

# Preconditioning for Scalable Gaussian Process Hyperparameter Optimization

Jonathan Wenger<sup>1,2,3</sup> Geoff Pleiss<sup>3</sup> Philipp Hennig<sup>1,2</sup> John P. Cunningham<sup>3</sup> Jacob R. Gardner<sup>4</sup>

## Abstract

Gaussian process hyperparameter optimization requires linear solves with, and log-determinants of, large kernel matrices. Iterative numerical techniques are becoming popular to scale to larger datasets, relying on the conjugate gradient method (CG) for the linear solves and stochastic trace estimation for the log-determinant. This work introduces new algorithmic and theoretical insights for preconditioning these computations. While preconditioning is well understood in the context of CG, we demonstrate that it can also accelerate convergence and reduce variance of the estimates for the log-determinant and its derivative. We prove general probabilistic error bounds for the preconditioned computation of the log-determinant, log-marginal likelihood and its derivatives. Additionally, we derive specific rates for a range of kernel-preconditioner combinations, showing that up to exponential convergence can be achieved. Our theoretical results enable provably efficient optimization of kernel hyperparameters, which we validate empirically on large-scale benchmark problems. There our approach accelerates training by up to an order of magnitude.

## 1. Introduction

Gaussian processes (GPs) are a theoretically well-founded and powerful probabilistic model (Rasmussen & Williams, 2006). However, conditioning a GP on data is often computationally prohibitive for large datasets. This problem is amplified when optimizing kernel hyperparameters. Gradient-based optimization requires repeated evaluation of the log-marginal likelihood  $\mathcal{L}$  and its derivatives. These computations both have cubic complexity in the size  $n$  of the data.

<sup>1</sup>University of Tübingen <sup>2</sup>Max Planck Institute for Intelligent Systems, Tübingen <sup>3</sup>Columbia University <sup>4</sup>University of Pennsylvania. Correspondence to: Jonathan Wenger <jonathan.wenger@uni-tuebingen.de>.

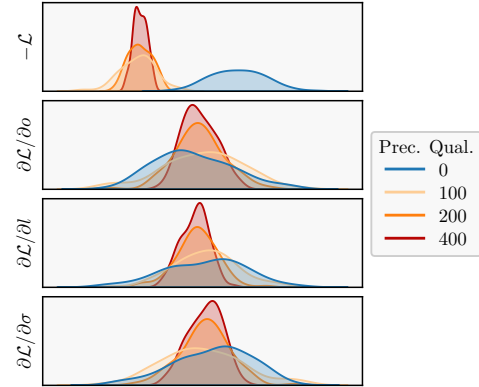


Figure 1: *Preconditioning reduces not only bias but also variance in stochastic approximations to the log-marginal likelihood  $\mathcal{L}$  and its derivatives.* GP hyperparameter optimization for large datasets requires cheap estimates of  $\mathcal{L}$  and its gradient. Preconditioning makes these more precise and less noisy as is shown here for increasing preconditioner quality on the “Elevators” dataset using a Matérn( $\frac{3}{2}$ ) kernel.

**GP Inference via Matrix-Vector Multiplication** Recently, Krylov methods (Golub & van Loan, 2013), based on iterative matrix-vector multiplication with the kernel matrix, have become popular for GP inference (Murray, 2009; Cutajar et al., 2016; Ubaru et al., 2017; Dong et al., 2017; Gardner et al., 2018; Wang et al., 2019). Methods primarily relying on matrix-vector products are advantageous. They can leverage structure in the kernel matrix (Nocedal & Wright, 2006), and importantly, they make effective use of modern hardware and parallelization (Gardner et al., 2018; Wang et al., 2019; Charlier et al., 2021). When optimizing GP hyperparameters one needs to repeatedly solve linear systems with, and compute log-determinants of, the kernel matrix. Both can be done using Krylov methods. The linear systems are solved via the conjugate gradient method (CG) (Hestenes & Stiefel, 1952), which reduces the cost of kernel matrix solves from  $\mathcal{O}(n^3)$  to  $\mathcal{O}(n^2m)$  for  $m$  iterations. The log-determinant can be approximated via stochastic trace estimation (STE) (Hutchinson, 1989) combined with another Krylov method, the Lanczos algorithm (Lanczos, 1950) (see (Ubaru et al., 2017)). Its derivative may also be estimated via STE combined with CG (Gardner et al., 2018).

**Challenges with this Approach** Despite the advantages of combining Krylov methods with stochastic trace estimation, there are considerable challenges in practice. These essentially reduce to bias and variance of the numerical approximations. First, the convergence of CG depends on the conditioning of the kernel matrix, which can grow rapidly with  $n$  (e.g. for the RBF kernel). Many iterations may be needed to achieve a desired error, and stopping the solver early can result in biased solutions (Potapczynski et al., 2021). However, if a preconditioner – i.e. an approximation of the kernel matrix – is available, convergence can be accelerated substantially (Golub & van Loan, 2013). Second, stochastic approximations of the log-determinant and its derivative introduce variance into hyperparameter optimization. While the estimates are unbiased (assuming sufficient Krylov iterations), variance can significantly slow down optimization. Reducing variance either requires further approximation at the cost of more bias (Artemev et al., 2021), or a larger number of samples  $\ell$  which only reduces error at a rate of  $\mathcal{O}(\ell^{-\frac{1}{2}})$  (Avron & Toledo, 2011). Now, while preconditioning is known to accelerate CG, it has not yet been explored for stochastic trace estimation in this context.

**Contributions** We demonstrate that, with only a small algorithmic modification, preconditioning can be exploited for highly efficient log-determinant estimation, and in turn GP hyperparameter optimization. We show that

- (a) *preconditioning reduces variance* – or equivalently accelerates convergence – of the stochastic estimate of the log-determinant and its derivative (Theorem 1).

We leverage this result, illustrated in Figure 1, to prove

- (b) *stronger theoretical guarantees* for the computation of the log-determinant (Theorems 2 and 3) and log-marginal likelihood (Theorem 4) than previously known (Ubaru et al., 2017; Gardner et al., 2018) and a *novel error bound for the derivative* (Theorem 5).

To make these general results concrete, we derive

- (c) *specific rates for important combinations of kernels and preconditioners* (Table 1), making preconditioner choice for GP inference rigorous rather than heuristic.

Finally, using our approach, we empirically observe

- (d) *up to twelvefold speedup in training* of GP regression models applied to large-scale benchmark problems with up to  $n \approx 325,000$  datapoints.

## 2. Background

We want to infer a map  $h : \mathcal{X} \rightarrow \mathcal{Y}$  from an input space  $\mathcal{X} \subset \mathbb{R}^d$  to an output space  $\mathcal{Y} \subset \mathbb{R}$ , given a dataset  $\mathbf{X} \in \mathbb{R}^{n \times d}$  of  $n$  training inputs  $\mathbf{x}_i \in \mathbb{R}^d$  and outputs  $\mathbf{y} \in \mathbb{R}^n$ .

### 2.1. Gaussian Processes

A stochastic process  $f \sim \mathcal{GP}(\mu, k)$  with mean function  $\mu$  and kernel  $k$  is called a *Gaussian process* if  $\mathbf{f} = (f(\mathbf{x}_1), \dots, f(\mathbf{x}_n))^\top \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{K})$  is jointly Gaussian with mean  $\boldsymbol{\mu}_i = \mu(\mathbf{x}_i)$  and covariance  $\mathbf{K}_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$ . Assuming  $\mathbf{y} | \mathbf{f} \sim \mathcal{N}(\mathbf{f}, \sigma^2 \mathbf{I})$ , the posterior distribution for test inputs  $\mathbf{x}_*$  is also Gaussian with

$$\begin{aligned} \mathbb{E}[\mathbf{f}_*] &= \boldsymbol{\mu}(\mathbf{x}_*) + k(\mathbf{x}_*, \mathbf{X}) \hat{\mathbf{K}}^{-1} \mathbf{y}, \\ \text{Cov}(\mathbf{f}_*) &= k(\mathbf{x}_*, \mathbf{x}_*) - k(\mathbf{x}_*, \mathbf{X}) \hat{\mathbf{K}}^{-1} k(\mathbf{X}, \mathbf{x}_*), \end{aligned}$$

where  $\hat{\mathbf{K}} = \mathbf{K} + \sigma^2 \mathbf{I}$ .

**Hyperparameter Optimization** The computational bottleneck when optimizing kernel hyperparameters  $\boldsymbol{\theta}$  is the repeated evaluation of the *log-marginal likelihood*

$$\begin{aligned} \mathcal{L}(\boldsymbol{\theta}) &= \log p(\mathbf{y} | \mathbf{X}, \boldsymbol{\theta}) \\ &= -\frac{1}{2} (\mathbf{y}^\top \hat{\mathbf{K}}^{-1} \mathbf{y} + \underbrace{\log \det(\hat{\mathbf{K}})}_{=\text{tr}(\log(\hat{\mathbf{K}}))} + n \log(2\pi)) \quad (1) \end{aligned}$$

and its *derivative* with respect to the hyperparameters

$$\frac{\partial}{\partial \boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}) = \frac{1}{2} \mathbf{y}^\top \hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \boldsymbol{\theta}} \hat{\mathbf{K}}^{-1} \mathbf{y} - \frac{1}{2} \text{tr} \left( \hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \boldsymbol{\theta}} \right). \quad (2)$$

Computing (1) and (2) via a Cholesky decomposition has complexity  $\mathcal{O}(n^3)$  which is prohibitive for large  $n$ . In response, many methods for approximate GP inference were developed (Titsias, 2009; Hensman et al., 2013; Wilson et al., 2015; Wilson & Nickisch, 2015; Pleiss et al., 2018).

In contrast, we aim for tractable *numerically exact* GP inference in the large-scale setting (Gardner et al., 2018; Wang et al., 2019). To achieve this, we focus on efficient computation of the log-determinant in (1) and its derivative in (2) (see Section 3). This allows us to theoretically (Section 4) and empirically (Section 5) accelerate GP hyperparameter optimization.

### 2.2. Numerical Toolbox for Inference

We will use the following established numerical techniques.

**Stochastic Trace Estimation (STE)** The trace  $\text{tr}(\mathbf{A})$  of a matrix can be approximated by drawing  $\ell$  independent random vectors  $\mathbf{z}_i$  with  $\mathbb{E}[\mathbf{z}_i] = \mathbf{0}$  and  $\text{Cov}(\sqrt{n} \mathbf{z}_i) = \mathbf{I}$  and computing *Hutchinson’s estimator* (Hutchinson, 1989)

$$\tau_\ell^{\text{STE}}(\mathbf{A}) = \frac{n}{\ell} \sum_{i=1}^{\ell} \mathbf{z}_i^\top \mathbf{A} \mathbf{z}_i \approx \text{tr}(\mathbf{A}). \quad (3)$$

Here, we additionally assume the random vectors are normalized<sup>1</sup>  $\mathbf{z}_i = \tilde{\mathbf{z}}_i / \|\tilde{\mathbf{z}}_i\|_2$  and that  $\sqrt{n}(\mathbf{z}_1, \dots, \mathbf{z}_\ell)^\top \in \mathbb{R}^{\ell n}$

<sup>1</sup>Normalization is necessary for Lanczos quadrature (see Golub & Meurant, 2009, Chap. 7.2).

satisfies the convex concentration property<sup>2</sup> (see Definition S1). These assumptions are fulfilled by Rademacher-distributed random vectors with entries  $\{+1, -1\}$ .<sup>3</sup> Evaluating the log-marginal likelihood (1) requires computing  $\text{tr}(\log(\hat{\mathbf{K}}))$ . To use Hutchinson’s estimator, we need to efficiently compute  $\ell$  quadratic terms  $\{z_i^\top \log(\hat{\mathbf{K}}) z_i\}_{i=1}^\ell$ .

**Stochastic Lanczos Quadrature (SLQ)** Given a matrix function  $f$ , one can approximate bilinear forms  $z^\top f(\hat{\mathbf{K}}) z$  using quadrature (Golub & Meurant, 2009, Chap. 7). The nodes and weights of the quadrature rule can be computed efficiently via  $m$  iterations of the *Lanczos algorithm* (Lanczos, 1950) (or equivalently via CG (Gardner et al., 2018)). The combination with Hutchinson’s estimator

$$\tau_{\ell, m}^{\text{SLQ}}(f(\hat{\mathbf{K}})) \approx \tau_\ell^{\text{STE}}(f(\hat{\mathbf{K}})) \approx \text{tr}(f(\hat{\mathbf{K}})), \quad (4)$$

is called *stochastic Lanczos quadrature* (Ubaru et al., 2017).

To compute the linear solves  $v \mapsto \hat{\mathbf{K}}^{-1} v$  with the kernel matrix in (1) and (2), we use the *conjugate gradient method*.

**Conjugate Gradient Method (CG)** CG (Hestenes & Stiefel, 1952) is an iterative method for solving linear systems with symmetric positive definite matrix. It is particularly suited for large systems since it is matrix-free, and relies primarily on matrix-vector multiplication with  $\hat{\mathbf{K}}$ .

**Preconditioning** It is well-known that CG can be accelerated via a symmetric positive definite preconditioner  $\hat{\mathbf{P}} \approx \hat{\mathbf{K}}$ , by solving an equivalent linear system with matrix  $\hat{\mathbf{P}}^{-\frac{1}{2}} \hat{\mathbf{K}} \hat{\mathbf{P}}^{-\frac{1}{2}} \approx \mathbf{I}$  (Trefethen & Bau, 1997). CG’s convergence is then determined by the condition number

$$\kappa := \kappa(\hat{\mathbf{P}}^{-\frac{1}{2}} \hat{\mathbf{K}} \hat{\mathbf{P}}^{-\frac{1}{2}}) \ll \kappa(\hat{\mathbf{K}}) = \frac{|\lambda_{\max}(\hat{\mathbf{K}})|}{|\lambda_{\min}(\hat{\mathbf{K}})|}. \quad (5)$$

Suppose the approximation quality of a sequence of preconditioners  $\{\hat{\mathbf{P}}_\ell\}_\ell$  indexed by  $\ell$  is given by<sup>4</sup>

$$\|\hat{\mathbf{K}} - \hat{\mathbf{P}}_\ell\|_F \leq \mathcal{O}(g(\ell)) \|\hat{\mathbf{K}}\|_F. \quad (6)$$

If  $g(\ell) \rightarrow 0$  quickly, a small amount of precomputation can significantly accelerate CG, since by Lemma S4

$$\kappa \leq (1 + \mathcal{O}(g(\ell))) \|\hat{\mathbf{K}}\|_F^2. \quad (7)$$

Preconditioners must be cheap to obtain and allow efficient linear solves  $v \mapsto \hat{\mathbf{P}}^{-1} v$ .<sup>5</sup> As an example, *diagonal-plus-low-rank* preconditioners  $\hat{\mathbf{P}}_\ell = \sigma^2 \mathbf{I} + \mathbf{L}_\ell \mathbf{L}_\ell^\top$ , with  $\mathbf{L}_\ell \in \mathbb{R}^{n \times \ell}$ , admit linear solves in  $\mathcal{O}(n\ell^2)$  via the matrix inversion lemma.

<sup>2</sup>Concentration enables us to prove probabilistic error bounds.

<sup>3</sup>We conjecture they also hold for vectors  $\tilde{z}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ .

<sup>4</sup>The use of  $\ell$  for the number of random vectors and the preconditioner sequence is deliberate. Setting them to the same value enables variance reduction as we prove in Theorem 1.

<sup>5</sup>While CG (and Lanczos) assume a symmetric pos. definite matrix, both can be implemented using only  $\hat{\mathbf{P}}^{-1}$ , not  $\hat{\mathbf{P}}^{-\frac{1}{2}} \hat{\mathbf{K}} \hat{\mathbf{P}}^{-\frac{1}{2}}$ .

### 3. Log-Determinant Estimation

Our goal is to compute  $\log \det(\hat{\mathbf{K}}) = \text{tr}(\log(\hat{\mathbf{K}}))$  and its derivative via matrix-vector multiplication. As described, we can use stochastic trace estimation to do so. Now assume we additionally have access to a preconditioner  $\hat{\mathbf{P}} \approx \hat{\mathbf{K}}$ . As we will show, we can then not just accelerate the convergence of CG, but *also* more efficiently compute the forward and backward pass for the log-determinant.

By the properties of the matrix logarithm we can decompose the log-determinant into a *deterministic approximation* based on the preconditioner and a *residual trace* computed via stochastic trace estimation.<sup>6</sup> It holds by Lemma S3, that

$$\log \det(\hat{\mathbf{K}}) = \log \det(\hat{\mathbf{P}}_\ell) + \underbrace{\text{tr}(\log(\hat{\mathbf{K}}) - \log(\hat{\mathbf{P}}_\ell))}_{=\Delta_{\log}} \quad (8)$$

where  $\text{tr}(\Delta_{\log}) = \text{tr}(\log(\hat{\mathbf{P}}_\ell^{-\frac{1}{2}} \hat{\mathbf{K}} \hat{\mathbf{P}}_\ell^{-\frac{1}{2}}))$  and we assume  $\log \det(\hat{\mathbf{P}}_\ell)$  is efficient to compute. Equation (8) has two crucial benefits we can exploit. First and foremost, the faster  $\hat{\mathbf{P}}_\ell \rightarrow \hat{\mathbf{K}}$ , i.e.  $g(\ell) \rightarrow 0$ , the less the stochastic approximation of  $\text{tr}(\Delta_{\log})$  affects the estimate. Since its contribution to the overall error decreases the better  $\log \det(\hat{\mathbf{P}}_\ell)$  approximates  $\log \det(\hat{\mathbf{K}})$ , significantly fewer random vectors are needed to achieve a desired error with high probability. Second, we can now run Lanczos on the preconditioned matrix accelerating its convergence. As we will show later, one can also exploit (8) for the backward pass.

#### 3.1. Variance-reduced Stochastic Trace Estimation

This intuitive argument for the log-determinant also holds generally, assuming a similar decomposition exists.

**Theorem 1** (Variance-reduced Stochastic Trace Estimation)

Let  $\hat{\mathbf{K}}, \hat{\mathbf{P}}_\ell \in \mathbb{R}_{\text{spd}}^{n \times n}$ ,  $\Delta_f \in \mathbb{R}^{n \times n}$  and  $f : \mathbb{R}_{\text{spd}}^{n \times n} \rightarrow \mathbb{R}$  such that  $\text{tr}(f(\hat{\mathbf{K}})) = \text{tr}(f(\hat{\mathbf{P}}_\ell)) + \text{tr}(\Delta_f)$ , and define the estimator  $\tau_* = \text{tr}(f(\hat{\mathbf{P}}_\ell)) + \tau_\ell^{\text{STE}}(\Delta_f)$ . Now, assume there exist  $c_\Delta > 0$  and  $g : \mathbb{N} \rightarrow (0, \infty)$  such that

$$\|\Delta_f\|_F \leq c_\Delta g(\ell) \|f(\hat{\mathbf{K}})\|_F. \quad (9)$$

Then there exists  $c_z > 0$  dependent on the choice of random vectors, such that, if  $\ell \geq c_z \log(\delta^{-1})$ , it holds with probability  $1 - \delta \in [\frac{1}{2}, 1)$  that

$$|\tau_* - \text{tr}(f(\hat{\mathbf{K}}))| \leq \varepsilon_{\text{STE}} \|f(\hat{\mathbf{K}})\|_F. \quad (10)$$

where for  $C_1 = c_\Delta \sqrt{c_z}$  the relative error is given by

$$\varepsilon_{\text{STE}}(\delta, \ell) = C_1 \sqrt{\log(\delta^{-1})} \ell^{-\frac{1}{2}} g(\ell). \quad (11)$$

*Proof.* See Section S2. □

<sup>6</sup>Similar approaches have been suggested by Adams et al. (2018); Meyer et al. (2021). Our work is notably different in that it a) uses preconditioning, b) also considers the backward pass and c) gives stronger theoretical guarantees.

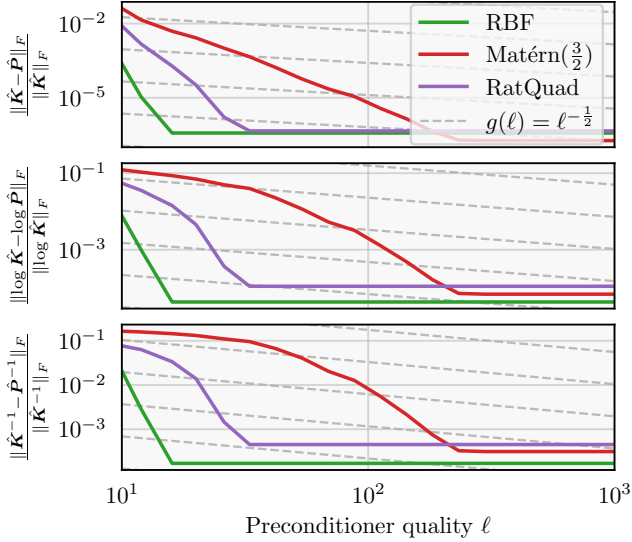


Figure 2: *Relative error of matrix functions.* For analytic functions  $f(\hat{\mathbf{P}}_\ell) \rightarrow f(\hat{\mathbf{K}})$  at the asymptotic rate of the preconditioner  $\hat{\mathbf{P}}_\ell \rightarrow \hat{\mathbf{K}}$ . Here, we use an incomplete Cholesky preconditioner on a synthetic dataset ( $n = 1,000$ ).

