf{f}}$ of the posterior over latent variables and offers an approximate log-marginal likelihood in the form:

$$\log[\hat{p}(\mathbf{y} | \boldsymbol{\theta}, X)] = -\frac{1}{2} \log |B| - \frac{1}{2} \hat{\mathbf{f}}^\top K^{-1} \hat{\mathbf{f}} + \log[p(\mathbf{y} | \hat{\mathbf{f}})]$$

which poses the same computational challenges as the regression case. Once again, we therefore seek an alternative way to learn kernel parameters by stochastic gradient optimization based on computing unbiased estimates of the gradient of the approximate log-marginal likelihood. This is complicated further by the inclusion of an additional ‘‘implicit’’ term accounting for the change in the solution given by the Laplace approximation for a change in $\boldsymbol{\theta}$. The full derivation of the gradient is rather lengthy and is deferred to the supplementary material. Nonetheless, it is worth not-

ing that the calculation of the exact gradient involves trace terms similar to the regression case that cannot be computed for large n , and we unbiasedly estimate these using the stochastic approximation of the trace.

3. Preconditioning Kernel Matrices

Here we consider choices for kernel preconditioners, and for the sake of clarity we focus on preconditioners for $K_{\mathbf{y}}$. Unless stated otherwise, we shall consider standard *left* preconditioning, whereby the original problem of solving $K_{\mathbf{y}}\mathbf{z} = \mathbf{v}$ is transformed by applying a preconditioner, P , to both sides of this equation. This formulation may thus be expressed as $P^{-1}K_{\mathbf{y}}\mathbf{z} = P^{-1}\mathbf{v}$.

3.1. Nyström type approximations

The Nyström method was originally proposed to approximate the eigendecomposition of kernel matrices (Williams & Seeger, 2000); as a result, it offers a way to obtain a low rank approximation of K . This method selects a subset of $m \ll n$ data (inducing points) collected in the set U which are intended for approximating the spectrum of K . The resulting approximation is $\hat{K} = K_{XU}K_{UU}^{-1}K_{UX}$ where K_{UU} denotes the evaluation of the kernel function over the inducing points, and K_{XU} denotes the evaluation of the kernel function between the input points and the inducing points. The resulting preconditioner $P = K_{XU}K_{UU}^{-1}K_{UX} + \lambda I$ can be inverted using the matrix inversion lemma

$$P^{-1}\mathbf{v} = \lambda^{-1} \left[I - K_{XU} (K_{UU} + K_{UX}K_{XU})^{-1} K_{UX} \right] \mathbf{v},$$

which has $\mathcal{O}(m^3)$ complexity.

3.1.1. FULLY AND PARTIALLY INDEPENDENT TRAINING CONDITIONAL

The use of a subset of data for approximating a GP kernel has also been utilized in the fully and partially independent training conditional approaches (FITC and PITC, respectively) for approximating GP regression (Candela & Rasmussen, 2005). In the former case, the prior covariance of the approximation can be written as follows:

$$P = K_{XU}K_{UU}^{-1}K_{UX} + \mathbf{diag} (K - K_{XU}K_{UU}^{-1}K_{UX}) + \lambda I.$$

As the name implies, this formulation enforces that the latent variables associated with U are taken to be completely conditionally independent. On the other hand, the PITC method extends on this approach by enforcing that although inducing points assigned to a designated block are conditionally dependent on each other, there is no dependence between points placed in different blocks:

$$P = K_{XU}K_{UU}^{-1}K_{UX} + \mathbf{bdiag} (K - K_{XU}K_{UU}^{-1}K_{UX}) + \lambda I.$$

For the FITC preconditioner, the diagonal resulting from the training conditional can be added to the diagonal noise matrix, and the inversion lemma can be invoked as for the Nyström case. Meanwhile, for the PITC preconditioner, the noise diagonal can be added to the block diagonal matrix, which can then be inverted block-by-block. Once again, matrix inversion can then be carried out as before, where the inverted block diagonal matrix takes the place of λI in the original formulation.

3.2. Approximate factorization of kernel matrices

This group of preconditioners relies on approximations to K that factorize as $\hat{K} = \Phi\Phi^\top$. We shall consider different ways of determining Φ such that P can be inverted at a lower cost than the original kernel matrix K . Once again, this enables us to employ the matrix inversion lemma, and express the linear system:

$$P^{-1}\mathbf{v} = (\Phi\Phi^\top + \lambda I)^{-1}\mathbf{v} = \lambda^{-1} [I - \Phi(I + \Phi^\top\Phi)^{-1}\Phi^\top] \mathbf{v}.$$

We now review a few methods to approximate the kernel matrix K in the form $\Phi\Phi^\top$.

3.2.1. SPECTRAL APPROXIMATION

The spectral approach uses random Fourier features for deriving a sparse approximation of a GP (Rahimi & Recht, 2008). This approach for GPs was introduced in (Lázaro-Gredilla et al., 2010), and relies on the assumption that stationary kernel functions can be represented as the Fourier transform of non-negative measures. As such, the elements of K can be approximated as follows:

$$\hat{K}_{ij} = \frac{\sigma_0^2}{m} \phi(\mathbf{x}_i)^\top \phi(\mathbf{x}_j) = \frac{\sigma_0^2}{m} \sum_{r=1}^m \cos [2\pi \mathbf{s}_r^\top (\mathbf{x}_i - \mathbf{x}_j)].$$

In the equation above, the vectors \mathbf{s}_r denote the *spectral points* (or *frequencies*) which in the case of the RBF kernel can be sampled from $\mathcal{N}(\mathbf{0}, \frac{1}{4\pi^2}\Lambda)$, where $\Lambda = [1/l_1^2, \dots, 1/l_n^2]$. To the best of our knowledge, this is the first time such an approximation has been considered for the purpose of preconditioning kernel matrices.

3.2.2. PARTIAL SVD

Another factorization approach that we consider in this work is the partial singular value decomposition (SVD) method (Golub & Van Loan, 1996). The SVD method factorizes the original kernel matrix K into $A\Lambda A^\top$, where A is a unitary matrix and Λ is a diagonal matrix of singular values. Here, we shall consider a variation of this technique called *randomized truncated SVD* (Halko et al., 2011), which constructs an approximate low rank SVD factorization of K using random sampling to accelerate computations.

3.2.3. STRUCTURED KERNEL INTERPOLATION (SKI)

Some recent work on approximating GPs has exploited the fast computation of Kronecker matrix-vector multiplications when inputs are located on a Cartesian grid (Gilboa et al., 2015). Unfortunately, not all datasets meet this requirement, thus limiting the widespread application of Kronecker inference. To this end, SKI (Wilson & Nickisch, 2015) is an approximation technique which exploits the benefits of the Kronecker product without imposing any requirements on the structure of the training data. In particular, a grid of inducing points, U , is constructed, and the covariance between the training data and U is then represented as $K_{XU} = WK_{UU}$. In this formulation, W denotes a sparse interpolation matrix for assigning weights to the elements of K_{UU} . In this manner, a preconditioner exploiting Kronecker structure can be constructed as $P = WK_{UU}W^T + \lambda I$. If we consider $V = W/\sqrt{\lambda}$, we can rewrite the (inverse) preconditioner as $P^{-1} = \lambda^{-1}(VK_{UU}V^T + I)^{-1}$. Since this can no longer be solved directly, we solve this (inner-loop) linear system using the CG algorithm (all within one iteration of the outer-loop PCG). For badly conditioned systems, although the complexity of the required matrix-vector multiplications is now much less than $\mathcal{O}(n^2)$, the number of iterations to solve linear systems involving the preconditioner is potentially very large, and could diminish the benefits of preconditioning.