Notice that Theorem 1 assumes that the sequence  $\{f(\hat{\mathbf{P}}_\ell)\}_\ell$  approximates  $f(\hat{\mathbf{K}})$  sufficiently fast with  $\ell$  in (9). Intuitively, if  $\hat{\mathbf{P}}_\ell \rightarrow \hat{\mathbf{K}}$  quickly, one might expect the same for  $f(\hat{\mathbf{P}}_\ell) \rightarrow f(\hat{\mathbf{K}})$  under certain conditions on  $f$ . Indeed, one obtains the same asymptotic rate  $g(\ell)$  of the preconditioner  $\hat{\mathbf{P}}_\ell$  for the approximation of  $f(\hat{\mathbf{K}})$  by  $f(\hat{\mathbf{P}})$  (see Proposition S1). This is illustrated in Figure 2. Therefore, the error of the variance-reduced stochastic trace estimate is determined by the quality  $g(\ell)$  of the preconditioner.

**Comparison of Theorem 1 and Existing Results** Consider the case where  $f = \text{id}$ . If we are not using a preconditioner, i.e.  $\hat{\mathbf{P}} = \mathbf{0}$  and thus  $c_\Delta = g(\ell) = 1$ , we recover the well-known convergence rate  $\mathcal{O}(\ell^{-\frac{1}{2}})$  of Hutchinson’s estimator (Avron & Toledo, 2011; Roosta-Khorasani & Ascher, 2015). If instead we choose a randomized low rank approximation as a preconditioner with  $g(\ell) = \ell^{-\frac{1}{2}}$ , then Theorem 1 recovers the convergence rate  $\varepsilon_{\text{STE}} \in \mathcal{O}(\ell^{-1})$  of HUTCH++ (Meyer et al., 2021; Persson et al., 2021; Jiang et al., 2021) as a special case. However, as we will show, using preconditioning one can achieve polynomial – even exponential – convergence rates for common kernels. Such a drastic improvement is possible since neither variants of Hutchinson’s make any assumptions about the kernel matrix, whereas preconditioners are designed to leverage structure.

### 3.2. Forward Pass

We can now analyze the error of the preconditioned stochastic log-determinant estimate. Combining Theorem 1 with Lanczos quadrature error analysis, the following holds.

### Theorem 2 (Error Bound for $\log \det(\hat{\mathbf{K}})$ )

Let  $f = \log$ ,  $\Delta_{\log} = \log(\hat{\mathbf{P}}^{-\frac{1}{2}} \hat{\mathbf{K}} \hat{\mathbf{P}}^{-\frac{1}{2}})$  and assume the conditions of Theorem 1 hold. Then, with probability  $1 - \delta$ , it holds for  $\tau_*^{\log} = \log(\det(\hat{\mathbf{P}})) + \tau_{\ell, m}^{\text{SLQ}}(\Delta_{\log})$ , that

$$\left| \tau_*^{\log} - \log \det(\hat{\mathbf{K}}) \right| \leq (\varepsilon_{\text{Lanczos}} + \varepsilon_{\text{STE}}) \|\log(\hat{\mathbf{K}})\|_F,$$

where the individual errors are bounded by

$$\varepsilon_{\text{Lanczos}}(\kappa, m) \leq K_1 \left( \frac{\sqrt{2\kappa+1}-1}{\sqrt{2\kappa+1}+1} \right)^{2m} \quad (12)$$

$$\varepsilon_{\text{STE}}(\delta, \ell) \leq C_1 \sqrt{\log(\delta^{-1})} \ell^{-\frac{1}{2}} g(\ell) \quad (13)$$

and  $K_1 = \frac{5\kappa \log(2(\kappa+1))}{2\|\log(\hat{\mathbf{K}})\|_F \sqrt{2\kappa+1}}$ .

*Proof.* See Section S3.2.  $\square$

### Corollary 1

Assume the conditions of Theorem 2 hold. If the number of random vectors  $\ell$  satisfies (11) with  $\varepsilon_{\text{STE}} = \frac{\varepsilon}{2}$  and we run

$$m \geq \frac{\sqrt{3}}{4} \sqrt{\kappa} \log(2K_1 \varepsilon^{-1}) \quad (14)$$

iterations of Lanczos, then it holds that

$$\mathbb{P} \left( \left| \tau_*^{\log} - \log \det(\hat{\mathbf{K}}) \right| \leq \varepsilon \|\log(\hat{\mathbf{K}})\|_F \right) \geq 1 - \delta.$$

*Proof.* See Section S3.2.  $\square$

We note two major improvements over the bound by Ubaru et al. (2017, Corollary 4.5). First, the number of Lanczos steps now depends on the condition number  $\kappa$  of the preconditioned matrix, implying faster convergence. Second, depending on the preconditioner quality  $g(\ell)$ , we need significantly fewer random vectors by Theorem 1.

### 3.3. Backward Pass

By differentiating through (8), we obtain a decomposition into a *deterministic approximation* based on the preconditioner and a *residual trace* for the backward pass. For  $\Delta_{\text{inv}\partial} = \hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \theta} - \hat{\mathbf{P}}^{-1} \frac{\partial \hat{\mathbf{P}}}{\partial \theta}$ , we have

$$\frac{\partial}{\partial \theta} \log \det(\hat{\mathbf{K}}) = \text{tr}(\hat{\mathbf{P}}^{-1} \frac{\partial \hat{\mathbf{P}}}{\partial \theta}) + \text{tr}(\Delta_{\text{inv}\partial}), \quad (15)$$

Therefore the stochastic trace estimator

$$\tau_{\ell, m}^{\text{SCG}}(\Delta_{\text{inv}\partial}) \approx \text{tr}(\Delta_{\text{inv}\partial}) \quad (16)$$

requires solves  $\mathbf{z}_i^\top \mathbf{K}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \theta} \mathbf{z}_i$  and  $\mathbf{z}_i^\top \hat{\mathbf{P}}^{-1} \frac{\partial \hat{\mathbf{P}}}{\partial \theta} \mathbf{z}_i$ . The former can be computed with  $m$  iterations of preconditioned CG, while the latter is simply a solve with the preconditioner. Note that the deterministic term  $\text{tr}(\hat{\mathbf{P}}^{-1} \frac{\partial \hat{\mathbf{P}}}{\partial \theta})$  is efficient to calculate for many types of preconditioners. For example, if  $\hat{\mathbf{P}}$  is a diagonal-plus-low-rank preconditioner it can be computed in  $\mathcal{O}(n\ell^2)$  (see Section S3.3). Using Theorem 1, we obtain a probabilistic error bound for the derivative estimate.

**Theorem 3** (Error Bound for  $\text{tr}(\hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \theta})$ )

Let  $f(\hat{\mathbf{K}}) = \hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \theta}$ ,  $\Delta_{\text{inv}\partial} = \hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \theta} - \hat{\mathbf{P}}^{-1} \frac{\partial \hat{\mathbf{P}}}{\partial \theta}$  and assume the conditions of Theorem 1 hold. If we solve  $\hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \theta} \mathbf{z}_i$  with  $m$  iterations of preconditioned CG, initialized at  $\mathbf{0}$  or better, then it holds with probability  $1 - \delta$  for  $\tau_*^{\text{inv}\partial} = \text{tr}(\hat{\mathbf{P}}^{-1} \frac{\partial \hat{\mathbf{P}}}{\partial \theta}) + \tau_{\ell, m}^{\text{SCG}}(\Delta_{\text{inv}\partial})$ , that

$$|\tau_*^{\text{inv}\partial} - \text{tr}(\hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \theta})| \leq (\varepsilon_{\text{CG}'} + \varepsilon_{\text{STE}}) \|\hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \theta}\|_F,$$

where the individual errors are bounded by

$$\varepsilon_{\text{CG}'}(\kappa, m) \leq K_2 \left( \frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1} \right)^m \quad (17)$$

$$\varepsilon_{\text{STE}}(\delta, \ell) \leq C_1 \sqrt{\log(\delta^{-1})} \ell^{-\frac{1}{2}} g(\ell) \quad (18)$$

and  $K_2 = 2\sqrt{\kappa(\hat{\mathbf{K}})}n$ .

*Proof.* See Section S3.3.  $\square$

**Corollary 2**

Assume the conditions of Theorem 3 hold. If the number of random vectors  $\ell$  satisfies (11) with  $\varepsilon_{\text{STE}} = \frac{\varepsilon}{2}$ , and we run

$$m \geq \frac{1}{2} \sqrt{\kappa} \log(2K_2 \varepsilon^{-1}) \quad (19)$$

iterations of CG, then

$$\mathbb{P} \left( \left| \tau_*^{\text{inv}\partial} - \text{tr}(\hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \theta}) \right| \leq \varepsilon \|\hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \theta}\|_F \right) \geq 1 - \delta.$$

*Proof.* See Section S3.3.  $\square$

## 4. Efficient GP Hyperparameter Optimization

Having established an efficient way to compute the forward and backward pass for the log-determinant, we can use these results to accelerate GP hyperparameter optimization by fully exploiting preconditioning not just for the linear solves, but *also* for the log-determinant and its derivative.

### 4.1. Log-Marginal Likelihood

We obtain a bound on the log-marginal likelihood by combining Theorem 2 with standard CG convergence analysis.

**Theorem 4** (Error Bound for the log-Marginal Likelihood)

Assume the conditions of Theorem 2 hold and we solve  $\hat{\mathbf{K}}\mathbf{u} = \mathbf{y}$  via preconditioned CG initialized at  $\mathbf{u}_0$  and terminated after  $m$  iterations. Then with probability  $1 - \delta$ , the error in the estimate  $\eta = -\frac{1}{2}(\mathbf{y}^\top \mathbf{u}_m + \tau_*^{\log} + n \log(2\pi))$  of the log-marginal likelihood  $\mathcal{L}$  satisfies

$$|\eta - \mathcal{L}| \leq \varepsilon_{\text{CG}} + \frac{1}{2}(\varepsilon_{\text{Lanczos}} + \varepsilon_{\text{STE}}) \|\log(\hat{\mathbf{K}})\|_F,$$

where the individual errors are bounded by

$$\varepsilon_{\text{CG}}(\kappa, m) \leq K_3 \left( \frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1} \right)^m \quad (20)$$

$$\varepsilon_{\text{Lanczos}}(\kappa, m) \leq K_1 \left( \frac{\sqrt{2\kappa+1}-1}{\sqrt{2\kappa+1}+1} \right)^{2m} \quad (21)$$

$$\varepsilon_{\text{STE}}(\delta, \ell) \leq C_1 \sqrt{\log(\delta^{-1})} \ell^{-\frac{1}{2}} g(\ell) \quad (22)$$

for  $K_3 = \sqrt{\kappa(\hat{\mathbf{K}})} \|\mathbf{y}\|_2 \|\mathbf{u}_0 - \mathbf{u}\|_2$ .

*Proof.* See Section S4.1.  $\square$

### 4.2. Derivative of the Log-Marginal Likelihood

Similarly, we can leverage Theorem 3 for the derivative.

**Theorem 5** (Error Bound for the Derivative)

Assume the conditions of Theorem 3 hold and we solve  $\hat{\mathbf{K}}\mathbf{u} = \mathbf{y}$  via preconditioned CG initialized at  $\mathbf{0}$  or better and terminated after  $m$  iterations. Then with probability  $1 - \delta$ , the error in the estimate  $\phi = \frac{1}{2}(\mathbf{u}_m^\top \frac{\partial \hat{\mathbf{K}}}{\partial \theta} \mathbf{u}_m - \tau_*^{\text{inv}\partial})$  of the derivative of the log-marginal likelihood  $\frac{\partial}{\partial \theta} \mathcal{L}$  satisfies

$$|\phi - \frac{\partial}{\partial \theta} \mathcal{L}| \leq \varepsilon_{\text{CG}}(\kappa, m) + \frac{1}{2}(\varepsilon_{\text{CG}'} + \varepsilon_{\text{STE}}) \|\hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \theta}\|_F,$$

where the individual errors are bounded by

$$\varepsilon_{\text{CG}}(\kappa, m) \leq K_4 \left( \frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1} \right)^m \quad (23)$$

$$\varepsilon_{\text{CG}'}(\kappa, m) \leq K_2 \left( \frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1} \right)^m \quad (24)$$

$$\varepsilon_{\text{STE}}(\delta, \ell) \leq C_1 \sqrt{\log(\delta^{-1})} \ell^{-\frac{1}{2}} g(\ell) \quad (25)$$

for  $K_4 = 6\kappa(\hat{\mathbf{K}}) \max(\|\mathbf{u}\|_2, \|\mathbf{u}\|_2^3) \|\frac{\partial \hat{\mathbf{K}}}{\partial \theta}\|_2$ .

*Proof.* See Section S4.2.  $\square$

### 4.3. Preconditioner Choice

Our theoretical convergence results fundamentally depend on how quickly the preconditioner approximates the kernel matrix, either directly via  $g(\ell)$ , or indirectly via the condition number improvement (7). This leaves the question which preconditioners should be chosen in practice and what rates  $g(\ell)$  they attain. In Table 1, we give an extensive list of kernel-preconditioner combinations with associated rates (see Section S5 for proofs). This includes the commonly used RBF and Matérn( $\nu$ ) kernels for which the incomplete Cholesky (Kershaw, 1978) and QFF (Mutný & Krause, 2018) preconditioners result in exponential and polynomial convergence rates, respectively. For STE in this context this is a substantial improvement over the rate of Hutchinson's estimator  $\mathcal{O}(\ell^{-\frac{1}{2}})$  (Avron & Toledo, 2011; Roosta-Khorasani & Ascher, 2015; Skorski, 2021) and HUTCH++ with  $\mathcal{O}(\ell^{-1})$  (Meyer et al., 2021; Persson et al., 2021; Jiang et al., 2021). Depending on the problem this can mean a difference of tens vs. thousands of random vectors. To the best of our knowledge, for the use of CG in GP inference, only the one-dimensional RBF kernel and incomplete Cholesky preconditioner have been previously analyzed theoretically (Gardner et al., 2018). In contrast, Table 1 gives convergence rates for arbitrary  $d$ -dimensional kernels and multiple preconditioners. In fact, our results also apply to any kernel

Table 1: *Error rates for combinations of kernels and preconditioners.* The rate  $g(\ell)$  measures how fast a sequence of preconditioners  $\{\hat{P}_\ell\}_\ell$  approaches the kernel matrix  $\hat{K}$  constructed from data  $\mathbf{X} \in \mathbb{R}^{n \times d}$ . Thus it determines both the convergence speed of Krylov methods and the preconditioned stochastic trace estimator. Or, equivalently, the faster  $g(\ell) \rightarrow 0$  the fewer CG iterations  $m$  and random vectors  $\ell$  are needed to approximate the log-marginal likelihood and its gradient.

Kernel	$d$	Preconditioner	$g(\ell)$	Condition	Proof
RBF	1	incomplete Cholesky	$\exp(-c\ell)$	for some $c > 0$	Proposition S2
RBF	$\mathbb{N}$	QFF	$\exp(-b\ell^{\frac{1}{d}})$	for some $b > 0$ if $\ell^{\frac{1}{d}} > 2\gamma^{-2}$	Proposition S3
Matérn( $\nu$ )	$\mathbb{N}$	incomplete Cholesky	$\ell^{-(\frac{2\nu}{d}+1)}$	$2\nu \in \mathbb{N}$ and maximin ordering	(Schäfer et al., 2021)
Matérn( $\nu$ )	1	QFF	$\ell^{-(s(\nu)+1)}$	where $s(\nu) \in \mathbb{N}$	Proposition S4
mod. Matérn( $\nu$ )	$\mathbb{N}$	QFF	$\ell^{-\frac{s(\nu)+1}{d}}$	where $s(\nu) \in \mathbb{N}$	Proposition S4
any	$\mathbb{N}$	truncated SVD	$\ell^{-\frac{1}{2}}$		Proposition S5
any	$\mathbb{N}$	randomized SVD	$\ell^{-\frac{1}{2}} + \mathcal{O}(\ell^{\frac{1}{4}} s^{-\frac{1}{4}})$	w/ high probability for $s$ samples	Proposition S6
any	$\mathbb{N}$	randomized Nyström	$\ell^{-\frac{1}{2}} + \mathcal{O}(\ell^{\frac{1}{4}} s^{-\frac{1}{4}})$	w/ high probability for $s$ samples	Proposition S7
any	$\mathbb{N}$	RFF	$\ell^{-\frac{1}{2}}$	w/ high probability	Proposition S8
additive	$\mathbb{N}$	any	$dg(\ell)$	all summands have rate $g(\ell)$	Lemma S5
any	$\mathbb{N}$	any approximation	$g(\ell)$	$\exists$ uniform convergence bound	Lemma S6

approximation with a uniform convergence bound (such as RFF (Rahimi et al., 2007)). All the while for many, e.g. diagonal-plus-low-rank preconditioners, the amount of pre-computation needed *amortizes with more data*, i.e. the cost of preconditioning becomes negligible the larger the dataset.

#### 4.4. Algorithms

The above leads to Algorithms 1 and 2 computing  $\mathcal{L}$  and  $\frac{\partial}{\partial \theta} \mathcal{L}$  for GP hyperparameter optimization.<sup>7</sup> Our algorithms are similar to those presented in prior work by Cutajar et al. (2016); Ubaru et al. (2017); Gardner et al. (2018), yet crucially they leverage preconditioning for **faster CG convergence and variance reduction** of the log-determinant estimate and its derivative. In the following,  $\text{CG}(\hat{K}, \mathbf{y}, \hat{P}, m)$  denotes a CG solve of  $\hat{K}\mathbf{u} = \mathbf{y}$  with preconditioner  $\hat{P}$  run for  $m$  iterations. Here, we equivalently use CG instead of Lanczos, as suggested by Gardner et al. (2018).

**Computational Complexity** Algorithm 1 has complexity  $\mathcal{O}(n^2 m \ell + P_{\log \det})$  and Algorithm 2 has complexity  $\mathcal{O}((n^2 m + P_{\text{solve}})\ell + P_{\text{tr inv } \partial})$ , where  $P_{(\cdot)}$  denotes the cost of an operation with the preconditioner.<sup>8</sup> Assuming  $m, \ell \ll n$ , this is asymptotically faster than Cholesky-based inference with complexity  $\mathcal{O}(n^3)$ . Due to the reduction to matrix-vector multiplication, if  $\mathbf{v} \mapsto \hat{K}\mathbf{v}$  is more efficient than  $\mathcal{O}(n^2)$  (e.g. for structured or sparse matrices) the complexity reduces further. Finally, the **for**-loops are embarrassingly parallel, giving additional speedup in practice.

<sup>7</sup>While presented sequentially for clarity, in practice one would presample all random vectors and run a single call of (parallelized) CG with multiple right hand sides, as in (Gardner et al., 2018).

<sup>8</sup>For diagonal-plus-low-rank preconditioners, such as the incomplete Cholesky,  $P_{\text{solve}}$ ,  $P_{\log \det}$ , and  $P_{\text{tr inv } \partial}$  are in  $\mathcal{O}(n\ell^2)$  by the matrix inversion and determinant lemmas.

#### Algorithm 1: log-Marginal Likelihood

**Input:**  $\mathbf{y}$  (labels),  $\hat{K}$  (kernel matrix),  $\hat{P}$  (preconditioner),  $\ell$  (# of random STE vectors),  $m$  (# of CG iterations)

- 1 **procedure** LOGMARGLIKELIHOOD( $\mathbf{y}, \hat{K}, \hat{P}, \ell, m$ )
- 2  $\mathbf{u} \leftarrow \text{CG}(\hat{K}, \mathbf{y}, \hat{P}, m)$   $\triangleright \approx \hat{K}^{-1}\mathbf{y}$
- 3  $\tau_{\hat{P}}^{\log} \leftarrow \log \det(\hat{P})$
- 4 **for**  $i = 1, \dots, \ell$  **do**
- 5  $\mathbf{z}_i \leftarrow \tilde{\mathbf{z}}_i / \|\tilde{\mathbf{z}}_i\|_2$  for rand. vector  $\tilde{\mathbf{z}}_i$
- 6  $\mathbf{T} \leftarrow \text{CG}(\hat{K}, \mathbf{z}_i, \hat{P}, m)$   $\triangleright$  equiv. to LANCZOS
- 7  $[\mathbf{W}, \boldsymbol{\lambda}] \leftarrow \text{EIGENDECOMP}(\mathbf{T})$   $\triangleright \mathbf{T}$  tridiagonal
- 8  $\omega_j \leftarrow (\mathbf{e}_j^\top \boldsymbol{\omega}_j)^2$  for  $j = 0, \dots, m$   $\triangleright$  quad. weights
- 9  $\gamma_i \leftarrow \sum_{j=0}^m \omega_j \log(\lambda_j)$   $\triangleright \approx \mathbf{z}_i^\top \boldsymbol{\Delta}_{\log \mathbf{z}_i}$
- 10  $\tau_*^{\log} \leftarrow \tau_{\hat{P}}^{\log} + \frac{n}{\ell} \sum_{i=1}^{\ell} \gamma_i$   $\triangleright \approx \log \det(\hat{K})$
- 11 **return**  $-\frac{1}{2}(\mathbf{y}^\top \mathbf{u} + \tau_*^{\log} + n \log(2\pi))$   $\triangleright \approx \mathcal{L}(\boldsymbol{\theta})$

#### Algorithm 2: Derivative of the log-Marginal Likelihood

**Input:**  $\mathbf{y}$  (labels),  $\hat{K}$  (kernel matrix),  $\hat{P}$  (preconditioner),  $\ell$  (# of random STE vectors),  $m$  (# of CG iterations),  $\frac{\partial \hat{K}}{\partial \theta} / \frac{\partial \hat{P}}{\partial \theta}$  (functions for computing kernel / preconditioner derivatives)

- 1 **procedure** DERIVATIVE( $\mathbf{y}, \hat{K}, \frac{\partial \hat{K}}{\partial \theta}, \hat{P}, \frac{\partial \hat{P}}{\partial \theta}, \ell, m$ )
- 2  $\mathbf{u} \leftarrow \text{CG}(\hat{K}, \mathbf{y}, \hat{P}, m)$   $\triangleright \approx \hat{K}^{-1}\mathbf{y}$
- 3  $\tau_{\hat{P}}^{\text{inv } \partial} \leftarrow \text{tr}(\hat{P}^{-1} \frac{\partial \hat{P}}{\partial \theta})$
- 4 **for**  $i = 1, \dots, \ell$  **do**
- 5  $\mathbf{z}_i \leftarrow \tilde{\mathbf{z}}_i / \|\tilde{\mathbf{z}}_i\|_2$  for rand. vector  $\tilde{\mathbf{z}}_i$
- 6  $\mathbf{w}_i \leftarrow \text{CG}(\hat{K}, \frac{\partial \hat{K}}{\partial \theta} \mathbf{z}_i, \hat{P}, m)$   $\triangleright \approx \hat{K}^{-1} \frac{\partial \hat{K}}{\partial \theta} \mathbf{z}_i$
- 7  $\tilde{\mathbf{w}}_i \leftarrow \hat{P}^{-1} \frac{\partial \hat{P}}{\partial \theta} \mathbf{z}_i$
- 8  $\gamma_i \leftarrow \mathbf{z}_i^\top (\mathbf{w}_i - \tilde{\mathbf{w}}_i)$   $\triangleright \approx \mathbf{z}_i^\top \boldsymbol{\Delta}_{\text{inv } \partial} \mathbf{z}_i$
- 9  $\tau_*^{\text{inv } \partial} \leftarrow \tau_{\hat{P}}^{\text{inv } \partial} + \frac{n}{\ell} \sum_{i=1}^{\ell} \gamma_i$   $\triangleright \approx \text{tr}(\hat{K}^{-1} \frac{\partial \hat{K}}{\partial \theta})$
- 10 **return**  $\frac{1}{2}(\mathbf{u}^\top \frac{\partial \hat{K}}{\partial \theta} \mathbf{u} - \tau_*^{\text{inv } \partial})$   $\triangleright \approx \frac{\partial}{\partial \theta} \mathcal{L}(\boldsymbol{\theta})$

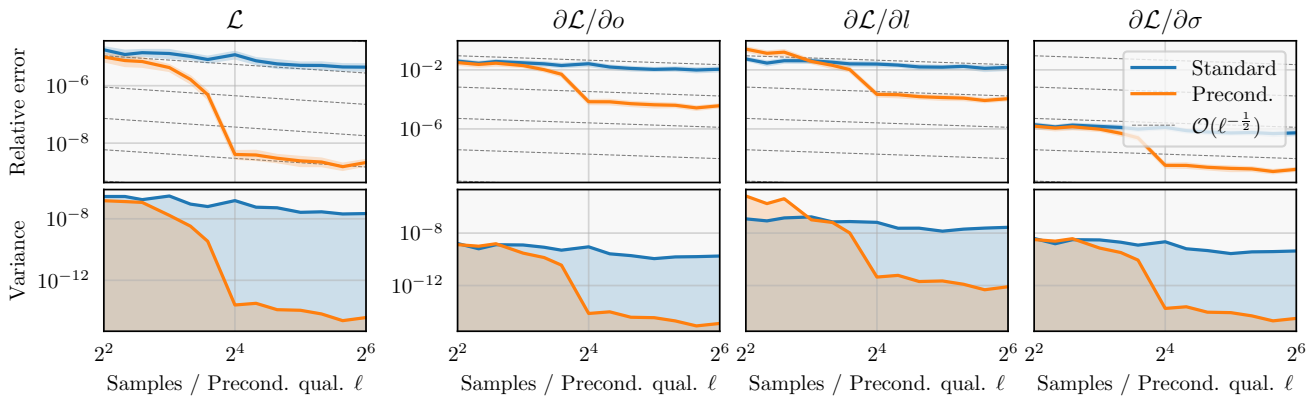


Figure 3: *Bias and variance of the estimators for the log-marginal likelihood  $\mathcal{L}$  and its derivatives.* The relative error and variance decrease faster with the number of random vectors  $\ell$  when using a preconditioner  $\hat{P}_\ell$ . The decrease rate  $\mathcal{O}(\ell^{-\frac{1}{2}}g(\ell))$ , determined by the preconditioner, significantly improves upon the standard Hutchinson’s rate  $\mathcal{O}(\ell^{-\frac{1}{2}})$ .

#### 4.5. Related Work

Krylov methods have been used for GP inference since the work of Gibbs (1997). While these methods were primarily relegated to structured GPs that afford fast matrix-vector products (Cunningham et al., 2008; Saatçi, 2012; Wilson & Nickisch, 2015), they have seen growing use as a general purpose method, especially when coupled with parallel hardware (Murray, 2009; Cutajar et al., 2016; Gardner et al., 2018; Wang et al., 2019; Artemev et al., 2021). Some recent works propose variance-free (but biased) estimates of the GP log-determinant (e.g. Artemev et al., 2021), though many works compute this term by combining STE (Hutchinson, 1989; Bekas et al., 2007; Avron & Toledo, 2011; Roosta-Khorasani & Ascher, 2015) with SLQ (Golub & Meurant, 2009; Ubaru et al., 2017; Dong et al., 2017; Gardner et al., 2018; Cortinovis & Kressner, 2020). Our work builds on ideas for variance-reduced STE (Adams et al., 2018; Meyer et al., 2021; Persson et al., 2021; Jiang et al., 2021), but, by leveraging preconditioning, requires significantly fewer random vectors than existing approaches. When applied to GP hyperparameter optimization, we obtain stronger theoretical guarantees for the forward pass than previously known (Ubaru et al., 2017) and novel guarantees for the backward pass. Finally, our results on preconditioners for kernel matrices (Table 1) give a rigorous foundation to their proposed use by Cutajar et al. (2016).

## 5. Experiments

We validate our theoretical findings empirically via GP hyperparameter optimization on synthetic and benchmark datasets with and without preconditioning. We find that

- (a) *preconditioning reduces bias and variance* in the forward and backward pass, which results in
- (b) *less noisy search directions* and *fewer log-likelihood*

*and gradient evaluations for the line search.*

This allows the use of rapidly converging optimizers, and

- (c) *accelerates training significantly.*

**Experimental Setup** We consider a one-dimensional synthetic dataset of  $n = 10,000$  iid standard normal samples, as well as a range of UCI datasets (Dua & Graff, 2017) with training set sizes ranging from  $n = 12,449$  to  $326,155$  (see Table 2). All experiments were performed on single NVIDIA GPUs, a GeForce RTX 2080 and Titan RTX, respectively. We perform GP regression using an RBF and Matérn( $\frac{3}{2}$ ) kernel with output scale  $o$ , lengthscales  $l_j$  – one per input dimension – and noise  $\sigma^2$ . Hyperparameters were optimized with L-BFGS using an Armijo-Wolfe line search and early stopping via a validation set. We use an incomplete Cholesky preconditioner throughout. An implementation of our method is available as part of GPYTORCH (Gardner et al., 2018).<sup>9</sup>

**Preconditioning reduces bias & variance in  $\mathcal{L}$  and  $\frac{\partial}{\partial\theta}\mathcal{L}$**  Figure 3 shows the relative error of the marginal log-likelihood and its derivatives on synthetic data. Already for  $\ell \geq 16$  random samples *bias and variance are reduced by several orders of magnitude*. We observe exponential decrease and then a return to the standard Hutchinson’s rate of  $\mathcal{O}(\ell^{-\frac{1}{2}})$ . After  $\ell = 16$  iterations the algorithm computing the preconditioner has reached a specified tolerance and terminates, invalidating the approximation quality assumption (6) for  $\ell \geq 16$ . Similar observations hold for the Matérn and RatQuad kernel (see Table S1 and Figure S1). As predicted by Theorem 1 and illustrated by Figure 3, the variance reduction is determined by the preconditioner. For higher dimensions, the rate  $g(\ell)$  generally slows (see Table 1), which in turn reduces the bias and variance reduction achieved by

<sup>9</sup>[github.com/cornelius-gp/gpytorch](https://github.com/cornelius-gp/gpytorch)

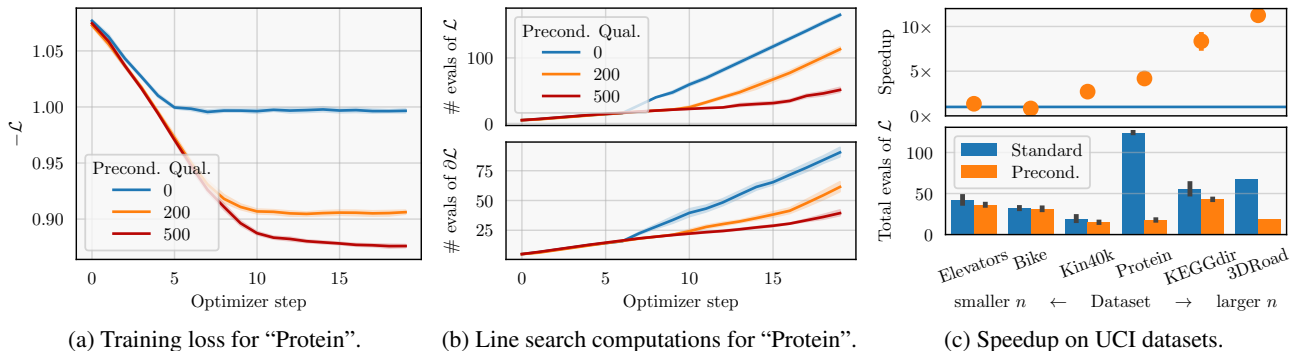


Figure 4: *Preconditioning reduces noise and in turn accelerates hyperparameter optimization.* Variance reduction improves optimization via better search directions and fewer evaluations of  $\mathcal{L}$  and  $\frac{\partial}{\partial \theta_i} \mathcal{L}$  for the line search. (a) Training loss and (b) model evaluations for line search decrease with preconditioner size, as shown for the “Protein” dataset. (c) The reduction in loss (and gradient) evaluations and of noise in the gradients results in an order of magnitude speedup on UCI datasets.

Table 2: *Hyperparameter optimization on UCI datasets.* GP regression using a Matérn( $\frac{3}{2}$ ) kernel and incomplete Cholesky preconditioner of size 500 with  $\ell = 50$  random samples. Hyperparameters were optimized with L-BFGS for at most 20 steps using early stopping. All results, but “3DRoad”, are averaged over 10 runs. Differences  $\geq 1$  standard deviation in bold.

Dataset	$n$	$d$	$-\mathcal{L}_{\text{train}} \downarrow$		$-\mathcal{L}_{\text{test}} \downarrow$		RMSE $\downarrow$		Runtime (s)	
			Standard	Precond.	Standard	Precond.	Standard	Precond.	Standard	Precond.
Elevators	12,449	18	0.4647	<b>0.4377</b>	0.4021	0.4022	0.3484	0.3482	53	<b>39</b>
Bike	13,034	17	-0.9976	-0.9985	-0.9934	-0.9877	0.0446	0.0454	<b>31</b>	37
Kin40k	30,000	8	-0.3339	<b>-0.4332</b>	-0.3141	-0.3135	<b>0.0929</b>	0.0949	187	<b>45</b>
Protein	34,297	9	0.9963	<b>0.9273</b>	0.8869	0.8835	0.5722	<b>0.5577</b>	893	<b>43</b>
KEGGdir	36,620	20	-0.9501	<b>-1.0043</b>	-0.9459	-0.9490	0.0861	0.0864	1450	<b>174</b>
3DRoad	326,155	3	0.7733	<b>0.1284</b>	1.4360	<b>1.1690</b>	0.2982	<b>0.1265</b>	82,200	<b>7306</b>

our method (see Table S1). However, on real datasets we still see strong variance reduction via our method, possibly since real data often lies on a low-dimensional manifold.

**Preconditioning accelerates hyperparameter optimization** On datasets from the UCI repository, we find that preconditioning results in lower training loss  $-\mathcal{L}(\theta)$  (illustrated in Figure 4(a)) on almost all datasets and essentially identical generalization error (see Table 2). Reducing stochasticity via preconditioning significantly lowers the number of  $\mathcal{L}$  and  $\frac{\partial}{\partial \theta} \mathcal{L}$  evaluations for the line search during optimization (see Figure 4(b)) and results in less noisy search directions. In fact, the noise in the loss and gradients caused by stochastic trace estimation previously necessitated the use of slower converging, but more noise-robust optimizers (Wang et al., 2019), such as Adam (Kingma & Ba, 2015). As these experiments show, our variance-reduced estimators make the use of L-BFGS possible, which significantly outperforms Adam (c.f. Table 2 and Table S2). These combined effects due to preconditioning *accelerate training up to twelvefold*, as Figure 4(c) shows. We observe that the speedup increases with the size of the dataset. This is partly explained by the amortization of the cost of computing and applying the preconditioner with increasing  $n$ .

## 6. Conclusion

One might reasonably hope that structural knowledge about the kernel matrix can accelerate GP hyperparameter optimization. Preconditioning is a way to encode and exploit such structure. As we showed, it can be used to great effect – not only for the solution of linear systems – but importantly *also* for stochastic approximation of the log-determinant and its derivative. Our convergence results combined with the rates for kernel-preconditioner pairs in Table 1 rigorously explain why preconditioning has been observed empirically to be so effective for large-scale GP inference (Cutajar et al., 2016; Gardner et al., 2018; Wang et al., 2019).

In fact, our work implies that software packages for GPs, which make use of Krylov methods for inference, should not use a fixed preconditioner. Instead, the preconditioner should be automatically chosen based on the specified model. While we derive a range of such kernel and preconditioner combinations, it is likely that better preconditioners exist for certain kernels or types of data. Other scientific fields invest substantial research effort into the design of preconditioners, e.g. for PDEs (Saad, 2003). Our work strongly suggests that, similarly, developing specialized preconditioners is a promising approach to scale Gaussian processes.



## Acknowledgements

JW and PH gratefully acknowledge financial support by the European Research Council through ERC StG Action 757275 / PANAMA; the DFG Cluster of Excellence “Machine Learning - New Perspectives for Science”, EXC 2064/1, project number 390727645; the German Federal Ministry of Education and Research (BMBF) through the Tübingen AI Center (FKZ: 01IS18039A); and funds from the Ministry of Science, Research and Arts of the State of Baden-Württemberg. JW is grateful to the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for support. GP and JPC are supported by the Simons Foundation, McKnight Foundation, the Grossman Center, and the Gatsby Charitable Trust.

The authors would like to thank Marius Hobbhahn, Lukas Tatzel and Felix Dangel for helpful feedback on an earlier version of this manuscript.

## References

- Adamczak, R. A note on the Hanson-Wright inequality for random vectors with dependencies. *Electronic Communications in Probability*, 20:1 – 13, 2015.
- Adams, R. P., Pennington, J., Johnson, M. J., Smith, J., Ovadia, Y., Patton, B., and Saunderson, J. Estimating the spectral density of large implicit matrices. *arXiv preprint*, 2018. URL <http://arxiv.org/abs/1802.0345>.
- Artemev, A., Burt, D. R., and van der Wilk, M. Tighter bounds on the log marginal likelihood of Gaussian process regression using conjugate gradients. In *International Conference on Machine Learning (ICML)*, 2021.
- Avron, H. and Toledo, S. Randomized algorithms for estimating the trace of an implicit symmetric positive semi-definite matrix. *Journal of the ACM (JACM)*, 58(2):1–34, 2011.
- Bekas, C., Kokiopoulou, E., and Saad, Y. An estimator for the diagonal of a matrix. *Applied numerical mathematics*, 57(11–12):1214–1229, 2007.
- Charlier, B., Feydy, J., Glaunès, J. A., Collin, F.-D., and Durif, G. Kernel operations on the GPU, with autodiff, without memory overflows. *Journal of Machine Learning Research (JMLR)*, 22(74):1–6, 2021. URL <http://jmlr.org/papers/v22/20-275.html>.
- Cortinovis, A. and Kressner, D. On randomized trace estimates for indefinite matrices with an application to determinants. *arXiv preprint*, 2020. URL <http://arxiv.org/abs/2005.10009>.
- Cunningham, J. P., Shenoy, K. V., and Sahani, M. Fast Gaussian process methods for point process intensity estimation. In *International Conference on Machine Learning (ICML)*, 2008.
- Cutajar, K., Osborne, M., Cunningham, J., and Filippone, M. Preconditioning kernel matrices. In *International Conference on Machine Learning (ICML)*, 2016.
- Dong, K., Eriksson, D., Nickisch, H., Bindel, D., and Wilson, A. G. Scalable log determinants for Gaussian process kernel learning. *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 6330–6340, 2017.
- Drineas, P., Mahoney, M. W., and Cristianini, N. On the Nyström method for approximating a Gram matrix for improved kernel-based learning. *Journal of Machine Learning Research (JMLR)*, 6(12), 2005.
- Drineas, P., Kannan, R., and Mahoney, M. W. Fast Monte Carlo algorithms for matrices II: Computing a low-rank approximation to a matrix. *SIAM Journal on Computing (SICOMP)*, 36:158–183, 2006.
- Dua, D. and Graff, C. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- Eckart, C. and Young, G. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3):211–218, 1936.
- Gardner, J. R., Pleiss, G., Bindel, D., Weinberger, K. Q., and Wilson, A. G. GPyTorch: Blackbox matrix-matrix Gaussian process inference with GPU acceleration. *Advances in Neural Information Processing Systems (NeurIPS)*, 2018:7576–7586, 2018.
- Gibbs, M. Bayesian Gaussian processes for classification and regression. *University of Cambridge, Cambridge*, 1997.
- Golub, G. H. and Meurant, G. *Matrices, moments and quadrature with applications*, volume 30. Princeton University Press, 2009.
- Golub, G. H. and van Loan, C. F. *Matrix Computations*. JHU Press, fourth edition, 2013.
- Harbrecht, H., Peters, M., and Schneider, R. On the low-rank approximation by the pivoted Cholesky decomposition. *Applied numerical mathematics*, 62(4):428–440, 2012.
- Hensman, J., Fusi, N., and Lawrence, N. D. Gaussian processes for big data. In *Proceedings of the 29th Conference on Uncertainty in Artificial Intelligence (UAI)*, pp. 282–290, 2013.
- Hestenes, M. R. and Stiefel, E. Methods of conjugate gradients for solving linear systems. *Journal of Research of the National Bureau of Standards*, 49, 1952.
- Hutchinson, M. F. A stochastic estimator of the trace of the influence matrix for Laplacian smoothing splines. *Communications in Statistics-Simulation and Computation*, 18(3):1059–1076, 1989.
- Jiang, S., Pham, H., Woodruff, D. P., Qiuyi, and Zhang. Optimal sketching for trace estimation. *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- Kasiswathan, S. P. and Rudelson, M. Restricted isometry property under high correlations. *arXiv preprint*, 2019. URL <http://arxiv.org/abs/1904.05510>.
- Kershaw, D. S. The incomplete Cholesky—conjugate gradient method for the iterative solution of systems of linear equations. *Journal of Computational Physics*, 26(1):43–65, 1978.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *International Conference on Learning Representations (ICLR)*, 2015.