3.3. Other approaches

3.3.1. BLOCK JACOBI

An alternative to using a single subset of data involves constructing local GPs over segments of the original data (Snelson & Ghahramani, 2007). An example of such an approach is the *Block Jacobi* approximation, whereby the preconditioner is constructed by taking a block diagonal of K and discarding all other elements in the kernel matrix. In this manner, covariance is only expressed for points within the same block, as $P = \mathbf{bldiag}(K_{\mathbf{y}} + \lambda I)$. The inverse of this block diagonal matrix is computationally cheap (also block diagonal). However, given that a substantial amount of information contained in the original covariance matrix is ignored, this choice is intrinsically a rather crude approach.

3.3.2. REGULARIZATION

An appealing feature shared by the aforementioned preconditioners (aside from SKI) is that their structure enables us to directly solve $P^{-1}\mathbf{v}$. An alternative technique for constructing a preconditioner involves adding a positive regularization parameter, δI , to the original kernel matrix, such that $P = K_{\mathbf{y}} + \delta I$ (Srinivasan et al., 2014). This follows from the fact that adding noise to the diagonal of $K_{\mathbf{y}}$ makes it better-conditioned, and the condition number is expected

to decrease further as δ increases. Nonetheless, for the purpose of preconditioning, this parameter should be tuned in such a way that P remains a sensible approximation of $K_{\mathbf{y}}$. As opposed to the previous preconditioners, this is an instance of *right* preconditioning, which has the following general form $K_{\mathbf{y}}P^{-1}(P\mathbf{x}) = \mathbf{v}$.

Given that it is no longer possible to evaluate $P^{-1}\mathbf{v}$ analytically, this linear system is solved yet again using CG, such that a linear system of equations is solved at every outer iteration of the PCG algorithm. Due to the potential loss of accuracy incurred while solving the inner linear systems, a variation of the standard PCG algorithm, referred to as *flexible* PCG (Notay, 2000), is used instead. Using this approach, a re-orthogonalization step is introduced such that the search directions remain orthogonal even when the inner system is not solved to high precision.

4. Comparison of Preconditioners

In this section, we provide an empirical exploration of these preconditioners in a practical setting. We begin by considering three datasets for regression from the UCI repository (Asuncion & Newman, 2007), namely the Concrete dataset ($n = 1030, d = 8$), the Power Plant dataset ($n = 9568, d = 4$), and the Protein dataset ($n = 45730, d = 9$). In particular, we evaluate the convergence in solving $K_{\mathbf{y}}\mathbf{z} = \mathbf{y}$ using iterative methods, where \mathbf{y} denotes the labels of the designated dataset, and $K_{\mathbf{y}}$ is constructed using different configurations of kernel parameters.

With this experiment, we aim to assess the quality of different preconditioners based on how many matrix-vector products they require, which, for most approaches, corresponds to the number of iterations taken by PCG to converge. The convergence threshold is set to $\epsilon^2 = n \cdot 10^{-10}$ so as to roughly accept an average error of 10^{-5} on each element of the solution.

For every variation, we set the parameters of the preconditioners so as to have a complexity lower than the $\mathcal{O}(n^2)$ cost associated with matrix-vector products; by doing so, we can assume that the latter computations are the dominant cost for large n . In particular, for Nyström-type methods, we set $m = \sqrt{n}$ inducing points, so that when we invert the preconditioner using the matrix inversion lemma, the cost is in $\mathcal{O}(m^3) = \mathcal{O}(n^{3/2})$. Similarly, for the Spectral preconditioner, we set $m = \sqrt{n}$ random features. For the SKI preconditioner, we take an equal number of elements on the grid for each dimension; under this assumption, Kronecker products have $\mathcal{O}(dn^{\frac{d+1}{d}})$ cost (Gilboa et al., 2015), and we set the size of the grid so that the complexity of applying the preconditioner matches $\mathcal{O}(n^{3/2})$, so as to be consistent with the other preconditioners. For the Regularized approach, each iteration needed to apply the preconditioner