- Kittaneh, F. On Lipschitz functions of normal operators. *Proceedings of the American Mathematical Society*, 94(3):416–418, 1985.
- Lanczos, C. *An iteration method for the solution of the eigenvalue problem of linear differential and integral operators*. United States Government Press Office Los Angeles, CA, 1950.
- Ledoux, M. *The concentration of measure phenomenon*. Number 89. American Mathematical Society, 2001.
- Meyer, R. A., Musco, C., Musco, C., and Woodruff, D. P. Hutch++: Optimal stochastic trace estimation. In *Symposium on Simplicity in Algorithms (SOSA)*, pp. 142–155. SIAM, 2021.
- Mitrinovic, D. S. and Vasic, P. M. *Analytic inequalities*, volume 1. Springer, 1970.
- Murray, I. Gaussian processes and fast matrix-vector multiplies. In *Numerical Mathematics in Machine Learning Workshop (ICML)*, 2009.
- Mutny, M. and Krause, A. Efficient high dimensional Bayesian optimization with additivity and quadrature Fourier features. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 9005–9016, 2018.
- Nocedal, J. and Wright, S. *Numerical optimization*. Springer Science & Business Media, 2006.
- Persson, D., Cortinovis, A., and Kressner, D. Improved variants of the Hutch++ algorithm for trace estimation. *arXiv preprint*, 2021. URL <http://arxiv.org/abs/2109.10659>.
- Pleiss, G., Gardner, J., Weinberger, K., and Wilson, A. G. Constant-time predictive distributions for Gaussian processes. In *Proceedings of the International Conference on Machine Learning (ICML)*, pp. 4114–4123, 2018.
- Potapczynski, A., Wu, L., Biderman, D., Pleiss, G., and Cunningham, J. P. Bias-free scalable Gaussian processes via randomized truncations. In *International Conference on Machine Learning (ICML)*, 2021.
- Rahimi, A., Recht, B., et al. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2007.
- Rasmussen, C. E. and Williams, C. K. I. *Gaussian Processes for Machine Learning*. The MIT Press, 2006.
- Roosta-Khorasani, F. and Ascher, U. Improved bounds on sample size for implicit matrix trace estimators. *Foundations of Computational Mathematics*, 15(5):1187–1212, 2015.
- Saad, Y. *Iterative methods for sparse linear systems*. Society for Industrial and Applied Mathematics (SIAM), 2003.
- Saatçi, Y. *Scalable inference for structured Gaussian process models*. PhD thesis, University of Cambridge, 2012.
- Samson, P.-M. Concentration of measure inequalities for Markov chains and  $\phi$ -mixing processes. *The Annals of Probability*, 28(1):416–461, 2000.
- Schäfer, F., Sullivan, T., and Owhadi, H. Compression, inversion, and approximate PCA of dense kernel matrices at near-linear computational complexity. *Multiscale Modeling and Simulation*, 19(2):688–730, 2021.
- Skorski, M. Modern analysis of Hutchinson’s trace estimator. In *55th Annual Conference on Information Sciences and Systems (CISS)*, pp. 1–5. IEEE, 2021.
- Sriperumbudur, B. K. and Szabó, Z. Optimal rates for random Fourier features. *Advances in Neural Information Processing Systems (NeurIPS)*, 2015.
- Titsias, M. Variational learning of inducing variables in sparse Gaussian processes. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, pp. 567–574. PMLR, 2009.
- Trefethen, L. N. and Bau, D. *Numerical Linear Algebra*. Society for Industrial and Applied Mathematics (SIAM), 1997.
- Ubaru, S., Chen, J., and Saad, Y. Fast estimation of  $\text{tr}(f(A))$  via stochastic Lanczos quadrature. *SIAM Journal on Matrix Analysis and Applications*, 38(4):1075–1099, 2017.
- Wang, K. A., Pleiss, G., Gardner, J. R., Tyree, S., Weinberger, K. Q., and Wilson, A. G. Exact Gaussian processes on a million data points. *Advances in Neural Information Processing Systems (NeurIPS)*, 32, 2019.
- Wilson, A. and Nickisch, H. Kernel interpolation for scalable structured Gaussian processes (KISS-GP). In *Proceedings of the 32nd International Conference on Machine Learning (ICML)*, 2015.
- Wilson, A. G., Dann, C., and Nickisch, H. Thoughts on massively scalable Gaussian processes. *arXiv preprint*, 2015. URL <http://arxiv.org/abs/1511.01870>.

This supplementary material is structured as follows. Section S1 contains background on Krylov methods, such as known convergence results. Section S2 contains the main result and proof for variance-reduced stochastic trace estimation. Section S3 gives proofs for the forward and backward pass of the approximation to the log-determinant. In turn, Section S4 contains the error bounds for the log-marginal likelihood and its derivative. Error rates for specific preconditioners are given in Section S5 and finally, additional experimental results can be found in Section S7.

References referring to sections, equations or theorem-type environments within the supplement are prefixed with ‘S’, while references to, or results from, the main paper are stated as is.

<b>S1 Background on Krylov Methods</b>	<b>11</b>
S1.1 Conjugate Gradient Method . . . . .	11
S1.2 Lanczos Algorithm . . . . .	12
S1.3 Stochastic Lanczos Quadrature . . . . .	12
<b>S2 Stochastic Trace Estimation</b>	<b>13</b>
<b>S3 Log-Determinant Estimation</b>	<b>15</b>
S3.1 Approximation of a Matrix Function . . . . .	15
S3.2 Approximation of the Log-Determinant . . . . .	16
S3.3 Approximation of the Derivative of the Log-Determinant . . . . .	17
<b>S4 GP Hyperparameter Optimization</b>	<b>19</b>
S4.1 Approximation of the Log-Marginal Likelihood . . . . .	19
S4.2 Approximation of the Derivative of the Log-Marginal Likelihood . . . . .	20
<b>S5 Preconditioning</b>	<b>22</b>
S5.1 Additive Kernels . . . . .	23
S5.2 Kernels with a Uniformly Converging Approximation . . . . .	23
S5.3 Incomplete Cholesky Decomposition . . . . .	24
S5.4 Quadrature Fourier Features (QFF) . . . . .	24
S5.5 Truncated Singular Value Decomposition . . . . .	25
S5.6 Randomized Singular Value Decomposition . . . . .	26
S5.7 Randomized Nyström Method . . . . .	26
S5.8 Random Fourier Features (RFF) . . . . .	27
<b>S6 Technical Results</b>	<b>27</b>
<b>S7 Additional Experimental Results</b>	<b>27</b>
S7.1 Synthetic Data . . . . .	27
S7.2 UCI Datasets . . . . .	29

## S1. Background on Krylov Methods

### S1.1. Conjugate Gradient Method

**Theorem S6** (Convergence Rate of Preconditioned CG (Trefethen & Bau, 1997))

Let  $\mathbf{A}, \mathbf{P} \in \mathbb{R}^{n \times n}$  be symmetric positive definite. The error of the conjugate gradient method with preconditioner  $\mathbf{P}$  after  $m \in \mathbb{N}$  steps is given by

$$\|\mathbf{x}_k - \mathbf{x}\|_{\mathbf{A}} \leq 2 \left( \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1} \right)^m \|\mathbf{x}_0 - \mathbf{x}\|_{\mathbf{A}} \tag{S26}$$

and in euclidean norm by

$$\|\mathbf{x}_k - \mathbf{x}\|_2 \leq 2\sqrt{\kappa(\mathbf{A})} \left( \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1} \right)^m \|\mathbf{x}_0 - \mathbf{x}\|_2 \tag{S27}$$

where  $\kappa = \kappa(\mathbf{P}^{-\frac{1}{2}} \mathbf{A} \mathbf{P}^{-\frac{1}{2}})$  is the condition number of the preconditioned system matrix.

*Proof.* Preconditioned CG is equivalent to running CG on the transformed problem

$$\tilde{\mathbf{A}} \tilde{\mathbf{x}} = \mathbf{P}^{-\frac{1}{2}} \mathbf{A} \mathbf{P}^{-\frac{1}{2}} \tilde{\mathbf{x}} = \mathbf{P}^{-\frac{1}{2}} \mathbf{b}$$

with the substitution  $\tilde{\mathbf{x}} = \mathbf{P}^{\frac{1}{2}} \mathbf{x}$ . By [Trefethen & Bau \(1997\)](#), the convergence rate of CG on the problem is given by

$$\|\tilde{\mathbf{x}}_m - \tilde{\mathbf{x}}\|_{\tilde{\mathbf{A}}} \leq 2 \left( \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1} \right)^m \|\tilde{\mathbf{x}}_0 - \tilde{\mathbf{x}}\|_{\tilde{\mathbf{A}}}$$

The first equation follows by recognizing that

$$\|\tilde{\mathbf{x}}_m - \tilde{\mathbf{x}}\|_{\tilde{\mathbf{A}}}^2 = (\tilde{\mathbf{x}}_m - \tilde{\mathbf{x}})^\top \mathbf{P}^{-\frac{1}{2}} \mathbf{A} \mathbf{P}^{-\frac{1}{2}} (\tilde{\mathbf{x}}_m - \tilde{\mathbf{x}}) = (\mathbf{x}_m - \mathbf{x})^\top \mathbf{A} (\mathbf{x}_m - \mathbf{x}) = \|\mathbf{x}_m - \mathbf{x}\|_{\mathbf{A}}^2.$$

Now it holds by the min-max principle, that

$$\sqrt{\lambda_{\min}(\tilde{\mathbf{A}})} \|\mathbf{x}_k - \mathbf{x}\|_2 \leq \|\mathbf{x}_m - \mathbf{x}\|_{\mathbf{A}} \leq 2 \left( \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1} \right)^m \|\mathbf{x}_0 - \mathbf{x}\|_{\mathbf{A}} \leq 2 \sqrt{\lambda_{\max}(\mathbf{A})} \left( \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1} \right)^m \|\mathbf{x}_0 - \mathbf{x}\|_2.$$

□

### Corollary S3

Let  $\varepsilon \in (0, 1]$ , then preconditioned CG has relative error  $\|\mathbf{x}_k - \mathbf{x}\|_{\mathbf{A}} \leq \varepsilon \|\mathbf{x}_0 - \mathbf{x}\|_{\mathbf{A}}$  after

$$m \geq \frac{\sqrt{\kappa}}{2} \log(2\varepsilon^{-1}) \tag{S28}$$

iterations, where  $\kappa$  is the condition number of the preconditioned system matrix. In euclidean norm  $\|\cdot\|_2$  relative error  $\varepsilon$  is achieved after

$$m \geq \frac{\sqrt{\kappa}}{2} \log\left(2\sqrt{\kappa(\mathbf{A})}\varepsilon^{-1}\right) \tag{S29}$$

iterations.

*Proof.* It holds by [Lemma S7](#) and the assumption on the number of iterations  $m$ , that

$$2 \left( \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1} \right)^m \leq 2 \exp\left(-\frac{2}{\sqrt{\kappa}} m\right) \leq 2 \exp\left(-\log\left(\frac{2}{\varepsilon}\right)\right) = \varepsilon$$

Using [Theorem S6](#) proves the statement. The proof for the euclidean norm is analogous. □

## S1.2. Lanczos Algorithm

The *Lanczos algorithm* ([Lanczos, 1950](#)) is a Krylov method, which for a symmetric matrix  $\mathbf{A} \in \mathbb{R}^{n \times n}$  iteratively builds an approximate tridiagonalization

$$\mathbf{A} \approx \tilde{\mathbf{Q}} \tilde{\mathbf{T}} \tilde{\mathbf{Q}}^\top$$

where  $\tilde{\mathbf{Q}} \in \mathbb{R}^{n \times m}$  orthonormal and  $\tilde{\mathbf{T}} \in \mathbb{R}^{m \times m}$  tridiagonal. For an initial probe vector  $\mathbf{b} \in \mathbb{R}^n$ , Gram-Schmidt orthogonalization is applied to the Krylov subspace basis. The orthogonalized vectors form  $\tilde{\mathbf{Q}}$ , while the Gram-Schmidt coefficients form  $\tilde{\mathbf{T}}$ . This low-rank approximation becomes an exact tridiagonalization  $\mathbf{A} = \mathbf{Q} \mathbf{T} \mathbf{Q}^\top$  for  $m = n$ . The Lanczos process is often used to compute (approximate) eigenvalues and eigenvectors, which is done by computing an eigendecomposition of the tridiagonal matrix  $\tilde{\mathbf{T}}$  at cost  $\mathcal{O}(m^2)$ . The tridiagonal matrix  $\tilde{\mathbf{T}}$  can also be formed by running CG on the linear system  $\mathbf{A} \mathbf{x} = \mathbf{b}$  and by collecting the step lengths  $\alpha_i$  and conjugacy corrections  $\beta_i$  used in the solution and search direction updates ([Saad, 2003](#), Section 6.7.3).

## S1.3. Stochastic Lanczos Quadrature

One can approximate  $\text{tr}(f(\mathbf{A}))$  for symmetric positive definite  $\mathbf{A}$  via *stochastic Lanczos quadrature* (SLQ) ([Golub & Meurant, 2009](#); [Ubaru et al., 2017](#)) by combining Hutchinson's estimator with quadrature and the Lanczos algorithm. It holds that

$$\text{tr}(f(\mathbf{A})) \approx \tau_\ell^{\text{STE}}(f(\mathbf{A})) = \frac{n}{\ell} \sum_{i=1}^{\ell} \mathbf{z}_i^\top f(\mathbf{A}) \mathbf{z}_i \approx \frac{n}{\ell} \sum_{i=1}^{\ell} I_m^{(i)} = \tau_{\ell, m}^{\text{SLQ}}(f(\mathbf{A}))$$

The quadratic terms  $\mathbf{z}_i^\top f(\mathbf{A}) \mathbf{z}_i$  are approximated by quadrature  $I_m^{(i)}$  where the weights and nodes of the quadrature rule are computed via  $m$  iterations of the Lanczos algorithm. For the log-determinant the following bound for the error incurred by Lanczos quadrature holds.

**Corollary S4** (Section 4.3 of Ubaru et al. (2017))

Let  $\mathbf{A} \in \mathbb{R}^{n \times n}$  be symmetric positive definite with condition number  $\kappa = \kappa(\mathbf{A})$ . Then it holds that

$$|\tau_\ell^{\text{STE}}(\log(\mathbf{A})) - \tau_{\ell,m}^{\text{SLQ}}(\log(\mathbf{A}))| \leq K \left( \frac{\sqrt{2\kappa+1} - 1}{\sqrt{2\kappa+1} + 1} \right)^{2m} \quad (\text{S30})$$

where  $K = \frac{5\kappa \log(2(\kappa+1))}{2\sqrt{2\kappa+1}}$ .

## S2. Stochastic Trace Estimation

**Definition S1** (Convex Concentration Property (Ledoux, 2001))

Let  $\mathbf{x} \in \mathbb{R}^n$  be a random vector. We say  $\mathbf{x}$  has the *convex concentration property* (c.c.p.) with constant  $K \in \mathbb{R}$  if for every 1-Lipschitz convex function  $\phi : \mathbb{R}^n \rightarrow \mathbb{R}$ , we have  $\mathbb{E}[\phi(\mathbf{x})] < \infty$  and for every  $t > 0$ ,

$$\mathbb{P}[|\phi(\mathbf{x}) - \mathbb{E}[\phi(\mathbf{x})]| \geq t] \leq 2 \exp\left(-\frac{t^2}{K^2}\right).$$

Some common examples of random vectors having the c.c.p. are

- random vectors with independent and almost surely bounded entries  $|x_i| \leq 1$ , where  $K = 2\sqrt{2}$  (Samson, 2000);
- Gaussian random vectors  $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \Sigma)$ , where  $K^2 = 2\|\Sigma\|_2$  (Kasiviswanathan & Rudelson, 2019); and
- random vectors which are uniformly distributed on the sphere  $\sqrt{n}S^{n-1}$ , where  $K = 2$  (Kasiviswanathan & Rudelson, 2019).

**Remark S1**

Note, that since Rademacher random vectors have iid entries  $\{+1, -1\}$ , they satisfy  $\|\tilde{z}_i\|_2 = \sqrt{n}$ . In particular, it holds that  $\sqrt{n}z_i = \tilde{z}_i$ . Therefore the random vectors  $\sqrt{n}z_i$  and  $\mathbf{z}' = \sqrt{n}(z_1, \dots, z_\ell)^\top \in \mathbb{R}^{\ell n}$  all have independent entries bounded by 1 and thus satisfy the convex concentration property with  $K = 2\sqrt{2}$ .

**Theorem S7** (Hanson-Wright Inequality for Random Vectors with the Convex Concentration Property (Adamczak, 2015))

Let  $\mathbf{A} \in \mathbb{R}^{n \times n}$  and  $\mathbf{x} \in \mathbb{R}^n$  a zero-mean random vector with the convex concentration property with constant  $K$ . Then for all  $t > 0$ , it holds that

$$\mathbb{P}(|\mathbf{x}^\top \mathbf{A} \mathbf{x} - \mathbb{E}[\mathbf{x}^\top \mathbf{A} \mathbf{x}]| \geq t) \leq 2 \exp\left(-c \min\left(\frac{t^2}{\|\mathbf{A}\|_F^2}, \frac{t}{\|\mathbf{A}\|_2}\right)\right) \quad (\text{S31})$$

where  $c = c(K) > 0$  is a constant only dependent on the distribution of the random vectors.

*Proof.* By Theorem 2.5 of Adamczak (2015) we have

$$\mathbb{P}(|\mathbf{x}^\top \mathbf{A} \mathbf{x} - \mathbb{E}[\mathbf{x}^\top \mathbf{A} \mathbf{x}]| > t) \leq 2 \exp\left(-\frac{1}{c'} \min\left(\frac{t^2}{2K^4\|\mathbf{A}\|_F^2}, \frac{t}{K^2\|\mathbf{A}\|_2}\right)\right)$$

where  $c' > 0$  is a universal constant. Now it holds that

$$\min\left(\frac{t^2}{2K^4\|\mathbf{A}\|_F^2}, \frac{t}{K^2\|\mathbf{A}\|_2}\right) \geq \frac{1}{K^2 \max(2K^2, 1)} \min\left(\frac{t^2}{\|\mathbf{A}\|_F^2}, \frac{t}{\|\mathbf{A}\|_2}\right)$$

Choosing  $c = \frac{1}{c'K^2 \max(2K^2, 1)}$  concludes the proof. □

**Lemma S1**

Let  $\mathbf{A} \in \mathbb{R}^{n \times n}$  and  $\ell \in \mathbb{N}$ . Consider  $\ell$  random vectors  $\tilde{z}_i \in \mathbb{R}^n$  with zero mean and unit covariance, such that for  $\mathbf{z}_i = \tilde{z}_i / \|\tilde{z}_i\|_2$  the stacked random vector  $\sqrt{n}(z_1, \dots, z_\ell)^\top \in \mathbb{R}^{\ell n}$  has the convex concentration property.<sup>10</sup> Then there exists  $c_z > 0$  such that if  $\ell \geq c_z \log(\delta^{-1})$ , then Hutchinson's trace estimator  $\tau_\ell^{\text{STE}}$  satisfies

$$\mathbb{P}(|\tau_\ell^{\text{STE}}(\mathbf{A}) - \text{tr}(\mathbf{A})| \leq \sqrt{c_z \log(\delta^{-1}) \ell^{-1}} \|\mathbf{A}\|_F) \geq 1 - \delta.$$

<sup>10</sup>See Remark S1 for an explanation why this is satisfied for Rademacher random vectors.

*Proof.* Note that the proof strategy used here is the same as in Meyer et al. (2021, Lemma 2) with a different assumption on the distribution of the random vectors. To begin, define

$$\mathbf{A}' = \begin{pmatrix} \mathbf{A} & 0 & \dots & 0 \\ 0 & \mathbf{A} & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & \mathbf{A} \end{pmatrix} \in \mathbb{R}^{\ell n \times \ell n} \quad \text{and} \quad \mathbf{z}' = \sqrt{n} \begin{pmatrix} \mathbf{z}_1 \\ \mathbf{z}_2 \\ \vdots \\ \mathbf{z}_\ell \end{pmatrix} \in \mathbb{R}^{\ell n}.$$

By assumption the random vector  $\mathbf{z}'$  has the convex concentration property and therefore Theorem S7 holds. We obtain

$$\mathbb{P}(|(\mathbf{z}')^\top \mathbf{A}' \mathbf{z}' - \mathbb{E}[(\mathbf{z}')^\top \mathbf{A}' \mathbf{z}']| \geq t) \leq 2 \exp\left(-c \cdot \min\left(\frac{t^2}{\|\mathbf{A}'\|_F^2}, \frac{t}{\|\mathbf{A}'\|_2}\right)\right). \quad (\text{S32})$$

Now, we have  $(\mathbf{z}')^\top \mathbf{A}' \mathbf{z}' = n \sum_{i=1}^{\ell} \mathbf{z}_i^\top \mathbf{A} \mathbf{z}_i = \ell \tau_\ell^{\text{STE}}(\mathbf{A})$  and

$$\begin{aligned} \mathbb{E}[(\mathbf{z}')^\top \mathbf{A}' \mathbf{z}'] &= n \sum_{i=1}^{\ell} \mathbb{E}[\mathbf{z}_i^\top \mathbf{A} \mathbf{z}_i] = n\ell \mathbb{E}[\text{tr}(\mathbf{z}_i^\top \mathbf{A} \mathbf{z}_i)] = n\ell \mathbb{E}[\text{tr}(\mathbf{A} \mathbf{z}_i \mathbf{z}_i^\top)] \\ &= n\ell \text{tr}(\mathbf{A} \mathbb{E}[\mathbf{z}_i \mathbf{z}_i^\top]) = n\ell \text{tr}(\mathbf{A} \text{Cov}(\mathbf{z}_i)) = \ell \text{tr}(\mathbf{A} \text{Cov}(\sqrt{n} \mathbf{z}_i)) = \ell \text{tr}(\mathbf{A}) \end{aligned}$$

Therefore by setting  $t = \sqrt{\frac{\log(2\delta^{-1})}{c}} \ell \|\mathbf{A}\|_F$ , we obtain

$$\begin{aligned} \mathbb{P}\left(\ell |\tau_\ell^{\text{STE}}(\mathbf{A}) - \text{tr}(\mathbf{A})| \geq \sqrt{\frac{\log(2\delta^{-1})}{c}} \ell \|\mathbf{A}\|_F\right) \\ \leq 2 \exp\left(-c \cdot \min\left(\frac{\log(2\delta^{-1})}{c} \frac{\ell \|\mathbf{A}\|_F^2}{\|\mathbf{A}'\|_F^2}, \sqrt{\frac{\log(2\delta^{-1})}{c}} \ell \frac{\|\mathbf{A}\|_F}{\|\mathbf{A}'\|_2}\right)\right) \end{aligned}$$

Further, it holds that  $\|\mathbf{A}'\|_F^2 = \ell \|\mathbf{A}\|_F^2$  and  $\|\mathbf{A}'\|_2 = \|\mathbf{A}\|_2$ , thus we have

$$= 2 \exp\left(-\min\left(\log(2\delta^{-1}), \sqrt{c \log(2\delta^{-1})} \ell \frac{\|\mathbf{A}\|_F}{\|\mathbf{A}\|_2}\right)\right).$$

Now assume  $\ell \geq \frac{1}{c} \log(2\delta^{-1})$ . Then since  $\|\mathbf{A}\|_2 \leq \|\mathbf{A}\|_F$ , the minimum is given by

$$\min\left(\log(2\delta^{-1}), \sqrt{c \log(2\delta^{-1})} \ell \frac{\|\mathbf{A}\|_F}{\|\mathbf{A}\|_2}\right) = \log(2\delta^{-1}).$$

Further setting  $c_z = 2c^{-1}$ , it holds that  $\ell \geq c_z \log(\delta^{-1}) = 2 \log(\delta^{-1}) c^{-1} \geq \log(2\delta^{-1}) c^{-1}$  since  $0 < \delta \leq \frac{1}{2}$ . Combining the above we obtain

$$\mathbb{P}\left(\ell |\tau_\ell^{\text{STE}}(\mathbf{A}) - \text{tr}(\mathbf{A})| \geq \sqrt{c_z \log(\delta^{-1})} \ell \|\mathbf{A}\|_F\right) \leq 2 \exp(-\log(2\delta^{-1})) = \delta,$$

which is equivalent to

$$\mathbb{P}\left(|\tau_\ell^{\text{STE}}(\mathbf{A}) - \text{tr}(\mathbf{A})| \leq \sqrt{c_z \log(\delta^{-1})} \ell^{-1} \|\mathbf{A}\|_F\right) \geq 1 - \delta.$$

This proves the statement. □

**Theorem 1** (Variance-reduced Stochastic Trace Estimation)

Let  $\hat{\mathbf{K}}, \hat{\mathbf{P}}_\ell \in \mathbb{R}_{\text{spd}}^{n \times n}$ ,  $\Delta_f \in \mathbb{R}^{n \times n}$  and  $f : \mathbb{R}_{\text{spd}}^{n \times n} \rightarrow \mathbb{R}$  such that  $\text{tr}(f(\hat{\mathbf{K}})) = \text{tr}(f(\hat{\mathbf{P}}_\ell)) + \text{tr}(\Delta_f)$ , and define the estimator  $\tau_* = \text{tr}(f(\hat{\mathbf{P}}_\ell)) + \tau_\ell^{\text{STE}}(\Delta_f)$ . Now, assume there exist  $c_\Delta > 0$  and  $g : \mathbb{N} \rightarrow (0, \infty)$  such that

$$\|\Delta_f\|_F \leq c_\Delta g(\ell) \|f(\hat{\mathbf{K}})\|_F. \quad (9)$$

Then there exists  $c_z > 0$  dependent on the choice of random vectors, such that, if  $\ell \geq c_z \log(\delta^{-1})$ , it holds with probability  $1 - \delta \in [\frac{1}{2}, 1)$  that

$$|\tau_* - \text{tr}(f(\hat{\mathbf{K}}))| \leq \varepsilon_{\text{STE}} \|f(\hat{\mathbf{K}})\|_F. \quad (10)$$

where for  $C_1 = c_\Delta \sqrt{c_z}$  the relative error is given by

$$\varepsilon_{\text{STE}}(\delta, \ell) = C_1 \sqrt{\log(\delta^{-1})} \ell^{-\frac{1}{2}} g(\ell). \quad (11)$$

*Proof.* By assumption  $|\tau_* - \text{tr}(f(\hat{\mathbf{K}}))| = |\tau_\ell^{\text{STE}}(\Delta_f) - \text{tr}(\Delta_f)|$ . By Lemma S1 it holds with probability  $\geq 1 - \delta$ , that

$$\begin{aligned} |\tau_\ell^{\text{STE}}(\Delta_f) - \text{tr}(\Delta_f)| &\leq \sqrt{c_z \log(\delta^{-1})} \ell^{-1} \|\Delta_f\|_F \\ &\leq \sqrt{c_z \log(\delta^{-1})} c_{\Delta_f} \ell^{-\frac{1}{2}} g(\ell) \|f(\hat{\mathbf{K}})\|_F && \text{Assumption (9).} \\ &= \varepsilon_{\text{STE}}(\delta, \ell) \|f(\hat{\mathbf{K}})\|_F \end{aligned}$$

This concludes the proof. □

### Corollary S5

Let  $\varepsilon \in (0, 1]$  be a desired error. If the conditions of Theorem 1 hold and the number of random vectors  $\ell$  satisfies

$$\ell^{\frac{1}{2}} g(\ell)^{-1} \geq C_1 \varepsilon^{-1} \sqrt{\log(\delta^{-1})}, \quad (S33)$$

then it holds that

$$\boxed{\mathbb{P}(|\tau_* - \text{tr}(\mathbf{A})| \leq \varepsilon \|\mathbf{A}\|_F) \geq 1 - \delta.}$$

*Proof.* Follows from Theorem 1 given (S33). □

## S3. Log-Determinant Estimation

### S3.1. Approximation of a Matrix Function

#### Lemma S2 (Lipschitz Continuity)

Let  $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{n \times n}$  be symmetric. Assume  $f : \Omega \rightarrow \mathbb{R}$  is globally Lipschitz continuous with Lipschitz constant  $L > 0$  on the combined spectrum  $\Omega = \lambda(\mathbf{A}) \cup \lambda(\mathbf{B}) \subset \mathbb{R}$ , then there exists  $c_p > 0$  such that

$$\|f(\mathbf{A}) - f(\mathbf{B})\|_p \leq c_p L \|\mathbf{A} - \mathbf{B}\|_p, \quad (S34)$$

where  $\|\cdot\|_p$  denotes any matrix norm. In particular  $c_2 = 1$  and  $c_F = \sqrt{n}$ .

*Proof.* Since  $\mathbf{A}, \mathbf{B}$  are symmetric, they are normal. By Kittaneh (1985), it holds that

$$\|f(\mathbf{A}) - f(\mathbf{B})\|_2 \leq L \|\mathbf{A} - \mathbf{B}\|_2.$$

The result now follows by equivalence of norms on finite dimensional spaces. For the Frobenius norm we have  $\frac{1}{\sqrt{n}} \|\mathbf{M}\|_F \leq \|\mathbf{M}\|_2 \leq \|\mathbf{M}\|_F$ , and therefore  $c_F = \sqrt{n}$ . □

#### Proposition S1

Let  $\hat{\mathbf{K}} \in \mathbb{R}^{n \times n}$  be symmetric positive definite and assume  $f$  is analytic in a domain containing the spectrum  $\lambda(\hat{\mathbf{K}})$ . Let  $\{\hat{\mathbf{P}}_\ell\}_\ell$  be a sequence of preconditioners with approximation quality (6). Then it holds that

$$\|f(\hat{\mathbf{K}}) - f(\hat{\mathbf{P}}_\ell)\|_F \leq c(n, \hat{\mathbf{K}}, f) g(\ell) \|f(\hat{\mathbf{K}})\|_F \quad (S35)$$

where  $c(n, \hat{\mathbf{K}}, f) = \frac{L \|\hat{\mathbf{K}}\|_F}{c_{f(\lambda)}}$ ,  $L > 0$  is the Lipschitz constant of  $f$  and  $c_{f(\lambda)} = \max\{\min_i |f(\lambda_i(\hat{\mathbf{K}}))|, \frac{\max_i |f(\lambda_i(\hat{\mathbf{K}}))|}{\sqrt{n}}\}$ .

*Proof.* It holds that

$$\|f(\hat{\mathbf{K}})\|_F = \sqrt{\sum_{i=1}^n f(\lambda_i)^2} \geq \begin{cases} \sqrt{n \min_i f(\lambda_i)^2} = \sqrt{n} \min_i |f(\lambda_i)| \\ \|f(\hat{\mathbf{K}})\|_2 = \sigma_{\max}(f(\hat{\mathbf{K}})) = \sqrt{\lambda_{\max}(f(\hat{\mathbf{K}})^2)} = \max_i |f(\lambda_i)| \end{cases}$$

and therefore  $\|f(\hat{\mathbf{K}})\|_F \geq \sqrt{n}c_{f(\lambda)}$ . Since  $f$  is analytic and therefore Lipschitz, it holds that

$$\begin{aligned} \|f(\hat{\mathbf{K}}) - f(\hat{\mathbf{P}}_\ell)\|_F &\leq L\sqrt{n}\|\hat{\mathbf{K}} - \hat{\mathbf{P}}_\ell\|_F && \text{Lemma S2} \\ &\leq L\sqrt{n}g(\ell)\|\hat{\mathbf{K}}\|_F && \text{Preconditioner quality (6)} \\ &= L\sqrt{n}g(\ell)\frac{\|\hat{\mathbf{K}}\|_F}{\|f(\hat{\mathbf{K}})\|_F}\|f(\hat{\mathbf{K}})\|_F \\ &\leq L\sqrt{n}g(\ell)\frac{\|\hat{\mathbf{K}}\|_F}{\sqrt{nc_{f(\lambda)}}}\|f(\hat{\mathbf{K}})\|_F \\ &\leq \frac{L\|\hat{\mathbf{K}}\|_F}{c_{f(\lambda)}}g(\ell)\|f(\hat{\mathbf{K}})\|_F. \end{aligned}$$

This proves the claim.  $\square$

### S3.2. Approximation of the Log-Determinant

**Lemma S3** (Decomposition of the log-determinant)

For  $\hat{\mathbf{K}}, \hat{\mathbf{P}} \in \mathbb{R}^{n \times n}$  symmetric positive definite, it holds that

$$\log \det(\hat{\mathbf{K}}) = \log \det(\hat{\mathbf{P}}) + \text{tr}(\log(\hat{\mathbf{K}}) - \log(\hat{\mathbf{P}})) \quad (\text{S36})$$

$$= \log \det(\hat{\mathbf{P}}) + \text{tr}\left(\log\left(\hat{\mathbf{P}}^{-\frac{1}{2}}\hat{\mathbf{K}}\hat{\mathbf{P}}^{-\frac{1}{2}}\right)\right). \quad (\text{S37})$$

*Proof.* Note that for symmetric positive definite matrices  $\mathbf{A}, \mathbf{B}$ , the matrix logarithm satisfies the following

$$\begin{aligned} \log \det(\mathbf{A}) &= \text{tr}(\log(\mathbf{A})), \\ \text{tr}(\log(\mathbf{AB})) &= \text{tr}(\log(\mathbf{A})) + \text{tr}(\log(\mathbf{B})), \\ \log(\mathbf{A}^{-1}) &= -\log(\mathbf{A}). \end{aligned}$$

Using the above properties, we obtain

$$\begin{aligned} \log \det(\hat{\mathbf{K}}) &= \text{tr}(\log(\hat{\mathbf{P}}\hat{\mathbf{P}}^{-1}\hat{\mathbf{K}})) \\ &= \text{tr}(\log(\hat{\mathbf{P}})) - \text{tr}(\log(\hat{\mathbf{P}})) + \text{tr}(\log(\hat{\mathbf{K}})) \\ &= \log \det(\hat{\mathbf{P}}) + \text{tr}(\log(\hat{\mathbf{K}}) - \log(\hat{\mathbf{P}})) \end{aligned}$$

Now since  $\hat{\mathbf{P}}^{-1}\hat{\mathbf{K}}$  and  $\hat{\mathbf{P}}^{-\frac{1}{2}}\hat{\mathbf{K}}\hat{\mathbf{P}}^{-\frac{1}{2}}$  are similar, they have the same determinant. Therefore we have

$$\text{tr}(\log(\hat{\mathbf{K}}) - \log(\hat{\mathbf{P}})) = \text{tr}(\log(\hat{\mathbf{P}}^{-1}\hat{\mathbf{K}})) = \log \det\left(\hat{\mathbf{P}}^{-\frac{1}{2}}\hat{\mathbf{K}}\hat{\mathbf{P}}^{-\frac{1}{2}}\right).$$

This completes the proof.  $\square$

**Theorem 2** (Error Bound for  $\log \det(\hat{\mathbf{K}})$ )

Let  $f = \log$ ,  $\Delta_{\log} = \log\left(\hat{\mathbf{P}}^{-\frac{1}{2}}\hat{\mathbf{K}}\hat{\mathbf{P}}^{-\frac{1}{2}}\right)$  and assume the conditions of Theorem 1 hold. Then, with probability  $1 - \delta$ , it holds for  $\tau_*^{\log} = \log(\det(\hat{\mathbf{P}})) + \tau_{\ell, m}^{\text{SLQ}}(\Delta_{\log})$ , that

$$\left|\tau_*^{\log} - \log \det(\hat{\mathbf{K}})\right| \leq (\varepsilon_{\text{Lanczos}} + \varepsilon_{\text{STE}})\|\log(\hat{\mathbf{K}})\|_F,$$

where the individual errors are bounded by

$$\varepsilon_{\text{Lanczos}}(\kappa, m) \leq K_1 \left(\frac{\sqrt{2\kappa+1}-1}{\sqrt{2\kappa+1}+1}\right)^{2m} \quad (12)$$

$$\varepsilon_{\text{STE}}(\delta, \ell) \leq C_1 \sqrt{\log(\delta^{-1})}\ell^{-\frac{1}{2}}g(\ell) \quad (13)$$

and  $K_1 = \frac{5\kappa \log(2(\kappa+1))}{2\|\log(\hat{\mathbf{K}})\|_F \sqrt{2\kappa+1}}$ .



*Proof.* Using the decomposition (8), we have

$$\begin{aligned}
 |\tau_*^{\log} - \log \det(\hat{\mathbf{K}})| &= \left| \tau_{\ell, m}^{\text{SLQ}}(\Delta_{\log}) - \text{tr}(\Delta_{\log}) \right| \\
 &\leq \left| \text{tr}(\Delta_{\log}) - \tau_{\ell}^{\text{STE}}(\Delta_{\log}) \right| + \left| \tau_{\ell}^{\text{STE}}(\Delta_{\log}) - \tau_{\ell, m}^{\text{SLQ}}(\Delta_{\log}) \right| \\
 &= \underbrace{\left| \text{tr}(\log(\hat{\mathbf{K}})) - (\text{tr}(\log(\hat{\mathbf{P}}))) + \tau_{\ell}^{\text{STE}}(\Delta_{\log}) \right|}_{e_{\text{STE}}} + \underbrace{\left| \tau_{\ell}^{\text{STE}}(\Delta_{\log}) - \tau_{\ell, m}^{\text{SLQ}}(\Delta_{\log}) \right|}_{e_{\text{Lanczos}}}.
 \end{aligned}$$

Now the individual absolute errors are bounded as follows. By the error bound for stochastic trace estimation in Theorem 1, we have

$$e_{\text{STE}} \leq \varepsilon_{\text{STE}}(\delta, \ell) \|\log(\hat{\mathbf{K}})\|_F = C_1 \sqrt{\log(\delta^{-1})} \ell^{-\frac{1}{2}} g(\ell) \|\log(\hat{\mathbf{K}})\|_F$$

and by Corollary S4, it follows that

$$e_{\text{Lanczos}} \leq K \left( \frac{\sqrt{2\kappa+1}-1}{\sqrt{2\kappa+1}+1} \right)^{2m} = K_1 \left( \frac{\sqrt{2\kappa+1}-1}{\sqrt{2\kappa+1}+1} \right)^{2m} \|\log(\hat{\mathbf{K}})\|_F.$$

□

### Corollary 1

Assume the conditions of Theorem 2 hold. If the number of random vectors  $\ell$  satisfies (11) with  $\varepsilon_{\text{STE}} = \frac{\varepsilon}{2}$  and we run

$$m \geq \frac{\sqrt{3}}{4} \sqrt{\kappa} \log(2K_1 \varepsilon^{-1}) \quad (14)$$

iterations of Lanczos, then it holds that

$$\mathbb{P} \left( \left| \tau_*^{\log} - \log \det(\hat{\mathbf{K}}) \right| \leq \varepsilon \|\log(\hat{\mathbf{K}})\|_F \right) \geq 1 - \delta.$$

*Proof.* By assumption Theorem 1 is satisfied and therefore  $\varepsilon_{\text{STE}} = \frac{\varepsilon}{2}$  with probability  $1 - \delta$ . Now for the error of Lanczos it holds by Theorem 2 in combination with Lemma S7, that

$$\begin{aligned}
 \varepsilon_{\text{Lanczos}} &\leq K_1 \left( \frac{\sqrt{2\kappa+1}-1}{\sqrt{2\kappa+1}+1} \right)^{2m} \\
 &\leq K_1 \exp\left(-\frac{4}{\sqrt{2\kappa+1}} m\right) && \text{Lemma S7} \\
 &\leq K_1 \exp\left(-\frac{\sqrt{3\kappa}}{\sqrt{2\kappa+1}} \log(2K_1 \varepsilon^{-1})\right) && \text{By assumption (14).} \\
 &\leq K_1 \exp(-\log(2K_1 \varepsilon^{-1})) \\
 &= \frac{\varepsilon}{2}
 \end{aligned}$$

The result now follows by Theorem 2. □

### S3.3. Approximation of the Derivative of the Log-Determinant

**Computation of  $\text{tr}(\hat{\mathbf{P}}^{-1} \frac{\partial \hat{\mathbf{P}}}{\partial \theta})$**  Algorithm 1 and Algorithm 2 primarily rely on matrix-vector multiplication, except for computation of  $\tau_{\hat{\mathbf{P}}}^{\text{inv}\partial} = \text{tr}(\hat{\mathbf{P}}^{-1} \frac{\partial \hat{\mathbf{P}}}{\partial \theta})$ . Efficient computation of this term depends on the structure of  $\hat{\mathbf{P}}^{-1}$ . If  $\hat{\mathbf{P}}$  is the pivoted-Cholesky preconditioner, or any other diagonal-plus-low-rank preconditioner  $\sigma^2 \mathbf{I} + \mathbf{L}_{\ell} \mathbf{L}_{\ell}^{\top}$ , we can rewrite this term using the matrix inversion lemma

$$\begin{aligned}
 \text{tr} \left( \hat{\mathbf{P}}^{-1} \frac{\partial \hat{\mathbf{P}}}{\partial \theta} \right) &= \sigma^{-2} \text{tr} \left( \frac{\partial \hat{\mathbf{P}}}{\partial \theta} \right) - \sigma^{-2} \text{tr} \left( \mathbf{L}_{\ell} (\sigma^2 \mathbf{I} + \mathbf{L}_{\ell}^{\top} \mathbf{L}_{\ell})^{-1} \mathbf{L}_{\ell}^{\top} \frac{\partial \hat{\mathbf{P}}}{\partial \theta} \right) \\
 &= \sigma^{-2} \sum_{i=1}^n \frac{\partial \hat{\mathbf{P}}_{ii}}{\partial \theta} - \sigma^{-2} \left( (\mathbf{L}_{\ell} (\sigma^2 \mathbf{I} + \mathbf{L}_{\ell}^{\top} \mathbf{L}_{\ell})^{-1}) \circ \left( \frac{\partial \hat{\mathbf{P}}}{\partial \theta} \mathbf{L}_{\ell} \right) \right) \mathbf{1}, \quad (\text{S38})
 \end{aligned}$$

where  $\circ$  denotes elementwise multiplication. The second term requires  $\ell$  matrix-vector multiplies with  $\frac{\partial \hat{\mathbf{P}}}{\partial \theta}$  and  $\mathcal{O}(n\ell^2)$  additional work. The first term is simply the derivative of the kernel diagonal which will take  $\mathcal{O}(n)$  time. We note that similar efficient procedures exist for other types of preconditioners, such as when  $\hat{\mathbf{P}}^{-1}$  has banded structure.