Figure 1. Comparison of preconditioners for different settings of kernel parameters. The lengthscale l and the noise variance λ are shown on the x and y axes respectively. The top figure indicates the number of iterations required to solve the corresponding linear system using CG, whilst the bottom part of the figure shows the rate of improvement (negative - blue) or degradation (positive - red) achieved by using PCG to solve the same linear system. Parameters and results are reported in \log_{10} . Symbols added to facilitate reading in B/W print.

tioner requires one matrix-vector product, and we add this to the overall count of such computations. For this preconditioner, we add a diagonal offset δ to the original matrix, equivalent to two orders of magnitude greater than the noise of the process. In general, although the complexity of PCG is indeed no different from that of CG, we emphasize that experiencing a 2-fold or 5-fold (in some cases even an order of magnitude) improvement can be very substantial when plain CG takes very long to converge or when the dataset is large.

We focus on an isotropic RBF variant of the kernel in eq. 1, fixing the marginal variance σ^2 to one. We vary the length-scale parameter l and the noise variance λ in \log_{10} scale. The top part of fig. 1 shows the number of iterations that the standard CG algorithm takes, where we have capped the number of iterations to 100,000.

The bottom part of the figure reports the improvement offered by various preconditioners measured as

$$\log_{10} \left(\frac{\# \text{ PCG iterations}}{\# \text{ CG iterations}} \right).$$

It is worth noting that when both CG and PCG fail to converge within the upper bound, the improvement will be marked as 0, i.e. neither a gain or a loss within the given bound. The results plotted in fig. 1 indicate that the low-

rank preconditioners (PITC, FITC and Nyström) achieve significant reductions in the number of iterations for each dataset, and all approaches work best when the lengthscale is longer, characterising smoother processes. In contrast, preconditioning seems to be less effective when the lengthscale is shorter, corresponding to a kernel matrix that is more sparse. However, for cases yielding positive results, the improvement is often in the range of an order of magnitude, which can be substantial when a large number of iterations is required by the CG algorithm.

The results also confirm that, as alluded to in the previous section, Block Jacobi preconditioning is generally a poor preconditioner, particularly when the corresponding kernel matrix is dense. The only minor improvements were observed when CG itself converges quickly, in which case preconditioning serves very little purpose either way.

The regularization approach with flexible conjugate gradient does not appear to be effective in any case either, particularly due to the substantial amount of iterations required for solving an inner system at every iteration of the PCG algorithm. This implies that introducing additional small jitter to the diagonal does not necessarily make the system much easier to solve, whilst adding an overly large offset would negatively impact convergence of the outer algorithm. One could assume that tuning the value of this

parameter could result in slightly better results; however, preliminary experiments in this regard yielded only minor improvements.

The results for SKI preconditioning are similarly discouraging at face value. When the matrix K_y is very badly conditioned, an excessive number of inner iterations are required for every iteration of outer PCG. This greatly increases the duration of solving such systems, and as a result, this method was not included in the comparison for the Protein dataset, where it was evident that preconditioning the matrix in this manner would not yield satisfactory improvements. Notwithstanding that these experiments depict a negative view of SKI preconditioning, it must be said that we assumed a fairly simplistic interpolation procedure in our experiments, where each data point was mapped to nearest grid location. The size of the constructed grid is also hindered considerably by the constraint imposed by our upper bound on complexity. Conversely, more sophisticated interpolation strategies or even grid formulation procedures could possibly speed up the convergence of CG for the inner systems. In line with this thought, however, one could argue that the preconditioner would no longer be straightforward to construct, which goes against our innate preference towards easily derived preconditioners.

5. Impact of preconditioning on GP learning

One of the primary objectives of this work is to reformulate GP regression and classification in such a way that preconditioning can be effectively exploited. In section 2, we demonstrated how preconditioning can indeed be applied to GP regression problems, and also proposed a novel way of rewriting GP classification in terms of solving linear systems (where preconditioning can thus be employed). We can now evaluate how the proposed preconditioned GP techniques compare to other state of the art methods.