**Theorem 3** (Error Bound for  $\text{tr}(\hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \theta})$ )

Let  $f(\hat{\mathbf{K}}) = \hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \theta}$ ,  $\Delta_{\text{inv}\partial} = \hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \theta} - \hat{\mathbf{P}}^{-1} \frac{\partial \hat{\mathbf{P}}}{\partial \theta}$  and assume the conditions of Theorem 1 hold. If we solve  $\hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \theta} \mathbf{z}_i$  with  $m$  iterations of preconditioned CG, initialized at  $\mathbf{0}$  or better, then it holds with probability  $1 - \delta$  for  $\tau_*^{\text{inv}\partial} = \text{tr}(\hat{\mathbf{P}}^{-1} \frac{\partial \hat{\mathbf{P}}}{\partial \theta}) + \tau_{\ell, m}^{\text{SCG}}(\Delta_{\text{inv}\partial})$ , that

$$|\tau_*^{\text{inv}\partial} - \text{tr}(\hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \theta})| \leq (\varepsilon_{\text{CG}'} + \varepsilon_{\text{STE}}) \|\hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \theta}\|_F,$$

where the individual errors are bounded by

$$\varepsilon_{\text{CG}'}(\kappa, m) \leq K_2 \left( \frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1} \right)^m \quad (17)$$

$$\varepsilon_{\text{STE}}(\delta, \ell) \leq C_1 \sqrt{\log(\delta^{-1})} \ell^{-\frac{1}{2}} g(\ell) \quad (18)$$

and  $K_2 = 2\sqrt{\kappa(\hat{\mathbf{K}})}n$ .

*Proof.* Using the decomposition (15), we have

$$\begin{aligned} \left| \tau_*^{\text{inv}\partial} - \text{tr} \left( \hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \theta} \right) \right| &= \left| \tau_{\ell, m}^{\text{SCG}}(\Delta_{\text{inv}\partial}) - \text{tr}(\Delta_{\text{inv}\partial}) \right| \\ &\leq \left| \text{tr}(\Delta_{\text{inv}\partial}) - \tau_{\ell}^{\text{STE}}(\Delta_{\text{inv}\partial}) \right| + \left| \tau_{\ell}^{\text{STE}}(\Delta_{\text{inv}\partial}) - \tau_{\ell, m}^{\text{SCG}}(\Delta_{\text{inv}\partial}) \right| \\ &= \underbrace{\left| \text{tr} \left( \hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \theta} \right) - \left( \text{tr} \left( \hat{\mathbf{P}}^{-1} \frac{\partial \hat{\mathbf{P}}}{\partial \theta} \right) + \tau_{\ell}^{\text{STE}}(\Delta_{\text{inv}\partial}) \right) \right|}_{e_{\text{STE}}} + \underbrace{\left| \tau_{\ell}^{\text{STE}}(\Delta_{\text{inv}\partial}) - \tau_{\ell, m}^{\text{SCG}}(\Delta_{\text{inv}\partial}) \right|}_{e_{\text{CG}}}. \end{aligned}$$

Now the individual absolute errors are bounded as follows. By the error bound for stochastic trace estimation in Theorem 1, we have

$$e_{\text{STE}} \leq \varepsilon_{\text{STE}}(\delta, \ell) \left\| \hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \theta} \right\|_F = C_1 \sqrt{\log(\delta^{-1})} \ell^{-\frac{1}{2}} g(\ell) \left\| \hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \theta} \right\|_F.$$

Now, let  $\mathbf{w}_i = \hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \theta} \mathbf{z}_i$ ,  $\tilde{\mathbf{w}}_i = \hat{\mathbf{P}}^{-1} \frac{\partial \hat{\mathbf{P}}}{\partial \theta} \mathbf{z}_i$  and  $\mathbf{w}_{m, i} \approx \mathbf{w}_i$  be the solution computed via preconditioned CG with  $m$  iterations. Then we have by Theorem S6, that

$$\begin{aligned} e_{\text{CG}} &= \left| \frac{n}{\ell} \sum_{i=1}^{\ell} \mathbf{z}_i^{\top} (\mathbf{w}_{m, i} - \tilde{\mathbf{w}}_i - (\mathbf{w}_i - \tilde{\mathbf{w}}_i)) \right| \\ &\leq \frac{n}{\ell} \sum_{i=1}^{\ell} \underbrace{\|\mathbf{z}_i\|_2}_{=1} \|\mathbf{w}_{m, i} - \mathbf{w}_i\|_2 \\ &\leq \frac{n}{\ell} \sum_{i=1}^{\ell} 2\sqrt{\kappa(\hat{\mathbf{K}})} \left( \frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1} \right)^m \|\mathbf{w}_{0, i} - \mathbf{w}_i\|_2 && \text{CG convergence by Theorem S6.} \\ &\leq 2n\sqrt{\kappa(\hat{\mathbf{K}})} \left( \frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1} \right)^m \|\mathbf{w}_i\|_2 && \text{CG initialized at } \mathbf{w}_{0, i} = \mathbf{0} \text{ or better.} \\ &\leq K_2 \left( \frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1} \right)^m \left\| \hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \theta} \mathbf{z}_i \right\|_F \\ &\leq K_2 \left( \frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1} \right)^m \left\| \hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \theta} \right\|_F \end{aligned}$$

This completes the argument.  $\square$

**Corollary 2**

Assume the conditions of Theorem 3 hold. If the number of random vectors  $\ell$  satisfies (11) with  $\varepsilon_{\text{STE}} = \frac{\varepsilon}{2}$ , and we run

$$m \geq \frac{1}{2}\sqrt{\kappa} \log(2K_2\varepsilon^{-1}) \quad (19)$$

iterations of CG, then

$$\mathbb{P}\left(\left|\tau_*^{\text{inv}\partial} - \text{tr}(\hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \theta})\right| \leq \varepsilon \|\hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \theta}\|_F\right) \geq 1 - \delta.$$

*Proof.* By assumption Theorem 1 is satisfied and therefore  $\varepsilon_{\text{STE}} = \frac{\varepsilon}{2}$  with probability  $1 - \delta$ . Now for the error of CG, it holds by Theorem 3 in combination with Lemma S7, that

$$\begin{aligned} \varepsilon_{\text{CG}} &\leq K_2 \left(\frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1}\right)^m \\ &\leq K_2 \exp\left(-\frac{2m}{\sqrt{\kappa}}\right) && \text{Lemma S7} \\ &\leq K_2 \exp(-\log(2K_2\varepsilon^{-1})) && \text{Assumption (19)} \\ &\leq \frac{\varepsilon}{2}. \end{aligned}$$

The result now follows by Theorem 3.  $\square$

## S4. GP Hyperparameter Optimization

### S4.1. Approximation of the Log-Marginal Likelihood

**Theorem 4** (Error Bound for the log-Marginal Likelihood)

Assume the conditions of Theorem 2 hold and we solve  $\hat{\mathbf{K}}\mathbf{u} = \mathbf{y}$  via preconditioned CG initialized at  $\mathbf{u}_0$  and terminated after  $m$  iterations. Then with probability  $1 - \delta$ , the error in the estimate  $\eta = -\frac{1}{2}(\mathbf{y}^\top \mathbf{u}_m + \tau_*^{\log} + n \log(2\pi))$  of the log-marginal likelihood  $\mathcal{L}$  satisfies

$$|\eta - \mathcal{L}| \leq \varepsilon_{\text{CG}} + \frac{1}{2}(\varepsilon_{\text{Lanczos}} + \varepsilon_{\text{STE}}) \|\log(\hat{\mathbf{K}})\|_F,$$

where the individual errors are bounded by

$$\varepsilon_{\text{CG}}(\kappa, m) \leq K_3 \left(\frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1}\right)^m \quad (20)$$

$$\varepsilon_{\text{Lanczos}}(\kappa, m) \leq K_1 \left(\frac{\sqrt{2\kappa+1}-1}{\sqrt{2\kappa+1}+1}\right)^{2m} \quad (21)$$

$$\varepsilon_{\text{STE}}(\delta, \ell) \leq C_1 \sqrt{\log(\delta^{-1})} \ell^{-\frac{1}{2}} g(\ell) \quad (22)$$

for  $K_3 = \sqrt{\kappa(\hat{\mathbf{K}})} \|\mathbf{y}\|_2 \|\mathbf{u}_0 - \mathbf{u}\|_2$ .

*Proof.* It holds by assumption that

$$\begin{aligned} |\eta - \mathcal{L}| &= \frac{1}{2} \left| \mathbf{y}^\top \mathbf{u}_m + \tau_*^{\log} - (\mathbf{y}^\top \hat{\mathbf{K}}^{-1} \mathbf{y} + \log \det(\hat{\mathbf{K}})) \right| \\ &\leq \frac{1}{2} \left( \underbrace{|\mathbf{y}^\top \mathbf{u}_m - \mathbf{y}^\top \mathbf{u}|}_{e_{\text{CG}}} + \underbrace{|\tau_*^{\log} - \log \det(\hat{\mathbf{K}})|}_{e_{\text{SLQ}}} \right). \end{aligned}$$

For the error of CG when solving  $\hat{\mathbf{K}}\mathbf{u} = \mathbf{y}$ , we have by Theorem S6

$$e_{\text{CG}} = |\mathbf{y}^\top \mathbf{u}_m - \mathbf{y}^\top \mathbf{u}| \leq \|\mathbf{y}\|_2 \|\mathbf{u}_m - \mathbf{u}\|_2 \leq \|\mathbf{y}\|_2 2\sqrt{\kappa(\hat{\mathbf{K}})} \left(\frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1}\right)^m \|\mathbf{u}_0 - \mathbf{u}\|_2 = 2\varepsilon_{\text{CG}},$$

and for the absolute error in the log-determinant estimate via preconditioned stochastic Lanczos quadrature, we obtain by Theorem 2, that

$$e_{\text{SLQ}} \leq (\varepsilon_{\text{Lanczos}} + \varepsilon_{\text{STE}}) \|\log(\hat{\mathbf{K}})\|_F.$$

This proves the statement.  $\square$

**Corollary S6**

Assume the conditions of Theorem 4 hold. If the number of random vectors  $\ell$  satisfies (11) with  $\varepsilon_{\text{STE}} = \varepsilon$ , and we run  $m \geq \max(m_{\text{CG}}, m_{\text{Lanczos}})$  iterations of CG and Lanczos, where

$$m_{\text{CG}} \geq \frac{1}{2} \sqrt{\kappa} \log(2K_3 \varepsilon^{-1}), \quad (\text{S39})$$

$$m_{\text{Lanczos}} \geq \frac{\sqrt{3}}{4} \sqrt{\kappa} \log(K_1 \varepsilon^{-1}), \quad (\text{S40})$$

then it holds that

$$\mathbb{P}(|\eta - \mathcal{L}| \leq \varepsilon(1 + \|\log(\hat{\mathbf{K}})\|_F)) \geq 1 - \delta.$$

*Proof.* We begin with the error of CG, it holds by Theorem 4 in combination with Lemma S7, that

$$\begin{aligned} \varepsilon_{\text{CG}} &\leq K_3 \left( \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1} \right)^m \\ &\leq K_3 \exp\left(-\frac{2m}{\sqrt{\kappa}}\right) && \text{Lemma S7} \\ &\leq K_3 \exp(-\log(2K_3 \varepsilon^{-1})) && \text{Assumption (S39)} \\ &= \frac{\varepsilon}{2}. \end{aligned}$$

Now for the error of the estimate of the log-determinant. By assumption Theorem 1 is satisfied and therefore  $\varepsilon_{\text{STE}} = \varepsilon$  with probability  $1 - \delta$ . For the error of Lanczos, it holds by Theorem 4, that

$$\begin{aligned} \varepsilon_{\text{Lanczos}} &\leq K_1 \left( \frac{\sqrt{2\kappa + 1} - 1}{\sqrt{2\kappa + 1} + 1} \right)^{2m} \\ &\leq K_1 \exp\left(-\frac{4}{\sqrt{2\kappa + 1}} m\right) && \text{Lemma S7} \\ &\leq K_1 \exp\left(-\frac{\sqrt{3\kappa}}{\sqrt{2\kappa + 1}} \log(K_1 \varepsilon^{-1})\right) && \text{Assumption (S40).} \\ &\leq K_1 \exp(-\log(K_1 \varepsilon^{-1})) \\ &= \varepsilon \end{aligned}$$

The result now follows by Theorem 4. □

**S4.2. Approximation of the Derivative of the Log-Marginal Likelihood**

**Theorem 5** (Error Bound for the Derivative)

Assume the conditions of Theorem 3 hold and we solve  $\hat{\mathbf{K}}\mathbf{u} = \mathbf{y}$  via preconditioned CG initialized at  $\mathbf{0}$  or better and terminated after  $m$  iterations. Then with probability  $1 - \delta$ , the error in the estimate  $\phi = \frac{1}{2}(\mathbf{u}_m^\top \frac{\partial \hat{\mathbf{K}}}{\partial \theta} \mathbf{u}_m - \tau_*^{\text{inv}\partial})$  of the derivative of the log-marginal likelihood  $\frac{\partial}{\partial \theta} \mathcal{L}$  satisfies

$$\left| \phi - \frac{\partial}{\partial \theta} \mathcal{L} \right| \leq \varepsilon_{\text{CG}}(\kappa, m) + \frac{1}{2}(\varepsilon_{\text{CG}'} + \varepsilon_{\text{STE}}) \|\hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \theta}\|_F,$$

where the individual errors are bounded by

$$\varepsilon_{\text{CG}}(\kappa, m) \leq K_4 \left( \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1} \right)^m \quad (23)$$

$$\varepsilon_{\text{CG}'}(\kappa, m) \leq K_2 \left( \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1} \right)^m \quad (24)$$

$$\varepsilon_{\text{STE}}(\delta, \ell) \leq C_1 \sqrt{\log(\delta^{-1})} \ell^{-\frac{1}{2}} g(\ell) \quad (25)$$

for  $K_4 = 6\kappa(\hat{\mathbf{K}}) \max(\|\mathbf{u}\|_2, \|\mathbf{u}\|_2^3) \|\frac{\partial \hat{\mathbf{K}}}{\partial \theta}\|_2$ .

*Proof.* It holds that

$$\begin{aligned} \left| \phi - \frac{\partial}{\partial \theta} \mathcal{L} \right| &= \frac{1}{2} \left| \mathbf{u}_m^\top \frac{\partial \hat{\mathbf{K}}}{\partial \theta} \mathbf{u}_m - \tau_*^{\text{inv}\partial} - \left( \mathbf{y}^\top \hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \theta} \hat{\mathbf{K}}^{-1} \mathbf{y} - \text{tr} \left( \hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \theta} \right) \right) \right| \\ &\leq \underbrace{\left( \mathbf{u}_m^\top \frac{\partial \hat{\mathbf{K}}}{\partial \theta} \mathbf{u}_m - \mathbf{u}^\top \frac{\partial \hat{\mathbf{K}}}{\partial \theta} \mathbf{u} \right)}_{e_{\text{CG}}} + \underbrace{\left| \tau_*^{\text{inv}\partial} - \text{tr} \left( \hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \theta} \right) \right|}_{e_{\text{CGSTE}}} \end{aligned}$$

Now by Theorem 3, we have

$$e_{\text{CGSTE}} \leq (\varepsilon_{\text{CG}'} + \varepsilon_{\text{STE}}) \left\| \hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \theta} \right\|_F.$$

For the absolute error of the quadratic term, it holds that

$$\begin{aligned} e_{\text{CG}} &= \left| \|\mathbf{u}\|_{\frac{\partial \hat{\mathbf{K}}}{\partial \theta}}^2 - \|\mathbf{u}_m - \mathbf{u} + \mathbf{u}\|_{\frac{\partial \hat{\mathbf{K}}}{\partial \theta}}^2 \right| \\ &\leq \left| \|\mathbf{u}\|_{\frac{\partial \hat{\mathbf{K}}}{\partial \theta}}^2 - (\|\mathbf{u}_m - \mathbf{u}\|_{\frac{\partial \hat{\mathbf{K}}}{\partial \theta}} + \|\mathbf{u}\|_{\frac{\partial \hat{\mathbf{K}}}{\partial \theta}})^2 \right| \\ &= \|\mathbf{u}_m - \mathbf{u}\|_{\frac{\partial \hat{\mathbf{K}}}{\partial \theta}} + 2\|\mathbf{u}_m - \mathbf{u}\|_{\frac{\partial \hat{\mathbf{K}}}{\partial \theta}} \|\mathbf{u}\|_{\frac{\partial \hat{\mathbf{K}}}{\partial \theta}} \\ &\leq \left\| \frac{\partial \hat{\mathbf{K}}}{\partial \theta} \right\|_2 (\|\mathbf{u}_m - \mathbf{u}\|_2^2 + 2\|\mathbf{u}_m - \mathbf{u}\|_2 \|\mathbf{u}\|_2) \\ &\leq \left\| \frac{\partial \hat{\mathbf{K}}}{\partial \theta} \right\|_2 g(\|\mathbf{u}_m - \mathbf{u}\|_2) (1 + 2\|\mathbf{u}\|_2) \end{aligned}$$

for  $g(t) = \max(t, t^2)$ . Now it holds by Theorem S6 and monotonicity of  $g$ , that

$$\leq \left\| \frac{\partial \hat{\mathbf{K}}}{\partial \theta} \right\|_2 g \left( 2\sqrt{\kappa(\hat{\mathbf{K}})} \left( \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1} \right)^m \|\mathbf{u}_0 - \mathbf{u}\|_2 \right) (1 + 2\|\mathbf{u}\|_2)$$

Since for  $a \leq 1$ , it holds that  $g(at) \leq ag(t)$  and for  $a > 1$ :  $g(at) \leq a^2g(t)$ , we have

$$\begin{aligned} &\leq \left\| \frac{\partial \hat{\mathbf{K}}}{\partial \theta} \right\|_2 4\kappa(\hat{\mathbf{K}}) \left( \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1} \right)^m g(\|\mathbf{u}_0 - \mathbf{u}\|_2) (1 + 2\|\mathbf{u}\|_2) \\ &\leq \left\| \frac{\partial \hat{\mathbf{K}}}{\partial \theta} \right\|_2 4\kappa(\hat{\mathbf{K}}) \left( \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1} \right)^m 3 \max(\|\mathbf{u}\|_2, \|\mathbf{u}\|_2^3) \\ &= 2K_4 \left( \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1} \right)^m \end{aligned}$$

where we used that CG was initialized at  $\mathbf{u}_0 = \mathbf{0}$  or better. □

### Corollary S7

Assume the conditions of Theorem 5 hold. If the number of random vectors  $\ell$  satisfies (11) with  $\varepsilon_{\text{STE}} = \varepsilon$ , and we run  $m \geq \max(m_{\text{CG}}, m_{\text{CG}'})$  iterations of CG and Lanczos, where

$$m_{\text{CG}} \geq \frac{1}{2} \sqrt{\kappa} \log(2K_4 \varepsilon^{-1}), \quad (\text{S41})$$

$$m_{\text{CG}'} \geq \frac{1}{2} \sqrt{\kappa} \log(K_2 \varepsilon^{-1}), \quad (\text{S42})$$

then it holds that

$$\mathbb{P} \left( \left| \phi - \frac{\partial}{\partial \theta} \mathcal{L}(\boldsymbol{\theta}) \right| \leq \varepsilon (1 + \|\mathbf{K}^{-1} \frac{\partial \mathbf{K}}{\partial \theta}\|_F) \right) \geq 1 - \delta.$$

*Proof.* We begin with the error of CG's estimate of  $\mathbf{y}^\top \hat{\mathbf{K}}^{-1} \frac{\partial \hat{\mathbf{K}}}{\partial \theta} \hat{\mathbf{K}}^{-1} \mathbf{y}$ , it holds by Theorem 5 in combination with Lemma S7, that

$$\begin{aligned} \varepsilon_{\text{CG}} &\leq K_4 \left( \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1} \right)^m \\ &\leq K_4 \exp\left(-\frac{2m}{\sqrt{\kappa}}\right) && \text{Lemma S7} \\ &\leq K_4 \exp(-\log(2K_4\varepsilon^{-1})) && \text{Assumption (S41)} \\ &= \frac{\varepsilon}{2}. \end{aligned}$$

Now for the error of the stochastic trace estimator. By assumption Theorem 1 is satisfied and therefore  $\varepsilon_{\text{STE}} = \varepsilon$  with probability  $1 - \delta$ . For the error of CG used in the stochastic trace estimate, we obtain by Theorem 5

$$\begin{aligned} \varepsilon_{\text{CG}'} &\leq K_2 \left( \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1} \right)^m \\ &\leq K_2 \exp\left(-\frac{2m}{\sqrt{\kappa}}\right) && \text{Lemma S7} \\ &\leq K_2 \exp(-\log(K_2\varepsilon^{-1})) && \text{Assumption (S42)} \\ &= \varepsilon. \end{aligned}$$

The result now follows by Theorem 5. □

## S5. Preconditioning

**Lemma S4** (Condition Number and Preconditioner Quality)

Let  $\hat{\mathbf{K}}, \hat{\mathbf{P}}_\ell \in \mathbb{R}^{n \times n}$  symmetric positive-definite such that (6) holds and assume that there exists  $c > 0$  such that  $c \geq \max(\lambda_{\min}(\hat{\mathbf{P}}), \lambda_{\min}(\hat{\mathbf{K}}))$  for all  $\ell$ . Then it holds that

$$\kappa = \kappa(\hat{\mathbf{P}}_\ell^{-\frac{1}{2}} \hat{\mathbf{K}} \hat{\mathbf{P}}_\ell^{-\frac{1}{2}}) \leq (1 + \mathcal{O}(g(\ell)) \|\hat{\mathbf{K}}\|_F)^2 \quad (\text{S43})$$

*Proof.* Part of the strategy for this proof is adapted from Gardner et al. (2018). First note, that the matrices  $\hat{\mathbf{P}}^{-1} \hat{\mathbf{K}}, \hat{\mathbf{K}} \hat{\mathbf{P}}^{-1}$  and  $\hat{\mathbf{P}}_\ell^{-\frac{1}{2}} \hat{\mathbf{K}} \hat{\mathbf{P}}_\ell^{-\frac{1}{2}}$  are similar, and thus have the same eigenvalues. Now, we have:

$$\begin{aligned} \kappa(\hat{\mathbf{P}}_\ell^{-\frac{1}{2}} \hat{\mathbf{K}} \hat{\mathbf{P}}_\ell^{-\frac{1}{2}}) &= \frac{\lambda_{\max}(\hat{\mathbf{P}}^{-1} \hat{\mathbf{K}})}{\lambda_{\min}(\hat{\mathbf{K}} \hat{\mathbf{P}}^{-1})} = \|\hat{\mathbf{P}}^{-1} \hat{\mathbf{K}}\|_2 \|\hat{\mathbf{P}} \hat{\mathbf{K}}^{-1}\|_2 \\ &= \|\hat{\mathbf{P}}^{-1}(\hat{\mathbf{K}} - \hat{\mathbf{P}} + \hat{\mathbf{P}})\|_2 \|(\hat{\mathbf{P}} - \hat{\mathbf{K}} + \hat{\mathbf{K}}) \hat{\mathbf{K}}^{-1}\|_2 \\ &= \|1 + \hat{\mathbf{P}}^{-1}(\hat{\mathbf{K}} - \hat{\mathbf{P}})\|_2 \|1 - \hat{\mathbf{K}}^{-1}(\hat{\mathbf{K}} - \hat{\mathbf{P}})\|_2 \end{aligned}$$

Applying Cauchy-Schwarz and the triangle inequality, we obtain:

$$\begin{aligned} &\leq (1 + \|\hat{\mathbf{P}}^{-1}\|_2 \|\hat{\mathbf{K}} - \hat{\mathbf{P}}\|_2) (1 + \|\hat{\mathbf{K}}^{-1}\|_2 \|\hat{\mathbf{K}} - \hat{\mathbf{P}}\|_2) \\ &\leq (1 + c \|\hat{\mathbf{K}} - \hat{\mathbf{P}}\|_F)^2 \\ &\leq (1 + \mathcal{O}(g(\ell)) \|\hat{\mathbf{K}}\|_F)^2 \end{aligned}$$

□

Note since typically  $\hat{\mathbf{P}} = \sigma^2 \mathbf{I} + \mathbf{P}$  with  $\lambda_{\min}(\mathbf{P}) \approx 0$  and  $\lambda_{\min}(\hat{\mathbf{K}}) \leq \sigma^2 + \lambda_{\min}(\mathbf{K})$ , where  $\lambda_{\min}(\mathbf{K})$  small for most kernels, usually  $c \approx \sigma^2$ .

### S5.1. Additive Kernels

#### Lemma S5 (Additive Kernels)

Let  $k(\mathbf{x}, \mathbf{y}) = \sum_{j=1}^d k_j(\mathbf{x}_j, \mathbf{y}_j)$  be an additive kernel and  $\{(\hat{\mathbf{P}}_\ell)_j\}_{j=1}^d$  a set of preconditioners indexed by  $\ell$ , such that for all  $j = 1, \dots, d$ , we have

$$\|\hat{\mathbf{K}}_j - (\hat{\mathbf{P}}_\ell)_j\|_F \leq c_j g(\ell) \|\hat{\mathbf{K}}_j\|_F. \quad (\text{S44})$$

Then it holds for  $\hat{\mathbf{P}}_\ell = \sum_{j=1}^d (\hat{\mathbf{P}}_\ell)_j$  and  $c = \max_j c_j$  that

$$\|\hat{\mathbf{K}} - \hat{\mathbf{P}}_\ell\|_F \leq cdg(\ell) \|\hat{\mathbf{K}}\|_F \quad (\text{S45})$$

*Proof.* It holds by assumption, that

$$\begin{aligned} \|\hat{\mathbf{K}} - \hat{\mathbf{P}}_\ell\|_F &\leq \sum_{j=1}^d \|\hat{\mathbf{K}}_j - (\hat{\mathbf{P}}_\ell)_j\|_F && \text{Cauchy-Schwarz} \\ &\leq \sum_{j=1}^d c_j g(\ell) \|\hat{\mathbf{K}}_j\|_F \\ &\leq \max_j c_j g(\ell) \sum_{j=1}^d \sqrt{\sum_{i=1}^n \lambda_i(\hat{\mathbf{K}}_j)^2} \\ &\leq cg(\ell) \sum_{j=1}^d \sqrt{\sum_{i=1}^n \lambda_i(\hat{\mathbf{K}})^2} && \mathbf{A}, \mathbf{B} \text{ spd} \implies \lambda_i(\mathbf{A}) \leq \lambda_i(\mathbf{A} + \mathbf{B}) \\ &\leq cg(\ell)d \|\hat{\mathbf{K}}\|_F \end{aligned}$$

□

### S5.2. Kernels with a Uniformly Converging Approximation

#### Lemma S6 (Preconditioner Quality from Uniform Convergence)

Let  $\Omega \subset \mathbb{R}^d$  be the data domain and  $k(\cdot, \cdot)$  a positive-definite kernel such that for all  $\mathbf{x} \in \Omega$  it holds that  $k(\mathbf{x}, \mathbf{x}) \leq \sigma^2$ . Let  $P_\ell(\cdot, \cdot)$  be a kernel approximation, such that a uniform convergence bound of the form

$$\sup_{\mathbf{x}, \mathbf{y} \in \Omega} |k(\mathbf{x}, \mathbf{y}) - P_\ell(\mathbf{x}, \mathbf{y})| \leq g(\ell) c_{\text{unif}}(d, \Omega, k) \quad (\text{S46})$$

holds. Then it holds for the preconditioner  $\hat{\mathbf{P}}_\ell = \sigma^2 \mathbf{I} + P_\ell(\mathbf{X}, \mathbf{X})$ , that

$$\|\hat{\mathbf{K}} - \hat{\mathbf{P}}_\ell\|_F \leq g(\ell) c(d, \Omega, k, n) \|\hat{\mathbf{K}}\|_F, \quad (\text{S47})$$

where  $c(d, \Omega, k, n) = \frac{\sqrt{n}}{\sigma^2} c_{\text{unif}}(d, \Omega, k)$ .

*Proof.* It holds that

$$\begin{aligned}
 \|\hat{\mathbf{K}} - \hat{\mathbf{P}}_\ell\|_F &= \|\mathbf{K} - \mathbf{P}_\ell\|_F = \sqrt{\sum_{i,j=1}^n (k(\mathbf{x}_i, \mathbf{x}_j) - P_\ell(\mathbf{x}_i, \mathbf{x}_j))^2} \\
 &\leq \sqrt{\sum_{i,j=1}^n (g(\ell)c_{\text{unif}}(d, \Omega, k))^2} && \text{uniform convergence bound (S46)} \\
 &= ng(\ell)c_{\text{unif}}(d, \Omega, k) \\
 &= g(\ell)\sqrt{n} \left( \sum_{i=1}^n \frac{k(\mathbf{x}_i, \mathbf{x}_i)^2}{k(\mathbf{x}_i, \mathbf{x}_i)^2} \right)^{\frac{1}{2}} && k \text{ bounded} \\
 &\leq g(\ell) \frac{\sqrt{n}}{o^2} \left( \sum_{i=1}^n \mathbf{K}_{ij}^2 \right)^{\frac{1}{2}} \\
 &= g(\ell)c(d, \Omega, k, n)\|\mathbf{K}\|_F \\
 &\leq g(\ell)c(d, \Omega, k, n)\|\hat{\mathbf{K}}\|_F
 \end{aligned}$$

□

### S5.3. Incomplete Cholesky Decomposition

#### Proposition S2 (Incomplete Cholesky Approximation Quality)

Let  $k(\mathbf{x}, \mathbf{y}) = k(\|\mathbf{x} - \mathbf{y}\|)$  be a stationary kernel with output scale  $o^2 = k(0) > 0$ . Assume the kernel matrix spectrum decays at least exponentially, i.e.  $\lambda_i(\mathbf{K}) \leq c \exp(-bi)$  for  $c > 0$  and  $b > \log(4)$ . Then the incomplete Cholesky preconditioner  $\hat{\mathbf{P}}_\ell = \sigma^2 \mathbf{I} + \mathbf{L}_\ell \mathbf{L}_\ell^\top$  satisfies

$$\|\hat{\mathbf{K}} - \hat{\mathbf{P}}_\ell\|_F \leq \frac{c\sqrt{n}}{o^2} \exp(-b'\ell) \|\mathbf{K}\|_F \quad (\text{S48})$$

where  $b' = b - \log(4) > 0$ .

*Proof.* First note that since  $k(\cdot, \cdot)$  is a positive definite kernel, the choice  $o^2 = k(0)$  is no restriction. Now, it holds that

$$\begin{aligned}
 \lambda_i &\leq c \exp(-bi) = c \exp(-(b' + \log(4))i) = c \exp(-b'i) 4^{-i} \\
 \iff 4^i \lambda_i &\leq c \exp(-b'i).
 \end{aligned}$$

Therefore by Theorem 3.2 of Harbrecht et al. (2012), we have  $\text{tr}(\mathbf{K} - \mathbf{L}_\ell \mathbf{L}_\ell^\top) \leq cn \exp(-b'\ell)$ . Now it holds since  $\mathbf{K} - \mathbf{L}_\ell \mathbf{L}_\ell^\top$  positive definite, that

$$\|\hat{\mathbf{K}} - \hat{\mathbf{P}}_\ell\|_F = \|\mathbf{K} - \mathbf{L}_\ell \mathbf{L}_\ell^\top\|_F \leq \text{tr}(\mathbf{K} - \mathbf{L}_\ell \mathbf{L}_\ell^\top) \leq cn \exp(-b'\ell),$$

and with  $\mathbf{K}_{ii} = o^2$  that

$$n = \sqrt{n} \left( \sum_{i=1}^n \frac{\mathbf{K}_{ii}^2}{o^4} \right)^{\frac{1}{2}} \leq \frac{\sqrt{n}}{o^2} \left( \sum_{i,j=1}^n \mathbf{K}_{ij}^2 \right)^{\frac{1}{2}} = \frac{\sqrt{n}}{o^2} \|\mathbf{K}\|_F.$$

This concludes the argument. □

### S5.4. Quadrature Fourier Features (QFF)

#### Proposition S3 (QFF Approximation Quality)

Assume  $\Omega = [0, 1]^d$ ,  $k$  a kernel with Fourier transform  $p(\omega) = \exp\left(-\frac{1}{2} \sum_{j=1}^d \omega_j^2 \gamma_j^2\right)$  such that Assumption 1 of Mutn̄y & Krause (2018) is satisfied and let  $\hat{\mathbf{P}}_\ell = \sigma^2 \mathbf{I} + \mathbf{P}_\ell$ , where  $\mathbf{P}_\ell$  is the QFF approximated kernel matrix. Let  $\ell^{\frac{1}{d}} > \frac{2}{\gamma^2}$ , then for  $b = \frac{1}{2}(\log(4) - 1)$ , it holds that

$$\|\hat{\mathbf{K}} - \hat{\mathbf{P}}_\ell\|_F \leq c(d, n, k) \exp\left(-b\ell^{\frac{1}{d}}\right) \|\hat{\mathbf{K}}\|_F. \quad (\text{S49})$$



*Proof.* By Theorem 1 of [Mutn̄y & Krause \(2018\)](#), replacing  $\ell = (2\bar{m})^d$  it holds that

$$\begin{aligned}
 \sup_{\mathbf{x}, \mathbf{y} \in \Omega} |k(\mathbf{x}, \mathbf{y}) - P_\ell(\mathbf{x}, \mathbf{y})| &\leq d2^{d-1} \sqrt{\frac{\pi}{2}} \left( \frac{e}{2\gamma^2 \ell^{\frac{1}{d}}} \right)^{\frac{1}{2} \ell^{\frac{1}{d}}} \\
 &\leq c(d) \left( \frac{e}{4} \right)^{\frac{1}{2} \ell^{\frac{1}{d}}} && \ell^{\frac{1}{d}} > 2\gamma^{-2} \\
 &= c(d) \exp\left( \frac{1}{2} \ell^{\frac{1}{d}} - \log(4) \frac{1}{2} \ell^{\frac{1}{d}} \right) \\
 &= c(d) \exp\left( -\frac{1}{2} (\log(4) - 1) \ell^{\frac{1}{d}} \right) \\
 &= c(d) \exp\left( -b \ell^{\frac{1}{d}} \right)
 \end{aligned}$$

Now by Lemma S6, we have

$$\|\hat{\mathbf{K}} - \hat{\mathbf{P}}_\ell\|_F \leq c(d, n) \exp\left(-b \ell^{\frac{1}{d}}\right) \|\hat{\mathbf{K}}\|_F, \quad (\text{S50})$$

where  $c(d, n) = \sqrt{nc}(d)$  by Assumption 1 of [Mutn̄y & Krause \(2018\)](#), which assumes  $k(\mathbf{x}, \mathbf{y}) \leq 1$ . □

**Proposition S4** (General QFF Approximation Quality)

Assume  $\Omega = [0, 1]^d$ ,  $k$  a kernel with Fourier transform  $p(\omega)$  such that Assumption 1 of [Mutn̄y & Krause \(2018\)](#) is satisfied and  $f_\delta(\phi) = p(\cot(\phi)) \frac{\cos(\delta \cot(\phi))}{\sin(\phi)^2}$  is  $(s-1)$ -times absolutely continuous. Let  $\hat{\mathbf{P}}_\ell = \sigma^2 \mathbf{I} + \mathbf{P}_\ell$ , where  $\mathbf{P}_\ell$  is the QFF approximated kernel matrix. Then it holds that

$$\|\hat{\mathbf{K}} - \hat{\mathbf{P}}_\ell\|_F \leq c(d, \Omega, k, n) \ell^{-\frac{s+1}{d}} \|\hat{\mathbf{K}}\|_F. \quad (\text{S51})$$

*Proof.* By Theorem 4 of [Mutn̄y & Krause \(2018\)](#), replacing  $\ell = (2\bar{m})^d$ , it holds that

$$\begin{aligned}
 \sup_{\mathbf{x}, \mathbf{y} \in \Omega} |k(\mathbf{x}, \mathbf{y}) - P_\ell(\mathbf{x}, \mathbf{y})| &\leq d2^{d-1} \frac{(s+2)^{s+1}}{s!} \max_{\delta \in \Omega} \text{TV}(f_\delta^{(s)}) 2^{s+1} \ell^{-\frac{s+1}{d}} \\
 &= c(d, \Omega, k) \ell^{-\frac{s+1}{d}}
 \end{aligned}$$

Now by Lemma S6, we have

$$\|\hat{\mathbf{K}} - \hat{\mathbf{P}}_\ell\|_F \leq c(d, \Omega, k, n) \ell^{-\frac{s+1}{d}} \|\hat{\mathbf{K}}\|_F. \quad \square$$

**Remark S2** (Modified Matérn Kernel)

Proposition S4 is satisfied for example for the modified Matérn( $\nu$ ) kernel, defined via its spectral density

$$p(\omega) = \prod_{j=1}^d \frac{1}{(1 + \gamma_j^2 \omega_j^2)^{\nu + \frac{1}{2}}}. \quad (\text{S52})$$

See the appendix of [Mutn̄y & Krause \(2018\)](#) for a detailed definition and motivation.

**S5.5. Truncated Singular Value Decomposition**

**Proposition S5** (Truncated SVD Approximation Quality)

Let  $\mathbf{K}$  be a kernel matrix and  $\mathbf{P}_\ell = \mathbf{V}_\ell \mathbf{\Lambda}_\ell \mathbf{V}_\ell^\top$  its truncated singular value decomposition consisting of the eigenvectors of the largest  $\ell$  eigenvalues. Then it holds for  $\hat{\mathbf{P}} = \sigma^2 \mathbf{I} + \mathbf{P}_\ell$ , that

$$\|\hat{\mathbf{K}} - \hat{\mathbf{P}}_\ell\|_F \leq c(n) \ell^{-\frac{1}{2}} \|\hat{\mathbf{K}}\|_F. \quad (\text{S53})$$

where  $c(n) = \sqrt{n}$ .

*Proof.* Since the optimal rank- $\ell$  approximation in Frobenius norm is given by the truncated SVD (Eckart & Young, 1936), we have for  $\lambda_1(\mathbf{K}) \geq \dots \geq \lambda_n(\mathbf{K})$ , that

$$\|\hat{\mathbf{K}} - \hat{\mathbf{P}}_\ell\|_F^2 = \|\mathbf{K} - \mathbf{P}_\ell\|_F^2 = \sum_{i=\ell+1}^n \lambda_i(\mathbf{K})^2$$

Now since  $\lambda_{\ell+1}(\mathbf{K}) \leq \frac{1}{\ell} \sum_{i=1}^{\ell} \lambda_i(\mathbf{K}) \leq \frac{1}{\ell} \text{tr}(\mathbf{K})$ , we have

$$\leq \lambda_{\ell+1}(\mathbf{K}) \sum_{i=\ell+1}^n \lambda_i(\mathbf{K}) \leq \frac{1}{\ell} \text{tr}(\mathbf{K})^2 = \frac{1}{\ell} \|\boldsymbol{\lambda}\|_1^2 \leq \frac{n}{\ell} \|\boldsymbol{\lambda}\|_2^2 = \frac{n}{\ell} \|\mathbf{K}\|_F^2 \leq \frac{n}{\ell} \|\hat{\mathbf{K}}\|_F^2$$

where  $\boldsymbol{\lambda} = (\lambda_1(\mathbf{K}), \dots, \lambda_n(\mathbf{K}))^\top \in \mathbb{R}^n$ . □

### S5.6. Randomized Singular Value Decomposition

**Proposition S6** (Randomized SVD Approximation Quality)

Let  $\mathbf{K}$  be a kernel matrix and  $\mathbf{P}_\ell = \mathbf{H}_\ell \mathbf{H}_\ell^\top \mathbf{K}$  its randomized singular value decomposition constructed via the LINEARTIMESVD algorithm (Drineas et al., 2006) with  $s \in \mathbb{N}$  samples drawn according to probabilities  $\{p_i\}_{i=1}^n$ . Then for  $\hat{\mathbf{P}} = \sigma^2 \mathbf{I} + \mathbf{P}_\ell$ , it holds with probability  $1 - \delta$ , that

$$\|\hat{\mathbf{K}} - \hat{\mathbf{P}}_\ell\|_F \leq c(n)(\ell^{-\frac{1}{2}} + \mathcal{O}(\ell^{\frac{1}{4}} s^{-\frac{1}{4}})) \|\hat{\mathbf{K}}\|_F, \quad (\text{S54})$$

where  $c(n) = \sqrt{n}$ .