To this end, in this section, we empirically report on the generalization ability of GPs as a function of the time taken to optimize parameters θ and compute predictions. In particular, for each of the methods featured in our comparison, we iteratively run the optimization of kernel parameters for a few iterations and predict on unseen data, and assess how prediction accuracy varies over time for different methods.

The analysis provided in this report is inspired by (Chalupka et al., 2013), although we do not propose an *approximate* method to learn GP kernel parameters. Instead, we put forward a means of accelerating the optimization of kernel parameters *without any approximation*². Given the predictive mean and variance for the

²The one proviso to this statement is that, for non-Gaussian likelihood, stochastic gradients target the approximate log-marginal likelihood obtained by the Laplace approximation.

N_{test} test points, say m_{*i} and s_{*i}^2 , we report two error measures, namely the Root Mean Square Error, $\text{RMSE} = \sqrt{\frac{1}{N_{\text{test}}} \sum_{i=1}^{N_{\text{test}}} (m_{*i} - y_{*i})^2}$, along with the negative log-likelihood on the test data, $-\sum_{i=1}^{N_{\text{test}}} \log[p(y_{*i} | m_{*i}, s_{*i}^2)]$, where y_{*i} denotes the label of the i th of N_{test} data points. For classification, instead of the RMSE we report the error rate of the classifier.

We can make use of stochastic gradients for GP models to optimize kernel parameters using off-the-shelf stochastic gradient optimization algorithms. In order to reduce the number of parameters to tune, we employ ADAGRAD (Duchi et al., 2011) – an optimization algorithm having a single step-size parameter. For the purpose of this experiment, we do not attempt to optimize this parameter, since this would require additional computations. Nonetheless, our experience with training GP models indicates that the choice of this parameter is not critical: we set the step-size to one.

Fig. 2 shows the two error measures over time for a selection of approaches. In the figure, PCG and CG refer to stochastic gradient optimization of kernel parameters using ADAGRAD, where linear systems are solved with PCG and CG, respectively. In view of the results obtained in our comparison of preconditioners, we decide to proceed with the Nyström preconditioning method. Furthermore, we construct the preconditioner with $m = 4\sqrt{n}$ points randomly selected from the input data at each iteration, such that the overall complexity of the PCG method matches plain CG. For these methods, stochastic estimates of trace terms are carried out using $N_r = 4$ random vectors. The baseline CHOL method refers to the optimization of kernel parameters using the L-BFGS algorithm, where the exact log-marginal likelihood and its gradient are calculated using the full Cholesky decomposition of K_y or B .

Alongside these approaches for optimizing kernel parameters without approximation, we also evaluate the performance of approximate GP methods. For this experiment, we chose to compare against approximations found in the software package GPstuff (Vanhatalo et al., 2013), namely the fully and partial independent training conditional approaches (FITC, PITC), and the sparse variational GP (VAR) (Titsias, 2009). In order to match the computational cost of CG/PCG, which is in $\mathcal{O}(n^2)$, we set the number of inducing points for the approximate methods to be $n^{2/3}$.

All methods are initialized from the same set of kernel parameters, and the curves are averaged over 5 folds (3 for the Protein and EEG datasets). For the sake of integrity, we ran each method in the comparison individually on a workstation with Intel Xeon E5-2630 CPU having 16 cores and 128GB RAM. We also ensured that all methods reported

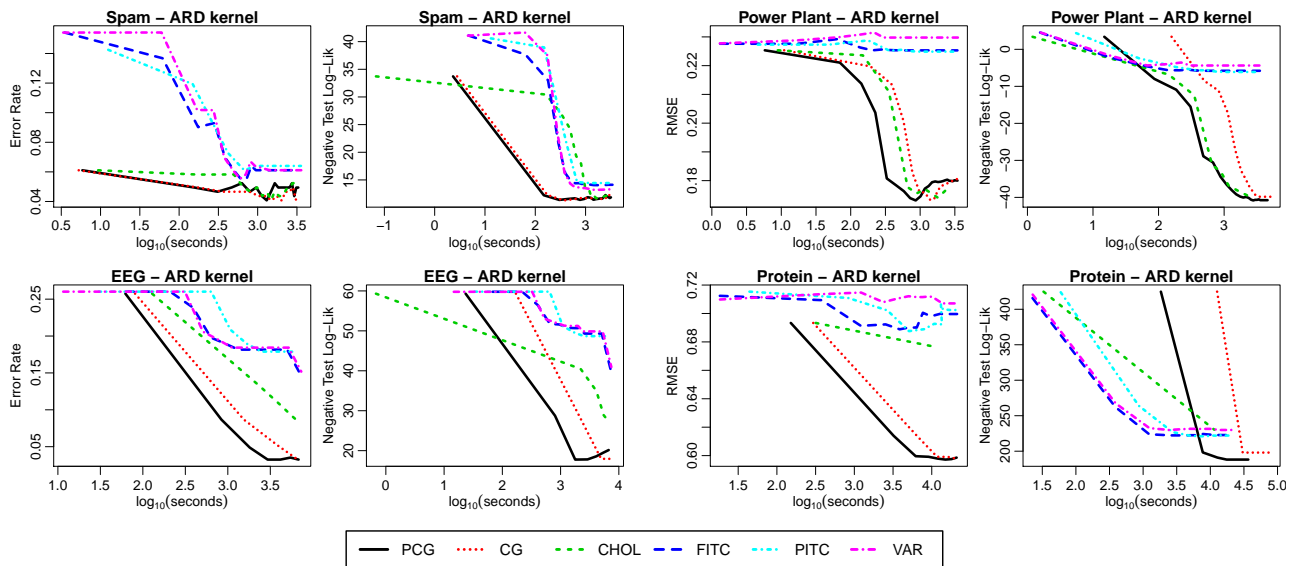


Figure 2. RMSE and negative log of the likelihood on \sqrt{n} held out test data over time. GP models employ the ARD kernel in eq. 1. GP classification: Spam dataset ($n = 4601, d = 57$) and EEG dataset ($n = 14979, d = 14$). GP regression: Power Plant dataset ($n = 9568, d = 4$) and Protein dataset ($n = 45730, d = 9$). Curves are averaged over multiple repetitions.

in the comparison used optimized linear algebra routines exploiting the multi-core architecture. This diligence for ensuring fairness gives credence to our assumption that the timings are not affected by external factors other than the actual implementation of the algorithms. The CG, PCG and CHOL approaches have been implemented in R; the fact that the approximate methods were implemented in a different environment (GPstuff is written in Matlab/Octave) and by a different developer may cast some doubt on the correctness of directly comparing results. However, we believe that the key point emerging from this comparison is that preconditioning feasibly enables the use of iterative approaches for optimization of kernel parameters in GPs, and the results are competitive with those achieved using popular GP software packages.

For the reported experiments, it was possible to store the kernel matrix K for all datasets, making it possible to compare methods against the baseline GP where computations use Cholesky decompositions. We stress, however, that iterative approaches based on CG/PCG can be implemented without the need to store K , whereas this is not possible for approaches that attempt to factorize K exactly. It is also worth noting that for the CG/PCG approach, calculating the log-likelihood on test data requires solving one linear system for each test point; this clearly penalizes the speed of these methods given the set-up of the experiment, where predictions are carried out every fixed number of iterations.

6. Discussion and Conclusions

Careful attention to numerical properties is essential in scaling machine learning to large and realistic datasets. Here we have introduced the use of preconditioning to the implementation of kernel machines, specifically, prediction and learning of kernel parameters for GPs. Our novel scheme permits the use of any likelihood that factorizes over the data points, allowing us to tackle both regression and classification. We have shown robust performance improvements, in both accuracy and computational cost, over a host of state-of-the-art approximation methods for kernel machines. Notably, our method is exact in the limit of iterations, unlike approximate alternatives. We have also shown that the use of PCG is competitive with exact Cholesky decomposition in modestly sized datasets, when the Cholesky factors can be feasibly computed. When data and thus the kernel matrix grow large enough, Cholesky factorization becomes unfeasible, leaving PCG as the optimal choice.