*Proof.* By Drineas et al. (2006, Theorem 4) it holds with probability  $1 - \delta$ , that

$$\|\hat{\mathbf{K}} - \hat{\mathbf{P}}_\ell\|_F^2 = \|\mathbf{K} - \mathbf{P}_\ell\|_F^2 \leq \|\mathbf{K} - \mathbf{K}_\ell\|_F^2 + c(p_i, \delta) \ell^{\frac{1}{2}} s^{-\frac{1}{2}} \|\mathbf{K}\|_F^2$$

By the same argument as in the proof of Proposition S5 for the error of the optimal rank- $\ell$  approximation  $\|\mathbf{K} - \mathbf{K}_\ell\|_F^2$ , we obtain

$$\begin{aligned} &\leq \frac{n}{\ell} \|\mathbf{K}\|_F^2 + c(p_i, \delta) \ell^{\frac{1}{2}} s^{-\frac{1}{2}} \|\mathbf{K}\|_F^2 \\ &\leq n(\ell^{-1} + \mathcal{O}(\ell^{\frac{1}{2}} s^{-\frac{1}{2}})) \|\hat{\mathbf{K}}\|_F^2 \end{aligned}$$

This completes the proof. □

### S5.7. Randomized Nyström Method

**Proposition S7** (Randomized Nyström Approximation Quality)

Let  $\mathbf{K}$  be a kernel matrix and  $\mathbf{P}_\ell = \mathbf{C} \mathbf{W}_\ell^\dagger \mathbf{C}^\top$  its randomized Nyström approximation constructed via Algorithm 3 of Drineas et al. (2005) with  $s \in \mathbb{N}$  columns drawn according to probabilities  $p_i = \frac{\mathbf{K}_{ii}^2}{\sum_{i=1}^n \mathbf{K}_{ii}^2}$ . Then for  $\hat{\mathbf{P}} = \sigma^2 \mathbf{I} + \mathbf{P}_\ell$ , it holds with probability  $1 - \delta$ , that

$$\|\hat{\mathbf{K}} - \hat{\mathbf{P}}_\ell\|_F \leq c(n)(\ell^{-\frac{1}{2}} + \mathcal{O}(\ell^{\frac{1}{4}} s^{-\frac{1}{4}})) \|\hat{\mathbf{K}}\|_F, \quad (\text{S55})$$

where  $c(n) = \sqrt{n}$ .

*Proof.* By Drineas et al. (2005, Theorem 3) it holds with probability  $1 - \delta$ , that

$$\|\hat{\mathbf{K}} - \hat{\mathbf{P}}_\ell\|_F = \|\mathbf{K} - \mathbf{P}_\ell\|_F \leq \|\mathbf{K} - \mathbf{K}_\ell\|_F + c(\delta) \ell^{\frac{1}{4}} s^{-\frac{1}{4}} \sum_{i=1}^n \mathbf{K}_{ii}^2$$

By the same argument as in the proof of Proposition S5 for the error of the optimal rank- $\ell$  approximation  $\|\mathbf{K} - \mathbf{K}_\ell\|_F$ , we obtain

$$\begin{aligned} &\leq n^{\frac{1}{2}}\ell^{-\frac{1}{2}}\|\mathbf{K}\|_F^2 + c(\delta)\ell^{\frac{1}{4}}s^{-\frac{1}{4}}\sum_{i=1}^n \mathbf{K}_{ii}^2 \\ &\leq n^{\frac{1}{2}}\ell^{-\frac{1}{2}}\|\mathbf{K}\|_F^2 + c(\delta)\ell^{\frac{1}{4}}s^{-\frac{1}{4}}\sqrt{\sum_{i=1}^n \mathbf{K}_{ii}^2}\|\mathbf{K}\|_F \\ &\leq n^{\frac{1}{2}}(\ell^{-\frac{1}{2}} + \mathcal{O}(\ell^{\frac{1}{4}}s^{-\frac{1}{4}}))\|\hat{\mathbf{K}}\|_F \end{aligned}$$

This completes the proof. □

### S5.8. Random Fourier Features (RFF)

#### Proposition S8 (RFF Approximation Quality)

Let  $k(\mathbf{x}, \mathbf{y}) = k(\mathbf{x} - \mathbf{y})$  be a positive-definite kernel with compact data domain  $\Omega \subset \mathbb{R}^d$ . Let  $\mathbf{P}_\ell = \mathbf{Z}_\ell \mathbf{Z}_\ell^\top$  be the random Fourier feature approximation (Rahimi et al., 2007), where  $\mathbf{Z}_\ell \in \mathbb{R}^{n \times \ell}$ . Then for  $\hat{\mathbf{P}}_\ell = \sigma^2 \mathbf{I} + \mathbf{P}_\ell$  it holds with probability  $1 - \delta$ , that

$$\|\hat{\mathbf{K}} - \hat{\mathbf{P}}_\ell\|_F \leq c(d, \Omega, k, \delta, n)\ell^{-\frac{1}{2}}\|\hat{\mathbf{K}}\|_F. \quad (\text{S56})$$

*Proof.* By Theorem 1 of Sriperumbudur & Szabó (2015) with probability  $1 - \delta$ , we obtain the uniform convergence bound

$$\sup_{\mathbf{x}, \mathbf{y} \in \Omega} |k(\mathbf{x}, \mathbf{y}) - P_\ell(\mathbf{x}, \mathbf{y})| \leq c_{\text{unif}}(d, \Omega, k, \delta)\ell^{-\frac{1}{2}}. \quad (\text{S57})$$

Now, applying Lemma S6 completes the proof. □

## S6. Technical Results

### Lemma S7

Let  $x \in \mathbb{R}$  such that  $x > 1$ , then it holds that

$$\frac{x-1}{x+1} \leq \exp\left(-\frac{2}{x}\right). \quad (\text{S58})$$

*Proof.* By Mitrinovic & Vasic (1970, Section 3.6.18) it holds for  $y > 0$ , that

$$\frac{\log(y+1)}{y} \geq \frac{2}{2+y}$$

Substituting  $y = \frac{x+1}{x-1} - 1$  we obtain

$$\begin{aligned} \log\left(\frac{x+1}{x-1}\right) &\geq \frac{2\left(\frac{x+1}{x-1} - 1\right)}{1 + \frac{x+1}{x-1}} = \frac{2(x+1-x+1)}{2x} = \frac{2}{x} \\ &\iff \frac{x+1}{x-1} \geq \exp\left(\frac{2}{x}\right) \end{aligned}$$

This proves the claim. □

## S7. Additional Experimental Results

### S7.1. Synthetic Data

We report bias and variance of the stochastic estimators for the log-marginal likelihood and its derivatives for the exponentiated quadratic, Matérn( $\frac{3}{2}$ ) and rational quadratic kernel on a synthetic dataset of size  $n = 10,000$  with varying dimensionality  $d \in \{1, 2, 3\}$  in Table S1. Using  $\ell = 128$  random samples with a preconditioner of the same size bias and variance are reduced by several orders of magnitude across different kernels. Note, that the variance reduction tends to decline with dimensionality, even though this is not necessarily universal across kernels. We show the bias and variance reduction for increasing number of random samples, respectively preconditioner size in Figure S1.

Preconditioning for Scalable GP Hyperparameter Optimization

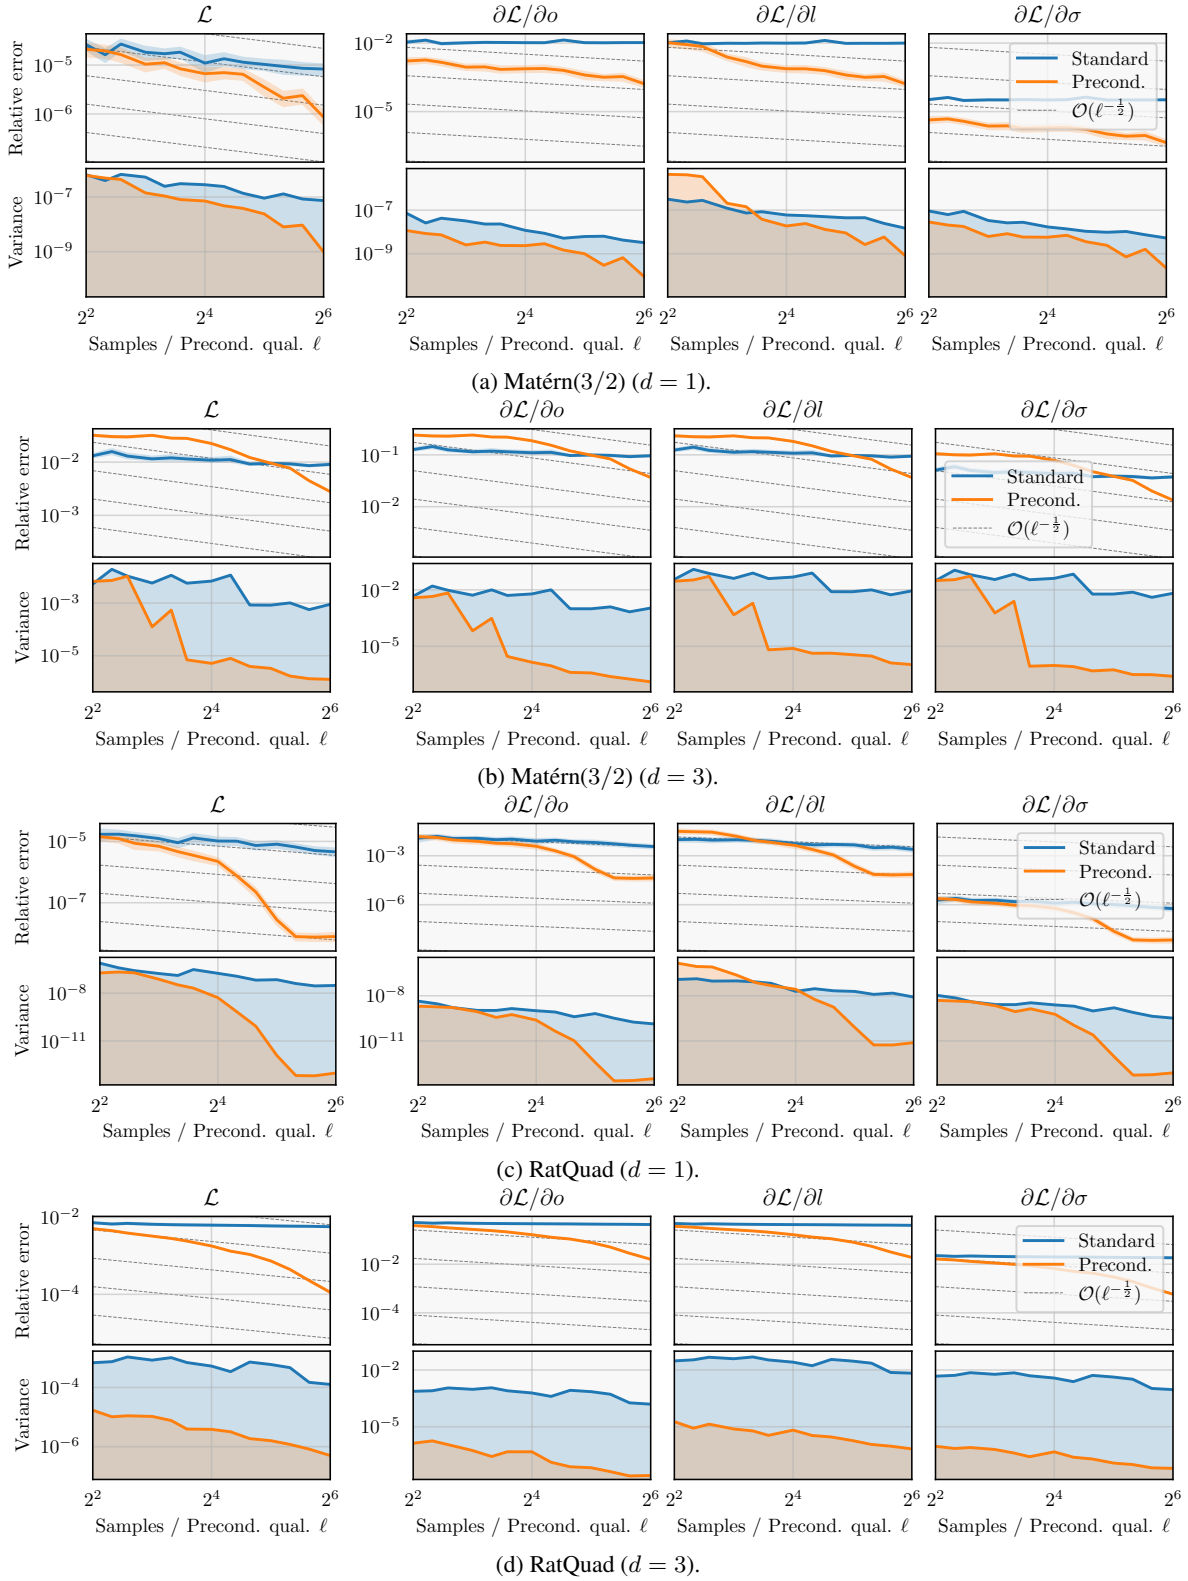


Figure S1: *Bias and variance decrease on synthetic datasets for different kernels.* Relative error and variance of the stochastic estimators of the log-marginal likelihood and its derivative for increasing number of random vectors, equivalently preconditioner size. Experiments were performed for different kernels on a synthetic dataset of size  $n = 10,000$  with dimension  $d \in \{1, 2, 3\}$ . Plots show mean and 95% confidence intervals for the relative error computed over 25 repetitions.

**Preconditioning for Scalable GP Hyperparameter Optimization**

Table S1: *Bias and variance reduction for different kernels.* Bias and variance of the stochastic estimators for the log-marginal likelihood and its derivative(s) computed for synthetic data ( $n = 10,000, \sigma^2 = 10^{-2}$ ) with 25 repetitions.

Kernel	$d$	Prec. Qual.	$\mathcal{L}$		$\partial\mathcal{L}/\partial o$		$\partial\mathcal{L}/\partial l$		$\partial\mathcal{L}/\partial\sigma$	
			Bias	Var.	Bias	Var.	Bias	Var.	Bias	Var.
Matérn(3/2)	1	0	3e-4	3e-8	9e-4	1e-9	2e-3	8e-9	2e-3	2e-9
		128	9e-6	4e-11	4e-6	8e-12	1e-5	7e-11	7e-6	2e-11
	2	0	3e-1	3e-3	4e-1	4e-3	1	3e-2	9e-1	3e-2
		128	3e-4	5e-8	7e-5	3e-9	2e-4	3e-8	1e-4	7e-9
	3	0	3e-1	6e-4	4e-1	7e-4	1	6e-3	8e-1	4e-3
		128	7e-3	7e-7	2e-2	3e-8	5e-2	4e-7	3e-2	8e-8
RBF	1	0	1e-4	1e-8	1e-5	4e-11	1e-4	3e-9	2e-5	1e-10
		128	5e-8	1e-15	3e-8	4e-16	7e-7	2e-13	4e-8	1e-15
	2	0	3e-3	7e-8	5e-3	2e-9	8e-2	1e-7	1e-2	2e-9
		128	1e-6	1e-12	5e-7	2e-13	1e-5	1e-10	8e-7	5e-13
	3	0	1e-1	2e-7	2e-1	4e-8	2	9e-7	4e-1	8e-9
		128	3e-4	5e-8	5e-5	2e-9	7e-4	3e-7	8e-5	4e-9
RatQuad	1	0	1e-4	1e-8	1e-5	1e-10	1e-4	5e-9	2e-5	3e-10
		128	3e-7	6e-14	2e-7	2e-14	2e-6	4e-12	2e-7	4e-14
	2	0	3e-2	2e-7	4e-2	3e-8	3e-1	2e-7	1e-1	7e-9
		128	8e-5	3e-9	2e-5	2e-10	2e-4	1e-8	3e-5	5e-10
	3	0	2e-1	2e-4	3e-1	3e-4	2	1e-2	8e-1	2e-3
		128	4e-4	1e-7	2e-4	1e-8	4e-3	4e-7	4e-4	4e-8

Table S2: *Hyperparameter optimization using Adam.* GP regression using a Matérn( $\frac{3}{2}$ ) kernel and incomplete Cholesky preconditioner of size 500 with  $\ell = 50$  random samples. Hyperparameters were optimized with Adam for at most 20 steps using early stopping via a validation set.

Dataset	$n$	$d$	Prec. Size	$-\mathcal{L}_{\text{train}} \downarrow$	$-\mathcal{L}_{\text{test}} \downarrow$	RMSE $\downarrow$	Runtime (s)
Elevators	12,449	18	500	0.4803	0.4593	0.3684	109
Bike	13,034	17	500	0.2265	0.3473	0.2300	64
Kin40k	30,000	8	500	0.4392	-0.1200	0.0982	159
Protein	34,297	9	500	0.9438	0.9319	0.5681	92
KEGGdir	36,620	20	500	-1.0070	-1.0390	0.0810	239

**S7.2. UCI Datasets**

For the experiments we conducted on UCI datasets, we report the full experimental results with their deviation across 10 runs in Table S3. Test errors with and without preconditioning did not differ by more than two standard deviations. However, model evaluations of the optimizer were significantly reduced when using a preconditioner of size 500, leading to substantial speedup. Note, that the experiment on the “3DRoad” dataset was only carried out once due to the prohibitive runtime without preconditioning.

We used the L-BFGS optimizer in our experiments due to its favorable convergence properties. As an ablation experiment we compared to the Adam optimizer as sometimes used for its robustness to noise, when using stochastic approximations of the log-marginal likelihood (Gardner et al., 2018; Wang et al., 2019). We find that with preconditioning optimization with L-BFGS significantly outperformed optimization with Adam, both in terms of training and test error, except for the “KEGGdir” dataset (cf. Table 2 and Table S2). Additionally, L-BFGS converged faster across all experiments. This shows that variance reduction via preconditioning makes the use of second-order optimizers not only possible, but preferred for GP hyperparameter optimization when using stochastic approximations.

Table S3: *Hyperparameter optimization on UCI datasets*. GP regression using a Matérn( $\frac{3}{2}$ ) kernel and incomplete Cholesky preconditioner of size 500 with  $\ell = 50$  random samples. Hyperparameters were optimized with L-BFGS for at most 20 steps using early stopping via a validation set. All results, but “3DRoad”, are averaged over 10 runs.

Dataset	$n$	$d$	Prec.	Qual.	Opt. Steps	Model evals.		Runtime (s)		Speedup		$-\mathcal{L}_{\text{train}} \downarrow$		$-\mathcal{L}_{\text{test}} \downarrow$		RMSE $\downarrow$	
						mean	std	mean	std	mean	std	mean	std	mean	std	mean	std
Elevators	12,449	18	0	500	19	42.2	7.5	53.0	7.7	1.0	0.0	0.4647	0.0035	0.4021	0.0102	0.3484	0.0052
						36.4	1.2	39.2	0.8	1.4	0.2	0.4377	0.0030	0.4022	0.0097	0.3482	0.0053
Bike	13,034	17	0	500	19	32.3	1.3	30.6	1.1	1.0	0.0	-0.9976	0.0114	-0.9934	0.0137	0.0446	0.0040
						31.4	2.1	37.1	1.5	0.8	0.0	-0.9985	0.0186	-0.9877	0.0180	0.0454	0.0032
Kin40k	30,000	8	0	500	8	19.0	4.4	186.5	68.0	1.0	0.0	-0.3339	0.0018	-0.3141	0.0016	0.0929	0.0014
						15.0	0.4	44.6	1.3	2.7	0.1	-0.4332	0.0055	-0.3135	0.0038	0.0949	0.0015
Protein	34,297	9	0	500	15	124.2	0.4	892.6	19.1	1.0	0.0	0.9963	0.0031	0.8869	0.0081	0.5722	0.0065
						17.8	0.6	42.5	5.0	4.2	0.4	0.9273	0.0044	0.8835	0.0053	0.5577	0.0079
KEGGdir	36,620	20	0	500	19	55.8	11.7	1450.3	253.4	1.0	0.0	-0.9501	0.0087	-0.9459	0.0343	0.0861	0.0043
						42.9	0.7	173.7	2.9	8.4	1.5	-1.0043	0.0093	-0.9490	0.0312	0.0864	0.0042
3DRoad	326,155	3	0	500	9	68.0		82,200.0		1.0		0.7733		1.4360		0.2982	
						19.0		7306.0		11.3		0.1284		1.1690		0.1265	