One of the key features of a PCG implementation is that it does not require storage of any $\mathcal{O}(n^2)$ objects. We plan to extend our implementation to compute the elements of K on the fly in one case, and in another case store K in a distributed fashion (e.g. in TensorFlow/Spark). Furthermore, while we have focused on solving linear systems, we can also use preconditioning for other iterative algorithms involving the K matrix, e.g., those to solve $\log(K)\mathbf{v}$ and $K^{1/2}\mathbf{v}$ (Chen et al., 2011), as is often useful in estimating marginal likelihoods for probabilistic kernel models like GPs.

Acknowledgements

KC and MF are grateful to Pietro Michiardi and Daniele Venzano for assisting the completion of this work by providing additional computational resources for running the experiments. JPC acknowledges support from the Sloan Foundation, The Simons Foundation (SCGB#325171 and SCGB#325233), and The Grossman Center at Columbia University.

References

- Anitescu, M., Chen, J., and Wang, L. A Matrix-free Approach for Solving the Parametric Gaussian Process Maximum Likelihood Problem. *SIAM Journal on Scientific Computing*, 34(1):A240–A262, 2012.
- Ashby, S. F. and Falgout, R. D. A Parallel Multigrid Preconditioned Conjugate Gradient algorithm for Groundwater Flow Simulations. *Nuclear Science and Engineering*, 124(1):145–159, 1996.
- Asuncion, A. and Newman, D. J. UCI machine learning repository, 2007. URL <http://archive.ics.uci.edu/ml>.
- Candela, J. Q. and Rasmussen, C. E. A Unifying View of Sparse Approximate Gaussian Process Regression. *Journal of Machine Learning Research*, 6:1939–1959, 2005.
- Chalupka, K., Williams, C. K. I., and Murray, I. A framework for evaluating approximation methods for Gaussian process regression. *Journal of Machine Learning Research*, 14(1):333–350, 2013.
- Chen, J., Anitescu, M., and Saad, Y. Computing $f(A)b$ via Least Squares Polynomial Approximations. *SIAM Journal on Scientific Computing*, 33(1):195–222, 2011.
- Chen, K. *Matrix Preconditioning Techniques and Applications*. Cambridge Monographs on Applied and Computational Mathematics. Cambridge University Press, 2005.
- Davies, A. *Effective Implementation of Gaussian Process Regression for Machine Learning*. PhD thesis, University of Cambridge, 2014.
- Duchi, J., Hazan, E., and Singer, Y. Adaptive Subgradient Methods for Online Learning and Stochastic Optimization. *Journal of Machine Learning Research*, 12:2121–2159, 2011.
- Filippone, M. and Engler, R. Enabling scalable stochastic gradient-based inference for Gaussian processes by employing the Unbiased Linear System Solver (ULISSE). In Blei, D. and Bach, F. (eds.), *Proceedings of The 32nd International Conference on Machine Learning*, volume 37 of *JMLR Proceedings*, pp. 1015–1024, 2015.
- Filippone, M., Zhong, M., and Girolami, M. A comparative evaluation of stochastic-based inference methods for Gaussian process models. *Machine Learning*, 93(1):93–114, 2013.
- Flaxman, S., Wilson, A., Neill, D., Nickisch, H., and Smola, A. Fast Kronecker inference in Gaussian processes with non-Gaussian likelihoods. In Blei, D. and Bach, F. (eds.), *Proceedings of The 32nd International Conference on Machine Learning*, volume 37 of *JMLR Proceedings*, pp. 607–616, 2015.
- Gibbs, M. N. *Bayesian Gaussian processes for regression and classification*. PhD thesis, University of Cambridge, 1997.
- Gilboa, E., Saatchi, Y., and Cunningham, J. P. Scaling Multidimensional Inference for Structured Gaussian Processes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(2):424–436, 2015.
- Golub, G. H. and Van Loan, C. F. *Matrix computations*. The Johns Hopkins University Press, 3rd edition, 1996.
- Halko, N., Martinsson, P. G., and Tropp, J. A. Finding Structure with Randomness: Probabilistic Algorithms for Constructing Approximate Matrix Decompositions. *SIAM Review*, 53(2):217–288, 2011.
- Kuss, M. and Rasmussen, C. E. Assessing Approximate Inference for Binary Gaussian Process Classification. *Journal of Machine Learning Research*, 6:1679–1704, 2005.
- Lázaro-Gredilla, M., Quinonero-Candela, J., Rasmussen, C. E., and Figueiras-Vidal, A. R. Sparse Spectrum Gaussian Process Regression. *Journal of Machine Learning Research*, 11:1865–1881, 2010.
- Murray, I., Adams, R. P., and MacKay, D. J. C. Elliptical slice sampling. In Teh, Y. W. and Titterton, D. M. (eds.), *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9 of *JMLR Proceedings*, pp. 541–548, 2010.
- Nickisch, H. and Rasmussen, C. E. Approximations for Binary Gaussian Process Classification. *Journal of Machine Learning Research*, 9:2035–2078, 2008.
- Notay, Y. Flexible Conjugate Gradients. *SIAM Journal on Scientific Computing*, 22(4):1444–1460, 2000.
- Rahimi, A. and Recht, B. Random features for large-scale kernel machines. In Platt, J. C., Koller, D., Singer, Y., and Roweis, S. T. (eds.), *Advances in Neural Information Processing Systems 20*, pp. 1177–1184, 2008.
- Rasmussen, C. E. and Williams, C. *Gaussian Processes for Machine Learning*. MIT Press, 2006.

- Robbins, H. and Monro, S. A Stochastic Approximation Method. *The Annals of Mathematical Statistics*, 22:400–407, 1951.
- Schölkopf, B. and Smola, A. J. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, Cambridge, MA, USA, 2001.
- Snelson, E. and Ghahramani, Z. Local and global sparse Gaussian process approximations. In Meila, M. and Shen, X. (eds.), *Proceedings of the 11th International Conference on Artificial Intelligence and Statistics*, volume 2 of *JMLR Proceedings*, pp. 524–531, 2007.
- Srinivasan, B. V., Hu, Q., Gumerov, N. A., Murtugudde, R., and Duraiswami, R. Preconditioned Krylov solvers for kernel regression, 2014. arXiv:1408.1237.
- Stein, M. L., Chen, J., and Anitescu, M. Difference Filter Preconditioning for Large Covariance Matrices. *SIAM Journal on Matrix Analysis Applications*, 33(1):52–72, 2012.
- Titsias, M. K. Variational Learning of Inducing Variables in Sparse Gaussian Processes. In Dyk, D. A. and Welling, M. (eds.), *Proceedings of the 12th International Conference on Artificial Intelligence and Statistics*, volume 5 of *JMLR Proceedings*, pp. 567–574, 2009.
- Vanhatalo, J., Riihimäki, J., Hartikainen, J., Jylänki, P., Tolvanen, V., and Vehtari, A. GPstuff: Bayesian modeling with Gaussian processes. *Journal of Machine Learning Research*, 14(1):1175–1179, 2013.
- Williams, C. K. I. and Seeger, M. W. Using the Nyström method to speed up kernel machines. In Leen, T. K., Dietterich, T. G., and Tresp, V. (eds.), *Advances in Neural Information Processing Systems 13*, pp. 682–688, 2000.
- Wilson, A. and Nickisch, H. Kernel Interpolation for Scalable Structured Gaussian Processes (KISS-GP). In Blei, D. and Bach, F. (eds.), *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *JMLR Proceedings*, pp. 1775–1784, 2015.